

# FRAMING TRADESPACE EXPLORATION TO IMPROVE SUPPORT FOR MULTIPLE-STAKEHOLDER DECISION MAKING

by  
Matthew Edward Fitzgerald

S.M. Aeronautics and Astronautics – Massachusetts Institute of Technology, 2012  
S.B. Aerospace Engineering – Massachusetts Institute of Technology, 2010

Submitted to the Department of Aeronautics and Astronautics  
in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy in Aeronautics and Astronautics

at the  
Massachusetts Institute of Technology  
June 2016

© 2016 Massachusetts Institute of Technology. All rights reserved.

Signature of Author.....  
Department of Aeronautics and Astronautics  
April 27, 2016

Certified by.....  
Daniel E. Hastings  
Cecil and Ida Green Education Professor of Aeronautics and Astronautics and Engineering Systems  
Thesis Chair

Certified by.....  
Adam M. Ross  
Research Scientist, Engineering Systems  
Lead Research Scientist, Systems Engineering Advancement Research Initiative  
Thesis Advisor

Certified by.....  
Joseph M. Sussman  
JR East Professor of Civil and Environmental Engineering and Engineering Systems  
Thesis Committee Member

Certified by.....  
Olivier de Weck  
Professor of Aeronautics and Astronautics and Engineering Systems  
Thesis Committee Member

Accepted by.....  
Paulo C. Lozano  
Associate Professor of Aeronautics and Astronautics  
Chair, Graduate Program Committee



# FRAMING TRADESPACE EXPLORATION TO IMPROVE SUPPORT FOR MULTIPLE-STAKEHOLDER DECISION MAKING

by  
Matthew Edward Fitzgerald

Submitted to the Department of Aeronautics and Astronautics on April 27, 2016 in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy in Aeronautics and Astronautics

## Abstract

As modern engineering projects increase in size and complexity, they have also tended to increase the number of people affected, thus expanding the set of involved stakeholders. The majority of research in tradespace exploration (TSE), as a paradigm for solving complex design problems, has focused on the analysis of the space of alternatives with the goal of uncovering design choices that are optimal or near-optimal. These designs feature desirable combinations of attributes for a given system stakeholder, including technical attributes, cost, and, more recently, -ilities. Less tradespace research has been devoted to the multi-stakeholder problem, in which there are multiple parties with different desired attributes, who must agree on a single design selection in order to proceed with development. Many standard value-measuring techniques, such as utility theory, operate on individuals only and have been shown to break down when used to combine the preferences of groups.

Because of these limitations, multi-stakeholder tradespace exploration (MSTSE) has largely relied on the best practices for *individual* tradespace exploration, with all stakeholders using those methods in parallel. This parallel exploration has the goal of uncovering as many interesting or desirable alternatives as possible, empowering stakeholders to make an educated decision on how best to negotiate with their counterparts. The group decision problem, however, is not just a series of individual decisions and must incorporate interpersonal dynamics and psychological considerations of what makes a “good” decision, and what constitutes a “fair” solution in the minds of the participants.

This thesis describes a research effort to develop the foundations of MSTSE by incorporating fundamental insights from the negotiation and framing literatures. A literature review is used to show that TSE is naturally aligned with the goals of productive negotiation. The framing of data in MSTSE is confirmed, via controlled experiment, to have impacts on negotiation which can be controlled through the visualizations given to the participating stakeholders. A combination of practitioner interviews, analysis of procedures for modern systems engineering methods, and case studies (on aerospace and transportation infrastructure systems) is used to create recommendations for applying MSTSE and demonstrate the new types of insights that can be achieved by doing so, beyond those of prior analyses.

Thesis Chair: Daniel E. Hastings

Title: Cecil and Ida Green Education Professor of Aeronautics and Astronautics and Engineering Systems

Thesis Advisor: Adam M. Ross

Title: Research Scientist, Engineering Systems



*"If life is to be filled with choices, many of them difficult, one ought to have a method with which to approach the art of decision-making. But no, I have left something out. One always has such a method; it is merely the case that one is not always aware of it."*

- Steven Brust, The Phoenix Guards



## Acknowledgements

I will keep this section short only out of necessity – the number of people who have positively impacted my education and research in my ten years at MIT is too large to do justice with any sort of exhaustive accounting. Instead, I will take this space to call out only those who simply cannot be ignored.

First, thank you to my committee – Prof. Hastings, Dr. Ross, Prof. Sussman, and Prof. de Weck – for taking a chance on a thoroughly abstract idea and letting the research pivot (sometimes dramatically) in new and unfamiliar directions as I learned more about multi-stakeholder problems early on. The freedom that I had to search out ideas in a variety of disciplines and ponder their relevance to our “home” disciplines is not one that every student is afforded, but is exactly what I wanted out of my time spent in grad school.

For providing my home-away-from-home (or at least, down the street from home), a big thank you goes to all of SEARI, past and present. Thanks to the leadership, Dr. Donna Rhodes and Dr. Adam Ross, for supporting, guiding, and getting exposure for my research. I’ve enjoyed being here since 2009 – through a UROP, master’s degree, Ph.D., construction, heater floods, bathroom floods, and that one time there was an earthquake (which somehow did not cause any flooding). Thanks to the other students, too many to name, for friendship, advice, commiseration, and the occasional cover story. I hope to get the chance to work with you again.

Thanks to all the people who have helped me, in big and small ways, without getting much in return. That includes the students of the MIT Rail / High-Speed Rail Research Group, for giving me the crash course necessary for my case study. It includes Dr. Hugh McManus, for reaching way back into the past to help me sort out challenges in my other case. And it includes the MIT staff that have answered my questions and made the wheels turn on my graduate career, especially Fran Marrone, Beth Marois, and Marie Stuppard.

Finally, thanks to my family and friends for supporting me. I’m done – see you soon!



## Biographical Note

Matt Fitzgerald is from Dover, Massachusetts. He is the son of Edward and Susan Fitzgerald, and the younger brother of Michael Fitzgerald. Matt graduated from Dover-Sherborn Regional High School in 2006 and then went to MIT and stayed there for a decade, collecting three degrees on the way in 2010, 2012, and 2016. While at MIT, he studied Aeronautics and Astronautics with an emphasis on systems engineering – partially because airplanes are cool, partially just for the challenge. He hopes to continue working to improve not just the systems that engineers make but the ways in which they are designed.

Matt doesn't have any particularly notable accolades. He was runner-up for three senior superlatives in high school and won none, presumably because “most likely to come in second place” was not a category. Matt also lacks any wild hobbies or talents. He spends his free nights watching the Red Sox, jumping at a local concert, seeing a movie, or reading a book. Simple pleasures for a simple man. He hopes to have more time to do those things now that he is finished with school.

Finally, Matt would like to wish the best to any future readers of this thesis. Godspeed, and bring a snack.



# Table of Contents

<b>ABSTRACT</b> .....	<b>3</b>
<b>ACKNOWLEDGEMENTS</b> .....	<b>7</b>
<b>BIOGRAPHICAL NOTE</b> .....	<b>9</b>
<b>TABLE OF CONTENTS</b> .....	<b>11</b>
<b>LIST OF FIGURES</b> .....	<b>16</b>
<b>LIST OF TABLES</b> .....	<b>20</b>
<b>1 INTRODUCTION</b> .....	<b>23</b>
1.1 MOTIVATION.....	23
1.2 RESEARCH QUESTIONS .....	24
1.3 ANTICIPATED CONTRIBUTION .....	24
1.4 METHODOLOGY .....	25
1.4.1 <i>Literature Review</i> .....	25
1.4.2 <i>Theory Building</i> .....	26
1.4.3 <i>Proof of Concept via Controlled Experiment</i> .....	26
1.4.4 <i>Interface with Practice</i> .....	27
1.4.5 <i>Application to Case Studies</i> .....	27
1.5 THESIS OUTLINE .....	27
<b>2 LITERATURE REVIEW</b> .....	<b>31</b>
2.1 MULTI-STAKEHOLDER DECISION MAKING .....	31
2.1.1 <i>Single Stakeholder, Multiple Goals</i> .....	31
2.1.2 <i>Connected Stakeholders and Aggregation Methods</i> .....	33
2.1.3 <i>Independent, Cooperative Decision Makers</i> .....	36
2.1.4 <i>Negotiation Analysis and Practice</i> .....	37
2.2 TRADESPACE EXPLORATION .....	40
2.2.1 <i>Multi-Stakeholder Tradespace Exploration</i> .....	42
2.3 FRAMING .....	44
2.3.1 <i>“Macro” Framing</i> .....	44
2.3.2 <i>“Micro” Framing</i> .....	46
2.3.3 <i>Visual Analytics</i> .....	47
2.4 THE –ILITIES .....	48
<b>3 THEORY BUILDING FOR MSTSE FRAMING</b> .....	<b>51</b>
3.1 PRINCIPLED NEGOTIATION AND TSE .....	51
3.2 FRAMING IN TSE .....	56
3.3 FRAMING MISMATCH IN MSTSE.....	58
3.4 WORKING THEORY .....	59
3.4.1 <i>Macro Framing in MSTSE</i> .....	59
3.4.2 <i>Micro Framing in MSTSE</i> .....	60
3.4.3 <i>Macro and Micro Interaction</i> .....	61
3.5 IMPROVING MSTSE FRAMING.....	61
3.6 MODIFYING TRADESPACE VISUALIZATIONS .....	62

3.7	DISCUSSION .....	64
<b>4</b>	<b>TRADESPACE REDESIGN EXPERIMENT .....</b>	<b>67</b>
4.1	EXPERIMENT OVERVIEW AND OBJECTIVES .....	67
4.2	UNIT OF ANALYSIS .....	68
4.3	SUBJECT SELECTION AND ASSIGNMENT.....	68
4.4	TRIAL PROTOCOL.....	69
4.4.1	<i>Pre-task</i> .....	69
4.4.2	<i>Experimental Task</i> .....	70
4.4.3	<i>Post-task</i> .....	73
4.5	ANALYSIS .....	73
4.5.1	<i>Demographics</i> .....	73
4.5.2	<i>Questionnaire – Multiple Choice</i> .....	74
4.5.3	<i>Questionnaire – Short Answer and Open Response</i> .....	75
4.5.4	<i>Offer History and Outcomes</i> .....	79
4.5.5	<i>Task Observation</i> .....	85
4.5.6	<i>Interesting Anecdotes</i> .....	86
4.6	DISCUSSION .....	88
<b>5</b>	<b>DEVELOPING MSTSE VISUALIZATIONS WITH ITERATIVE PRACTITIONER FEEDBACK .....</b>	<b>91</b>
5.1	GUIDING PRINCIPLES .....	91
5.2	FURTHER IMPROVING THE NEGOTIATION TRADESPACE.....	92
5.3	VISUALIZING RELATIONSHIPS .....	96
5.4	TRADESPACE EXPLORATION AND STAKEHOLDER CORRELATION .....	98
5.5	ACCESSIBILITY OF PARETO SETS .....	103
5.6	INTERACTIVE PRACTITIONER INTERVIEWS.....	107
5.6.1	<i>Interview Structure</i> .....	107
5.6.2	<i>Participants</i> .....	109
5.6.3	<i>Personal Multi-stakeholder Experience</i> .....	110
5.6.4	<i>Response to TSE</i> .....	112
5.6.5	<i>Response to MSTSE</i> .....	113
5.6.6	<i>Reflection on Applicability</i> .....	115
5.7	DISCUSSION .....	116
<b>6</b>	<b>IMPACT OF PROBLEM STRUCTURE ON MSTSE .....</b>	<b>119</b>
6.1	STAKEHOLDERS .....	120
6.1.1	<i>Number</i> .....	120
6.1.2	<i>Outside Relationships</i> .....	122
6.1.3	<i>Representation</i> .....	123
6.2	PREFERENCES (BENEFITS/COSTS) .....	123
6.2.1	<i>Alignment</i> .....	124
6.2.2	<i>Single Stakeholder Value Models</i> .....	126
6.2.3	<i>Multiple Stakeholder Aggregation</i> .....	129
6.2.4	<i>Unoperationalizable Preferences</i> .....	130
6.2.5	<i>Metapreferences</i> .....	131
6.2.6	<i>Divisible Attributes</i> .....	132

6.3	ALTERNATIVES .....	134
6.3.1	<i>BATNA Type</i> .....	135
6.3.2	<i>BATNA Strength</i> .....	140
6.3.3	<i>Tradespace Size</i> .....	142
6.3.4	<i>Tradespace Completeness</i> .....	144
6.4	UNCERTAINTY .....	145
6.4.1	<i>Contextual Uncertainty</i> .....	145
6.4.2	<i>Model Uncertainty</i> .....	146
6.4.3	<i>Preference Uncertainty</i> .....	147
6.4.4	<i>The -ilities</i> .....	149
6.5	LOGISTICS .....	151
6.5.1	<i>Decision Maker Access</i> .....	152
6.5.2	<i>Schedule / Timeline</i> .....	153
6.5.3	<i>Private Information</i> .....	154
6.6	DISCUSSION .....	156
<b>7</b>	<b>RECOMMENDATIONS FOR PERFORMING MSTSE.....</b>	<b>159</b>
7.1	PROBLEM FORMULATION.....	160
7.2	MODELING / EVALUATION .....	163
7.3	EXPLORATION / ANALYSIS.....	165
7.4	FRAMING INFORMAL MSTSE WITHOUT NEGOTIATION.....	166
<b>8</b>	<b>MSTSE WITHIN SE METHODS.....</b>	<b>171</b>
8.1	THE CLIOS PROCESS.....	171
8.1.1	<i>CLIOS Stages</i> .....	171
8.1.2	<i>CLIOS Steps</i> .....	173
8.1.3	<i>Relation to SAM-PD</i> .....	177
8.1.4	<i>Summary</i> .....	178
8.2	USAF ANALYSIS OF ALTERNATIVES .....	179
8.2.1	<i>Overview</i> .....	179
8.2.2	<i>Tradespace Exploration in AoA</i> .....	181
8.2.3	<i>Multi-stakeholder in AoA</i> .....	182
8.2.4	<i>Using MSTSE in AoA</i> .....	183
8.2.5	<i>Summary</i> .....	185
8.3	INCOSE DECISION MANAGEMENT PROCESS .....	186
8.3.1	<i>Process Overview</i> .....	186
8.3.2	<i>Trade Study Framework</i> .....	188
8.3.3	<i>Summary</i> .....	192
8.4	RESPONSIVE SYSTEMS COMPARISON .....	193
8.4.1	<i>Stakeholder Participation</i> .....	194
8.4.2	<i>Multi-stakeholder Analysis</i> .....	197
8.4.3	<i>RSC Feedback</i> .....	200
8.4.4	<i>Summary</i> .....	201
8.5	DISCUSSION .....	201
<b>9</b>	<b>CASE STUDY – SATELLITE RADAR.....</b>	<b>203</b>
9.1	HISTORY .....	203

9.2	ORIGINAL TRADESPACE FORMULATION.....	205
9.2.1	<i>Context Space</i> .....	206
9.2.2	<i>Design Space</i> .....	207
9.2.3	<i>Evaluative Models (Performance and Cost)</i> .....	209
9.2.4	<i>Value Models</i> .....	209
9.3	INSIGHTS AND CONCLUSIONS OF ORIGINAL NEGOTIATION SESSION .....	212
9.4	NEGOTIATION STRUCTURE.....	215
9.5	MSTSE ANALYSIS.....	218
9.6	REVISED TRADESPACE FORMULATION AND MODELING.....	231
9.7	VALIDATING THE REVISED TRADESPACE.....	233
9.8	PASSIVELY ROBUST ALTERNATIVES .....	240
9.9	CHANGEABILITY-ENABLED ROBUST ALTERNATIVES .....	243
9.10	LEVERAGING CHANGEABILITY TO SATISFY MULTIPLE STAKEHOLDERS AT LOW COST.....	246
9.11	DISCUSSION .....	257
<b>10</b>	<b>CASE STUDY – NORTHEAST CORRIDOR.....</b>	<b>263</b>
10.1	NORTHEAST CORRIDOR OVERVIEW .....	263
10.2	CREATING THE NEC TRADESPACE.....	264
10.2.1	<i>Needs and Value Models</i> .....	265
10.2.2	<i>Design Space</i> .....	269
10.2.3	<i>Evaluative Models (Performance and Cost)</i> .....	270
10.2.4	<i>Contexts</i> .....	272
10.3	NEGOTIATION STRUCTURE.....	273
10.4	NEGOTIATION SUCCESS CRITERIA .....	275
10.5	NEGOTIATION TRADESPACE AND FAVORITE ALTERNATIVES.....	277
10.6	STAKEHOLDER ALIGNMENT AND COALITIONS.....	280
10.7	FINDING MUTUALLY BENEFICIAL ALTERNATIVES.....	284
10.8	ADDING NEW ALTERNATIVES.....	289
10.9	BEYOND FOTE MSTSE .....	292
10.10	DISCUSSION .....	296
<b>11</b>	<b>DISCUSSION .....</b>	<b>302</b>
11.1	BENEFITS AND APPLICABILITY.....	302
11.2	FUTURE WORK .....	304
11.2.1	<i>Developing More Visualizations</i> .....	304
11.2.2	<i>Expanding Data Collection on Tradespace Framing</i> .....	305
11.2.3	<i>Culture of Negotiation in Engineering</i> .....	306
11.2.4	<i>Game Theory in MSTSE</i> .....	308
11.2.5	<i>Dividing Costs for Shared Systems</i> .....	309
11.2.6	<i>MSTSE for Detailed Design</i> .....	310
<b>12</b>	<b>CONCLUSION .....</b>	<b>312</b>
12.1	ANSWERING THE RESEARCH QUESTIONS .....	312
12.2	ADDITIONAL CONTRIBUTIONS .....	316
12.3	FINAL THOUGHTS .....	319
<b>13</b>	<b>GLOSSARY.....</b>	<b>322</b>

13.1	KEY TERMS.....	322
13.2	OTHER ACRONYMS .....	324
<b>14</b>	<b>REFERENCES.....</b>	<b>326</b>
<b>15</b>	<b>APPENDICES .....</b>	<b>340</b>
APPENDIX A	DESCRIPTION OF TRADESPACE REDESIGN EXPERIMENT QUESTIONNAIRE.....	340
APPENDIX B	NEC VALUE MODEL.....	343

## List of Figures

<i>Figure 1-1: The theory-building pyramid (Carlile and Christensen, 2004)</i> .....	26
<i>Figure 1-2: Research outline</i> .....	29
<i>Figure 2-1: Asymmetric perceived value around a reference point, according to Prospect Theory (Tversky and Kahneman, 1981)</i> .....	46
<i>Figure 3-1: "Midpoint" solutions in the design space do not always occupy a midpoint in the value space</i> .....	52
<i>Figure 3-2: "Gold plated" solutions combine favorable design features, but result in vastly higher costs</i> .....	52
<i>Figure 3-3: An example two-stakeholder tradespace, with both integrative and distributive decision elements</i> .....	53
<i>Figure 3-4: Engaging in distributive bargaining before integrative bargaining (via creation of many alternatives) limits the total amount of value available in the tradespace</i> .....	54
<i>Figure 3-5: Disconnect between controllable variables and preferences in engineering applications</i> .....	56
<i>Figure 3-6: Utility = 1 reference point, prior to plotting designs, with associated "loss" once designs are seen</i> .....	57
<i>Figure 3-7: Tradespace with plotted designs, "loss" evaluated with reference to Pareto front</i> ..	58
<i>Figure 3-8: Increasing complexity from problem definition to physical constraints to relational concerns</i> .....	59
<i>Figure 3-9: A cost-utility tradespace before (a) and after (b) being centered on the BATNA</i> .....	63
<i>Figure 3-10: Example re-centered, rotated, and color/transparency MSTSE visualization</i> .....	64
<i>Figure 4-1: The standard benefit-cost tradespace view (for Nat) given to the control group</i> .....	71
<i>Figure 4-2: The negotiation tradespace view (for Nat) given to the treatment group, annotated with the key differences from the standard tradespace</i> .....	72
<i>Figure 4-3: Slightly more pronounced "knee" to the Pareto front on Nat's tradespace (left) than Vic's (right)</i> .....	76
<i>Figure 4-4: Example "rational" responses on the control (left) and treatment (right) questionnaires</i> .....	77
<i>Figure 4-5: "Other" responses, displaying benefit-cost confusion</i> .....	77
<i>Figure 4-6: Nat (left) and Vic (right) tradespaces, with final selections marked</i> .....	79
<i>Figure 4-7: Nat's offer trajectories at minute 26 (control on left, treatment on right)</i> .....	82
<i>Figure 4-8: Vic's offer trajectories at minute 40 (control on left, treatment on right)</i> .....	83
<i>Figure 4-9: All offers by FPN and time</i> .....	84
<i>Figure 5-1: The negotiation tradespace, set to show the same stakeholder for axes and color/transparency</i> .....	93
<i>Figure 5-2: The negotiation tradespace, with color and transparency set to 4 different stakeholders. Darker regions are more opaque and favored by that stakeholder</i> .....	94
<i>Figure 5-3: Negotiation tradespace with a maximum display FPN of 10, 30, and 50, respectively</i> .....	94
<i>Figure 5-4: Negotiation tradespace with maximum display FPN of 10, without and with minimum opacity</i> .....	95
<i>Figure 5-5: An all-Quadrant-1 tradespace, rotated back to natural position and with color disabled, but still bolding the BATNA axes</i> .....	96

Figure 5-6: Increasing complexity in the multi-stakeholder problem demands consideration of more than alternatives .....	97
Figure 5-7: Correlation heatmap visualizing benefit correlation for five stakeholders, in low- and high-detail modes.....	99
Figure 5-8: Correlations on Benefit, Cost, and FPN for the Used Car stakeholders .....	100
Figure 5-9: Correlation Map, toggling metric used and invalid inclusion .....	101
Figure 5-10: Attribute Correlation impact visualization, comparing three attributes of one stakeholder against four other stakeholders.....	102
Figure 5-11: Venn diagram visualization of three stakeholder Pareto sets.....	104
Figure 5-12: Venn diagram changing from 0% to 10% fuzziness.....	105
Figure 5-13: Gridmap visualization of four stakeholder's Pareto sets .....	106
Figure 5-14: Gridmap changing from 0% to 10% fuzziness .....	107
Figure 6-1: Notional Utility-Utility axes showing different amounts of preference alignment .	124
Figure 6-2: Stakeholder 2 tradespace, with "optimal" trades against Stakeholder 1's value metric .....	128
Figure 6-3: Tradespace differences between full cost, even split cost, and uneven split costs ..	132
Figure 6-4: The Do Nothing BATNA results in all alternatives in Quadrant 1 .....	135
Figure 6-5: "Work Alone" tradespace identifies BATNA, "Work Together" tradespace highlights newly affordable alternatives.....	137
Figure 6-6: Estimating an unknown "other opportunity" BATNA .....	140
Figure 6-7: Clear 'gains' are limited by a strong BATNA, exacerbated by multiple strong BATNAs.....	141
Figure 6-8: Linked scatterplots of design variables, with stakeholder favorites marked in blue and red triangles.....	143
Figure 6-9: Model errors with different impacts depending on a design variable may misrepresent ordering.....	147
Figure 6-10: Change mechanism connects alternatives that are on each stakeholder's Pareto front.....	151
Figure 8-1: The Stages (left) and Steps (right) of the CLIOS Process (Sussman et al., 2009) ..	174
Figure 8-2: Mapping the CLIOS and SAM-PD Processes (Mostashari, 2005).....	178
Figure 8-3: USAF system lifecycle outline, indicating AoA as pre-Milestone-A (Office of Aerospace Studies, 2013).....	180
Figure 8-4: INCOSE Decision Management Process "context diagram" (INCOSE, 2012) .....	187
Figure 8-5: INCOSE Decision Management Process v4.0 draft (Cilli and Parnell, 2014).....	189
Figure 8-6: RSC flowchart (Rhodes, Ross, and Hastings, 2008) .....	193
Figure 8-7: RSC vs. the System Acquisition Lifecycle (Rhodes, Ross, and Hastings, 2010).....	195
Figure 8-8: Example tradespaces where some "compromises" are "fair" and some are not ...	199
Figure 9-1: Negotiation tradespace for Stakeholder E, colored for each other stakeholder.....	219
Figure 9-2: Negotiation tradespace for Stakeholder N, colored for each other stakeholder.....	219
Figure 9-3: Negotiation tradespace for Stakeholder R, colored for each other stakeholder.....	220
Figure 9-4: Satellite Radar Pareto Venn diagram .....	221
Figure 9-5: Satellite Radar Pareto Venn diagram - 0.1% fuzzy .....	222
Figure 9-6: Satellite Radar gridmap - 0.1% fuzzy.....	223
Figure 9-7: Satellite Radar negotiation tradespaces with 0.1% fuzzy Pareto set intersections marked.....	225
Figure 9-8: Satellite Radar utility function correlations.....	227

<i>Figure 9-9: Satellite Radar FPN correlations.....</i>	<i>227</i>
<i>Figure 9-10: Satellite Radar FPN correlations - Scatterplot mode .....</i>	<i>228</i>
<i>Figure 9-11: Attribute correlations for Stakeholder E.....</i>	<i>229</i>
<i>Figure 9-12: Attribute correlations for Stakeholder N.....</i>	<i>230</i>
<i>Figure 9-13: Attribute correlations for Stakeholder R.....</i>	<i>230</i>
<i>Figure 9-14: Geolocation Accuracy Single Attribute Utility function for Stakeholder N, relaxing requirement from 100 m to 200 m.....</i>	<i>235</i>
<i>Figure 9-15: Negotiation tradespaces for context #72 in revised Satellite Radar .....</i>	<i>236</i>
<i>Figure 9-16: Stakeholder utility correlations for original and revised tradespaces.....</i>	<i>237</i>
<i>Figure 9-17: Stakeholder FPN correlations for original and revised tradespaces.....</i>	<i>237</i>
<i>Figure 9-18: Pareto Venn diagram (0.1% fuzzy) for revised tradespace.....</i>	<i>238</i>
<i>Figure 9-19: Yield histograms for E (left), R (center), and N (right) across the context space. ....</i>	<i>240</i>
<i>Figure 9-20: Combining the accessibility matrices of different types of orbital burns.....</i>	<i>244</i>
<i>Figure 9-21: Satellite Radar design groups, by cost and minimum 0.1% efNPT, colored by <math>\Delta V</math> cost to change stakeholders .....</i>	<i>247</i>
<i>Figure 9-22: Satellite Radar design groups, by cost and minimum 0.1% efNPT, colored by <math>\Delta V</math> cost expended in response to a context shift.....</i>	<i>248</i>
<i>Figure 9-23: Satellite Radar design groups, by cost and sum of 0.1% efNPTs, colored by <math>\Delta V</math> cost to change stakeholders.....</i>	<i>250</i>
<i>Figure 9-24: Satellite Radar design groups, by cost and minimum 0.1% efNPT, colored by <math>\Delta V</math> cost to change stakeholders, reduced burn strategy.....</i>	<i>251</i>
<i>Figure 9-25: Satellite Radar design groups, by cost and minimum 0.1% efNPT, colored by <math>\Delta V</math> cost expended in response to a context shift, reduced burn strategy.....</i>	<i>252</i>
<i>Figure 9-26: Mission success rate for designs of interest with and without orbit changes .....</i>	<i>255</i>
<i>Figure 9-27: Minimum 0.1% efNPT for designs of interest, "prevent failure" vs. "maximize efficiency".....</i>	<i>255</i>
<i>Figure 9-28: Mockup of a "Stakeholder 1" bubble plot, with vectors indicating relative position of Stakeholders 2 (red) and 3 (green).....</i>	<i>261</i>
<i>Figure 10-1: A generic model-centric decision making framework, capable of supporting TSE .....</i>	<i>265</i>
<i>Figure 10-2: NEC Stakeholder clustering (Sussman et al., 2015; Moody, 2016), with additional annotations in red indicating the chosen MSTSE stakeholders.....</i>	<i>266</i>
<i>Figure 10-3: USDOT negotiation tradespace in base context .....</i>	<i>278</i>
<i>Figure 10-4: Negotiation tradespaces for all stakeholders in the base context, with "favorite" bundles circled.....</i>	<i>279</i>
<i>Figure 10-5: NEC stakeholder utility correlations without (left) and with (right) invalid bundles included.....</i>	<i>280</i>
<i>Figure 10-6: NEC stakeholder cost correlation .....</i>	<i>281</i>
<i>Figure 10-7: NEC stakeholder FPN correlation, with three main coalitions highlighted.....</i>	<i>282</i>
<i>Figure 10-8: Coalitions identified by clustering analysis, compared to FPN correlation (with invalid bundles).....</i>	<i>283</i>
<i>Figure 10-9: Gridmap of base context, showing Bundle 6 with 9/10 stakeholders and none other above 5 .....</i>	<i>285</i>
<i>Figure 10-10: 10% fuzzy gridmap of base context, Bundle 5 is first to 10/10 stakeholders .....</i>	<i>285</i>
<i>Figure 10-11: Negotiation tradespaces for all stakeholders in chosen contexts, with notable changes from base circled .....</i>	<i>288</i>

<i>Figure 10-12: USDOT negotiation tradespace in preferred context, with new bundles.....</i>	<i>290</i>
<i>Figure 10-13: Comparison of notional Air Industry tradespace if engaged during MSTSE vs. after agreement on preferred bundle .....</i>	<i>295</i>
<i>Figure 10-14: Private Consortiums tradespace: piecewise-HSR bundles (3/4) dominate all-over-HSR bundles (5/6) due to unanticipated model interaction.....</i>	<i>299</i>
<i>Figure 12-1: Negotiation experiment offers by FPN and time, showing change in initial reference point and search direction based on treatment .....</i>	<i>314</i>
<i>Figure 12-2: Fixed gains from changeability (height above diagonal) provide greater proportional improvement to low cost solutions (smaller circles).....</i>	<i>316</i>
<i>Figure 12-3: The "negotiation tradespace" and its key modifications from the standard tradespace.....</i>	<i>317</i>
<i>Figure 12-4: Relationship visualizations, deemphasizing the individual alternatives relative to most classic TSE visualizations .....</i>	<i>318</i>
<i>Figure 12-5: Correlation plots, identifying possible stakeholder coalitions for Satellite Radar (left) and Northeast Corridor (right).....</i>	<i>319</i>

## List of Tables

<i>Table 2-1: Human Endeavor Characteristics (Lu et al., 2007)</i> .....	35
<i>Table 4-1: Demographics summary</i> .....	74
<i>Table 4-2: Questionnaire between groups t-test p-values</i> .....	75
<i>Table 4-3: Preferred Regions of Tradespace</i> .....	76
<i>Table 4-4: Final Agreement Information</i> .....	80
<i>Table 5-1: Interview Participants - Field and Experience</i> .....	110
<i>Table 6-1: Summary of problem structure with impact on MSTSE</i> .....	157
<i>Table 7-1: Summary of MSTSE recommendations, chapters to refer to for more detail on each, and applicability for informal MSTSE</i> .....	170
<i>Table 9-1: Satellite Radar context space description</i> .....	206
<i>Table 9-2: Satellite Radar design space description</i> .....	208
<i>Table 9-3: Walker constellation lookup table</i> .....	208
<i>Table 9-4: Satellite Radar attributes able to be chosen by stakeholders</i> .....	210
<i>Table 9-5: Questionnaire results for Stakeholder E</i> .....	211
<i>Table 9-6: Questionnaire results for Stakeholder N</i> .....	211
<i>Table 9-7: Questionnaire results for Stakeholder R</i> .....	211
<i>Table 9-8: Satellite Radar initial favorite designs (adapted from Ross et al., 2010a)</i> .....	213
<i>Table 9-9: Satellite Radar group selections; asterisks indicate the method of finding them - * compromise Pareto set, ** “gold plated” (adapted from Ross et al., 2010a)</i> .....	214
<i>Table 9-10: Structural features of Satellite Radar tradespace</i> .....	217
<i>Table 9-11: Shared design variables of Pareto set intersections</i> .....	224
<i>Table 9-12: Revised Satellite Radar design space description</i> .....	232
<i>Table 9-13: Revised Satellite Radar context space description</i> .....	233
<i>Table 9-14: Shared design variables of revised Pareto set intersections (bold indicates same as original)</i> .....	239
<i>Table 9-15: Highest NPT alternative(s) for each stakeholder, and highest maximin NPT across stakeholders</i> .....	241
<i>Table 9-16: Highest 0.1% fNPT alternatives for each stakeholder, and highest maximin fNPT across stakeholders</i> .....	242
<i>Table 9-17: Highest 0.1% efNPT alternatives for each stakeholder, and highest maximin efNPT across stakeholders</i> .....	245
<i>Table 9-18: Designs of interest trading cost and minimum 0.1% efNPT</i> .....	249
<i>Table 9-19: Minimum 0.1% efNPT for designs of interest under different change strategies ...</i>	253
<i>Table 9-20: Mission success rate (stakeholder average) for designs of interest under different change strategies</i> .....	254
<i>Table 10-1: Summary of NEC stakeholder value models</i> .....	268
<i>Table 10-2: NEC design space, comprised of six bundles (Sussman et al., 2015)</i> .....	269
<i>Table 10-3: Cost division for Private+State pool</i> .....	271
<i>Table 10-4: NEC context variables</i> .....	272
<i>Table 10-5: Structural features of NEC tradespace</i> .....	274
<i>Table 10-6: NEC stakeholders mandatory for some bundles</i> .....	277
<i>Table 10-7: Bundles within 5% of the Pareto front in the base context</i> .....	286
<i>Table 10-8: NEC stakeholders’ preferred contexts (noting changes from the base)</i> .....	287
<i>Table 10-9: Bundles within 5% of the Pareto front in stakeholder-chosen contexts</i> .....	288
<i>Table 10-10: Creation of new alternatives by shifting funding sources</i> .....	290

<i>Table 10-11: Number of stakeholders with each bundle at different levels of efficiency, in preferred contexts .....</i>	<i>291</i>
<i>Table 15-1: US Dept. of Transportation Utility function .....</i>	<i>344</i>
<i>Table 15-2: Amtrak and Rail Agencies Utility function.....</i>	<i>344</i>
<i>Table 15-3: Congress Utility function .....</i>	<i>345</i>
<i>Table 15-4: Northern States Utility function .....</i>	<i>345</i>
<i>Table 15-5: Southern States Utility function .....</i>	<i>346</i>
<i>Table 15-6: EPA and Landowners Utility function .....</i>	<i>346</i>
<i>Table 15-7: Private Consortiums Utility function .....</i>	<i>347</i>
<i>Table 15-8: Suppliers and Labor Utility function.....</i>	<i>347</i>
<i>Table 15-9: Highway System Utility function.....</i>	<i>347</i>
<i>Table 15-10: Travelers Utility function.....</i>	<i>348</i>



# 1 Introduction

## 1.1 Motivation

Engineering projects have been experiencing a consistent trend over the previous decades, increasing steadily in cost and complexity as engineers continue to squeeze more and more performance out of their systems. One attempt to combat this trend has been an increased emphasis on affordability in recent years (Carter 2010a, 2010b), representing a cultural shift in acknowledgment that the benefits of optimizing performance are frequently not worth the corresponding increases in cost – one that mirrors the well-known “80-20” heuristic, where the last 20% of performance may bring with it 80% of the total cost. An alternative strategy to managing size and complexity has been to involve additional stakeholders in the design process. The combined resources of multiple stakeholders, including money, expertise, and authority, among others, may enable the resultant system to provide significantly more benefit than would be achieved by any party alone or all parties independently. Other interested parties are often not difficult to find, as the increased size and impact of large engineering systems often affects large groups of people. Frequently, the reason for seeking additional stakeholders is also implicitly or explicitly tied to affordability, because working together is heuristically less expensive than working alone, either due to sharing costs or economies of scale.

Regardless of the reason behind engaging in a multi-stakeholder project, the design process itself becomes more challenging when there is no longer a unilateral decision making authority. The needs and interests of participating stakeholders must be met and balanced in order to reach an agreement, but the lack of a unilateral decision making authority adds a new dimension to the problem solving tasks of designing an engineering system. This problem can become intractable when the interests of the different parties are unaligned, such that the traditional single-stakeholder-optimal solutions are significantly different. For example, the USAF’s attempts to build a satellite radar system, started and stopped on multiple occasions throughout its history, were specifically cancelled in 2005 by reduced funding due to incompatibility between the desired CONOPs of participating divisions (Satellite Today, 2008; Braganca, 2011; Los Angeles AFB, 2011). Project cancellations such as this represent a significant potential loss in value, particularly when it can be shown that the space of alternatives presents many opportunities for mutual benefit to all stakeholders (Ross et al., 2009). Thus, techniques for improving the multi-stakeholder design process are of apparent interest to the engineering community.

Tradespace exploration (TSE) has emerged in recent years as a powerful tool for solving complex design problems (Ross and Hastings, 2005; Stump et al., 2009). The methods and techniques developing around TSE have commonly been geared towards finding optimal or near-optimal solutions, with high performance and low cost. Current TSE research has placed a significant emphasis on the –ilities, design properties that deliver value over time in ways other

than mission performance, trying to exploit TSE's natural ability to capture and explain complex interrelationships between design variables and uncertainty. Considerably less research effort has targeted multi-stakeholder tradespace exploration (MSTSE), which has, to this point, largely entailed the execution of individual TSE by each stakeholder in parallel, only coming together to discuss at fixed intervals (Ross et al., 2010a). Structuring the interactions of MSTSE in this way was driven partially by the inability of many preference- and value-measuring techniques, such as utility theory, to be successfully aggregated across multiple people.

In contrast with the idea of simply aggregating stakeholder preferences, this research seeks to address the challenge of MSTSE by incorporating insights from negotiation theory, framing, and engineering practice to restructure the underlying individual-TSE framework of MSTSE to better support the goals and activities present in productive multi-stakeholder bargaining and decision making. This approach has the potential to improve both the likelihood of reaching an agreement (avoiding unnecessary project cancellations) and increase the design quality and stakeholder satisfaction with the outcome of MSTSE. Furthermore, this research is necessary in order to create a strong foundation for future MSTSE research to build on, rather than relying on extrapolation from single-stakeholder TSE research.

## **1.2 Research Questions**

This research set out to answer the following research questions, which were identified as crucial for the future development and practice of MSTSE:

- RQ1. Are the principles of tradespace exploration (TSE) fundamentally aligned with those of complex, sociotechnical negotiations?
- RQ2. Has the evolution of multi-stakeholder tradespace exploration (MSTSE), as an offshoot of single-stakeholder TSE, resulted in unintentional framing effects impacting decision making, and can those effects be controlled?
- RQ3. How can MSTSE be effectively incorporated into a design process, such that it best complements the tasks required by practicing engineers and the needs of decision makers?
- RQ4. Can –ilities contribute to MSTSE as a potential avenue for creating mutual value and breaking impasses?

## **1.3 Anticipated Contribution**

This research is firmly located within the TSE field. The main contribution of this research should be a solid foundation for the practice of MSTSE, with emphasis on how it differs from traditional TSE. In support of this goal – and because MSTSE is an emerging area of tradespace research and lacks a substantial existing literature base – this research will primarily emphasize systems with cooperative stakeholders who are willing to participate in good-faith negotiation. This represents a scoping decision to focus on the “basics” of multi-stakeholder

decision making, before future research can develop MSTSE for use in competitive environments.

Answering the research questions will involve a systematic analysis of the principles and assumptions of the tradespace paradigm and how they map into the social domain of multi-stakeholder decision making. Correspondingly, any limitations of existing MSTSE formulations or practices discovered in this way can be targeted with prescriptive recommendations designed to augment the ability of MSTSE to support mutually beneficial decisions, using activities and visualizations created specifically for use with multiple stakeholders. These recommendations will be targeted to practicing systems engineers interested in deploying TSE on multi-stakeholder problems and communicating the results of technical design to non-technical stakeholders. However, technically-minded or involved stakeholders may also be interested in using MSTSE as a negotiation aid, given its potential to “close the loop” on iterative design by increasing the amount of available information with which stakeholders can make decisions.

In terms of specific lessons, emphasis will be placed on the appropriateness of framing for the overarching and intermediate decision tasks that are a part of MSTSE. Additionally, the recommendations will be tailored to fit within existing systems engineering methods and frameworks, so as to minimize the inertial barriers to adoption and maximize the potential benefit in application. Finally, the use of –ilities to resolve differences in stakeholder needs or beliefs has been left largely unexplored in favor of the emerging development of rigorous –ility valuation metrics for single stakeholders; this research hopes to demonstrate their potential to not only provide value in the face of uncertainty but also to improve performance across different stakeholders’ interests simultaneously.

## **1.4 Methodology**

To investigate these research questions and achieve the desired contribution, the following steps were taken: 1) literature review, 2) theory building, 3) proof of concept via controlled experiment, 4) interface with practice, and 5) application to case studies. These steps are now described.

### **1.4.1 Literature Review**

This research is positioned as a synthesis of modern systems engineering, specifically the field of tradespace exploration, with fundamental insights from negotiation and behavioral economics research. The literature review was conducted in depth on TSE, covering its foundations and recent developments in multi-stakeholder value modeling and decision support. The review of the other topics was performed with the goal of capturing key insights from the main body of research (rather than current, developing topics) that have solidified enough to be applied in the systems engineering domain. The literature review is the primary tool for addressing RQ1. Analysis of the fundamental activities and ideology underlying TSE provides

the necessary information to determine if it is well-suited to address the challenges of multi-stakeholder problems as identified by research in negotiation and conflict resolution.

### 1.4.2 Theory Building

The process of theory building typically involves trips up and down the “theory-building pyramid”, as depicted in Figure 1-1, usually beginning with inductive theory (Carlile and Christensen, 2004). In this research, the general motivation combined with the observations of previous research in the literature review was used to inductively model a working theory of how stakeholders interact with information to solve multi-stakeholder problems in the tradespace. The further development of MSTSE presents an opportunity to influence this interaction and improve it. In particular, this inductive theory was used to partially answer RQ2, forming a hypothesized model for the impact of framing on decision making in TSE. Then, from that model, promising ways in which to control the negative effects of that framing were deduced.

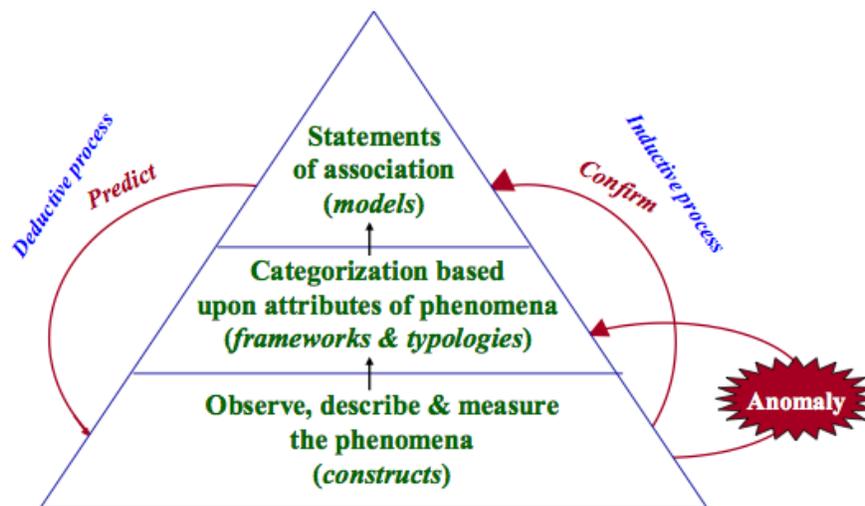


Figure 1-1: The theory-building pyramid (Carlile and Christensen, 2004)

### 1.4.3 Proof of Concept via Controlled Experiment

A controlled experiment was used to inductively confirm the expected framing impacts of the previous subsection and answer the second half of RQ2. Subjects were randomly divided into two-person teams and assigned to either a control group, with access to standard, previously-deployed tradespace exploration and visualization tools, or a treatment group, which had modified tools developed using the ideas of the theory building activity. Resource limitations inhibit the ability for strong normative arguments to result from the experiment, but the treatment group displayed statistically significant improvements in their understanding of the group problem, among other outcomes predicted by the theory. Additionally, the resulting observational data from the experiment (including the identification of anomalies) was used to make more prescriptive recommendations for the practice of MSTSE.

#### **1.4.4 Interface with Practice**

In the spirit of actionable prescriptive research, this work intends to position itself as more than a theoretical contribution. RQ3 targets the practical application of MSTSE by acknowledging the challenging realities of systems engineering, in order to ensure the applicability and feasibility of any insights that are developed during this research. To this end, practitioners were consulted, via semi-structured interviews, about their interactions with TSE and MSTSE in their careers as well as their current desires for analysis tools. We also collected practitioners' responses to new tradespace visualizations designed specifically to support MSTSE by controlling framing and conflict related issues, and used their feedback to iteratively improve those visualizations. Similarly, end-to-end design processes that incorporate TSE, such as the Responsive Systems Comparison method (RSC) and the USAF Analysis of Alternatives (AoA), were examined in detail for potential mismatches in objectives when applied to a multi-stakeholder problem (Ross et al., 2008; US Air Force, 2010). Specific recommendations for how best to deploy MSTSE within these methods were derived from this analysis, as a means of reducing the barriers to entry for "real world" application of MSTSE.

#### **1.4.5 Application to Case Studies**

The recommendations of the research were applied to reconstructed case studies in the domains of space architecture and transportation infrastructure, two areas with considerable multi-stakeholder challenges. Space systems, due to the high cost of launching mass into orbit, are often highly driven towards performance optimality and yet also attract different stakeholders with the promise of shared satellite uptime and distributed costs. Transportation infrastructure by its very nature frequently crosses jurisdictional boundaries and requires the approval of multiple governmental and private authorities, often with conflicting interests. The demonstration of this research using case studies serves to clarify its application for potential adopters. The study of specific cases also provides more concrete grounds for the discussion and analysis of -ilities in MSTSE, assisting the investigation of RQ4. Additionally, the case studies are necessary to support validation, which is often challenging for subjective human-in-the-loop techniques such as tradespace exploration due to the challenge of isolating the impact of the technique itself when seeking empirical performance validity (Pederson et al., 2000). To address this challenge, the insights generated by MSTSE were directly compared against prior analyses on the same cases.

### **1.5 Thesis Outline**

The following chapters of this thesis cover the results of enacting the above plan for this research. These can be roughly divided into six sections:

- Chapter 2 presents a background overview, encompassing key results from the existing literature that are used to build up the working theory of how stakeholders respond to

TSE artifacts and methods. This literature review mainly focuses on three domains: tradespace exploration, multi-stakeholder decision making, and framing.

- Chapter 3 combines the insights of the presented literature into a proposed theory of how framing drives stakeholder perception of complex design problems, and how these perceptions can be guided to more accurately represent the given problem.
- The third section describes the work performed to develop the theory based on the literature. Chapter 4 explains the formulation and results of a controlled experiment to test the underlying theory of framing in MSTSE on a sample negotiation task. Chapter 5 describes the creation of additional visualizations specifically targeting the multi-stakeholder problem and the results of interactive interview sessions exposing practitioners to an integrated tradespace negotiation environment. Chapter 6 discusses the structure of a multi-stakeholder problem and features that may demand alternative strategies for MSTSE.
- Chapter 7 collects the insights of the previous chapters and synthesizes them into a set of prescriptive recommendations for preparing and conducting MSTSE, providing guidelines for potential adopters. This includes both a full negotiation form and an informal analysis form which a single stakeholder or team of engineers can use.
- The next chapters demonstrate the application of the insights of the research. Chapter 8 analyzes existing systems engineering methods for their compatibility with multi-stakeholder problems and the best use of MSTSE within them. Chapters 9 and 10 cover the case studies performed to demonstrate the application of the MSTSE formulation and associated visualizations to real-world examples of multi-stakeholder engineering systems: Satellite Radar and the Northeast Corridor.
- Finally, Chapters 11 and 12 feature discussion and conclusions of the research as a whole.

Figure 1-2 shows a visual outline of this research plan, with the relevant chapter numbers included.

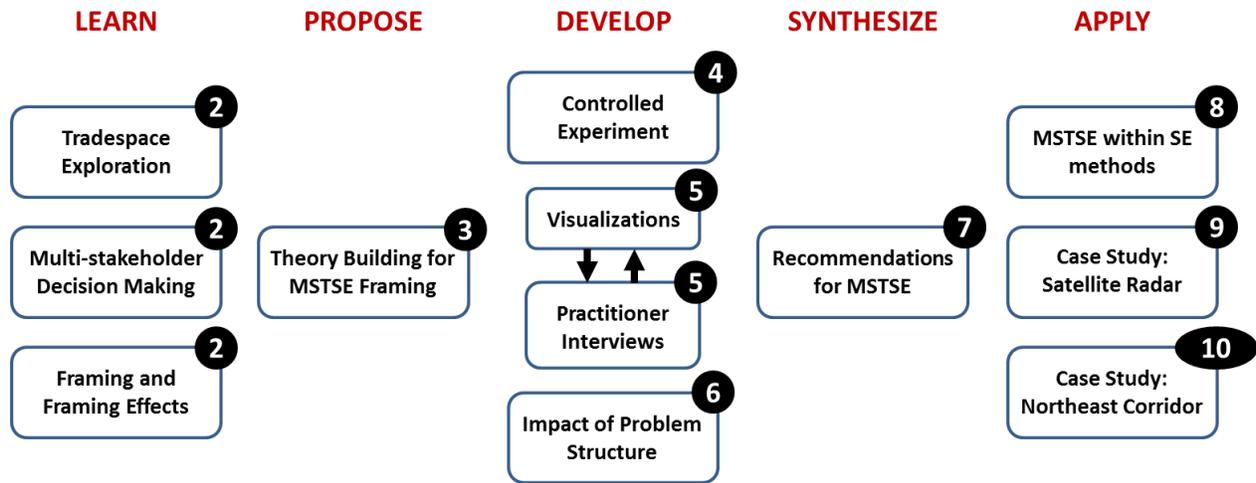


Figure 1-2: Research outline



## **2 Literature Review**

This research sits at the intersection of three main research areas: multi-stakeholder decision making (including the fields of both engineering and negotiation), tradespace exploration, and framing. The following subsections will provide an overview of the relevant literature in these areas. The final section is a short introduction to the –ilities, as they have the potential to provide additional means with which to create mutually beneficial solutions to multi-stakeholder engineering problems.

### ***2.1 Multi-Stakeholder Decision Making***

Decision problems with multiple goals have particular importance to engineers, as physical constraints and limitations often place the desired outcomes of projects in conflict. The nature of these conflicts can vary in practice. For example, conflicts may occur for one decision maker between competing interests or between multiple subsystem engineers tasked with representing their group’s needs. Conflict can also arise between completely distinct stakeholders (e.g. agencies, policymakers) with competing interests but a desire to work together and reap the benefits of combining resources (e.g. money, votes) to create a more capable system. Each of these types of conflict has been targeted by researchers seeking to create tools and processes to assist in the identification and selection of system designs that are acceptable to all parties and generate as much mutual value as possible. However, a variety of factors affect when techniques are applicable and when they are not, including what definition of “value” is appropriate and the availability of data to accurately portray preferences. The following subsections will briefly review some of the key research outcomes in this area and then discuss their applicability to the particular problem type this research is interested in solving.

#### **2.1.1 Single Stakeholder, Multiple Goals**

Before discussing multiple-stakeholder analysis, it can be useful to consider the related problem of a single stakeholder dealing with multiple dimensions of benefit. Analysts often desire to combine these different sources of value into a single metric. Techniques of this type would be described as “basic values” models by Fischhoff (1991) as opposed to “articulated values” models, indicating that they are defining value using a combination of basis functions, rather than allowing stakeholders to assess alternatives holistically and assign values to them directly. Benefits of this approach include scalability to large numbers of alternatives and traceability to sources or reasoning behind value assessments, leading to their popularity in engineering applications. Aggregation of different sources of value into a single value score allows the ranking of alternatives, typically with the intent to select the choice with the highest ranking.

The debate over what technique should be used to perform this aggregation has become somewhat ideological, depending considerably on the backgrounds of the debaters and the problem areas in which they are most familiar. This research has no intent to resolve these

trenchant debates; therefore an exhaustive accounting of multi-criteria decision making methods is not necessary to accomplish our goals. Unfortunately, some popular techniques (e.g. AHP, QFD, Pugh charts) stretch or defy the following categorization, but for narrative simplicity, we will divide this field of the literature into two normative theories: cost-benefit analysis (monetized) and utility theory (non-monetized).

Utility theory originated as a normative decision theory from Von Neumann and Morgenstern and has been developed into a variety of prescriptive techniques by many others, most notably Keeney and Raiffa with Value Focused Thinking (Von Neumann and Morgenstern, 1953; Keeney, 1992; Keeney and Raiffa, 1993). Utility, in its most rigorous form, encodes the value of an item as equivalent to the mathematical expectation of a lottery between other items when the stakeholder is indifferent between the first item and the lottery. This creates a continuous metric of value, typically scaled between zero and one, where zero is the worst alternative considered and one is the best. Other formulations establish zero as the minimally acceptable choice and one as a maximally satisfied condition, where the stakeholder is indifferent to additional performance. Multi-attribute utility theory (MAUT) combines the utilities of individual dimensions of value into a single utility term, with the prescriptive conclusion that stakeholders are likely to be indifferent between equally scored alternatives. Utility theory has proven useful for many engineering problems in which empirical data is difficult or impossible to obtain, but is often critiqued for the difficulty involved in eliciting preferences from stakeholders and with questions of the applicability of value statements based on lotteries to deterministic problems.

Cost-benefit analysis is also a normative theory that has been developed into a variety of different prescriptive techniques, perhaps most famously Real Options Analysis (Brach, 2003; Adner and Levinthal, 2004). These techniques are characterized by the attempt to assign monetary value to the positive and negative outcomes of decisions. These monetary terms are considered commensurate and thus able to be aggregated into a single value metric, occasionally also called “utility” if the accounting procedure follows utility theory axioms (Hazelrigg, 1998). For many applications, these are discounted if they occur in the future and rolled up into a present value, reflecting the opportunity costs of investment, using the techniques of Discounted Cash Flow (DCF) and Net Present Value (NPV). Stakeholder preferences are accounted for by analyzing their willingness-to-pay for different alternatives. The preferred means of performing these calculations is empirically: observing patterns of behavior and purchases to establish stakeholder’s revealed preferences. When this is not possible, interviews are performed, asking what tradeoffs between alternatives stakeholders would be willing to take (Viscusi et al., 2005). Cost-benefit analysis has seen the most success in fields such as product design, where empirical data on the behavior of large groups of customers is obtainable. However, it faces serious criticism when applied in areas where the monetization of value is deemed inappropriate or discount rates are undefined (e.g. scientific progress), or where empirical data is unavailable (e.g. large projects with no close comparisons and only one target buyer). Cost discounting has also

been accused of systematically deferring irreversible costs on disenfranchised future people (Ackerman and Heinzerling, 2002).

### 2.1.2 Connected Stakeholders and Aggregation Methods

Many large engineering design efforts feature tightly connected stakeholders. These stakeholders are responsible for advocating for their own interests but are ultimately subject to the will of the entire group and lack the ability to withdraw from the design process. As an example, consider a project organized into subsystem design teams: each team must satisfy their own needs, but they all *must* come to an agreement because they work for the same company/organization. The most common means of addressing multi-stakeholder decisions of this type in the systems engineering field is by simply aggregating the *requirements* of each stakeholder – the first “technical process” covered by the INCOSE Systems Engineering Handbook (2012) is one for Stakeholder Requirements Definition, which is then followed by Requirements Analysis. Requirements represent “characteristics or constraints” on the resulting system, and are often divided into two types: *threshold* requirements (minimum acceptable) and *objective* requirements (desired).

Structuring a design process around requirements is often called “requirements engineering” and has various procedures and analysis techniques (van Lamsweerde, 2009). Notionally, aggregating requirements between stakeholders allows the resulting system to be designed as if there were a single master stakeholder – one with a very long list of needs. However, requirements engineering also has drawbacks. For example, Axiomatic Design (Suh, 1998), a branch of requirements engineering that maps functional requirements into design parameters and process variables, recommends “zigzagging” between the different levels in order to decompose the requirements and architect the system such that “information content” (which can be understood to be roughly similar to complexity) is minimized. This direct allocation of function to specific elements of form can lead to elegant solutions but prevents effective exploration of designs that create value through emergent synergies between different elements of their form. Other limitations of requirements engineering that have challenged systems engineers include:

- Limited ability to encode value beyond binary/trinary interpretations of meeting the threshold and objective requirements
- Requirements may drive “lock-in” on specific areas of the solution space too early in the design process
- Arbitrarily or heuristically set requirements (or hierarchies for the requirements) can cause downstream disruption in the design process when they must be changed

And specifically for aggregating multiple stakeholders’ requirements:

- Unequal number of requirements for different stakeholders indirectly (and usually unintentionally) weights the resources dedicated to them in detailed design, with possible negative consequences
- Aggregation limits the ability to examine tradeoffs between different needs, as they become “lumped together”

Utility theory is often put forth as a strong alternative to requirements analysis given its ability to address many of the above challenges and the natural correspondence between threshold/objective requirements and the minimum/maximum value endpoints of a single-attribute utility function. Many utility-oriented researchers have approached the multi-stakeholder problem with the goal of some sort of aggregate utility maximization between stakeholders (Kusiak and Wang, 1994; Bahler et al., 1995; Scott and Antonsson, 1996). This research thread has extended to the point of complete automation: using algorithms to generate “optimal” or Pareto-efficient solutions, or applying the principles of game theory to simulate negotiation (Chen et al., 2004; Gatti and Amigoni, 2005; Romanhuki et al., 2008). Recent work has adapted dynamic programming to directed systems of systems; this allows individual stakeholders to optimize their own non-dimensional value but relies on the existence of a “centralized” stakeholder with the authority to set “transfer prices” that impose exchange rates between stakeholders, constraining them to act in a way that optimizes an aggregate capability (Fang and DeLaurentis, 2015).

Though conceptually clear and intuitively appealing, aggregating the utility functions of different stakeholders ignores the fact that there is no universal, ratio scale of utility that can be compared across individuals. Aggregations of multiple people’s utilities sacrifice the normative benefits of utility theory, so the results of such techniques should be taken with caution and treated less as “optimal” designs and more as “interesting” designs.

The key assumption that multi-stakeholder aggregation relies on is that the different stakeholders are closely tied, characterized by extensive information exchange, iterated communication, and shared “high level” objectives (such as those of their employer). This assumption allows designers to skirt the conclusions of Arrow’s Theorem (Arrow, 1963). Roughly worded, the theorem states that there is no system for combining multiple rank-orderings into an aggregate ordering without violating at least one commonly accepted axiom of rationality on group decisions, including transitivity and non-dictatorship, among others. In this context: there is no way to aggregate multiple stakeholder utility functions without having the resulting function display some irrational choice behaviors when compared to the group of functions individually. However, tightly connected stakeholders can arguably avoid this conclusion by virtue of not needing to obey all of those axioms, as their individual interests are ultimately secondary to the group as a whole (Scott and Antonsson, 2000).

When using monetary (or other assumed-additive) value functions, the challenges for aggregating stakeholder value do not disappear. Perhaps the most famous result of mechanism design<sup>1</sup> is the VCG mechanism, which maximizes the social good of a decision between multiple stakeholders by making each stakeholder pay the externality they impose on the other stakeholders (Vickrey, 1961; Clarke, 1971; Groves, 1973). The VCG mechanism's most well-known form is the second-price auction: the auction system is a means of aggregating information about each stakeholder's preferences and outputting a single solution. The second-price auction is incentive-compatible: there is no incentive for a stakeholder to lie about their valuation of the item being auctioned. However, the optimality of the final result *does* depend on the accuracy of the reported valuations. Given the difficulty inherent in estimating value for complex systems, stakeholders often must iteratively update their value statements as new information becomes available. This can be viewed as an unreliable oracle, which reports a value subject to some unknown noise. Unfortunately, regardless of incentive-compatibility, achieving a tight approximation of optimality requires an exponential number of queries to an unreliable oracle, limiting the practicality of such a method (Hassidim and Singer, 2015). This result limits the applicability of voting or bidding mechanisms for multi-stakeholder decision making on complex systems.

The field of collaborative engineering would refer to this high degree of interaction between stakeholders as either *collaborative* or *cooperative* (as opposed to the less-connected *coordinated*), and indeed collaborative negotiation has received significant attention in recent years (Lu et al., 2007; Shahan and Seepersad, 2009). Table 2-1 highlights the key differences between these categories of human interactions.

**Table 2-1: Human Endeavor Characteristics (Lu et al., 2007)**

	Stakeholder	Resource	Goal	Task Structure
<b>Coordination</b>	Large Community	Limited and Exchanged	Multiple & Competing	Pre-defined, same layer in hierarchy, uni-direction
<b>Cooperation</b>	Mid-size Group	Limited and Shared	Multiple & Private	Pre-defined, across layers in hierarchy, bi-direction
<b>Collaboration</b>	Small Team	Limited, Shared, Complementary	Single & Common	Undefined, non-hierarchical, multi-direction

The most striking difference between these three types of behavior lies in the nature of the goals of the participating stakeholders. Collaboration requires a shared goal, while

---

<sup>1</sup> Mechanism design is also known as the reverse-game-theory problem, in which the game (“mechanism”) is designed to achieve a certain outcome from the players. Decision making rules such as voting or auctioning are common topics in this area.

cooperation entails different goals, and coordination has (potentially competing) hidden goals. Collaborative engineering focuses (naturally) on collaboration, with research threads extending into detailed design with collaborative design software and information sharing. However, collaboration is not a *superior* form of partnership, as one might assume due to its higher degree of interaction and alignment of goals. Rather, it is an artifact of the nature of the project and the relationship of the participating stakeholders. Many stakeholders do not engage in collaboration, but rather work together with much more limited information exchange and distinct preferences. This shift towards coordination is particularly apparent when stakeholders are *not* tightly connected but rather independent, as we will discuss now.

### **2.1.3 Independent, Cooperative Decision Makers**

The literature discussed the previous section was targeted at tightly connected stakeholders with similar overarching goals (and typically shared oversight), but engineering negotiations also can take place between independent but coordinating *decision makers*: stakeholders with a measure of control over the ultimate design selection and with the ability to withdraw from the decision making process if they so desire. These types of decisions are more complicated than those between connected stakeholders, as attention to interpersonal conflict and resolution is required to ensure completion. Using Table 2-1, this type of endeavor would be seen as either coordination or cooperation but not collaboration, as each decision maker is responsible for supporting their own goals. As such, the aggregation of their preferences will likely lead to unacceptable results according to Arrow's Theorem, due to the lack of shared high-level objectives. Additionally, automated or algorithmic negotiations are typically not appropriate for these problems, as decision makers are less likely to abdicate responsibility to a "black box" in charge of creating and dividing value between the parties. Because of this, analysis of a decision problem between multiple independent decision makers must avoid aggregation and keep their interests separate. This problem type is the main area of interest for this research, so for the remainder of this thesis the word "stakeholder" may be assumed to mean this type of independent decision maker capable of either supporting or rejecting any possible design, unless stated otherwise.

Not all independent stakeholders are equally influential in a decision: it may be that only a subset of the stakeholders affected by a system is allowed to vote on the action to take or have veto power over unacceptable solutions. Stakeholder analysis is the practice of identifying the main players in a given decision problem, particularly for the purpose of feeding into other analyses and guiding "participatory, consensus-building" activities such as negotiation (Schmeer, 1999). The stakeholder saliency framework is often used to assess the position of each stakeholder in a given decision by categorizing them according to their power, legitimacy, and urgency (Mitchell et al., 1997). Stakeholders who possess more of these characteristics are deemed more salient and thus more influential on the ultimate decision. Stakeholders with all three characteristics are considered "definitive" and possess enough leverage to guarantee their satisfaction, providing *de facto* veto power. Another more quantitative method of assessing the

relative importance of stakeholders is with a stakeholder value network (Cameron, 2007; Feng et al., 2012; Feng et al., 2010). This method structures the problem with “value flows” weighted by importance and urgency, which transmit value between stakeholders. From there, the most critical value cycles can be identified and the stakeholder set under consideration can be scoped to the desired level of detail by removing some number of less-impactful stakeholders.

Utility theory has been applied to resolve conflicts between independent stakeholders, modeling preferences and offering up potential compromises one at a time using a combination of joint maximization and equality to address interpersonal friction (Sycara, 1988). This application of utility theory is prescriptive, demonstrating effectiveness at persuading stakeholders to accept compromises. Due to the more open-ended nature of these problems, the explicit inclusions of insights from conflict resolution and negotiation theory (which are discussed in the next section) are most prevalent in research targeting independent stakeholders, layering them into process-oriented systems engineering frameworks (Mostashari, 2005). Utility theory has also been used to show that strictly self-interested behavior (“hill-climbing”) can lead to inefficient solutions compared to those achievable by group-focused strategy (Klein et al., 2003).

For applications to profit-driven business, variants of cost-benefit analysis have been employed to model interactions between decision makers as well. Theory W and Win-Win negotiation were combined into the Value Based Theory of Systems Engineering (VBTSE) to describe the relationship between developers, customers, and users (referred to as “success critical stakeholders”) in a software engineering project, each with their own “win” conditions that must be satisfied for the system to succeed (Horowitz et al., 1999; Boehm and Jain, 2007). VBTSE also draws from some insights of utility theory and decision theory, resulting in some similarities to the work of Mostashari in the setup, but informs much of the actual negotiation using cost-benefit analysis to evaluate time streams of cash flow. These financial streams of different business plans and customer willingness-to-pay are presented as the main tools for evaluation and decision making, limiting the applicability of VBTSE to problems in which the basic assumptions of cost-benefit analysis are acceptable.

#### **2.1.4 Negotiation Analysis and Practice**

The field of negotiation has a rich literature behind it that remains largely untapped by engineering research; this is undoubtedly related to the engineering literature’s focus on tightly connected, collaborative stakeholders. However, as this research moves into the realm of independent, cooperative or coordinated decision makers, the insights of the negotiation literature are quite relevant due to the increased importance of social choice and reduced ability to aggregate stakeholders and/or “solve” the decision problem. Bazerman et al. (2000) provide an excellent recap of the history of negotiation analysis and the rise of behavioral decision theory in the field, while making a compelling argument that psychological characteristics have a significant impact on negotiations. Readers are referred to their work for a complete perspective

on the history of negotiation and the individual studies that support it, but here we will call out their collection of the main experimental results of the application of decision theory and bounded rationality to negotiation:

- Concessions are more palatable when framed positively rather than negatively
- Negotiators are often inappropriately affected by anchoring and available information
- Overconfidence is common among negotiators
- Mutually beneficial tradeoffs are often missed and negotiators assume that their interests are incompatible with their counterparts' interests
- Common negative behaviors include escalating conflict, ignoring alternative perspectives, and reactively devaluing offers and concessions of the opponent

The first two bullets will be addressed in further detail in the framing section of this review. Much of that research was galvanized by the work of Raiffa (1982), who followed up twenty years later with an excellent guide to the analysis of negotiation conflicts (Raiffa, 2002). In particular, Raiffa extensively discusses the differences in approach between *distributive* problems, in which a fixed amount of 'value' must be split between the participants, and *integrative* problems, in which mutual benefit can be achieved due to differing preferences. Techniques for creating and dividing value are discussed, with the caveat that different conceptions of fairness exist and must be agreed upon by the participants in order to be used without further negotiation.

Raiffa also subscribes to the principle of "Full, Open, and Truthful Exchange" (FOTE), which asks negotiators to communicate with honesty and refrain from strategic deception or omission. FOTE is an ideal – though perhaps rarely achievable in full, any amount of information sharing about personal interests has been shown to increase the effectiveness with which negotiators discuss and identify opportunities for mutual gains when solving problems with integrative issues. Strategic deception is consistently superior to FOTE bargaining only for purely distributive issues, frequently referred to as "zero-sum games". Complex problems, including those in the engineering domain, are typically mostly integrative and are very rarely zero-sum due to the diverse sets of goals between stakeholders; thus FOTE bargaining has the potential to improve the negotiation outcome in our domain of interest. FOTE can be a "tough sell" to disputing parties who perceive (accurately or not) that their decision is zero-sum, but is a desirable ideal that can be supported – again, at least in part – by stakeholders with existing positive relationships or productive partnerships.

The work of Fisher and Ury (1991) has become a touchstone for the practice of what they refer to as "principled" negotiation, which builds heavily on the insights of behavioral analysis. With a central commitment to negotiation as a means of reaching consensus, they lay out the principled negotiation as the best way to get there while encouraging the necessary goodwill and confidence necessary for an agreement to last in the long term. One of the main features of

principled negotiation is an emphasis on *interests* over *positions*, a key step towards reducing the impulse to engage in *positional bargaining*: the sequential, back-and-forth, offer-counteroffer negotiation strategy that has been shown to result in high tension and low satisfaction due to its adversarial nature. These recommendations are echoed in the work of Brett (1991), with additional discussion of alternative negotiation tactics depending on whether the ultimate goal is mutual, coalitional, or individual gain. However, she notes that beginning with a mutual value approach is typically the most effective negotiation tactic, using individual goals only to counteract the urge to accept a low-quality but agreeable solution.

Building the positive relationships necessary for mutual principled negotiation can be assisted through the use of multi-stakeholder dialogues (MSDs): meetings of stakeholders and interested parties to discuss issues at hand between them (Susskind et al., 2003). Negotiation, as we have been using the word, could perhaps be viewed as a subset of MSD, which includes efforts to build relationships, share information, set agendas, brainstorm solutions, and build consensus. This suggests that bringing stakeholders together before attempts at direct negotiation or consensus-building can be a productive lead-in. This is particularly important in multiparty negotiations – negotiation between more than two participants – in which there are more inter-stakeholder relationships to manage (Crump and Glendon, 2003). Within Crump and Glendon’s breakdown of the major areas of multiparty negotiation, this research is most aligned with that of organizational negotiation (as opposed to international negotiation or public disputes), which emphasizes positive strategic alliances more than distributive value negotiations.

The use of Joint Fact Finding (JFF) has also represented one of the key prescriptive recommendations of the negotiation literature. JFF is a cooperative effort between multiple stakeholders to establish objective facts upon which to base a negotiation, and can be categorized as an MSD with an interest in information sharing and brainstorming (Ehrmann and Stinson, 1999; Ozawa, 1991; Lee, 1994). JFF serves to preemptively defuse potential conflicts arising over disagreements regarding science or forecasting, a common breakdown point for negotiations dealing with uncertain future scenarios, while also often improving stakeholders’ understanding of the both scientific and nonscientific elements of the decision (Karl et al., 2007). JFF has become a critical component of many successful negotiations in different fields, with roots in public policy. Indeed, JFF is a key recommendation of the Water Diplomacy Framework of Islam and Susskind (2013), which seeks to assist the creation and management of the complex, sociotechnical issue of water networks, which share many similarities with engineering systems. In this context JFF demonstrates its value in helping to create a productive environment for negotiation and enable not only integrative thinking but also the insertion of – ilities as a means of managing different expectations of uncertain futures.

JFF has a straightforward connection to the task of modeling the system, as the objective consensus can be used to help guide the creation of a trusted model. Processes that consist of

stakeholders participating in the construction of a model are sometimes collectively referred to as Collaborative Modeling, consisting of a variety of specific methods including Group Model Building, Mediated Modeling, and Computer-Aided Negotiation (Langsdale et al., 2013). In line with JFF, one of the key conclusions derived from this research is that keeping the model transparent and accessible is capable of building trust amongst the participants. Note that the process of Collaborative Modeling begins with “framing the decision”, which acknowledges the importance attached to reaching agreement on what exactly the problem being solved *is* (Antunes et al., 2006). The “best practice” for the model itself is to simulate alternatives in support of the decision process. The work of Czaika (2015) has expanded on this role of models as boundary objects *inside* the negotiation by experimentally verifying the impact of model use on negotiations. In particular, advantageous changes to solution quality and stakeholder satisfaction were observed when: (1) using a model (over not using a model), (2) having stakeholders create a model (over being given a model), and (3) using the model to test alternatives (over simply verifying agreements). These experiments also demonstrated (via self-reporting) that models encourage stakeholders to consider *other* stakeholders’ perspectives when making decisions without unduly influencing their own interests. A majority of negotiators changed their own stated priorities regardless of model usage, illustrating that the challenge of accurately reporting value is present in a negotiation setting in addition to more traditional engineering tasks.

## **2.2 Tradespace Exploration**

Tradespace exploration (TSE) is a modern engineering technique that explores a design space by enumerating and evaluating a large number of potential designs, including apparently sub-optimal designs (however “optimality” may be defined), with the understanding that certain valuable behaviors may not be captured by a given value metric (Ross and Hastings, 2005). The “tradespace” itself is the set of designs considered during this process, and is typically viewed with a two-dimensional scatterplot of benefits and costs, which offers a concise visualization of the two main decision-driving features of the design process. As a paradigm for solving complex design problems, the majority of research in tradespace exploration has focused on the analysis of the space of alternatives with the goal of uncovering design choices that are optimal or near-optimal (Daskilewicz and German, 2009; Ross et al., 2004). The central benefits of TSE over alternative methods include an improved grasp of performance tradeoffs and relationships between variables in the design space. These benefits are most relevant to conceptual, “top-down” design problems between independent stakeholders, rather than to the “bottoms-up” specificity of problems between tightly connected stakeholders.

As an offshoot of the “design by shopping” paradigm, TSE has been found to be particularly useful for the design of complex engineering systems with multiple dimensions of benefit (Balling, 1999). These systems are difficult to optimize and rarely intuitive, resulting in a system where stakeholders may not know what they want until seeing their potential choices. In Multi-Attribute Tradespace Exploration (MATE), the often chosen benefit metric is a multi-

attribute utility function, created as a combination of different performance attributes which are rated from zero, defined as minimally acceptable, to one, at which point no benefit is gained from performing better (Ross et al., 2004). The different attributes represent the different ways in which a system delivers benefit to the stakeholder. For example, a car might have attributes for top speed, acceleration, and turning radius; each attribute delivers some utility which can be combined into an overall utility for a design. MATE has been effectively used to investigate the preferences of stakeholders and to find attractive designs in line with, or superior to, those found using traditional point-design engineering techniques, without the cost of algorithmic optimization (Ross et al., 2010b).

Performing tradespace exploration with utility theory rather than cost-benefit analysis is typically beneficial for analyzing complex sociotechnical systems. Cost-benefit analysis is an extremely powerful tool for finding optimal designs but depends heavily on the reliability of the data used to model both cost and benefit. Cost modeling for complex systems is challenging and rarely precise enough to ensure confidence in a single “optimum” solution out of potentially many thousands of possible solutions in a tradespace. Complex sociotechnical systems can also be far removed from typical assumptions of economic analysis, potentially featuring only one relevant buyer or seller (e.g. the government), the influence of mutual “you scratch my back, I scratch yours” types of relationships, and value propositions that are not always easily converted into monetary terms (e.g. –ilities). Though the analysis of some complex systems will be amenable to the adoption of cost-benefit assumptions, analysts under conditions like the above must question whether or not willingness-to-pay is possible to extract from stakeholders, or if the results of such an effort would be representative (Ross et al., 2010c). Additionally, negotiations are often accompanied with changes in stakeholder preferences as new information is revealed, which would necessitate frequent reassessment of willingness-to-pay (Curhan et al., 2004; Czaika, 2015). Because of the emphasis on exploration over optimization and the desire to foster situational awareness for participating stakeholders, TSE is benefitted from using utility theory over cost-benefit analysis. The ability to break off and consider different dimensions of multi-attribute utility also gives the potential to account for modified preferences on the fly during a negotiation.

TSE is most commonly contrasted against optimization and Multidisciplinary Design Optimization (MDO) due to its emphasis on investigating the complete span of the design space rather than only optimal and near optimal solution (Ross and Hastings, 2005). Value-Driven Design (VDD), a growing strategy for formulating systems engineering decisions on the basis of improved objective functions, was developed with the intention of incorporating MDO as a key decision support tool (Collopy et al., 2012). However, the large number of requirements that must be imposed on a stakeholder to create a value function capable of being optimized has been identified as a significant barrier to its application on complex systems (Mesmer et al., 2013). The initial ambiguity of these requirements, and the frequency with which they change later in the design process, necessitate multi-level optimization schemes that iteratively redefine

requirements and can be difficult to converge. The ability for TSE to incorporate a wide variety of objectives – and to change requirements during exploration by allowing for interactive tuning of utility functions – is a potentially valuable simplification of this process in addition to the knowledge gained by considering a larger breadth of designs. This is also in line with the findings of Nutt (2008), which show that following “discovery” and intelligence-gathering type activities lead to superior decision making than “idea imposition,” which is characterized by focus on the form of few, specific solutions.

### **2.2.1 Multi-Stakeholder Tradespace Exploration**

Tradespace exploration is an attractive technique, at least on the surface, for problems in which there are multiple decision makers or stakeholders with different value propositions or preferences defining their utility functions. This is because, while a point-design or independent optimization study would result in a single preferred design for each stakeholder, the tradespace allows each stakeholder to find desirable designs and compare interests with each other across a common set. The presence of designs that are evaluated for all stakeholders also assists in the finding of mutually satisfactory choices: those designs that perform well for each set of preferences but may not be the optimal point choice for any one. However, given the relative youth of TSE as a design paradigm, there has been relatively little research explicitly devoted to the multi-stakeholder extension.

Previous methods for negotiation between stakeholders using a TSE study have largely relied on the best practices for *individual* TSE, with all stakeholders using those methods in parallel. For example, careful consideration of the tradespace may reveal that Stakeholder A has dramatically fewer acceptable designs because his requirement on a single attribute is very restrictive. Upon relaxing this requirement, new mutually attractive options between Stakeholder A and the other stakeholders may appear (Ross et al., 2010a). This practice was conducted on a case-by-case basis, but with little regard for the fundamental differences between single- and multi-stakeholder problems. An attempt to formalize the steps involved into a coherent process has the potential to unlock more consistent success in finding solutions that are desirable to all participating stakeholders.

Note that one commonly suggested approach for finding compromises between stakeholders using multi-attribute utility is to simply combine their utility functions into a single, aggregate decision maker utility, as discussed in section 2.1.2. This recommendation has even been given specifically for MSTSE (Garber et al., 2015); however, the same caution with respect to Arrow’s Theorem is merited (Arrow, 1963). Without a clear hierarchical relationship amongst the stakeholders, such as those that might be found among tightly connected subsystem stakeholders, this result provides a compelling argument in favor of avoiding attempts at achieving compromise by combining their various preferences in a super-preference set (Scott and Antonsson, 2000). This is consistent with the classification of multi-criteria decision making

methods as “attention directing tools” whose assumptions and limitations should be regularly questioned rather than unswerving problem solvers (Roman et al., 2004).

MSTSE’s interactive and exploratory nature is geared towards a paradigm of directing stakeholders to investigate interesting designs and understand their underlying preferences. When explored by multiple stakeholders together, the tradespace becomes a “boundary object” used to support not only data-to-stakeholder communication but also inter-stakeholder communication (Carlile, 2002). Specifically, Carlile states that effective boundary objects both establish “shared syntax” between stakeholders and “a concrete means for individuals to specify and learn about their differences and dependencies.” This is a challenge for complex engineering systems, which typically span many fields of expertise, and is a reason behind the focus on *value* (a shared concept) rather than *performance* (often domain-centric) in the TSE literature. It has been demonstrated that effective design is possible with a large degree of specialist knowledge and that, in fact, specialist and trans-specialist knowledge are substitute goods in many conditions (Postrel, 2002). The potential role of MSTSE as a communication enabler between different specialists can enable this type of design for complex systems without the need for a stakeholder with knowledge of the “complete” system – a position that often doesn’t exist.

Golkar and Crawley (2014) describe an example problem in which two “stakeholders” – scientists and engineers – must design a space mission together, but the value statement for each group is unclear. They deploy a variant of the Delphi method, a consensus-building technique that iteratively elicits preferences from a group of experts (Dalkey, 1969; Linstone and Turoff, 1975), to structure the development of a utility function for each group, reaching convergence to a true, or perhaps at least accepted, group value model after a few iterations. The Delphi method successfully resolves the *internal* debate for each stakeholder in this problem, but does not address the subsequent design problem between the scientists and engineers. This illustrates the jump in complexity from negotiating between tightly connected stakeholders (within-group) and independent stakeholders (between groups). Golkar and Crawley use tradespace exploration of the two utility functions to find the best design alternatives for both groups in the form of a set of Pareto efficient solutions, but they do not attempt to determine which alternative should be selected either analytically or by consulting the experts used in the Delphi method. Their transition from the Delphi method to TSE follows a similar logic to this research, but makes no effort to expand or improve TSE in response to the specific challenges of the multi-stakeholder problem.

The work of Dwyer uses tradespace exploration to assess the impact of sharing authority with other stakeholders, specifically in the form of the additional cost needed to field a system capable of satisfying multiple parties (Dwyer et al., 2014; Dwyer and Szajnfarter, 2014). Shared authority, or jointness, is contrasted against asymmetrical roles such as acquisition agents in an analogy similar to the empowered decision maker vs. connected stakeholders distinction

discussed earlier in this chapter. The proximity of the individual Pareto fronts of two stakeholders on a utility vs. utility scatterplot is used as an indicator of the synergy of their missions, with the conclusion that sharing authority with non-synergistic stakeholders is only possible through large expense. This carries increased risk of failure due to lack of intermediate-value solutions if funding is later reduced. Analysis of this type is most appropriate for a single stakeholder looking to gain insight into a potential multi-stakeholder opportunity, in contrast to the shared, interactive nature of most other multi-stakeholder tasks discussed here. As such, this method is more suited to strategic planning than to attempts at FOTE negotiation.

## **2.3 Framing**

The concept of *framing* has been used in many different ways, to describe many different ideas. In their most basic sense, all the uses of framing share one key feature: the understanding that contextual factors impact human perception and thus human action. The wide scope of framing can sometimes lead to confusion when discussing its implications. Here, we will divide the relevant literature into categories we refer to as “macro” or “micro” framing issues. The following subsections provide an introduction to these kinds of framing and examples of research into their impact on decision making.

### **2.3.1 “Macro” Framing**

“Macro” framing is related to issues of writ-large beliefs and perspectives. Perhaps the most famous science-oriented discussion of framing is that of Kuhn (1962) on the subject of scientific revolutions. Kuhn describes the progress of science as one of prevailing paradigms that are upset by revolutions in favor of new paradigms. Revolutions are often characterized by heated debate between the scholars of the different paradigms, who frequently have difficulty communicating because they are figuratively speaking different languages. The paradigms can be viewed as frames (or perhaps lenses in this analogy) through which people ‘see’ an issue. Competing paradigms can make normative arguments in completely different directions, as the norms to which they appeal do not necessarily align. This can affect negotiations even at a mechanical level. For example, there is evidence that *differences* in outcome goal orientation and process goal orientation, two types of mental framing positively correlated with high-value negotiation results, can negatively impact the quality of negotiation outcomes. When measuring each type of goal orientation present in negotiations, similar levels of both key types of goal orientation resulted in better negotiation outcomes than simply having high absolute levels of goal orientation (Katz-Navon and Goldschmidt, 2009).

Macro framing can also be applied directly to the nature of a decision making activity. Nutt (1999) found that framing decisions generally as searches for “wanted” properties of a solution rather than as “problems” to be eliminated resulted in higher chances of success. His recommendations for improving decision making include focusing on objectives, testing many options, and working with many people: a short list of macro framing topics directly relevant to MSTSE. Nutt also refers to the process of moving from consideration of specific solutions to

more general analysis as “asking reframing questions.” Interestingly, one of Nutt’s examples is to contrast a request to “lower costs” as superior to asking for answers to the “root cause of the cost problem.” With the rise of the Better Buying Power initiative, lowering costs is a predominant macro frame in modern defense acquisition, a multi-stakeholder engineering problem of interest to this research, (Kendall, 2011, 2014).

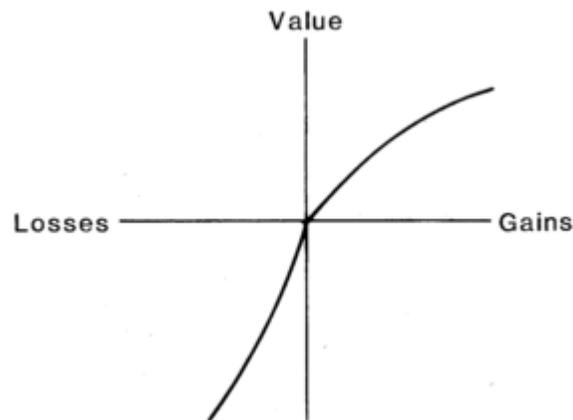
Schon and Rein (1994) also approach the issue of interpersonal conflict through framing, specifically targeting the realm of policy creation. They stress the importance of “frame reflection”: deliberately considering the differing frames of each actor as a preliminary step to effective policy design. Moreover, each actor can balance multiple frames, both rhetorical and action-oriented, that operate on different levels. At the highest level, “metacultural frames” are heuristic frames with highly engrained societal norms. They provide an example of the use of metaphors like “sickness versus health” to justify actions such as urban renewal, which may be in direct conflict with a frame of “family and culture” on the same issue. These metacultural frames influence lower level “institutional” frames (general norms and actions of an institution) and down to “policy” frames (the framing of a particular issue). Though couched in the language of policy, due largely to the prominent role that ideology plays in political debate and policy creation, Schon and Rein’s work can be applied to any field with multi-party conflict over issues more fundamental than objective fact. Frame reflection allows participants in the conflict to examine not only where their own beliefs come from but also those of their counterparts. Though Schon and Rein rightly acknowledge the risk of relativist paralysis (e.g. questioning the objective validity of norms can lead to failure to act), they provide many examples of frame reflection by key actors resolving entrenched conflicts by clarifying the decision criteria for each party to the other. Recently, attempts to deliberately teach and incorporate reflection into the systems engineering practice have shown promising early returns (Muller, 2015; based on Schon, 1983). Reflection was self-reported as useful for both learning and implementation, but communication barriers with colleagues without systems engineering training was thought to limit its overall impact.

More generally, macro framing is often a subset of personal philosophy. Of particular interest to this research is the issue of fairness or equality, as it has considerable bearing on the evaluation of outcomes in group problem solving. Raiffa (2002) points out that there are many credible definitions of fairness, which can have dramatic impacts on the “fairest” solution for a given problem. In order to prevent self-interested “gaming” of the system, he recommends that participants in a negotiation agree in advance on an objective criterion of fairness. This is similar to the concept of the “veil of ignorance” central to Rawls’ *Theory of Justice* (1971) because, without knowledge of how it affects one’s own well-being, a person will likely choose what they truly believe to be fair. An alternative, more pragmatic, view of the same idea lies in the game theoretic heuristic that the best way to avoid “gaming” is to make the game too complex for players to discern what will improve their outcome. Macro framing can also be an

influencing factor in the creation of metapreferences for decision makers, such as a favoring of passively robust systems over actively changeable systems.

### 2.3.2 “Micro” Framing

“Micro” framing concerns the impact that problem-specific display of information can have on decision making. The most prominent results in this field include bounded rationality (Simon, 1957) and Prospect Theory (Kahneman and Tversky, 2000), which are descriptive decision theories that delineate how humans may attempt to act rationally but do not succeed. Bounded rationality refers to the inability of humans to accurately analyze complex problems and find optimal solutions, instead relying on heuristics to reduce the cost of deliberation. Prospect Theory is an empirically derived theory describing the nature of many common deviations from rationality. It states that people make decisions by comparing outcomes to a specified *reference point*. Outcomes are judged as differences from the reference point and are therefore perceived as “gains” or “losses.” Reference points are created from available information and are reinforced by anchoring, the observed bias that humans display towards information they are shown first, regardless of its ultimate relevance. Changing a reference point, once established, usually requires a deliberate effort (Tversky and Kahneman, 1974). Perceived value around the reference point is asymmetric, resulting in a higher impact of losses over gains as pictured in Figure 2-1. It has also been found that decision making in the losses domain is more stressful and more likely to lead to irrational or regretted behavior (Gelfand et al., 2004; Gonzalez et al., 2005).



**Figure 2-1: Asymmetric perceived value around a reference point, according to Prospect Theory (Tversky and Kahneman, 1981)**

Another common bias covered in Kahneman and Tversky’s work is the availability bias, which describes the human bias towards information that is accessible (Tversky and Kahneman, 1974). In this way, information that is provided or readily recalled is implicitly assumed to be more important than hidden or forgotten information. Other biases they cover include

insensitivity to probability, misconception of chance, and improper grasps of regression and representativeness.

The phrase *framing effect* is often used in the context of micro framing to denote an observable change in behavior derived only from changes in framing, usually with regards to whether the outcome is characterized as a gain or a loss. For example, switches between positive and negative (gains / losses) phrasing have resulted in dramatic changes in decisions, with people tending to strongly avoid losses over seeking gains (for examples and analysis, see Tversky and Kahneman, 1981; Levin et al., 1998). Prospect Theory also incorporates the observation that decision makers often distort probabilities – overestimating the likelihood of rare occurrences and underestimating more frequent occurrences – leading to high sensitivity around zero-probability and one-probability captured by a curve known as the  $\pi$  function. This bias for certainty has led people to be largely characterized as risk averse for gains, preferring a certain gain to a higher expectation uncertain gain, and risk seeking for losses, preferring a chance at no loss to a guaranteed loss.

Other topics in micro framing include the effects of detailed deliberation and expert opinion. Some research has suggested that extensive consideration of preferences can lead to behavior that deviates from expert opinion and leads to decreased satisfaction in decision outcomes (Wilson and Schooler, 1991). Excessive time spent developing a numerical value model, often without seeing the impacts immediately, effectively codifies the *estimation* as *truth* when more satisfaction would be gained by allowing future changes in response to emergent insight. This is a strong argument against overtaxing decision makers, and relates to the theory that the expertise of ‘experts’ is in fact dependent on a stable frame for them to leverage (Shanteau, 1992).

Finally, the concept of two-path information processing, a theory originally developed in the 1980s by the Elaboration Likelihood Model (Petty and Cacioppo, 1986) and the Heuristic-Systematic model (Chaiken et al., 1989) and recently popularized by Kahneman (2011), outlines two main ways in which humans perceive information and make decisions: heuristically and systematically (in ELM parlance, peripherally and centrally). Heuristic thinking is fast, developed over time and through intuition, allowing people to rapidly assimilate new information that they can fit into an existing mental frame. Systematic thinking is the more in-depth, analytical thought that promotes new learning but requires more effort on the part of the decision maker. The framing of a problem has an impact on which path a decision maker uses, depending largely on how familiar the situation is to them.

### **2.3.3 Visual Analytics**

Here we will briefly address the topic of visual analytics as it relates to framing. Humans are highly visual creatures, and the field of visual analytics has recently developed around the topic of assisting human-in-the-loop computerized analysis for both learning and decision

making (Chang et al., 2010). Visual analytics has been applied in the domain of engineering design and modeling, to assist designers when approaching complex systems with large amounts of multi-dimensional data (Mavris et al., 2010). The insights of visual analytics have potential value for this research, as MSTSE is a data-intensive technique and the exploration paradigm is driven largely by visualizations of that data that in turn create particular framings of the information. Incorporating the insights of visual analytics also has the potential to increase stakeholder buy-in to MSTSE, by better illustrating the achievable benefits of the technique and lessening the occasionally “black box” nature of complex design in a tradespace (Lotov et al., 2004).

The main concepts this research will consider are those of the four-part “visual analytics paradigm” of Keim et al. (2008): analyze first, show the important, zoom/filter and analyze further, and details on demand. Roughly worded, these describe key elements of assisting interactive learning:

- *Analyze first* – Reduce burden on human by performing some analysis in the background
- *Show the important* – Direct attention to the most salient information
- *Zoom/filter and analyze further* – Iterative learning with gradually increasing detail
- *Details on demand* – Allow human to reveal the details originally hidden to reduce complexity when requested

Incorporating these concepts with MSTSE visualizations can help to properly frame the process to best support both learning and negotiation goals, by reducing complexity and directing attention to productive areas of the tradespace.

Note that despite the shared interest in mental responses to interfaces and therefore framing, the goals of this research are distinctly separate from the field of cognitive engineering. Cognitive engineering has positioned itself as a primarily descriptive research field, largely disavowing normative and/or prescriptive views of decision making such as utility theory and information processing (Endsley et al., 2007). Cognitive engineering is concerned more with capturing what decisions people make, rather than why they make them or how to improve them. That said, cognitive engineering has influenced the field of visual analytics, specifically in the methods with which new interfaces are tested and evaluated (Greitzer et al., 2011).

## **2.4 The –ilities**

The –ilities are system properties that provide value in the presence of uncertainty, through means other than static mission performance (McManus and Hastings, 2006). For example, flexibility can allow a system to change in response to shifts in the operational context, while survivability can prevent point disturbances from disrupting value delivery. Many research endeavors have sought to define and quantify various –ilities in systems with the goal of intentionally (rather than accidentally) incorporating –ilities into the design process as enablers

of contingent value (McManus et al., 2007; Saleh, 2009; Beesemyer, 2012). Part of the challenge related to including –ilities in engineering systems lies in the fact that the build and carrying costs of –ility enablers are typically paid up front and are easily quantified, while the benefits are uncertain and delayed (Fitzgerald and Ross, 2012). Still, the –ilities have become an important consideration in modern systems engineering for their ability to support traditionally point-optimized designs that would otherwise be susceptible to failure from a change in context. Recent years have seen the beginning of government-sponsored research efforts and agendas devoted to flexibility (Deshmukh et al., 2010), affordability (Carter, 2010a, 2010b), and resilience (Spero et al., 2014), among others.

This literature review should provide the background necessary to proceed to the next chapter, in which we will cover this research’s working theory on how framing impacts the perception of multi-stakeholder problems in tradespace exploration.



### 3 Theory Building for MSTSE Framing

In this chapter, we will discuss the relationship between tradespace exploration, multi-stakeholder decision making (particularly negotiation), and framing. In the process, we will explain why TSE is a *conceptually* appropriate – and even desirable – technique for use in a negotiation setting. Finally, we propose a working theory for how the most prevalent visualization in TSE, the benefit-cost scatterplot, may negatively impact the *application* of MSTSE through poor framing and hypothesize how to improve that framing.

#### 3.1 Principled Negotiation and TSE

Though tradespace exploration was originally tabbed for potential application to multi-stakeholder negotiation through observation of its “coincidental” success in dealing with challenges overlapping those of multi-agent problem solving (Ross et al., 2010a), it is important to address the question of TSE’s ultimate ability to support the underlying theory of productive, mutually beneficial negotiation. If TSE is unable to accommodate the agenda of positive interpersonal negotiation, then its ability to address the matters of structural complexity and multi-objective problem solving will be severely limited by strained relationships and dissatisfied stakeholders. Returning to the “principled negotiation” of Fisher and Ury (1991), they identify four key aspects of negotiation leading to valuable and amicable agreements: separate the people from the problem, focus on interests not positions, invent options for mutual gain, and insist on using objective criteria. These four concepts turn out to be very well aligned with the overarching objectives of TSE, lending strong credibility to the possibility of using MSTSE as a negotiation tool.

First, the purpose of separating people from the problem is to disentangle relationships from problem solving. All negotiations occur in two domains with potential value: the substantive, problem solving domain and the relational domain. Positive relationships carry value by enabling future mutually beneficial interaction (in addition to supporting common societal norms favoring courtesy and teamwork). TSE, as a means of encapsulating a complex problem and separating it from personalities, allows the participating stakeholders to maintain a positive interpersonal relationship even if the substantive issues are difficult to reconcile between them. On a similar level, TSE has a very strong focus on interests over positions. By enumerating many thousands of possible design solutions and exploring them through the use of utility functions, TSE is communicating directly with the interests of all parties rather than allowing discussion to center on specific, individual alternatives. This reduces the likelihood that stakeholders become entrenched on specific solutions (positions) and engage in positional bargaining. Positional bargaining in the design space, without referring back to interests, can lead to a variety of common suboptimal solutions. One such type of solution is the “midpoint” solution that divides the design space but may unequally divide value or result in little to no value at all, as in Figure 3-1. Alternatively, “gold plated” solutions may seek to provide maximum performance/utility to all parties by combining their preferred features, but result in

high costs that reduce overall value, as in Figure 3-2. The importance of focusing on interests lies in the ability of the value space to identify potential similarities between stakeholders and opportunities for mutual value.

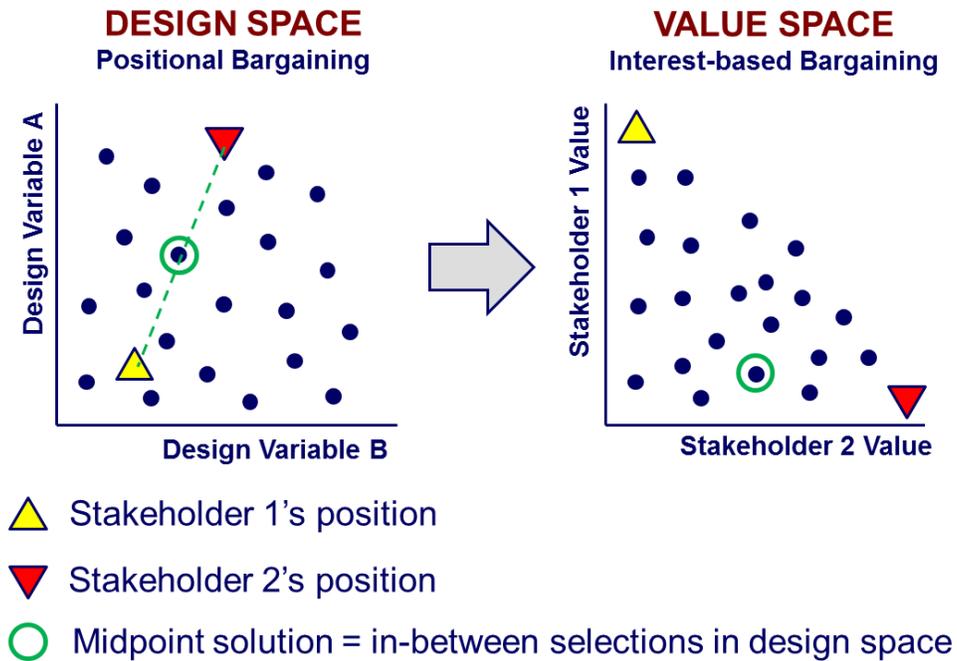


Figure 3-1: "Midpoint" solutions in the design space do not always occupy a midpoint in the value space

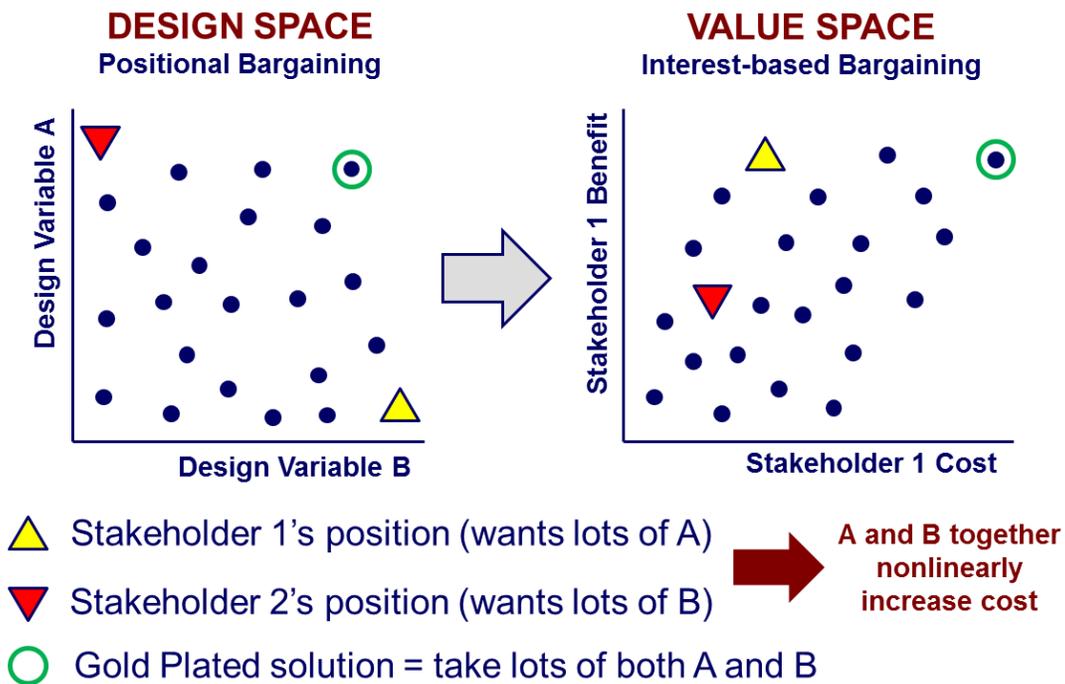


Figure 3-2: "Gold plated" solutions combine favorable design features, but result in vastly higher costs

TSE has the strongest correspondence with principled negotiation in the area of inventing options for mutual gain. The creative assembling of packages that have value for all stakeholders is the most important means of breaking the “fixed pie” assumption, whereby negotiators may assume that a problem is distributive when it is in fact integrative. Most negotiations require both integrative and distributive thinking – a purely integrative problem would imply the existence of a shared preferred solution and thus not much negotiation – thus “integrative” usually implies both types while “distributive” implies a zero-sum game with no mutual benefit available. Figure 3-3 shows these two decision modes on an example two-stakeholder tradespace.

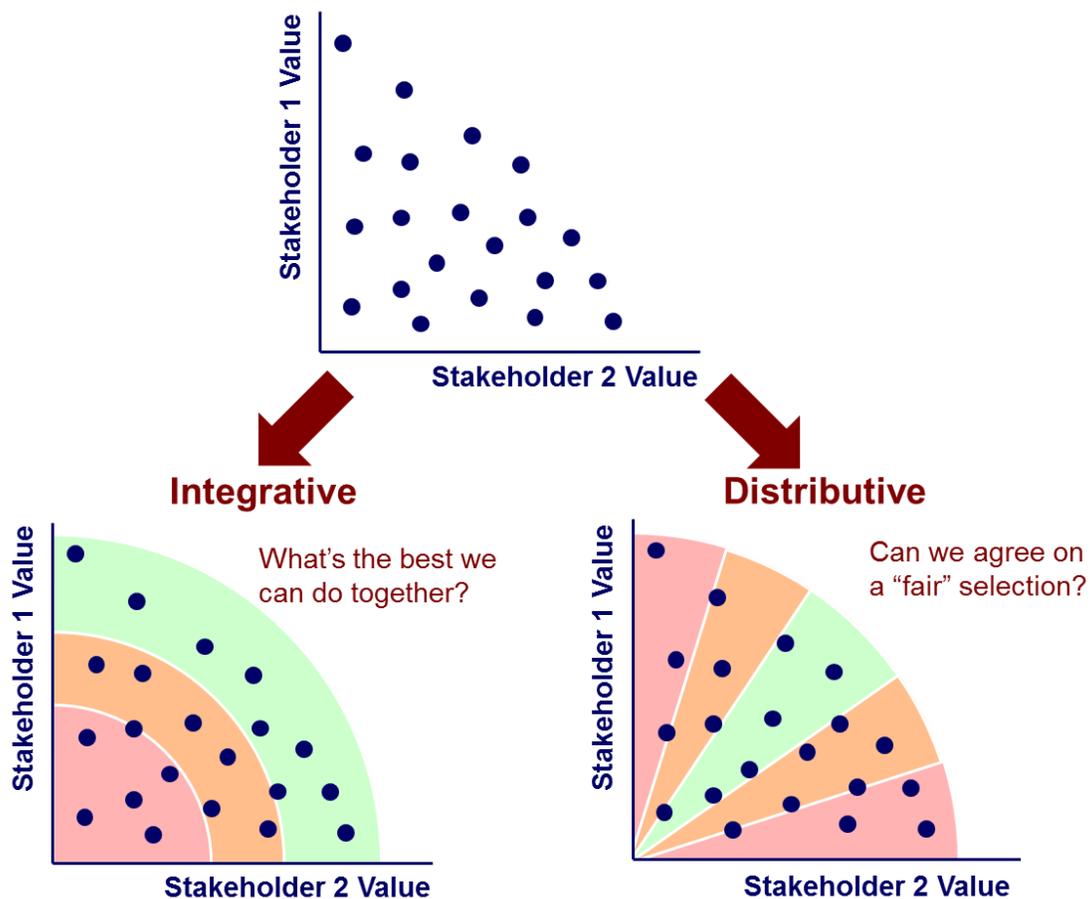
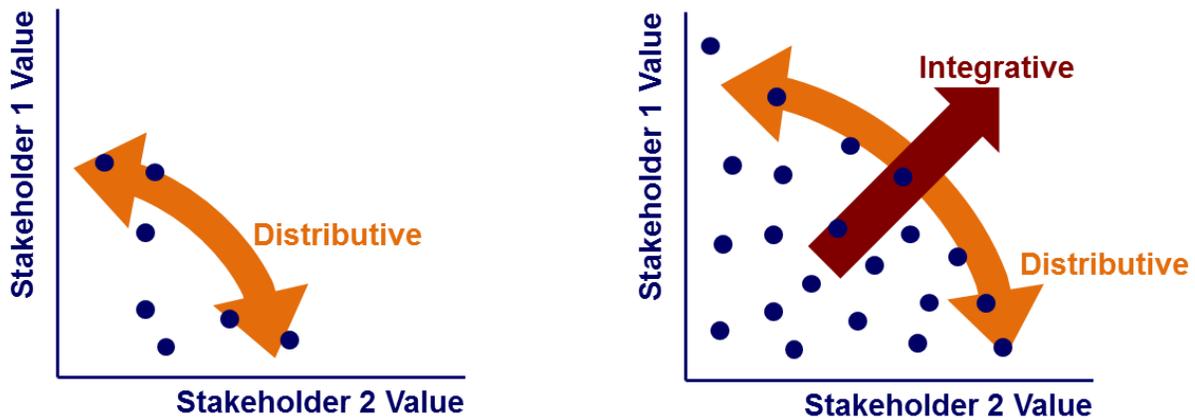


Figure 3-3: An example two-stakeholder tradespace, with both integrative and distributive decision elements<sup>2</sup>

<sup>2</sup> This graphic may make the MSTSE problem seem trivial (select a design in both “green” areas), but this simplifies the true MSTSE problem by giving each stakeholder a one-dimensional value function. Remember: in TSE, each stakeholder generally has separate benefit and cost dimensions of value and is potentially uncertain what tradeoff is best for themselves, preventing this type of visualization. We only use this image to demonstrate how the concepts of integrative and distributive bargaining appear in a tradespace of alternatives.

Evaluating and exploring many potential solutions is critical to solving the integrative problem as well as possible. TSE, by procedurally generating thousands of potential solutions, often via a full-factorial combination of all design variables when possible, creates many options as a matter of course. TSE also allows for (and encourages) iterative expansion of the design space if more variables or choices are discovered. This process has a high probability of uncovering mutual-value creating opportunities, if they exist. Because distributive bargaining inherently puts stakeholders' interests at odds with each other, it is the more likely type of bargaining to disrupt negotiations and should be conducted after creating integrative solutions. Stakeholders that engage in the "value claiming" of distributive bargaining prematurely may end negotiations before high-value solutions are found, as in Figure 3-4. Hence, a key benefit of applying TSE in the negotiation domain (over, say, point-design studies) is the emphasis on establishing a large number of potential solutions at the beginning of the analysis, increasing the likelihood of "good" alternatives being available when distributive bargaining begins.



**Figure 3-4: Engaging in distributive bargaining before integrative bargaining (via creation of many alternatives) limits the total amount of value available in the tradespace**

Finally, the use of objective criteria is fundamental to engineering and engineering design, so the last principle is not as likely to be an issue when applied to TSE as it might in other fields. However, subjective judgements and 'expert opinion' are still common decision drivers in some engineering projects. TSE seeks to operate on an objective dataset and strongly encourages more rigorous criteria when available, both to increase confidence in the results and to reduce the stakeholder effort necessary to subjectively assess thousands of alternatives (or more). The creation of an objective dataset can also be accomplished jointly when possible, in accordance with the principles of Joint Fact Finding or a Collaborative Modeling technique, further reducing the potential for he-said-she-said arguments of opinion.

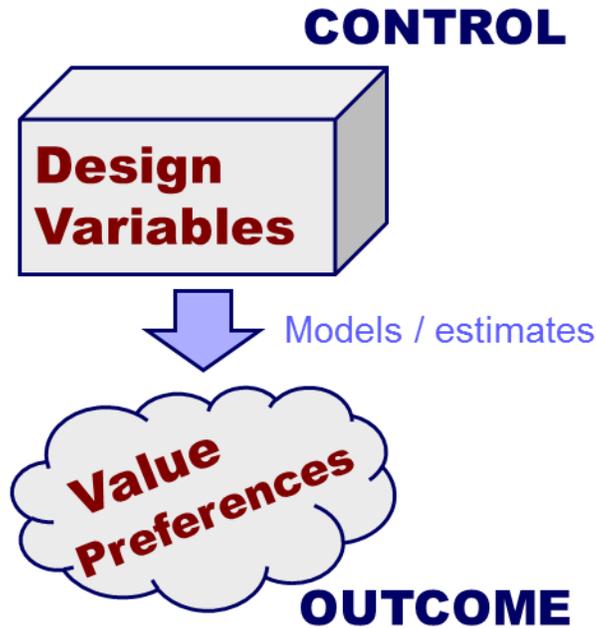
It is also worth mentioning that the FOTE principle of Raiffa (2002) is also compatible with TSE. Typically, no information is hidden in TSE; early applications of MSTSE extended this principle to multiple stakeholders by allowing full and complete information sharing between all parties (Ross et al., 2010a). Though some applications may need to develop

workarounds due to proprietary information (or even competitive interests on the farthest end of the collaboration-coordination spectrum), many engineering projects are conducted in good faith cooperation between stakeholders where information is shared freely. In this case, TSE should be well suited to communicating this information.

Thus, overall it appears that TSE does, in fact, have a considerably tight alignment with the mantra of principled negotiation, providing significant reason to believe that MSTSE has the potential to be a powerful technique for assisting engineering design tasks that must balance multiple interests. None of this is to suggest that this combination of negotiation and engineering will be as easy as taking the recommendations of the negotiation literature and simply applying them directly. A key difference between most negotiations and the complex system design of interest to this research is the degree of control over substantive issues available to the negotiators/designers. In many negotiations, participants are directly negotiating on the issues that provide value; for example, if a union is negotiating with an employer, the resulting agreement will directly set the pay, vacation, and benefits that each party has preferences on. However, engineering design typically controls low-level form variables which only *indirectly* determine the value-generating attributes. Using a car as an example, designers can directly negotiate the horsepower and wheel size, but derive value from the acceleration and top speed that result from these decisions<sup>3</sup>. This disconnect between the control and outcome, filtered through models and estimates as pictured in Figure 3-5, can obfuscate value and is a complexity not accounted for in much of the negotiation literature. Though our car example is relatively simple, a more complex system may rely on models that can display unanticipated or emergent behavior (for example, an agent-based simulation), which further confounds the ability to predict the impacts of available decisions in the value space. This limits the applicability of some types of prescriptive negotiation advice, which depend on exact knowledge of the value associated with each controllable variable.

---

<sup>3</sup> Some stakeholders may actually have preferences on controllable design variables such as horsepower (form) instead of acceleration (function), but we will assume (for now) that our hypothetical TSE stakeholder is following the principles of Value Driven Design (Collopy et al., 2012) and does not hold form-dependent preferences.



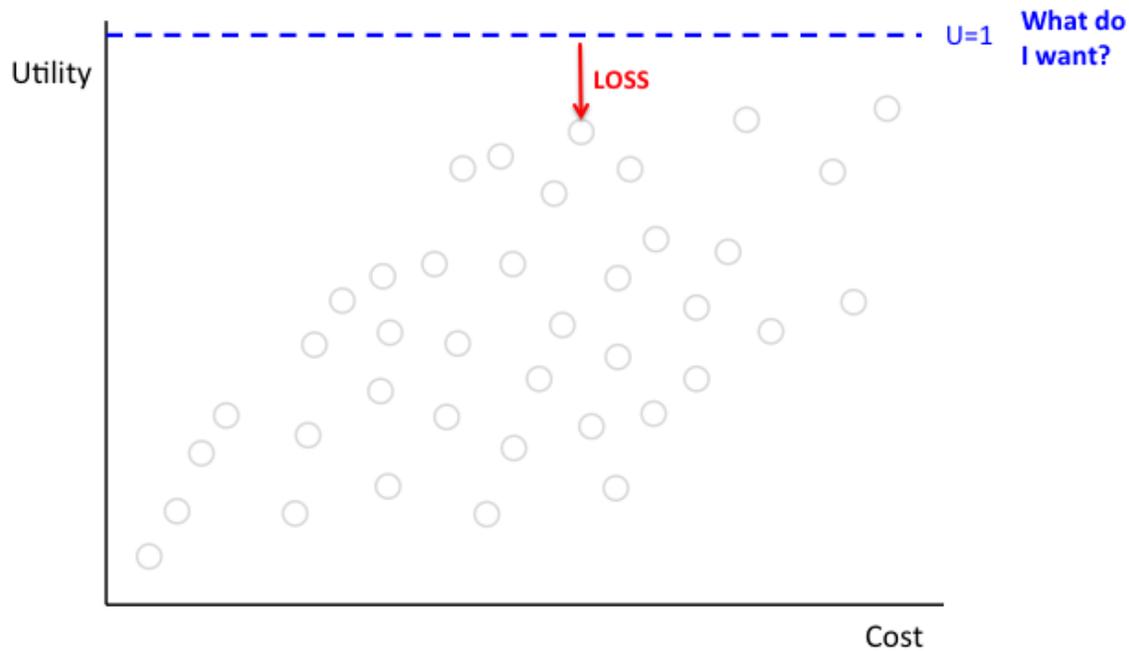
**Figure 3-5: Disconnect between controllable variables and preferences in engineering applications**

The benefit-cost scatterplot is the key TSE tool for overcoming this control-outcome disconnect, by allowing for the ready comparison of all possible alternatives on value axes. However, we must be careful when applying the standard TSE activities and visualizations to MSTSE. As the problem changes, so too must its framing in order to prevent stakeholders from utilizing design heuristics developed for individual design problems and individual TSE in the more complex multi-stakeholder domain.

### **3.2 Framing in TSE**

Framing is a pervasive concept, occurring constantly even without intention, which is why adopted techniques (such as TSE visualizations in MSTSE) should be examined carefully for unintentional framing effects. During standard, individual-stakeholder TSE, the stakeholder’s utility function is usually elicited before the design points in the tradespace are evaluated, yet analysis of these stated preferences is held until after the stakeholder is shown the cost-utility tradespace plot (Ross and Hastings, 2005). One of the reasons for this is that experts in TSE know the risk of the stakeholder fixating on the “utility = 1” condition. With no information about physical constraints (as represented by evaluated design points), it is natural for the stakeholder to fixate on reaching maximum utility, or the complete satisfaction of all of their stated preferences, since the act of considering the utility function revolves around the question of “What do I want?” However, it is extremely common for no designs to achieve maximum utility and it is important to prevent the stakeholder from establishing that as their reference point. The use of such a reference point would make other feasible designs in the

tradespace look significantly worse by comparison, appearing as losses from the unreasonably high reference, as in Figure 3-6.I



**Figure 3-6: Utility = 1 reference point, prior to plotting designs, with associated "loss" once designs are seen**

Instead, the benefit-cost tradespace scatterplot is typically the first view shown, illustrating feasible solutions for varying combinations of cost and utility and changing the operative question from “What do I want?” to “What can I get?” People comfortable with TSE are accustomed to this visualization, and this research hypothesizes that the available reference point from this view is the Pareto front, as depicted in Figure 3-7. The Pareto front is the set of designs with the most efficient cost-utility tradeoffs and is readily apparent on the cost-utility plot, as it is on the edge of the mass of design alternatives. This is a strong choice of reference point for individual TSE, as the stakeholder is empowered to select any of those Pareto efficient designs if he desires. Correspondingly, many TSE visualizations and analysis techniques focus heavily on the Pareto front (Ross and Hastings, 2005; Stump et al., 2009), as it represents the “best” solutions available for a single decision-maker’s defined sources of benefit and cost. These techniques were originally designed to help stakeholders find alternatives with the most value *for them*. In a single stakeholder problem, this is reasonable because the ultimate goal of the design task is to find the design choice that best satisfies the stakeholder’s perception of value. The visualization tools accurately present a framing match between the macro frame of the individual problem solving effort and the micro frame of the data visualization.

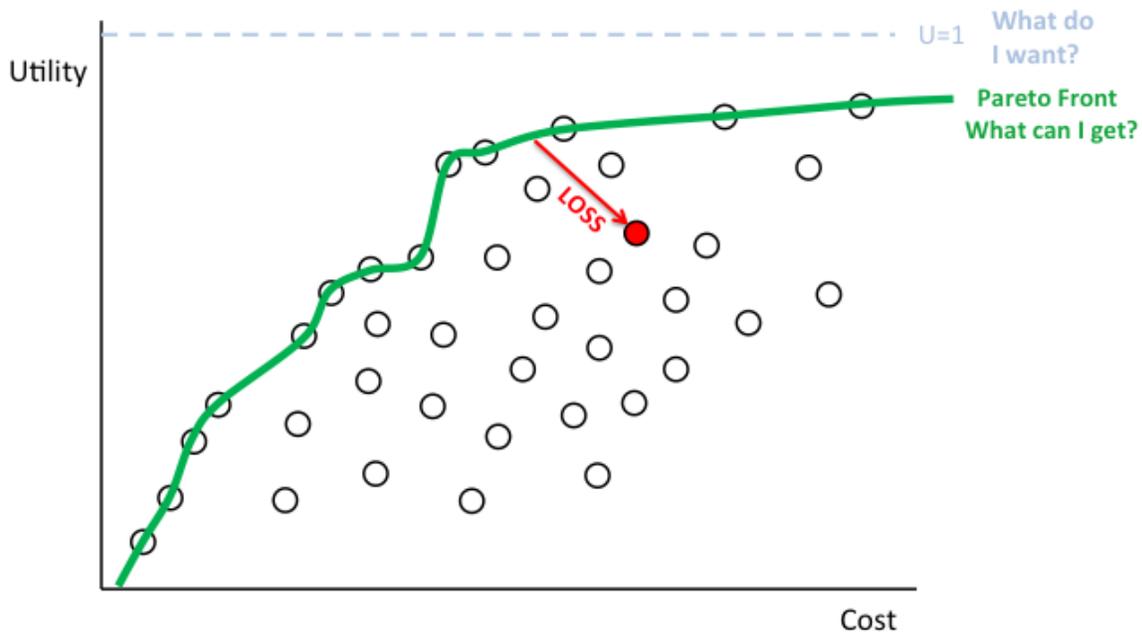


Figure 3-7: Tradespace with plotted designs, "loss" evaluated with reference to Pareto front

### 3.3 Framing Mismatch in MSTSE

As discussed in the literature review, limitations in the ability of most common value-quantifying techniques, including utility theory, to be aggregated across groups of people has caused MSTSE to largely rely on the best practices for individual tradespace exploration, with all stakeholders using those methods in parallel. In this context, individual TSE was performed with the expressed goal of finding specific design points to use as focusing points for later discussion and was conducted using a series of standard tradespace views and analyses, including individual benefit-cost tradespace plots, design variable comparisons, variable sensitivity plots, and animated views of context uncertainty (Ross et al., 2004; Ross et al., 2010a). By encouraging stakeholders to spend time on the problem individually, existing MSTSE procedure has implicitly encouraged each stakeholder to reflect on the same operative question as TSE: “What can I get?” This question ignores the importance of the other participants and their needs during MSTSE. The ultimate goal of MSTSE is to find a mutual agreement, and searching for individually optimal or near-optimal solutions does not necessarily support this goal as it ignores the relational aspects of the problem. The group decision problem is not just a series of individual decisions and is heavily affected by interpersonal dynamics and psychological considerations of what makes a “good” compromise and what constitutes a “fair” solution.

In MSTSE between independent decision makers, no single stakeholder has the ability to dictate decisions based on only their own value. Allowing a stakeholder to establish their individual Pareto front as a reference point is optimistic, since the actual solution cannot be better than that reference (with respect to their own utility and cost) and will almost certainly be

worse, driven away from individual optimality by the realities of physical constraints and the divergent needs of different stakeholders. If the individual Pareto front becomes the reference point for one stakeholder, their interactions with other stakeholders will operate exclusively in the “losses” frame, which leads to more stressful decision making and an increased chance that one or more stakeholders will simply leave the negotiations (Gonzalez et al., 2005). Since this applies equally to every party, feasible system alternatives may appear as losses to each stakeholder, potentially blocking the acceptance of a mutually beneficial solution that could otherwise be agreed upon. Therefore, the view best suited to individual TSE, due to its information density and ability to show tradeoffs in what are typically the most important decision criteria, risks being counterproductive when used in MSTSE without regard for controlling the resulting framing of the problem. The transition from TSE to MSTSE is one of increasing complexity, which demands increased sophistication from the associated tools in order to accommodate and visualize relational constraints, as shown notionally in Figure 3-8. If the utility-cost tradespace under the individual value problem is extrapolated for use in the multi-stakeholder problem, the micro framing of the activity underrepresents the complexity of the problem, potentially leading participants to adopt suboptimal or counterproductive behaviors in the context of the larger problem solving effort.

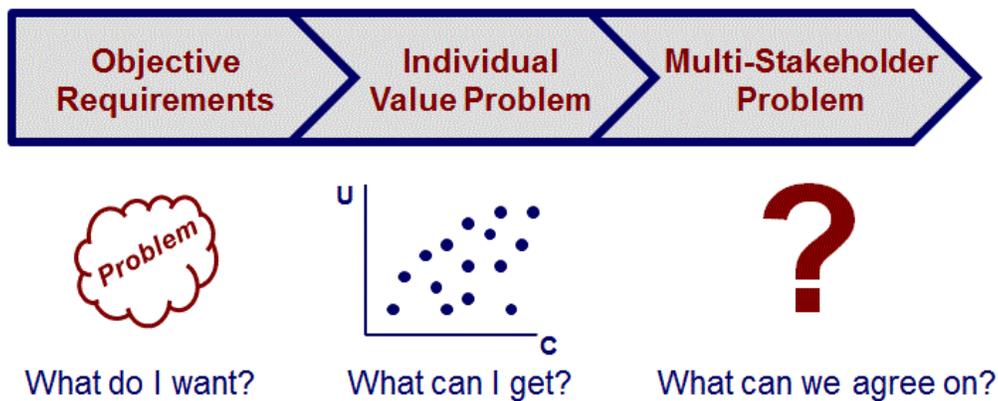


Figure 3-8: Increasing complexity from problem definition to physical constraints to relational concerns

### 3.4 Working Theory

This research operates under a working theory of the ways in which framing impacts participants in MSTSE under the current practices. The following two subsections explain the current view of the role of both macro and micro framing in MSTSE.

#### 3.4.1 Macro Framing in MSTSE

Macro framing is partially a function of the individual stakeholders who choose to participate. Personal experience, philosophy, and perspective all affect the ways in which stakeholders interact with problems and other stakeholders. These are not features to be controlled or changed, but rather to be guided to productive ends. Without explicit recognition

of budding ideological impasses between stakeholders driven by competing macro frames, it is possible for any joint activity, including MSTSE, to be permanently derailed or cancelled due to miscommunications and the inability to “speak the same language”. However, since this type of frame is essentially open-ended and highly application specific, with effectively infinite different ways to approach a given problem, it is impractical to suggest that unilateral action or advice could account for the challenge. In the spirit of productive, mutually beneficial negotiation, encouraging an open discourse where motivations and philosophies are made clear is the most likely way to expose and address the general frames that stakeholders are operating under.

The large impact of personal history on macro framing does *not*, however, mean that the macro frame is not influenced by the activities of MSTSE. Tasks stakeholders are asked to consider or perform can be implicitly associated with particular problem solving frames. To reiterate a previous example, an activity such as “find the cost-utility Pareto front” given to each stakeholder individually (a common TSE tactic frequently extended to MSTSE) can serve to reinforce a macro frame of value claiming by implicitly emphasizing the importance of maximizing one’s own objectives. Macro framing is driven by high-level problem formulation and approach. Addressing this type of macro framing can be accomplished by examining the systems engineering processes surrounding applications of MSTSE: checking that any macro frames promoted by these methods are as close in-line with the principles of negotiation as possible, in order to promote positive interaction between the participating stakeholders. Additionally, certain features of a given problem will naturally impact the macro frames with which stakeholders interact (e.g. positive vs. negative existing relationships), and these can be used to customize MSTSE in appropriate, problem-specific ways.

### **3.4.2 Micro Framing in MSTSE**

Micro framing, in obvious contrast to macro framing, is driven by the particular data representations and visualizations present in MSTSE. These constructs have considerable impact on what stakeholders see and consider during their analysis. Accurate reference points should be emphasized, and relevant data should be accessible at all times. Unfortunately, the current implementation of MSTSE as just TSE with more than one person has led to the use of techniques designed to solve the less-complicated single-stakeholder problem, such as individual cost-utility plots and Pareto front solving. These techniques were not originally intended for use on problems so heavily influenced by social factors and thus have little to no consideration for establishing group-oriented reference points or data types. This can lead to loss of situational awareness and a reversion to strictly self-interested behavior when a mutual group-based approach would be superior.

Direct management of the micro frame is possible through the control of the data and views of the data shown to stakeholders during MSTSE, giving considerable leverage to improve the framing with new MSTSE visualizations and tools. Of particular interest is the ability to encourage systematic thinking over heuristic thinking, as heuristics derived from experience with

individual TSE will again favor individualistic behavior, possibly leading to discord or failure to reach agreement. The notional reason for applying visual analytics to MSTSE lies in its ability to promote systematic thinking and capture the group dynamic of the problem in ways superior to existing heuristic visualizations.

### **3.4.3 Macro and Micro Interaction**

The macro and micro frames also can have a form of “weakest link” relationship, where if one frame is off, the other can be dragged down. If a stakeholder approaches MSTSE with a macro frame that is highly confrontational and individualistic, they will likely choose to utilize a micro frame, in the form of a particular visualization for example, that matches their outlook (if they are able to do so). Alternatively, if only individualistic visualizations are available, a stakeholder’s macro frame may be slowly pushed into a similar mindset in order to reduce cognitive dissonance: an extension of the micro/macro framing mismatch discussed in Section 3.3. Thus it is critical that both types of framing be addressed to the best ability within each application of MSTSE.

## **3.5 Improving MSTSE Framing**

We (as developers and researchers of MSTSE) are limited in our ability to manipulate the personally held beliefs and perceptions that affect macro framing, as discussed above. Embedding new activities into the standard TSE process (or calling them out explicitly where they may previously be implicit) has the potential to engender a positive multi-stakeholder problem solving and negotiation mindset from the beginning. For example, a Joint Fact Finding (JFF) venture inserted at the beginning of the process to develop a shared understanding of the design problem and a common, objective ground on which to negotiate has been demonstrated to reduce interpersonal friction and miscommunication in negotiations.

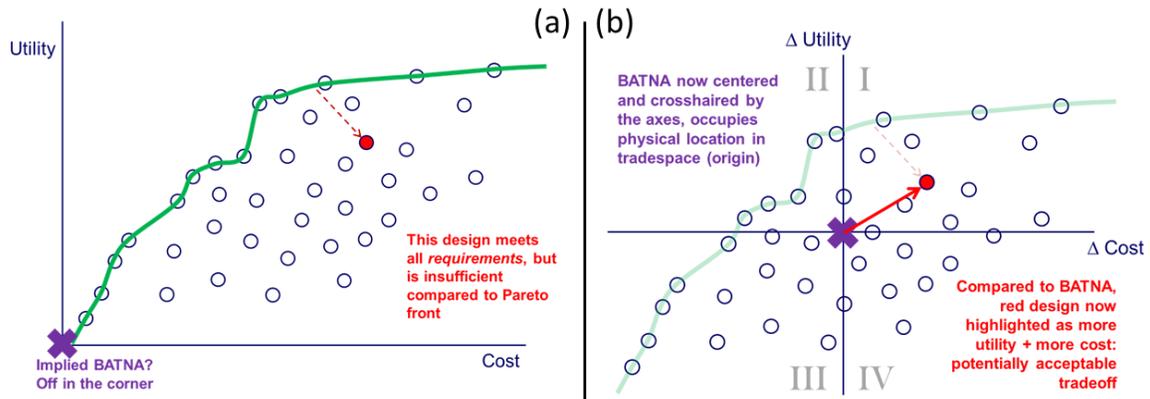
The *best alternative to a negotiated agreement* (BATNA) (Fisher et al., 1991) represents an interesting concept that straddles the differences between macro and micro framing. In the negotiation literature, negotiators are encouraged to compare potential agreements directly to their BATNA, as it represents a boundary of value that defines when leaving the negotiation is the best decision. The BATNA is the recommended reference point as it is the point at which *perceived* gains and losses map into *actual* gains and losses. Establishing a best alternative is not a standardized practice in TSE, despite occasional attempts to do so (Richards et al., 2008), as generally all of the potential solutions are included in the defined design space. Some TSE endeavors are conducted to explore potential changes to some baseline or currently existing system, where failure to adopt a new solution would default into the original design as the best alternative. Other stakeholders may simply have a rough understanding of the utility achievable for a given amount of money via means outside of the design space (although this is harder with utility than with monetized benefits, given its abstract nature). However, if there are truly no alternatives, then the best alternative is simply “no solution”, which creates no utility for zero cost. Which of these conditions applies should be determined up front by each stakeholder when

engaging in MSTSE, and used as a BATNA. This activity – the explicit analysis and consideration of the BATNA, performed as a setup step to MSTSE – represents an attempt to influence the macro frame of the negotiation by grounding it firmly in reality. When the tradespace exploration begins, it would then ideally be designed to set the micro frame of the problem in such a way that the BATNA is reinforced as the reference point through the data and visualizations given to and used by the stakeholders.

In addition to emphasizing a proper reference point, MSTSE should also theoretically benefit from increasing the availability of information relevant to the group problem that is absent in traditional TSE. *Availability* is often used as a heuristic in human decision making, based on the common but sometimes erroneous assumption that the information that is available is the information that is important (Tversky and Kahneman, 1974). Of particular relevance here is the fact that the standard TSE cost-utility tradespace does not show each stakeholder any indicators of their counterpart's value, creating a micro frame that completely leaves out information relevant to the group problem and implicitly suggests that it is not important. Without this key information, an MSTSE exercise will struggle to evolve past positional bargaining, the offer-counteroffer negotiation strategy that is a natural approach to feeling out a partner's interests. Positional bargaining is considered a degenerate strategy to be avoided, as it frequently fails to identify solutions that are mutually beneficial (Fisher et al., 1991; Raiffa, 2002). By making each stakeholder aware of the others' interests and estimates of their value up front, MSTSE may reduce the time wasted and negative feeling associated with the offer of a potential solution by one party that is clearly unacceptable to one or more other parties.

### **3.6 Modifying Tradespace Visualizations**

As discussed previously, the default view of traditional TSE is the individual cost-utility tradespace, as depicted in Figure 3-9a. Where is the BATNA in this visualization? With no additional information, it is generally implied that this is a zero-cost zero-utility BATNA, placing it in the lower left corner, away from the focus of attention in the plot: the design points. However, we want to draw attention to the real BATNA (as recommended to be clarified early in the design process) and encourage its use as a reference point. An easy way to accomplish this is to re-center the visualization, changing the axes to *differences* in cost and utility from the BATNA, placing the BATNA at the center of the plot and on the origin of the axes, as shown in Figure 3-9b. With a more immediate reference point in the origin than the Pareto front (because it is directly indicated on the plot), designs like the one highlighted in red that may have been identified as inferior to Pareto efficient alternatives in the original plot are more visibly identifiable as tradeoffs in cost and utility that may be preferable to leaving the negotiation with no agreement.



**Figure 3-9: A cost-utility tradespace before (a) and after (b) being centered on the BATNA**

Part of the challenge in redesigning the benefit-cost tradespace for MSTSE lies in breaking the bad habits of individual value claiming by stakeholders who can confidently engage with tradespaces: a developed heuristic path of information processing. Presenting the same information in a new way can prevent people from leaping to comfortable, predefined conclusions such as a commitment to choosing a design on the Pareto front, in favor of more systematic thought. Even a simple change, such as the rotation of the tradespace axes, could force stakeholders to more carefully consider the presented information, including the new information introduced by the updated MSTSE visualization. A forty-five degree clockwise rotation has the added benefit of making “up” the desired direction of movement (lower cost, higher utility) instead of “up-left”, which is marginally more intuitive.

Tradespaces are often augmented with a color-axis, displaying additional objective information such as design variables, in order to assist stakeholders in mentally categorizing areas of the tradespace and understanding the relationships and tradeoffs within the problem. With the introduction of additional stakeholders, color presents one way of visualizing the new dimensions relevant to decision making. One possible implementation of this could color designs by the quadrant they occupy of the BATNA-centered tradespace of another stakeholder<sup>4</sup>. This makes information about four distinct categories of designs immediately available: superior designs (quadrant II), tradeoff designs with higher cost and utility or lower cost and utility (I and III, respectively), and inferior designs (IV), as can be seen in Figure 3-9b. Furthermore, additional plot dimensions such as transparency or size could be used to indicate a measure of distance from the Pareto front such as Fuzzy Pareto Number (FPN) (Fitzgerald and Ross, 2012).

<sup>4</sup> Note that this visualization suggestion allows only one extra stakeholder to be considered, though MSTSE is in general not restricted to only two participants. The number of intuitive, clearly visualizable dimensions in a tradespace/scatterplot is limited (for example, adding a third axis to a 2D visualization makes it more difficult to read), suggesting that visualizations that show each possible alternative as a design point will be forced to accommodate more than two stakeholders in other ways (e.g. toggling, or through stacked/linked graphs).

This can be used to both visually indicate the approximate efficiency of a design for a negotiation partner and also deemphasize the individual Pareto front further, by either fading out or shrinking the designs that are not also efficient or near-efficient for the other party. A notional tradespace matching Figure 3-9 but also incorporating rotation, color, and transparency is shown in Figure 3-10. For the sake of clarity, we will refer to this as the “negotiation tradespace” (in contrast to the classic benefit-cost tradespace) for the remainder of this research.

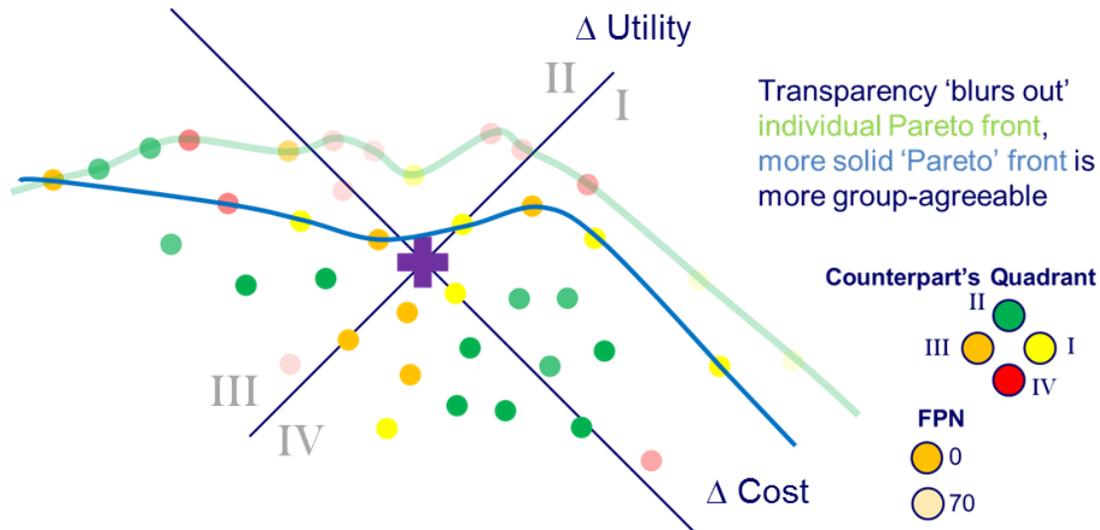


Figure 3-10: Example re-centered, rotated, and color/transparency MSTSE visualization

### 3.7 Discussion

This chapter has attempted to outline a working theory on the differences between TSE and MSTSE, specifically with regards to how the concepts of principled negotiation and framing impact the current practice of MSTSE. The core tenets of principled negotiation (separating people from problem, interests over positions, creating alternatives, and using objective information) were shown to have strong correspondence with the underlying mechanics of TSE, suggesting that TSE could be a useful technique to support such negotiations. However, the realities of framing at both the micro and macro levels suggest that traditional TSE visualizations and activities, such as the benefit-cost tradespace and Pareto front solvers, are likely to prime stakeholders to engage in counterproductive negotiation strategies and bias themselves towards negative impressions of feasible, beneficial solutions.

Using only these principles, some modifications to the classic benefit-cost tradespace have been presented here. These changes are hypothesized to improve the effectiveness of the tradespace in a multi-stakeholder environment by increasing the availability of important negotiation-specific information related to each stakeholder's BATNA. Should the implementation of these recommendations improve MSTSE negotiations, we would have

evidence in favor of the working theory. The following chapter describes the implementation of these modifications and their subsequent testing in a controlled experiment.



## 4 Tradespace Redesign Experiment

This chapter describes a controlled experiment that was conducted in order to test the working theory of reference points in the benefit-cost tradespace. In summary, the experiment involved asking subjects to complete a short two-person negotiation task, assisted by tradespace exploration software. Subjects were separated into either the control group, who used traditional TSE visualization tools, or the treatment group, who were given access to the “negotiation tradespace” of the previous chapter instead of the classic benefit-cost tradespace. Theoretically, it was anticipated that the treatment group would display improved grasp of the group problem, along with the potential for improved solution quality, increased speed of deliberation, and higher degree of satisfaction, among other outcomes. Results of the experiment confirmed the main hypothesis but were unable to determine if solution quality improved due to the small sample size and “easy” nature of the problem making it easy to solve even with traditional tools. Other observations provided interesting unexpected insights as well.

### 4.1 Experiment Overview and Objectives

This experiment was designed to investigate two research questions (identified here as EQ for “experiment questions” to avoid confusion with this thesis’ larger research questions):

- EQ1. In what ways do framing effects caused by visualization tools impact participants’ perceptions when conducting multi-stakeholder tradespace exploration?
- EQ2. In what ways do these framing effects impact the group outcomes of multi-stakeholder tradespace exploration?

It was hypothesized that the current practice of multi-stakeholder tradespace exploration systematically reinforces an unrealistically optimistic individual reference point, resulting in detrimental negotiation behaviors due to misinterpretation of gains and losses by the participants. To investigate this, we used an experiment, comparing the results of a demonstrative tradespace exploration exercise with two participants between a control group using traditional tools and a treatment group using alternative tools designed to reduce the impact of the hypothesized framing effects.

An experiment was the method of choice to target these questions for a variety of practical reasons. First, TSE as a design paradigm is rising in popularity but has far from the sort of widespread user base that would support a descriptive survey or ethnographical approach, particularly the multi-stakeholder variant which is notionally intended to support the negotiation of high-level decision makers rather than rank-and-file engineers. A small-N case study of real world applications of MSTSE would be more appropriate, but is hampered by the sensitive nature of the decisions made with MSTSE, limiting the majority of complex sociotechnical system case studies to be performed on reconstructed data sets, which may lose key features of the original framing.

Fortunately, an experiment is well positioned to achieve the desired goals of this research. Negotiations with role-playing participants are a popular subset of serious games for research purposes, often incorporating controlled access to a computer model with open or unstructured communication between the parties (see Mayer, 2009 for a history of gaming research; or Grogan, 2014 for a recent example). Although mental processes are difficult to control experimentally, impeding our ability to control emotional or historical macro framing, micro framing such as that created by visualizations can be controlled effectively through the use of the tools given to experimental subjects. If the TSE visualizations that have been used in MSTSE can be modified to reduce the influence of their hypothesized framing, subjects given access to these new tools should display differences from a control group in their self-reported perception of the problem, and differences in the task outcomes (agreement, speed, etc.) can be attributed to these effects.

## ***4.2 Unit of Analysis***

The unit of analysis for the EQ1 is the individual. Framing effects are known to operate at an individual, perceptual level. Each participant in a multi-stakeholder tradespace exploration activity supplies their own interpretation of the data and is therefore subject to potential framing effects. Accordingly, we will quantify the effects of framing on each subject in the experiment individually, despite them working in groups, by comparing exit questionnaire responses between the control and treatment groups. Though there is a possibility of "group-level" framing effects impacting MSTSE, we leave the investigation of those to future research. Additional questions were included to capture the subjects' responses to the tools and their negotiation counterparts, as strong responses to factors such as these may interact with the effects of framing on reference points in unanticipated ways.

The unit of analysis for the EQ2 is the group: each pair of subjects randomly assigned to work together. Because MSTSE (and thus the experimental task) is a highly interactive group task, it is most meaningful to consider the outcome of the task at the group level rather than breaking it down to the level of individual performance. Outcomes of the pairs were used to assess if there were any effects on task performance and solution quality associated with the presence of unintentional or detrimental framing effects.

## ***4.3 Subject Selection and Assignment***

The subjects were volunteers taken from MIT undergraduate and graduate engineering programs, restricted to students with at least 2 full years of experience in engineering. Volunteers were sought via mass email and other public announcement media. Though the sample population is one of convenience, raising concerns for external validity, engineering students are a reasonable approximation of the most common participants in real applications of MSTSE: technical stakeholders with an engineering background. Generalizability is an obvious concern; however, many technical stakeholders in industry and government come from

engineering education backgrounds, similar to that of MIT, that influence how they approach analytic problem solving tasks. Students with some experience in engineering analysis (hence, the restriction to a minimum of 2 years of engineering education) should roughly replicate those behaviors, giving reason to believe that insights from this experiment can apply elsewhere. Additionally, it has been demonstrated that students display comparable behavior to professionals when role-playing in a negotiation, though this has relied on “training” that may not be applicable to our population (Herbst and Schwarz, 2011). Regardless, generalizations of the results of the experiment should still be made carefully, particularly when applied to non-technical stakeholders.

Thirty-two subjects volunteered for the experiment and twenty-six completed the task, for thirteen total trials. A larger sample size would be desirable to increase significance; however, the goals of EQ1 and EQ2 are concerned more with identifying the direction of impact caused by framing, rather than quantifying an exact effect such as “percent time saved”, which does not require as large of a sample. This experiment could be repeated in the future with a larger sample size to capture more quantitative insights into solution quality and speed.

All volunteers were randomly assigned into pairs and each pair randomly assigned to the control or treatment group. Assigned pairs were contacted together in order to coordinate a time in which both subjects could attend the experiment for a minimum of one hour. Fully random assignment from the volunteer population was used to promote internal validity for the experiment. Though the sample size for this experiment was limited due to resource constraints, problems were not anticipated with regards to potentially uneven distribution of demographics between the test groups. Previous experience with TSE is almost certainly the most relevant demographic variable, but the potential to avoid framing effects due to higher levels of comfort analyzing the data is likely counterbalanced by an increased tendency to fall into individualistic habits developed from familiarity performing single-stakeholder tradespace exploration.

## **4.4 Trial Protocol**

The following subsections describe the protocol which the subjects were asked to complete during their session.

### **4.4.1 Pre-task**

Subjects were presented with a description of the experiment as "an investigation of multi-stakeholder tradespace exploration", along with information of all of the relevant, minimal risks and will be asked to sign a standard consent form. Data collected for the experiment is confidential; all collected data has been attached only to an identifying code for each participating group that is not tied to subject’s real names or identifiable personal information.

After this, subjects were given a short document describing the basics of TSE, along with a description of the problem they will solve and their roles in the case. Previous research has

suggested that the use of a short instructional video can train people with no previous tradespace experience to find solutions for moderately complex problems that approach similar quality to the solutions of TSE experts (Wolf et al., 2011). Because the tradespace for this experiment was heavily simplified, it is believed that a brief document covering the use of the tools and the basic goals of tradespace exploration and an accompanying walkthrough by a researcher was sufficient for training, though future research should verify this assumption. Subjects in the treatment group received a slightly longer document that included example plots and corresponding explanations that describe the new visualizations designed for the experiment, though with any references to framing effects removed in order to limit potential 'pretest' bias. A researcher was present for the entire experiment to answer any functional questions (e.g., "How do I do this?") but remind subjects that all decision making processes and ultimate decisions are acceptable and up to them (e.g., deflecting any "What should I...?" questions).

#### 4.4.2 Experimental Task

The problem the subjects engaged in solving was a notional tradespace exploration between two roommates, **Nat** and **Vic**, who are deciding on a car to purchase together. In order to keep the trials within a reasonable length and not unduly load the subjects, the example case was stripped of technical complexity in favor of analytical simplicity, allowing participants to focus on interacting with each other and the data. Additionally, by making the case simple enough that "good" solutions can be found even with the control tradespace, we eliminate the possible counterargument that our control is a "strawman" that would not be a realistic application of TSE. The downside of this decision was the risk that the (relatively) subtle differences between the control and treatment groups would be unable to generate significant differences in task performance. Future experiments could explore the effects of increased technical complexity, though they will likely require more sophisticated and experienced subjects in order for results to be meaningful.

The characters in the problem have clearly defined costs and benefits: one has agreed to pay the purchase price and cares about the car's reliability rating (on a scale of 1-10), and the other has agreed to pay for all future gas (thus "paying" for gallons per mile) and desires a car with as high a top speed as possible. Note that the "benefit" attributes were used directly, rather than through a utility function. This allowed the subjects to supply their own subjective utility without requiring an explanation of the mathematics underlying utility theory<sup>5</sup>. The tradespace was constructed of 100 design points, corresponding to different available cars, each graded on all four of the relevant value dimensions. Subjects visualized the data using VisLab, a MATLAB® software package designed for interactive tradespace exploration. VisLab was

---

<sup>5</sup> This is besides the fact that using a utility function when there is only one dimension of benefit does not reduce complexity and it would potentially compromise the experiment by resulting in different behaviors based whether or not a given subject agreed with the chosen utility function.

developed in approximately three six-month periods (with sporadic work in between) at the Systems Engineering Advancement Research Initiative (SEARI) at MIT, and was augmented specially for this experiment.

The experimental treatment differentiating the two groups was in the control of the VisLab "dashboard", which was customized for each group to allow for access to different views of the data. The control group was given access to the basic tradespace exploration tools and views that have been used in traditional individual TSE and previous MSTSE explorations. The standard VisLab tradespace viewer is shown in Figure 4-1. The treatment group had access to the same tools, but with the 'default' tradespace viewer replaced by the negotiation tradespace: a modified version of the benefit-cost tradespace designed to emphasize the BATNA over the individual Pareto front, while requiring coloring and transparency to be used as indicators of their counterpart's value, shown in Figure 4-2. The BATNA for each role and its associated performance in each attribute was explicitly defined in the role-play materials.

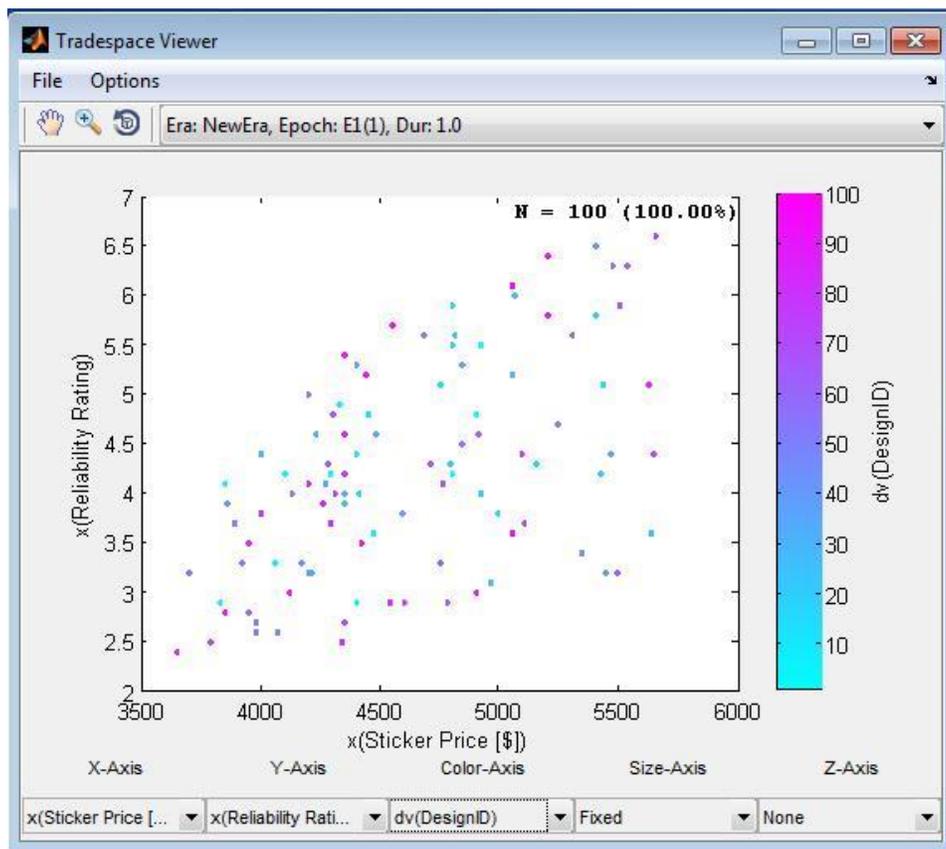
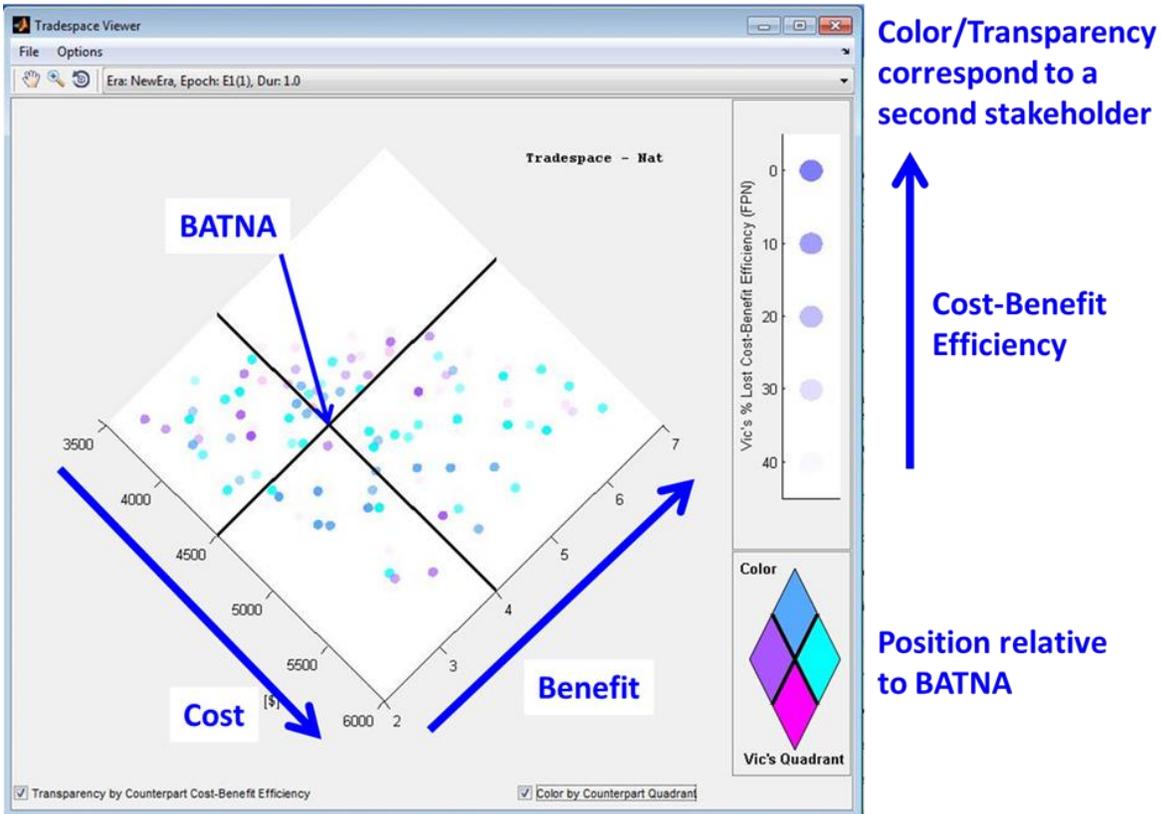


Figure 4-1: The standard benefit-cost tradespace view (for Nat) given to the control group



**Figure 4-2: The negotiation tradespace view (for Nat) given to the treatment group, annotated with the key differences from the standard tradespace<sup>6</sup>**

The other VisLab tools, comprising basic TSE functionality, that were available to the subjects were:

- *Favorites Manager* – Stores designs of interest as ‘favorites’ and plots them in the tradespace with a special marker that is customizable in shape, color, and size. Favorites can also be batches of multiple designs.
- *Filter Tool* – Allows for identification of batches of designs sharing one or more logical properties (e.g., cost less than X). Batches may be saved directly as favorites for further analysis in the tradespace viewer.
- *Comparison Tool* – Places specified designs side-by-side in a table displaying all of their associated attributes in the database. Also allows for the specification of a ‘baseline’ design, then coloring the table entries for other designs based on whether they are higher or lower than the baseline.

<sup>6</sup> Note that the coloring scheme has changed from our mockup in the previous chapter. This was for two reasons: (1) using red and green would render the graph illegible to any red-green colorblind individuals, and (2) we did not want to imply that quadrant I (previously, yellow) was “better” than quadrant III (orange), as might be heuristically assumed with a green-yellow-orange-red color spectrum, which could result in unintentional bias in the agreements.

Subjects were allowed to utilize the available tools and communicate with each other in any way they wanted for thirty minutes, at which point, if they were not finished already, they were prompted that they had ten more minutes to either agree on a design choice or choose to leave the negotiation if they felt that they would be better off with their BATNAs. During the experiment, observations of the verbalized interaction between the subjects, including tone, offers, counteroffers, and final decision, were recorded along with the times at which they occurred.

#### **4.4.3 Post-task**

Immediately after concluding the MSTSE task, the subjects were given a short questionnaire, which was developed and revised through pilot testing on four groups with varying familiarity with the VisLab software and TSE. The questionnaire began with some simple demographic questions, asking for potentially pertinent information such as age, gender, education, and previous exposure to tradespace exploration.

The second part of the questionnaire was designed to capture the participants' impression of the problem (particularly the number of mutually beneficial and mutually acceptable solutions), their partner (in attitude and cooperativeness), and the position for success that they were placed in (with regards to tools and understanding of their own and their partner's problem). These questions and the specific concepts they were designed to target are discussed in Appendix A. Questions #1-4 were short answer questions, while Questions #5-25 were multiple-choice, utilizing a 7-point Likert-type scale. An open response field for generic feedback on the task was also included at the end, as user comments had the potential to inform future refinement of the computer interface or to identify an unforeseen complication and allow mid-experiment correction in order to save resources.

The questionnaire was followed by a short debriefing of the purpose of the experiment, including a document explaining the concepts of framing and its application to tradespace exploration. Subjects were allowed to ask any questions about the experiment or leave at their leisure.

### **4.5 Analysis**

The following subsections will discuss the results of the MSTSE experiment with respect to the original experiment questions. As this experiment was largely exploratory, the direction of the analysis was partially emergent based on both the overall research agenda and the most salient sources of information that stemmed from the task.

#### **4.5.1 Demographics**

Table 4-1 lists a summary of the reported demographic information of the participants. Note that subjects were allowed to report multiple majors as they felt necessary, so they do not add to N=26, but each participant was verified to be in an engineering discipline even if cross

registered elsewhere. Given the goals of the research, the only potential concern raised by these demographics is the possibility that the subjects who performed their own research or development in design technique could be inappropriate partners for less experienced subjects. Further research could investigate the ramifications of unequal design familiarity on MSTSE.

**Table 4-1: Demographics summary**

<b>Age</b>	<b>Number of Responses</b>
<i>19-21</i>	10
<i>22-24</i>	8
<i>25-27</i>	1
<i>28+</i>	7
<b>Gender</b>	
<i>Male</i>	18
<i>Female</i>	8
<b>Current or Highest Degree</b>	
<i>BS</i>	13
<i>SM</i>	7
<i>PhD</i>	6
<b>Major</b>	
<i>Mechanical Engineering</i>	9
<i>Aeronautical and Astronautical Engineering</i>	2
<i>Computer Science and Electrical Engineering</i>	6
<i>Biological Engineering</i>	2
<i>Physics</i>	1
<i>Technology and Policy</i>	2
<i>Engineering Systems</i>	1
<i>No Response</i>	5
<b>Tradespace Exploration Experience</b>	
<i>Yes</i>	4
<i>No</i>	22
<b>Design Experience</b>	
<i>None</i>	2
<i>Novice</i>	9
<i>Intermediate</i>	9
<i>Expert</i>	0
<i>Research/Developer</i>	6

#### 4.5.2 Questionnaire – Multiple Choice

The closed-form part of the questionnaire that the subjects filled out upon completing the task proved to be the least interesting source of data for the experiment, as the small number of samples and subtle difference between the control and treatment groups limited the ability to draw statistically significant inferences from the data. Of the twenty-one questions answered on

a Likert-type scale, none showed a difference between the control and treatment groups significant at the  $p < 0.05$  level for a two-sided Student’s t-test. Testing for differences between the two roles in the case resulted in only one significant difference, corresponding roughly to the anticipated rate of spurious correlations. The p-values for the t-tests on each question are included in Table 4-2.

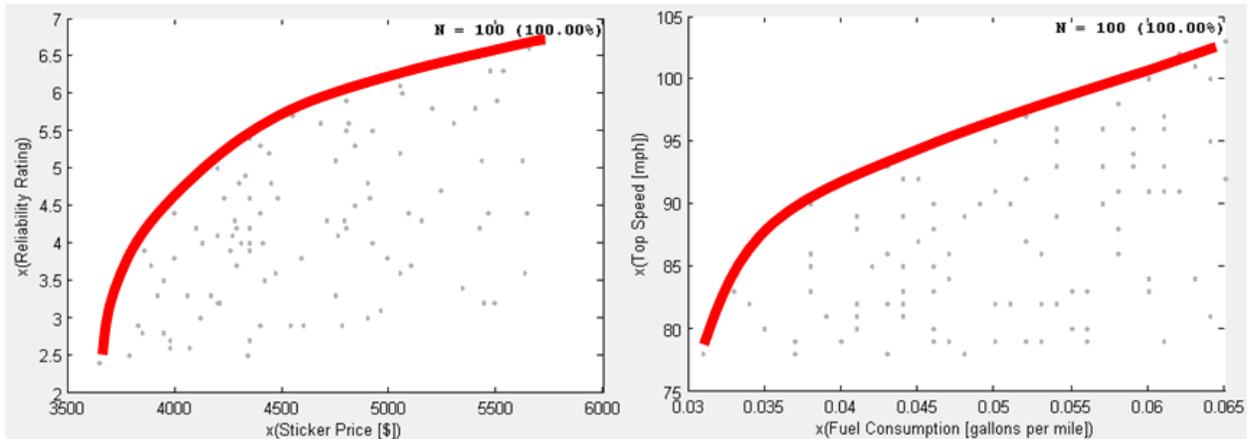
**Table 4-2: Questionnaire between groups t-test p-values**

<b>Question #</b>	<b>5</b>	<b>6</b>	<b>7</b>	<b>8</b>	<b>9</b>	<b>10</b>	<b>11</b>	<b>12</b>	<b>13</b>	<b>14</b>	
<b>Roles p-value</b>	0.73	0.41	0.87	0.53	0.87	0.59	0.59	0.91	0.87	0.41	
<b>Control-Treatment p-value</b>	0.80	0.98	0.57	0.44	0.46	0.25	0.95	0.92	0.58	0.44	
<b>Question #</b>	<b>15</b>	<b>16</b>	<b>17</b>	<b>18</b>	<b>19</b>	<b>20</b>	<b>21</b>	<b>22</b>	<b>23</b>	<b>24</b>	<b>25</b>
<b>Roles p-value</b>	0.92	0.64	0.03	0.12	0.13	0.09	0.55	0.17	0.40	0.26	0.17
<b>Control-Treatment p-value</b>	0.60	0.43	0.34	0.76	0.78	0.19	0.92	0.61	0.28	0.24	0.28

Given the lack of clear separation between the control and treatment groups, more detailed analysis of these results, including indexing the questions, was forgone in favor of focusing on the open-response forms of the questionnaire and the interesting intersections between the responses and the corresponding tasks. It is worth noting that despite the inconclusiveness of these results, a majority of the questions displayed the anticipated *directionality* of difference between the control and treatment group, including all of the questions in the blocks on “understanding the problem”, “problem difficulty”, and “tools satisfaction”. This lends credence to the belief that a larger sample size (or more exaggerated difference between the groups) could yield positive results in a future experiment.

### 4.5.3 Questionnaire – Short Answer and Open Response

The open response fields of the questionnaire provided a variety of different insights into the mindset and experience of the participants. Marginally statistically significant differences ( $p=0.077$ ) were observed between the two roles of the case for Question #1, which asked the participants to estimate how many alternatives were preferable to their BATNA. In this case, Vic reported more designs preferable to his BATNA than Nat. The case was intended to be as symmetric as possible without appearing trivial or manufactured, so this result is somewhat surprising. The benefit-cost tradespaces of the two roles had similar numbers of designs in each quadrant, so a difference between the two is likely an effect of the difference in *shape* of the Pareto front (as the first or second most prominent reference point, depending on the treatment of the group). Nat’s tradespace has a slightly more pronounced “knee” in the curve of Quadrant 2, which may be causing this outcome. Further research into the effect of tradespace shape on stakeholder perception is warranted before conclusions are drawn on this matter.



**Figure 4-3: Slightly more pronounced "knee" to the Pareto front on Nat's tradespace (left) than Vic's (right)**

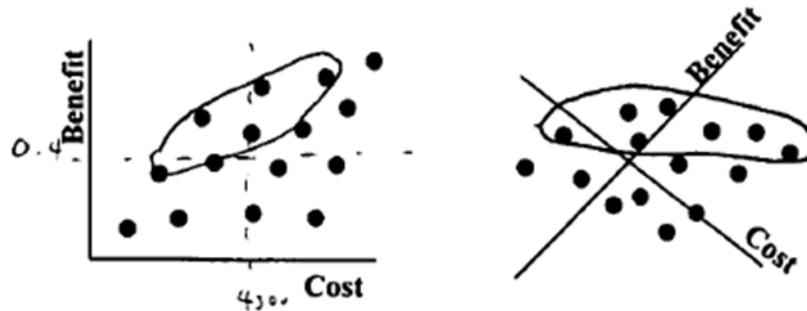
Question #3 asked the subjects to draw the region of the tradespace that they preferred to their BATNA on a supplied set of benefit-cost axes. Generally, it appears from the responses that the treatment groups had a better grasp of the BATNA, with twelve of thirteen responses indicating that the entirety of Quadrant 2 (containing all the strictly superior car choices) was preferable to the BATNA. Table 4-3 summarizes the results of this question by categorizing responses in terms of the areas highlighted. Note that one response in the treatment group was illegible and omitted from this summary.

**Table 4-3: Preferred Regions of Tradespace**

Condition	Control	Treatment
<b>Number of Responses</b>	12	13
<b>Areas Highlighted</b>		
<i>Quadrant 1</i>	3	0
<i>Quadrant 2</i>	6	12
<i>Quadrant 3</i>	0	0
<i>Quadrant 1 (Pareto Front Only)</i>	4	4
<i>Quadrant 2 (Pareto Front Only)</i>	3	0
<i>Quadrant 3 (Pareto Front Only)</i>	3	6
<i>Other</i>	1	1

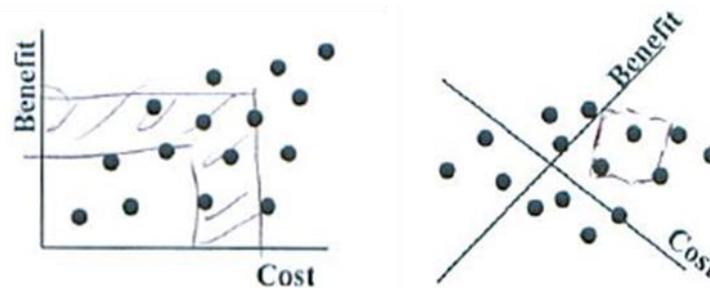
The preferred-regions question is very effective at evaluating the mindset of the subjects, because their responses can be compared to a canonically "rational" response. The "rational" response to this question would include all of Quadrant 2, where alternatives are strictly superior to the BATNA, and optionally the Pareto Front (and near-front) designs of Quadrants 1 and 3, where the tradeoff between cost and benefit is highly favorable. This pattern was indicated by five control subjects and twelve treatment subjects, which suggests a significant difference ( $p=0.0095$ ) between the control and treatment groups using Fisher's exact test. A typical

“rational” response is shown in Figure 4-4 for both the control and treatment groups (note that the BATNA axes were hand-drawn by the subject to describe their answer on the control side).



**Figure 4-4: Example "rational" responses on the control (left) and treatment (right) questionnaires**

The “other” responses, shown in Figure 4-5, are also interesting cases. The “other” from the control group displays considerable confusion over the relationship between benefit and cost, with a “band” of acceptableness shaped roughly like a ‘7’ sweeping out an arc in Quadrants 2, 1, and 4. The best-case scenario for the subject who submitted this answer was that they both (1) mirrored the cost axis and (2) misinterpreted the question and attempted to show their willingness to compromise by stepping back from the x-axis mirrored Pareto front into a more central location. During the task, this subject had considerable trouble utilizing filters to capture the idea of the BATNA, which apparently manifested again in the questionnaire here. The “other” from the treatment group, in contrast, indicated only a small region of Quadrant 1, reflecting a very strong preference for a specific region of the tradespace (though, again, possibly misinterpreting the question as to indicate where the best final solutions were actually located).



**Figure 4-5: "Other" responses, displaying benefit-cost confusion**

The results for Question #4, in which the subjects were asked to briefly describe the process they used to find a solution, proved to be relatively unenlightening. The recommended 2-3 sentences were not enough to cause the subjects to elaborate, with the vast majority of subjects simply reporting some variation of (1) finding good designs, (2) comparing the designs with their partner, and (3) selecting a choice acceptable to both parties. The observational data supplied much more useful and complete information than this question on the same topic.

Finally, subjects were also given an open response field at the end of the questionnaire to provide any comments that they wished to share. Seventeen of the twenty-six subjects provided feedback, with the majority of those comments comprised of positive remarks on the software or tradespace exploration and suggestions for features that could improve the experience. Following are a few choice quotes illustrating common responses and feature requests:

### **Control subjects:**

“Great tool - quantifying a BATNA and filtering options based on mine & my partner's parameters = super useful. I will try to use this in future decisions. Because I think it's so good at visually demonstrating what's decent for both parties.”

“It was relatively simple and straight forward all around. The only real difficulty I had was the sensitivity of the mouse when clicking a point. I think it would be useful to drag and highlight sections of the graph if possible.”

“It was a great interface that I'd love to be able to use in real life for similar things. Feature to make it easier such as select all or change all/edit selection or something like that could make it easier.”

“Making the dots on graphics larger (or larger on hover) would have made them easier to click. I feel like we didn't use the tool to its fullest capabilities. Having a mass-select option, and different tiers of favorites would have been useful”

### **Treatment subjects:**

“Giving feedback on solution Pareto Optimal of the 4 traits would be helpful (feel satisfying). The software is a bit cumbersome (lots of clicking) but does display good info.”

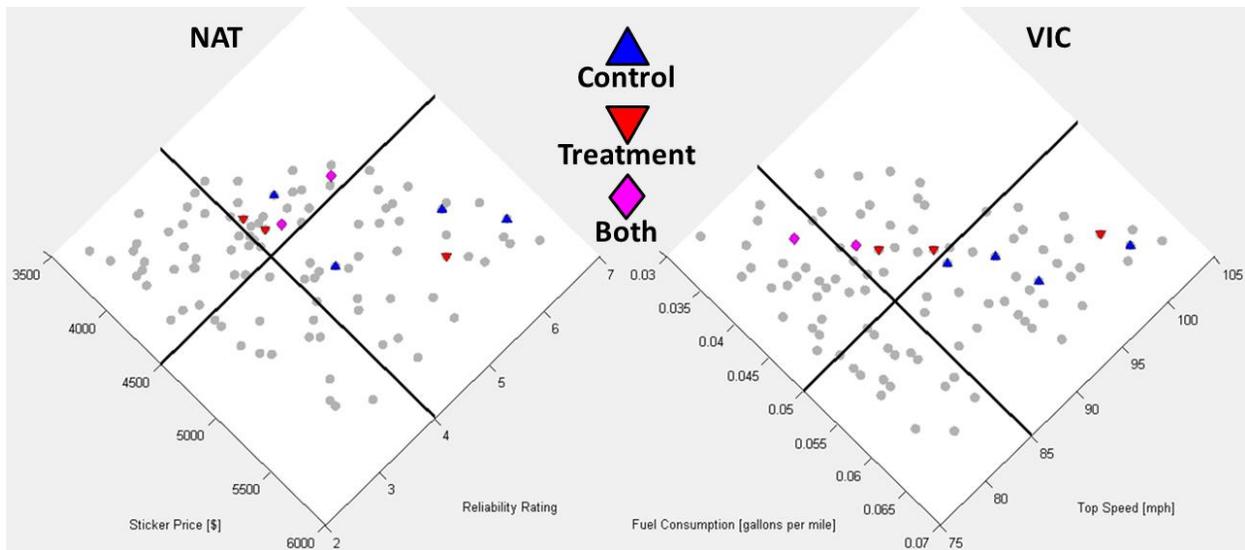
“I don't understand how we're supposed to selectively disclose information + negotiate if both of us can see each others' screens + preferences so easily + openly. I didn't really feel there was much to talk about since (rather unrealistically) my partner and I could see each other's benefits, costs, preferences, etc just by turning around and talking to each other.”

That two treatment subject comments are particularly interesting. The first comment demonstrates that the appeal of the Pareto front as a reference point persists even with the modified tradespace. This subject was one who indicated having prior experience with TSE before the experiment, thus this response may reflect the (anticipated) impact of macro framing based on familiarity with existing TSE concepts, “anchoring” them on those concepts even in a new negotiation setting. The second comment shows the subject's desire to bargain strategically by withholding information. Strategic behavior is certainly something that, when performed correctly, can improve individual outcomes in a negotiation setting; however, the setup for this experiment was designed explicitly to appeal to the Full, Open, and Truthful Exchange principle

for productive mutual bargaining. In certain applications, hidden information may be necessary or otherwise desirable, and those will likely require different visualizations and room setups in order to be most effective. This example, on the other hand, demonstrates that FOTE can be a “tough sell” to aggressive negotiators, and may require additional emphasis in the early phases of MSTSE to manipulate macro framing and promote stakeholder buy-in – changing the visualization is not sufficient on its own.

#### 4.5.4 Offer History and Outcomes

The final decisions of the groups are highlighted in the treatment tradespace view in Figure 4-6, with control-only selections in blue, treatment-only selections in red, and both-conditions selections in magenta. All thirteen trials ended in a successful agreement. Moreover, no statistically significant differences were found between the control and treatment groups in either the time in which agreement was reached, nor in solution quality as measured by the average and worst Fuzzy Pareto Number (FPN) of the chosen car, suggesting that (as expected) the problem was not difficult enough to cause groups to miss the “good” solutions in either condition. The final agreements of each trial and associated information are shown in Table 4-4.



**Figure 4-6: Nat (left) and Vic (right) tradespaces, with final selections marked**

**Table 4-4: Final Agreement Information**

Trial Number	Condition	Final Agreement Design ID	Time of Agreement (min)	Average FPN	Worst FPN	Both in Quad 2	FPN Minimax Solution	Modal Solution (Quad 2 minimax)	Num. Stakeholders in Quad 1
1	Control	34	13	6.8	12.0				1
3	Control	44	29	6.0	12.0				2
4	Control	42	29	5.6	8.8			X	0
6	Control	56	35	14.7	21.4				2
8	Control	26	23	1.7	2.9		X		2
12	Control	11	24	12.4	14.7	X			0
2	Treatment	42	29	5.6	8.8			X	0
5	Treatment	80	12	11.1	14.7	X			0
7	Treatment	11	33	12.4	14.7	X			0
9	Treatment	86	31	9.6	12.0	X			0
10	Treatment	42	31	5.6	8.8			X	0
11	Treatment	62	31	6.0	11.9				2
13	Treatment	80	14	11.1	14.7	X			0

Despite the lack of statistically significant conclusions to be pulled from the time and efficiency data, there are nevertheless multiple interesting features of the chosen solutions. Design IDs 11, 80, and 86 are the only three choices in this problem that occupy Quadrant 2 for both stakeholders, making them obvious choices for focal points of discussion as strictly-superior solutions to the BATNA for both participants. Four of seven treatment trials chose one of these designs, compared to only one of six control trials, suggesting a greater emphasis on this type of hill-climbing solution in the treatment group. The FPN minimax solution, Design 26, is the solution most likely to be obtained by a standard multi-objective optimization but was selected by only one group. The count of solutions in each quadrant is nine in Quadrant 1, fourteen in Quadrant 2, but only three in Quadrant 3, suggesting a bias in favor of higher-cost, higher-benefit tradeoffs than the inverse. In fact, this may understate that bias given that the three Quadrant 3 solutions were three separate trials that chose Design 42, which was the sixth-best FPN minimax solution and the best FPN minimax solution *with at least one stakeholder in Quadrant 2*, where most attention was directed. Very little other attention was paid to Quadrant 3, possibly indicating that stakeholders are more inclined to improve benefits than they are costs: future research could explore this in more detail, but it matches roughly with the observation of bias towards “gold plated” solutions in the real world, where money is thrown at a problem in order to get every stakeholder an increase in value.

More generally, over the course of the task, definite patterns in the designs offered by one subject to the other emerged, differentiating the control and treatment groups. As expected, control teams made more offers on or near the Pareto front, typically moving farther and farther off the front as time passed until an agreement was reached. Treatment teams were more likely

to begin near the center with the BATNA and the three dual-Quadrant-2 choices, and spiral out as necessary to test other options and look for more favorable tradeoffs. Figure 4-7 shows the trajectory of offers made by Nat on his tradespace for all thirteen trials as of minute 26, with control trials on the left and treatment trials on the right. The control tradespaces appear to have a considerably more exhaustive “search and destroy” pattern that spans the Pareto front, while most of the treatment tradespaces are clustered more heavily in the center of Quadrant 2. The same patterns hold when looking at Vic’s offers in Figure 4-8, though this time there is one highly clustered control subject and a few subjects that never attempted to reach the true Pareto front but rather started one “layer” of dominance in. The time chosen for the snapshot of each of these figures was simply the one that showed the pattern most clearly.

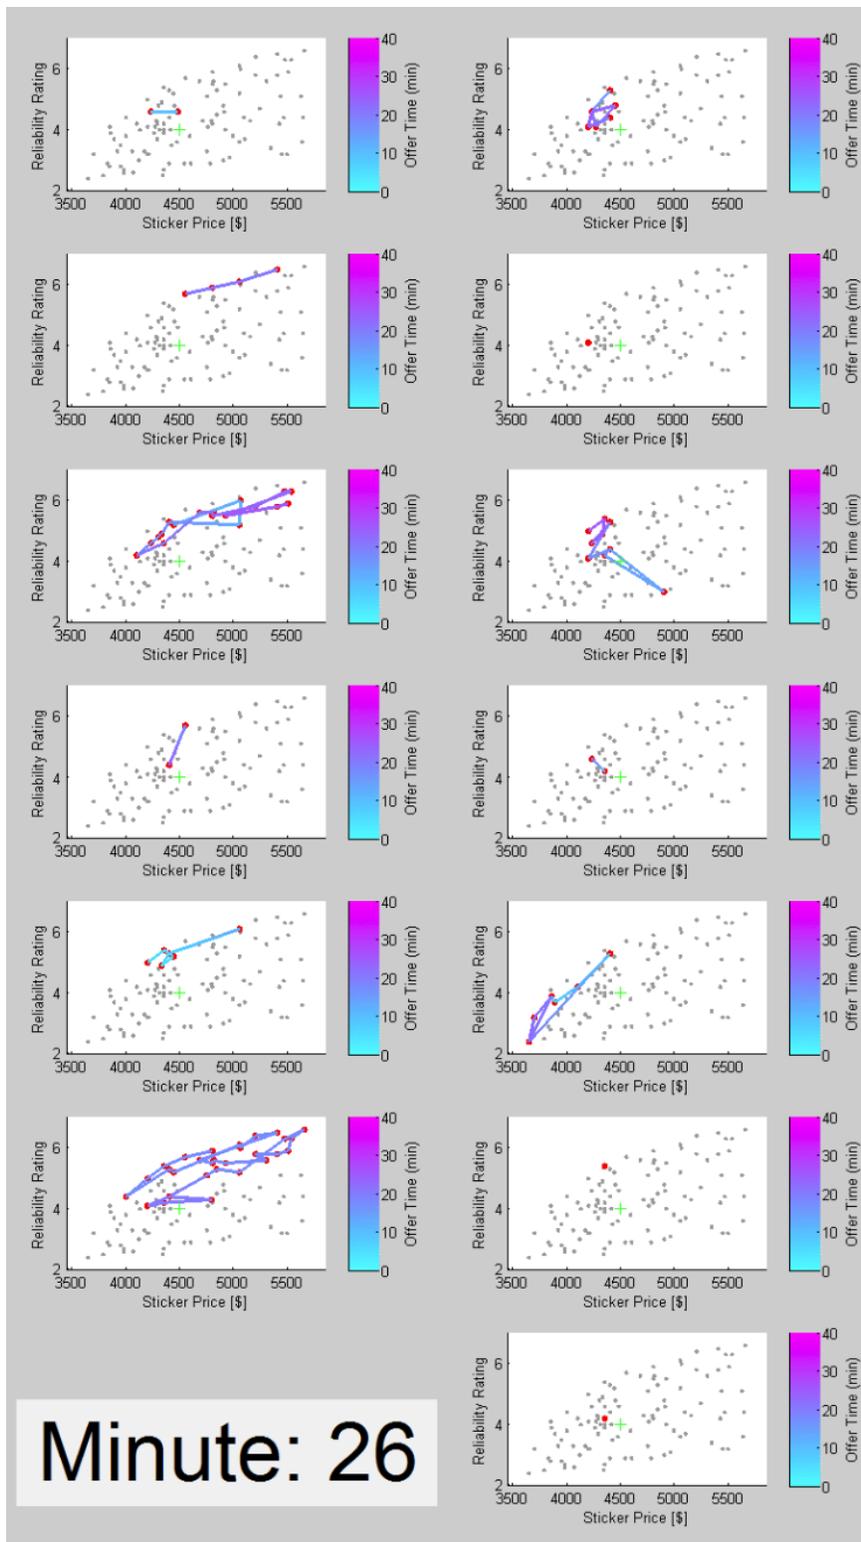


Figure 4-7: Nat's offer trajectories at minute 26 (control on left, treatment on right)

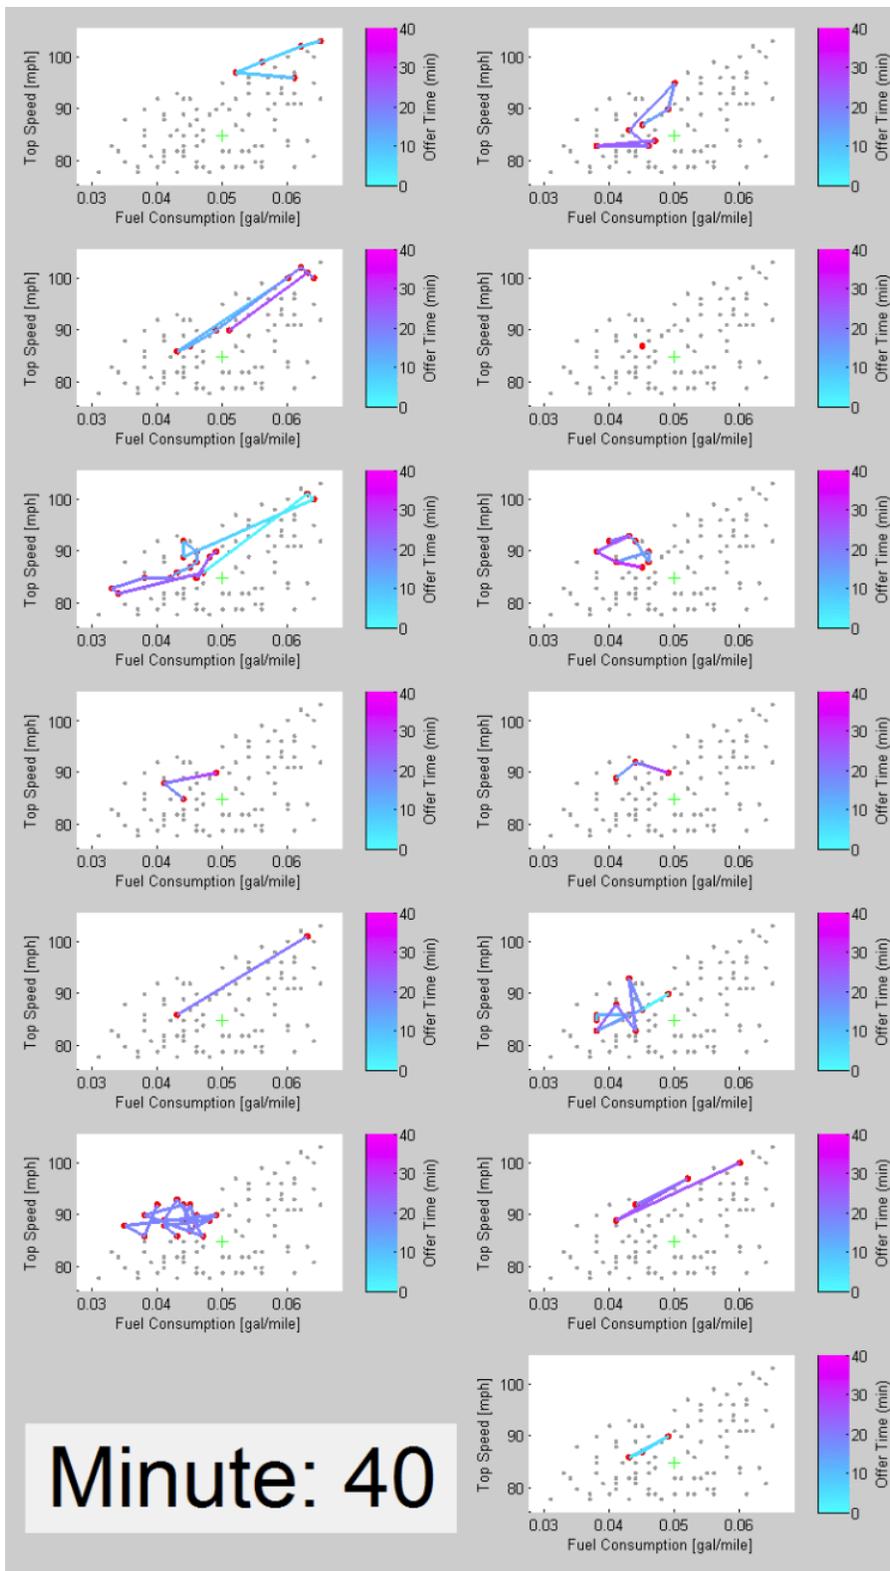


Figure 4-8: Vic's offer trajectories at minute 40 (control on left, treatment on right)

Figure 4-9 shows a plot of all offers in all trials by their FPN (for the offeror) and the time offered, separated by the treatment. The resulting least-squares linear regression of the two sets shows that each group is moving approximately one-quarter of a percent of cost-benefit efficiency per minute, but the control group becomes less efficient over time while the treatment group becomes more efficient. Both of these trends are significantly different from zero ( $p < 0.01$ ). This confirms the hypothesis that the control group generally works in the losses frame (from the Pareto front, inwards) and the treatment group in the gains frame (from the BATNA, outward). Note that the intercept of the best-fit line for the treatment group is at approximately FPN = 12, which is very near to the FPNs of the BATNAs. The gradually worse performance of the control group is also theoretically likely to increase the stress of the decision over time and decrease satisfaction in the final solution.

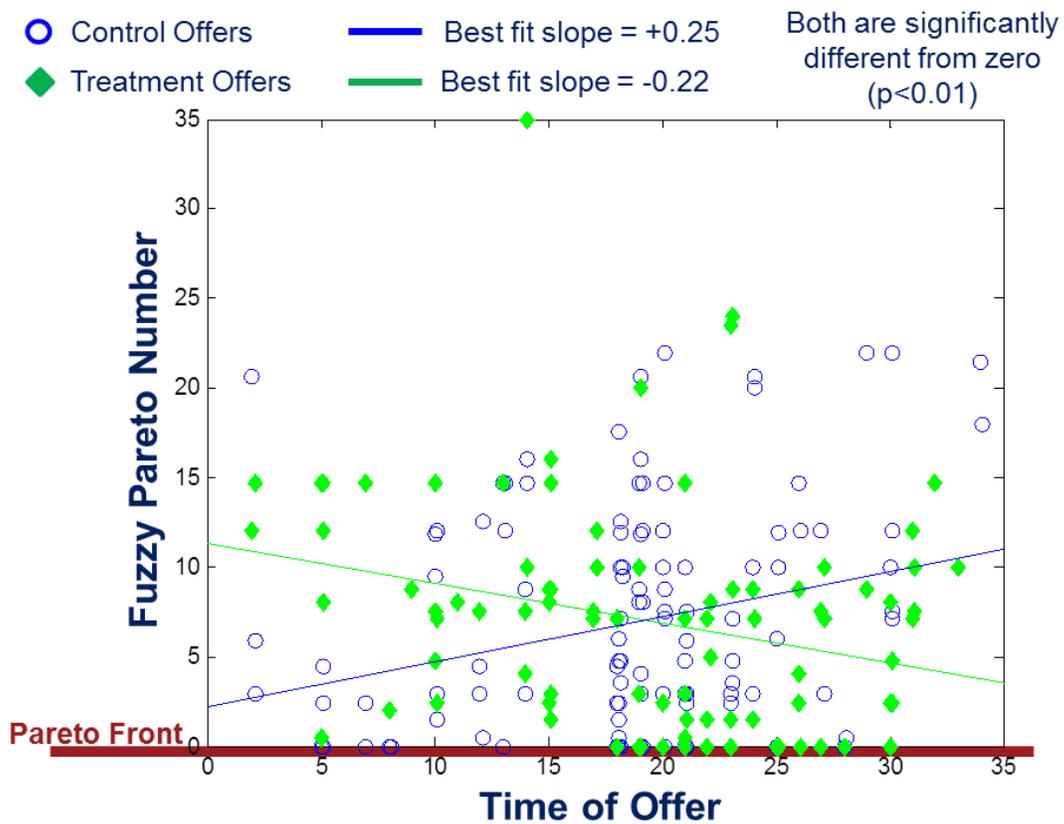


Figure 4-9: All offers by FPN and time

Unsurprisingly, the control and treatment groups were not strictly separated in terms of their preference for Pareto-focused versus Quadrant 2-focused exploration; rather, most trials mixed and matched some of both techniques. The overall trend is one that moves away from Pareto-focus after switching to the treatment, confirming the hypothesized effects of the new visualization. However, this does appear to come at a cost of a more clustered, local investigation of Quadrant 2 and less exploration of the tradeoff Quadrants 1 and 3, which runs

counter to the overarching goals of MSTSE and tradespace *exploration* in general. While a modified tradespace view like the treatment in this experiment may successfully un-anchor stakeholders from their Pareto front, application of the technique may require additional emphasis on tradeoffs and exploration in order to encourage stakeholders to adequately consider all the available design alternatives.

#### 4.5.5 Task Observation

Observational records of the on-screen behavior and verbal communication of the subjects were coded, to see if differences in negotiating behavior were noticeable between the control and treatment groups. The codes were:

- Pareto front focus
- Quadrant 2 focus
- Confusion over costs/benefits
- Discussion of BATNA
- Discussion of preferred tradeoffs
- Discussion of fairness
- Creation of a tentative agreement
- Working individually
- Positional Bargaining (back-and-forth)
- Appeal “outside the case”
- Exhaustive “search and destroy”
- Pressure for concession/agreement
- Treating problem like an optimization (maximize/minimize)
- Negativity about prospects of success
- Use of Filter Tool
- Use of Comparison Tool
- Use of a defined color/shaping scheme in Favorites Manager

Using a rank-sum test on the number of observed instances in each trial, significant differences ( $p < 0.05$ ) between control and treatment were found in the codes of Pareto front focus, Quadrant 2 focus, and negativity about prospects of success. Additionally, there was a significant difference between the Pareto focus and Quadrant 2 focus *within* both the control and treatment groups. This corroborates the observations made on the sequences of offers, with the control group focusing more on the Pareto front and the treatment group focusing more on Quadrant 2. Additionally, all expressions of negativity were limited to four control trials, with none in the treatment trials. This code was intended to capture statements similar to “This isn’t going to help; none of these will work for you,” as expressed by a subject during Trial 12, vocalizing a defeatist mentality towards the prospect of a productive agreement. This supports

the hypothesis that operating in the “losses” domain off of the Pareto front creates a negative environment for negotiation.

A few other codes are interesting enough to remark on despite have too small a sample size to reach statistical significance. Confusion over the directionality of costs and benefits occurred in four of six control trials, but only two of seven treatment trials. It is possible that a greater emphasis on the BATNA early in the MSTSE process forces stakeholders to engage and understand those concepts better up front. The same split was observed for viewing the problem as an optimization, by explicitly engaging in a maximize-benefit or minimize-cost activity for some amount of time. This can be considered to go hand-in-hand with a Pareto front focus, which is essentially a multidimensional optimization problem, but again can detract from understanding the BATNA as an appropriate reference point. Finally, treatment teams made eight appeals “outside the case” to only one for the control teams. This suggests that the treatment group was more inclined to view the problem creatively, citing information not included in the case description (most often gas prices and speed limits in this case) as justifications for making tradeoffs. This type of creative behavior is valuable in tradespace exploration, which is frequently iterative and supports the incremental improvement of models and assumptions, and it is encouraging to see that the treatment visualization appears to spark this mindset.

#### **4.5.6 Interesting Anecdotes**

One feature of discussion that arose multiple times was that of a sort of ‘activation energy’ required to make a decision. Different teams, in both conditions, had discussions about whether or not the cars in Quadrant 2 for both people (11, 80, and 86) were worth picking at all, despite being numerically superior to the BATNAs. To some extent this is unsurprising, since these three designs were deliberately set to be only *slightly* better than the BATNAs so that they would not be obviously superior to the other potential choices and thus make the task trivially easy. Despite this, the idea that working alone and taking the BATNA could be superior to these choices (usually citing simplicity as a rationale) was not intended, yet frequently mentioned. Teams 3, 6, 9 and 10 all had this discussion, and only Team 9 ended up choosing one of those designs. Often, the designs were talked down as “barely better” or with phrases like “I can’t really see the difference”, and accompanied by a plan to further explore Quadrants 1 and 3 in search of better tradeoffs. This may be an example of the idea of just-noticeable differences, which has recently been extended into the field of visual analytics (Harrison et al., 2014). This phenomenon may also be correlated with the observed tendency for the control group to explore Quadrants 1 and 3 more than the treatment group: the re-centered tradespace axes may make the differences more noticeable and thus (in this example) make the designs look better by comparison to the BATNA. Further research into the impact of tradespace visualizations on the perceived differentiation of alternatives could improve the display of information needed to make tradeoffs effectively in MSTSE.

Team 5 was fastest to reach agreement at only minute 12, which actually understates their true speed, as they had chosen their design at minute 8 and spent four minutes doing back-of-the-envelope estimates of driving distances and costs. During their debrief, Vic said that “this would have been a lot more difficult if there were no cars [that were in Quadrant 2 for both people]”, which Nat quickly agreed with, suggesting that they only spent the last four minutes because it felt odd to finish so quickly. The workflow of the team went from (1) filter on the four BATNA criteria, finding designs 11, 80, and 86, (2) look at the tradespace with the three designs highlighted, (3) bringing the three designs up in the comparison tool, and (4) choosing design 80 as the “compromise” between 11 and 86. No effort was made to investigate tradeoffs, as both subjects interpreted the task directly as “improve all attributes”. This is an excellent example of the impact of macro framing. The effect of having the negotiators agree completely on the underlying purpose and decision criteria of the negotiation led directly into a speedy agreement that, by their own admission, seemed trivially easy. This outlook on negotiation is something that can potentially be influenced by macro framing, implying that the setup phases of MSTSE should be leveraged to align the stakeholders’ understanding on high-level negotiation objectives.

Team 6, in the control group, exemplified the desire for negotiators to move beyond the traditional tradespace view. Early on, they acknowledged that comparing “regions” of the tradespace was bound to be ineffective when the axes upon which each person’s tradespace were plotted did not represent the same dimensions. Desperate to find more concrete ground upon which to negotiate, Vic suggested they combine and view both tradespaces on hybrid axes composed of Nat’s benefit and Vic’s cost (their chosen “preferred” axis). A large portion of their time (from minute 10 to minute 18) was spent using these axes as a stand-in for the ability to accurately portray some measure of counterpart satisfaction on personal benefit-cost axes. This also drove a considerable amount of confusion related to the BATNA, eventually requiring them to have another discussion at minute 23, more than halfway through their time, on what exactly it means to find a “better” solution to the problem. A more inclusive, group-oriented visualization like the treatment group’s tradespace view would likely have alleviated some of this confusion by allowing full consideration of the problem and not requiring certain attributes (like Nat’s cost and Vic’s benefit) to be cast aside, at least temporarily, in the name of simplicity.

On a related note, early communication was a common feature in most of the trials, with nine of the thirteen teams beginning to work together within three minutes of starting, typically by discussing attributes and/or BATNAs of both parties. Clear and open communication is desirable for productive mutual negotiation, so this is a positive observation that at the very least MSTSE does not appear to *encourage* people to work alone. Only Team 11 had a large delay before beginning to work together, choosing to spend twenty-one minutes working alone at the beginning, before speeding up to reach an agreement only ten minutes later, in a noticeable rush. During the debrief, Nat volunteered, unprompted, that he wished they had communicated earlier instead of waiting to “split the difference”, indicating that he seemed aware that hurrying had

forced the team to fall into the “midpoint trap” of positional bargaining. This further reinforces the importance of early communication, but is perhaps also another point in favor of tradespace exploration as a negotiation tool given that at least the subject was aware of the problem. Similarly, during the debrief for Team 13, Vic asked how focusing on designs near the Pareto front could be detrimental and Nat interrupted with a reasonable answer, correctly remarking that if the designs on the front are not likely to be accepted by the other party then spending time on them would be wasteful and possibly stressful. It can be considered another positive mark for MSTSE that a simple task such as this experiment could convey an unintuitive concept like that to an inexperienced participant.

Usage of color and transparency amongst the teams in the treatment condition varied, as there were no requirements on when, how, or even if those features should be used. The most common strategy for using them was to leave them off for general exploration, toggling them on only when looking to get information on the other stakeholder for a specific design or group of designs. Only Team 10 chose to leave both color and transparency on for the entire task, and they also displayed the strongest un-anchoring response to the Pareto front. After spending a considerable amount of time discussing color, including finding the 11/80/86 trifacta visually as the only Quadrant 2 and cyan-colored design points, discussion turned to creating more value than those designs could provide. In addition to both participants verbalizing a strong understanding of the different types of tradeoffs associated with each quadrant, Nat also on multiple occasions referred to designs as “on my Pareto [front]” that were simply *near* the front and dominated only by solutions unacceptable to Vic. Nat eventually traced a curve over the tradespace with his mouse, calling it “my Pareto”, while ignoring many of the more transparent and/or purple (Quadrant 4) designs, demonstrating a lack of fixation on the individually valuable but group-infeasible designs on the true Pareto front. It is possible that if the color and transparency settings were set to always-on for all treatment trials, other teams would have followed steps similar to Team 10 and displayed even less fixation on the Pareto front than they did with the toggle ability.

## **4.6 Discussion**

The purpose of this experiment was to explore the use of tradespace exploration in a negotiation setting, seeking evidence to support the working theory of the larger research plan and to uncover unanticipated effects or interactions with the modified visualizations that could be used to give additional direction to the next steps of the research. Even with the limited available sample size, statistically significant differences between the control and treatment groups were shown for (1) focus on Pareto front, (2) focus on Quadrant 2, and (3) ability to “rationally” answer which designs were preferable to the BATNA. Additionally, the treatment group was shown to begin near the BATNA and work “out” into the gains domain, while the control group typically began near the individual Pareto front and work “in” into the losses domain. All of these differences were in line with the working theory and suggest that new visualizations can potentially improve the practice of MSTSE. Considerable other evidence

supporting the theory was presented via qualitative analysis of the final agreements and offer trajectories, in addition to observation of the language and techniques employed by the subjects.

This experiment also brought the need for more work on the concept of macro framing MSTSE into sharp relief. The problem definition and setup for MSTSE have the potential to reframe the types of interactions between stakeholders, particularly with respect to fostering effective early communication of value statements and negotiation objectives. Also of key interest is the possibility of including additional tasks or features intended to support the exploration objectives of MSTSE and investigate high-value tradeoffs, preventing over-committal to strict improvement (Quadrant 2). Other potential future research topics that were identified during this experiment, including just-noticeable differences in TSE, effects of tradespace shape on perceived value, and mechanisms behind high-cost-high-benefit solution biases, may all be investigated at a later date but will not be explored further in this dissertation.

Overall, the negotiation experiment demonstrated many of the key features of the working theory of framing in MSTSE and provided quantitative evidence of framing effects with respect to the Pareto front and BATNA. With this credibility established, the following chapter will discuss efforts to refine and expand the set of visualizations for MSTSE, performed in conjunction with the practicing systems engineers that are most likely to apply them.



## 5 Developing MSTSE Visualizations with Iterative Practitioner Feedback

This chapter will describe this research’s attempts to expand on the results of the initial MSTSE experiment and recommend further improvements to the visualizations and analyses that accompany TSE. In addition to developing new data representations that leverage the MSTSE framework into useful multi-stakeholder insight, the feedback of professional engineers was sought through interactive group interviews. These two tasks were conducted concurrently, incorporating the practical insights of the interviews into the visualizations in order to target realistic tasks of interest to systems engineers.

The following subsections will demonstrate these new visualizations, as well as provide the reasoning behind their development and example insights that can be derived from their use. They have been implemented in the MIT SEAr Interactive Value-driven Tradespace Exploration and Analysis Suite (IVTea), the next-generation version of VisLab as used in the negotiation experiment. The figures used in this section come from sample datasets, chosen to clearly demonstrate the purpose of the given visualization without the need to provide extensive context. Analysis of the specific conclusions able to be drawn from each figure will be limited only to the extent necessary to clarify the intended purpose.

### 5.1 Guiding Principles

Before diving into the creation of new MSTSE visualizations, it is helpful to reiterate some of the key guiding principles of visual analytics and their hypothesized connection to MSTSE, on which we are building the justification for the new data views. Recall the “visual analytics paradigm” of Keim et al. (2008):

- *Analyze first* – Reduce burden on human by performing some analysis in the background
- *Show the important* – Direct attention to the most salient information
- *Zoom/filter and analyze further* – Iterative learning with gradually increasing detail
- *Details on demand* – Allow human to reveal the details originally hidden to reduce complexity when requested

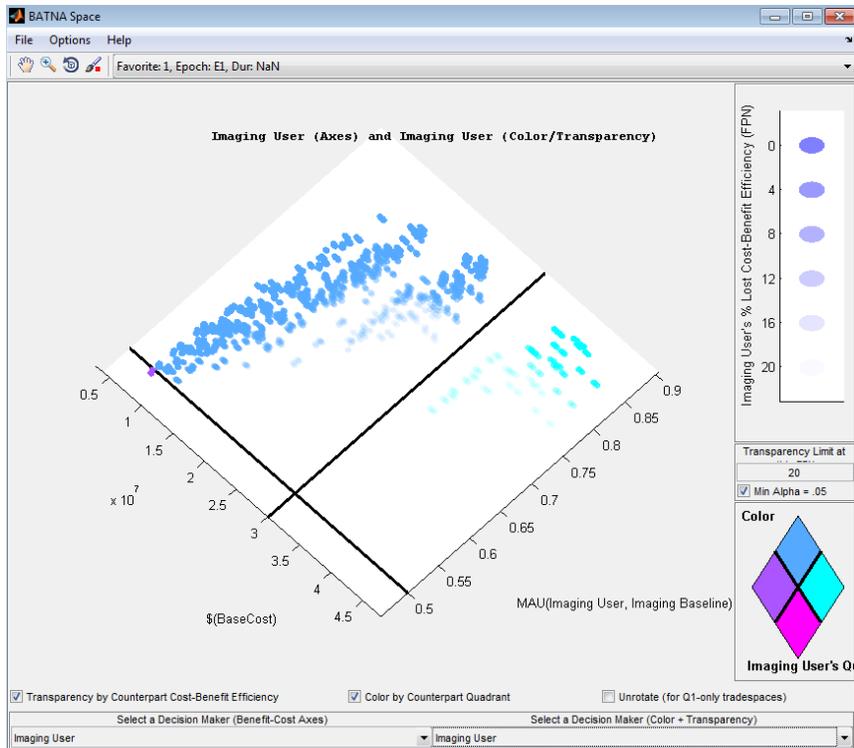
These four tenets are specifically geared towards creating *effective* visualizations: those that are capable of augmenting human ability to reach conclusions and make good decisions based on data. There is a fundamental tension in MSTSE between reducing the cognitive burden on the stakeholder by providing pre-analysis through clear, concise plots (the first and second bullets) and representing the full detail and complexity of a multi-stakeholder tradespace. This tension mandates the need for interactivity (the third and fourth bullets) that allows stakeholders to understand the underlying mechanics behind the high-level insights the visualizations can provide directly.

*Show the important* also ties directly into the concept of availability bias as discussed earlier: because people assume the most easily available information is the most important, we should strive to make the most important information easily available. We have already targeted this principle in the previous section by increasing the availability of information regarding the BATNA (for use as a reference point / decision anchor), but there may be more negotiation-oriented information types that are useful to highlight. In particular, we can use the concept of “question-driven” tradespace exploration (Ross et al., 2010b) to find new, desirable queries for the data in a multi-stakeholder problem. Visualizations that can rapidly communicate insights to key negotiation questions, such as “Which stakeholder am I most likely to agree with?”, without the need for extensive human analysis will improve the efficacy of MSTSE.

## **5.2 Further Improving the Negotiation Tradespace**

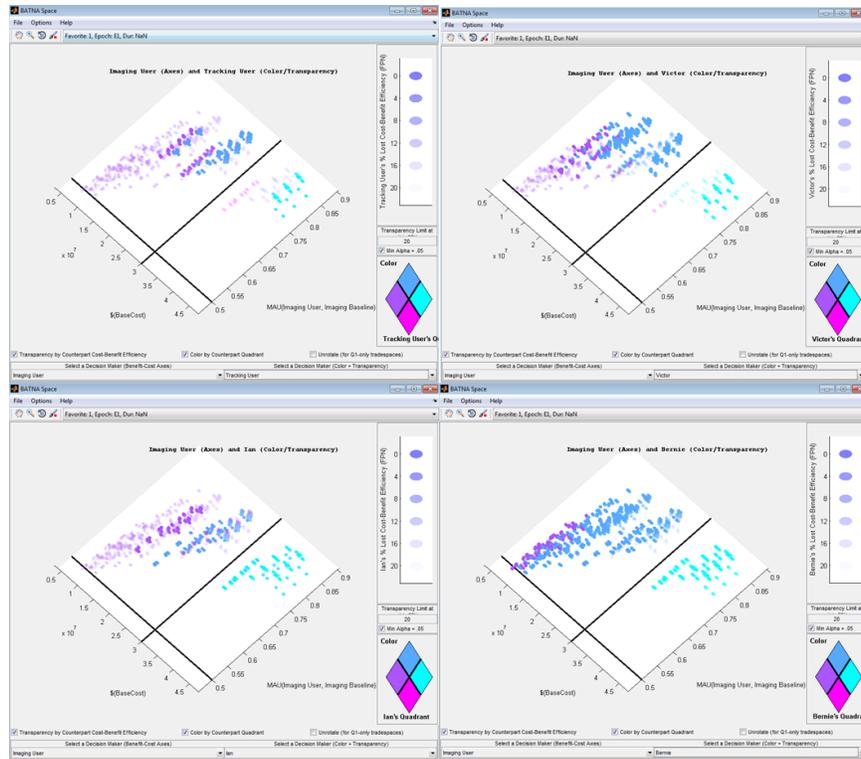
The “negotiation tradespace” – the modified benefit-cost tradespace used in our negotiation experiment – was created deductively from the principles of TSE, negotiation theory, and prospect theory, without significant testing or feedback before deployment. The results of its use in the experiment and application to test datasets highlighted particular areas for improvement, despite its success at improving awareness of negotiation concepts while exploring.

For example, one of the main types of feedback from the experiment subjects was concern over the amount of time it takes to grasp what the color and transparency dimensions are actually showing. This process could take up to 10 minutes, even in the experiment setting with a small set of alternatives and a limited other demands on each participant’s attention. In a real application, this learning curve could present a significant barrier to entry if a participating stakeholder does not have prior experience with MSTSE and is time-constrained. In order to accelerate understanding, the tradespace can also be set to show quadrant and efficiency (via color and transparency) for the *same* stakeholder if necessary, as shown in Figure 5-1.



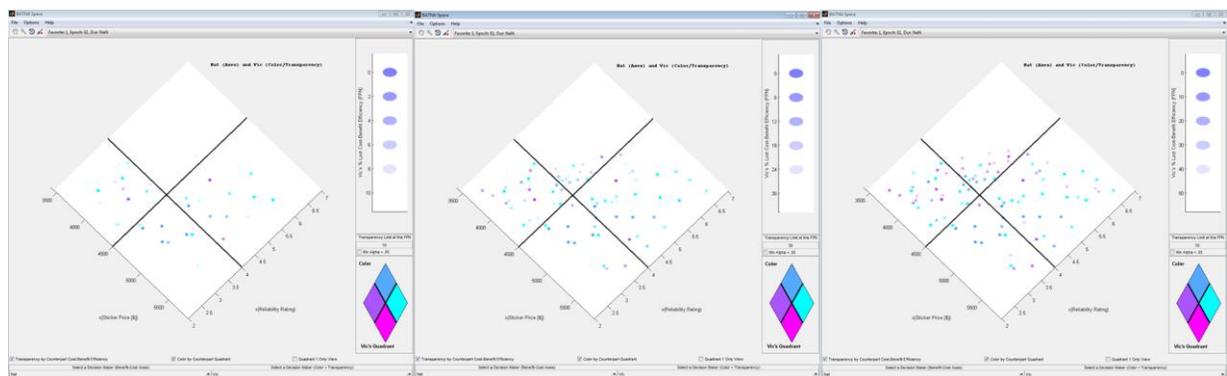
**Figure 5-1: The negotiation tradespace, set to show the same stakeholder for axes and color/transparency**

This view does not achieve the goals of increasing availability of multi-stakeholder information, but can be a useful training tool as it clearly shows design points fading out as they move away from the Pareto front and colored by which quadrant they occupy. Looking at this can communicate the purpose of the added dimensions before subsequently switching those dimensions to target another stakeholder. As discussed in the original description of the tradespace modifications, this representation is not effective at communicating information about more than two stakeholders at a time. However, the implementation of the stakeholder toggle on color/transparency allows for rapid comparison of regions of the tradespace that each stakeholder favors, as shown in Figure 5-2.



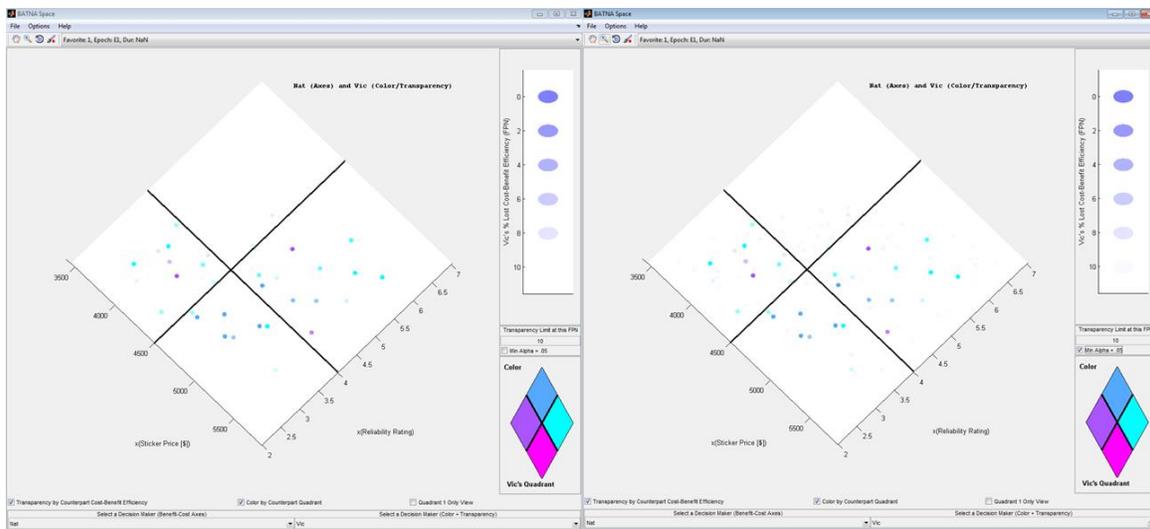
**Figure 5-2: The negotiation tradespace, with color and transparency set to 4 different stakeholders. Darker regions are more opaque and favored by that stakeholder.**

The use of the tradespace on a variety of different datasets also illuminated another outstanding problem: the relevant scope of designs of interest may change from case to case. In the original formulation, design points were fully “faded out” at an FPN of 40, or 40% removed from the Pareto front. For some cases, depending on the enumeration of the dataset and the similarity of the stakeholder preferences, this may be too tight or too loose of a bound: either hiding some potentially interesting alternatives or cluttering the tradespace with irrelevant solutions. In correspondence with the principle of *show the important*, the lower bound of efficiency in the tradespace can now be controlled, demonstrated in Figure 5-3, in order to find a case-appropriate visibility threshold.



**Figure 5-3: Negotiation tradespace with a maximum display FPN of 10, 30, and 50, respectively**

As the maximum display FPN is reduced, fewer design points appear on the plot, increasing attention to the subset of design points *most* efficient for the other stakeholders. The right-most tradespace above, with the highest display FPN, looks similar to a standard single-stakeholder tradespace but includes many highly visible, dominated (magenta) solutions in Quadrant 2, which may distract the negotiators from more productive alternatives. In some negotiations, it may make sense to start with a low maximum FPN and gradually increase the threshold over time in search of alternatives slightly worse for the counterpart stakeholder but slightly better for oneself. However, too low of a maximum FPN runs the risk of portraying the problem too negatively: in Figure 5-3, the left-most tradespace has eliminated nearly every alternative in Quadrant 2 by virtue of its very strict maximum FPN threshold, which could provoke a negative initial impression of the chances of reaching an agreement. Overall, it is important to strike a balance between decluttering the tradespace by focusing on “good” designs and still representing a full space of possible solutions, and the challenge lies in finding the case-specific “sweetspot”. This can be addressed in part by setting a minimum opacity, rendering even designs that do not pass the FPN threshold visible but highly faded, as in Figure 5-4, allowing the use of restrictive FPN thresholds while keeping all *valid* design choices visible.

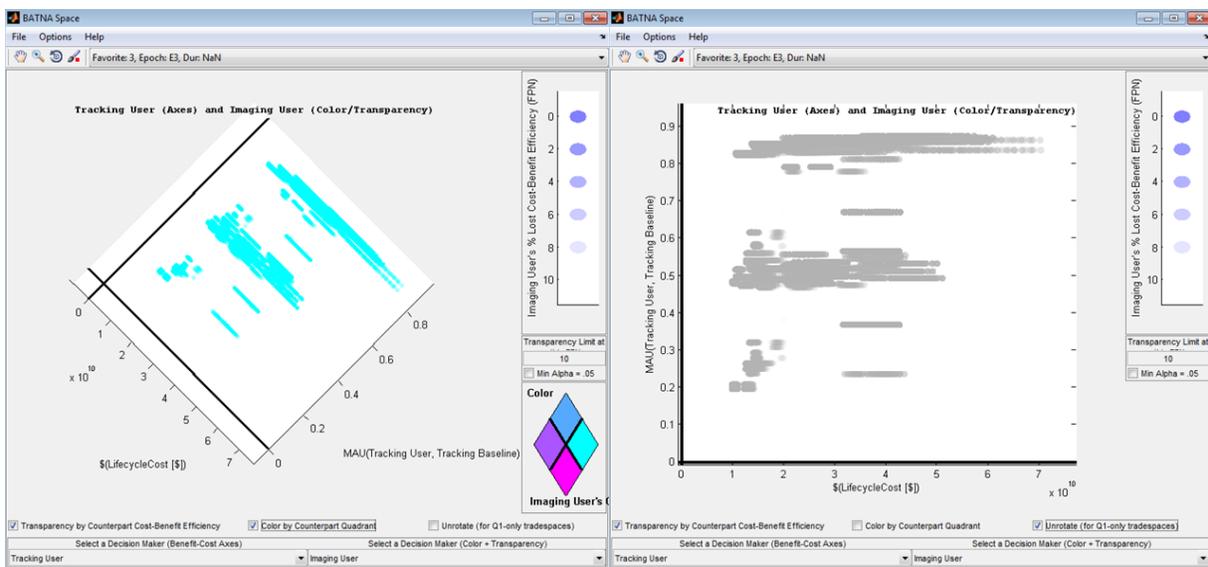


**Figure 5-4: Negotiation tradespace with maximum display FPN of 10, without and with minimum opacity<sup>7</sup>**

Finally, in some cases the data is not spread between all quadrants. Depending on the BATNA, it is possible to have the entire tradespace located in Quadrant 1, particularly when the BATNA is to simply do nothing (for more discussion of this particular circumstance, see 6.3.1). In this case, the rotation and coloring of the negotiation tradespace become unnecessarily complex, because the orientation naturally resembles a traditional single-stakeholder tradespace

<sup>7</sup> This is difficult to see when reduced in size for print, but the right tradespace has many more faded points: compare the bottom row of the transparency legend to see the difference between 0 and 5% opacity.

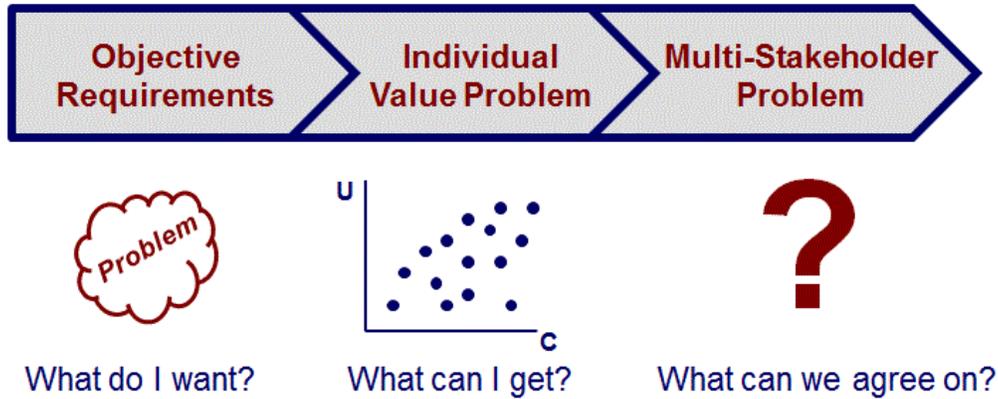
with the BATNA in the lower-left corner and all designs in the same quadrant. Color can be turned off, and the tradespace can now also be rotated back to its natural orientation while still highlighting the BATNA, as shown in Figure 5-5, if negotiators want to use a more traditional tradespace plot. This can again be considered a *show the important* improvement, as the adjustment hides irrelevant information. However, starting with the default rotated view and switching back to the unrotated view (rather than vice-versa) serves an important purpose by illustrating *why* that information is irrelevant. If all designs are in Quadrant 1 then the color dimension can be freed up to show a different type of data, such as a specific design variable as in classic TSE (Ross et al., 2010b). However, when it comes time to evaluate alternatives against the BATNA, it is beneficial to remember that they all represent higher-cost-higher-benefit choices.



**Figure 5-5: An all-Quadrant-1 tradespace, rotated back to natural position and with color disabled, but still bolding the BATNA axes**

### 5.3 Visualizing Relationships

The nature of the relationships between stakeholders is the operative information that MSTSE is interested in making available to support decision making. As discussed in chapter 3.3, the tradespace scatterplot, composed of many individual design alternatives, is effective at communicating the individual value problem of “What can I get?” but fails to provide relevant information about the multi-stakeholder problem of “What can we agree on?”. Alternatives are still the outcome of MSTSE, just as with any design process, but are no longer the complete problem, as understanding and navigating the relationships between stakeholders becomes necessary to successfully choose one alternative.



**Figure 5-6: Increasing complexity in the multi-stakeholder problem demands consideration of more than alternatives**

Part of the challenge in communicating information about relationships is that relationships are abstract, rendering them much more difficult to visualize effectively than design alternatives: hypothetical but concrete potential decisions that can be made. The modified negotiation tradespace supplements the standard tradespace with additional information but still filters that information through a “middleman” of design points – the properties of each design point as considered by another stakeholder, compared between designs. This places a cognitive burden on the stakeholder or analyst conducting the exploration to interpret that information and comprehend the stakeholder relationships. For example, looking at the same tradespace given to the treatment group in the used car experiment (Figure 4-2), it is possible that after a few minutes of careful inspection Nat would conclude that Vic’s interests are largely uncorrelated with his own, given the lack of an obvious pattern in the color or transparency of the design points. This conclusion is correct, but because it is through a layer of analysis there is a risk that it is missed or the stakeholder could lack confidence in his intuition. If information of this type – concerning the relationships between the stakeholders rather than the designs – could be communicated directly, these potential barriers could be eliminated. Relationships are critical to understanding the multi-stakeholder problem, thus visualizing them is an opportunity for MSTSE to contribute a new dimension of interactive tradespace analysis<sup>8</sup>.

The potential benefits of visualizing relationships directly should not be construed as a justification to replace or retire design-point-based visualizations. Ultimately, the goal of MSTSE is to reach an agreement to move forward with a single design, therefore the concrete alternatives remain fundamentally necessary artifacts of the group decision problem. That said, a

---

<sup>8</sup> This is particularly relevant for deploying MSTSE within a larger SE framework (covered in chapter 8), as these frameworks frequently cite the need to “involve” multiple stakeholders but provide minimal specifics on how to do so, implying that exploration of relationships is done qualitatively and separate from the main analysis. Quantitative analysis is more repeatable and less likely to reinforce unacknowledged internal biases, and thus is desirable to be deployed on relationships as well as designs.

plethora of design-centric visualizations exist in the TSE literature that can effectively communicate information about designs. To supplement those, quantifying and visualizing relationships can serve to offload some of the necessary analysis and interpretation of the data from stakeholder to computer. This supports the principles of *analyze first* and *show the important*, leaving the stakeholders to focus on the tasks that they are uniquely qualified to do: decide who to work with and what designs they can both agree on.

#### **5.4 Tradespace Exploration and Stakeholder Correlation**

As used in the motivating example above, the correlation of needs between stakeholders is a useful insight in a multi-stakeholder problem. Stakeholders with positively correlated needs are likely to share many preferred designs, making them natural bedfellows in an agreement. Stakeholders with negatively correlated needs are less likely to reach agreement and if a potential agreement exists, it is likely removed from both parties' Pareto fronts. Knowing this information at the beginning of the negotiation can provide important context for the remainder of the discussion. Additionally, for problems with more than two stakeholders, correlations are directly relevant to the formation of coalitions, which can reduce the complexity of the negotiation and/or lead to productive side-agreements among stakeholders. For more discussion on the potential impact of the nature of alignment in stakeholder needs and how this can influence a negotiation, see section 6.2.1.

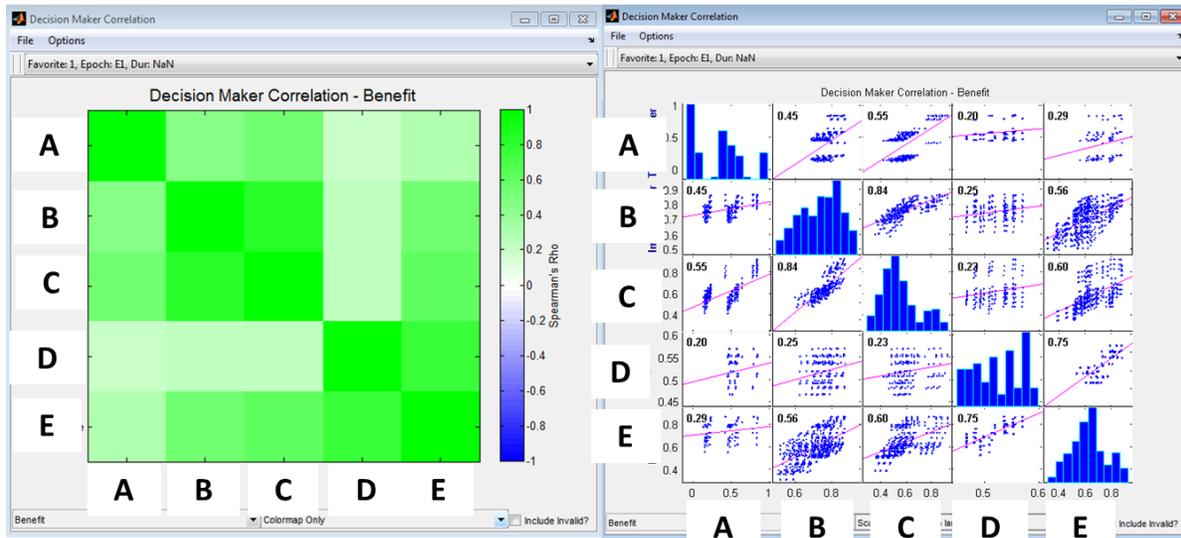
The data underlying the tradespace is sufficient for supporting a calculation of the correlation between stakeholders. Each design is evaluated on many dimensions, including the benefit and cost for each stakeholder, enabling correlation without any additional modeling effort<sup>9</sup>. A dedicated visualization for correlation data would ideally be able to rapidly condense the information relating all of the stakeholders at once (as opposed to two at a time in the negotiation tradespace) to provide a high-level overview while still able to be expanded into detailed analysis upon stakeholder request. To do this, we can use a *correlation map*, as pictured in two forms in Figure 5-7.

The left plot is a heatmap of correlation coefficients: each square in the grid shows the correlation for the given x- and y-axis stakeholders colored from green (correlation of 1, interests fully aligned) to white (uncorrelated) to blue (correlation of -1, interests fully divergent). Note that the diagonal is green, as each stakeholder is fully aligned with himself. This plot provides information about the alignment of each stakeholder's preferred performance attributes as represented by their value model. Correlations such as this can be effective indicators of

---

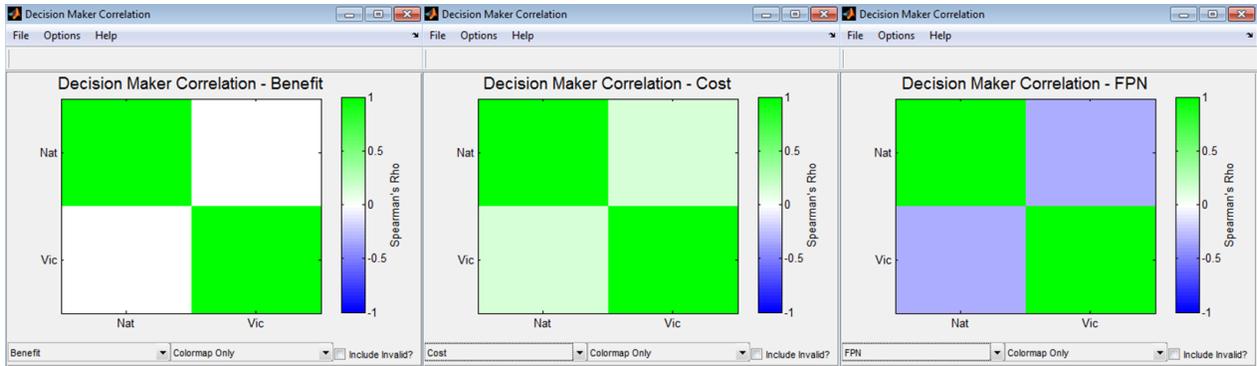
<sup>9</sup> Note that it is important to use a correlation technique that is valid for the input data. For example, in many MSTSE applications, it may be of interest to calculate correlations between stakeholders with multi-attribute utility benefit functions. Because utility measures are ordinal, a basic linear correlation is an inappropriate choice that would unfairly punish strong nonlinearities in preference. A Spearman's Rho rank correlation is a valid choice, and is the preferred correlation metric in this work.

promising coalitions of stakeholders; in this example, stakeholders B and C have very high correlation as do D and E, appearing as small clusters of green. For large numbers of stakeholders, clustering algorithms could be utilized to find these groups. If *detail on demand* should become desirable, the heatmap can be switched into a slightly more information-dense visualization as shown in the right plot, where the correlation values are displayed numerically in the upper-left corner of each square, over a scatterplot of each stakeholder's scores. Additionally, the line of best fit is plotted for each scatterplot in magenta and the diagonal entries are density scatterplots, displaying the overall distribution of value in the tradespace for each stakeholder.



**Figure 5-7: Correlation heatmap visualizing benefit correlation for five stakeholders, in low- and high-detail modes**

These plots can be toggled to show correlations on different dimensions, which can occasionally reveal unexpected or counterintuitive dynamics in the tradespace. For example, consider Figure 5-8, which shows the correlations between Nat and Vic of the used car negotiation on the dimensions of benefit, cost, and FPN (the efficiency-based metric used for transparency in the negotiation tradespace). Their benefits are nearly perfectly uncorrelated and their costs are slightly positively correlated, which would seem to imply that the negotiation should be relatively painless: a matter of finding a design that by chance happens to be good for both of them. However, Nat and Vic have negatively correlated FPN, implying that designs with high efficiency for one typically have low efficiency for the other. This is despite the fact that both of the components of efficiency are at worst neutrally correlated individually: designs with similar benefit ranks for both stakeholders are more likely to have very different cost ranks and vice-versa, explaining why the negotiation experiment was not trivially easy. This type of insight would be nearly impossible to glean from simply looking at the tradespace of design points, and is vastly more available in a non-design-centric visualization.



**Figure 5-8: Correlations on Benefit, Cost, and FPN for the Used Car stakeholders**

One challenge to the representation of correlations in the tradespace is the existence of designs that are invalid for one or more stakeholders. These designs fail to meet minimum requirements for one or more value-generating attributes and, depending on the value model in use, often have undefined benefit or cost, preventing their ability to be correlated. By default, the above visualizations ignore any designs invalid for either stakeholder when calculating their correlation, which *shows the important* information on feasible agreements for both stakeholders. However, a design that is valid for Stakeholder A and not Stakeholder B *still appears in A's benefit-cost tradespace*, and thus may be of interest to A for agreements with other stakeholders, and B's dislike of it should not be ignored. Thus, more *detail on demand* is available in the form of a toggle to include invalid designs, set to a worst-case value of the appropriate metric. Including invalid designs in the correlation calculations can range from little effect to considerable effect. Stakeholders who see their correlations increase or decrease can attribute that effect to agreement or disagreement, respectively, on which designs are invalid. Figure 5-9 shows an example of the range of effect this change can make. Though the inclusion of invalid designs has little evident impact on the correlation of benefit between the stakeholders, the correlation of FPN is dramatically affected. Stakeholders A and D greatly increase in correlation with each other (due to agreement on which designs fail to meet requirements) but decrease relative to stakeholder B. Stakeholders C and E also decrease in correlation; inspection of the high-detail view reveals that designs invalid for C are all very favorable for E.

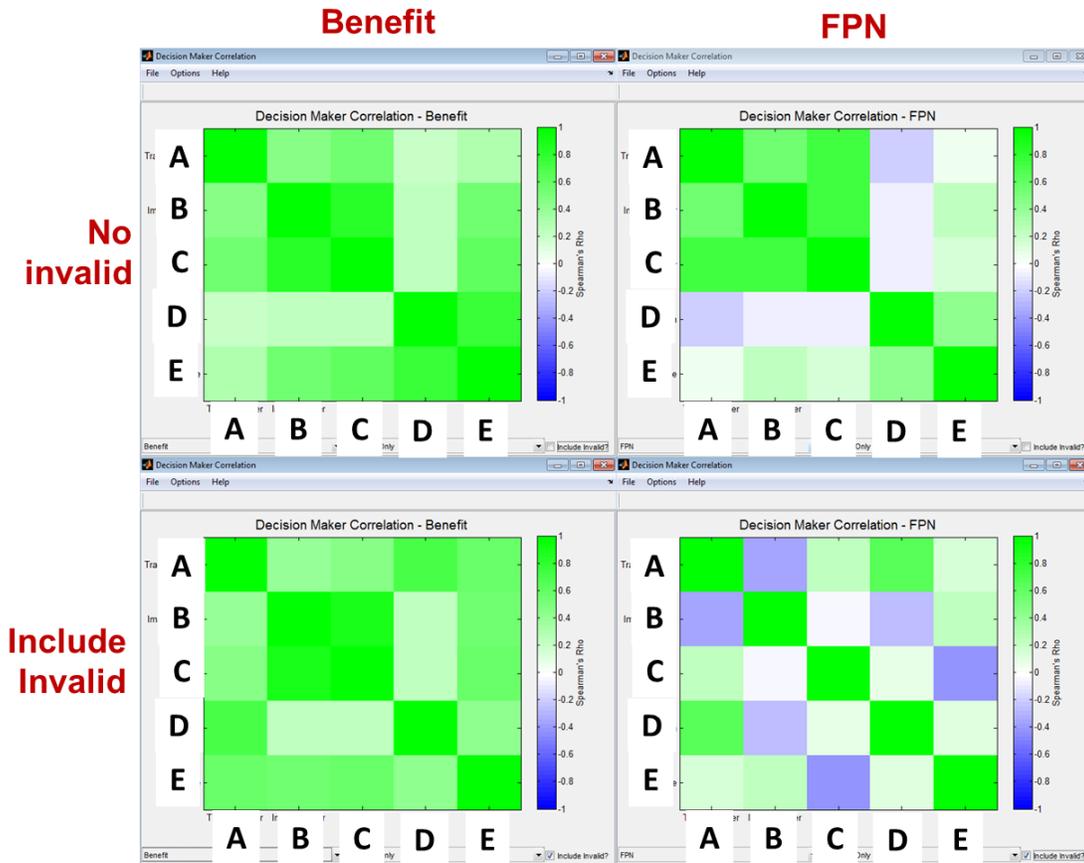
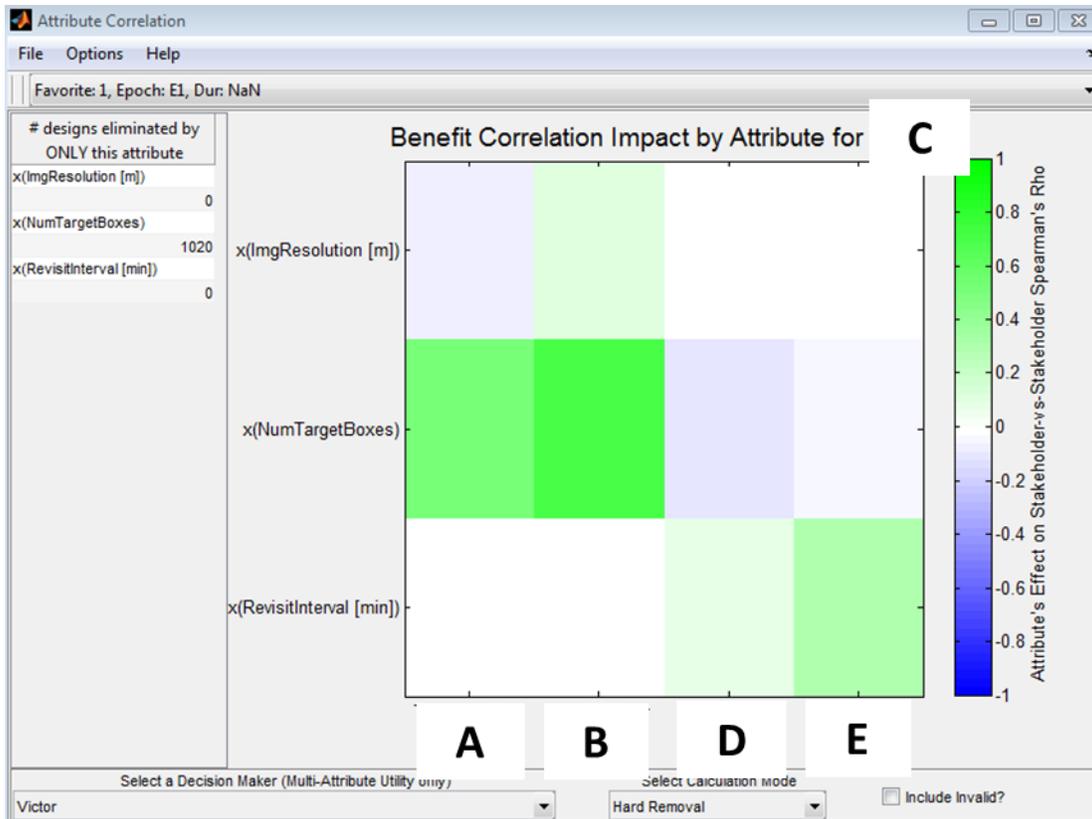


Figure 5-9: Correlation Map, toggling metric used and invalid inclusion

MSTSE can also support the analysis of correlation impacts of individual performance attributes, another variety of *zoom and analyze further*. Figure 5-10 is an example of a detailed stakeholder-specific visualization of attribute correlations for a multi-attribute utility function. In this case, C is the chosen stakeholder, his value-generating attributes are on the y-axis, and the other stakeholders are on the x-axis. The color of each grid square is the impact that the given attribute has on C's correlation with the other stakeholder (as measured by comparison to what the correlation *would* be were that attribute removed). Here we see that the key driver of C's strong alignment with B is through his preference for the NumTargetBoxes attribute, which increases their overall correlation by approximately 0.7. At the same time, this preference is what separates C from the D-E cluster, as indicated by the light blue coloring. A detailed breakdown like this reveals the causality underneath the preference alignments for each stakeholder at a level unable to be shown by the aggregated functions.



**Figure 5-10: Attribute Correlation impact visualization, comparing three attributes of one stakeholder against four other stakeholders**

The attribute correlation view can also be modified for increased *detail on demand* in a variety of ways, including:

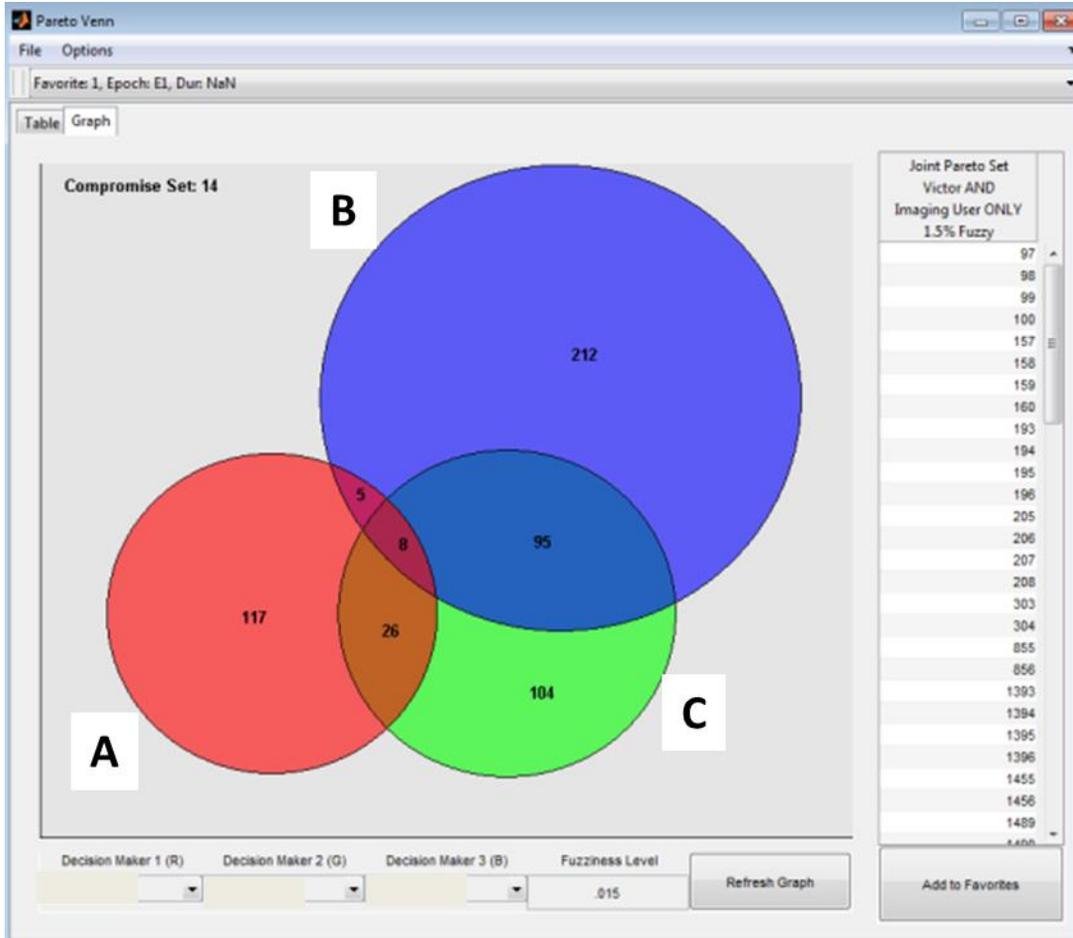
- Inclusion of invalid designs
- Data in left table indicates which attributes are most responsible for invalidating designs
- Selection of correlation model (including options for rescaling the utility swing weights in order to preserve complementary/substitute goods behavior, or simply viewing single-attribute correlations)

Overall, this type of correlation information, provided accessibly in the form of visualization, has the potential to emphasize relationships in the multi-stakeholder tradespace in a way that was previously subject to stakeholder or analyst intuition or qualitative analysis. Understanding the correlation in the value models of stakeholders strongly matches the principled negotiation tenet of “focus on interests, not positions”, and can serve as a high-level reminder of the reason for negotiating in the first place: some stakeholders want the same things and the task at hand is to find an alternative that is agreeable to all parties.

## 5.5 Accessibility of Pareto Sets

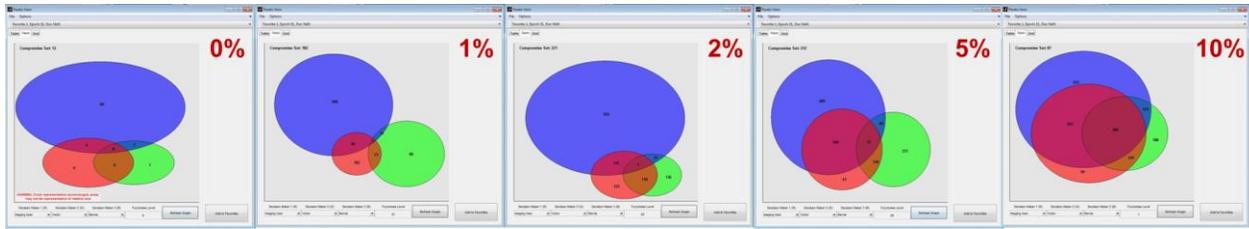
The importance of the Pareto front in tradespace exploration is unquestionable. It encapsulates the most common understanding of value (maximizing benefit-at-cost) and helpfully designates a smaller set of alternatives most worthy of the attention of the stakeholder. In this sense, all Pareto set visualizations would both *analyze first* and *show the important*. However, as previously discussed, the individual Pareto front lacks additional important information about the larger group problem and is an inappropriate reference point for decision making in a negotiation. This puts Pareto sets in a tricky position: stakeholders should and do care about them, but they need to be presented in a way that reinforces the idea that they exist within the context of a negotiation with other stakeholders. With this in mind, rather than calculating Pareto sets separately and allowing each stakeholder to take detailed views of their own preferred alternatives, they should be calculated and displayed together in a way that reveals their interrelationships. This increases the accessibility of information about other stakeholder's needs while also providing new information about the group problem.

An obvious choice for viewing the similarities and differences of Pareto sets is with a Venn diagram, the classic visualization of overlapping sets. It is familiar to most people and offers a simple comparison of different sets. Figure 5-11 is an example of this type of visualization applied to Pareto sets for MSTSE, in this case for three stakeholders. It clearly shows the relative size of each set and intersection, in both area and the number placed in the center (which can be clicked to expand into a list of the pertinent alternatives on the right for *details on demand*). This provides quick access to not only the individual sets, “joint” intersection, and “compromise” set (see 8.4.2) of classic Pareto set analysis, but also the complementary sets: good for some stakeholders but not others. This information can clarify how the excluded stakeholder influences the direction of the group or joint solutions, and may become relevant for identifying when deals between a subset of stakeholders are more valuable than for the entire group.



**Figure 5-11: Venn diagram visualization of three stakeholder Pareto sets**

The Pareto set Venn diagram offers additional useful visual feedback when live updating for changes in either the underlying utility functions or, more directly, the fuzziness setting at the bottom of the window. This can provide instant feedback on the impacts of varying preferences, particularly any nonlinearities in the ways the intersections change. Figure 5-12 shows an example of a Venn diagram changing from 0% to 10% fuzziness. This shows that up until about ~5% fuzziness, the blue stakeholder's Pareto set starts larger and grows faster in size than the other two, yet it is the red stakeholder who is somewhat of the “bridge” between the blue and green stakeholders, who share very few efficient alternatives. After 5%, this pattern changes and the other two sets grow faster – mostly by sharing alternatives with the blue stakeholders, whose “solo” region shrinks in (relative) size while the intersecting regions grow. This information would be very difficult to extract by simply comparing lists of Pareto sets and fuzzy Pareto sets.



**Figure 5-12: Venn diagram changing from 0% to 10% fuzziness**

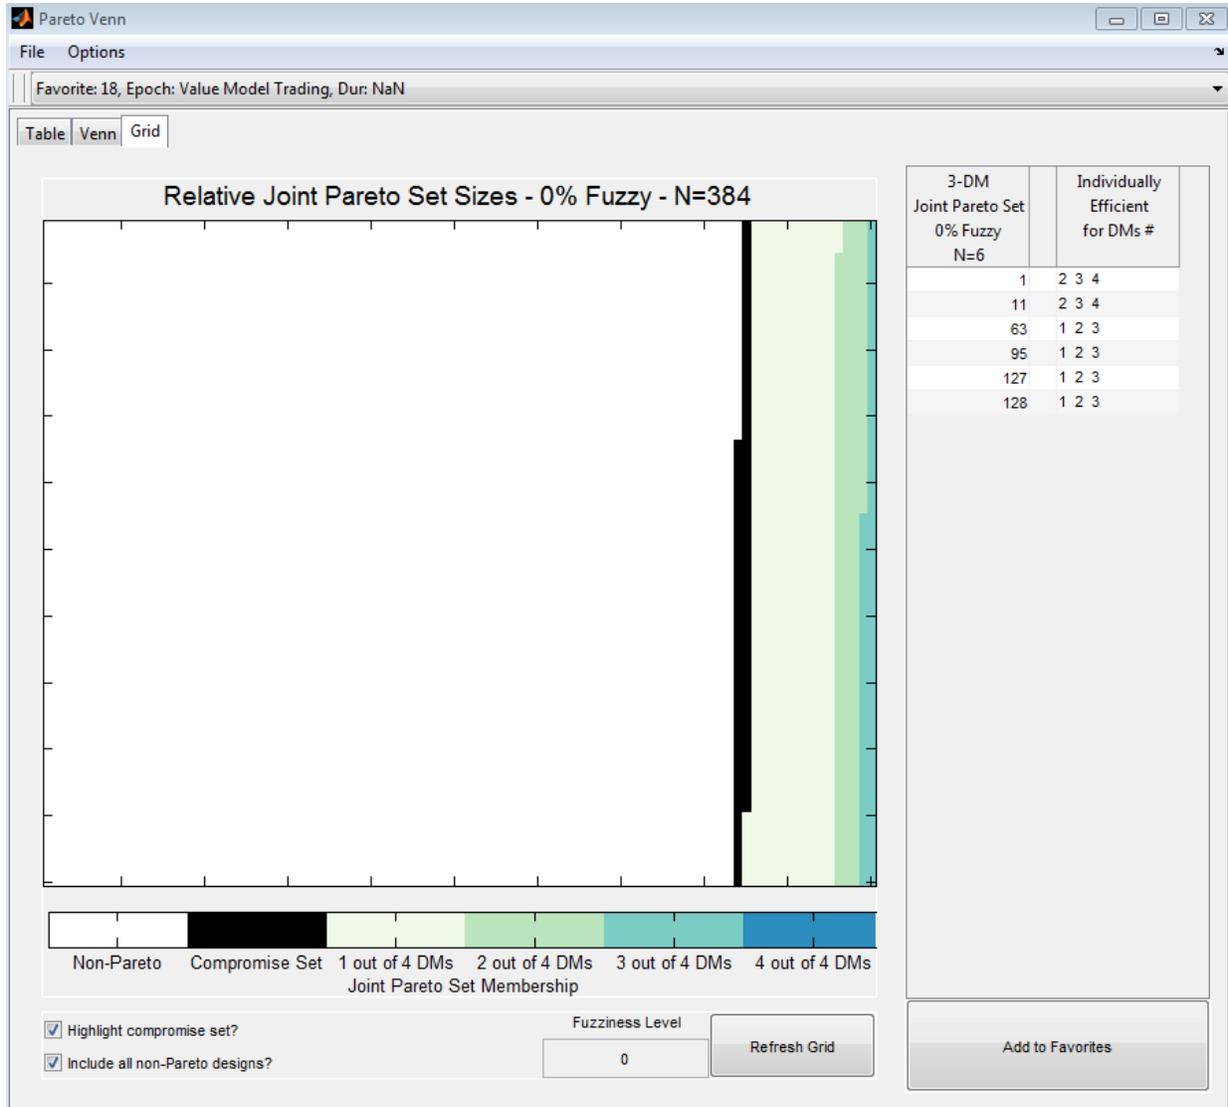
There are two key drawbacks to the Venn diagram representation<sup>10</sup>:

- Depending on the sizes of the different subsets, it may be impossible for the areas of each to be proportional to the correct size. This happens mostly for degenerate sets with very few entries and therefore unfortunately becomes more likely the less fuzziness is used. This problem is mitigated by putting the size of the set in its center but still can reduce the affordance of the visualization.
- It can only show each possible intersection for up to three stakeholders. This is better than the negotiation tradespace, which could only show two stakeholders at once, but is still not scalable to many-party negotiation. Toggling between stakeholders remains an option to get all of them on screen but not simultaneously.

What kind of visualization could show proportional set sizes and be scalable to any number of stakeholders? This can be accomplished by removing one aspect of the Venn diagram, sets identified by *specific* stakeholder, and instead focusing only on the *number* of stakeholders who share a given alternative in their Pareto set. This removes the need to visualize overlaps in sets and therefore allows guaranteed proportional representation for a fully general number of sets. Figure 5-13 is an example implementation of this for four stakeholders, which we are calling a “gridmap”. The complete area of the visualization represents the size of the design space and each block of color corresponds to a group of designs with each belonging to a given *number* of individual Pareto sets (belonging to any subset of the stakeholders). This gives situational awareness of the size of the sets relative to the number of alternatives under consideration in the tradespace, while also *showing the important* by drawing more attention to the generally more interesting designs – the color darkens for the alternatives with more stakeholders in agreement. Each block of color can be queried on click for *details on demand* with a list of the relevant designs and which stakeholders like them. For large design spaces with many inefficient designs, the white region can occupy nearly the entire visualization. In this case there are options to *zoom, filter, and analyze further* by cutting out the white (not efficient for

<sup>10</sup> At least when using only ellipses – algorithmically creating a Venn diagram using general shapes for an unconstrained number of stakeholders is much more difficult to implement effectively, but is a possible further improvement.

anyone) and black (“compromise” efficient) regions and displaying only the designs which actually reside on a Pareto front: the same group of designs contained in the original Venn diagrams.



**Figure 5-13: Gridmap visualization of four stakeholder's Pareto sets**

Again, fuzziness is also a customizable parameter for a gridmap visualization. Playing with it is not as effective as the Venn diagram at communicating the changes in specific inter-stakeholder relationships as the area of interest is expanded, but it can be instructive to see how quickly the entire design space is included in at least one Pareto set. This is especially true for large numbers of stakeholders – as the number of different dimensions under consideration increases, it becomes harder and harder to be dominated in all of them. This is shown in Figure 5-14, where little more than one-sixth of the tradespace is efficient for someone at 0% fuzziness, increasing to over half at 2% and about nine-tenths at 10%. It is also clear from this

visualization that the first jointly efficient design across all stakeholders (appearing as blue on the plot) appears somewhere between 5% and 10%, a range which could be easily narrowed down by live exploration of the fuzziness parameter.

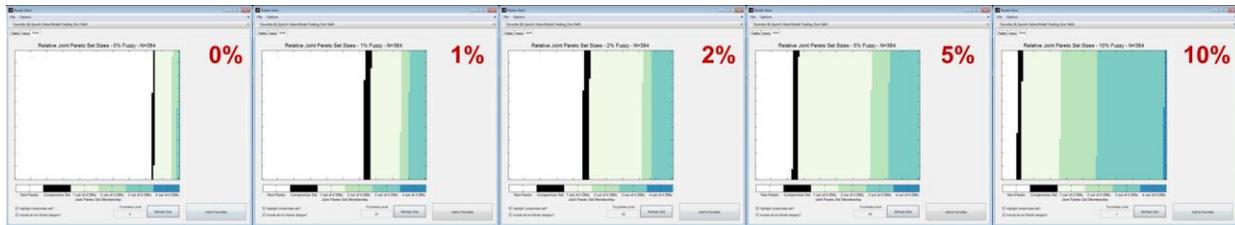


Figure 5-14: Gridmap changing from 0% to 10% fuzziness

## 5.6 Interactive Practitioner Interviews

We have just detailed a variety of improved and innovative visualizations, designed to lead to more effective MSTSE practice. Though there is certainly room for future improvements and new types of representation, they have been presented in their final form as of this point. It is easier to demonstrate their key features and benefits this way than by explicitly demonstrating the significant number of gradual revisions and modifications that occurred over the course of this research. As stated in the outline of the research plan, these revisions were driven partially through interacting and receiving feedback from practicing engineers, in order to ground the development of MSTSE in some of the realities of professional engineering, particularly systems engineering.

To collect this feedback, we conducted a series of semi-structured interviews, guided by the IVTea software suite, concurrently with the development of the new MSTSE visualizations. Semi-structured interviews are often the research method of choice for situations in which there will be a single opportunity to interview the subjects and a limited time for interacting, as they keep the main purpose of the research in focus while allowing tangents based on any emergent information revealed in the interview (Bernard, 2006). In addition to capturing the interviewees' responses regarding the visualizations, discussion was guided to capture the macro frames with which professional engineers approach the concepts of MSTSE and their assessment of its potential for use in the workforce. The identification of these influential perceptions is useful for further prescription on the larger, procedural aspects of MSTSE geared towards improving its applicability and, perhaps more importantly, *palatability* to potential "real-world" adopters. The ability of semi-structured interviews to capture a range of experiences shared by a group, rather than emphasizing individual history, is important for the purpose of finding shared macro frames (Dicicco-Bloom and Crabtree, 2006).

### 5.6.1 Interview Structure

Participants were scheduled in groups of two or three to come to the MIT Systems Engineering Advancement Research Initiative Tradespace Exploration Lab (SEArI TSElab) for a

two hour session. Upon arrival, each participant read and signed an interview consent form, which included optional consent to be recorded and quoted. The semi-structured sessions were divided into four parts, with an estimated time allotment to each:

### **1. (30min) Discussion of personal and professional experience with multiple stakeholders**

The first part of each interview was dedicated to a conversation specifically targeting the participants themselves, without the use of any of the tradespace exploration tools. The goal of the conversation was to extract the initial perspectives of the participants on the purpose and practice of multi-stakeholder design and/or negotiation in their career. Specific attention was paid to relevant experience “in the field,” particularly if those experiences were related to the use of tradespace exploration or any similar design paradigms focusing on the comparison of alternatives. Such experiences provide context to the way each participant views the role that design technique plays at the intersection of different stakeholder interests, which clarifies later reactions to the tools. Some of the question prompts used to guide discussion in this part of the interview include:

- What is your position and experience?
- Describe your experience with tradespace exploration or similar techniques (trade studies, analysis of alternatives, etc.).
- Are you comfortable with negotiation in day-to-day life? At work?
- Do you feel that negotiating between different stakeholders is a part of your job?
- What tools do you use regularly when solving these types of problems?
- What types of decisions are you making?
- How often does this happen in person as opposed to asynchronously or remotely?

### **2. (15min) Introduction to the IVTea tradespace exploration software and example case**

This part of the interview was set aside to familiarize the participants with the IVTea software and problem that they will be using to guide further discussion. It consisted of a short review of the tools and how they work, accompanied by the same TSE “cheat sheet” given to the subjects in the controlled experiment of Chapter 4. The case used as an example was a miniaturized version of Satellite Radar dataset, using about 10% of the design space and only 8 contexts (see Chapter 9 for more detail on this case). This dataset was chosen to represent a more complex, realistic design task than the Used Car dataset of the experiment in response to the increased sophistication of the participants, while still remaining a manageable amount of data for a time-limited session. The participants were not asked to make any type of group decision, in order to keep the focus of the exploration on the accompanying interview questions. However, they were each asked to choose a stakeholder in the dataset for roleplay, replicating the experience of having different preferences on the alternatives while exploring.

### **3. (60min) Case roleplay with guided discussion about technique feasibility and applicability, for both classic TSE and new MSTSE concepts**

Most of the interview was spent guiding a roleplay through the Satellite Radar data with the IVTea software. This part featured a walkthrough of the different visualizations within IVTea, allowing the participants to play with and explore the data. This included both traditional TSE visualizations and tools (benefit-cost scatterplot, filtering and marking designs, etc.) as well as the MSTSE visualizations described earlier in this chapter. The walkthrough asked participants to try out the different tools simultaneously and comment on how they would choose to use them on a real problem and what insights they would hope to gain. The question prompts for this part of the interview were geared towards extracting the perceived familiarity, purpose, and usefulness of each visualization:

- Is this similar to a visualization you currently / previously have used at work?
- How would you utilize this tool in your job? Would you use it in conjunction with another?
- What do you think are the key takeaways from this visualization?
- Do you think this would be helpful for reaching an agreement during a negotiation between stakeholders?
- Does this give you a different impression of the problem than you had from the previous visualizations?

### **4. (15min) Reflection on experience and final feedback**

The interviews concluded with a short reflection on the overall experience and any final thoughts or feedback from the participants. This discussion was intended to capture any “big picture” thoughts that may not have been pertinent to any specific part of the walkthrough, particularly those related to engaging with another stakeholder in a cooperative environment or adapting MSTSE to the needs of the workplace.

- Did you feel like the software was capable of demonstrating your own value?
- Did you feel like the software was keeping you informed of other stakeholders’ values?
- Would you utilize an integrated set of multi-stakeholder tradespace tools such as this in your job? Do you think it would be received well?
- How might you try to convince another stakeholder to work with you using multi-stakeholder tradespace exploration?

## **5.6.2 Participants**

Ten practitioners were participants over four interview sessions. These participants were volunteers, approximately half of whom were solicited individually from the extended professional networks of researchers at MIT SEArI and the other half were collected via

snowball sampling from the first half. None of the participants had prior interaction or involvement with this research. All participants consented to be recorded, but we will be keeping personal identifiers confidential as not all participants agreed to be identified or quoted. Most participants were systems engineers and most had multiple years of experience, targeting the main demographic of TSE and SE practice. However, the additional perspectives provided by non-systems-engineers and less experienced professionals enabled some diversity in the interview discussions. A summary of the participants’ fields of expertise and experience in that area is shown in Table 5-1.

**Table 5-1: Interview Participants - Field and Experience**

<b>Field</b>	<b>Experience</b>
Acquisitions	3 years
Aerospace	9 years
Design / Product development	30 years
Systems engineering	20 years
Systems engineering	15 years
Systems engineering	10 years
Systems engineering	9 years
Systems engineering	9 years
Systems engineering	1 year
Systems engineering	6 months

### 5.6.3 Personal Multi-stakeholder Experience

By far, the most prominent feedback in this part of the interview was that negotiation, at least at a low-level, is a constant part of the systems engineer’s job. It should be acknowledged that this particular result is likely due at least in part to self-selection, as the volunteers to participate in a session could naturally be expected to have an interest in the subject matter of multi-stakeholder decisions. However, given that systems engineers often enter the field after transitioning mid-career from a different role, the examples of “on the job” negotiation varied dramatically with their varying expertise and demonstrated the relevance of negotiation knowledge across a broad spectrum of circumstances. One participant went so far as to declare “everything is a negotiation,” when his list went on longer than he originally planned. Some of the examples given included:

- How to assign “effort” or deploy workforce amongst different research and design activities with different priorities to different customers
- Designing testing protocols during prototyping while balancing the performance for different use cases
- Committee-based decision making for flat hierarchical organizations
- Directing operations to support different objectives for different stakeholders that cannot be performed simultaneously

Notably, none of the examples provided specifically involved negotiating the design of a shared system *with* the involved stakeholders – the participants used the word “balancing” regularly, to indicate that their job as systems engineers specifically requires “balancing” the needs of multiple stakeholders *outside* the design loop. As such, they ascribed considerable importance to reporting activities, which were the critical points at which stakeholders would potentially intervene to influence the overall process in their own favor.

Despite this agreement on the importance of reporting, there was no clear consensus on whether reporting and subsequent decisions are more often made asynchronously or face-to-face, with approximately equal discussion on both sides. Asynchronous decision making driven by digitally-distributed reports and email voting was considered acceptable for small, daily activities and useful for its ease. Face-to-face negotiation, often as a culmination of presentations/reports, was considered rarer and for more important decisions, though some of the more experienced participants insisted that face-to-face negotiation was strictly necessary for decision making despite a trend away from it due to its higher logistic cost. Participants showed some awareness of the value of face-to-face negotiation with stakeholders, generally agreeing that asynchronous decision making carried additional risks, such as a perception that “backdoor” dealings could influence design too late in the process and that design could devolve into “chasing problems” and fixing them one at a time without a coherent plan.

A range of techniques for informing multi-stakeholder decisions were mentioned. Half of the participants cited TSE as a popular technique, though some considered themselves “beginners” in its application and others noted that the biggest barrier to the application of TSE lies in convincing stakeholders that it is worth the time and effort. Excel and/or spreadsheet analysis enabling few-alternative trade studies through simplified performance models was the most common suggestion, with the additional indication that for multi-stakeholder problems the stakeholders were often simply weighted and aggregated in the spreadsheet by “follow[ing] the money” invested in the project. A variety of techniques for identifying functional connections between variables and objectives were also mentioned – including affinity diagrams, House of Quality, and Quality Functional Deployment – under these schemes, each stakeholder would occupy a header/row and design decisions would be traced to the positive/negative effects on each. SysML and other more detailed modeling tasks were cited only twice, with some participants saying that they were rarely feasible on the types of problems they were familiar with due to a lack of supporting data on which to base the models. Interestingly absent from this list of techniques is multi-objective optimization. When prompted, the participants generally suggested that this was due to a similar lack of reliable models on which to optimize, but some expressed an additional layer of distrust in the ability to optimize effectively at the system level.

Finally, it is important to note that human elements were cited as a significant driver of decision making, in a number of different ways. One participant pointed out that they have to be comfortable using a variety of techniques because “it comes down to what [stakeholders] are

comfortable with,” and another thought that the “biggest variable is personality [of the stakeholders]” in influencing choice of method, as the work must appeal to the stakeholders. Another expressed that stakeholders are “afraid” of the large influence that personalities and politics can have on a decision, since a clash at that level can derail an investment, resulting in large decisions being locked in early with only small trades still available at the negotiation stage. Knowledge of these challenges was considered critical to pushing a project to successful completion: as one participant put it, “if a trade study could be done by a computer, I’d let a computer do it.” However, the primary response to these challenges was not to involve the stakeholders in order to generate buy-in, but rather to separate the engineering from the political process and seek incontrovertible evidence that a given solution was in everyone’s best interest.

#### **5.6.4 Response to TSE**

Given that most participants had at least a passing familiarity with TSE, the classic TSE visualizations and tools portion of the IVTea walkthrough prompted less discussion of the tools themselves and more of the benefits and challenges of applying TSE. As was emergent in the initial discussion, the participants all agreed that TSE would be a “useful” technique; however, each interview session sparked opinions that TSE was ultimately limited by (1) unavailability of models even amounting to “medium” fidelity for large systems early in their lifecycle, and/or (2) the culture or relationship between engineers and stakeholders. Specifically, culture was considered a “barrier” to TSE due to a prevalent unwillingness of stakeholders to participate in either model creation (such as eliciting a utility function) or exploration, preferring instead to have engineers “put in values that seem reasonable” and then report results. As such, having experience conducting TSE and a supporting toolset was deemed “necessary but not sufficient” for its widespread adoption. As a counterpoint, one participant believed that the structure provided by a TSE toolset such as IVTea, along with its apparent visualization benefits, could slowly nudge the culture towards TSE by disincentivizing the “ad hoc” combinations of models and analysis common in his experience.

With respect to the range of tools available to systems engineers, TSE (when possible) was considered an upgrade capable of replacing spreadsheet analysis or trade studies, providing more detailed insights on a larger set of alternatives. Unsurprisingly, TSE was not considered a replacement for detailed modeling techniques such as SysML or CAD modeling. In fact, the difficulty of integrating such a model with the large number of alternatives necessary to support TSE was identified as one weakness of TSE compared to trade studies, which can more feasibly support the higher level of detail without overwhelming the user. Despite this, the large set of alternatives was considered particularly invaluable for building “trust” or “confidence” in the underlying models. Even without stakeholders participating in TSE, building model trust was identified as necessary for justifying design decisions, particularly for the commonly low fidelity models used by systems engineers, often without supporting data against which their credibility can be assessed (National Aeronautics and Space Administration, 2008). One interview session

was particularly interested in the opportunity to use TSE to respond to trust “threats” by responsively discretizing key locations of the design space in more detail.

Usage of the benefit-cost scatterplot conformed to the insights of the experiment: most participants chose to go immediately to the Pareto front for their value functions and attempt to characterize those designs and find shared preferred alternatives between stakeholders. Upon finding no shared alternatives, the discussion typically turned back towards individual design decisions (i.e. at what altitude should the satellite be located?) and the search for shared patterns between stakeholders, leveraging the tradespace in a manner similar to a House of Quality. It should be noted that, given the time constraint and the emphasis on interview discussion, this exploration was mostly cursory, and thus perhaps more indicative of the easiest or most accessible way of using a tradespace visualization than the workflow the participants would use in reality. One two-person session took an interesting secondary approach after finding no jointly efficient solutions by using a utility-vs-utility scatterplot between the two stakeholders, finding clustered “regions” of designs in that view, and then identifying the lowest cost solution in each region. When asked for clarification, they suggested that this was the preferred approach in their work: minimizing cost for given levels of performance across multiple stakeholder missions. This action is analogous to picking alternatives from a 3-dimensional Pareto front and deliberately spreading them around the complete range of the dimensions, but would not be possible using this particular method with more than two stakeholders or without the existence of a single, shared cost metric.

Beyond the general impressions of TSE, two main points were raised as shortcomings of the classic TSE visualizations when working with multiple stakeholders. First, alternatives that one stakeholder is indifferent towards (i.e., occupying the same position in the benefit-cost tradespace) were noted to be key “opportunities” for agreement by leveraging sub-optimization amongst them for the other participants. Yet, with points “overlapping” in the tradespace, there is no indication that such an opportunity is present without resorting to filtered lists of alternatives. Another participant criticized the benefit-cost tradespace for having too much extraneous information to be a proper base visualization of the problem, arguing that if the key challenge was to work between stakeholders then the analysis should “start with goals:” relating the stakeholders at a more fundamental level before diving into detailed alternative comparison. These points were used as segues into the new MSTSE visualizations, where they could be addressed more directly.

### **5.6.5 Response to MSTSE**

The BATNA-centered negotiation tradespace received mixed feedback from the interviewees. On the positive side, multiple participants appreciated the clear distinction between good, bad, and tradeoff alternatives for each stakeholder, as represented by the quadrants. The concept of the BATNA itself was new for many participants and there was a positive response to grounding analysis against a fixed, achievable decision. When discussing

what can constitute a BATNA, one group was excited by the possibility of using current assets as the BATNA and enumerating alternatives as combinations of additional assets. Additionally, the group that discussed “overlapping” alternatives in the standard tradespace agreed that the use of transparency to partially hide alternatives unattractive to other stakeholders was able to address that concern. One group described it as getting “down to where you want to be, quicker” by reducing the amount of positional bargaining between points to find good solutions across stakeholders. All interviewees expressed that they would be willing to use the negotiation tradespace over the standard tradespace for a multi-stakeholder problem

The negotiation tradespace was criticized by every group for being too complex at first exposure, presenting a “barrier to entry” that makes using it difficult at first and might turn away potential adopters. Additionally, the people most likely to be turned away would be the stakeholders themselves, further reducing the likelihood that they would be active participants in exploration. Multiple groups suggested that a tradespace expert would be needed to serve in a facilitator’s role if stakeholders were to explore themselves. The possibility of engineers exploring as “proxy” decision makers was also entertained: some participants felt that stakeholder participation was too critical to pass up, while others felt that the “proxy” role was a natural fit for systems engineers in their current job description. Alternatively, two groups recommended “bucketing” alternatives in order to simplify the appearance of the tradespace – categorizing them by their relative appeal to different stakeholders.

That request was met through the Pareto Venn diagram, which the interviewees agreed was a much simpler visualization of important information, communicated at a level that would appeal to stakeholders. One group felt that using the Venn diagram (with fuzziness corresponding to model uncertainty) was the fastest way to screen the design space, reduce its size, and iterate with higher fidelity models. Recommendations from early interview sessions were partially responsible for the creation of the gridmap representation of this information, which was also received positively by later groups. As an unanticipated side effect, these visualizations prompted questions of “prioritizing” or otherwise weighting stakeholders in an attempt to get a unidimensional ranking of alternatives. Upon discussion of the drawbacks (and mathematical liberties) of such techniques, the interviewees seemed comfortable leaving the interests separate; however, the initial response was concerning from the perspective of how to deploy such analysis in the field without making such messaging clear. One interesting comment arising from this discussion concerned the use of fuzziness in Pareto set calculation to indicate model uncertainty versus “willingness to compromise.” These two concepts each have a natural correspondence in fuzziness that makes them difficult to separate (i.e. it may be unclear if a design is acceptable due to uncertainty or generosity).

Finally, the stakeholder correlation visualizations were also mostly positively received, notably for addressing the “start with goals” complaint from the benefit-cost tradespace. The attribute correlation was praised for its ability to easily identify “pain points” between

stakeholders – the issues that divide their interests, as represented by driving lower correlation – because it is a regular task for systems engineers and a valuable means of simplifying a complex problem into a small set of tradeoffs on key attributes. One group criticized this visualization because they wanted correlations driven by design variables rather than attributes since it would lead more directly to actionable decisions; i.e., “do all stakeholders agree that a larger antenna is better?” When discussing what such a visualization would look like, they agreed that the heatmap-style visualization would not work on more than two stakeholders. Additionally, one participant expressed concern that a negative correlation for one stakeholder could lead to them being categorized as “difficult” and then largely ignored for the remainder of the negotiation.

### 5.6.6 Reflection on Applicability

Overall, each group agreed that the visualizations in IVTea were capable of demonstrating multiple value propositions across a shared set of alternatives, would likely be received well by their coworkers, and that they would be interested in deploying similar techniques in the field – always with the caution that the challenge of implementing TSE makes it not feasible for some design tasks. The interviewees were most concerned with the difficulties they foresaw in getting stakeholders to participate, driven largely by the visual complexity necessary to capture the intricacies of a multi-stakeholder problem. A variety of suggestions for “marketing” MSTSE were given and they serve as an informative overview of the big-picture benefits that the interviewees saw in the technique:

- Emphasize “convergence” in insight across the different visualizations to build trust in decisions, rather than allowing stakeholders to ascribe any counterintuitive insights to model errors and therefore revert to their gut instinct
- Make it explicit that MSTSE highlights relevant information to support stakeholder decision making and does *not* find the “right” answer – again, stakeholders will not trust black-box algorithms over their gut instinct
- Demonstrate ease of iterating on design space by making incremental decisions on unanimously approved design variables
- Find the right programs to use MSTSE on, with the budget and schedule to support detailed analysis during conceptual design
- Sell the reduced “overhead” for visualization through reuse when consistently using a structured technique like MSTSE
- Figure out a way to integrate MSTSE with modeling in order to prevent garbage-in-garbage-out situations

However, the drawbacks to MSTSE mentioned throughout the interview were also brought up again here:

- Not “practical” to assume stakeholders will participate, requiring a backup plan able to be performed entirely by engineers

- Complexity barriers when stakeholder do choose to participate
- Training costs of creating TSE experts to facilitate MSTSE
- Unavailability of supporting models for many projects

Additional requested features for the IVTea software included:

- “Bucketing” alternatives to provide a reduced-complexity version of the negotiation tradespace
- Allow the fuzziness setting of Pareto set calculation to be set on a per-attribute basis to reflect differences in model uncertainty
- Separate fuzziness into parameters for uncertainty and “willingness to compromise”

The interview sessions all ended positively, with criticism for MSTSE couched in interest for streamlining its application in the professional environment. Generally, the current status of MSTSE was considered a first step towards wider adoption. One participant wrapped up his thoughts with: “No doubt that in thirty years, this is how people are going to do these kinds of problems” and the remaining question was how to use trust-building to accelerate progress to that point. Another specifically called out the benefit of exploratory analysis that allows stakeholders to learn more about the system without committing to an agreement using a rhetorical question: “If you *could* do this, why would you not?”

## **5.7 Discussion**

The comments provided by the engineers participating in the interviews were invaluable to the progression of the research. Specifically with regards to the visualizations, their feedback helped to refine the controls and appearance of each while also justifying the value and importance of backing tradespace analysis out to the inter-stakeholder relationship level when designing to meet multiple sets of needs. Additionally, their comments also supported other parts of the research plan for developing MSTSE: particularly the need to integrate MSTSE in model creation, which is a part of the larger systems engineering methods described in Chapter 8.

The interviews also provided significant insight into the macro framing with which these systems engineers approach design tasks. The stark separation between engineering and stakeholders was a recurring theme, resulting in a design process where engineers build analyses, submit reports, and iterate when (frequently) necessary. This divide is both a significant barrier to the application of MSTSE amongst stakeholders with no experience actively participating in the design process and also an opportunity for MSTSE to provide an accessible entry-point for stakeholders to clearly communicate their needs to engineers and reduce costly iteration. It is beyond the scope of this research to determine if this situation is present across the entire field, as we make no assertions that this sample of engineers is representative. However, the

messaging from the interviewees was completely consistent with regards to their interactions with stakeholders, suggesting that there is at least a niche in the systems engineering community in which MSTSE has the potential to contribute value to the design process.

Regardless, it is clear that the biggest challenge for the implementation of MSTSE in the near future will lie in convincing stakeholders to participate in what is, on the surface, a traditional systems engineering task. MSTSE without active stakeholder participation, potentially through the use of systems engineers as “proxy” decision makers in an informal negotiation, was positively received by the interviewees and is discussed further in Chapters 6.5.1 and 7.4. Fortunately, the self-described “balancing” act that systems engineers use now when solving multi-stakeholder problems is similar to how this problem would be approached without stakeholders driving the exploration. This can make MSTSE an asset to the existing job description of many systems engineers, creating additional data and visualizations to support the analysis they already need to perform at the intersection of different stakeholder interests.

We have now described some of the tools available to systems engineers interested in performing MSTSE, as well as covered the perception of MSTSE as a potential asset to early concept design in practice. The following chapter will describe how problem structure can impact the effectiveness of MSTSE or specific visualizations, and how systems engineers could look to mitigate these impacts.



## 6 Impact of Problem Structure on MSTSE

Returning to the work of Bazerman et al. (2000) on the history of negotiation, much of the research in the field in the 1960s and 1970s centered on the analysis of situational or structural variables that described the problem, such as power dynamics, deadlines, and constituencies. These techniques eventually lost favor to the behavioral decision theories, which offered more effective practical advice on the grounds that people do not obey consistent standards of rationality and thus how they choose to interpret the problem has a greater explanatory effect on their behavior than the problem's "actual" structure. Obviously, this research has been influenced considerably by the behavioral perspective through its concern with framing but a return to the basics of structural analysis can potentially increase the effectiveness of MSTSE by providing additional prescriptive guidance. A smaller domain ("engineering negotiation" as opposed to "negotiation") should allow for tighter prescription. Additionally, an appeal to structure and rationality will likely be more appreciated among the target audience of engineers than among all negotiators.

The benefits of analyzing the structure of a multi-stakeholder design problem are twofold. First, the usefulness of any given activity within MSTSE can be strongly dependent on the nature of the problem itself. As a simple example, consider a case in which the stakeholders themselves cannot participate in MSTSE and instead send proxies: in this situation, exploring the potential addition or subtraction of value-generating attributes from a utility function has reduced impact because of the inability to rapidly incorporate updated stakeholder preferences. The structural elements of an MSTSE case will naturally emphasize certain actions and deemphasize others, and thus understanding the structure will allow for more focused, efficacious negotiating. Second, matching the structure perceived by participants to the real structure should result in stronger decision making, in the same way that matching the perceived reference point to the actual BATNA results in stronger interpretation of gains and losses. A common example of this in the negotiation literature revolves around convincing negotiators to drop preconceived notions of zero-sum bargaining when mutual gains are available, as it leads to superior outcomes (Raiffa, 2002). By looking at structure, we can ensure that MSTSE participants are not just solving problems but solving the right problems.

This work is not intended to be a complete taxonomy of the structural dimensions of multi-stakeholder negotiation, but rather to serve as an outline of important considerations to make when preparing for and conducting MSTSE. For clarity, it is helpful to split these considerations into five categories:

1. Stakeholders – who is participating and how are they related?
2. Preferences (Benefits/Costs) – what interests are being satisfied?
3. Alternatives – what choices or decisions can be made?
4. Uncertainty – what is unknown and how can the situation change over time?
5. Logistics – how is the negotiation to be conducted?

The following subsections will cover these categories in more detail. The analysis will focus mostly on the “forward path” – going from assessing the problem structure to tailoring an instance of MSTSE to best suit that structure. The “reverse path” of using MSTSE methods or visualizations in order to uncover that structure (presumably to then use that insight on the forward path) will also be discussed for characteristics that are not immediately apparent from the basic problem formulation (e.g. correlations between values). Also, note that many of these considerations are interrelated and therefore this chapter will have many references to other relevant subsections within the chapter, indicated in brackets (e.g. [6.1]).

## **6.1 Stakeholders**

Stakeholders are the core component of MSTSE. They are responsible for the final outcome, with a selected course of action requiring the agreement of some combination of stakeholders. Understanding who is “at the table” (and, relatedly, who *should* be at the table) is paramount to a successful negotiation. Identifying and collecting the proper stakeholders is a critical task in problem formulation that, if done properly, should generate a structure that should be considered when going forward with MSTSE. In this section, we will assume that the set of stakeholders has been properly assembled. The key structural dimensions we will consider are the *number* of participating stakeholders, the *outside relationships* that they have with each other, and the *representation* of a constituency for which they are responsible.

### **6.1.1 Number**

Though many example negotiations are conducted between two stakeholders, in general there can be any number of participants with a vested interest in a negotiation. Unsurprisingly, the minimum of two stakeholders is the simplest case both to conceptualize and to visualize, hence its frequent use in example problems. Working with more than two stakeholders brings some interesting challenges for clearly representing all the relevant value dimensions while exploring the tradespace. In some cases, the nature of a stakeholder’s value model may allow for it to be combined with others or reduced to fewer dimensions; these cases are covered in the section on preferences [6.2.3]. This section will consider the exploration of three or more distinct stakeholders with their own benefit and cost functions.

When dealing with more than two stakeholders, the benefit-cost tradespace scatterplot is limited in its ability to demonstrate the complete problem space. If the benefit-cost axes are fixed to one stakeholder in order to maintain similarity to single-stakeholder TSE, there are only a few other dimensions that can be leveraged effectively, among them size, shape, color, and transparency. This research has already demonstrated the use of color and transparency to increase the availability of value information about *one* other stakeholder. Using size and shape to display another stakeholder (really, two stakeholders with any two dimensions each) is possible but will make the visualization cluttered and difficult to read. Of particular concern is placing the same *type* of information in two different dimensions, such as making color indicate quadrant for one stakeholder and size indicate the quadrant for another. This inconsistency runs

counter to the principles of affordance in designing interfaces for accurate human perception (Gibson, 1979; Norman, 1988).

One approach to extending this type of visualization to more than two stakeholders is to simply generate multiple tradespaces with color and transparency tuned to a different stakeholder in each. A similar strategy would be to divide each data point into a section for each stakeholder, for example putting one stakeholder on the left half and another on the right and assigning the appropriate color/transparency to each half. This removes the need to compare multiple tradespaces by increasing information density, which may make the information difficult to read in tradespaces with large numbers of alternatives (or otherwise small markers for each alternative). These solutions are manageable for a small group of 3 to 4 stakeholders, where the burden of cross-referencing between the tradespaces is minimal and it is feasible to divide markers and remain readable.

For larger groups of stakeholders, it is impractical to attempt to visualize every relevant dimension at once. This presents a challenge unique to MSTSE compared to TSE: the lack of a “home” visualization that summarizes the complete problem. Standard TSE often treats the benefit-cost tradespace as the “home” off of which other analyses are performed and their insights cross-referenced back to the “home” view; for example, modifying a value function and seeing the impact on the tradespace. To some extent this is misleading, as even in single-stakeholder TSE the benefit-cost tradespace shows only a fraction of the complete information in the problem, which the other analyses supplement with increased detail. A similar paradigm is possible in MSTSE for a small number of stakeholders, but loses effectiveness with larger numbers that preclude an effective “home” that can summarize the value for every stakeholder at once. This places an increased demand on the participant to actively search for alternatives that are potentially attractive to the group while simultaneously limiting the ability of the BATNA-centric tradespace to provide consistent framing of the problem.

In order to preserve the classic benefit-cost view and BATNA-centric framing of the negotiation tradespace, a divide-and-conquer approach could be utilized. Searching for alternatives that are attractive to two stakeholders can yield interesting intermediate solutions, particularly if the stakeholder pairings are potentially viable coalitions [6.2.1]. These alternatives can then be brought to the rest of the stakeholders in search of additional backing. A different approach would be to focus on visualizations and analyses that identify functional drivers of value in the design space, such as carpet plots (German et al., 2013; O’Neill, 2013). Understanding interactions between design variables and value models can lead to stronger intuition for what a “good” alternative would look like for each stakeholders and enable effective detailed exploration of promising areas of the design space.

### 6.1.2 Outside Relationships

The outside relationships between the participating stakeholders in an MSTSE task will have an impact on how the ensuing exploration is performed. As discussed in the literature review, this research has focused on the use of MSTSE as support for independent decision makers, but they are not the only potential participants. Connected stakeholders, such as subsystem design teams working within a company, may choose to use MSTSE as a means to learn about their problem space and they would therefore follow the basic recommendations (while retaining the option to aggregate their preferences and optimize with minimal loss [6.2.3]). On the other hand, some negotiations involve both mission-critical decision makers, who must agree on a course of action and hold veto power over any alternative, and involved stakeholders such as interest groups, who are invited to participate and express their opinion but do not ultimately control the outcome. This is common in negotiations with far-reaching social impacts, where many people are affected by the decision and the relevant decision makers have an interest in their satisfaction with the result, either altruistically or with respect to their influence in other areas.

Critical decision makers should remain the core of the MSTSE activities. Because each critical decision maker must be satisfied with the agreement regardless of the presence of other stakeholders, utilizing MSTSE visualizations and analyses that limit consideration to the decision makers provides a simplified view of the overall problem that retains its most important features. The decision makers can then choose to blend in the values of whichever stakeholders they themselves are interested in satisfying. The method of incorporating other stakeholders can vary depending on the degree of importance they are ascribed, from low (e.g. used to break ties of indifference only) to medium (added as a component to the decision maker's multi-attribute value function) to high (treated as critical by that decision maker). The ability to customize the impact of each stakeholder is an advantage of MSTSE over other multi-stakeholder analysis techniques that may only leverage assigned importance weights. The minor stakeholders can be encouraged to explore the tradespace on their own, focusing mostly on their own relation to the decision makers, and make recommendations to the decision makers with whom they have influence.

Another relationship between stakeholders that has been assumed to this point is that they are naturally either collaborative or cooperative: operating either as cordial partners with a desire for mutual success or independent agencies with no preference on the other party's outcome. These are considered the prime candidates for the use of MSTSE, given the significant degree of mutual investment and interaction it requires. However, it is possible that decision makers who are naturally competitors may seek to use MSTSE as a means to coordinate and find win-win alternatives. A basic example would be contract negotiations between a customer and a supplier, where cost for the customer is benefit for the supplier but the other value dimension is independent for each. The competitive scenario brings with it a large barrier to the Full, Open, and Truthful Exchange principle because each party can reasonably expect that holding private

information will give them an advantage. Using methods to circumvent private information [6.5.3] can still yield useful exploration outcomes, particularly with the use of a neutral mediator for competitors. However, competitive stakeholders will naturally make reaching an agreement more challenging, as they nominally have opposing interests relative to the other participants. If reaching an agreement is of paramount importance (i.e. taking the BATNA is considered a failure), problem formulation should identify if competitive stakeholders can feasibly be excluded from the MSTSE negotiation in order to build momentum with cooperative stakeholders only. Unofficial or nonbinding MSTSE sessions do not necessarily need to incorporate or invite every stakeholder, particularly if the competitive stakeholders are not system-critical decision makers.

### **6.1.3 Representation**

Stakeholders are often responsible for representing the interests of a constituency that cannot be present at the negotiation. Legislators are a common example of this role, as they typically represent populations too large and diverse to attend the negotiation as minor stakeholders [6.1.2]. Stakeholders may be limited in their ability to make rapid or unilateral decisions when acting as representation. This can result in costly delays in response to changes in the problem formulation such as creative expansion of the tradespace or emergent needs, which may have to be referred back to the constituency for evaluation.

This situation concisely summarizes the reasons that having empowered, informed decision makers as the participants in MSTSE is valuable. With representation only, either the process is subject to delays when communicating back with the constituency or the representative must make spot judgements that may not accurately reflect the backing interests or success criteria for the system. There is no effective way around this problem without the ability to guarantee that the problem formulation is 100% correct and will not change, which hinders the ability to leverage any emergent insights arising from TSE and MSTSE. Nevertheless, if minimizing delays is desired, the MSTSE activities that encourage creative rework of the problem become less useful. This includes exploration of value functions (which may need to be amended in response to other stakeholders) and the design space (which may need to be expanded if high value regions are near the edges). Ultimately, this results in a “back to basics” type of tradespace exploration, focusing entirely on the value space. As a top-down representation of the current problem formulation, the benefit-cost tradespace can be considered complete information when the problem formulation is fixed, removed from potentially changing value functions or alternatives, and thus becomes an even more powerful insight-generating view.

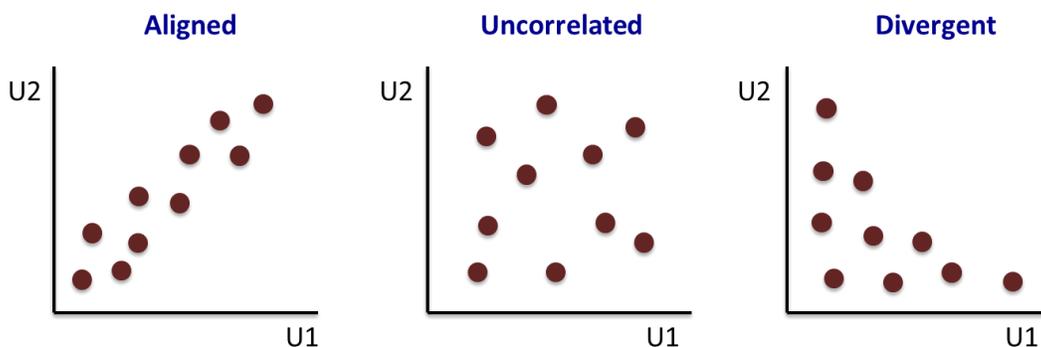
## **6.2 Preferences (*Benefits/Costs*)**

Preferences represent the interests of the participating (and occasionally non-participating) stakeholders. Typically, when performing tradespace exploration, preferences are encoded using a value model: a functional transformation from design attributes into a value

scale. These value models can range from non-dimensional multi-attribute functions (e.g. multi-attribute utility, analytic hierarchy process) to monetizations (e.g. cost-benefit analysis) to simple one-to-one mappings (e.g. measures of effectiveness). Regardless of their functional form, the output of these value models is a numerical representation of preference for a given stakeholder. This provides a mathematical means for both comparing design alternatives for a single stakeholder, as in classic TSE, and comparing the relative attractiveness of designs between stakeholders in MSTSE. Because the ultimate goal of MSTSE is to identify design alternatives that are desirable for all stakeholders, their stated preferences are important tools towards determining the *alignment* of interests for the participants. In addition to this top-level goal present in all MSTSE endeavors, we will also discuss some corner cases where preference structures allow for simpler analysis or demand deeper understanding.

### 6.2.1 Alignment

The alignment of preferences for stakeholders participating in MSTSE is, superficially at least, a significant determinant of the ease of reaching an agreement. If everyone’s interests are aligned, design alternatives that are attractive to one stakeholder will also likely be attractive to the others: not necessarily to the same degree, but at least not unappealing in overall performance. In the opposite situation, when preferences are strongly divergent, finding mutually beneficial alternatives will require more deviation from individually preferred choices. Figure 6-1 shows notional plots with utility axes for Stakeholder 1 on the x-axis and Stakeholder 2 on the y-axis, demonstrating a variety of different alignments. Keep in mind that these plots show only *benefit* for *two* stakeholders: in problems with separated benefits and costs or more than two stakeholders, using a plot of this type to justify a “best” selection (on the U1-U2 Pareto front most likely) would leave out multiple dimensions of important value information.



**Figure 6-1: Notional Utility-Utility axes showing different amounts of preference alignment**

Alignment of preferences can be roughly estimated with a basic understanding of attributes of interest. For example, we could expect that two stakeholders who both care about top speed (among other, different things) would have some alignment due to sharing an objective attribute. Similarly, if a group of performance attributes are known to be directly related through the underlying physical model of the system (e.g. weight and volumetric displacement for a

ship), we could predict alignment between stakeholders who value any subset of those attributes. Though these are effective rules of thumb for informing general discussion, treating stakeholder alignment in this way relegates it to an abstract concept. Alignment can be quantified directly using statistical correlation methods, providing not only a numerical measure to compare and contrast but also a considerably higher level of detail. The correlation visualizations of section 5.4 can be used within MSTSE to gain a deeper, quantitative understanding of whether and why stakeholders agree or disagree.

The alignment of preferences for the stakeholders, understood either generally or with the assistance of MSTSE, can inform positive negotiation tactics. Again, the case in which all the stakeholders have well-aligned preferences is in some sense the “easy” case, as their interests point them towards similar desirable locations in the tradespace. However, if reaching an agreement proves difficult (perhaps through unoperationalized preferences or metapreferences [6.2.4 - 6.2.5]), returning attention to the positive overall alignment can be a useful reminder that the parties are in fact “on the same team” and should be working together to find a solution that works for all. In this circumstance, a Utility-Utility plot like Figure 6-1, or a higher-dimensional variant for more stakeholders, is in fact a good starting point for finding other promising designs of interest.

In contrast, if stakeholder preferences are strongly divergent, the Utility-Utility plot should be avoided as it only serves as a negative reminder of the lack of alternatives that are highly valuable to all parties. In this case maintaining focus on BATNA-centric visualizations is advisable, as if any mutually beneficial solutions exist they are unlikely to be found near individually optimal areas of the tradespace. Instead they would be in the “body” of the tradespace: increasing the importance that the appropriate reference point be emphasized in order to properly evaluate solution quality. Alternatively, the problem of divergent preferences could be tackled directly if it is believed that the preferences themselves are possibly inaccurate representations of true value and the true preferences may be more aligned [6.4.3].

Preference alignment also offers a potentially useful opportunity for problem simplification through grouping stakeholders. Stakeholders with strong correlation in their value models are prime candidates for coalitions (Mitchell, Agle, and Wood, 1997). In a game theoretic interpretation of a negotiation, coalitions present a united front (in MSTSE: choosing a preferred alternative together) in order to improve their bargaining power. Although it may be a winning strategy to form a large coalition to bully other stakeholders or otherwise control the negotiation, we will set this tactic aside for now in the name of continuing to support the tenets of principled negotiation (Fisher, Ury, and Patton, 1991). Coalitions can still provide useful organization to large MSTSE efforts with many stakeholders, however, by reducing the overarching complexity of the problem. If four stakeholders have two strongly aligned pairs, it is possible that reducing the negotiation down to a two-coalition negotiation can break deadlocks created by all four parties vying for their own interests. Coalitions can choose to either

aggregate the preferences of their members into a single coalition-level value model [6.2.3] or simply negotiate as one unit, relegating satisfaction of each member to a side task.

Additionally, emergent coalitions can sometimes present strong evidence in favor of searching for solutions, either within or external to the enumerated tradespace, that specifically require only the stakeholders in the coalition. For example, if stakeholders A, B, and C are engaging in MSTSE but it becomes apparent that A and B are strongly aligned but C is not, it may be in the best interest of all parties for A and B to work with only each other (if such a decision is possible). This allows the A-B coalition to focus on finding potential high-value solutions for themselves, preventing the divergent interests of C from reducing their own value in an ill-advised attempt to be inclusive. Similarly, this condition may be a strong indicator that C should look for other potential stakeholders with whom to do business. Of course, in some domains all stakeholders are mandatory participants in the final agreement (e.g. legislators, licensors), and it is not feasible to drop or remove a subset of them in search of higher value.

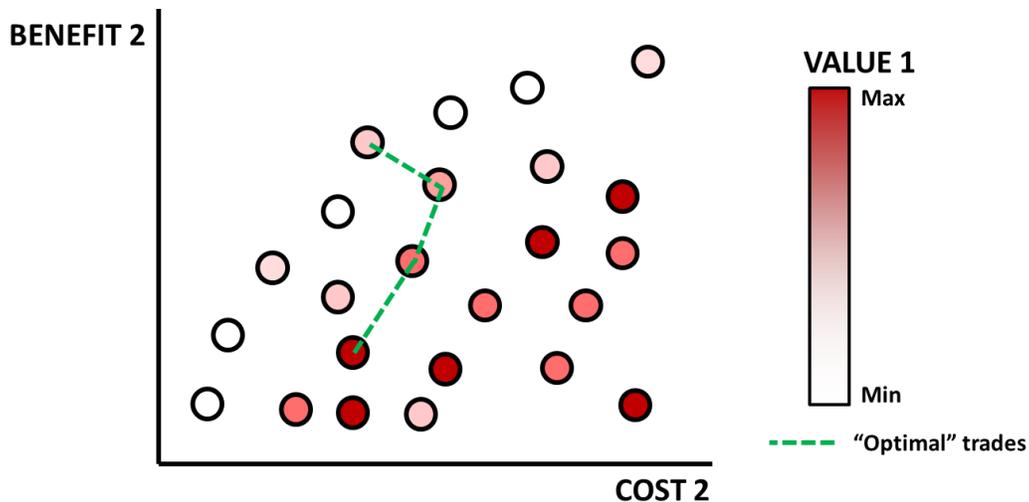
## **6.2.2 Single Stakeholder Value Models**

Stakeholder preferences are almost always plural: rarely do stakeholders invested in system design care about only one attribute. Analysis of multiple attributes is considerably more difficult than analysis of one, but is essential to accurately capture the value of an engineered system. The literature on multi-attribute decision making and value modeling is well-developed and widely used in the engineering field despite considerably varying strengths and weaknesses between the different methods (Hazelrigg, 2003; Ross, Rhodes, and Fitzgerald, 2015). This research assumes that the potential practitioners of MSTSE are capable of selecting an appropriate value model and able to simplify the value of a system for a stakeholder down to two dimensions: benefit and cost. These are the cornerstones of many tradespace exploration activities, particularly the standard MATE benefit-cost tradespace (Ross et al., 2004). Each stakeholder participating in MSTSE should be able to, with some assistance and potentially variable degree of fidelity, create benefit and cost value functions, aggregating into these two dimensions if they care about multiple attributes. With that established as the baseline situation, there are three deviations that are worth examining: (1) a stakeholder is unable (or unwilling) to define their preferences effectively as benefit and cost, (2) a stakeholder desires (and is able) to aggregate their benefit and cost preferences together into a complete, one-dimensional value function, and (3) a stakeholder naturally experiences only benefit *or* cost from the system.

The first situation presents additional challenges related to the feasibility of quickly evaluating the full tradespace. Without a defined benefit and/or cost function, a stakeholder must mentally perform the tradeoff calculations between multiple attributes when evaluating alternatives, adding considerably to their cognitive load. Additionally, the absence of benefit/cost functions precludes the use of many TSE metrics, including Fuzzy Pareto Number and Pareto Trace (Fitzgerald and Ross, 2012; Ross, Rhodes, and Hastings, 2009), that are useful for evaluating alternatives on a scale comparable between stakeholders: an important criterion

for any metric that may be used as a “fairness” metric in a negotiation. Without the ability to present a meaningful benefit-cost tradespace, additional emphasis should be placed on conducting pairwise comparisons between the BATNA and specific designs of interest (rather than the complete tradespace) in order to manage cognitive load. Accurately comparing alternatives to the BATNA is of paramount importance to making good negotiation decisions, and should therefore be conducted carefully without the assistance of individual benefit and cost metrics that allow multiple alternatives to be assessed as better or worse simultaneously. This will necessarily limit the number of designs able to be fully investigated, lowering the effective size of the tradespace and making some of the “small tradespace size” suggestions useful [6.3.3]. Overall, this particular limitation demands a back-to-basics approach to design evaluation and eliminates the option for more complex types of analysis.

On the other hand, sometimes the structure of the case is such that stakeholder has the means to combine benefit and cost into a single, all-encompassing value metric. Arguably the most common form of this action comes from monetization: in applications such as commercial product development where costs and benefits can be effectively and appropriately monetized, the total value of a project can be expressed as benefits minus costs (or perhaps benefit-to-cost ratio). Occasionally, other value models will be used to similar effect, for example an MAU or AHP function with “acquisition cost” as an attribute. Generally, combining benefits and costs is not advised for TSE as it aggregates away the ability to visualize the important high-level tradeoffs between benefit and cost that is a main strength of TSE and reduces the ability for stakeholders to act on unexpressed preferences in those tradeoffs (Ross and Hastings, 2005). However, the main appeal of this aggregation is that it simplifies the overall problem, negating the need for that stakeholder to perform any tradeoffs at all. Consider the reduction in complexity of a two-person negotiation in which Stakeholder 1 has combined their benefit and cost functions. Stakeholder 2’s tradespace can now easily incorporate Stakeholder 1’s value with a color scale (or another dimension), and the “optimal” tradeoffs would be the most preferred Stakeholder 1 alternatives at each fuzziness level of Stakeholder 2’s tradespace, as in Figure 6-2.



**Figure 6-2: Stakeholder 2 tradespace, with "optimal" trades against Stakeholder 1's value metric**

This makes one-dimensional value models better suited for a variety of other techniques, optimization foremost among them. When a stakeholder's value can be represented with a single number, it becomes considerably easier to optimize it against any number of other objectives than if it were still a benefit-cost tuple. The single value metric has a single dominant solution (i.e. the highest value) while tuples require the consideration of all non-dominated alternatives, greatly increasing the computational complexity and cost of optimization. In this situation, optimization can be leveraged to quickly and effectively find designs of interest even within extremely large tradespaces. However, it should be emphasized that the aggregation of benefit and cost in order to access optimization should be performed with caution, as generally the ramifications of optimization run counter to the goals of TSE, as it "black boxes" more of the exploration step by limiting stakeholder involvement. If optimization is used, it should be used explicitly in tandem with traditional TSE views (for example, tracing the path of the algorithm), using the optimization as a tool to further the participants' understanding of the functional relationships within the tradespace and between stakeholders.

Finally, in some cases one or more stakeholders may naturally have a one-dimensional value model by virtue of experiencing *only* benefit or *only* cost from the potential system. A common example of this phenomenon is the presence of a stakeholder who suffers externalities of the decision, such as an environmental organization dealing with the secondary impacts of a large infrastructure project. These stakeholders do not have a benefit-cost tradeoff to consider, though they may still balance multiple attributes in their chosen dimension with a multi-attribute function. The presence of a benefit-only or cost-only stakeholder has similar impact to the presence of an aggregated-benefit-and-cost stakeholder. By limiting their interests to a single value function, these stakeholders are prime candidates for the use of optimization as an aid to MSTSE. If all participating stakeholders are one-dimensional, for example if the military derives benefit from a project while Congress pays the costs, then a natural Pareto front of non-

dominated alternatives exists. These designs are promising candidates for selection and negotiation should begin with them.

### 6.2.3 Multiple Stakeholder Aggregation

Aggregating the preferences of multiple stakeholders into a combined set of preferences is a common tactic in many multi-stakeholder decision making methods (Gatti and Amigoni, 2005; Chen et al., 2004; Romanhuki et al., 2008). Its popularity is driven by the corresponding simplicity with which it treats the resulting problem: by combining all of the stakeholders into one, the problem has been reduced into a traditional single stakeholder problem which can be approached with any standard decision techniques. Unfortunately, there is no way in which to make this combined preference set without violating common axioms of rationality, as famously proved by Arrow (1963). Nevertheless, many researchers (and the practitioners that use their methods) choose to acknowledge but set aside this normative argument in favor of using an aggregated value model in order to have a simple means of generating group rankings of alternatives to be used as a prescriptive aid to decision making.

Despite the fact that group preference sets and corresponding value models are often accepted in prescriptive design, their simplicity hides the basic truth that the process of aggregation *necessarily* results in the loss of detail. In this context, aggregating the preferences of multiple stakeholders into a single function is a form of *pre-negotiation*. The balance of the preferences of the different stakeholders, through their importance weights and the aggregation function itself, is fixed. The function is doing the negotiation between the different interests and outputting the “best” alternatives. If any negotiation is to occur between the stakeholders, it must take the form of modifying the function rather than debating and selecting alternatives based on their relative merits for each individual<sup>11</sup>. This an unnecessary abstraction: why negotiate a function to choose an alternative, even though it is already known a rational group function doesn’t exist, if there is a viable means to negotiate on the alternatives themselves? It can seem daunting to find the best choice amongst thousands to millions of alternatives, favoring the use of a group ranking, but MSTSE seeks to supply the structure necessary to enable the consideration of that large number of alternatives without resorting to aggregation. The alternatives are realistic design concepts with corresponding design variables and performance attributes, grounding the discussion in objective data rather than a subjective aggregation function. MSTSE should ideally be conducted with the “subjective math” of value modeling limited entirely to individuals, allowing each stakeholder to define their own subjective values uncontested and negotiating *between* them rather than negotiating to *combine* them.

---

<sup>11</sup> Note that negotiation is not precluded by the use of an aggregated function. In fact, it likely *will* happen if the function gives a “best” alternative that is not immediately attractive to all parties and the differences must be settled by returning to the individual value propositions of each stakeholder. If that has to happen, the aggregation function served little more purpose than a contentious screening metric for potentially interesting alternatives.

With all of that understood, there are at least two structural considerations of a multi-stakeholder problem that merit the consideration of an aggregated preference set. First, as previously discussed, the presence of stakeholder coalitions can make aggregating preferences an attractive means of simplifying the problem without too significant a loss in information. Stakeholders with interests similar enough that tradeoffs between them are slight may be able to combine their preferences with minimal loss of information and reduce the number of value models that need to be visualized/explored, while still retaining the central advantages of MSTSE when negotiating between coalitions. Second, in some circumstances, stakeholders that are tightly connected (e.g. multiple teams internal to a single company) can be aggregated rationally on axiomatic grounds (Scott and Antonsson, 2000). In this case, their preferences can be combined due to the logical structure defining their relationships even if those preferences are significantly different. The end result is functionally identical to a coalition of stakeholders with a collective benefit and cost value function. Note that if *all* participating stakeholders are combined in this way, the problem is reduced to a standard TSE activity with a single benefit-cost tradespace. The negotiation has still occurred in the definition of the collective value model but all standard, single-stakeholder TSE techniques can be deployed on the result.

#### **6.2.4 Unoperationalizable Preferences**

Operationalization is the act of defining the metric used to measure an abstract concept (Babbie, 2009). Stakeholders sometimes have preferences on attributes that are difficult or impossible to operationalize. In fact, supplying values for unoperationalizable variables is one of the most important roles of “subject matter experts” and “senior decision makers” as distinct from regular analysts: these people have highly specific, internalized knowledge that often can’t be modeled or otherwise calculated without their direct involvement. These attributes are typically of the “soft” variety (e.g. political capital, public approval) though they may occasionally be complex or cutting-edge physical attributes beyond the abilities of a model to measure. The challenge presented by unoperationalized attributes is that they cannot readily be incorporated into value models, because their measurement is either undetermined (there is no universally accepted, objective measurement) or impossible (simulations lack the required fidelity or the underlying theory is unknown). There are two courses of action for quantifying an unoperationalized variable to be used in a value model:

1. Assess with qualified outside source
2. Assess internally by the interested stakeholder

The first option gives the variable the most objective possible position, in accordance with the general negotiation principle of using objective data (Fisher, Ury, and Patton, 1991). Having a trusted, official organization or consultant assess the attribute for the alternatives in the tradespace gives the data credibility and eliminates the chance of personal bias of an involved stakeholder unduly influencing the results. When this is not possible, such as if the attribute is a matter of personal opinion, the stakeholder must evaluate the alternatives himself. This results in

the same functional end state but introduces a subjective element that can be contested by other stakeholders.

For an attribute to be included in a value model, *each* alternative must be evaluated for that attribute. This means that larger tradespaces have more difficulty incorporating unoperationalized variables in the value model, because they require considerably more human input in order to evaluate every alternative. If assessing the attribute for all of the alternatives in the tradespace is infeasible, the attribute must be left out of the model and become a metapreference that is assessed only for alternatives identified as designs of interest [6.2.5].

### **6.2.5 Metapreferences**

Metapreferences are preferences held by the stakeholders participating in MSTSE that are not captured by the assigned value model that calculates the value they credit to each alternative. Metapreferences are a reality that must be accounted for, as constructed value models are ultimately an approximation of the true value statement of a stakeholder. The “eyeball test” and “gut reaction” are common phrases that simply indicate that stakeholders have evaluative criteria that remain unexpressed by the output of the numerical value model. Sometimes the metapreference is a result of inability to include all relevant performance attributes in a benefit or cost function, such as when an attribute is left unoperationalized [6.2.4] or has simply been omitted as a consequence of preference uncertainty [6.4.3]. In other situations, metapreferences can be held for properties of the system relative to the space of alternatives, such as a preference for low-cost, low-benefit solutions over high-cost, high-benefit solutions. Even further than that, some stakeholders may have preferences on the means of value delivery; for example, a stakeholder may decide between two alternatives with nearly identical benefit and cost by favoring a passive solution over a changeable solution or a short-term, low-risk solution over a longer lasting solution that is susceptible to uncertainty [6.4.4].

Given the wide range of potential forms that metapreferences can take, it is difficult to suggest particular MSTSE visualizations or activities that broadly incorporate them in the decision making process. The difficulty is compounded by the fact that metapreferences are occasionally unknown to the stakeholders *a priori*, but rather are revealed over the course of the exploration of the tradespace through the gradual realization of additional evaluative criteria with which each stakeholder is weighing the alternatives. Recording metapreferences as they are revealed, either in the problem formulation as preferences uncaptured by the value model or later on while negotiating, is an important step towards keeping an open dialogue between the participants, heading off potential disagreements due to unstated preferences and keeping all parties informed of each other’s interests. In cases when metapreferences dominate the discussion, the details of the value models become less instructive because the models themselves are less indicative of how the stakeholders are actually evaluating alternatives. For this reason, exploration and visualization of the value models themselves (e.g. visualizing a utility function) and their relationships (e.g. stakeholder correlations) should be correspondingly

de-emphasized in favor of other activities that more directly target the given metapreferences (e.g. –ilities or uncertainty analysis).

## 6.2.6 Divisible Attributes

In some cases, attributes may be divisible between stakeholders. Consider a two-stakeholder system in which both stakeholders A and B use a shared lifecycle cost model to estimate the total cost (in dollars) of any given alternative, and differ only in their benefit functions. Each stakeholder would have a benefit-cost tradespace on the same cost axis, but the assessed cost is *total* cost when in fact that cost will be shared amongst the stakeholders. If the cost share is evenly split between the two, each tradespace would look the same but with a linear transformation on the cost axis (individual cost is one-half the total lifecycle cost). However, in the absence of other constraints, there is nothing preventing the stakeholders from *unevenly* splitting the cost of a given alternative as a function of the relative value for each stakeholder.

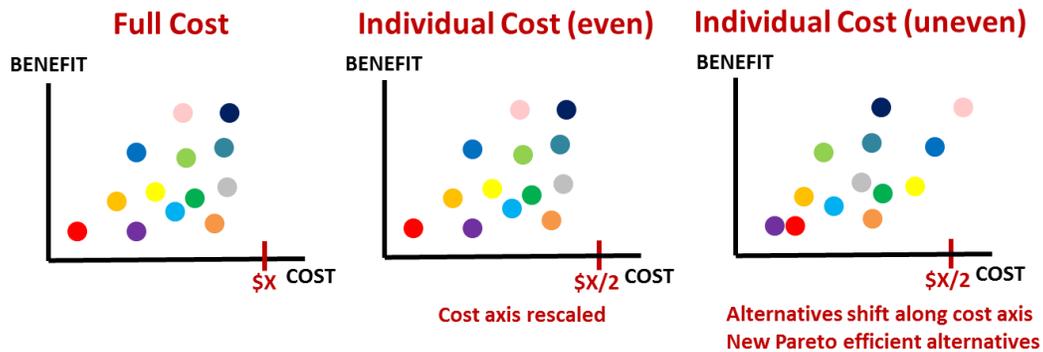


Figure 6-3: Tradespace differences between full cost, even split cost, and uneven split costs

This presents an interesting wrinkle to standard tradespace exploration in that there is effectively an even larger set of alternatives than those enumerated and evaluated by the performance models, defined by additional negotiation-driven agreements introduced in the jump to MSTSE. Monetary costs are the prime example of this concept by virtue of money’s liquidity, but other types of attributes can be shareable as well, including benefits such as “mission uptime” which can be divided between different stakeholder missions. There are two basic approaches to capturing these alternatives. First, an additional design variable can be introduced explicitly into the problem formulation, increasing the size of the tradespace and allowing for the negotiation to trade on the division of the attribute. To match the previous example, a variable “fraction of cost paid by Stakeholder A” could be added, with a discrete enumeration of {0.33, 0.5, 0.67} to capture a range of cost splits from one-third to two-thirds the total cost for each stakeholder. This approach is straightforward, requires little additional work, and potentially will introduce more “fair” solutions to the tradespace by more appropriately splitting costs or benefits among the stakeholders when total value is uneven.

However, the addition of a new design variable multiplies the total size of the tradespace. With unlimited computing, this is a non-issue, but for practical application it may be impossible or undesirable to further increase the number of alternatives evaluated. Additionally, for every newly interesting design added to the tradespace, a correspondingly *less* interesting design is also added (e.g. if the 0.33 split is more fair, the 0.67 split is almost certainly unacceptable), resulting in a waste of computational resources. This is the key motivation for the second approach for divisible attributes: assigning an appropriate or optimal division with a model or algorithm. This approach “solves” the problem of dividing the attribute according to a specified rule. The rule could be a simple addition to the basic performance/cost model, such as “Each stakeholder will pay a proportion of total lifecycle cost equal to the ratio of their mission uptime to total uptime of all stakeholders”. Alternatively, splits can be performed algorithmically with respect to the entire tradespace. Returning to our continuing example, let’s say Stakeholders A and B agree to use Fuzzy Pareto Number as a measure of design efficiency and wish to split the system cost such that the FPN is equal for each of them<sup>12</sup>. An algorithm can progress in the following way:

1. Determine cost split versus current Pareto Front that results in equal FPN for each stakeholder for a given alternative
2. Repeat (1) for each alternative in the tradespace
3. Recalculate Pareto Front
4. Repeat (1-3) until split proportions converge between iterations within a set threshold

This results in a tradespace where each alternative has the same FPN for both Stakeholder A and B. The divisible dimension of cost has been leveraged to make each alternative “fair” with regards to benefit-cost efficiency, in the process rearranging the design points by sliding them along the cost axis and generating a new Pareto Front. These new Pareto Front alternatives are all jointly Pareto efficient and thus are strong candidates for negotiation and eventual selection. The power of a fully divisible attribute to create mutually beneficial solutions is considerable, especially when (as in this example) it represents the entire benefit or cost function for a stakeholder. Divisible attributes that are only a part of a multi-attribute benefit or cost function can be similarly leveraged but without the guarantee that an even split exists, as it may not have the necessary weight to balance the FPNs.

If this algorithmic response to divisible attributes is deemed superior to the additional design variable approach, the negotiation between the stakeholders occurs in the definition of the target criterion of the algorithm. Equalizing FPNs is only one potential definition of a fair

---

<sup>12</sup> It is a mathematical certainty that at least one cost split exists for which this is true for every alternative. Because FPN varies continuously and any reduction in cost/FPN for one stakeholder increases cost/FPN for another, the FPNs will eventually intersect. Finding the solution for two stakeholders is straightforward, but it applies to any number of stakeholders. There is no a priori way to select a “best” cost split if multiple equal-FPN divisions exist, though common sense might suggest that the split closest to even is the fairest.

outcome, and each stakeholder may have a different conception of what is fair. Making assumptions explicit early in the problem formulation is another supporting tenet of principled negotiation, and thus reaching an understanding of what is acceptable for all parties is a desirable goal. If no agreement on a target criterion for the algorithm can be reached, the design variable approach can be used as a fallback to prevent impasse by delaying discussion of fairness until tangible alternatives are considered.

All of the preceding discussion on divisible attributes presupposes that more than one stakeholder holds a preference on the given attribute. We have already used monetary costs as an example of a divisible attribute, but imagine a manufacturer-customer negotiation where the manufacturer naturally cares about production costs and the customer about operations costs. These two attributes of the system are both monetary, and therefore divisible, but fundamentally experienced by only one stakeholder each. To leverage divisibility in this situation, some sort of *exchange rate* needs to be established, capable of transforming one attribute into an equal value of another. In this example, the two stakeholders might agree to value a production-cost-dollar as 1.7 operation-cost-dollars since the latter cost is delayed and therefore has less opportunity cost. Other possible means of establishing an exchange rate include using a nonlinear schedule (e.g. a 2:1 exchange rate up to X production dollars, then 1:1 after) or defining the exchange rate on an alternative-by-alternative basis within the model based on other system attributes (e.g. expected lifetime may impact the exchange between production and operations costs). Creating an exchange rate is an act of negotiation itself, as at its most basic level it involves multiple stakeholders agreeing to a single valuation of two or more attributes<sup>13</sup>. This may mean it is impossible to establish an exchange rate that all parties agree upon, negating the ability to use divisible attributes as “side payments” to improve an alternative. This task belongs in the problem formulation: if multiple divisible attributes are identified in the formulation, an attempt can be made to agree upon an exchange rate. If so, the algorithmic redistribution of the attributes for each alternative can be used before beginning the tradespace exploration. If not, it is better to ignore the divisibility and instead return to the “base” case where each stakeholder’s values are distinct and separate rather than engage in a drawn out negotiation of the rate.

### **6.3 Alternatives**

Alternatives are the available choices for the stakeholders to negotiate over. This includes the design points in the tradespace and the BATNAs for each stakeholder: the no-

---

<sup>13</sup>It is possible to let each stakeholder define their own exchange rate, and then perform MSTSE looking for alternatives that are satisfactory to each party using each *other* exchange rate, imitating the classic Zone of Possible Agreement (ZOPA) (Raiffa, 2002). This leaves the negotiation of the final exchange rate to the end, after design selection, with the understanding that some exchange rate between the endpoints must be acceptable. However, this risks expending effort on MSTSE to choose a design that may ultimately be discarded by the additional exchange rate negotiation at the end. It also increases the complexity of the MSTSE task (using multiple exchange rates while exploring) and becomes particularly unwieldy with more than two stakeholders.

agreement alternatives. Understandably, the set of available alternatives in a given decision problem can impact the course of the negotiation. The BATNA influences each stakeholder’s bargaining position and the attractiveness of other alternatives; thus, its *type* and *quality* relative to the tradespace are of significant import. The construction of the tradespace itself around those BATNAs creates the “playing field” of the negotiation, which can proceed differently depending on both the *size* and *completeness* of the alternative set.

### 6.3.1 BATNA Type

Stakeholders come into negotiations with a wide variety of positions, and for that reason BATNAs come in many different forms. The type or source of the BATNA can – and should – impact the goals of the stakeholder, changing the types of MSTSE activities that are most pertinent. Here, we will go over some key BATNA types (do-nothing, build preferred alternative alone, use existing system, and other opportunities) and the logical responses to them.

#### 1. Do-nothing (exploratory negotiation)

Some negotiations are performed strictly for exploratory purposes. The participating stakeholders may agree to perform MSTSE in order to see what options are available if they chose to work together, with the understanding that the potential resulting system is evaluated against the alternative of doing nothing (distinct from *continuing with an existing system*, covered later). The do-nothing BATNA has zero cost and zero benefit. For most value models, every alternative will have positive cost and positive benefit, placing the entire tradespace in Quadrant 1 on BATNA-centered axes, as in Figure 6-4. An example exception would be a monetized benefit function, which could become negative for some alternatives depending on the defined zero points and support of the underlying function, but we will focus on the all-Quadrant-1 format, which best captures the idea of working together at some cost to achieve some benefit. Using the BATNA-centered negotiation tradespace can serve as a consistent reminder that the alternatives being discussed are all tradeoffs of the same type: more cost, more benefit. Each stakeholder must decide for themselves if the added benefit of a potential design is worth the cost.

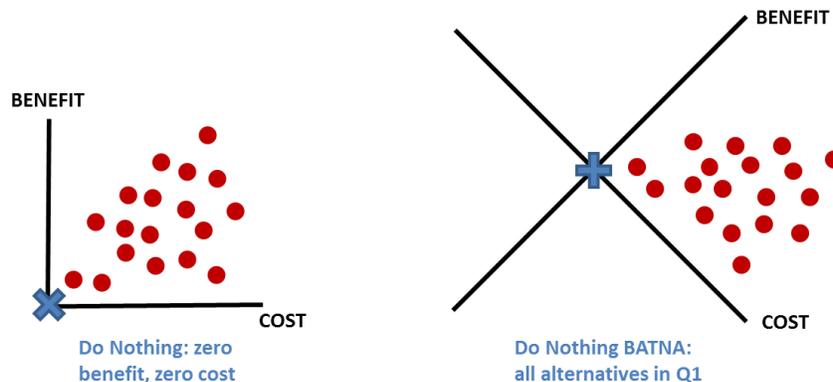


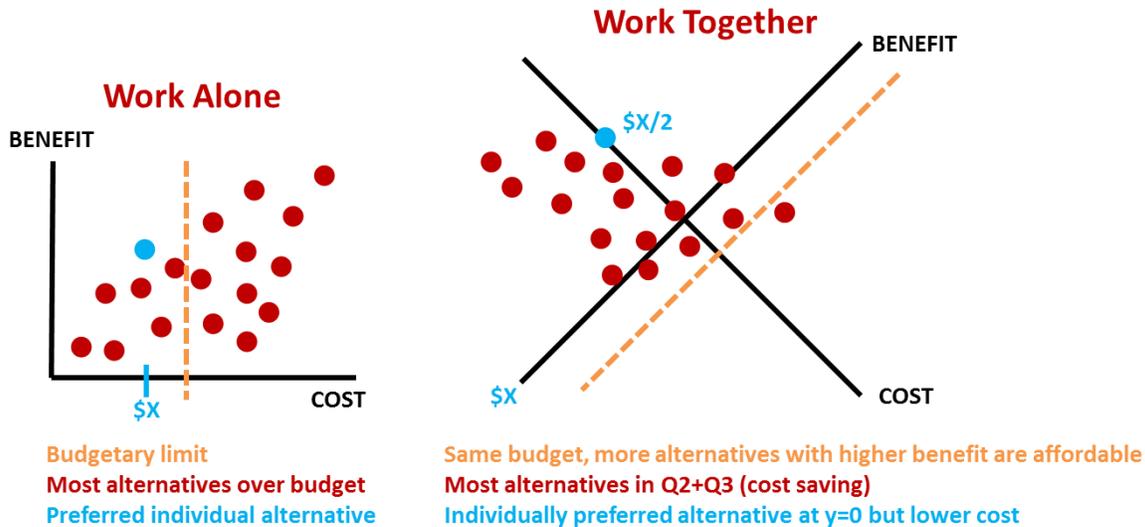
Figure 6-4: The Do Nothing BATNA results in all alternatives in Quadrant 1

A defining feature of exploratory analysis is that it is purely hypothetical, lowering the perceived cost of failure and freeing up creativity and intuition in the design process in a way conducive to learning (similar to the “four freedoms of play” theory of Osterweil, 2010). Extending this to negotiation, we could expect similar benefits, where the “best” alternative is now the best that can be agreed upon. In a multi-stakeholder context, the nature of exploratory analysis offers a relatively low social cost of withdrawal. The agreement can be made if it looks good, and if not the parties can simply amicably separate. Without a pressing need to continually evaluate the alternatives in the tradespace against a BATNA in order to maintain situational awareness of the negotiation, MSTSE approaches standard TSE in terms of the individual tasks stakeholders perform. A high emphasis on exploration and evaluating tradeoffs is the preferred approach, as these tactics are the best at revealing the underlying structure of the problem and educating stakeholders about the relationships between design decisions and value. The fact that all the alternatives are located in Quadrant 1 also makes traditional visualizations of the tradespace more appropriate than in the general negotiation case, because the standard axes of absolute benefit and absolute cost also function as differences from the BATNA when the BATNA is located at (0,0). This can be seen on the left side of Figure 6-4, and is why the negotiation tradespace was extended to include an option to “un-rotate” the tradespace as a result of feedback from the practitioner interview of chapter 5.6.

The do-nothing BATNA is also the only type of BATNA that can realistically give multiple stakeholders the *same* bargaining position. Other BATNAs are individual solutions that involve only one stakeholder and are evaluated on their personal benefit and cost value functions, making it exceedingly unlikely that more than one stakeholder has a BATNA in the same relative position in the tradespace. If all of the participants in MSTSE have a do-nothing BATNA, this gives them all the same initial position and removes the possibility of one stakeholder leveraging a stronger BATNA against the others. Cases of this type are strong candidates for algorithmic approaches to “fairness” such as using side payments to balance efficiency [6.2.6], as the stakeholders have less claim for preferential treatment when their positions are *a priori* equal.

## **2. Build preferred alternative alone**

Some negotiations are conducted between stakeholders whom otherwise would be performing similar design tasks, searching for the opportunity to save cost by sharing the resulting system. For example, science missions conducted in orbit can be launched on their own satellite, or stakeholders can negotiate the design of the satellite in order to accommodate multiple missions and share production and launch costs. Sometimes in these cases, the BATNA for each stakeholder is their individually most preferred alternative in the tradespace at its *full* cost, rather than the reduced portion of the cost. This creates an interesting opportunity to use individual TSE on the dataset as a means of identifying the BATNA during problem formulation.



**Figure 6-5: "Work Alone" tradespace identifies BATNA, "Work Together" tradespace highlights newly affordable alternatives**

Sharable costs imply divisibility in at least one attribute [6.2.6], which is what opens up the possibility of balancing the costs between the stakeholders. If the costs are to be split in some fixed rate across alternatives, then the tradespace can be examined as if it were a single-stakeholder TSE exploration for the purposes of the identification of the BATNA. Then, the negotiation tradespace will be cost-shifted, with all alternatives reducing in cost equally but the BATNA remaining in the same place. This is pictured in Figure 6-5, where the preferred cyan design defines the BATNA in the individual tradespace on the left and shifts to half its original cost in the negotiation tradespace on the right. Other designs that were previously unaffordable (to the right of the dotted line indicating the budget) become affordable in the negotiation tradespace and possibly superior to the original preferred design.

If the cost-sharing scheme is not fixed across alternatives, then the work-alone tradespace should use a full-cost axis for the identification of the BATNA, which may rearrange the relative position (and therefore appeal) of some designs. Unless the BATNA is an exceptionally low-cost alternative or multiple orders of magnitude of cost are considered, we can expect that the considerable majority of alternatives in the negotiation tradespace will occupy Quadrants 2 and 3 in this case. Saving cost is the driving force of the negotiation, and the alternatives considered will heavily lean towards that half of the tradespace, driven by the full-cost BATNA.

If individual TSE is used to locate the BATNA, it is likely that the BATNA will be on the Pareto front of the full-cost tradespace. It is important to retire the full-cost tradespace as soon as the MSTSE activity begins, as this could unintentionally reinforce anchoring on the Pareto front in the shared-cost tradespace. There should be a clear divide between the BATNA-finding task and the actual negotiation, as the top-level goal has changed from “maximize individual value” to “find a mutually beneficial solution”. It is also worth increasing emphasis on the

consideration of other participants' BATNAs, especially if the participating stakeholders have a cordial relationship or other long-term interests in working together. Failure to reach an agreement may be acceptable for one stakeholder but cause considerable distress in the budget for another, potentially even reaching disastrous consequences for some parties if no alternatives are affordable to them at full-cost. Drawing attention to this fact can delay premature withdrawal from the negotiation by stakeholders with viable BATNAs if they hold any value in the relationship; this is applicable regardless of BATNA type but a common scenario for this kind, in which the desired solutions are bigger and better than what can be implemented separately.

### **3. Existing system**

This situation arises when a stakeholder is already performing/accomplishing the targeted objectives, but is looking to acquire a new, superior system. A simple example of this would be a person who is interested in buying a car to improve his work commute over his current solution of riding public transportation. To some extent this is the most "obvious" type of BATNA, as the stakeholder is already using it. Existing systems are relatively easy to assign as BATNAs because the relevant performance criteria are typically already known, making it trivial to calculate the benefit/cost position of the BATNA. Given that the stakeholder's current experience is with the BATNA, it also should be a natural reference point for decision making (more so than a hypothetical BATNA). Thus, we can expect that the BATNA-centric techniques and visualizations presented in this research will be more immediately accessible and desirable for stakeholders with this type of BATNA. As demonstrated in our negotiation experiment in chapter 4, this increased stakeholder focus on the BATNA results in reduced attachment to individual Pareto fronts, but can lead to hill-climbing strategies that fail to explore tradeoff opportunities in Quadrants 1 and 3 of the BATNA-centered tradespace. Care should be taken to reinforce the idea that the "replace existing system" task can be completed successfully without necessarily improving in both benefit and cost.

The enumeration of the tradespace has a considerable impact on the type of discussion that will happen with an existing system BATNA. Because existing systems can vary dramatically in quality from nearly unacceptable to excellent, it is possible that the majority of alternatives in the tradespace could be in any of the four quadrants. If the stakeholder wants to focus on either lower cost solutions or higher benefit solutions, the chosen levels of the design variables should reflect this by targeting those areas of the tradespace. The relative quality of the BATNA to the alternatives is further discussed in the following section [6.3.2].

One variant of the existing system BATNA occurs when the existing system is a member of the tradespace. This occurs when the design task is a *modification* rather than a *replacement* of the existing system, and it is logical to simply include the current status of the system as an alternative. Large-scale projects such as infrastructures are common examples of design tasks that are revisited over time for modification. In this case, the design variables that define the alternatives also fully describe the BATNA and it is therefore easy to include in the tradespace

enumeration. The inclusion of the BATNA as an alternative may further anchor the stakeholders on the BATNA, with the corresponding positive and negative impacts. Additionally, the nature of designing for incremental improvements is such that small changes are often attractive, making the area of the tradespace immediately around the BATNA of primary interest.

#### **4. Other opportunity**

Occasionally, the BATNA for a particular stakeholder is to simply direct their resources to another opportunity entirely. As a notional example, the Navy and Air Force may be negotiating how to design and share the costs of a new reconnaissance UAV, while the Navy holds a BATNA of building themselves a new surface combatant. The limited budget of the Navy precludes the ability to build both systems, despite the fact that they accomplish completely separate objectives, and they will choose to work alone on a new ship if the UAV negotiations fail to satisfy them.

The “other opportunity” BATNA type is the most difficult to capture. There are potentially infinite opportunities available to stakeholders, requiring strict scoping in order to make the evaluation of those BATNAs manageable. Complex design tasks (like the ship design in the above example) may require considerable analysis on their own to determine the best alternative within them, delaying the negotiation; or if the analysis expense is deemed unwarranted just to use as a BATNA in a negotiation, the BATNA is forced to be a rough approximation of value. Additionally, because the other opportunity targets a different set of objectives (if it didn’t, it would simply be “build preferred alternative alone”) the value of the BATNA must be “translated” into the benefit and cost functions used to evaluate the tradespace. This can be challenging when the other opportunity cannot be evaluated using the operationalized attributes of the negotiated system. There is a considerable advantage to using universal value models (e.g. monetization and cost-benefit analysis) if the appropriate BATNA for a stakeholder is a distinctly different opportunity, as the value of the BATNA can presumably also be reduced to these terms in order to serve as a reference point for designs in the tradespace, regardless of its intrinsic performance attributes. If monetization is undesirable, the BATNA’s location in the tradespace can be approximated by searching for alternatives for which the stakeholder is indifferent between them and the BATNA and estimating a box-range (minimum and maximum cost and benefit), as shown in Figure 6-6.<sup>14</sup>

---

<sup>14</sup> Alternatively, the use of utility theory as a value model can provide a useful baseline for simply estimating BATNA value using the idea of lottery equivalents, similar to the way a utility function is traditionally elicited. For example, the stakeholder can be posed binary questions such as “Would you prefer to pursue another opportunity or to spend \$X for a 50% chance at Y utility and 50% chance at 0 utility?”. Intelligently sequencing these questions can provide an effective description of the stakeholder’s estimates of probable benefit-at-cost for the “other opportunity” BATNA, and then the “best” benefit-cost pairing can be picked as BATNA by the stakeholder. Ultimately however this method is not recommended: lottery equivalents can be difficult to understand for beginners and the best case outcome is no more detailed than a regular stakeholder estimate.

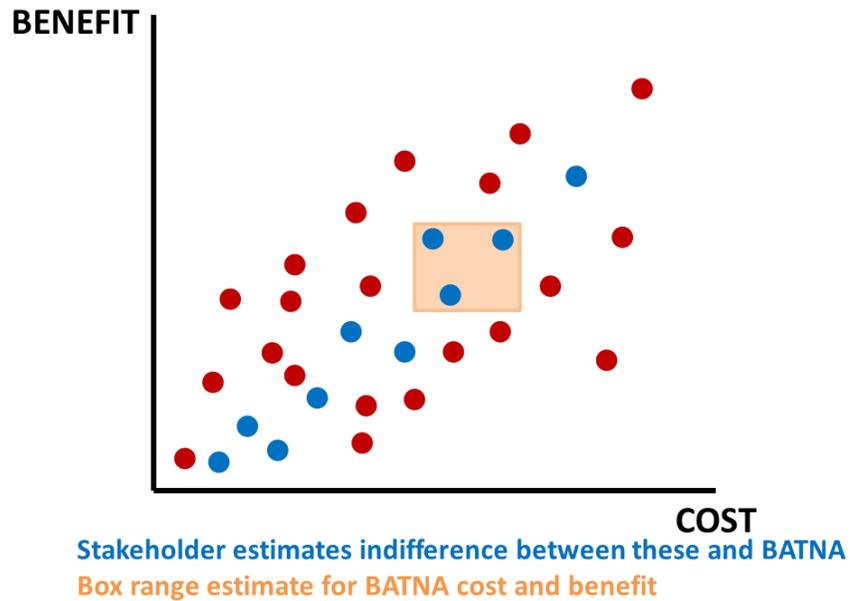
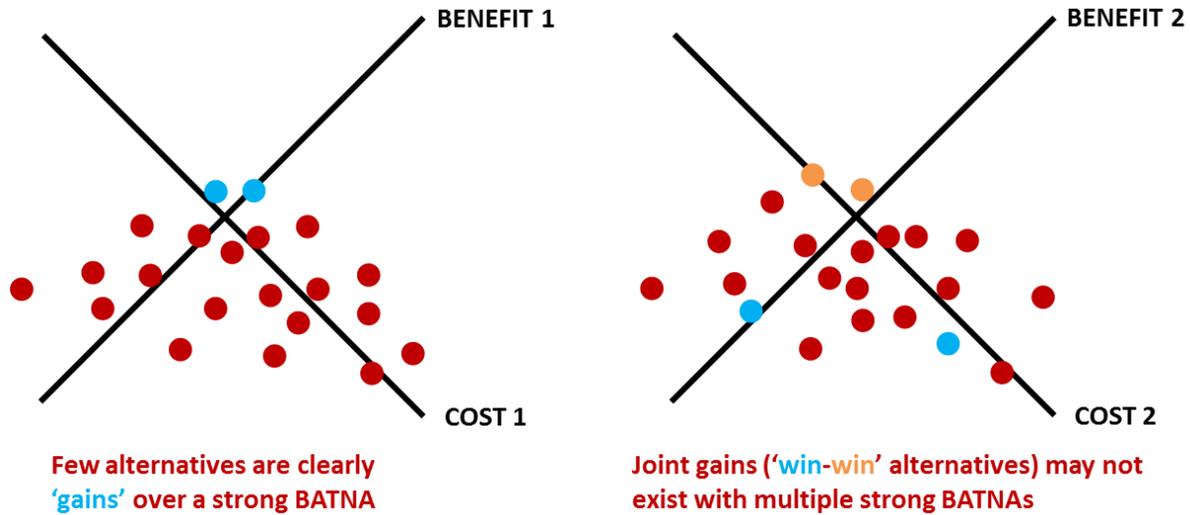


Figure 6-6: Estimating an unknown "other opportunity" BATNA

### 6.3.2 BATNA Strength

The relative strength of a BATNA compared to the alternatives in the negotiation tradespace has obvious impacts on the perceived value of the alternatives, particularly when using MSTSE visualizations designed to emphasize the BATNA. Returning to the language of Prospect Theory, stakeholders with poor BATNAs will see the majority of the tradespace as “gains” to be had from negotiation, while stakeholders with strong BATNAs will see most alternatives as “losses”. Efforts to emphasize the BATNA in tradespace visualizations attempt to align *perception* of gains and losses with *reality*, but the *reality* itself can vary from case to case.

Because the goal of MSTSE is to locate mutually beneficial alternatives in the tradespace, most attention will be paid to “gains” alternatives. At the most basic level, an extremely good BATNA can reduce the effective size of the tradespace by rendering only a small fraction of the designs “gains”, as shown in Figure 6-7.



**Figure 6-7: Clear 'gains' are limited by a strong BATNA, exacerbated by multiple strong BATNAs**

While forcing stakeholders to actively consider “losses” is not likely to be a productive strategy because those alternatives won’t simply *become* attractive due to the lack of “gains”, it is important to maintain situational awareness of the entire tradespace if the benefits of TSE such as insight generation are to be captured. Cases with strong BATNAs resulting in few “gains” alternatives should refrain from analyses resulting in early “zoom-in” on designs of interest (Pareto optimality, best performers in individual attributes, etc.), instead starting with analysis of the tradespace as a whole (functional relationships, stakeholder correlations, etc.) as to not preemptively focus on the smaller subset of designs that are good for the strong-BATNA stakeholder. By exploring the tradespace at large, it may become apparent that additional alternatives can be enumerated that would expand the tradespace further into the gains domain, increasing the number of attractive alternatives for the more difficult to please stakeholders. Even if no additional alternatives can be identified, the substance of the discussion between stakeholders can build into the discussion of mutually desirable designs rather than beginning with individually desirable designs and then combining analyses and “sacrificing” value.

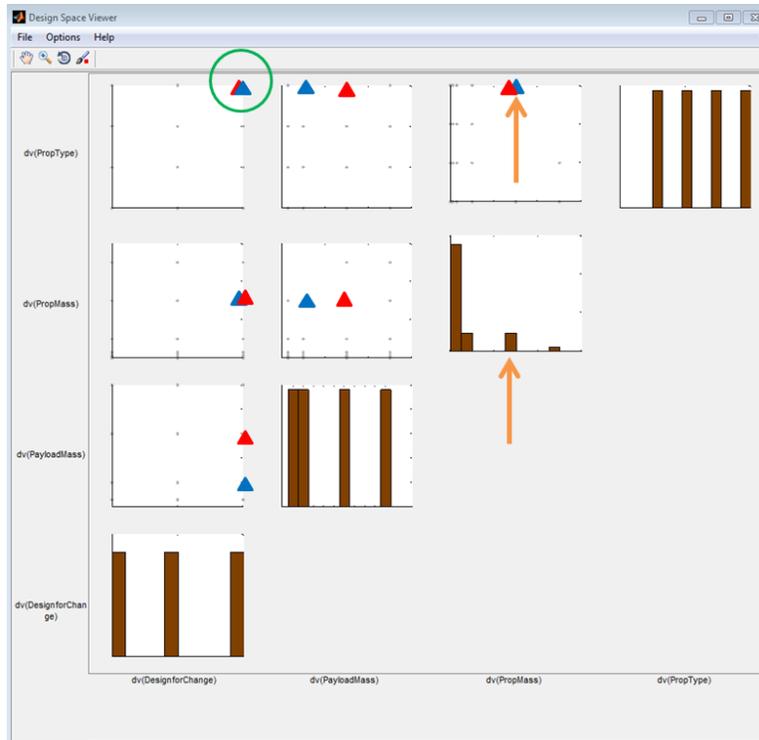
On the other hand, a poor BATNA for a given stakeholder will make the majority of the tradespace attractive, making working in the positive “gains” domain with standard human exploration techniques relatively easy. If some stakeholders have strong BATNAs, these weak-BATNA stakeholders are not likely to be the “dealbreakers” in the negotiation and, if desired, this characteristic can be leveraged with algorithmic exploration techniques. Although human exploration of the tradespace is still encouraged, the high value of most alternatives (relative to the BATNA) means that the output of an optimization is likely to be satisfactory for these stakeholders even if the optimizer is operating on other value models. For example, as long as their value models are not strongly divergent, it is likely that one of the five to ten “best” designs for a strong-BATNA stakeholder generates considerable value for the weak-BATNA stakeholder as well. The use of tactics such as this can quickly narrow down the tradespace to highly

attractive alternatives, but must be performed carefully so as not to cast these stakeholders as “second-class citizens” of the negotiation. It does, however, accurately capture the power dynamics in a traditional negotiation between stakeholders with BATNAs of highly different quality and can be used when egos are not likely to be roadblocks to agreement.

### **6.3.3 Tradespace Size**

The size of a tradespace (its number of evaluated alternatives) is arguably its most critical characteristic when used in a negotiation setting, where exploring many options is considered to be a success-driving criterion (Fisher, Ury, and Patton, 1991). The sizes of tradespaces can vary dramatically from case to case depending on the desired scoping, available information, and computing power, among other factors. For this reason, attempting to designate a specific minimum number of alternatives for conducting MSTSE is ultimately unlikely to be a useful outcome. Rather, the tradespace should be sized appropriately to the problem, as it would be regardless of the intent to negotiate. This will limit the possibility of overextension: creating unnecessary design variables or discretizing them too finely, strictly for the purpose of increasing tradespace size, which can dilute the informational value of the tradespace. Once the tradespace has been created, the stakeholders or analysts can consider tailoring the MSTSE process if there are particularly few or many alternatives.

With fewer than approximately 20 alternatives, the label of “tradespace exploration” begins to lose some of its meaning. Applying tradespace exploration techniques to a design space this small resembles a more traditional “trade study” between individually interesting point alternatives (Parnell, Cilli, and Buede, 2014; Cilli and Parnell, 2014). In fact, tradespaces this small are often enumerations of specific, purchasable systems rather than bottoms-up design tasks. For example, a small tradespace can be constructed by comparing a subset of the different commercially available UAVs even though a tradespace enumerated by discretizing the main aircraft design variables (span, chord length, airfoil, thrust, etc.) could easily number in the hundreds of thousands. Given that the key strengths of TSE lie in its ability to generate insight across large and diverse design spaces, MSTSE may not be the correct choice of negotiation support if the design task is well described by a number of alternatives small enough to be conceivably evaluated and compared manually. However, scaling up into a larger design space and utilizing MSTSE to explore it is a potential means of breaking deadlock in a negotiation that is overly constrained by a limited set of alternatives. This is also true for negotiations based on small tradespaces of a hundred to a few thousand alternatives. Visualizations that provide information regarding the composition of the design space, like Figure 6-8, can be used to identify target areas for further enumeration of alternatives, either by finding gaps in the enumeration scheme or locating “edges” of the design space where high value alternatives are grouped.



Shared region of interest on boundary of design space in PropType and DesignforChange dimensions  
 Current enumeration on preferred PropMass level is limited

**Figure 6-8: Linked scatterplots of design variables, with stakeholder favorites marked in blue and red triangles**

As tradespaces become large, considering millions of alternatives or even higher orders of magnitude at once, the chief concern is that of computational and cognitive limitations. Though it is feasible with modern supercomputing to evaluate trillions of alternatives, visualizing that information effectively remains an open challenge (Thomas and Cook, 2006; Keim et al., 2008). Many traditional tradespace views like the benefit-cost scatterplot are not feasible with this quantity of alternatives, limited by screen area and pixel density to show smaller amounts of data. Additionally, the matter of responsiveness is also critical in a negotiation setting, where interaction between stakeholders is expected to generate numerous “questions” that are investigated in the tradespace. Generally, interactive software can degrade user experience if the latency in visualizations is over just one second (Bryson, 2011; Liu and Heer, 2014), which can be a daunting proposition when handling such large amounts of data. Addressing these issues of speed and representation is a topic of current research in visual analytics, and future improvements in both visualization theory and technology may ultimately make TSE of trillions of alternatives an attractive approach, but currently it is feasibly limited to millions without stretching the capability of existing technology. Even millions of alternatives can present cognitive challenges to stakeholders, particularly those with little to no experience with TSE. Limiting activities that identify specific designs of interest can help to limit

opportunities for “paralysis by analysis”, especially for large tradespaces that could have hundreds of designs or more called out as valuable by screening metrics, focusing instead on identifying broad characteristics of mutually beneficial alternatives (regions of the design space, important attributes, etc.) before narrowing down near the end. Alternatively, the tradespace can be reduced in size to make it more easily managed, using techniques such as Design Value Mapping (Richards et al., 2009) to identify the design variables with the least impact and remove them or reduce their discretization.

#### **6.3.4 Tradespace Completeness**

Here we will call a tradespace *complete* if it contains *all* possible alternatives. Most tradespaces are not complete, and it is difficult to guarantee completeness due to the possibility of accidentally omitting relevant design variables in the problem formulation. Additionally, the discretization of continuous design variables by definition precludes the ability to explore all possible alternatives, making tradespaces composed entirely of discrete variables the only candidates for true completeness. Practically, these concerns are addressed through due diligence in problem formulation, ensuring that the proper design variables are included and discretized at the appropriate level of fineness. Even with an acceptable enumeration of the design space, completeness can be compromised further by computational limitations. If the evaluative performance model(s) for the system are computationally intensive, it may not be feasible to evaluate the full-factorial enumeration of the tradespace, which increases in size combinatorically with the number of design variables and the number of levels of each variable (Ross, 2006). Random sampling or Design of Experiments methods such as Latin hypercubes (which guarantee some amount of even distribution of alternatives throughout the design space) can be used to downselect the design space to the desired size, at the cost of completeness.

The appeal of a complete tradespace is that the best alternative in the tradespace is the best alternative of all, addressing one of the perceived weaknesses of TSE relative to optimization: the ability to guarantee that the best solution has been found. This is compounded when in a negotiation setting and “create and explore many options” is a specific goal. Frustration with an incomplete tradespace is understandable during a negotiation in which no alternatives that are near-optimal for all stakeholders can be found. If it is feasible to postpone negotiations in order to allow additional alternatives to be evaluated, that may be a desirable strategy. Changes in the tradespace should be monitored carefully between sessions if the evaluated alternatives are expanded, both to prevent confusion and to make sure that negotiation is not unduly delayed. For example, if a 10% sample of the design space is increased to 25% to hopefully include better alternatives yet no *new* alternatives are identified as equal or better choices to the previous best alternatives, the tradeoff of spending additional time completing the tradespace is not likely to generate a considerably better design.

## 6.4 Uncertainty

Uncertainty is a primary concern in the design and operation of complex engineering systems. These systems tend to have long lifetimes and considerable exposure to uncontrollable, *contextual uncertainty*. This drives a need to consider multiple potential operating environments for the resulting system, each of which can significantly shuffle the relative performance of different alternatives. There are other relevant sources of uncertainty present in the MSTSE formulation that must be accounted for as well. *Model uncertainty* can prevent the accurate evaluation of individual alternatives, particularly when designing systems with no existing comparables against which the performance model can be validated. *Preference uncertainty* is also a concern when stakeholders are not confident in their value statements and may choose to iteratively refine their estimates. Finally, this section will also cover the role that *the -ilities* can play in a negotiation setting, as they relate strongly to the ability of a given solution to operate effectively in the face of uncertainty.

### 6.4.1 Contextual Uncertainty

When engineered systems are subjected to different operating scenarios over the course of their lifetime, they are exposed to contextual uncertainty: the inability to control the environment in which the system operates, or to predict with certainty what that environment will be. When contextual uncertainty is included in a design process it is often modeled through its impact on model parameters outside of the control of the designers, such as “targets of interest” for an observational satellite or “emissions tax” for a power plant. Context can also impact the value models of the stakeholders, for example by prioritizing different attributes depending on the political climate, and the value of the BATNA may also change. Accounting for contextual uncertainty is a critical design challenge of modern systems engineering. For this research, we will consider the use of Epoch-Era Analysis (Ross and Rhodes, 2008) to model contextual uncertainty due to its strong connections to the TSE literature, but it is likely that alternative methods are feasible as well. Regardless of its implementation, finding mutually beneficial designs across contextual uncertainty is a considerable challenge. Just as many designs may be good for a single stakeholder but few are good for all stakeholders, many designs may be good in a single context but few are good in all contexts. When considered together, multi-stakeholder and multi-context decision making is likely to have even fewer desirable solutions.

When setting up MSTSE using EEA, each possible context is instantiated separately, with associated model results for each alternative. For a small uncertainty space it may be feasible to look at *each* corresponding benefit-cost tradespace and associated visualizations, attempting to negotiate using all of the available information to find alternatives that perform well for all stakeholders in each context. However, for large numbers of possible future contexts the use of multi-epoch analysis metrics will be necessary to summarize performance and identify designs of interest (Fitzgerald and Ross, 2012). Part of the challenge of negotiating across

contextual uncertainty is the impact of metapreferences [6.2.5] on performance across the uncertainty space. For example, one stakeholder may prefer alternatives with the best average cost-benefit performance while a more risk-averse stakeholder prefers alternatives with the best worst-case benefit. Different multi-epoch metrics will be desirable for screening and locating potentially interesting solutions depending on these preferences. As mentioned in the section on metapreferences, explicit recording of the acknowledged decision making criteria outside the constructed value models can serve to minimize confusion between stakeholders about their interests.

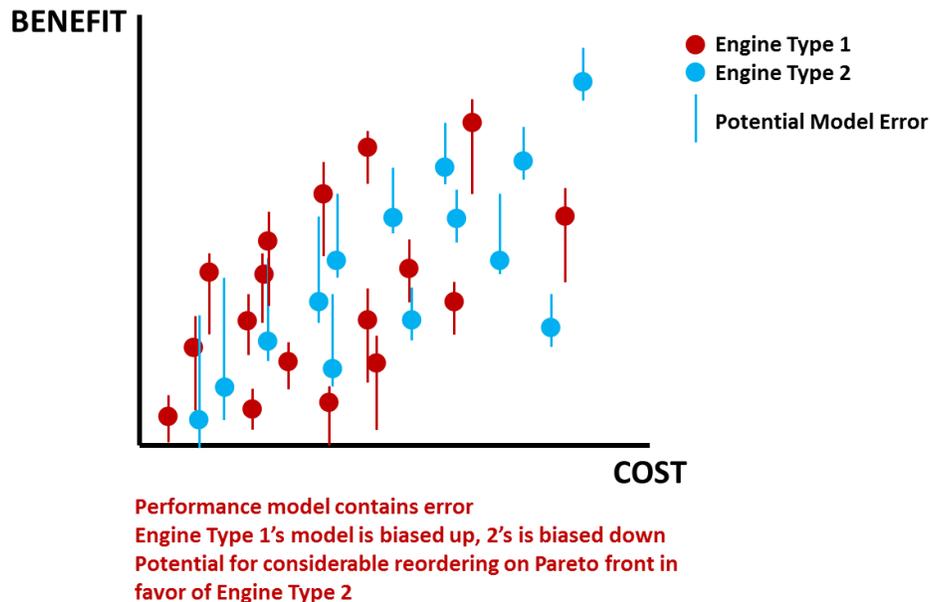
The central lesson of contextual uncertainty is that a good solution in one context may not be good in another. Therefore, single-epoch, traditional TSE visualizations must be used with caution. Unduly focusing on a single, baseline epoch may lead to attachment on the “best” solution in that context only to find that it is not acceptable in another. If the participating stakeholders identify their favorite alternative in one epoch, it should *not* be used as a reference point against which alternatives in other epochs are judged. This can be particularly troublesome if the chosen alternative remains good for some stakeholders but not others, as the lucky group has considerably less motivation to explore for other, better designs in the new epoch. For this reason, continual emphasis of the BATNA is appropriate in order to prevent lock-in on intermediate solutions. If possible, single-epoch views keyed to different epochs should be displayed simultaneously in small groups to emphasize the variability in performance across contexts.

## 6.4.2 Model Uncertainty

Model uncertainty, or evaluative uncertainty, is a problem endemic to early conceptual design. When alternatives are evaluated using models, typically constructed from some combination of deterministic equations and partially randomized simulations, those models are likely to have a margin of error. This error can be difficult to detect and correct, particularly for systems with no existing comparables against which the results can be validated. As a result, the performance of each alternative in the tradespace, as quantified by the calculated value-generating system attributes, is considered uncertain.

Because of the degree of uncertainty associated with early modeling, methods like TSE often deemphasize the *specific* numbers output by the model (which are better analyzed later on by higher fidelity models and/or prototypes) in favor of focusing on the *relative* outputs of different alternatives. This is well matched to the “focus on interests, not positions” objective of principled negotiation, as it couches the discussion in the language of better/worse rather than analyzing individual alternatives. Additionally, it lessens the negative impact of sources of systematic error in the models that affect each alternative in the same way. Of greater concern is the impact of model uncertainties that differentially affect alternatives based on their associated design variable levels, such as in Figure 6-9. For example, a categorical variable such as “propulsion type” might have considerably more error (or error in a different direction) for an

undeveloped nuclear option than for a standard bipropellant engine, which could falsely order the performance of those alternatives. This possibility is illustrated in Figure 6-9, where two different engine types have bias in opposite directions that possibly favors Type 1 (red) inappropriately. This can be compounded by the presence of nonlinearities in value models that strongly differentiate alternatives with similar performance.



**Figure 6-9: Model errors with different impacts depending on a design variable may misrepresent ordering**

Cooperative tasks in problem formulation, such as Joint Fact Finding and other joint modeling activities, increase in importance in the presence of model uncertainty. When models are known to be less than 100% accurate, the door is opened for stakeholders to dispute the results of the model in an attempt to further their own interests. It is critical to the success of the negotiation that the objective metrics used to assess the alternatives are accepted by all stakeholders. Joint modeling both exploits the individual knowledge of each stakeholder and encourages buy-in for the results. Correspondingly, more time should be spent developing and validating the performance models when uncertainty is believed to be high, proceeding to negotiation using the results only when all stakeholders are comfortable with the level of uncertainty. While negotiating, activities and visualizations that target specific design attributes such as a table view or logical filtering should be withheld in favor of complete value functions. Stakeholders should also be encouraged to engage with concepts such as fuzzy Pareto fronts that specifically call out the inability to definitively rank similar designs.

### 6.4.3 Preference Uncertainty

The elicitation of preferences is a difficult process. Most *basic values* models require carefully constructed parameters to match the nuances of a stakeholder's needs, often captured through structured interviews (Keeney and Raiffa, 1993), on the assumption that the stakeholder

is unable to directly express their *articulated values* (Fischhoff, 1991). This translation from complete, articulated values into a highly structured model of basic values is rarely a perfect match. Previous TSE research has noted the ability of the tradespace to demonstrate the consequences of a stated value model, resulting in gradual revision of stated preferences to be more in line with true preferences (16.89 Space Systems Engineering, 2002). From this we can conclude that at any given time a set of elicited preferences is only an *estimator* of true preferences, and thus is subject to change. This issue is compounded for multi-stakeholder problems, where it has been demonstrated that stakeholders will deviate from or correct their stated preferences when confronted with other preferences. This has been shown both in a negotiation setting (Curhan, 2004) and in an engineering setting (Golkar, 2014) and is theorized to be the result of subconscious or conscious efforts to reduce cognitive dissonance arising from differing opinions. Ultimately, the stated preferences upon which negotiations are conducted must be viewed with some degree of uncertainty (distinct from changes in preference motivated by changes in the system's operational context [6.4.1]).

Preference uncertainty should ideally be addressed early in the MSTSE task, just as it is in single-stakeholder TSE. If preference uncertainty remains unresolved, alternatives that appear valuable using the assigned benefit and cost functions may not match a stakeholder's internal understanding of value, limiting the effectiveness of most TSE visualizations and analysis techniques. This challenge is made manageable by transparency in the way the value models are constructed, one of the key advantages to using a multi-attribute utility value model (Ross, Rhodes, and Fitzgerald, 2015). The most direct way to approach the resolution of differences between the true and stated value models is to pose additional elicitation interview questions. Elicitation schemes for MAU usually anticipate receiving rationally contradictory responses due to the difficulty of verbalizing preferences. These contradictions are resolved through a combination of further targeted questioning in that area and regression to find the closest valid utility curve. This is not likely to be a useful tactic *during* a negotiation session, however, as the utility function should be elicited *before* the session and the initial interview should nominally catch any issues that could be resolved by a time-consuming revisionary interview.

Instead, the updating of preferences can take the same approach used in previous TSE design tasks: allowing stakeholders to directly manipulate the value model parameters when they feel the results of a visualization or analysis are not representing their interests accurately. Because MSTSE is prescriptively oriented, the loss of the normative benefits of MAU derived from the structured elicitation process is not a dealbreaker. MAU, or any other value model, is used within TSE strictly as an aid to decision making, not as a designation of the "best" or "correct" answers (which would be dubious anyway given the acknowledgement of unavoidable preference uncertainty). Allowing stakeholders to freely edit their own utility functions can allow them to quickly and individually "close the loop" on inaccuracies in the original elicitation. This is particularly useful in the case when a stakeholder revises their preferences based on another stakeholder's preferences. For example, if two stakeholders each value "image

resolution” as a benefit attribute but have different minimum acceptable values, they may discuss what the intended use case of the system is (e.g. how big are the objects they are imaging) and one may realize he set too strict of a minimum requirement. Rather than conduct an entirely new elicitation interview, the stakeholder can simply adjust his value model, recreate any tradespace visualizations with the new function, and continue to explore and negotiate.

The freedom to manually edit value models opens up additional possibilities for *gaming*, by allowing stakeholders to attempt to strategically misrepresent their preferences without the elicitation process as a filter. During good-faith negotiations this is not an issue, but in some applications it may be necessary to monitor stakeholder behavior with regards to excessive changes in stated preferences and/or to prevent them from making changes without informing the group of their reasoning. Overall, however, gaming a constructed value model in this way is unlikely to be an effective negotiation strategy, because it is difficult to anticipate what changes will result in the best negotiation outcome for oneself, especially compared to a classic negotiation where articulated values are used directly (Raiffa 2002). If preferences do become heavily modified over the course of the MSTSE session, it may become reasonable to elicit a new value function from scratch in order to accurately capture the new perspective of the stakeholder (in the honest case) and to limit opportunities for gaming (in the dishonest case).

Because preference uncertainty should be addressed early, visualizations that support the identification of discrepancies between the internal and elicited value models should be presented to the stakeholders at the beginning of the session. The benefit-cost tradespace has previously been used to this effect (16.89 Space Systems Engineering, 2002), with the option to color the tradespace by particular design variables or performance attributes as an effective way to highlight patterns in value that run counter to the stakeholder’s expectations. Similarly, direct visualizations of the value models – for example, showing single-attribute utility curves and tradespaces for an MAU function – can directly call out the consequences of the value model such as overly restrictive requirements and, for interactive software such as IVTea, incorporate the controls for direct modification of the model. Activities that rely on accurate stated preferences, such as Pareto front solving and stakeholder correlations, should be withheld until stability in preferences is reached.

#### **6.4.4 The –ilities**

The –ilities are system properties that impact value delivery without directly providing value and are a common concern in large, complex systems with long lifecycles and exposure to contextual uncertainty. The definition, quantification, and valuation of –ilities is a matter of considerable interest and current research in the literature. This research does not seek to identify “correct” –ilities or –ility measures for use in MSTSE. However, given the complex systems of interest to this research, it is likely that stakeholders participating in MSTSE will be interested in designing the system to include some subset of the –ilities. This interest often manifests itself as a metapreference [6.2.5], such as a preference for flexible alternatives over

passively robust alternatives, and should be handled accordingly with the appropriate metrics being calculated and incorporated into the exploration as properties of each alternative. Some stakeholders may want to use an –ility metric as a component of their utility function (or other type of fixed-context value model), but this is typically discouraged on normative grounds: the benefits of –ilities are experienced over time and thus a “snapshot” of value such as a utility function should not include them (McManus et al., 2007).

Of more direct interest to this research is the potential for –ilities to function as additional leverage with which to find mutually beneficial solutions. This is a creative activity that will need to take a different form for any given –ility and case, but can in some circumstances yield additional alternatives that provide considerably more value for all stakeholders even without the presence of uncertainty. As an example, consider an imaging satellite with two stakeholders; Stakeholder A is interested in high-resolution images of particular areas of interest, while Stakeholder B is doing a low-resolution mapping of the Earth’s surface. Correspondingly, Stakeholder A favors low-altitude satellites and Stakeholder B favors high-altitude satellites, which may ultimately make no alternatives mutually beneficial. However, system changeability can bridge the gap between the two stakeholders. If the stakeholders acknowledge that their respective missions do not require 100% uptime, the system can be changed over time between alternatives that favor one mission over the other (for example, by including jets on the satellite that allow it to burn fuel and change orbit altitude). Change mechanisms connect alternatives in a tradespace into a *tradespace network* (Ross, 2006) of alternatives that can be switched for some cost. This ability can be used to satisfy diverse missions where otherwise impossible, as in this example, or to save costs compared to “gold plated” solutions that satisfy multiple missions but at a prohibitive cost. In this way, changeability can be used to establish agreements that allow for time-sharing of the system (a form of divisible benefit [6.2.6] not captured by a static value model) while letting the system be optimized for each stakeholder when they are in control. An example of this is shown in Figure 6-10, and is applied within the Satellite Radar case study in chapter 9.10.

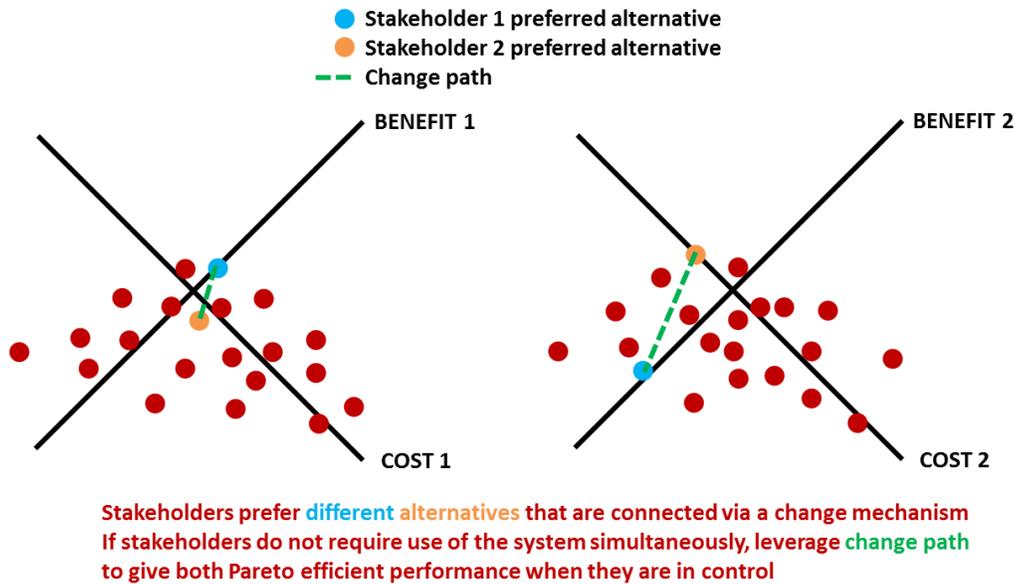


Figure 6-10: Change mechanism connects alternatives that are on each stakeholder's Pareto front

The exploration of –ilities as a part of TSE is an ongoing research area, so it is difficult to say what specific visualizations or activities could be conducted during an MSTSE session to best reveal valuable –ility possibilities, particularly for multiple stakeholders. For now, the bulk of the burden associated with designing for –ilities occurs in the problem formulation stage where the proper “hooks” are put into the model to evaluate the desired –ility metrics: e.g. the inclusion of design variables that enable change mechanisms to be used during operation. That means the stakeholders should have an idea of what –ilities they are interested in exploring during joint modeling activities. The knowledge of which –ilities are potentially *worth* exploring will normally be the domain of a stakeholder or subject matter expert, but a small list of universally common –ilities may be a useful tool to prompt discussion in the modeling phase on which –ilities to explore.

## 6.5 Logistics

Perhaps the most pragmatic consideration to be made when preparing to conduct MSTSE is the role of logistics in the process. Engineering design tasks are almost exclusively non-trivial, time-consuming activities, and the difficulties related to assembling the necessary talent and knowledge to solve those problems are compounded when multiple stakeholders become involved. Therefore, it is important to consider both the *schedule/timeline* for the MSTSE activities and the *access to decision makers* for the negotiation itself. Accommodations need to be made when decision makers are unavailable to participate actively and instead must send proxies. Furthermore, in some sensitive applications, the principle of Full, Open, and Truthful Exchange may not be feasible. In these cases, the presence of *private information* presents an interesting challenge to the basics of tradespace exploration.

### 6.5.1 Decision Maker Access

High level decision makers are often important people with many demands on their time and attention. As such, it may not be possible for them to participate in all phases of an MSTSE task, particularly those that involve meeting other decision makers with equally crowded schedules. This was identified as a key challenge for MSTSE during our practitioner interviews. Fortunately, critical problem formulation tasks that rely *specifically* on decision maker opinions (namely, the creation of benefit/cost functions) are conducted individually, allowing them to be performed at convenience. On the other hand, many MSTSE tasks that are ideally done with decision maker participation can replace decision makers with qualified proxies suffering only minimal loss of detail. For example, an experienced engineer or analyst will usually be capable of performing joint modeling exercises in place of a decision maker, achieving the same “trusted objective measures” outcome as long as the decision maker trusts their proxy.

Proxies are also potentially capable of participating in the negotiation session itself when decision makers cannot attend. It is important to keep three things in mind when considering the use of proxies in the negotiation. First, *proxies will not actually make decisions*. The suggestion to use a proxy does not eliminate the need for the “final call” to come from the decision maker. Proxies operate similarly to representatives [6.1.3], but ultimately are not expected to actually make decisions and thus have more freedom while exploring. With this in mind, *the job of the proxy is to serve as a stand-in that can represent decision maker interests in the negotiation while collecting and condensing information to pass back to the decision maker for the final decision*. The proxy is a cipher for the decision maker, strictly participating as if he were a decision maker in order to replicate the exploration of the tradespace that would occur if the decision maker he represents were present. Finally, *using proxies will always be inferior to direct decision maker participation*. Whenever possible, decision makers should be present during the negotiation. Beyond simply improving efficiency by removing a delayed feedback loop, the direct participation of decision makers is valuable because the decision makers are nominally better able to weigh alternatives using their knowledge of the “big picture” than their proxy is.

An example situation where proxies would be useful can help clarify their use. Imagine that the Air Force and Navy are working together on a UAV. This problem has (1) multiple very busy stakeholders such as 5-star generals, and (2) value propositions for each stakeholder that are complex enough to demand localized expertise. An Air Force analyst may be able to roughly replicate the interests of his commanding decision maker (especially with the assistance of a recorded utility function) but is not versed enough in Navy needs to effectively trade off the performance metrics of that value proposition or reevaluate requirements. Classically, the Air Force and Navy might each deploy analysts separately to characterize the *individual* design problem, and then pass that information back up to their respective decision makers, and then the decision makers would have to use their limited time to uncover the nuances of the *group* design problem themselves. If the two branches wanted their analysts to collect and condense

information about the group problem effectively, proxies for both the Air Force and Navy needs would have to work together. At this point, conducting an MSTSE negotiation session would be an effective way for the two proxies to explore the tradespace together, rather than individually. This allows for information about the multi-stakeholder problem to be pushed up to the senior decision makers, rather than requiring the decision makers to resolve that entire aspect of the problem themselves.

The use of proxies to negotiate is difficult in the presence of significant preference uncertainty [6.4.3]. If the decision makers themselves are uncertain what they want, it is unrealistic to expect that proxies will be capable of effectively conducting tradeoffs and identifying attractive alternatives in a way that approximates the decision makers. Additionally, the act of exploring itself is the main approach for resolving preference uncertainty and if the decision makers do not take part in the exploration, they may not be able to successfully refine their value functions to reduce uncertainty. Even if they did, the fact that their interpretation of the results and adjustment of preferences would occur offline instead of during the session would imply the need for a completely new session, causing costly repetition of the proxy negotiation. Therefore, in cases where proxies are a necessity, particular attention must be paid to resolve preference uncertainty *before* negotiation. This could take the form of a more detailed value elicitation, or a mockup single-stakeholder tradespace for each stakeholder to quickly review their value model and catch obvious discrepancies between the Pareto front and their intuition or mental model about the problem.

## 6.5.2 Schedule / Timeline

There is no canonically “correct” amount of time to spend on tradespace exploration, but the reality of using TSE or MSTSE in practice comes with demands for structuring and scheduling the design process. In line with the previous section, this is especially important for MSTSE due to the need to schedule a negotiation session between stakeholders. Ideally, enough time could be spent on each step (problem formulation, evaluation, exploration, etc.) to fully extract all of the relevant insights and make the best possible decision. For a given case, this may entail scheduling multiple sessions for Joint Fact Finding, Collaborative Modeling, and negotiation, depending on the difficulty inherent in reaching agreement (Ehrmann and Stinson, 1999; Mostashari, 2005). Focusing here on the negotiation, separating it into multiple sessions may require a few additional considerations. Though tentative or intermediate agreements can be useful strategies in a negotiation, the end of a session should not be used as pressure to accept one. Tentative agreements are likely to become reference points, and thus must be a solution each stakeholder is *actually* willing to accept in order to maintain situational awareness of gains and losses. This mandates the importance of record keeping between sessions, as the outcome of each session cannot be summarized in a single, current “best” alternative. The use of TSE software should be leveraged to save information from session to session. For example, IVTea will maintain lists of identified “favorites” from session to session, which can also be used to prevent rework by saving important batches of designs from calculations (logical filters, Pareto

sets, etc.). At the least, debrief notes should be taken at the end of each session to note conclusions that have been reached during the session and their rationale, in order to capture insights that may influence the final decision.

On the other hand, what happens when a restrictive timeline limits the ability to guarantee full completion of the analysis? For example, perhaps the decision makers can only attend a single, two hour session. Deadlines can function as “action-forcing events”, which may be imposed externally (e.g. a limited window to pursue an opportunity) or internally (e.g. a mutually agreed-upon deadline to force a decision) depending on the specific case (Watkins, 1998). This is tied in with the decision maker access problem of the previous section, so proxies could be deployed in other sessions if more sessions are required. Advantages of having the decision maker session *before* the proxy sessions include the resolution of preference uncertainty, identification of addressable issues with the models, and the possibility of finding either an obviously dominant solution or quickly determining that the project should be abandoned to save further effort. The advantage of having the decision maker session as the *final* session is that the results from previous analyst sessions can be leveraged to increase the value of the decision makers’ time and the face-to-face negotiation of the final decision can reduce the time delay of the communication.

Sometimes proxies will not be a viable method because the timeline mandates that a decision is made by a particular date. Generally, negotiation is not well suited to operating under a deadline; most discussion of deadlines concerns either the strategic setting of deadlines to pressure another party or efforts to remove or extend deadlines (Brett, 1991; Watkins, 1998). Agreement cannot be forced and takes time to develop. Thus, there is no catch-all way to run a time-limited negotiation. The best advice for avoiding slow negotiations is to avoid package bargaining and stick to principled negotiation with a focus on interests. Structuring the use of time in the negotiation with clear intermediate goals and decisions also supports this goal by increasing the perceived momentum of the negotiation (referred to as “progressive summarization method” in Spoelstra and Pienaar, 1999). If mutually beneficial solutions are difficult to find, it is likely that the BATNAs will end up being the final decision under time pressure. Activities should focus on high-level visualizations and analysis of value such as the benefit-cost tradespace and Pareto front solvers, as these tools communicate key decision criteria in the most compact and familiar way.

### **6.5.3 Private Information**

This research has supported the use of principled negotiation and its corollary, the principle of Full, Open, and Truthful Exchange (FOTE), in the practice of MSTSE. These are considered to be the proper methods for negotiating with a desire for mutual benefit. However, in some applications, the use of the FOTE principle is restricted. Sensitive or proprietary information may be unable to be shared between stakeholders, which may cause difficulties in

communication. There are three key potentially private components of TSE that can be addressed in different ways: performance models, value functions, and value outputs.

Performance models often contain proprietary information, developed internally by a stakeholder organization and not available to others. Fortunately, the models can be treated as “black boxes” with respect to the remainder of MSTSE. As long as the proper inputs are passed along, proprietary models can remain in the sole possession of their owners. The outputs from the proprietary model can be shared in the TSE database if possible, or kept private if they are also proprietary. The drawback to this method is that it does not allow for joint modeling or joint fact finding, opening up the possibility that a different stakeholder distrusts the outputs of the proprietary model. This may be unimportant if the attributes calculated by the model are solely valued by the owner stakeholder, unless the relationship is soured to the point of possible accusations of tampering or gaming of the results. If a public model that calculates the same attributes (presumably at a lower fidelity) is available, it can either be substituted in for the proprietary model or used as validation to build trust in the outputs for the other stakeholders.

Value functions are occasionally sensitive when organizations have security-related reasons not to divulge their interests or objectives. Again, the value function can be “black boxed” if desired, though with slightly more difficulty. First, all of the relevant value-generating attributes still need to be available, which means they must either be publicly shared or the private output of a proprietary model. The public option may raise difficulties while in the joint modeling phase, as the citation of a need to calculate a particular attribute is functionally a dead giveaway of its use in a value model. This may be acceptable if some value dimensions are not sensitive and others (calculated by proprietary models) are. If not, it is possible to have the sensitive stakeholder provide a list of attributes they intend to calculate with some subset of those attributes to be used in their value function, as a cover for which attributes are actually important. Beyond the calculation of value, though, are issues raised by negotiating with private value functions. A hidden value function has to be carefully managed, either by software or hardware, so that only the appropriate stakeholder can see it. This may drive communication difficulties by hiding the insights derived from visualizations of the value model from the group. Additionally, it also disables the ability to use visualizations that specifically compare the value models of multiple stakeholders, such as the attribute correlation heatmap. Instead, analysis will need to rely on the value function’s outputs only. Essentially, the effect of a private value function is the removal of traceability of *why* a given alternative is good or bad for the given stakeholder (though they retain that information themselves). Given the one-directionality of the information in this case, it may be preferable for all stakeholders to have private value functions, in order for the negotiation to be balanced.

Finally, the challenge of private value model outputs presents the greatest barrier to productive negotiation. In this case, knowledge of which alternatives a stakeholder finds good or bad are hidden from the other stakeholders, severely limiting their ability to leverage TSE to

improve the negotiation. If the sharing stakeholders trust the private stakeholder, as the only holder of all the relevant information, to drive the session fairly, they can do so. Otherwise, a solution must be negotiated using the inferior “dance of packages” method of alternating offers, which can lead to ineffective midpoint solutions or no solution at all. In this case, the benefits of MSTSE are experienced only by the subset of stakeholders that are able to share, but are severely limited if the withholding stakeholder has veto power. This scenario is another potential use for a mediator role within MSTSE. If the private values can be revealed to an independent third party mediator, the mediator can utilize full MSTSE visualizations to guide the participating stakeholders into discussing mutually beneficial alternatives, where otherwise the stakeholder would have to guess the needs of the other party.

## ***6.6 Discussion***

This chapter has discussed a variety of structural components of a given multi-stakeholder problem that would logically have impact on how MSTSE should be approached. When appropriate, we provided some simple suggestions based on our working theory of framing in MSTSE for mitigating any negative effects; these could, of course, be tested via experiment in future research. Table 6-1 summarizes the chapter as a whole, serving as a quick reference for what might be of interest to adopters of MSTSE when faced with a problem that does not match the basic assumptions used before this point. In the following chapter, these suggestions are synthesized with the other insights about MSTSE developed to this point in the thesis into a short set of recommendations for setting up and executing MSTSE.

**Table 6-1: Summary of problem structure with impact on MSTSE**

<b>Stakeholders</b>	Number	For more than two stakeholders, limits ability to show all information in one tradespace view. Emphasize functional relation of design variables to value, or break down into pairwise analysis.
	Outside Relationships	Focus on critical stakeholders to simplify. Scope out competitive stakeholders if possible to support FOTE.
	Representation	Limits ability to change problem formulation “on the fly”. De-emphasize activities that would modify value functions or design space.
<b>Preferences</b>	Alignment	With high alignment, focus on coalition formation to simplify negotiation. Low alignment makes feasible alternatives look worse: focus on comparing against the BATNA.
	Single Stakeholder Value Models	If stakeholder is unable to create value model, limits effective size of tradespace: focus on BATNA to reduce cognitive burden. If value model is one-dimensional, leverage optimization.
	Multiple Stakeholder Aggregation	Aggregate for tight coalitions or connected stakeholders who can deal amongst each other on the side, reduces complexity in the negotiation.
	Unoperationalizable Attributes	Introduces subjectivity to analysis that can be disputed in negotiation. Use a neutral outside source to assess if possible.
	Metapreferences	Limits the ability of value models to effectively describe the actual decision criteria. De-emphasize value models in favor of uncertainty analysis.
	Divisible Attributes	Adds an additional dimension of leverage to the negotiation. Use to create new alternatives that are more “fair” when possible.
<b>Alternatives</b>	BATNA Type	Impacts context of the negotiation decision for each stakeholder. Depending on type, leverage the tradespace in different ways to keep discussion positive.
	BATNA Quality	Impacts relative attractive of alternatives and likelihood of strictly superior choices for all stakeholders. Avoid allowing strong BATNAs to prematurely reduce the number of designs under consideration.
	Tradespace Size	Impacts ability to give attention to individual alternatives. With few designs, look to expand design space. With many, avoid comparing lists of favorites.
	Tradespace Completeness	Impacts confidence in final solution. Attempt to be as complete as possible but cease adding alternatives if they are not making new “best” choices.
<b>Uncertainty</b>	Contextual Uncertainty	Demands assessment of alternatives across range of operational conditions. Use BATNA to prevent lock-in on intermediate solutions.
	Model Uncertainty	Impacts confidence in final solution. Emphasize JFF to minimize disagreements; focus on <i>relative</i> over <i>absolute</i> performance.
	Preference Uncertainty	Limits ability of value models to predict true preferences. Resolve as early as possible, emphasize analysis of value models to correct any errors.
	The -ilities	Leverage to make creative solutions not described by static performance.
<b>Logistics</b>	Decision Maker Access	Limited access may force iteration. Difficult to revise value models: resolve preference uncertainty before using proxy negotiators.
	Schedule / Timeline	May rush decision. Focus on interests, maintain notes between sessions.
	Private Information	Creates asymmetry in power and threatens FOTE. “Black box” when possible or use a trusted/impartial mediator.



## 7 Recommendations for Performing MSTSE

This research has delved into the implications of framing on multi-stakeholder engineering design problems. We began by analyzing and improving the tools and visualizations used to communicate a multi-stakeholder problem to stakeholders. These conclusions were tested with controlled experiments and brought to practitioners to more closely couple them with current needs. Then we discussed the underlying structure of the problem itself and how it can impact the creation and exploration of the tradespace. This chapter collects the insights of those connected but distinct research paths into a concise set of recommendations for setting up and performing MSTSE, controlling framing to support rather than hinder negotiations. Note that MSTSE is not intended to be a step-by-step systems engineering process, but rather an analysis technique and negotiation aid that builds on the foundation of prior TSE literature and negotiation theory. Instead, the recommendations will be grouped according to a generic outline of a model-centric systems engineering process capable of supporting TSE. This should allow the recommendations to be incorporated into a variety of larger SE processes, accommodating personal familiarity or organizational standards and lowering the barriers to entry for applying MSTSE.

Thus, without restricting attention to a particular implementation, generally a TSE project will follow a procedure similar to this:

1. **Problem Formulation** – the structuring of the problem and scope of decision making. This normally includes the definition of the design space used to enumerate potential system alternatives, the context in which those systems will operate, and the stakeholders and value attributes used to assess them.
2. **Modeling/Evaluation** – the development and use of models for the purposes of evaluating the designs. Models can take many forms, which necessitates a selection of modeling technique(s) appropriate to the problem formulation. Creating models is itself nontrivially difficult and normally takes considerable effort without the benefit of reuse of previous models.
3. **Exploration/Analysis** – the attempt to curate insights from the model outputs. Stakeholders and analysts are both capable of performing this step, with different strengths and weaknesses. Exploration is typically intended to generate results capable of justifying a decision to select a given design alternative.

These three activities represent conceptually different tasks. For a given project, different personnel are sometimes responsible for the different activities. However, even if the same group of stakeholders and analysts work together through all three tasks, the impacts of framing manifest in different ways from task to task and can propagate from one to the next. As such, MSTSE practitioners should be cognizant of how their actions affect framing and perform the necessary diligence to ensure that framing does not misrepresent the underlying negotiation problem.

Using the generic three-step outline of a TSE procedure, the following subsections detail recommendations for controlling the framing of common TSE activities in order to support a successful MSTSE application. The recommendations are based on the conclusions of the prior sections, synthesizing the TSE, negotiation, and framing literatures with the work of this research. The success criterion of MSTSE lies in the ability to find and identify mutually beneficial alternatives, if they exist. To do that, the macro framing of the problem should be aligned with the tenets of principled negotiation as much as possible and the micro framing must accurately represent the value of the different alternatives. The recommendations included here are not intended to be exhaustive but rather instructive advice for potential adopters of MSTSE. Following these recommendations should improve the communication of preferences and needs between negotiators (a skill not developed or supported by classic TSE) and the value assessment of the alternatives by each negotiator (which is a different, more complex task than in classic TSE). This improves MSTSE, compared to traditional TSE with multiple stakeholder value statements, by reducing the likelihood of key failure modes at both the inter-stakeholder and stakeholder-data interfaces, limiting opportunities for negotiation breakdown driven by social conflict or misattribution of value.

## **7.1 Problem Formulation**

Problem formulation has a large impact on the resulting direction of a tradespace analysis. It defines the scope of the system to be analyzed, what factors are (and are not) under designer control, and the sources of value that are sought by the stakeholders. Unsurprisingly, the predominant impact of framing in this stage is likely to come from macro framing as the beliefs, perspectives, assumptions, and sometimes biases of the participants work their way into the problem. To address this challenge, communication becomes paramount: explicitly capturing some of the macro frames with which stakeholders and/or analysts are approaching the problem can allow for the identification and mitigation of potential future barriers to agreement before they become negotiation impasses.

**Capture macro frames.** Note that the objective of these efforts is not to *change* the macro frames with which stakeholders approach the problem, but to capture what they are. Practically, macro frames are developed by a lifetime of experience and opinion, and are difficult to change. More fundamentally, since MSTSE is positioned as a *prescriptive* rather than *normative* analysis technique, it is inappropriate to suggest that one macro frame is the “correct” frame to use (a normative argument). Rather, we are interested in knowing the macro frames favored by each stakeholder so that when *they* attempt to make a normative argument we can understand the frame leading them to make that argument and, hopefully, communicate it effectively to other stakeholders who do not share that frame. This is intended to prevent incidents of the stakeholders “talking past” each other by assuming others share their underlying assumptions.

For example, consider the *purpose* behind conducting MSTSE. This could be simply to explore and learn about the opportunity, or it could be an official means of making and justifying a funding decision with which to move forward. A part of the purpose for MSTSE could be in pursuit of other, extraneous benefits such as establishing a working relationship amongst the stakeholders, regardless of the final decision. An existing system could be upgraded or it could be completely replaced. Regardless of what it is, the purpose for designing a system and interacting with other stakeholders has a role in framing the ensuing discussion – yet the purpose of a given project is inherently subjective. It is possible that each stakeholder may approach MSTSE with a different understanding of why they are there and what is to be accomplished. Clarifying these high-level objectives early on during problem formulation supports Full, Open, and Truthful Exchange and increases the likelihood that any available opportunities that support all stakeholders’ goals are identified since now those goals are known to each participant.

There is no exhaustive list of macro frames to consider in this phase. The challenge of capturing macro frames is similar to that of capturing preferences for a value model: it is an open-ended question and is often so engrained that a person may not realize the assumption or be able to articulate it when prompted. In fact, strong metapreferences (discussed in 6.2.5) may impact the macro frame of decision making for a stakeholder, such as a relative desire for flexible alternatives over passive alternatives (even if they have the same total value according to their value model). Eliciting information of this kind is a subset of the “relationship building” and “information sharing” goals of multi-stakeholder dialogues (MSD; Susskind et al., 2003); therefore, implementing a formal MSD is a potential avenue for capturing macro frames when they are believed to be diverse between the participating stakeholders or unlikely to be revealed in the normal course of constructing a tradespace.

**Create many alternatives (if possible).** The central tenet of principled negotiation and the main justification for its natural correspondence with TSE, it is important to create a tradespace with a sufficient number of alternatives to ensure that any sources of integrative value are evaluated and thus visually available during exploration. Nominally, enumerating and evaluating as many (reasonable) designs as possible is a recommendation of standard TSE, but it is one that increases in importance in the transition to MSTSE. While a sparsely enumerated tradespace may generate enough information to sufficiently understand the “best” tradeoffs for a single stakeholder, the introduction of multiple stakeholders combinatorically increases the relevant number of tradeoffs that impact the design process: tradeoffs not only between benefit and cost for a given stakeholder but between benefit and cost *across* stakeholders as well. If the “best” alternative is not enumerated in single-stakeholder TSE it is often possible to examine the trades between design variables taking place on the Pareto front and use them to estimate the intermediate values that would put the design in the desired benefit-cost area. With more than one or two stakeholders, it rapidly becomes infeasible to make this estimation, increasing the importance that many alternatives are evaluated to reduce the “gaps” between them.

Similar to the point on macro framing above, there are no unilateral benchmarks for what constitutes a sufficiently large tradespace, requiring case-by-case attention to this recommendation. The boundary between trade studies and TSE is subjective – trade studies usually being on the order of 10 or fewer alternatives, while TSE is usually a minimum of hundreds – and it is similarly unclear exactly how many alternatives are necessary for MSTSE, partially because these benchmarks are distilled from community consensus and there is not enough of a body of MSTSE literature and case examples to support such a claim. At the very least, we can be confident that MSTSE should be conducted with at least as many alternatives as one would use for TSE on a similar one-stakeholder system and preferably more than that. If the design space is restricted through a Design of Experiments sampling method for computational cost concerns, expanding the design space to approach a full-factorial enumeration should be conducted in the background of the negotiations in case no agreement can be made on the initial set of alternatives. Only highly constrained systems (with very few feasible alternatives) should have less than a full-factorial enumeration and fewer than hundreds of alternatives.

**Record key elements of problem structure.** This activity is already a main component of problem formulation for TSE, which requires explicit accounting of the factors impacting the system and their assignment as variables in the tradespace: design variables, context variables, or performance attributes. However, the multi-stakeholder problem has additional structural elements on top of those from single-stakeholder tradespaces that can impact the best micro frames to use in later phases of MSTSE. A provisional accounting of these elements and how they impact MSTSE was detailed in chapter 6; most of the elements for a given case are able to be determined in the problem formulation phase, before any modeling or analysis. Explicitly capturing the structure during problem formulation can improve later analysis, as certain analysis types can become more or less relevant depending on these key features. For example, if some attributes of interest to the stakeholders are divisible at-will (e.g. manufacturing costs, which can be split between stakeholders as desired), these can be leveraged through additional analysis later by customizing or sub-optimizing a given alternative. On the other hand, negative pre-existing relationships between the stakeholders may limit the effectiveness of some types of exploration, particularly those that involve directly comparing desired alternatives.

**Determine each stakeholder’s BATNA.** The BATNA is a “key element of problem structure” from the previous point, but is critical enough to merit its own description. Because the BATNA is an important reference point with respect to the value of any of the design alternatives under consideration, it is imperative that the BATNA be captured as accurately as possible during problem formulation so that it can be incorporated in later framing activities. Failure to define and then leverage the BATNA during exploration reduces the situational awareness of the stakeholders and can weaken MSTSE’s ability to support a negotiation.

In some cases, the BATNA will be readily apparent, particularly if the stakeholder(s) have *no* viable alternatives to a negotiated agreement and must simply accept the status quo. However, in general this task requires careful thought and consideration just like the rest of

problem formulation. It can help to consider the variety of BATNA “types” introduced in chapter 6.3.1 in order to prompt brainstorming in multiple areas. As a reminder here, common BATNAs include the following:

- **Do-nothing** – if the MSTSE is strictly exploratory, inaction is likely the course of action should no agreement to proceed be made. Doing nothing typically carries zero cost and zero benefit.
- **Existing system** – for design tasks intended to improve or replace an existing system, the do-nothing alternative actually entails using the current system. This type of BATNA is one that commonly drives differences in stakeholders’ bargaining leverage, as some stakeholders may be much better off with the current system than others.
- **Build preferred alternative alone** – some projects seek agreement between multiple stakeholders to reduce the cost borne by each individual. If a stakeholder is capable of affording some or all of the alternatives by themselves, those alternatives become viable BATNAs (though at a higher cost than if they could agree to share one).
- **Other opportunity** – resources that are expended on the alternatives in the tradespace represent an opportunity cost in that they cannot then be spent on other projects, which may be more valuable. This type of BATNA is the most difficult to capture, as the number of other opportunities is potentially limitless, but this fact is true for all design tasks. Usually a small number of known viable or attractive opportunities can be considered without fear of missing drastically better choices.

Identifying the best alternative in each of these categories and then assigning the best of those as the BATNA can reduce the complexity of comparing all possibilities “outside” the tradespace at once. Sometimes it may be difficult to assess which of these choices is the “best” (and thus, the BATNA) at this point, because the evaluative model has yet to be created, particularly for the “build alone” choice. In that case, preserving the list of potential BATNAs and then choosing one after modeling but before exploration is feasible.

## **7.2 Modeling / Evaluation**

Engaging in the modeling of the system after completing a thorough problem formulation seems at first glance to be trivial: simply a matter of taking the defined design vectors and finding the right equations to calculate the desired performance attributes, subject to any influencing contextual parameters. However, the modeling task itself can also propagate cooperative versus individualistic framing implicitly into the exploration phase. When multiple stakeholders will be conducting the exploration, it is important to make sure that the modeling is satisfactory to all of them, which requires some additional management.

**Joint Fact Finding (JFF) and Collaborative Modeling.** Joint Fact Finding (Ozawa, 1991; Ehrmann and Stinson, 1999) is a valuable use of time in order to build trust in the data that exploration will be based on. It is difficult to reach consensus if some stakeholders dispute the

facts on which the decision will be based, making uncoordinated multi-person modeling activities a threat to productive negotiation. JFF seeks to establish credible and objective data, one of the foundations of principled negotiation (Fisher, Ury, and Patton, 1991), to use as the foundation for evaluation of alternatives and discussion of their relative merits. The outputs of JFF can be used to build the models that evaluate the alternatives in the tradespace. If possible, all efforts should be made to convene stakeholders prior to actual exploration in order to perform JFF in support of the modeling task. JFF also helps to establish a macro frame of cooperation *before* engaging in the negotiation itself, which can help preserve positive, mutually-beneficial bargaining in the face of any naturally developing competitiveness.

Having stakeholders participate in modeling according to one of the techniques in the Collaborative Modeling literature (Langsdale et al., 2013) is also superior to allowing one stakeholder to control the modeling process. These techniques are designed to develop models that promote consensus-building activities such as MSTSE; many of the “best practices” (identified by Langsdale et al.) common across the techniques reflect goals that are similar to principled negotiation and FOTE, including rapidly accommodation of new alternatives and transparency in software implementation. Different techniques in this field are often tailored for particular types of models: for example, Group Model Building and Mediated Modeling specifically target the use of systems dynamics models and would be a strong choice if systems dynamics was considered an appropriate and useful modeling technique for a given problem. In addition to building rapport and trust in the underlying model, participation in model creation has been shown to improve negotiation outcomes (Czaika, 2015). Combining the objective facts of JFF with Collaborative Modeling establishes the strongest foundation for subsequent negotiation, minimizing the potential for breakdowns due to disagreements over the outputs of the modeling phase.

**Private Information.** Not all models can be developed through JFF and Collaborative Modeling. If a stakeholder already possesses a model for a piece of the larger system, reusing that model can save time and effort. If they are willing to share that model (both how it works and its results) with the rest of the stakeholders as a part of a larger JFF effort, that is a valuable step in building rapport, in accordance with the principle of Full, Open, and Truthful Exchange. Some stakeholders may be reluctant to share models, but should be encouraged to do so for the above reasons. However, some models’ inner workings may depend on proprietary or classified information that the stakeholder is unable to share. In the case of a stakeholder unwilling or unable to reveal their models, two approaches can be taken, as discussed in chapter 6.5.3: the existing model can either be ignored in favor of a newly-created JFF model (if possible) or “black-boxed” so that other stakeholders can only see its outputs. A black-boxed model can be fully effective if its outputs only impact the value proposition of the stakeholder who owns it. If not, other stakeholders will need to trust that the model is accurate. If a public - but presumably lower fidelity - model is available, it can be used to help validate the black-boxed model and build trust.

### 7.3 Exploration / Analysis

Entering the exploration phase, the dominant framing concern shifts to micro framing: the actions the participating stakeholders are asked to perform and the way the data generated by the previous steps is presented. Macro framing still has a role to play in exploration however, specifically when weighing specific alternatives as potential final agreements.

**Emphasize the BATNA.** For a proper valuation of the designs in the tradespace, they must be valuated against the BATNA as a reference point. This provides the necessary perspective for determining the value of a design *as a multi-stakeholder agreement* rather than the typical, less-contextualized evaluations *in a vacuum* or *relative to other designs* commonly used in classic TSE activities. Taking classic TSE visualizations and intelligently incorporating a prominent indicator of the BATNA is a functional way of improving negotiation behavior, as demonstrated by the negotiation tradespace experiment described in chapter 4. Views designed to compare alternatives should include the BATNA as a “sticky” alternative, even in simple implementations such as tables of performance data.

**Limit strictly-individual analysis.** Activities should incorporate the value statements of multiple stakeholders as much as possible in order to consistently keep each participant aware of the “group” aspect of the negotiation problem. This can prevent fixation on alternatives that are very good for one stakeholder but not for others. In the negotiation tradespace, color and transparency accounted for the value of other stakeholders, and the resulting negotiations saw fewer exhaustive search patterns in favor of more direct paths to mutually-valuable solutions. If the participating stakeholders want to utilize a particular analysis of the tradespace using their own value, it should be replicated for other stakeholders and shown together. For example, the benefit-cost efficient solutions on the Pareto front are highly desirable for a given stakeholder, but should be calculated and presented relative to the Pareto fronts of the other stakeholders. This can be accomplished in multiple ways, including the use of Venn diagrams to illustrate overlap between specific stakeholders’ preferred alternatives and gridmaps to show the relative sizes of the regions of agreement for all stakeholders, as discussed in chapter 5.5. It is possible that a wide variety of “traditional” visualizations can be modified to accommodate multiple stakeholders while maintaining familiarity, which future research can explore.

**Analyze relationships.** The relationships between stakeholders in the value domain is a component of the multi-stakeholder tradespace that is not present in classic TSE, but is just as important as the evaluation of the alternatives directly. These relationships, whether or not they are analyzed, will affect the ways stakeholders interact and the designs that they might agree on; thus explicitly considering them is a powerful means of understanding the dynamics at play in the negotiation. Stakeholder relationships in the value domain can be quantified through the correlation of their value metrics at the holistic level (e.g. the correlation between Stakeholder A and Stakeholder B using their respective cost-benefit efficiencies) and displayed in a heat map for all stakeholders at once. In addition to revealing potential coalitions between stakeholders

with similar interests, explicitly showing positive correlations indicative of shared interests can be a useful reminder of the potential for mutual gains for stakeholders caught up in a distributive negotiation fallacy or fixated on individually-optimal alternatives.

Correlations can also be displayed on an interest-by-interest basis (e.g. the impact on the correlation of A and B's utility functions caused by A's interest in expected lifetime as a value metric). The resulting correlation data is combinatorically larger than at the holistic level but can be segmented to provide an intuitive breakdown of how one stakeholder relates to all of the others. This can be used to identify key "free" attributes that do not need to be traded between stakeholders and "pain points" that drive the differences in the value statements for each stakeholder.

**Allow stakeholders to change their mind.** Negotiation in MSTSE exposes each stakeholder to large amounts of information that they may not have previously known, particularly the preferences of other stakeholders which are not present in classic TSE. New information can change subjective assessments of value (Curhan et al., 2004) and invalidate parts of the original problem formulation, in addition to subtly changing the macro frames through which the stakeholders view the negotiation. Stakeholders should be encouraged to critically reassess their value statements during the negotiation. Adjusting the value functions to more closely align with a "new" reality can be performed during a session in order to accelerate the iterative design loop. Additionally, if the value function updates are convergent in a manner leveraged by other consensus-building techniques such as the Delphi method (Golkar and Crawley, 2014), these live updates have the potential to open up new regions of mutual value in the tradespace.

**Refer back to macro frames.** When discussing individual alternatives, effort should be made to refer back to the macro frames of each stakeholder (ideally captured during problem formulation). When a stakeholder refers to a design with a subjective assessment like "good", the first question should always be "Why?". Each stakeholder wants a "good" design, but each has different criteria for what is "good" that includes not only their reported value function but also the macro frames with which they choose to make decisions. For example, if Stakeholder A recommends an alternative as "good" on the grounds that it has high benefit for all parties, Stakeholder B can make a more intelligent counteroffer with less chance of sparking a debate over the definition of "good" if it is clear to all parties that he prefers low-cost, high-efficiency solutions over strictly high-benefit solutions.

#### ***7.4 Framing Informal MSTSE without Negotiation***

The original conception of MSTSE (Ross et al., 2010a) specifically envisioned the support of simultaneous, interactive negotiation between stakeholders interested in and capable of performing TSE. Analysis of the negotiation literature in this research confirmed the potential for TSE to function as an effective medium in direct, interpersonal negotiations on the basis of

their shared commitment to objective data and exploration of interests across many potential alternatives. However, not all tradespace exploration is conducted by stakeholders, usually due to a lack of time necessary to devote to TSE or insufficient technical knowledge. As we confirmed in our interviews, analysts and engineers are often responsible for exploration in addition to whatever roles they play in problem formulation and modeling. When analysts conduct the exploration activities, they typically disseminate a small set of insights, recommendations, or designs of interest back to the stakeholders for either a final decision or an iterative improvement of the tradespace.

Stakeholder participation issues only become more likely when attempting to deploy MSTSE. Lack of knowledge, time, or even a positive working relationship may make one or more stakeholders unable or unwilling to participate in an MSTSE session. Alternatively, a particular stakeholder may want to explore the tradespace on their own – or through an analyst – to gain understanding of the problem before engaging in whatever formal decision making process will be used with the other stakeholders. Given the benefits of having stakeholders engage in exploration (notably, the intuition and insights created by performing the exploration directly thereby removing a costly feedback loop), MSTSE without live negotiation will necessarily be less effective in its ability to capture emergent decision rationales and possible creative or “outside the box” solutions. However, the prevalence of this use case makes it worthwhile to consider: how should tradespace exploration be framed and conducted when there are multiple stakeholder value propositions to consider, but no live negotiations between them? More specifically, what parts of the MSTSE recommendations in this chapter change with only one driving stakeholder? We will call this variety of MSTSE “informal”, to indicate the lack of formal negotiation and actual decision making possible with only one stakeholder.

Problem formulation for informal MSTSE will be very similar to full MSTSE. At the most basic level, problem formulation is about defining the system and the decision that needs to be made. If we want to solve the same problem, regardless of the participation of other stakeholders, we need as similar of a problem formulation as possible. Thus, the same alternatives would be created for the tradespace and the problem would retain the same fundamental structure. It will be necessary to estimate the value functions for the non-participating stakeholders, which is a downgrade in information quality – but many applications of TSE have generated useful insights even while relying on estimated value models as long as the basic interests are known<sup>15</sup>. However, capturing the macro frames of non-participating stakeholders is less likely to be successful. It is possible that some macro frames could be

---

<sup>15</sup> See (Spaulding, 2003) for a description of various methods similar to TSE and MATE that use informally elicited and/or estimated value functions as a simplification for incorporating stakeholders who cannot (or will not) conduct a rigorous utility function elicitation. (Ross, 2006) also provides examples of insensitivity in the relative “goodness” of alternatives in a tradespace to deviations in curvature and weighting in utility function that might be the result of estimation.

inferred from available information (e.g. a public statement about budget cuts could logically suggest a metapreference to reduce costs) but for the most part the informal analysis will rely on the data in the tradespace only.

Conducting JFF or Collaborative Modeling as a part of the modeling phase is no longer possible under the informal condition, as it explicitly requires the active participation of all parties. This means that the driving stakeholder (and their analysts) are responsible for all facts, assumptions, and models used to create the tradespace. While this simplifies the activity by removing a potential source of disagreement in the modeling phase and returning it to the state of traditional TSE, the resulting tradespace product may not accurately represent how *other* stakeholders will evaluate the alternatives, lowering its ability to predict their future negotiation tactics<sup>16</sup>. Without using JFF to establish a shared objective ground before exploration of the alternatives, the conclusions based on exploration must be viewed with additional caution. There is no action that can be performed without input from the other stakeholders that mitigates this concern. However, given that JFF is an accepted negotiation practice, it is possible that the other stakeholders would be willing to conduct JFF prior to whatever non-MSTSE negotiation method will be used. In this case, informal MSTSE practitioners could feasibly base the tradespace on the output of an external JFF, provided that there is enough time between the JFF and the “real” negotiation to build and analyze the tradespace to extract whatever insights were desired. Thus, if JFF is agreeable but a full negotiation MSTSE is not, a stakeholder interested in conducting a side-analysis with MSTSE should push to organize JFF as far in advance as possible. A similar practical consideration is that, without full stakeholder participation, informal MSTSE is benefitted by iteration. If informal MSTSE is being deployed as an analysis technique in preparation for or simultaneous to the “official” decision making process, it is likely that more information will gradually become available that will impact the models used to support MSTSE. Iteratively refining the models and exploring the data can be used to provide gradually improved insight into the multi-stakeholder dynamics of the decision problem.

The most significant difference between full MSTSE and informal MSTSE with respect to the structural elements of chapter 6 is a large decrease in importance of the Logistics category. Questions of negotiation timeline and private information are no longer pertinent when only one stakeholder is contributing to the tradespace and its analysis. However, when beginning the exploration and analysis phase, the recommendations for dealing with limited decision maker access using proxies remain potentially relevant for informal MSTSE. Rather than each stakeholder sending a reliable proxy for their own interests in order to capture the benefits of MSTSE negotiation on a time-delayed feedback loop, a single stakeholder can have analysts role-play as the other stakeholders to conduct a “mock” negotiation. Role-play can be an

---

<sup>16</sup> This includes assumptions such as the modeling parameters or future scenarios they would choose to consider when evaluating the alternatives.

effective means of exploring or understanding the mechanics that govern the environment the player is in<sup>17</sup>, which is a key goal of simulating MSTSE for the benefit of a single stakeholder.

Even if role-play is not possible (for example, if one analyst is responsible for the entire analysis), informal MSTSE should still be approached from the perspective of *multiple individuals*. The new tradespace visualizations of this research can still be used outside of a negotiation setting to analyze the relationships between stakeholders. However, classic TSE visualizations such as the standard benefit-cost tradespace may *increase* in importance when conducting informal MSTSE, particularly if there is concern that the other stakeholders will be using naïve single-stakeholder techniques and not incorporating negotiation reference points such as the BATNA into their analysis. Considering these other plots may give a more realistic impression of how other stakeholders will view the alternatives in the tradespace, especially if they are performing their own individual TSE analysis in line with the current paradigm. Contrasting the most attractive alternatives between the classic TSE plots and the MSTSE plots for each stakeholder can help prepare the driving stakeholder for (1) what will likely be offered and (2) what counteroffers are most likely to be accepted by all parties.

This chapter has compiled the insights of this research into a small set of general recommendations for applying MSTSE, with or without the participation of stakeholders enabling face-to-face negotiation with the tradespace. Table 7-1 summarizes this information and includes the chapters which directly address each.

---

<sup>17</sup> See (Bartle, 1996) for a game-centric description of this type of role playing. It is serendipitous that Bartle refers to this type of player as an “explorer”. We believe the analogy to the use of role playing for exploring and understanding the “mechanics” of a negotiation is straightforward and part of the reason for the popularity of role play as a research method in the negotiation literature.

**Table 7-1: Summary of MSTSE recommendations, chapters to refer to for more detail on each, and applicability for informal MSTSE**

<b>Phase</b>	<b>Recommendation</b>	<b>Refer to:</b>	<b>Informal MSTSE</b>
<b>Problem Formulation</b>	Capture macro frames	3.4.1 and 5.6	All of these apply except for capturing macro frames of other stakeholders. Make best estimates for stakeholders' BATNAs and value models.
	Create many alternatives	3.1 and 6.3.3	
	Record key elements of problem structure	6	
	Determine each stakeholder's BATNA	3.5 and 6.3.1	
<b>Modeling / Evaluation</b>	Joint Fact Finding and Collaborative Modeling	2.1.4	Treat modeling as normal TSE
	Private information	6.5.3	
<b>Exploration / Analysis</b>	Emphasize the BATNA	3.6 and 4	Continue to use BATNA-centric visualizations and analyze relationships, but limit activities related to changing stakeholder value models without their participation.
	Limit strictly individual analysis	4 and 5.5	
	Analyze relationships	5.3 and 5.4	
	Allow stakeholders to change their mind	2.2 and 6.4.3	
	Refer back to macro frames	3.4.1 and 5.6	

The following chapters will demonstrate the application of these insights. The next chapter analyzes the potential to deploy MSTSE within larger systems engineering methods, considering their ability to support both tradespace exploration and the framing they provide for multi-stakeholder decisions. This is followed with the application of informal MSTSE on two case studies.

## 8 MSTSE within SE Methods

As previously mentioned, MSTSE is not positioned in this research as a standalone, end-to-end, systems engineering method. These methods typically feature step-by-step walkthroughs with defined inputs and outputs, gated documentation procedures, and recommended analysis techniques. MSTSE is an analysis technique that extends the foundation of the TSE and MATE literature to better accommodate the needs of multiple and distinct value statements. As such, MSTSE is likely to be deployed *within* a larger process when that process is applied to a multi-stakeholder problem. This chapter addresses the ability of MSTSE to be readily incorporated into a few modern SE methods as described by their primary sources. Topics of interest include the assignment of MSTSE activities to the appropriate steps of each process, the value-added by MSTSE to the standard process, and the identification of macro framing concerns that should be addressed before applying MSTSE (and possibly before doing any multi-stakeholder analysis). Specific insights resulting from this comparison may be limited to a particular method, but generic insights are likely applicable to a broad range of SE methods given their frequently considerable underlying similarity.

### 8.1 The CLIOS Process

The CLIOS Process is designed to assist in the study of (and decision making behind) Complex, Large-scale, Interconnected, Open, Sociotechnical (CLIOS) systems. These systems nearly always have multiple involved stakeholders, making MSTSE a potentially useful technique for improving communication and decision making between them. As a complete-lifecycle support process, spanning from definition and description to prescriptive analysis to continuous monitoring and support, the CLIOS Process can both provide information to assist in the setup of MSTSE and utilize MSTSE to better accomplish its analysis goals. Additionally, MSTSE should be easily included within the CLIOS Process, due to the latter's modular capacity to add supporting models and frameworks to its basic structure, referred to as "ornaments" on the larger process. In particular, tradespace exploration can be viewed as a specific variety of the recommended "tradeoff analysis", as it is a highly detailed form of tradeoff analysis geared towards systems (like CLIOS systems) that have many design variables and value-driving attributes. Additionally, MSTSE adds further support for the challenges associated with the presence of many stakeholders over standard tradespace exploration. The following breakdown is based largely on the CLIOS Process as outlined in the CLIOS User's Guide (Sussman et al., 2009) as well as a more recent update with revised terminology (Sussman et al., 2015).

#### 8.1.1 CLIOS Stages

The CLIOS Process is divided into three stages corresponding to three different conceptual tasks needed to understand, evaluate, and manage the system: (1) Representation, (2) Design, Evaluation, and Selection, and (3) Implementation and Adaptation. At a high level, the problem formulation of stage 1 will support the MSTSE activity while MSTSE will support the

goals of the CLIOS Process in stages 2 and 3. Here we will discuss the relationship between MSTSE and each of these stages, presented with their “key ideas”:

## **1. Representation**

- a. Understanding and visualizing the structure and behavior**
- b. Establishing preliminary goals**

The representation stage is analogous to the problem formulation step that feeds into MSTSE. Deploying MSTSE as a technique within the CLIOS Process should leverage the work done in this stage as a seed for the modeling tasks needed to create a tradespace. The “physical domain” will correspond to the evaluative model of potential systems and the “institutional sphere” will be used to identify the correct stakeholders who should be invited to participate as well as assist in the formulation of their value models. The task of “establishing preliminary goals” can also feed into the macro framing efforts of MSTSE by including the goals of each stakeholder as a participant in the negotiation *beyond* just their value-driving attributes, ensuring that shared decision making will not be derailed by unexpressed, mismatched assumptions about the purpose of negotiating.

## **2. Design, Evaluation, and Selection**

- a. Refining goals aimed at improvement of the CLIOS System**
- b. Developing bundles of strategic alternatives**

Stage 2 is where MSTSE can be most useful in supporting the goals of the CLIOS Process because MSTSE is positioned as a technique for helping stakeholders assess and select from many alternatives with complex value propositions. Computational (as opposed to diagrammatic) modeling occurs in this stage, driven by the resulting structure and insights of the representation stage. The level of fidelity used in the models can change (as necessary or as desired by the stakeholders) from case to case, and thus may require different levels of modeling effort: for example, in a more complex system such as a Distributed System of Systems it may become necessary to use the links of the institutional sphere within the evaluative model if the interactions of stakeholders dynamically affect system performance.

The frequent application of the CLIOS Process to infrastructure projects, which often must account for existing infrastructure as a key component of analysis, has given CLIOS a more detailed treatment for accommodating existing systems than generic tradespace exploration, which is more often applied to built-from-scratch systems and architectures. This is advantageous for MSTSE, as the determination of an improvement in the system requires the establishment of a baseline, current level of performance that implies a bare minimum BATNA for each involved stakeholder. Some or all of the involved stakeholders may have a superior BATNA corresponding to some form of action they are capable of taking unilaterally and these should also be accounted for before beginning exploration of the data. Each potential “bundle” of strategic alternatives would be a point in the tradespace of possible implementations, thus

making the design variables the individual decisions that comprise a bundle. The development of those bundles (plus any other designer-controlled variables that may have been deemed too low-level to be applied at the top-down bundle level) creates the design space of interest to the exploration. At this point, MSTSE can assist the decision makers in evaluating the strengths and weaknesses of the different bundles, providing prescriptive advice on the future direction of the system.

### **3. Implementation**

#### **a. Implementing bundles of strategic alternatives**

#### **b. Following-through – changing and monitoring the performance of the CLIOS system**

As a technique for selecting amongst alternatives, MSTSE has little input on the eventual implementation of the chosen bundle. However, integration with naturally tradespace-supporting techniques such as Epoch-Era Analysis (Ross and Rhodes, 2008) could provide a backbone for monitoring the system over time as the surrounding context changes. Given the logistical challenges associated with collecting decision makers for MSTSE, the monitoring task would likely have to fall upon a “watchdog” stakeholder, with some criteria set for when a decision mandating a new MSTSE session is needed. This allows MSTSE to support continued collaborative decision making over the lifetime of the system.

#### **8.1.2 CLIOS Steps**

The CLIOS Process is further divided into 12 steps identifying specific tasks that must be accomplished, as shown in Figure 8-1. MSTSE can leverage some steps and support the goals of others, which will be covered briefly here.

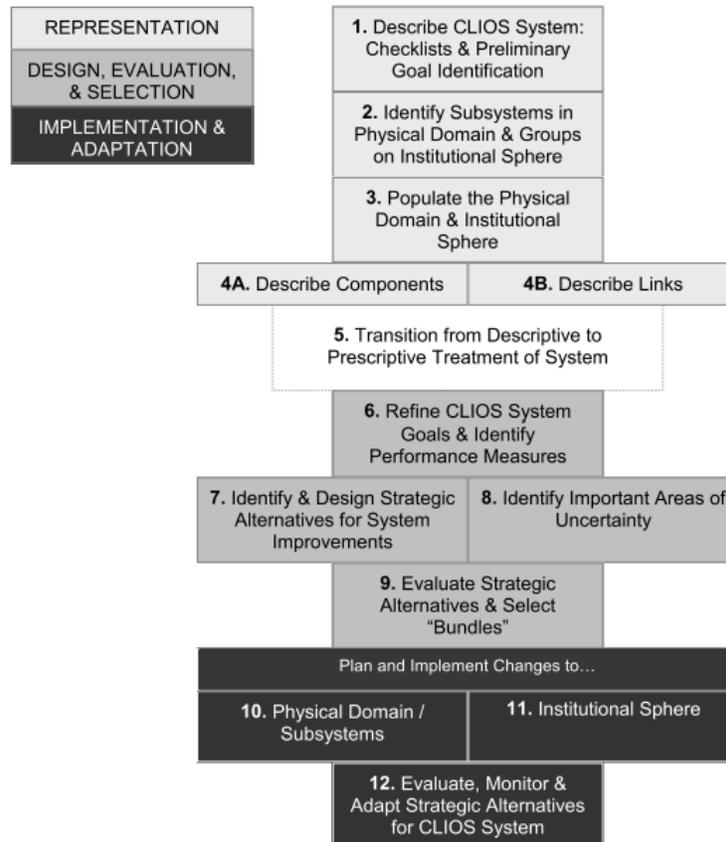


Figure 8-1: The Stages (left) and Steps (right) of the CLIOS Process (Sussman et al., 2009)

### 1. Describe CLIOS system: checklists and preliminary goal identification

The checklists and preliminary goal identification step of the CLIOS Process can serve to focus the value-driven aspects of MSTSE. The “opportunities, issues, and challenges” list will contain useful information regarding the positive and negative attributes of the system that will eventually drive value models in MSTSE. Additionally, if the system analysts are planning on utilizing MSTSE, the listing of preliminary CLIOS system goals is an excellent opportunity to engage stakeholders about their motives for participating in the design process *beyond* standard performance measures – macro frames that will impact the ultimate decision. Explicitly stating these motivations can improve communication and head off potential future conflicts arising from philosophical differences and unspoken metapreferences. For example, identifying that one stakeholder believes his job is to maximize system efficiency while another believes that his is to simply maximize performance will reduce the likelihood that they “talk past” each other when later debating what action to take. This is particularly important when planning to eventually discuss flexibility in the system, as these philosophies reflect not only on the initial design, but also the strategy for continued management and execution of options later in the system’s lifecycle (Fitzgerald, Ross, and Rhodes, 2012).

It is worth noting here that the CLIOS Process is intended to be iterative and, in that vein, it may be easier to begin the process alone (as a single interest group) and then iterate with additional stakeholders brought on. If using MSTSE, this would mean pursuing informal MSTSE for the first iterations and only convening stakeholders for face-to-face negotiation on later iterations. If that is the case, this step will be filled with “best guesses” for the goals of the other relevant parties on the first pass, with the intention to revise those guesses where necessary after seeking the input of other stakeholders.

**2. Identify subsystems in physical domain and groups on institutional sphere**

**3. Populate the physical domain and institutional sphere**

**4.**

**a. Describe components**

**b. Describe links**

These steps model the system of interest diagrammatically, which will eventually drive the models used to evaluate the design alternatives that populate the tradespace in steps 6-9. Eventually, the physical components will typically become either self-contained analysis modules (e.g. “traffic congestion” will be modeled by a discrete-event simulator) or design variables (e.g. “highway infrastructure” is a variable that can be built in different ways). Policy levers will almost exclusively be design variables, as they represent decisions that participating stakeholders are able to make, while regulatory constraints (e.g., policy that cannot be changed) will limit the set of valid or feasible alternatives. The links between components describe information flow between the models and links between stakeholders can be used to organize hierarchical value models, if necessary. Additionally, if a tradespace-oriented uncertainty framework such as Epoch-Era Analysis is going to be used, the external factors identified in this step will likely become the key uncertainties: treated as fixed in a given context but capable of changing and forcing a new tradespace to take shape.

**5. Transition from descriptive to prescriptive treatment of system**

This step is called out in the CLIOS Process to remind analysts and stakeholders of the difference between describing the system’s features and planning to improve the system’s value delivery. Tradespace exploration in general and Multi-Attribute Tradespace Exploration (MATE) in particular have frequently been positioned as prescriptive aids to decision making, aligning themselves with the overall objectives of this stage of the CLIOS process (Ross, 2006).

**6. Refine CLIOS system goals and identify performance measures**

MSTSE is explicitly intended to support decision making across multiple goals of multiple stakeholders. This step is where the value models for each stakeholder should be constructed. Multi-Attribute Utility Theory (MAUT) is recommended as a highly effective and customizable value model for use in tradespace exploration (Ross et al., 2010), but in application each

stakeholder is able to use whatever value model they choose, even if it is simply a one-to-one mapping of a single performance attribute. The explicit value models can build from the preliminary goals of Step 1 but do not need to copy them exactly, as the descriptive information of the previous steps may have resulted in modified preferences for one or more stakeholders.

The CLIOS User's Guide reminds the reader that frequently consensus will not be reached about what performance metrics should be used or what their priorities should be. MSTSE is an effective addition to the stable of "ornaments" in the CLIOS toolbox because it explicitly does not attempt to reach consensus on a single value model, but rather allows each stakeholder to control their own (Fitzgerald and Ross, 2013). It is thus a particularly effective technique for those cases in which stakeholders are too diverse to consolidate their preferences. However, MAUT is also effective at including shared goals (sometimes manifesting within the CLIOS Process as "common drivers"); because each attribute has its own utility curve, stakeholders can choose to ascribe the same single-attribute utility of a shared goal to their personal multi-attribute function.

## **7. Identify and design strategic alternatives for system improvements**

This step is analogous to the enumeration of the design space for a tradespace exploration. The design variables identified in the descriptive stage are now divided into discrete, actionable alternatives. For example, a highway may be built with two, three, or four lanes. This is a creative process and thus stakeholder input is encouraged, particularly for design variables corresponding to their own policy levers. The more feasible alternatives that are established (in line with the principled negotiation tenet of creating many options), the richer the eventual tradespace will be. Tradespaces often include a full-factorial enumeration of design variables when it is computationally feasible: each possible combination of levels forming an individual design vector that can be evaluated by the performance model. These vectors are the "bundles" of the CLIOS Process terminology, corresponding to a set of strategic alternatives enacted. If a full-factorial enumeration is not feasible due to demanding computational requirements, either a standard Design of Experiments (e.g. Latin Hypercube) sampling technique can be used or the participating stakeholders can hand-select bundles of interest for evaluation.

## **8. Identify important areas of uncertainty**

Uncertainty is often an important driver of performance in CLIOS systems and scenario planning is one of the recommended techniques for addressing uncertainty, often driven by the external factors identified in the representation stage. Tradespace exploration has a variety of extensions that can be employed to capture uncertainty including Epoch-Era Analysis, which often includes a form of narrative-driven lifecycle construction that has a direct analogue to scenario planning. Identifying areas of uncertainty can lead to the definition of context variables, and these context variables can be enumerated into a complete uncertainty space, in which a point (*epoch*) corresponds to a scenario with a fixed context and a resulting tradespace. Additionally,

tradespace exploration has employed real options analysis to quantify flexibility in systems (Shah et al., 2008) and new EEA-based techniques such as the Valuation Approach for Strategic Changeability (VASC) are emerging for tradespace-centric flexibility analysis (Fitzgerald, Ross, and Rhodes, 2012). Any of these tools can be included as desired, though the cognitive load imposed on the stakeholders to negotiate over many tradespaces may make a simpler sensitivity analysis to the uncertainty variables preferable.

### **9. Evaluate strategic alternatives and select “bundles”**

This is the step where the actual tradespace exploration of MSTSE should be performed. The evaluative models can be run, creating a tradespace of the “bundles” of strategic alternatives the stakeholders are capable of choosing. Following the tradespace exploration paradigm in this way incorporates, by default, any cross-effects and dependencies between different decisions the stakeholders can make as they are evaluated *as* bundles, rather than evaluated and then assembled into bundles. MSTSE can help the participating stakeholders explore the space of possibilities for the system and, hopefully, come to an agreement on a course of action. This process can include the added insights of EEA, VASC, or other tools to capture value derived from robustness and/or flexibility as the stakeholders desire.

### **10. Plan and implement changes to physical domain / subsystems**

### **11. Plan and implement changes to institutional sphere**

### **12. Evaluate, monitor, and adapt strategic alternatives for CLIOS system**

The first parts of these steps include making sure that the tradeoffs amongst the different performance attributes are such that all the participating stakeholders are satisfied and no parties feel marginalized. Using MSTSE in the evaluation stage should make this process smoother, as the negotiation setting should allow stakeholders to feel more engaged and empowered to make a decision that suits them, over simply maximizing a weighted, group value function. Regardless, reviewing the final decision to verify that it is both physically and institutionally feasible is still advisable, even with the use of MSTSE.

Maintaining the models used for evaluating the system and updating the input parameters as exogenous contextual variables change over time, assumptions are proven false, or new interactions arise can allow for a continually updating tradespace. One or more parties can monitor the quality of the implemented solution for each stakeholder, with the possibility of reconvening a new MSTSE session (or iterating through the CLIOS Process again) should it become apparent that it is possible to improve due to degrading performance or new alternatives.

### **8.1.3 Relation to SAM-PD**

Stakeholder-Assisted Modeling and Policy Design (SAM-PD) is a process extending the CLIOS system analysis framework with fundamentals of the consensus-building literature, in order to promote collaborative decision making between stakeholders (Mostashari and Sussman,

2004; Mostashari, 2005). SAM-PD emphasizes the role of stakeholder conflict assessment, Joint Fact Finding (Ehrmann and Stinson, 1999), and consensus-seeking negotiation as key steps towards generating goodwill and buy-in from participating stakeholders, all of which have been identified as valuable problem formulation and modeling activities for MSTSE in this research. In this sense, SAM-PD has a similar relationship to the CLIOS Process as MSTSE does to its precursor, MATE. MSTSE can be deployed during the SAM-PD step of Model-Based Consensus Seeking Negotiation. This step maps directly to Step 9 of the CLIOS process, as shown in Figure 8-2, which is where the “action” of tradespace exploration is recommended to be performed. The earlier steps of SAM-PD, particularly Stakeholder Conflict Assessment, can be effectively leveraged through MSTSE by using the more detailed breakdown of the values of each stakeholder to inform superior value models and better facilitate communication between stakeholders.

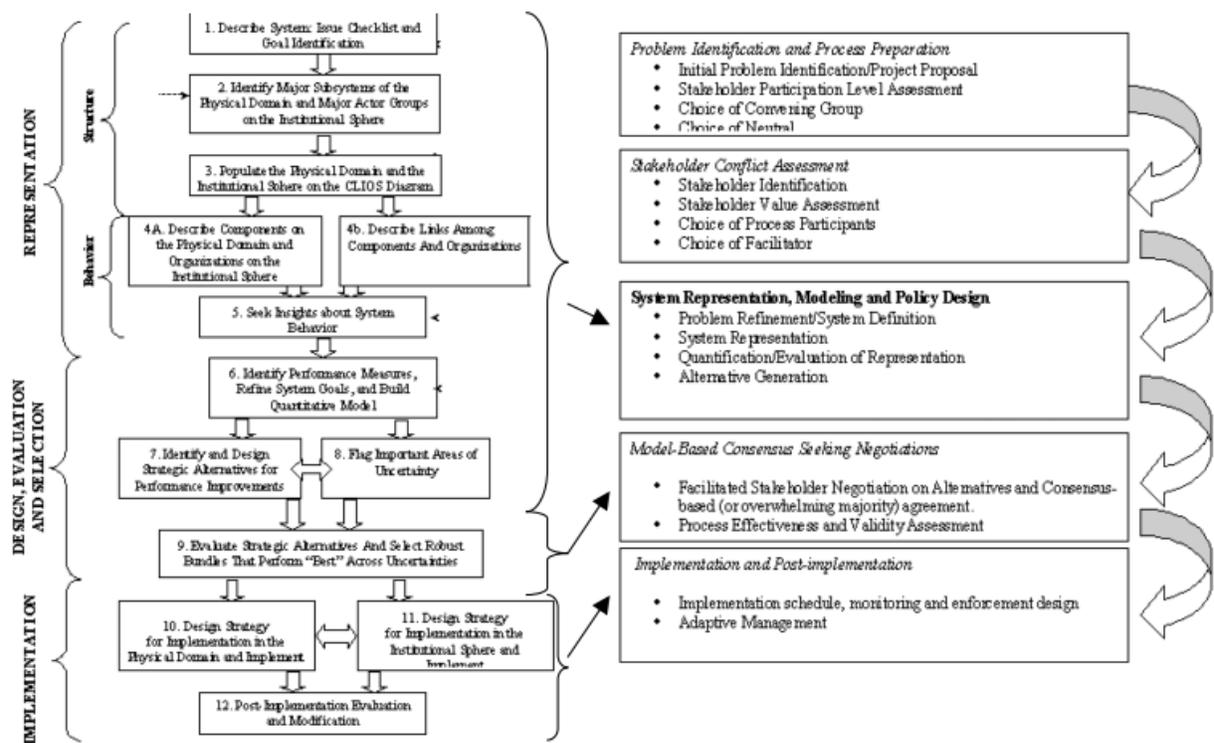


Figure 8-2: Mapping the CLIOS and SAM-PD Processes (Mostashari, 2005)

### 8.1.4 Summary

Overall, the CLIOS Process has an excellent correspondence with the principles of MSTSE. It has considerable top-down similarity with the phases of TSE as described in chapter 7: beginning with detailed problem formulation, moving to modeling, and then analyzing discrete potential decisions. Additionally, the approach for including the value statements of multiple stakeholders within the CLIOS Process is well-developed due to the expectation that CLIOS systems are rarely single stakeholder. The emphasis on prescriptive analysis of the

relationships between stakeholders is directly enhanced by the application of MSTSE, particularly when the SAM-PD extension is leveraged for consensus-building stakeholder involvement during the problem formulation.

The most significant discrepancy between the CLIOS Process and the typical application of MSTSE is in the description of the “bundles” that form the design space. In the CLIOS Process, the bundles usually number less than ten, as they are heavily vetted by the analysis team as feasible and potentially interesting alternatives. This is an accommodation to the relatively complex models necessary to encapsulate a complete CLIOS system, potentially requiring some degree of bundle-specific modeling in order to properly capture any complex interactive effects of policy levers. MSTSE prefers to operate on a larger set of alternatives, as it enables more effective pattern-finding and certain tradespace metrics lose relevance when applied to few alternatives. The small tradespace size is likely the biggest challenge for the application of MSTSE within the CLIOS process. However, in cases where a more traditional full-factorial tradespace *is* feasible, MSTSE becomes an even more valuable “ornament” for the CLIOS Process, as it can provide relevant guidance for the analysis of large spaces of alternatives that is not directly addressed by the main process.

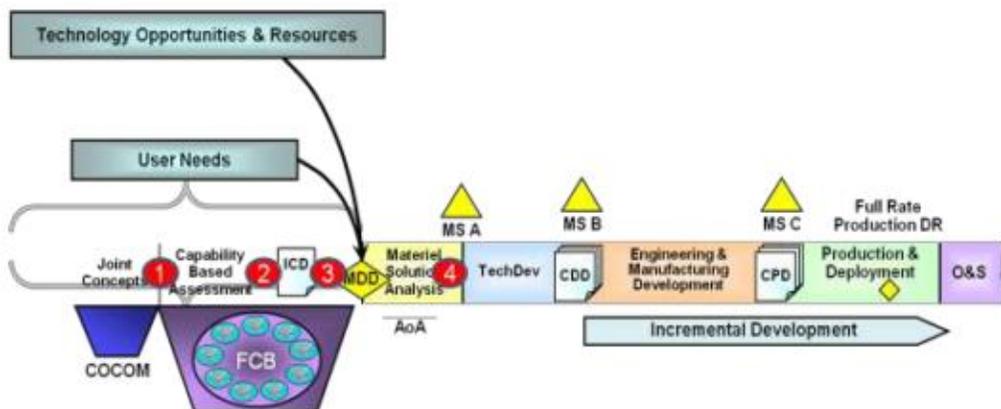
## **8.2 USAF Analysis of Alternatives**

Analysis of Alternatives (AoA) is a framework employed by the United States Air Force (USAF) and mandated for use when upgrading existing systems or acquiring new systems. The primary goal of AoA is to provide a structured means of due diligence, ensuring that the relevant design space is explored and that the proper information on design tradeoffs is available to stakeholders at critical decision points in the design process. AoA sits firmly within the tradespace exploration paradigm (Ross et al., 2010), with emphasis placed on understanding implicit tradeoffs of value-generating attributes, and uses much of the standard TSE terminology (including “tradespace” itself). This section will briefly overview the official use of AoA and discuss the implications of its procedures when applied to multi-stakeholder problems, particularly projects between branches of the military. The following analysis is based on the description of AoA contained in the Office of Aerospace (OAS) Studies AoA Handbook (2013), with additional insights derived from (US Air Force, 2010).

### **8.2.1 Overview**

AoA is not a process that can be easily outlined with a short set of ‘steps’ that should be performed either sequentially, simultaneously, or iteratively. Rather, it is a group of “best practices” for a number of officially-required activities, supporting the analysis goals of the task while remaining freeform enough to allow the participating engineering/design team to exercise their judgement on what techniques are most appropriate for their specific project, with oversight and assistance from support organizations such as the OAS. This is deliberate because, as the AoA handbook notes, “the exact questions and issues [for a project] change over time and are based on personalities and politics,” which precludes the ability to define a fixed structure for

interaction between participating stakeholders. The structure of AoA is provided largely through the surrounding set of output form documents that feed into the larger USAF acquisition timeline. AoA is a required activity in pre-Milestone-A analysis (as shown in Figure 8-3), intended to support conceptual, early-design decision making and guide later, detailed design into the most promising area of the tradespace; this is perfectly aligned with the intended purpose of MATE and MSTSE.



**Figure 8-3: USAF system lifecycle outline, indicating AoA as pre-Milestone-A (Office of Aerospace Studies, 2013)**

Within an AoA, there are five main artifacts: the study plan, analysis, progress briefing, final report, and final briefing. These artifacts are briefly described here.

### Study plan

Defines targeted capability gap and design space to resolve that gap. The study plan must be approved by the Air Force Requirements Review Group (AFRRG) and Requirements Oversight Committee (AFROC) in order to proceed with AoA.

### Analysis

- *Effectiveness analysis*
- *Cost analysis*
- *Risk analysis*
- *Sustainability analysis*
- *Alternative comparison*

The analysis within an AoA has four main categories: effectiveness, cost, risk, and sustainability. Effectiveness analysis covers the modeling of traditional system performance metrics such as speed or image resolution. These are typically the key performance parameters (KPPs) or Measures of Effectiveness (MoEs) used to define the capability gap. Cost analysis includes the modeling and analysis of development, acquisition, operations, and disposal costs. Risk analysis involves the use of the USAF Risk Assessment Framework (RAF) to

identify and evaluate areas of technological, manufacturing, and operational risk that may inhibit the ability of a given alternative to complete its mission. Sustainability analysis focuses on traditional -ility metrics, such as mean time to repair, quantifying the ability to satisfy reliability, availability, and maintainability goals, among others. The alternative comparison analysis is where top-performing alternatives in each of these categories are compared in order to investigate the tradeoffs between the categories. Analysis is conducted with respect to “threat” scenarios, representing different contexts the system may operate in.

#### Progress briefing

Creates an opportunity for support organizations and senior reviewers to provide additional steering to the AoA direction while analysis remains ongoing.

#### Final report

Repository for all activities performed during AoA, their outcomes, and the lessons learned.

#### Final briefing

Distillation of complete AoA results into the most relevant insights into the guiding questions of the analysis. This is the main output for communicating AoA findings to decision makers and thus must cover both the outcomes (e.g. favorable regions of the design space) and also the reasoning behind those conclusions (e.g. physical intuition for tradeoffs).

### **8.2.2 Tradespace Exploration in AoA**

Generally, all AoA tasks support tradespace exploration goals, though they do not necessarily take the common form of a large set of alternatives, evaluated simultaneously and viewed on standard benefit-cost axes. The four categories of analysis (benefit, cost, risk, and sustainability) can be considered analogous to the modeling tasks underlying TSE; the AoA handbook chapters on these tasks provide advice on selection of measures, modeling techniques, and other fundamental data *generating* activities. These are recommended to be performed with considerable human input, gradually refining the areas of interest, which ultimately limits the set of considered alternatives passed into the “compare alternatives” task to those that perform well in at least one of these four categories of metrics. Enumerating the design space in this way has pros and cons associated with it compared to an agnostic full-factorial or randomly-sampled design space. The positives include (1) increased focus on ‘interesting’ designs, (2) reduced computation time due to reduced number of alternatives, and (3) increased granularity of alternatives in interesting regions of the tradespace (Cilli and Parnell, 2014). The negatives include (1) reduced ability to visualize and develop insight on tradeoff patterns, (2) possibility of introducing human selection bias into the compared alternatives, and (3) difficulty finding interesting designs with nearly-optimal performance in multiple attributes (Parnell, Cilli, and Buede, 2014).

The alternative comparison task is where high-level exploration of the tradespace can occur. The AoA handbook generally focuses on examples of 2-6 alternatives of interest, but tradespaces with considerably more alternatives are not discouraged. Frequently in TSE projects, the effectiveness MoEs would be used as the basis for a multi-attribute value function (utility, AHP, etc.) to be the benefit axis in the benefit-cost tradespace view. The AoA handbook specifically advises *against* the use of multi-attribute functions in favor of analyzing the tradeoffs between all value-generating attributes directly in order to avoid the risk of unintentionally obscuring poor-to-infeasible performance in one or more objectives or to burden high-level decision makers with potentially zero familiarity with value modeling. Certainly, a detailed look at individual attributes is beneficial to develop greater understanding of the available tradeoffs in the design space; however, looking at multiple measures demands considerably higher cognitive load on the analyst or decision maker than using a single measure of benefit. A well-designed multi-attribute value function (such as a Keeney-Raiffa utility function) will adequately inform the analysts when any attribute is infeasible as opposed to undesirable, while allowing for easy, high level understanding of solution quality and access to detailed information about individual attributes on demand (Keeney and Raiffa, 1993). It is worth mentioning that, while the handbook rejects this approach, it is likely that multi-attribute value functions are still used by some analysts even if they are not included in the AoA final report.

### **8.2.3 Multi-stakeholder in AoA**

The AoA handbook offers relatively little instruction specifically concerning multi-stakeholder procedures. This may often be a non-issue since, as a USAF process, the primary stakeholder will frequently have the Air Force's strategic objectives as a value proposition. Other identified stakeholder types may include end users and enabling organizations. However, when the USAF participates in a design task determined to be of joint interest to the Department of Defense (meaning the targeted capability gap and system type are applicable across branches of the military), other participating branches become decision makers in the project. If the USAF is the project lead, the AoA process is still used but with a key difference in the study plan task. For joint projects, the study plan also requires the approval of the Joint Requirements Oversight Council (JROC). The JROC receives the proposed study plan and then "validate[s] the threshold and objective values" in the system requirements. This places the JROC in a position akin to that of an arbiter, unilaterally deciding the target system requirements after receiving input from the participating stakeholders. Once the JROC returns the approved study plan, AoA proceeds as normal, with instruction to take input from all stakeholders as the only concession to the multiple participating parties. The JROC-provided requirements are treated as a single value proposition, investigated using all the standard, single-stakeholder AoA analysis techniques.

During Materiel Development Decision (MDD), an analysis step pre-AoA, an interesting macro framing standard is held in the form of "preferred" solution types. Six types of solutions

are identified, with the mandate that less-preferred types must be accompanied by evidence justifying their selection over the superior types. In decreasing order of preference, they are:

1. Doctrine, Operations, Training, (existing) materiel, Leadership/education, Personnel, and Facilities Analysis (DOTmLPP-P) solutions
2. Commercial, off the shelf solutions
3. Modification of existing US or allied military system
4. Cooperative development with allied nations
5. Joint DoD program
6. DoD unique component development

Obviously, the preference is for easier and less expensive solutions over more expensive solutions, and for shared applicability over problem-specific solutions. This supports multi-stakeholder analysis by labeling it as a priority. However, the implicit emphasis on cost savings is even higher, reflecting an organization-wide priority on low cost solutions that places the burden of proof on joint solutions to be “worth the cost”. This priority is apparent across many other official sources, including the Better Buying Power initiative and interagency memoranda instructing “every effort to reduce the cost of products and services we acquire” (Kendall, 2011, 2014). A mandate on minimizing cost could potentially drive strong attachment to the Pareto front in multi-stakeholder problems, as each stakeholder must prove the cost-effectiveness of the solution.

Another interesting macro framing issue is that of information sharing. AoA is officially in support of the FOTE principle (Raiffa, 2002), stating that “frequent and open exchanges of ideas and data are essential to a successful AoA”. This stance should keep stakeholders open minded to incorporating new information and searching for mutual benefit instead of individual benefit.

#### **8.2.4 Using MSTSE in AoA**

The prevalence of TSE language and activities in the existing AoA process eases the adoption of MSTSE, as the underlying tradespace structure is preserved and the participants are likely already familiar with the basic concepts. The major hurdles to applying MSTSE within AoA are related to the proper handling of the various stakeholders in order to find the best mutually beneficial solutions in the tradespace. Focusing here on joint projects involving multiple military branches, there are nontrivial barriers to the application of many of the insights of negotiation theory and framing present in AoA that may hinder inclusion of MSTSE activities.

First, it is worth returning to the recommendation *against* the use of multi-attribute value functions. As previously discussed, the tradeoff inherent in this decision is typically one between ease of use (of a single function) versus increased detail (of individual attributes). The AoA recommendation, however, is made specifically with an eye towards the decision makers

who receive the final briefing, citing the additional “complexities” of value functions that are “difficult to understand and difficult to explain to decision makers”. This is somewhat of a false dichotomy since the use of a multi-attribute value function only creates a new metric: the individual attribute calculations are still performed and able to be analyzed directly if that becomes desirable. Aggregation functions can be deployed for stakeholders and/or analysts who desire them and are adequately informed of their strengths and weaknesses, while still allowing for analysis of individual attributes for less sophisticated stakeholders or for verification of specific capabilities. This issue increases in importance when engaging with multiple stakeholders. In addition to wanting to communicate with each stakeholder at a level of detail most comfortable to them, the number of individual attributes of interest increases with the number of participating stakeholders, increasing the challenge associated with adequately considering all of the attributes simultaneously. Reducing the problem to a single benefit and cost function for each stakeholder is a valuable simplification to the complexity of the multi-stakeholder problem and assists the communication of stakeholder-level value more easily and effectively than the separate attributes, while still allowing for deeper analysis when necessary.

TSE is often recommended to be performed by the decision makers themselves when possible, as the benefits of rapid exploration and feedback are lessened when proxy decision makers must do the analysis and then pass information up the chain of command. Obviously, the logistical realities of the DoD (e.g. multitude of projects, strict hierarchy) limit the ability for decision makers to spend the time required for TSE on all of the relevant projects concerning their interests. AoA is intended to be performed by analysts on behalf of decision makers: distilling the key insights from TSE into easily consumed chunks in order to enable decision makers to make the best use their limited time. Deploying MSTSE within AoA will likely require the use of analysts as proxy decision makers (as in chapter 6.5.1), even though the chain of command delay is exacerbated by the presence of multiple negotiating interests.

Hypothetically, each branch participating in a joint AoA could designate an analyst to serve as a proxy decision maker for the purposes of engaging in MSTSE activities. Ideally that analyst would be as familiar as possible with the preferences of the decision maker they represent, in order to minimize the need for time-consuming iteration between the AoA team and the decision authority. However, the role of the JROC in joint AoA precludes the ability for analysts to freely negotiate across the design space in search of mutual benefit. The JROC-approved study plan specifies threshold and objective values for the capabilities requested by each stakeholder. In some sense, this requirements list is a pre-negotiated outcome, as the JROC takes the submitted requirements of each stakeholder and balances or trades them off into a final, complete list that the program office is recommended to explore with single stakeholder AoA techniques. Ironically, this aggregation of the *distinct, individual preference sets* of each branch into a single block of requirements results in the same risk of lost situational awareness that the AoA handbook sought to avoid with its rejection of multi-attribute value functions. For example, with no explicit separation of the different stakeholders’ objectives, it may go

unnoticed until later design reviews that, for a given alternative, all the sacrificed threshold objectives belong to the same stakeholder and that all of the reached objective thresholds belong to another. This can result in acrimonious differences between stakeholders later in the design process and costly rework or cancellation. Additionally, without explicit attention paid to which stakeholders care about which attributes, it is possible to miss value-generating tradeoff opportunities, where low-importance capabilities for one stakeholder can be sacrificed for high-importance capabilities for another. MSTSE encourages consistent separation of preferences between stakeholders, as the interaction *at the stakeholder level* is of primary importance to finding and identifying a mutually agreeable solution (Fitzgerald and Ross, 2014).

AoA requires that the current capability levels in all KPP/MoE attributes be explicitly accounted, in order to support the accurate determination of capability gaps. This feeds directly into BATNA generation for each stakeholder, as the current capability-at-cost represents an alternative solution to any evaluated design points. Each stakeholder's BATNA will be at least as good as the currently implemented system(s) (Fisher, Ury, and Patton, 1991). The BATNAs may improve over this baseline level if potential non-joint system alternatives are included in the tradespace enumeration, an activity that is also supported by the AoA task of investigating multiple "solution categories" including modifying existing systems, increasing production of legacy systems, and use of commercial, off the shelf systems through the use of accompanying analysis methods such as DOTmLPF-P Analysis and MDD. Proper identification of the BATNA for use as a reference point in the tradespace is thus highly feasible within AoA. It should be noted that the AoA handbook has some discussion of "reference points" throughout, specifically with regards to the designated threshold requirements. Threshold requirements are indicated to "frame the decision space" rather than "disqualify" alternatives that fail to meet them. This is sound advice from a TSE perspective (alternatives should not be summarily discarded based on a single criterion) but also indicates that the threshold requirements should not be unduly emphasized because they could cause solutions that exceed baseline capability but do not reach threshold requirements to be incorrectly viewed as losses. A physical, achievable BATNA such as the baseline system remains a superior choice for reference point within the tradespace. The BATNA may change in value between "threat" scenarios: this can be explored through the use of tradespace uncertainty methods such as Epoch Era Analysis (EEA; Ross and Rhodes, 2008).

### **8.2.5 Summary**

MSTSE addresses a number of deficiencies in the handling of multiple stakeholders in joint AoA projects, including keeping their interests separate and using the BATNA as a reference point instead of requirements. With the underlying TSE structure already in place, it would seem to be a relatively easy move to include MSTSE activities during an AoA project, however there are key logistical barriers present in the USAF and DoD organizations that may limit the effectiveness of MSTSE. Some of these barriers can be feasibly overcome with stopgap measures, such as replacing decision maker participation with analysts who are thoroughly briefed in the high-level value propositions of their branch. Others are derived from problematic

policies seemingly arising from a fear of “getting it wrong”, such as JROC top-down requirements management or rejection of multi-attribute value functions, preventing stakeholders and analysts from exploring all possible avenues of creating mutual value. These policies were certainly created with the best of intentions and it is far beyond the scope of this research to criticize the many intricacies of DoD organization, but the rules of AoA as they currently exist likely have unanticipated consequences when multiple branches of the military seek to work together. In application, it may be that joint AoA tasks resemble MSTSE more closely than the AoA handbook suggests, as it would be foolish to assume that all AoA projects follow the handbook to the letter. With that understood however, revision to the official AoA procedure for joint projects could serve to improve messaging consistency and promote use of MSTSE techniques for finding the areas of the tradespace with the greatest potential value across all stakeholders rather than across the union of all stakeholder attributes.

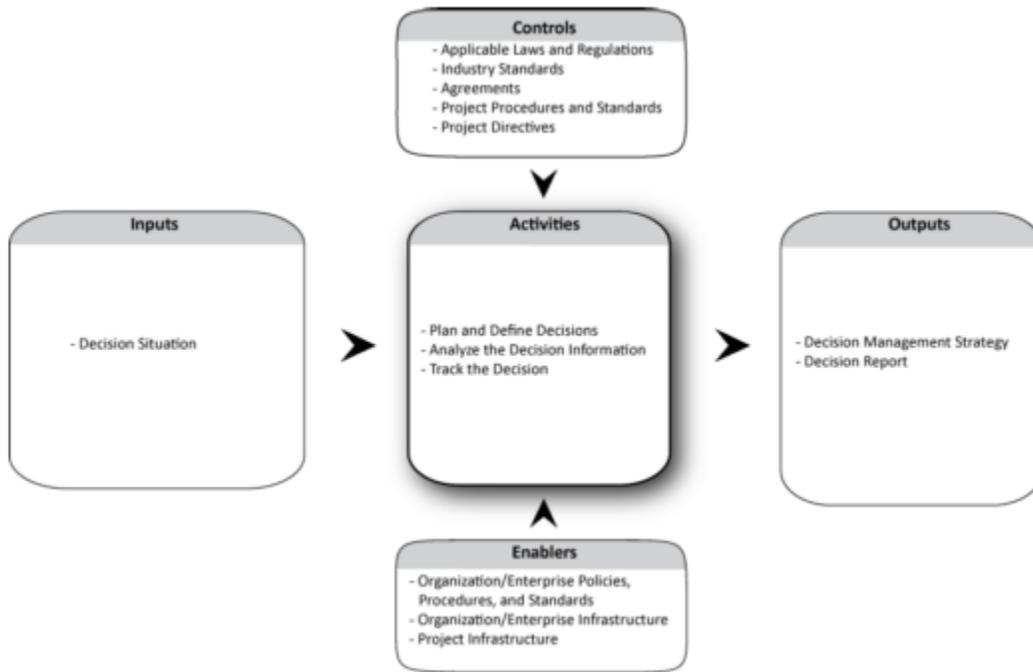
### **8.3 INCOSE Decision Management Process**

The International Council on Systems Engineering (INCOSE) is a community organization of systems engineers and researchers that offers a variety of services including conferences, continuing education, and certification. The INCOSE Systems Engineering Handbook (v3.2.2: INCOSE, 2012; v4.0: INCOSE, 2015) is a popular publication of INCOSE that regularly updates to include new research and best practices, serving as a reference for practitioners looking for guidance in implementing systems engineering methods and principles. Part of the SE Handbook is the INCOSE Decision Management Process, which was substantially revised for version 4.0 of the handbook, including additional detail on trade studies under the headline of a Systems Engineering Tradeoff Study Process Framework (Cilli and Parnell, 2014). This section will discuss the basic formulation of the larger Decision Management Process in its older and updated forms along with specifics of the recommendations for trade studies, interspersed with analysis of the readiness of this process to incorporate MSTSE.

#### **8.3.1 Process Overview**

The SE Handbook v3.2.2 states that “the purpose of the Decision Management Process is to select the most beneficial course of project action where alternatives exist”. It goes on to clarify that “beneficial” here can mean a number of different things depending on the project, including specified, desirable, or optimized. Given the large number and variety of decisions with multiple alternatives, the intended purview of the Decision Management Process is wide. Therefore, the process is generally discussed in broad terminology applicable to many domains and levels of abstraction, leaning slightly towards broad, strategic decisions. The “context diagram”, shown in Figure 8-4, illustrates the key inputs, enablers, controls, activities, and outputs of the process. The entries in this diagram are abstract, including “decision situation” as an input and “analyze the decision information” as an activity, reinforcing the role of the Decision Management Process as a high-level approach to tackling generic decisions. Definitions for these items can be found in the handbook if the reader is curious, but they are

generally self-explanatory. Of particular interest to this research is the “agreements” control, which implies the existence of multiple stakeholders who hold those agreements. In the role of a control, prior agreements may be responsible for defining the scope of the decision or apply constraints on the final decision outcome. With multiple stakeholders, the output “decision management strategy” is *also* an agreement, as it should involve each stakeholder and presumably satisfies all previous agreements.



**Figure 8-4: INCOSE Decision Management Process "context diagram" (INCOSE, 2012)**

The descriptions of the three activities in the Decision Management Process include many of the same goals and pieces of advice present in the negotiation and TSE literature. “Plan and define decisions” encourages the early determination of a “strategy” for decision making as well as the definition of measurable (objective) criteria to base the decision on. MSTSE could be considered a problem-solving strategy in itself, as a technique for investigating and making decisions. At a lower level, strategy may entail agreement on subjective fairness criteria used to identify the best alternatives. “Analyze the decision information” is the step that the actual exploration of the tradespace would reside within, directly supporting the directives of “involve all personnel [...] relevant to the decision” and “evaluate the consequences of alternative choices”. “Track the decision” refers to the use of proper documentation and supporting communication of the final decision.

The Systems Engineering Handbook then goes on to list a number of recommended decision making techniques. First among these is decision analysis, the longstanding field of research concerning decision making under uncertainty. Raiffa (1968) is identified as a key

contributor to decision analysis, whose later work on value modeling and negotiation has influenced this research significantly. The handbook draws particular attention to the ten principles of good decision making from Skinner (2009), many of which have immediate parallels to the concepts within MSTSE: most notably an explicit directive to “frame the real problem”, acknowledging the role that framing plays in decision making. Additionally, Skinner cites that good decision processes “develop creative [...] alternatives” and use experts to gather “reliable information”, which mirror the principled negotiation tenets of exploring many choices and using objective data. “Dialogue”, such as face-to-face negotiation, is also identified as a key component of learning and making actions clear.

Additional consideration is given to lean development and value of information modeling (as methods for optimally delaying commitment), decision trees and sensitivity analysis (as means of capturing uncertainty), and multi-attribute utility. These can be considered as the toolbox of methods that can support the Decision Management Process. The final and most detailed discussion is on trade studies, of which MSTSE can be considered a variant. The description of MSTSE in this research has heavily incorporated decision analysis and multi-attribute utility ideas, and therefore seems like a strong match for the overall objectives of the INCOSE process, especially if Epoch-Era Analysis is considered an acceptable alternative means of modeling and exploring uncertainty.

Rather than go into detail on the description of trade studies in v3.2.2 of the handbook, we will consider the elaborated Trade Study Framework of v4.0. This represents an update over the previous version, combining many of the older insights with newer ones and more tightly connecting the Decision Management Process with trade studies in the form of a new ten step framework. The following section will discuss the new framework and its connection to the larger process.

### **8.3.2 Trade Study Framework**

Version 4.0 of the Decision Management Process is presented as a ten step process designed to support trade studies. It is essentially a highly detailed breakdown of the “activities” box in the original context diagram. The outline of the process is shown in Figure 8-5, with the blue arrows corresponding to the ten steps and the green circle identifying key components of standard systems engineering tasks. The circle is meant to imply that the process is iterative, with trade studies subject to gradual refinement. MSTSE, as a type of trade study, has a similar structure to many of these steps, which will be covered here.



Figure 8-5: INCOSE Decision Management Process v4.0 draft (Cilli and Parnell, 2014)

## 1. Frame Decision

The first step is to frame the complete decision, including the determination of applicable baseline designs, available timeframe, and relevant stakeholders and decision makers. Specific emphasis is given to the consideration of the full system lifecycle, in order to capture a complete picture of whom and what should be included in the analysis. This type of holistic problem formulation is directly in line with the early steps of MSTSE, particularly in the mandate to collect the appropriate group of stakeholders. The use of the word “frame” as part of the name for this step implicitly acknowledges the importance of proper framing in problem solving, though without ever addressing it directly in the description. The principle is supported by the direction to consider baseline systems and full lifecycles, which are important steps towards the evaluation of a proper BATNA for each stakeholder.

Note that the Decision Management Process explicitly acknowledges the wide variety of potential multi-objective decision analysis (MODA) techniques that can be deployed to support decision making. Optimization and monetization are specifically called out but deemed difficult to implement for many of the same reasons identified in this research. Trade studies are identified as the most common decision analysis technique, again supporting the use of MSTSE. The potential benefit of iterating through multiple MODA techniques when time and resources

allow for additional analysis is also mentioned, as it allows for the comparison of results and an increase in decision confidence through the reconciliation of differences. This has been a matter of recent interest specifically within tradespace exploration, as TSE can leverage multiple value models and rapidly compare their outputs (Ross, Rhodes, and Fitzgerald, 2015).

## **2. Develop Objectives and Measures**

Step 2 maps into the MSTSE step of creating the appropriate value model(s). Recommended techniques here include stakeholder analysis, functional decomposition of the system, and the operationalization of targeted objectives. The Decision Management Process champions a combination of individual value functions and assessed swing weights for each objective, which matches the preferred use of multi-attribute utility theory for MSTSE, though both are able to incorporate a wide variety of potential value models. Note that there is no specific mention of keeping apart different stakeholders' objectives or the benefit and cost objectives for any given stakeholder (though benefit and cost do happen to be separated in the example of later steps). This runs the dual risk of premature compromising and loss of detail, respectively, through the over-aggregation of the relevant value metrics. Most applications of MSTSE within this framework should clearly indicate the benefits and costs of each stakeholder separately, unless specific circumstances are such that aggregation is necessary. Cases in which aggregation is *desirable* are typically better suited for alternative MODA techniques such as optimization (as discussed in chapter 6.2.3).

## **3. Generate Creative Alternatives**

This step is analogous to the creation of the design space in MSTSE, but is one of the biggest divides between it and the Decision Management Process. Although there is a call for a “useful decomposition of the physical elements of the system” that sounds similar to the identification of key design variables, the purpose it is identified for is to “quickly and accurately communicate differing design features” among fixed alternatives. This is in contrast to the common TSE and MSTSE practice of using the design variables to combinatorically enumerate a complete design space (or decision space, in this language). The example provided is one of a UAV decision between six commercially available options, and in fact the current handbook states that “between 4 and 7 reasonable alternatives” should be considered in a trade study. Though it is possible to perform TSE or MSTSE on a tradespace consisting of only a small number of alternatives, it is generally considered to be more effective at generating insights on larger design spaces. There is no specific example provided of a bottoms-up design problem that lends itself to a larger design space, as opposed to the supplied top-down acquisition problem. The general positioning of the Decision Management Process as a few-alternatives trade study framework has a trickle-down effect on the following steps that makes some recommendations infeasible on a larger set of alternatives.

4. **Assess Alternatives via Deterministic Analysis**
5. **Synthesize Results**

These two steps together form the basis of a standard tradespace exploration, with “assess alternatives” corresponding to creating and then running the evaluative models and “synthesize results” corresponding to the exploration and analysis of the output data, including the use of multi-attribute value functions. Example benefit-cost scatterplots and bar charts are provided that closely match those of MSTSE, though with fewer alternatives. However, two features of these steps are not feasible when scaled up to a larger design space. First, the assignment of performance scores to each alternative is recommended to be done by subject matter expert evaluators *interpreting* data from models or simulations with some documented rationale. This is transparently infeasible for even a relatively small tradespace with hundreds of alternatives, as the time cost of individually interpreting the results would be too great. When considering many alternatives, the models themselves must be validated by the experts rather than the model outputs. Similarly, the notional “home” view of the method (in contrast to the benefit-cost tradespace for TSE) is a “value scorecard” that represents the value in each performance objective for each alternative using a table view. This type of “spreadsheet engineering” is not conducive to summarizing key insights once the number of alternatives grows too large for them to be considered simultaneously and is best used only to compare designs identified as attractive solutions. The “value scorecard” also suffers from being difficult to read if it attempts to represent multiple different value functions of different stakeholders at the same time.

6. **Identify Uncertainty and Conduct Probabilistic Analysis**
7. **Assess Impact of Uncertainty**

The importance of uncertainty analysis has already been discussed considerably in this research, so it will not be repeated here other than to say that MSTSE can incorporate uncertainty analysis in a number of ways and this research prefers the use of Epoch-Era Analysis. The Decision Management Process, like most other systems engineering processes, has an explicit step targeting the analysis of uncertainty. In particular, it recommends the identification of key uncertainties and the subsequent assessment of their impacts to be done here, after the deterministic analysis. Types of uncertainty analysis identified in the guide include low-medium-high estimation, Monte Carlo simulation, tornado diagrams, decision trees, and sweeps of value weightings for sensitivity analysis, all of which are potential supplements or alternatives to Epoch-Era Analysis in MSTSE. It is worth noting that the identification of key uncertainties in Epoch-Era Analysis occurs during problem formulation in the first steps of the method, with the understanding that knowing what uncertainties exist can spur creative thinking in the creation of alternatives in the design space in order to withstand that uncertainty. Holding off on this setup until after deterministic analysis can lead to required but unnecessary iteration of the process if a key decision element was forgotten, when it would have been included should uncertainty have been discussed early in the process.

## 8. Improve Alternatives

This step is the initiation of an iterative loop: using the lessons learned from analysis to improve the set of alternatives and then re-analyze. Iteration can be a valuable tool for gradual learning and incorporating insights “on the fly” in order to reach a better final decision and thus it is useful for a process to iterate effectively, but it is costly in both time and effort. The Decision Management Process recommends creatively deploying Value Focused Thinking (Keeney, 1992) and morphological boxes (Zwicky, 1969) to generate new alternatives that improve on the original designs. The importance of this step is driven largely by the smaller design spaces that are generally recommended in the earlier steps. A full-factorial tradespace in MSTSE is essentially the same as a complete morphological box and thus is less likely to require iteration, though iteration is encouraged if additional design variables are revealed during the analysis.

## 9. Communicate Tradeoffs

## 10. Present Recommendation and Implementation Plan

The final two steps of the Decision Management Process are simply the documentation, presentation, and final decision for the problem at hand and the lessons learned through the analysis. The proper communication of relevant tradeoff insights provides rationale for the final decision that can be used to feed forward into the other INCOSE-recommended lifecycle processes included in the Systems Engineering Handbook.

### 8.3.3 Summary

The INCOSE Decision Management Process is an open-ended setup for conducting trade studies, intended specifically for deciding between small numbers of alternatives. To this end, it uses some techniques in both evaluation (expert opinion) and analysis (value scorecards) that can provide high detail for few alternatives but become infeasible or difficult to understand when applied to large numbers of alternatives. With respect to this weakness, MSTSE could best be employed within the process as a guideline for dealing with many thousands of alternatives, making the move from trade studies into tradespace exploration. The Decision Management Process also has very little to say on the topic of multiple stakeholders beyond the need to include all of the relevant stakeholders in the decision framing. No guidance is provided on how to manage a set of potentially diverse interests and many recommendations imply that the stakeholders should be aggregated. For example, the Systems Engineering Handbook says that “to achieve objectivity, consensus should be reached [on value function weights] before the alternative solutions have been identified.” While the individual attributes are ideally evaluated with an objective *measure*, the weights themselves are intended to capture subjectivity: the preferences of the participating stakeholder. There are no objectively correct weights to a value function, particularly to a value function that attempts to reconcile multiple preferences. An application of MSTSE within this trade study framework would be best served by relying on

MSTSE for all multi-stakeholder analysis and using the INCOSE process to structure the passing of information from step to step, including appropriate iteration.

## 8.4 Responsive Systems Comparison

The Responsive Systems Comparison Method (RSC) is a systems engineering process from the Systems Engineering Advancement Research Initiative, originally developed between 2008 and 2010. At a very high level, RSC can be viewed as a wrapper for the integration of Multi-Attribute Tradespace Exploration (MATE) and Epoch-Era Analysis (EEA) into a combined tradespace and uncertainty framework. In its original form, RSC is a seven-step process progressing from early problem formulation through modeling and analysis. Later research modified and restructured the analysis steps as additional tradespace analysis techniques were developed and refined, but the steps remain functionally similar. This section will focus on the original RSC formulation as presented in the first two reports, as these sources provide the most detail on the process as a whole and the early steps in particular (Rhodes, Ross, and Hastings, 2008, 2010).

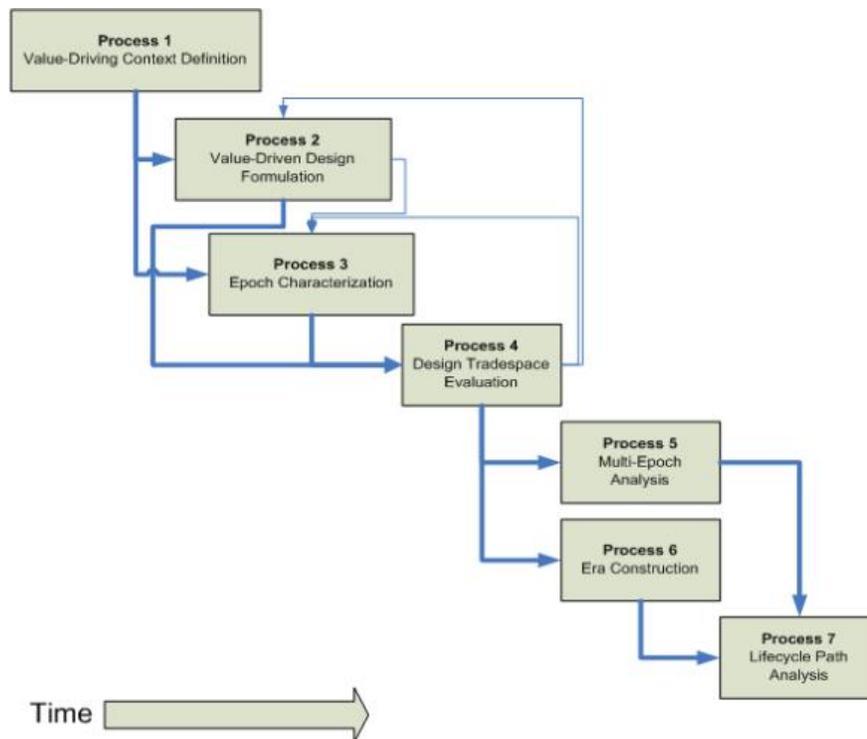


Figure 8-6: RSC flowchart (Rhodes, Ross, and Hastings, 2008)

A short description of each step of RSC is as follows:

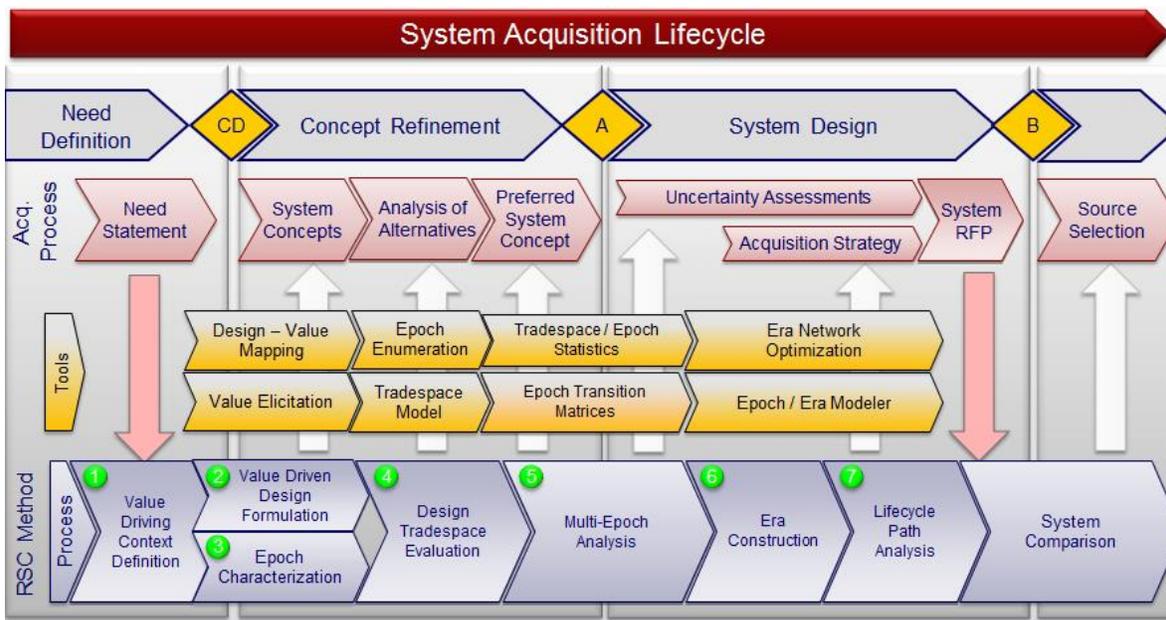
1. **Value-Driving Context Definition** – Identify and describe the high-level problem and associated goals of the system.

2. **Value-Driven Design Formulation** – Collect the need statements of all stakeholders and use them to develop a design space of solution concepts that target those needs.
3. **Epoch Characterization** – Identify contextual and needs-based uncertainties that affect the system and parameterize them for future analysis.
4. **Design Tradespace Evaluation** – Model the system and context in order to calculate and analyze the value of alternatives in the tradespace.
5. **Multi-Epoch Analysis** – Look across the range of uncertainty in order to find alternatives that are robust to changes in context and needs.
6. **Era Construction** – Construct plausible timelines of varying context and needs.
7. **Lifecycle Path Analysis** – Identify system designs and management strategies that deliver value across targeted eras.

This analysis of RSC will not go into detail on its recommendations for conducting tradespace exploration, which understandably form the bulk of the description of the seven steps. MATE is the largest influence on the tradespace analysis prescribed by this research, and so it would be highly repetitive to discuss the RSC tradespace formulation (and also slightly disingenuous to imply that similarities between RSC and MSTSE validate each other in any way). Instead, discussion will focus only on the treatment of the multi-stakeholder problem with RSC, which is stated to “assist and accelerate multi-stakeholder negotiations around cooperatively-developed complex systems”. Specifically of interest are the ways in which stakeholders are encouraged to participate in RSC and the prescribed analysis techniques for considering multiple value propositions. Note that Year Two of the development of RSC featured an early attempt at MSTSE as a spinoff of the development of real-time interactive TSE software, using a role-playing session with three decision makers (Rhodes, Ross and Hastings, 2010; Ross et al., 2010a); this was addressed in the literature review and will not be reprised here as it did not explicitly follow the steps of RSC.

#### **8.4.1 Stakeholder Participation**

RSC is noted as being developed “from the frame of reference of a government program manager” at the System Program Office (SPO) level, responsible for performing early conceptual design. Interestingly, RSC is scoped through Milestone B in the government acquisition timeline, making it nominally intended to be more in-depth than the pre-Milestone A Analysis of Alternatives process and plausible to carry the design all the way up to request for proposals and contracting.



**Figure 8-7: RSC vs. the System Acquisition Lifecycle (Rhodes, Ross, and Hastings, 2010)**

RSC is positioned as a technique for the engineers and analysts in the SPO, managed by an “RSC practitioner”, to do the necessary due diligence for acquisition and report those findings to the relevant stakeholders, rather than a primarily stakeholder-driven activity. This makes RSC closer to informal MSTSE [chapter 7.4] than the more-desirable version with full participation and negotiation of stakeholders. That said, RSC is also intentionally scalable to the “desires of decision makers”, which could entail participation in an active MSTSE negotiation session in any or all of the analysis steps (4-7)<sup>18</sup>.

Most stakeholder-related discussion is contained within the first two steps. Value-Driving Context Definition includes the identification of all key stakeholders in the system (including both decision makers and those impacted by the decision) and the creation of an enterprise stakeholder value network: an illustration of the connections between the involved stakeholders. The value network, as an extension to the brainstormed list of stakeholders featuring additional relationship information, has considerable similarities to the diagrammatic tasks of the CLIOS process on the institutional sphere (Sussman et al., 2009), and can be expanded to include more detailed value flows between stakeholders akin to the Stakeholder Value Network work of Cameron, Feng, et al. (Cameron, 2007; Feng et al., 2010, 2012). The RSC practitioner is advised to consider how each stakeholder will respond to uncertainty and

<sup>18</sup> More recent discussion of RSC and RSC-adjacent work (Schaffner et al., 2014) generally indicates an increased importance of stakeholder involvement than is reflected in the original documentation. This is perhaps a combination of both practicality in the original messaging (RSC would be easier to adopt and implement with less stakeholder participation, just like informal MSTSE) and the general trend amongst the systems engineering community that has recently seen more interest in working closely with stakeholders.

how those responses would affect the others, an activity that would certainly be improved through direct stakeholder participation. One concern is that the stakeholder value network is “centered” on the “main” system developer with others that “assist” around them, which may unduly emphasize a single party if multiple stakeholders have decision making authority on the complete system.

After identifying the stakeholders in step 1, Value-Driven Design Formulation involves interviewing the stakeholders. These interviews are intended to yield information necessary to develop the modeling artifacts of RSC, including the design space, relevant uncertainties, and desired attributes or utility functions. It is implied that these interviews are meant to be performed individually, in order to prevent conflation of the different interests at stake. A Joint Fact Finding effort, particularly for the design and context spaces, would be more aligned with the goals of MSTSE and individual sessions for the development of value models can be conducted on the side. JFF is also of interest with respect to the “record fixed and assumed system parameters” activity, as these assumptions may differ between stakeholders at first and require resolution (or assignment as a design/context variable if no consensus can be reached).

Upon reaching the Era Construction step, RSC recommends the use of domain experts to assemble a set of diverse, plausible, and representative potential eras for analysis. Era construction and analysis has been the subject of considerable research since the original RSC was reported (Schaffner et al., 2013; Schaffner, 2014), particularly in the realms of computational, many-era analysis which is relatively immune to human subjectivity. However, there is still a role for narrative-based exploration of particularly interesting or likely eras in order to focus discussion, and the stakeholders should be drawn in to this task in addition to domain experts. Final decisions are unlikely to be made until each stakeholder is satisfied with the explored eras, and including them in the initial construction both streamlines this process and provides an opportunity for domain experts to correct any held misconceptions about the interaction of context variables as they change over time.

Finally, Lifecycle Path Analysis requires one input that is not provided by the outputs of the previous steps: the desired “strategic objectives” for managing the system over time. These inform what the system will do when epochs change, such as “always maximize utility” or “improve slowly over time”. Little detail is provided on how to define these strategic objectives, but the most appropriate objectives will come from the stakeholders themselves, particularly the one(s) with managerial control over the resulting system, and thus they should be consulted actively at this step. The deployment of these objectives and “full program-level utility” on which to judge the successful implementation of these objectives treads the line of aggregating out key multi-stakeholder details from the problem, and relies on the SPO being framed as an objective assessor and balancer of the value of all stakeholders, similar to the current perception of the role of systems engineering that was a frequently raised topic in our practitioner interviews [chapter 5.6]. If no single stakeholder has managerial control, the determination of the

appropriate strategy is a negotiation in itself, not something that could be imposed by the SPO. This particular issue can be explored through analysis of the impacts of different strategies across the era space using simulation, similar to the Valuation Approach for Strategic Changeability (VASC; Fitzgerald, Ross, and Rhodes, 2012).

#### **8.4.2 Multi-stakeholder Analysis**

Regardless of the level of participation of decision makers, the RSC guidebook reminds analysts that all decision makers should be kept distinct, avoiding the negative consequences of premature aggregation. This caveat is presented in the description of MATE and the elicitation of stakeholder preferences, along with other reminders scattered throughout. However, some analysis recommendations within RSC skirt the multi-stakeholder problem by abstracting the stakeholders. For example, step 1 identifies the creation of a hypothetical “benevolent dictator” decision maker as a best practice capable of simplifying the later analysis, presumably by avoiding the conflict between multiple value propositions through the use of a hybrid value model. This recommendation does state that the benevolent dictator can be withdrawn later in order to extract more detail, but this could drive more losses-domain framing if the dictator’s solution is established as a reference point and must be discarded due to competing interests and a lack of jointly optimal solutions. At best, the benevolent dictator should remain strictly a thought experiment for devising system concepts and not used for analysis unless stakeholder aggregation *is* feasible for the given case.

No specific mention of BATNAs is included in the description of RSC. The closest is an aside on the nature of existing systems and infrastructure constraining the feasible decision space, but this is not related back into any form of analysis. The natural place to include the determination of BATNAs is in the “generate system concepts” activity of step 2, which asks practitioners to brainstorm all potential *types* of systems that could plausibly meet the identified needs of each stakeholder. This has considerable overlap with the brainstorming of viable negotiation alternatives (for example, a given simpler concept may not require the agreement of all stakeholders to pursue) and thus can be combined with BATNA formulation if the stakeholders also participate.

Design Value Mapping (DVM) is a recommended early analysis step in RSC that relates the anticipated impact of each design variable on each attribute of interest using a Quality Functional Deployment (QFD) type of description, typically on a 0-1-3-9 scale. The DVM can be used to iteratively improve both the design space, by adding new variables where necessary and trimming out variables with little impact, and the value space, by identifying attributes that are not actually affected by the system. Although not specifically mentioned in the RSC handbook, in some DVMs the attributes are subdivided into groups corresponding to each decision maker for problems with multiple stakeholders (Ross, 2006), and this should absolutely be performed. Categorizing the attributes in this way can illustrate the relative impact the system has on different stakeholders, which may influence their relative desire to complete an agreement

versus walking away. It can also serve to prevent a situation in which, for example, two attributes are to be trimmed out due to low relevance in the DVM, but those two attributes represent two of the three value sources for one stakeholder, severely hampering the representation of complexity in their value model.

RSC uses a number of tradespace metrics to quantify the value of a given system and its ability performance, such as Pareto Trace and Filtered Outdegree. These metrics can be deployed in multi-stakeholder analysis as long as the participants understand that each metric should be calculated for each stakeholder separately and these are not designed to identify good alternatives *across* stakeholders, which is not specifically addressed in the report. The identification of multi-stakeholder “compromise” designs is performed by calculating the multi-dimensional Pareto front across *all* of the benefit and cost functions of the stakeholders, and finding the alternatives that are present in that set but *not* in the individual benefit-cost Pareto set of any stakeholder. These designs are definitely interesting and should be considered, but their presentation in RSC is slightly problematic.

First and most importantly, as they are defined here, the “compromise” designs are not necessarily balanced between stakeholders: many alternatives in the “compromise” set will be highly unlikely agreements. For two stakeholders, a “compromise” design may lie closer to one’s Pareto front than the others, only to be discovered by looking at the individual tradespaces, as shown in Figure 8-8. These designs are noticeably “unfair” and the worse-off stakeholder will be unlikely to accept them as potential agreements. This problem is compounded when there are more than two stakeholders, as the “compromise” designs can favor any *subset* of stakeholders and therefore an even smaller fraction of the “compromise” set will be “fair”. It is worth keeping in mind that as the number of dimensions considered increases, more and more design points become “compromise” Pareto efficient because it becomes more difficult to be strictly dominated; thus with many stakeholders with different value models, this set of “compromise” alternatives can increase into a considerable fraction of the entire tradespace, limiting its effectiveness as a screening tool. This method can still be used to prompt discussion or test out a few group-efficient designs, but is probably best deployed on a pairwise or small-group basis between stakeholders with relatively similar value propositions, making it a natural choice when trying to build coalitions within the larger group of stakeholders.

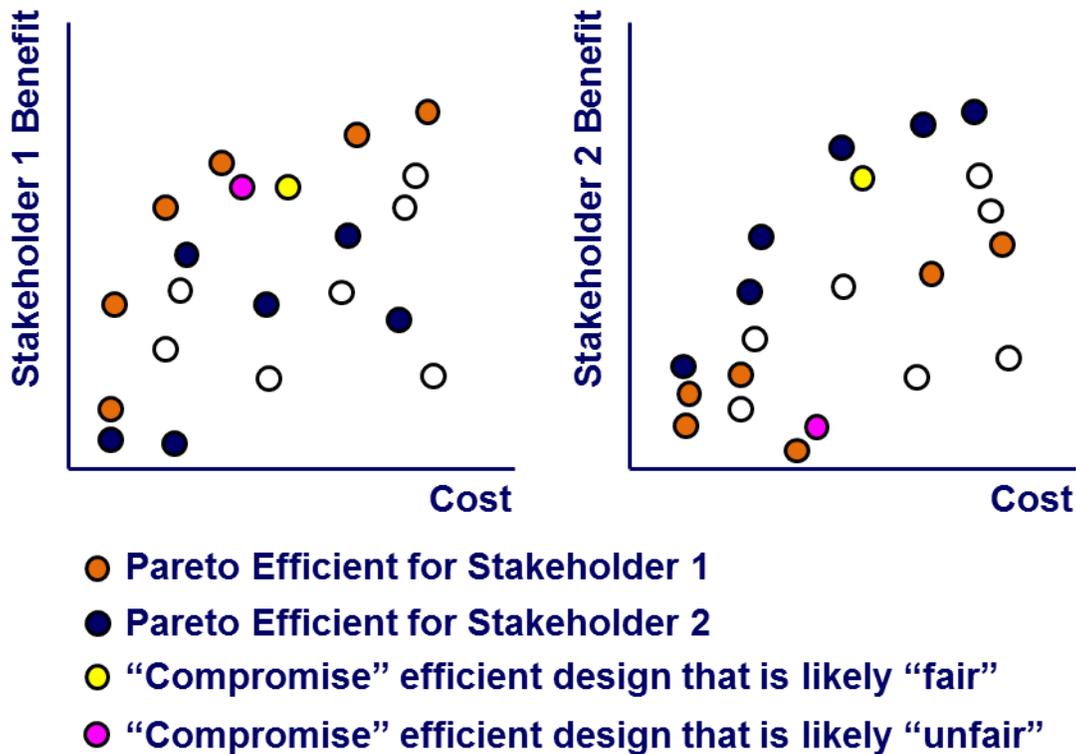


Figure 8-8: Example tradespaces where some “compromises” are "fair" and some are not<sup>19</sup>

Additionally, the use of the word “compromise” itself is misleading when describing these designs. It likely intends to capture the casual use of the word to mean *middle* or *between extremes* which is correct in that the designs are between the extremes of individual optimality. However, “compromise” has a much more specific primary definition, particularly when applied within the context of a negotiation, meaning “a settlement of differences by mutual concessions”. *Concessions* is the critical term here: a compromise is reached when both sides give something up for the sake of reaching an agreement. However, the stakeholders are not able to *concede* from their individual Pareto fronts as they *do not have* those alternatives, they are merely *potential* agreements. Value can only be *conceded* relative to the BATNA, an actual benchmark

<sup>19</sup> This concept is difficult to visualize (as are most >2-dimensional concepts), but these plots can be read as follows -- There are two stakeholders who have different benefit functions but use the same cost metric (the points in each plot are shifted only vertically). They share no Pareto efficient alternatives. The magenta and yellow circles both satisfy the RSC definition of “compromise” designs: neither is efficient for either stakeholder, but they are efficient across all three dimensions – we can see this here because the only design that dominates them in Stakeholder 1’s tradespace (fourth orange circle from the left) does not dominate them in Stakeholder 2’s tradespace. Note that the yellow circle is nearly efficient for both stakeholders: these are the “compromise” designs that this technique wants to find. However, the magenta circle is good for Stakeholder 1 and very bad for Stakeholder 2 and yet is also picked up by this method. The “compromise” Pareto set, though capable of finding “good” designs that are “between” the stakeholders’ individual optima, captures all the trades between them – including alternatives that are obviously “unfair” to one or more stakeholders.

of held value. The so-called “compromise” design set often includes many alternatives in Quadrants 1, 2, and 3 of the BATNA-centered negotiation tradespace, which represent possible *gains* in value (which, it is worth reiterating, is why these designs are interesting to each stakeholder in the first place). Referring to them as “compromises” reinforces the idea that the individual Pareto front is the reference point against which other alternatives are judged, instead of the BATNA, and that these designs are *losses* relative to it. It is possible that the terminology of “compromise” efficient alternatives is too engrained at this point to be changed, but a phrasing such as “cross-stakeholder efficient” or “group Pareto set” would be more accurate and could also support positive gains-frame evaluation of these designs by the participating stakeholders.

### 8.4.3 RSC Feedback

Part of the RSC research endeavor was the specific targeting of the method to the needs of the government, including that of the balancing of multiple stakeholders. As part of that research effort, a questionnaire was used to extract detail on the current state of practice. Some interesting feedback pertinent to multiple stakeholders and the use of tradespace exploration is presented here, much of which echoes comments received from our practitioner interviews:

- Decisions are framed through the presence of “internal” (engineering) and “external” (e.g., strategic, policy) decision makers. Communication between these types of decision makers is perceived as a significant challenge, with it heavily implied that “internal” decision makers are best qualified to make decisions effectively but are forced to reduce the sophistication of their analysis to participate. This feedback corroborates the similar insight from our own practitioner interviews. MSTSE negotiation sessions are positioned well to help bridge this gap by allowing for customizable complexity in representing/visualizing the relevant data. Face-to-face negotiation on the tradespace would also serve to alleviate frustration over having to communicate through “limited” means such as PowerPoint.
- Internal decision makers, who are actually engineers and analysts, are responsible for satisfying all of the external decision makers. Again, this ties into the perception that systems engineers are responsible for “balancing” the needs of strategic decision makers.
- Decision makers work in parallel rather than in sequence. Coordinating parallel but separate activities can be challenging, and would likely be simplified if at least some group decision making was performed through face-to-face negotiation.
- Alternatives are evaluated in small groups and sent off to external decision makers for consideration, but this reinforces form-dependent preferences in the guise of decision makers “having a specific system they would like to have”. TSE is generally considered to have superior connection to the principle of value-driven design than point alternative studies. Having stakeholders interact within MSTSE could help loosen form-dependent preferences by clearly illustrating value trades between different concepts.

- The models available for use in this phase of the acquisition timeline are considered to be of medium fidelity and more precise than accurate. These are ideal conditions for TSE, which is computationally challenging at high fidelity but often requires more detail than low fidelity and which values precision more than accuracy in order to maintain correct ordering of alternatives.
- Delays are often caused by external decision makers raising a “discrepancy” with the analysis or recommendation of the internal decision makers, leading to costly iteration. This is where Joint Fact Finding can serve to prevent post-hoc challenges to the validity of insights.

#### **8.4.4 Summary**

RSC, by virtue of relying on the same MATE foundations as MSTSE, is naturally well-suited to incorporate MSTSE features. Additional stakeholder involvement during the early problem formulation steps, including through Joint Fact Finding and BATNA definition, is the only major addition necessary to support interactive negotiation sessions on the tradespace during later steps. The MSTSE negotiation itself would likely span steps 4-7 if enough time was available, as single-epoch, multi-epoch, and era analyses are all potential sources of relevant decision making information on which to negotiate. The original scoping of RSC was set to the SPO level, resulting in some artifacts of both analysis and language that become unnecessary and/or inappropriate with more direct and structured participation from decision makers, including hypothetical “benevolent dictators” and not-actually-“compromise” designs.

### **8.5 Discussion**

This chapter has sought to provide some guidelines for the application of MSTSE within four SE methods, highlighting the strengths and weaknesses of each method with respect to enabling both tradespace exploration and multi-stakeholder decision making. Though each method was approached separately, it is clear that there are considerable similarities between them. Referring back to our generic outline of a decision making process from chapter 7, these methods all seek to coordinate problem formulation activities to “frame” the problem correctly and then proceed to modeling and exploration, which makes MSTSE a natural fit within them. Common pitfalls include preemptively aggregating stakeholder value models before exploration or preselecting a small number of alternatives for exploration, limiting the ability to create integrative, mutually beneficial trades.

Undoubtedly, these recommendations are driven in part by the minimal level of expected stakeholder involvement under the current paradigm – a practical consideration for methods seeking wider adoption. Requiring stakeholder participation is a barrier to entry for many applications (similar to requirements for specific engineering expertise when using other advanced techniques) and can reduce the feasibility of adopting such a method/technique when those resources are constrained. The logic behind the simplifications we saw in the SE methods

is justifiable. *With* stakeholder participation it can be reasonable to expect any number of stakeholders to *each* actively manage their *own* value models. However, *without* stakeholder participation, the sheer dimensionality of many value models and many alternatives drives engineers to pragmatically simplify by reducing the space under consideration by aggregating stakeholders and/or limiting the number of alternatives, allowing for use on a wider range of problems. Thus, the potential value-added by MSTSE to these methods is predominantly driven by the structured way to keep those interests separate and leverage a large design space, even when decision maker access is restricted and informal MSTSE must be used. The availability of a technique such as MSTSE may allow future SE methods to recommend large design spaces and separate value models without limiting their scope to problems in which all stakeholders are willing to participate.

The methods and frameworks for managing SE projects discussed here have, of course, been used in the past. We will proceed with the MSTSE analysis of two case studies that have previously been structured and analyzed using RSC and the CLIOS Process. New insights derived from MSTSE will demonstrate the types of knowledge at the inter-stakeholder level that can be missed without an approach explicitly targeted at multi-stakeholder problems.

## 9 Case Study – Satellite Radar

This chapter and the following chapter will discuss two case studies as demonstrations of (1) the use of this research’s recommendations for visualizing and conducting multi-stakeholder tradespace exploration and (2) the types of insights that can be gained when analyzing a case in this way, particularly those insights that are “hidden” when using traditional, single-stakeholder perspective analysis techniques. The realistic, concrete foundation of the cases should help to clarify the ideas and conclusions of the prior sections of this document in addition to supporting the validation of MSTSE as a feasible approach to the analysis of multi-stakeholder engineering design problems. As these cases will not include any real or staged simulation of negotiation proceedings, they will follow the recommendations of the informal mode of MSTSE described in chapter 7.4. The two case studies were selected to display a few desired criteria:

- Fundamentally multi-stakeholder (likely or certainly cannot be solved/built by one party)
- Non-trivially complex (tradespace exploration does not represent obviously unnecessary effort)
- Data and/or reports available from prior analyses (allows for comparison against previous insights)
- Structurally complementary (cases demonstrate applicability of research to differing problem types according to MSTSE structure, i.e. Chapter 6)

This chapter covers Satellite Radar: a case study on the potential for a constellation of Earth-facing satellites to provide 24/7 all-weather surveillance of the Earth surface. This system would have considerable value for defense- and intelligence-focused agencies of the United States (or another) government, as there is no other reasonable means of performing this task. The prohibitively high cost makes it very unlikely that the system could be funded without the approval of shared-use between those agencies. Satellite Radar has been the subject of prior development programs cancelled during multi-stakeholder design, providing ample evidence of the possible benefit to be gained through MSTSE.

### **The Satellite Radar case will be approached with the following key objectives:**

1. Demonstrate the use of MSTSE analysis applied to a technically complex system with many potential alternatives and relatively few stakeholders
2. Revisit prior TSE studies and stakeholders in order to capture changes in insight or outlook prompted by the new techniques
3. Identify the potential for -ilities to create opportunities for additional value by allowing transitions between designs that favor individual stakeholders

### **9.1 History**

The potential for an on-orbit, earth-facing radar system to deliver significant strategically valuable information to a variety of stakeholders interested in surveillance and intelligence missions has been a topic of interest since as early as the 1950s with the United States launch of

the Discoverer and, later, Lacrosse imaging satellites (Braganca, 2011). The high vantage point of satellites enables them to provide a wide area of line-of-sight coverage, avoid terrain masking of mountainous regions, and to remain largely protected from enemy attack (DeLap and Suhr, 1996; Corcoran, 2000). As technology matured, stakeholders discovered that a satellite radar system would be capable of performing multiple mission types, including imaging (via SAR - synthetic aperture radar), tracking (GMTI - ground-moving target identification), terrain mapping, and ocean surveillance (Congressional Budget Office, 2007). However, the large power required for radar to effectively spot a target from orbit as well as the constellation density necessary to ensure near-continuous visibility of that target result in a prohibitively expensive system – and one with limited opportunities for future upgrades due to its remote location (Richards, 2009). This extreme cost essentially requires the participation and agreement of multiple government agencies (stakeholders) in order to secure the necessary funding.

The United States government has attempted to field a modern satellite radar system with near-complete coverage of the earth on multiple occasions (Richards, 2009). Each program involved cooperation between multiple stakeholders and each ultimately resulted in cancellation. These programs include:

- Discoverer II (1998-2000). Cancelled due to cost and schedule uncertainty
- Space-Based Radar (2001-2005). Cancelled due to significant funding cut by Congress
- Space Radar (2005-2008). Cancelled due to lack of affordability

Looking slightly deeper into the formation and subsequent termination of each of these programs reveals the key role of multi-stakeholder interaction in the history of satellite radar. The Discoverer II cost and schedule uncertainty was considered largely a product of an inefficiently-handled collaboration between the Air Force (USAF), National Reconnaissance Office (NRO), and Defense Advanced Research Projects Agency (DARPA) on the program (Federation of American Scientists, 2000). As a result, the cancellation of Discoverer II prompted a more explicit division of responsibilities, including the NRO being funded specifically to develop enabling technology and the USAF being given executive power over all military space activities – including the new Space-Based Radar program (Tirpak, 2002). This attempt to reduce the amount of necessary collaboration backfired, as the decision to cut funding from Space-Based Radar was likely driven in part by a US Government Accountability Office report (2004) that cited a gap in the requirements of the defense and intelligence communities. As the USAF designed a satellite constellation to support military needs, the NRO was designing support technology for intelligence, and it was unclear if their combination would result in a desirable system for either party. Space Radar was then an attempt to reestablish a joint development process, with a specific call to make the program more affordable by integrating and reconciling the requirements of the defense and intelligence stakeholders (Los Angeles Air Force Base, 2011). Ultimately, however, these stakeholders were unable to do so and the project

was shelved due to the high cost of supporting the needs of both communities (*Satellite Today*, 2008).

Overall, the history of satellite radar in the United States is one of consistent failure of government agencies to work together to find a workable, mutually satisfactory system for a reasonable cost. This failure was nominally driven by the inability to reconcile the disparate requirements of the different mission types desired by the stakeholders (Braganca, 2011). When viewed as a strictly technical challenge, as it has been almost exclusively in the past, this may seem insurmountable: if the necessary constellation size and power to accomplish multiple missions is prohibitive, there may be no solution until technology improves (for example, with the discovery of lighter materials). However, this belies the possibility of the so-called “80% solution” for each mission, which though not individually optimal, may still be superior to the currently realized alternative: nothing at all. If such a solution exists, then the failure of the United States to build a satellite radar system is likely partially due to ineffective negotiation or group decision making by the participating agencies. It is possible that MSTSE, with its focus on finding solutions that are attractive across different value functions rather than attempting to consolidate value into a single metric or set of requirements, could help resolve this challenge.

## **9.2 Original Tradespace Formulation**

Part of the 2004 GAO report’s recommendations for improving the outlook for Space-Based Radar included the “strengthening of systems engineering applications” for proposing alternative design concepts. In the years after, satellite radar was a recurring example used in the development of the Responsive Systems Comparison (RSC) method at the MIT Systems Engineering Advancement Research Initiative, designed to help satisfy the interests of stakeholders across various mission contexts (Ross et al., 2008; Ross et al., 2009; Richards et al., 2009). In chapter 8.4, we covered the RSC framework and its blend of tradespace exploration and Epoch-Era Analysis and how it succeeds and fails in supporting multi-stakeholder analysis (Ross and Rhodes, 2008). Specifically of interest to this research is the offshoot of the RSC development outlining the beginnings of multi-stakeholder tradespace exploration, which also used a hypothetical “Satellite Radar System” as an example, though not intended to be representative of a specific real program (Ross et al., 2010a).

Because this case study has already received a full tradespace formulation, we can easily reuse it and compare the insights achieved from the original negotiation session to those that are accessible using the MSTSE recommendations outlined in this research on the same dataset<sup>20</sup>.

---

<sup>20</sup> Note that this will describe the dataset used *specifically* in (Ross et al., 2010a) as well as (Rhodes et al., 2010). As RSC and the Satellite Radar System case were developed in parallel, there are minor differences in the tradespace formulation between publications over time, particularly in the enumerations of certain design and context variables as well as the value models under consideration. (Rhodes et al., 2010) has a complete discussion of the Satellite Radar System model; the reader is referred to that source if they desire more detailed information.

This tradespace formulation supported a three-stakeholder negotiation in which each stakeholder had an informal value model elicitation and participated in a short tradespace exploration session together. The following sections will provide an overview of that tradespace to give the reader the necessary background to understand the insights and conclusions drawn from the data. More detail on the rationale behind the construction of these tradespaces can be found in the original sources.

### 9.2.1 Context Space

The first step of RSC is Value-Driving Context Definition: the identification of those parameters outside the control of the system designers that can impact the value delivery of the system. These contexts, when enumerated, form the basis of the variety of epochs to be analyzed through Epoch-Era Analysis within RSC, allowing designs to be evaluated across a range of potential usage scenarios. Table 9-1 lists the context variables and their enumerated levels for this case.

**Table 9-1: Satellite Radar context space description**

Context Variables	Levels	Enumerated Range	Units/Notes
Radar Technology	2	1, 2	1=TRL 9 only (low), 2=TRL 6 and above (med)
Communication Infrastructure Level	3	1, 2, 3	1=low: no backbone+AFSCN 2=mid: WGS+AFSCN 3=high: TSAT+AFSCN
Operation Plan	9	[1:9]	Nine different target sets
AISR Assets	3	None, Small, Large	Presence of Airborne, Intelligence, Surveillance, and Reconnaissance assets in some locations
Threat Environment	2	None, Jamming	
Utility Expectations	3	[1:3]	Three different utility functions

The full-factorial enumeration of this context space leads to 972 epochs in the dataset, but note that the “utility expectations” variable corresponds to a set of value models not used in the negotiation, so the effective size for the negotiation was only 324 epochs. The other variables correspond to the following parameters:

- **Radar technology.** Sets radar constants based on Technology Readiness Level (TRL). The settings are for a “low” level, counting only existing deployable technology at TRL 9, or a “medium” level, counting technologies down to TRL 6.
- **Communication infrastructure level.** Corresponds to the ability to communicate across other satellite “backbone” networks, reducing latency and power costs. WGS, the Wideband Global SATCOM, is a communication satellite network jointly operated by the United States and allied countries (Air Force Space Command, 2015). TSAT, the

Transformational Satellite System, was a proposed high-capacity, secure communications satellite network for the United States' defense and intelligence purposes. It was under development roughly concurrently with the Space Radar program but was also cancelled in 2009 before reaching the contracting phase as a part of a widespread Department of Defense budget adjustment (Gates, 2009)<sup>21</sup>. All contexts assume access to the Air Force Satellite Control Network (AFSCN) for tracking the position of the satellites.

- **AISR assets.** Corresponds to the existence of local Airborne Intelligence, Surveillance, and Reconnaissance assets in certain geographic locations for closer observation of targets by the complete surveillance system-of-systems.
- **Operations plan.** Specifies the targets of interests. Targets vary in latitude/longitude location, size (radar cross-section, smaller targets are harder to see), and speed (slower targets are harder to track). The enumeration of this variable was a random sampling of the complete space of these target variables.
- **Threat environment.** Specifies the existence of enemy jamming. If jamming is active, the signal strength of the satellites is effectively lowered by dissipation, nominally increasing the power requirements for communication.

### 9.2.2 Design Space

The second step of RSC is Value-Driven Design Formulation: the identification of those parameters *within* the control of the system designers that will define the system and determine its performance and cost. The most common goal of tradespace exploration is the selection of an alternative – a specific combination of design variable levels – with which to proceed into detailed design. Table 9-2 lists the design variables and their enumerated levels for this case, with additional detail on the Walker constellation orbit parameters in Table 9-3. This set of design variables was selected and curated by the RSC team through an iterative verification process with Design-Value Mapping, ensuring each impacted performance in at least one attribute and all attributes were covered (Richards et al., 2009).

---

<sup>21</sup> Because TSAT will no longer be built, using the TSAT context variable setting is strictly hypothetical: if TSAT (or an equivalently powerful communications network) *did* exist, how would satellite radar's performance be affected?

**Table 9-2: Satellite Radar design space description**

Design Variables	Levels	Enumerated Range	Units/Notes
Walker ID Number	8	[1:8]	see Walker Lookup Table
Orbit Altitude	3	[800, 1200, 1500] *10 <sup>3</sup>	m
Antenna Area	3	10, 40, 100	m <sup>2</sup>
Peak Transmit Power	3	1.5, 10, 20	kW
Radar Bandwidth	3	[0.5, 1, 2] *10 <sup>9</sup>	Hz
Comm Downlink	2	Backbone, Ground only	
Tactical Downlink	2	Yes, No	
Maneuver Package	3	Base, 2x Base, 4x Base	
Constellation Option	3	1, 2, 3	1=none, 2=long lead, 3=spares

**Table 9-3: Walker constellation lookup table**

Walker ID #	Inclination (deg)	# Satellites	# Planes
1	53	5	5
2	53	10	10
3	53	10	5
4	53	20	5
5	67	5	5
6	67	10	10
7	67	10	5
8	67	20	5

Most of the design variables are self-explanatory satellite parameters defining the orbit, altitude, antenna, power, and bandwidth of the system. The Comm Downlink specifies the ability to use the communications backbone (if available in the context) rather than communicating directly to a ground station. The Tactical Downlink allows communication directly to a user without using either a ground station or backbone. Maneuverability specifies the amount of monopropellant included in the satellite to enable thrusters to move the satellite. Finally, the Constellation Options are options to spend more money to procure additional satellite parts or complete spares, reducing the time and cost to build more satellites in the future.

The full-factorial enumeration of this design space creates 23,328 alternatives for consideration in the negotiation. Note that additional design variables for radar frequency (10 GHz), antenna type (AESA), and a “tugable” attachment point (included) were omitted from the table, as the final enumeration included only one level for each.

### 9.2.3 Evaluative Models (Performance and Cost)

The satellite radar alternatives were evaluated using a MATLAB script developed by the RSC team. This model was low-to-medium fidelity, based on physics and first-principles equations with basic satellite design parameters sourced largely from the third edition of Space Mission Analysis and Design (SMAD; Wertz and Larson, 1999). The goal for the models was to generate “insights for complex problems, not to provide accurate design details” (Ross et al., 2010a), in keeping with the learning goals of tradespace exploration. The code itself was divided into a number of different modules, evaluated consecutively. A list of the modules in evaluation order and a short description of each is included here, in order to provide perspective on how the performance attributes were calculated:

- **Constants.** Loads constant parameters for the evaluation. Includes checks on the current epoch for parameters which are affected by the operational context, such as the target set.
- **Orbit.** Calculates orbital period, satellite velocity, and eclipse times.
- **Radar.** Calculates a number of radar parameters, including swath, minimum detectable velocity, field of regard, and resolution.
- **Constellation.** Projects the field of vision for each satellite in the constellation onto latitude/longitude coordinates and propagates them through two days of orbit. Used to find the timing windows for which the system can and cannot see the targets of interest, then calculate attributes such as revisit interval, track life, and target acquisition time.
- **Processor, Communications, and Bus.** Based on the observation patterns, calculates necessary power and mass requirements for the on-board electronics and the resulting data latency and necessary launch vehicles.
- **Cost and Schedule.** Parametric models estimating recurring and non-recurring engineering costs as well as the likely timeline for the design, build, and test phases before launch.

### 9.2.4 Value Models

Three stakeholders participated in the negotiation, with aliases E, N, and R. The value models for the three stakeholders were multi-attribute utility functions informally elicited via a written questionnaire. The questionnaire allowed the stakeholders to choose the attributes that delivered value for their intended satellite radar mission from a list of 12 outputs of the evaluative models, categorized by their applicability to the nominal imaging (SAR) or tracking (GMTI) missions<sup>22</sup>. This list is shown in Table 9-4 with a short definition of each attribute. Then, the stakeholders were allowed to set the utility parameters for each attribute of minimum

---

<sup>22</sup> These attributes and their corresponding single-attribute utility curves were created in early development of RSC based on stakeholder discussion and analysis of the SAR and GMTI concept-of-operations, before the negotiation was conducted. For simplicity, the negotiation participants were not allowed to set the curves (other than the minimum/maximum endpoints) or choose attributes outside the set.

acceptable performance and maximum value threshold (the lower and upper bounds of the single-attribute utility curve, referred to here as the “requirement” and “objective” for each attribute), as well as the swing weight<sup>23</sup>. Note that, as part of a negotiation, the stakeholders were explicitly informed that these values were not final and could be changed during the exploration.

**Table 9-4: Satellite Radar attributes able to be chosen by stakeholders**

Name	Units	Mission Type	Definition
Minimum radar cross-section (RCS)	m <sup>2</sup>	Tracking	Smallest target area able to be detected
Minimum detectable velocity (MDV)	m/s	Tracking	Smallest velocity able to be detected
Number of target boxes	#	Tracking	200x200 km boxes tracked for movement in a single pass
Target acquisition time	min	Tracking	95 <sup>th</sup> percentile longest time gap to see target in a given location
Tracking life	min	Tracking	95 <sup>th</sup> percentile longest time target can be continuously seen
Tracking latency	min	Tracking	Time delay between target sighting and tracking data delivery
Revisit Interval	min	Imaging	Average time between imaging observations
Image Latency	min	Imaging	Time delay between target sighting and imaging data delivery
Image Resolution	m	Imaging	Separation necessary to distinguish two objects in image
Targets per pass	#	Imaging	Number of targets able to be imaged in a single pass
Field of regard	km <sup>2</sup>	Imaging	Area of Earth surface able to be seen at one time
Geolocation accuracy	m	Imaging	Absolute positioning error of imaging data

Table 9-5 lists the questionnaire results for stakeholder E. E’s primary interest was in a time-sensitive imaging mission, hence the emphasis on revisit interval and image latency.

---

<sup>23</sup> Swing weights are used to differentiate the contribution of each attribute to overall value, but are NOT importance weights. Importance weights suggest absolute evaluation: if A has twice the weight as B it is twice as important. Swing weights take into account the range of the attribute under consideration: e.g. if A is my most “important” attribute it may still have a low swing weight because the “importance” is manifested in a strict requirement (difficult to achieve Utility=0) and therefore little value is added above that point (Utility=1 only slightly better than Utility=0). See (Parnell et al., 2014) for discussion.

**Table 9-5: Questionnaire results for Stakeholder E**

<b>Attribute</b>	<b>Unit</b>	<b>Swing Weight</b>	<b>Requirement (Utility=0)</b>	<b>Objective (Utility=1)</b>
Revisit Interval	min	0.5	500	30
Targets Per Pass	#	0.35	1	100,000
Image Latency	min	0.15	720	60

Table 9-6 lists the questionnaire results for stakeholder N. N took a notional program manager role, with an interest spread across the imaging and tracking missions. He wanted to ensure the ability to view any target of interest (via minimum radar cross-section and detectable velocity) with high accuracy and low acquisition time. Note that N is the only stakeholder with swing weights that do not add up to 1, but rather 1.3. This means that N considers the attributes to be substitute goods – high performance in one attribute can effectively mitigate poor performance in another.

**Table 9-6: Questionnaire results for Stakeholder N**

<b>Attribute</b>	<b>Unit</b>	<b>Swing Weight</b>	<b>Requirement (Utility=0)</b>	<b>Objective (Utility=1)</b>
Geo-Location Accuracy	m	0.5	100	3
Target Acquisition Time	min	0.3	120	10
Tracking Latency	min	0.3	10	1
Minimum Radar Cross-Section	m <sup>2</sup>	0.1	1000	0.01
Minimum Detectable Velocity	m/s	0.1	50	5

Table 9-7 lists the questionnaire results for stakeholder R. R's had interest in both imaging and tracking objectives, but chose to limit the complexity of his utility function by using only three attributes. His preferences indicate an emphasis on large quantities of data through high weightings on resolution and number of targets.

**Table 9-7: Questionnaire results for Stakeholder R**

<b>Attribute</b>	<b>Unit</b>	<b>Swing Weight</b>	<b>Requirement (Utility=0)</b>	<b>Objective (Utility=1)</b>
Image Resolution	m	0.45	5	0.01
Number of Target Boxes	#	0.35	1	10
Revisit Interval	min	0.2	500	50

The multi-attribute utility function used to aggregate the attributes for each stakeholder was the Keeney-Raiffa multi-attribute utility function, which is as follows:

$$U(\hat{X}) = \frac{[\prod_{i=1}^n (K \cdot k_i \cdot U_i(X_i) + 1)] - 1}{K}, \text{ where } K = -1 + \prod_{i=1}^n (K \cdot k_i + 1)$$

In this equation,  $U$  is the multi-attribute utility,  $U_i$  is the single-attribute utility of  $X_i$ , the  $i$ -th attribute out of  $n$ , and  $K$  is a normalization constant that ensures that  $U$  will range between 0 and 1. The parameter  $k_i$  is the elicited swing weighting factor for  $i$ -th attribute.

All three stakeholders simply used the estimated program lifecycle cost as their cost metric. This shared total cost would be divided amongst the stakeholders should an agreement be made. The possibility of an uneven split of the money

### ***9.3 Insights and Conclusions of Original Negotiation Session***

This section will quickly overview the results of the RSC-driven negotiation described in Ross et al. (2010a). Procedurally, the three stakeholders were guided in their negotiation by following a five step process given to them by the RSC team:

- 1. Each stakeholder finds “good value” designs.** Work individually and select personal favorite alternatives.
- 2. Stakeholders compare designs and make “easy” compromises.** Compare personal favorites and look for shared designs or naturally easy choices.
- 3. Negotiations and relaxation of attribute requirements.** Explore value models and consider changing them to facilitate negotiation.
- 4. Finding better compromises.** Look for strictly-superior solutions to any acceptable solutions already found.
- 5. Sensitivity to time and change.** A deep-dive into the best designs with Epoch-Era Analysis

The stakeholders selected their individual favorite “good value” designs by manually exploring their benefit-cost tradespace in a single “baseline” context, with low radar technology, no communications backbone or AISR assets, and no jamming. The targets were medium/large targets in Iran and China. As expected, this exploration focused heavily on the Pareto front, with only R selecting any non-Pareto-efficient alternatives as personal favorites. The favorite selections, along with some key design variables and the different utilities (with ‘X’ indicating infeasible according to the utility function), are shown in Table 9-8.

**Table 9-8: Satellite Radar initial favorite designs (adapted from Ross et al., 2010a)**

Stakeholder	Design ID	Design Variables						"E" MAU	"N" MAU	"R" MAU	Cost (B\$)
		Altitude (km)	Num. Sats	Ant. Area (m <sup>2</sup> )	Bandwidth (GHz)	Power (kW)	Tac Comm				
E	3328	1500	10	10	2	10	No	0.72	X	0.62	9.6
	3436	1500	10	40	2	10	No	0.81	X	0.92	14.2
	3760	1500	20	40	2	10	No	0.86	X	0.92	20.8
N	471	800	10	40	1	1.5	Yes	0.66	0.53	0.58	13.3
	3387	1500	10	40	1	1.5	Yes	0.75	0.47	0.58	12.8
	6303	1200	20	40	1	1.5	Yes	0.81	0.57	0.58	18.4
R	424	800	10	10	2	20	No	0.56	X	0.78	12.0
	2692	1500	5	10	2	20	No	0.59	X	0.66	9.1
	2752	1500	5	40	1	10	No	0.60	X	0.80	10.8
	2788	1500	5	40	2	10	No	0.60	X	0.85	10.8
	3436	1500	10	40	2	10	No	0.81	X	0.92	14.2

From this data, the main insight obtained was the apparent differences in generally favored design variable settings between stakeholders. N prefers the lowest power and always takes tactical communications, which neither other stakeholder does. E and R have more similarity between their selections, generally favoring the highest altitude and power settings and even both choosing design number 3436 from their respective Pareto fronts. R is also the only stakeholder to consider 5-satellite constellations.

After sharing favorites with the group, N was found to have all of E and R's favorites marked as unacceptable. With no favorites or Pareto-efficient designs shared between all three stakeholders, the multi-dimensional compromise Pareto set was used to generate suggestions agreeable to each participant<sup>24</sup>. Additional solutions acceptable to the entire group were found by making "easy" compromises: adding bandwidth and tactical communication to all individually-efficient designs (as relatively low-cost ways to add value for E and N, respectively) and also finding a "gold plated" solution with high cost and high benefit for all stakeholders. Stakeholders E and R also picked two designs that they liked without N. This set of group selections is shown in Table 9-9.

---

<sup>24</sup> We discussed the "compromise" Pareto set in chapter 8.4, with the conclusion that this choice of language is detrimental to negotiation by implying that stakeholders must sacrifice something to arrive at these solutions, when the only sacrifices are the hypothetical Pareto efficient alternatives that are not feasible for the group as a whole. This belief still applies – however we will continue to use the term in this chapter in order to prevent confusion while discussing the original source.

**Table 9-9: Satellite Radar group selections; asterisks indicate the method of finding them - \* compromise Pareto set, \*\* “gold plated” (adapted from Ross et al., 2010a)**

Stakeholder	Design ID	Design Variables						“E” MAU	“N” MAU	“R” MAU	Cost (B\$)
		Altitude (km)	Num. Sats	Ant. Area (m <sup>2</sup> )	Bandwidth (GHz)	Power (kW)	Tac Comm				
All	3423	1500	10	40	2	1.5	Yes	0.75	0.47	0.63	12.8
	3435	1500	10	40	2	10	Yes	0.81	0.47	0.92	14.3
	*6027	1200	10	40	2	10	Yes	0.80	0.50	0.92	14.6
	*519	800	10	40	2	10	Yes	0.68	0.53	0.88	15.8
	**6349	1500	20	40	2	10	Yes	0.87	0.57	0.92	22.3
E+R	2627	1500	5	10	2	10	Yes	0.56	X	0.60	8.3
	2787	1500	5	40	2	10	Yes	0.60	X	0.85	10.9

Negotiation of attribute requirements focused on N, whose utility function eliminated the most alternatives. Ultimately, N was willing to temporarily lower his requirements on target acquisition time and minimum detectable velocity to continue exploring, returning some low cost designs to his tradespace and also making E and R’s two designs feasible for himself. For the “better compromises” step, no designs superior to those identified in Table 9-9 were found.

Finally, the Epoch-Era Analysis treatment of the designs of interest was limited to a single scenario era: a sequence of epochs (contexts) extending forward in time after the original “baseline” context. This era consisted of five epochs with incremental changes in context. The changes were as follows:

- “Baseline” – low radar technology, no communications infrastructure, no AISR assets, no jamming, medium/large targets in Iran and China
- Add communications backbone availability
- Target change to small/medium targets in North Korea and China
- Radar technology increase
- Reset targets to baseline

Though most designs stayed on or near the Pareto front across the era for each stakeholder, indicating that they are potentially strong choices for a passively value robust system, the expensive “gold plated” design (6349) saw considerable fluctuation in value. E had 6349 stay on his Pareto front and gain significant utility over the other group selections upon entering the third epoch, with smaller targets. N had 6349 stay near the Pareto front with relatively little movement. Most surprisingly, R saw the utility of 6349 *decrease* in the small-target epochs, to the point of it becoming dominated by a large fraction of the tradespace.

Overall, the insights of the original Satellite Radar tradespace negotiations were very design-focused. Favorite designs were identified, exploring was done in and around individual and compromise Pareto sets, and key designs of interest were viewed across a small range of

contexts. Given the one hour time limit for the session, it would have been unrealistic to expect a large amount of detailed analysis, and a focus on specific alternatives is often an attractive means for stakeholders to simplify and rapidly make decisions – despite its inferiority to more open-ended analysis (Nutt, 1999). The extent of analysis directly targeting multi-stakeholder value in this negotiation was limited to:

- Calculation of a joint Pareto set, which proved to be empty. A jointly Pareto efficient design would be the “easiest” agreement, and the lack of these indicates the difficulty of the problem and the need for negotiation.
- Identification of design variables that have a relatively small effect on cost but create considerable utility for one stakeholder, without harming the others. This includes bandwidth (for E) and tactical communication (for N). These were found indirectly, by considering the lists of favorite designs and inferring the pattern correctly, but could also have been found with little effort by coloring the benefit-cost tradespace by those variables.
- Analysis of N’s aggressive requirements, which were found to eliminate large parts of the tradespace, particularly in the low-cost region. Though no “real” decision was made as a result of these negotiations, the ability to rapidly edit the utility function to evaluate the impact of requirements and explore new areas of the tradespace was identified as a key functionality of tradespace exploration.

It is also worth highlighting the fact that prior, detailed analysis of the Satellite Radar tradespace in a non-negotiation setting had also identified designs 3435 and 6027 as jointly efficient across the tracking and imaging missions (Rhodes et al., 2008). Thus, the participating stakeholders were clearly accessing some of the same insights on “good” designs in the live negotiation session, despite using personalized utility functions.

Stakeholder feedback from the session indicated two main goals for future development: (1) further progress the “intuition” building visualizations and activities that are a part of TSE, and (2) directly “support negotiation” as a necessity for multi-stakeholder design (SEArI, 2010). These goals provide justification for much of this research – the search for better ways to incorporate negotiation and multi-stakeholder insight into the tradespace. Additionally, one specific output request was for an “easier way to identify key issues in agreement and disagreement among [stakeholders]”, which is directly addressed by our new attribute correlation visualization, confirming that our approach is addressing open needs of real stakeholders in this domain.

## **9.4 Negotiation Structure**

As a first step, we need to identify the Best Alternative to a Negotiated Agreement (BATNA) for each stakeholder. The original study did not explicitly develop the BATNAs for each of the stakeholders, and there is not much available information about their individual

resources on which to base a BATNA. In terms of fundamental objectives, this Satellite Radar case is modeling a 24/7 surveillance mission for both imaging and tracking targets. There are currently fielded systems capable of imaging and tracking, so in that sense some kind of “existing system” BATNA seems likely. As previously mentioned, the United States has already placed a number of radar satellites in orbit, including three still-operating Lacrosse satellites. However, regardless of the (classified) radar performance of these satellites, the three-satellite constellation is smaller than any design in the enumerated tradespace and therefore highly unlikely to provide enough coverage to produce acceptable target acquisition times and revisit intervals for true 24/7 surveillance. Since all three stakeholders care about one of those attributes, timeliness is necessary for the intended use case. Radar-carrying airborne alternatives such as AWACS or JSTARS are also unlikely to provide acceptable performance. The low vantage point of AISR assets (relative to satellites) can provide high-quality observations, but with much smaller fields of regard and limited range. This requires military presence to create staging areas, an expensive proposition and one that is also unlikely to be enforceable by a single stakeholder. Additionally, staging is not possible in certain unfriendly locations: creating unobservable zones near the very targets of most interest to the defense and intelligence communities.

Overall, it seems that nothing short of a full Satellite Radar constellation is capable of providing acceptable performance for this mission. In this case, the BATNA for each stakeholder should be “do nothing,” as the observational need will go unmet. Even if an alternate solution *were* able to satisfy the system requirements, it would do so with significantly worse performance and lower cost, placing the BATNA in the same relative location as “do nothing.” Thus we can be confident that “do nothing” for zero cost and zero benefit will at least present an accurate reference point, and will put all of the alternatives in Quadrant 1 of the negotiation tradespace.

It is possible that one or more stakeholders could take an “other opportunity” BATNA and get some utility for their money. However, given that these are nominally government stakeholders and the cost of Satellite Radar far exceeds discretionary funding, any “other opportunity” would require a completely separate authorization from Congress. In other words, “do nothing” in this case does not imply that the money *not* spent on Satellite Radar is simply retained by the stakeholders – the money would never be procured in the first place. Therefore “do nothing” is particularly appropriate here: if the Satellite Radar design negotiations fail, no system will be bought by any stakeholder<sup>25</sup>.

---

<sup>25</sup> Looking at this problem with a larger scope, it could be argued that the total funding from Congress earmarked for defense in a given budget will always be limited and thus there *is* an “other opportunity” gained to build another system should Satellite Radar fail to resolve. This observation is fair, but is indirect enough that it would be difficult to justify using the hypothetical successful completion of a different future system design as a reference point against which to evaluate Satellite Radar (never mind the extreme difficulty in estimating that future system’s value

Now it is helpful to consider the structure of the problem with respect to the negotiation. Table 9-10 contains a list of some key features of the tradespace based on the categories of chapter 6. Note that the logistics category is not directly relevant to this case study, as from this point on it will be an informal MSTSE application without actual face-to-face negotiation between the stakeholders.

**Table 9-10: Structural features of Satellite Radar tradespace**

<b>Category</b>	<b>Key Features</b>	<b>Consequences</b>
Stakeholders	<i>Number</i> – 3 <i>Relationship</i> – Frequent collaborators	Toggle pairwise visualizations; less likely to aggressively negotiate
Preferences	<i>Metapreferences</i> – Resilience <i>Divisible Attributes</i> – Cost	Investigate –ility driven value against uncertainty; consider uneven split of costs
Alternatives	<i>BATNA Type</i> – Do nothing <i>Tradespace Size</i> – 69,984 (medium)	Low cost of withdrawal socially, consider 2-stakeholder agreements; do not restrict analysis to a small set of designs early on
Uncertainty	<i>Contextual Uncertainty</i> – Epoch-Era <i>Preference Uncertainty</i> – Requirements negotiable	Analyze across contexts; check excessively restrictive attributes
Logistics	Informal MSTSE	No logistic challenges

The structure of the Satellite Radar case is one that promotes pairwise stakeholder analysis methods such as the negotiation tradespace. With three stakeholders it is reasonable to toggle between different tradespace visualizations corresponding to different comparisons. Additionally, the low effort necessary to “pursue” the BATNA and generally collaborative relationship between defense and intelligence stakeholders can make the negotiation a relatively cordial exploratory activity, opening the possibility of two-stakeholder agreements without necessarily generating hard feelings.

The other main structural component of interest here is the presence of uncertainty. Contextual uncertainty has been modeled with Epoch-Era Analysis, providing a useful stable of analysis techniques for exploring exogenous impacts on the system. In addition to this, the stakeholders likely have metapreferences for *resilient* designs over single-context optimal designs. Resilience is currently a point of emphasis for defense-related acquisitions and is typically treated as a result of value sustainment across contextual uncertainty, often driven by one or more –ilities. Both passively robust and actively changeable designs can be considered

---

without an equally elaborate modeling and simulation effort, or the possibility that it too could require multiple stakeholders).

for this task. Finally, as demonstrated in the original analysis, there is enough preference uncertainty present in this case that the performance requirements are negotiable. Normally, without active stakeholder participation in informal MSTSE, it is dubious to modify their elicited utility functions. However, based on the activities of the prior negotiation session, checking attribute requirements that eliminate many alternatives and examining their real-world justification is both acceptable to these stakeholders and a potential way to create more valuable solutions.

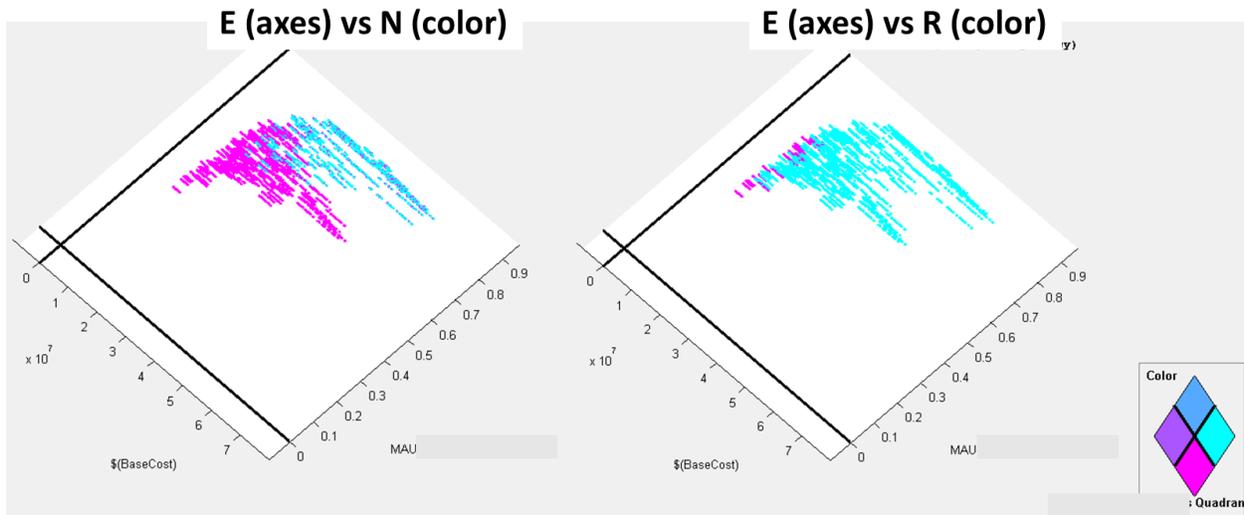
## **9.5 MSTSE Analysis**

This section will demonstrate how the multi-stakeholder visualizations described in chapter 5 can replicate the analysis conducted in the original negotiation session but in greater detail, leading to additional insights. Unless otherwise specified, the images in this section were generated on the same data in the “baseline” context of the prior analysis.

To begin, we can look at the negotiation tradespaces for each stakeholder. As gathered from the BATNA, every valid alternative is located in Quadrant 1. However, coloring by quadrant can quickly reveal regions of each stakeholder’s tradespace where *other* stakeholders find invalid designs. Figure 9-1 shows this for E, with colors for N on the left tradespace and R on the right tradespace (these can be quickly toggled in live application). Cyan indicates a valid design, in Quadrant 1, and magenta an invalid design<sup>26</sup>. Clearly, a large fraction of E’s tradespace is infeasible for N, particularly a large cloud mostly in the low-cost, low-benefit area corresponding to designs with small (10 m<sup>2</sup>) antennas or few (5-10) satellites. In contrast, R has only a small number of infeasible designs appear, and these have nearby valid designs on the Pareto front.

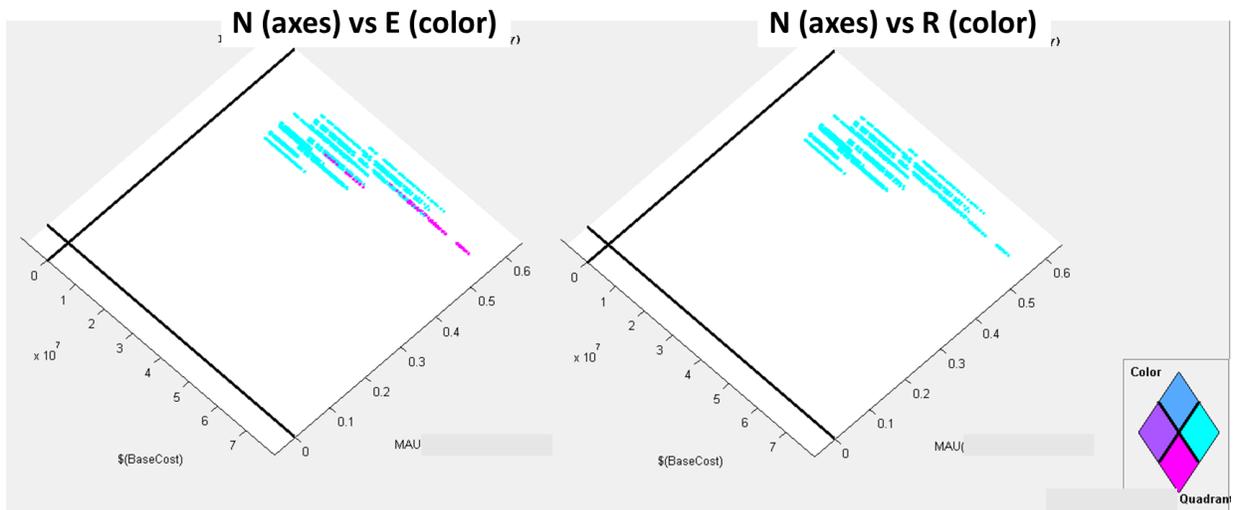
---

<sup>26</sup> Magenta corresponds to Quadrant 4, which is *technically* incorrect as invalid designs have undefined utility and thus are not “in” the tradespace at all. However, the analogy to an alternative strictly worse than the BATNA is clear enough that the same purpose is achieved.



**Figure 9-1: Negotiation tradespace for Stakeholder E, colored for each other stakeholder**

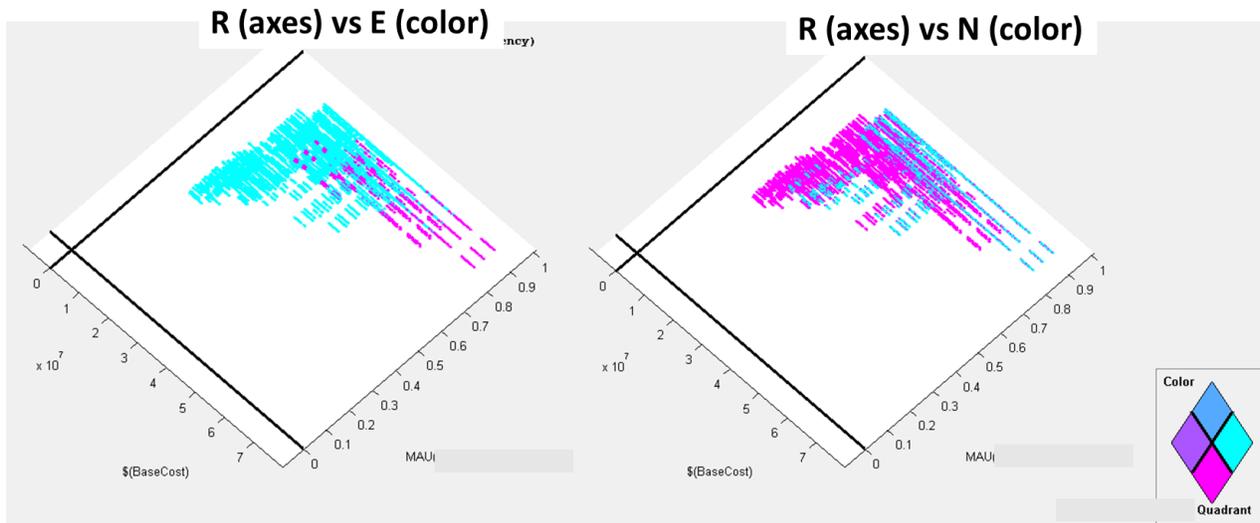
Figure 9-2 shows the same tradespaces for N. As we just saw, N has a much smaller yield of feasible designs. Fortunately, the requirements of the other stakeholders do not have significant ramifications on N’s tradespace. All of these alternatives are also feasible for R, and only a few dominated solutions are invalid for E.



**Figure 9-2: Negotiation tradespace for Stakeholder N, colored for each other stakeholder**

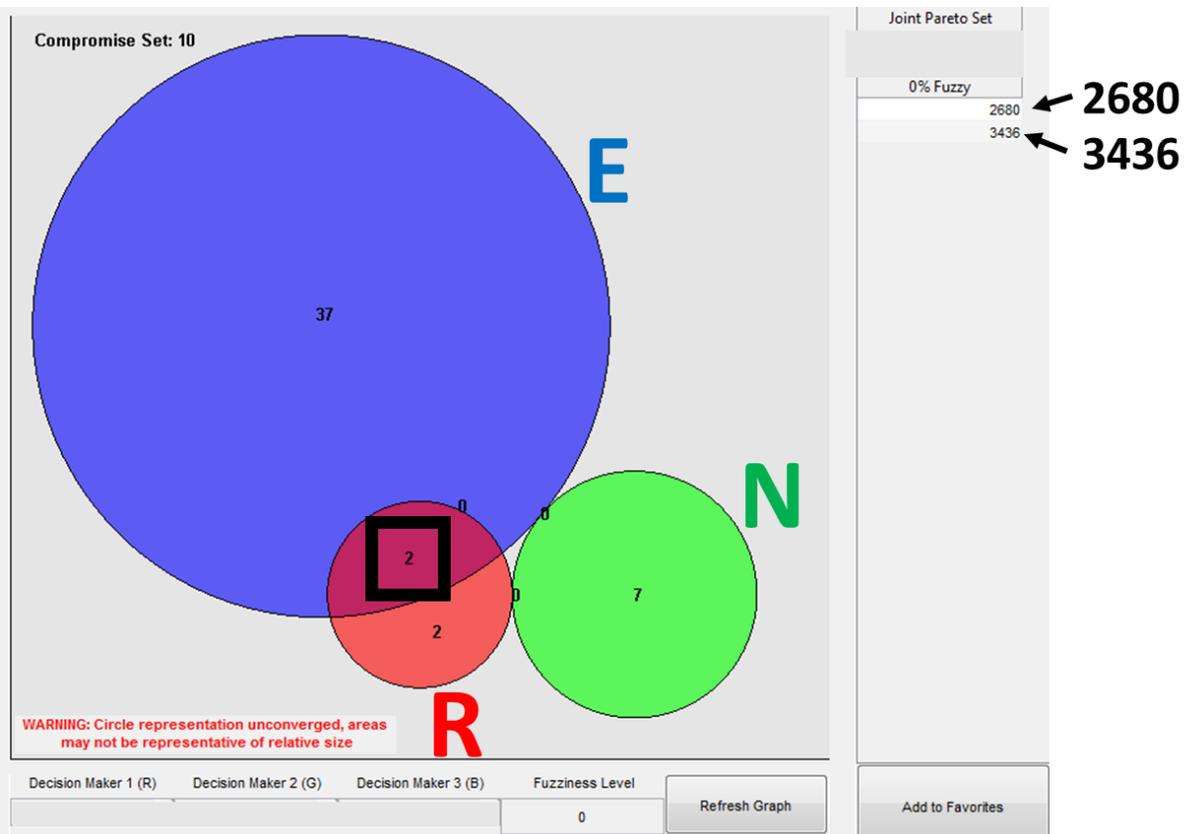
Figure 9-3 repeats this for R. Stakeholder E eliminates a fairly large number of alternatives, but they are located either on the expensive tail end of the Pareto front or in the middle of the tradespace. Again, the large number of infeasible designs for N are clearly visible, but this time the cloud of infeasible points is much more closely overlapping with the feasible ones. The only clear region of shared acceptable designs is at the flat top of R’s tradespace – a large number of designs undifferentiated in performance but with a large range in cost. This

suggests that stakeholders R and N will have the most difficulty finding shared low-cost and efficient designs.



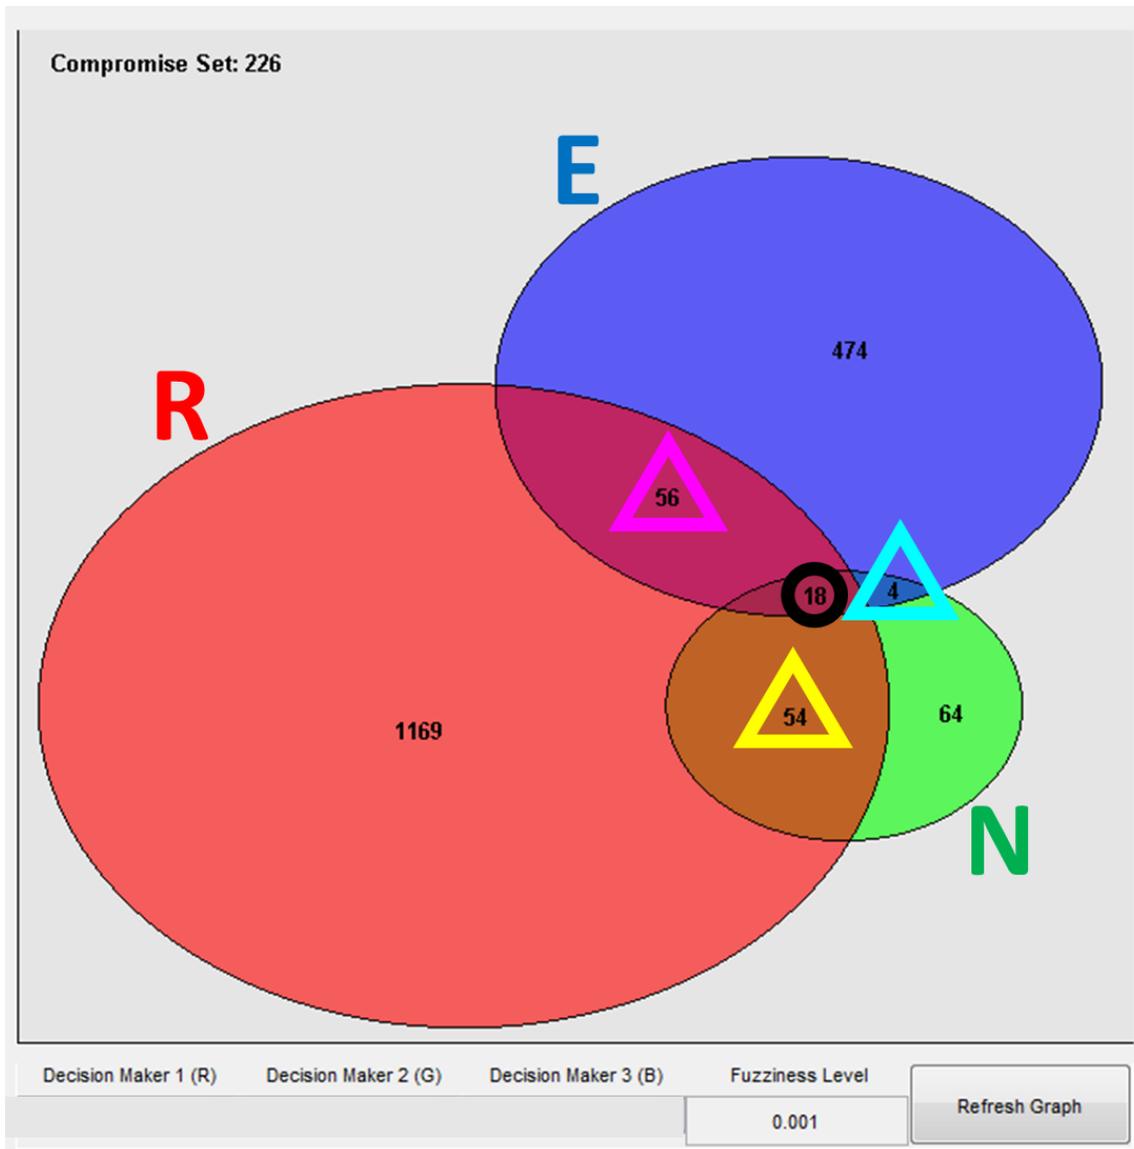
**Figure 9-3: Negotiation tradespace for Stakeholder R, colored for each other stakeholder**

Moving on, we can use the new visualizations to rapidly perform the Pareto set analysis of the three stakeholders. Figure 9-4 shows the Pareto Venn diagram for the three stakeholders. We can quickly see that only R and E share any Pareto efficient designs: 2680 and 3436. 3436 was found to be a shared favorite of R and E in the original live negotiation. 2680 is the same as 2679, a design not found until later on as a result of the design variable “easy” compromises – just without the tactical communication favored by N. Viewing the Pareto sets this way drastically reduces the “busywork” associated with checking between design lists and immediately identifies these potentially interesting opportunities.



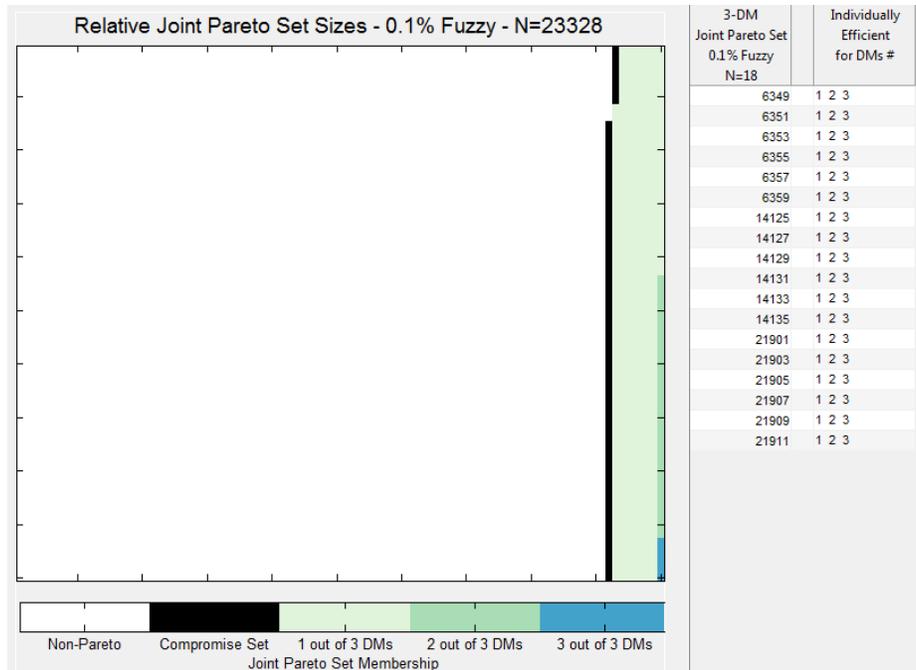
**Figure 9-4: Satellite Radar Pareto Venn diagram**

Perhaps more importantly, the Pareto Venn diagram is particularly effective when comparing fuzzy Pareto sets with many more designs. Figure 9-5 shows the same diagram but with a 0.1% fuzziness buffer added. Despite the very small fuzziness added, the large increase in the size of these sets is immediately noticeable, particularly for R and E who move over 1,000 and 500 efficient designs, respectively: far too many to effectively compare manually. This implies there are many near-efficient designs for each stakeholder – particularly at the “top” of the tradespace where many alternatives have nearly-equal utility due to maximizing the value of some attributes. We also now see that there is overlap for all stakeholders, including 18 jointly-fuzzy-efficient designs across all three. This information is much more detailed than that provided by the compromise Pareto set alone in the original analysis, as it finds “good” designs for each pair of stakeholders rather than lumping all stakeholders together. Each intersection has been marked with a symbol that will carry over into the following figures: black circles for the jointly efficient designs, and differently colored triangles for each pair.



**Figure 9-5: Satellite Radar Pareto Venn diagram - 0.1% fuzzy**

Before looking at these designs in more detail, consider the gridmap representation of the same 0.1% fuzzy Pareto sets in Figure 9-6. Recall that the specific goal of the gridmap was to generalize the Venn diagram for more than three stakeholders, which is not a concern for this case. As such, the additional insights we can draw from this visualization are very limited. The one advantage of this view is that it gives a sense of relative scale between the Pareto sets and the design space at large. The 0.1% fuzzy sets, though perhaps larger than expected, still only comprise around seven percent of the complete tradespace. Adjusting the fuzziness can be used to customize the “depth” to which the stakeholders wish to look for good solutions.



**Figure 9-6: Satellite Radar gridmap - 0.1% fuzzy**

We can explore the design variables of the elements in the different Pareto set intersections to discover patterns not only for single stakeholders but for possible joint and pairwise solutions. Table 9-11 shows the design variables in common between the elements of each intersection, with the final row clarifying any patterns in the other variables. For example, upon investigation, the 18 designs in the joint fuzzy Pareto set are all very similar: featuring a 5 plane, 20 satellite, 53 degree inclination, and 1200 km altitude orbit, 40 m<sup>2</sup> antenna, 2 GHz bandwidth, and 10 kW peak power. The set is defined by these features, with no preference on the remaining three variables. Note that one design in this set is 6349 – the “gold plated” solution found in the original negotiation. This makes sense, as the high cost serves to make sure the interests of all three stakeholders are satisfied, but now we have greater insight into the key features of this type of system.

Table 9-11: Shared design variables of Pareto set intersections

	Joint	N-E	R-N	R-E
# Planes	5	5	5	5
# Satellites	20	-	20	-
Inclination	53 deg	53 deg	53 deg	-
Altitude	1200 km	1500 km	1200 km	-
Antenna Area	40 m <sup>2</sup>	-	-	-
Bandwidth	2 GHz	-	2 GHz	-
Peak Power	10 kW	1.5 kW	-	-
Comm Downlink	<i>any</i>	No	<i>any</i>	-
Tactical Downlink	Yes	Yes	Yes	-
Maneuverability	<i>any</i>	Base	<i>any</i>	-
Constellation Option	<i>any</i>	None	<i>any</i>	-
Notes		10 sats + 40m <sup>2</sup> + any bandwidth OR 20 sats + 100m <sup>2</sup> + 2 GHz	40 m <sup>2</sup> + 20 kW OR 100 m <sup>2</sup> + 10 kW OR 100 m <sup>2</sup> + 20 kW	No clear patterns, but includes all Joint w/o Tactical Downlink

For the two-stakeholder intersections, the patterns are more elaborate. The R-N pair is very similar to the joint set, expanding to include all combinations of 40 or 100 m<sup>2</sup> antennae with 10 or 20 kW power. This suggests that R and N are more willing to scale up from the joint solutions than E, who does not prefer these designs. The N-E pair on the other hand would rather increase altitude to 1500 km, and is willing to take fewer satellites and lower power to do so. Interestingly, this pair also has a single scale up design using a 100 m<sup>2</sup> antenna as well. Finally, the R-E intersection has only 5 orbital planes as a shared feature. The majority of the set consists of designs with 1500 km altitude, 10 kW power, and 2 GHz bandwidth but these and the remaining design variables are all traded with no clear patterns. The only obvious group of designs in the set are 18 designs that match the joint set but without Tactical Downlink.

We can take these insights and verify them against the negotiation tradespace again. This time, we will use the negotiation tradespace rotated into its Quadrant-1-only mode and without color, but we will set a threshold of 5 FPN on the transparency scale. As such, we will only see designs that are “good” for the second stakeholder as gray circles, with the fuzzy Pareto set intersections marked in circles (joint) and triangles (pairs). The tradespaces are shown in Figure 9-7.

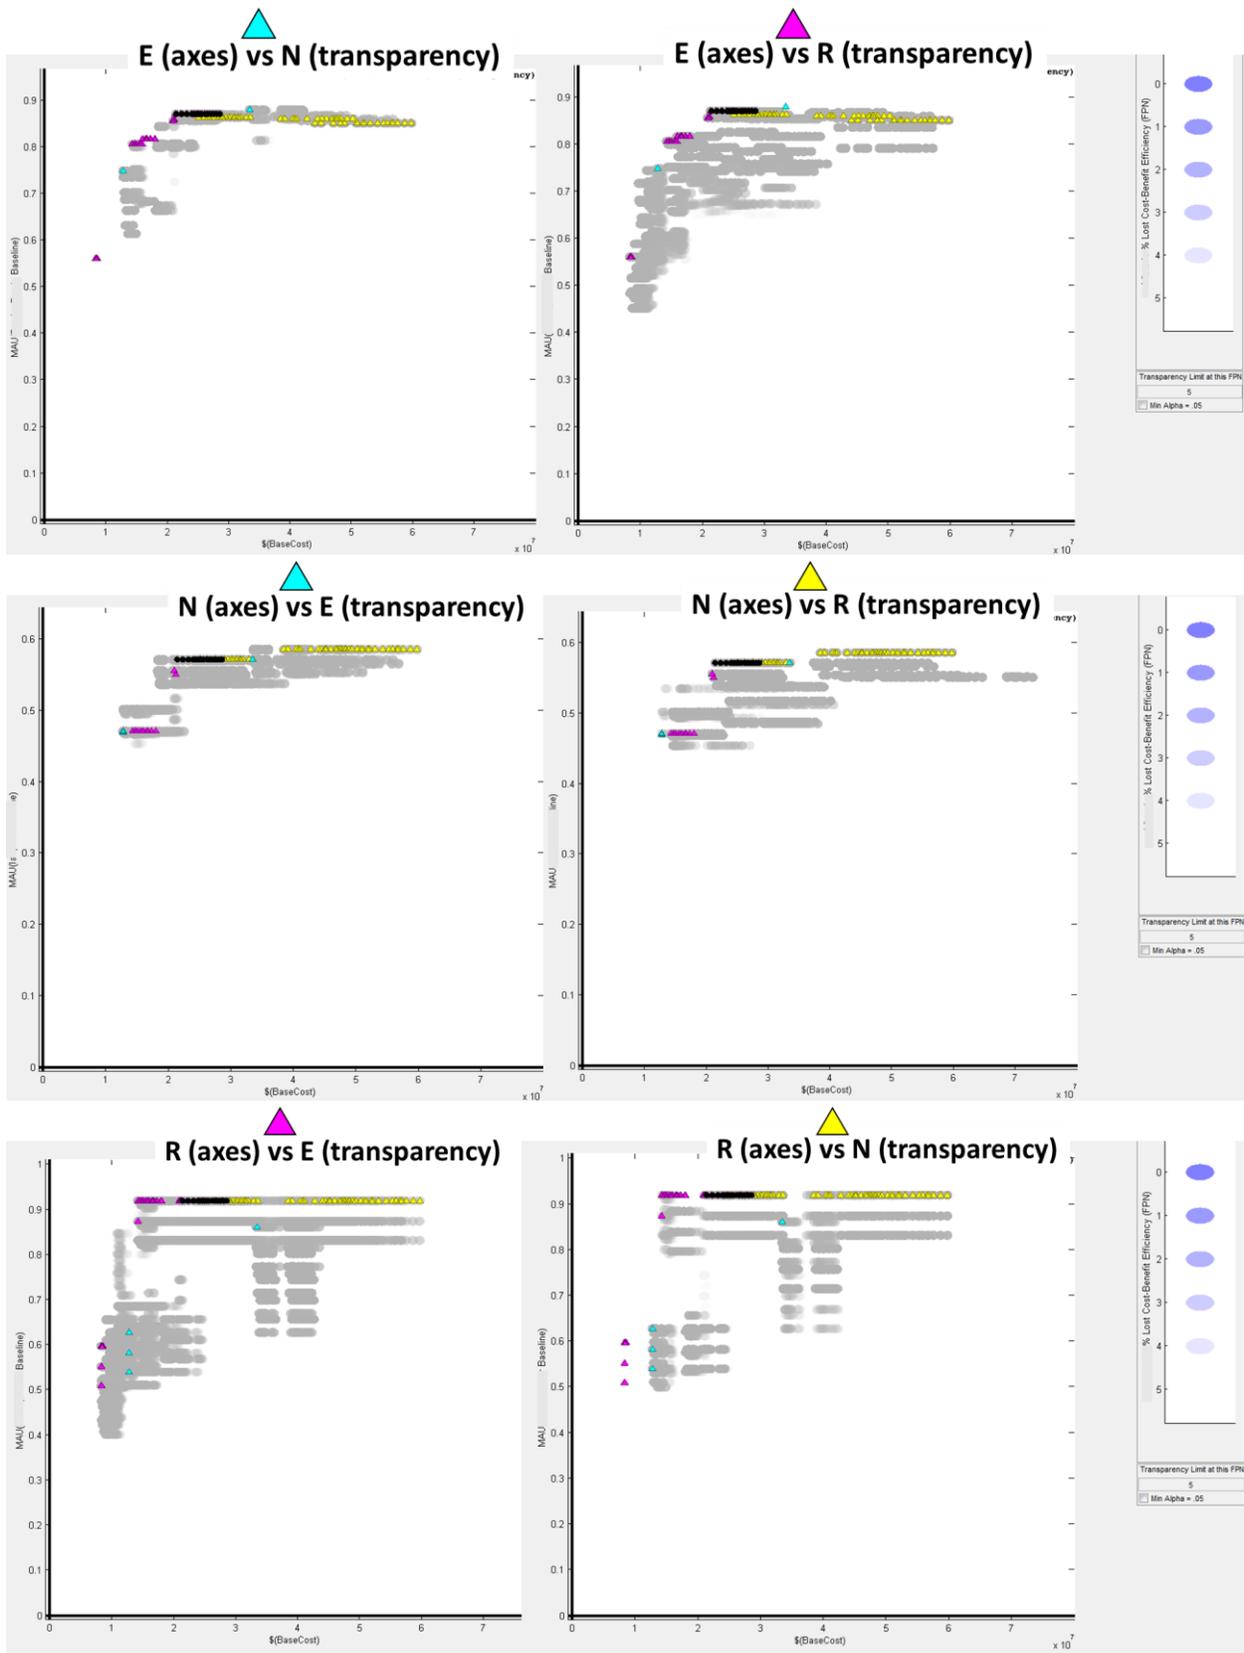


Figure 9-7: Satellite Radar negotiation tradespaces with 0.1% fuzzy Pareto set intersections marked

These tradespaces corroborate the insights from the design variables. The joint set (black) has high utility for all three stakeholders, without significant distinction between them. The R-N pair (yellow) scales up the joint solutions into more expensive designs, though the increase in benefit is slight and experienced mostly by N. The N-E pair (cyan) has distinct lower-cost and higher-cost solutions. The background points can also provide additional insight. For example, the tradespaces with N's FPN as the transparency clearly show that the low-cost R-E solutions (magenta) are invalid for N, as they do not have gray points behind them. However, looking at the R vs. N tradespace in the bottom-right shows that, if those invalid points are removed from consideration, the low-cost N-E solutions located to their right would also be among the preferred low-cost solutions for R. Also note the areas of semi-faded points near the "knee" of the Pareto front in the N vs. R and R vs. E tradespaces – these designs would be included in the yellow and magenta sets, respectively, if we chose to increase the "depth" of our fuzzy Pareto fronts. At that point, we would likely find additional design patterns if we reexamined the design variables in those clusters.

Moving from design-centric analysis to relationship analysis, Figure 9-8 shows the utility function correlations between the three stakeholders, excluding invalid designs on the left and including them on the right. Looking at the left plot, there is a high correlation visible between N and E; however, this decreases when including the invalid designs. This means that, E and N rank the designs in N's tradespace very similarly but, when including the many designs that N rules out, the gap between them grows – i.e. E prefers many of the designs N eliminates to the designs he does not. Overall, it is valuable from a negotiation perspective to note that this entire visualization is green, suggesting that the interests of all participants are more aligned than not. Even if we had done no previous analysis, this would suggest that it is highly likely that alternatives that are good for all three stakeholders do exist.

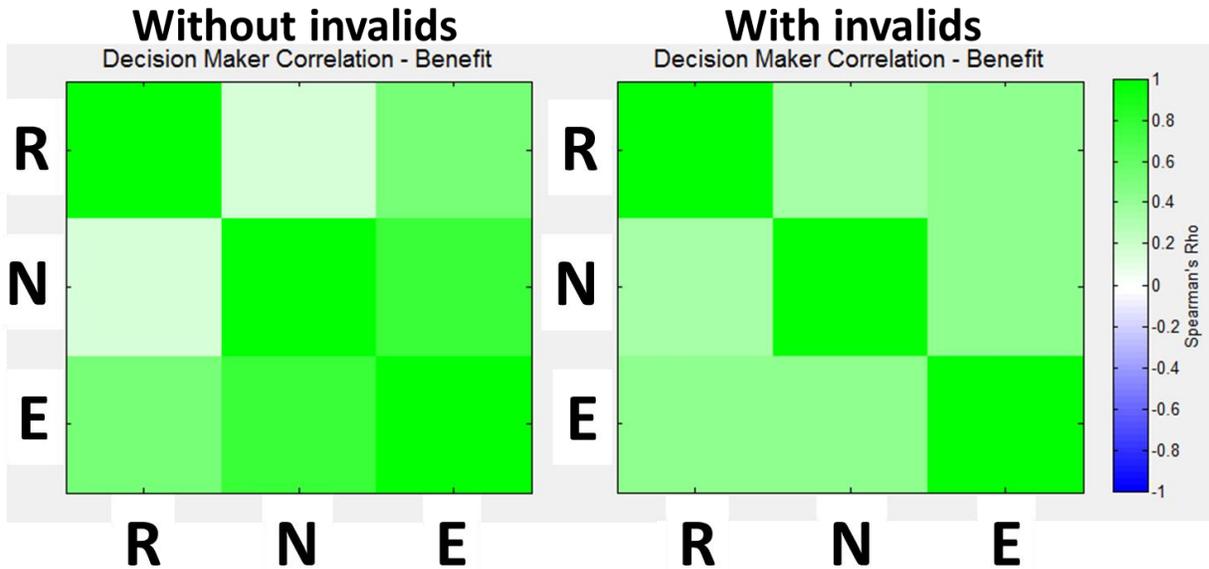


Figure 9-8: Satellite Radar utility function correlations

To illustrate a point, consider the same correlation plots but calculated with FPN instead of utility, shown in Figure 9-9. Though not identical, it is clear that the patterns in these plots are the same as those in the utility plots. This is to be expected when each stakeholder is using the same cost function, as differences in utility are then the only driver of differences in FPN. The Spearman Rho correlations change slightly because FPN correlation is driven disproportionately by the Pareto efficient designs (which define FPN, in addition to being correlated themselves), rather than equally for all designs in the case of utility.

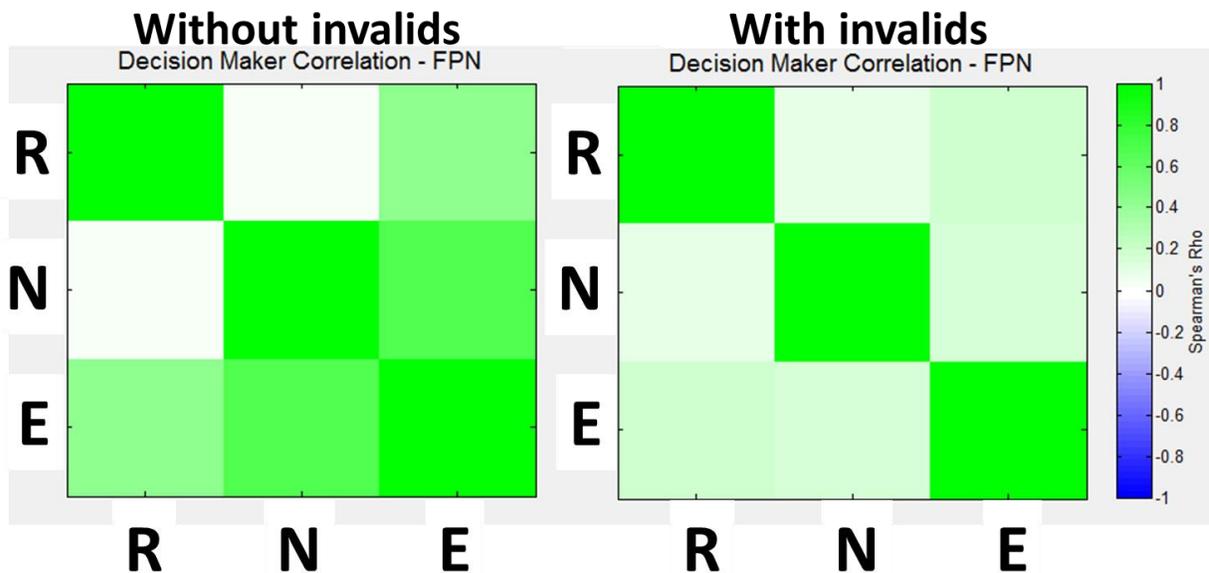


Figure 9-9: Satellite Radar FPN correlations

Interestingly, the FPN correlations do reveal additional insights when viewed in their scatterplot form, as in Figure 9-10. First, the red circle highlights the area of lowest (that is, best) FPN for R and N. Unlike the other scatterplots, this area is mostly empty except for a single point very close to zero-zero. That point is actually 72 overlapping points: the 72 designs that were in the joint and R-N fuzzy Pareto sets. However, based on this plot we know that increasing the fuzziness will increase the set sizes for the other intersections, but NOT joint or R-N. If one of those 72 designs is not selected as the final agreement, we can expect that R and N will be the most difficult to reconcile when searching deeper into the tradespace. This meshes perfectly with the insight from Figure 9-3, which predicted that R would have only a small area of efficient designs after accommodating N's requirements. Also note that this visualization can explain why the Pareto Venn diagram showed E as clearly having the most Pareto efficient solutions with no fuzziness, but being immediately surpassed by R when including fuzziness. The density histograms show that E's tradespace is densest on the Pareto front, decreasing as FPN increases, while R's tradespace peaks in density just off of the Pareto front.

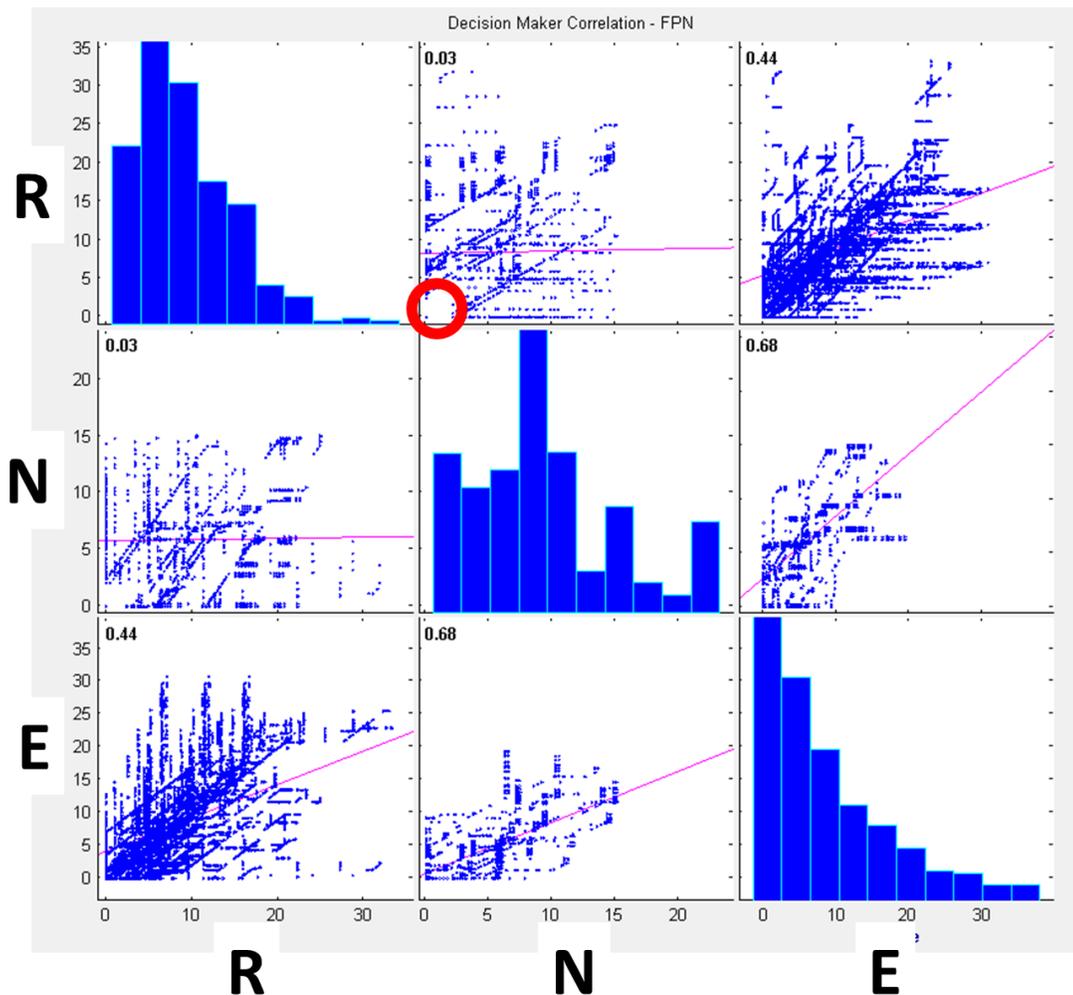


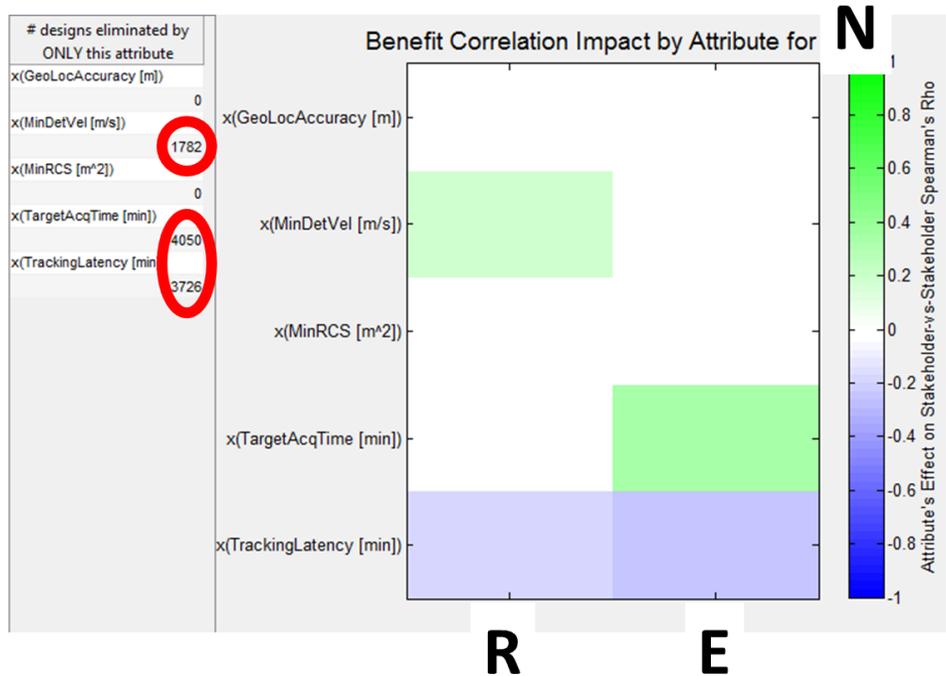
Figure 9-10: Satellite Radar FPN correlations - Scatterplot mode

Finally, we can consider the correlation impacts of each attribute in the stakeholders' utility functions. With this analysis, we can quickly identify "free" issues and divisive ones. Each of the following figures uses scaled correlation and includes invalid designs. Figure 9-11 shows the breakdown for E. In his case, Targets Per Pass is most agreeable to the others and can likely be maximized. On the other hand, Imaging Latency lowers correlation with R and N in addition to having E's strictest requirement, eliminating 1,188 designs by itself.



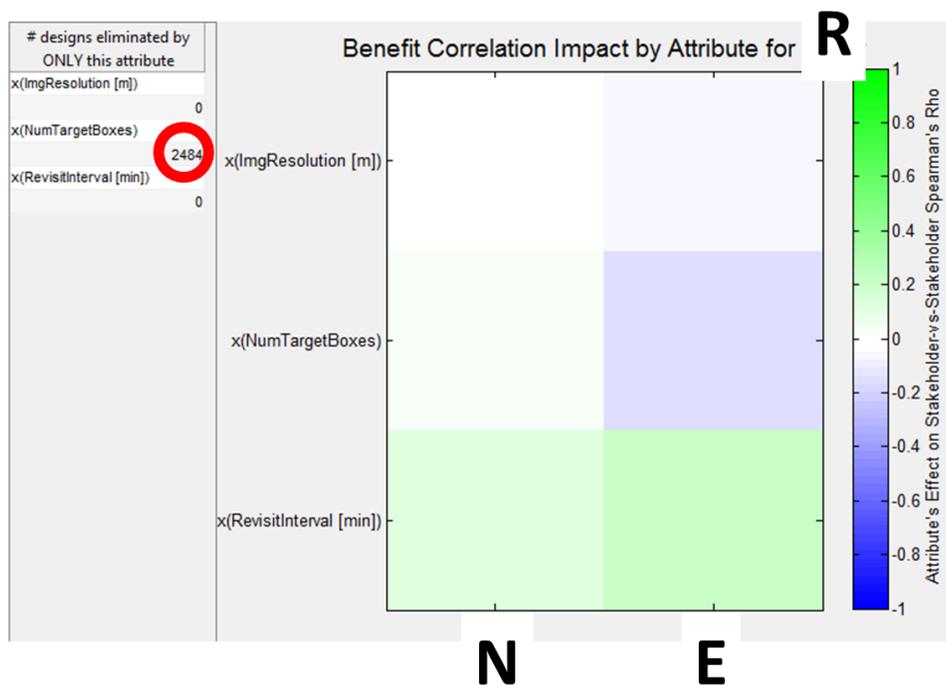
**Figure 9-11: Attribute correlations for Stakeholder E**

Figure 9-12 shows the same for N. In his case, there are multiple dangerous requirements, resulting in the large number of invalid designs that we have already noted. In particular, Tracking Latency eliminates 3,726 designs by itself and is negatively correlated with both other stakeholders. Additionally, this figure is a good example of one impact of having many attributes in a utility function: Geolocation Accuracy and Minimum RCS both have no impact on stakeholder correlation due to their low swing weights. Fortunately, neither of these attributes have requirements responsible for eliminating alternatives – if they did, those requirements should be questioned, as the attribute has little impact on value anyway.



**Figure 9-12: Attribute correlations for Stakeholder N**

Finally, Figure 9-13 shows that Revisit Interval is the most agreeable objective for R, and that Number of Target Boxes is his only preference to eliminate alternatives (in addition to positioning him slightly against E).



**Figure 9-13: Attribute correlations for Stakeholder R**

Overall, the level of detail provided by the multi-stakeholder visualizations is considerably higher than that achieved with the standard tools in the original negotiation session. In addition to finding the original insights, we were able to:

- Identify regions of the tradespace widely removed by requirements (i.e. N eliminating all alternatives with 10 m<sup>2</sup> antennas and most with 10 or fewer satellites)
- Separate out the patterns in efficient designs pairwise rather than by the group as a whole, to gain a sense of how each pair would choose to deviate from the group if working together (i.e. N and E would trade some variables down in order to push altitude higher)
- Predict areas of interest if search depth increases and which stakeholders will be least likely to be satisfied by such a move (R and N)
- Quantify the correlation between the stakeholders and capture the impact that designs marked invalid by requirements have on their alignment
- Measure the impact on alignment driven by each utility attribute, identifying issues that are in agreement across stakeholders (Targets Per Pass and Revisit Interval) and those that drive disagreement (Imaging and Tracking Latency, Number of Target Boxes, Resolution)

## ***9.6 Revised Tradespace Formulation and Modeling***

While preparing to analyze the satellite radar dataset with MSTSE, it became apparent that issues with the evaluative model were manifesting in some strange and counterintuitive insights from the tradespace. Further examination revealed that the original RSC satellite radar model contained a few programmatic errors, as well as an underlying structure that did not support the assumptions inherent in EEA and was preventing effective analysis across uncertainty. The prior section analyzed the original data in a static context with the new MSTSE visualizations anyway, in order to provide a direct comparison against the original TSE and negotiation activities. However, a refresh of the evaluative model was necessary in order to generate more detailed insights into the case. The updates to the satellite radar model were as follows:

- **Error handling.** The different modules contained error flags for different impossible or invalid situations, however not all of the flags were propagated back to the main script. This was corrected, resulting in some previously-feasible designs failing a mission. Additionally, certain errors resulted in the termination of the evaluation midway through completion in order to save time. However, these errors were sometimes applicable to *only one of tracking or imaging* – meaning that if allowed to finish, the design could potentially be an interesting solution for another stakeholder. This is indicative of an assumption that a solution *must* satisfy all stakeholders, denying the possibility that

smaller coalitions could build together or use one of these designs as a discussion point in the negotiation. Therefore, these preemptive terminations were disabled.

- **Geolocation accuracy.** The geolocation calculations had unit conversion errors that were corrected.
- **Radar swath.** The radar swath code was working incorrectly on the corner case when the minimum detectable velocity of the satellite was greater than the target velocity. In this case, the swath should be nonexistent as the satellite can never track the target. This fix corrected an unintuitive result where, for some missions, large antennas were reported as having worse performance than small antennas.
- **Mass.** The mass of a satellite was originally calculated in the bus module as a function of the operations, including the operational context. This would cause a design to change mass between epochs (via alternative sizing of batteries and solar arrays, for example), conceptually changing the actual design build and preventing accurate analysis of performance of a single physical system across the epoch space. To correct this, the mass calculated by the model was treated as a mass *requirement*, necessary to complete the mission in the current epoch. The maximum mass for each design across the epoch space was recorded and a new design variable was added corresponding to the system’s mass budget: the fraction of the maximum necessary mass actually budgeted into the satellite design. Designs with insufficient mass to meet the mass requirement for a given epoch were declared infeasible in that epoch<sup>27</sup>. Thus, building a lighter satellite will result in lower cost, but risk failure in more challenging mission contexts. The mass budget was enumerated with the levels [1, 0.75, 0.5], tripling the number of alternatives to 69,984. The resulting design space is shown in Table 9-12.

**Table 9-12: Revised Satellite Radar design space description**

Design Variables	Levels	Enumerated Range	Units/Notes
Walker ID Number	8	[1:8]	see Walker Lookup Table
Orbit Altitude	3	[800, 1200, 1500] *10 <sup>3</sup>	m
Antenna Area	3	10, 40, 100	m <sup>2</sup>
Peak Transmit Power	3	1.5, 10, 20	kW
Radar Bandwidth	3	[0.5, 1, 2] *10 <sup>9</sup>	Hz
Comm Downlink	2	Backbone, Ground only	
Tactical Downlink	2	Yes, No	
Maneuver Package	3	Base, 2x Base, 4x Base	
Constellation Option	3	1, 2, 3	1=none, 2=long lead, 3=spares
Mass Budget	3	1, 0.75, 0.5	fraction of budget used

<sup>27</sup> A more elaborate modeling effort could attempt to calculate the “partial performance” of the mission when the satellite is undersized/underpowered, but that level of detail is not commensurate with the rest of the model.

In addition to these changes, it was also discovered that the target constructor corresponding to the Operations Plan context variable was not working as intended. Rather than simply fix this problem, Operations Plan was dropped as a variable in favor of a more transparent combination of three context variables: target size, speed, and location. This should make tradeoffs between designs as a function of the target of interest more clear. Additionally, the AISR context variable was also removed (set to “none”), because this research is not particularly interested in the systems-of-systems features of the satellite-AISR network. The presence of AISR in a given location would almost completely supplant the contribution of the satellites, to the detriment of our ability to differentiate the different satellite alternatives. The resulting context space is shown in Table 9-13, which results in a full-factorial enumeration of 324 contexts.

**Table 9-13: Revised Satellite Radar context space description**

Context Variables	Levels	Enumerated Range	Units/Notes
Radar Technology	2	1, 2	1=TRL 9 only (low), 2=TRL 6 and above (med)
Communication Infrastructure Level	3	1, 2, 3	1=low: no backbone+AFSCN 2=mid: WGS+AFSCN 3=high: TSAT+AFSCN
Threat Environment	2	None, Jamming	
Target Size	3	5, 15, 100	m <sup>2</sup>
Target Location	3	Russia, Iran, Venezuela	
Target Speed	3	20, 45, 100	m/s

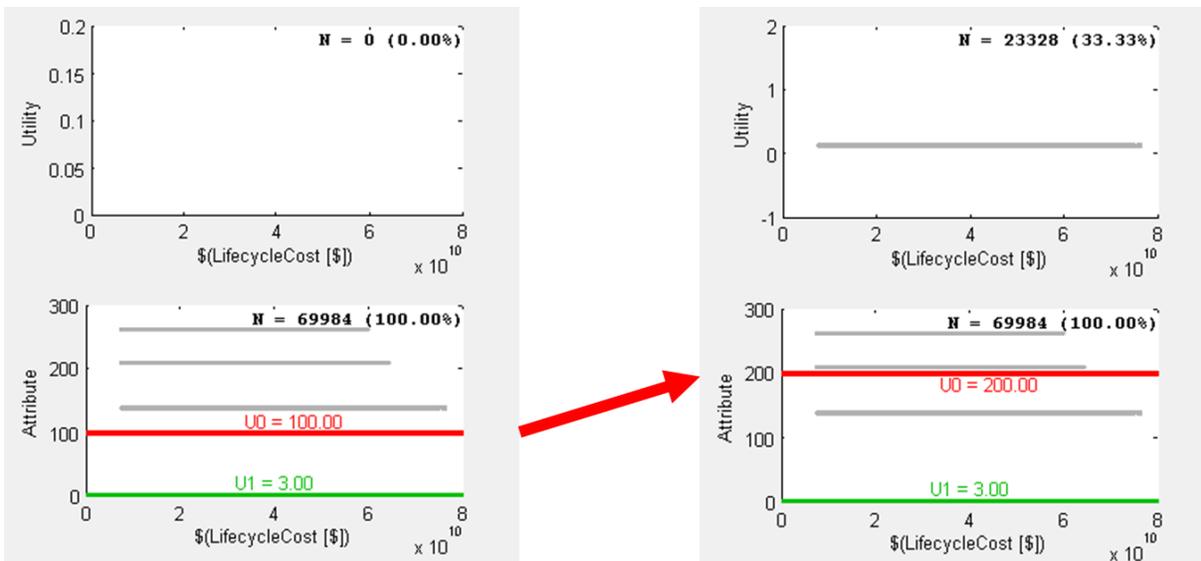
## 9.7 Validating the Revised Tradespace

With the new model formulation, we can now evaluate the alternatives against uncertainty, as represented by the context space. However, before leveraging tradespace metrics to identify designs of interest across multiple contexts, we should validate the adjustment in modeling: making sure that it does not have a drastic or degenerate impact on the tradespace and, when it is different, the differences are realistic. Of particular interest are the attribute requirements in the utility functions for each stakeholder, as they will remove designs from the tradespace and limit the set of alternatives that we can negotiate over. In this case, one requirement is notably harsh after incorporating the improved modeling results: N’s requirement on Geolocation Accuracy. The adjustment to the geolocation module in the evaluative model decreased the accuracy of each alternative, to the point where *none* are capable of meeting N’s requirement of 100 meters at any time (despite previously *all* designs meeting this requirement). With a heavily constraining impact such as this, it is worthwhile to engage the stakeholder on the source of the requirement, so that if the requirement is negotiable the negotiation doesn’t come to an end when N sees that no designs are feasible. Because this is an informal MSTSE analysis, N

is not participating; however, he was amenable to exploring reduced requirements in the original session and thus we should consider the practicality of sub-100-meter accuracy.

Satellite accuracy is primarily a function of altitude, antenna length, and pointing accuracy. 100 meters is generally considered a “very high” accuracy for orbital imaging – according to Grace et al. (2004) the approximate antenna length necessary to achieve 100 meter accuracy at the altitude range considered in this tradespace (800 – 1500 km) is also on the order of 100 meters. Compared to the antenna lengths in the tradespace (2 – 20 m, depending on antenna area), it is unsurprising that that degree of accuracy is not achieved by any of the alternatives. Looking to similar systems, WindSat (US Naval Research Laboratory, 2016) operates at the lower end of our altitude range (840 km) and has a geolocation accuracy of approximately 1,000 meters, though its weather observation mission is understandably less stringent than potential military or intelligence applications (Purdy et al., 2006). Commercial optical satellite systems which offer sub-100-meter accuracy, such as GeoEye, WorldView, and QuickBird (Ozcanli et al., 2014; Smiley, 2009), are located at lower altitudes to combat attenuation: roughly 500 – 700 km. Thus, it appears that the 100 meter accuracy requirement was set unrealistically for the design space under consideration, short of modeling at an increased level of detail able to capture design trades enabling superior pointing accuracy. Without direct stakeholder access to N, we are unable to confirm if there is a specific mission task that requires such a high accuracy. In his absence we can still perform “what if” analysis with a relaxed requirement, similar to what he agreed to in the original negotiation, though on a new attribute.

Figure 9-14 shows N’s single attribute utility function for Geolocation Accuracy before and after a relaxation (left to right). The revised tradespace model returns accuracies between approximately 150 and 250 meters, which are reasonable considering the system examples described above – certainly more realistic than the sub-100-meter results of the original model. The red line in the bottom plot indicates the requirement: the worst acceptable accuracy. On the left, no designs meet the threshold as indicated by the empty scatterplot in the middle. On the right, the requirement has been moved to 200 meters, and one third of the alternatives become feasible. These are the 800 km altitude designs, which have the best accuracy available in the tradespace. Again, without N’s participation, it is difficult to justify a *specific* new requirement level. However, because the accuracy is so heavily striated by altitude, we do not have to worry about the sensitivity of our results to small changes in this requirement: the set of designs under consideration will be identical for any requirement between approximately 140 meters and 200 meters. Making this change allows us to continue exploration of this multi-stakeholder problem without essentially pre-eliminating N from any possible agreement. As such, all figures and analysis from this point on will use this modified utility function for N. N now effectively requires an 800 km altitude system, which will have an impact on the alternatives that are most closely shared between stakeholders.



**Figure 9-14: Geolocation Accuracy Single Attribute Utility function for Stakeholder N, relaxing requirement from 100 m to 200 m**

For the remainder of our validation, we will simply confirm that the model adjustments have not significantly changed the model’s outputs nor the apparent relationships between the stakeholders: if these are the same, we can be confident that any new insights we generate can still be effectively compared against the insights of the original analysis<sup>28</sup>. To do so, we can briefly look at a single context most similar to the original “baseline” context. We can hold the unchanged context variables the same (radar technology and communication infrastructure on low, no jamming) and set the new target variables to most closely match the original baseline target (large and fast target, in Iran). In the new enumeration this corresponds to context #72, which we can compare directly to the prior analysis. Figure 9-15 shows the negotiation tradespaces for each pair of stakeholders in this context. These plots clearly display the same patterns as those generated under the original models: N’s requirements eliminate far more alternatives than the other stakeholders’, including the entire band of lowest-cost alternatives, while E’s few eliminated alternatives are dominated, high-cost designs. In this case, R does not eliminate any alternatives. Additionally, through the similarity in the correlation plots for utility and FPN, shown in Figure 9-16 and Figure 9-17, it is clear that the three stakeholders retain similar relationships: positively correlated across the board, with E and N naturally more aligned than the other pairs. If there are any significant changes, it appears that R and N have more aligned interests than originally indicated, particularly when considering FPN.

<sup>28</sup> This is much less stringent than fully validating the entire model from scratch, but since the original model was itself previously validated and our main goal in this research is not to decide on the “best” system design with our new model but rather to compare the original TSE application and the new MSTSE application, it is sufficient for the analysis we want to perform (particularly since the model changes were minor).

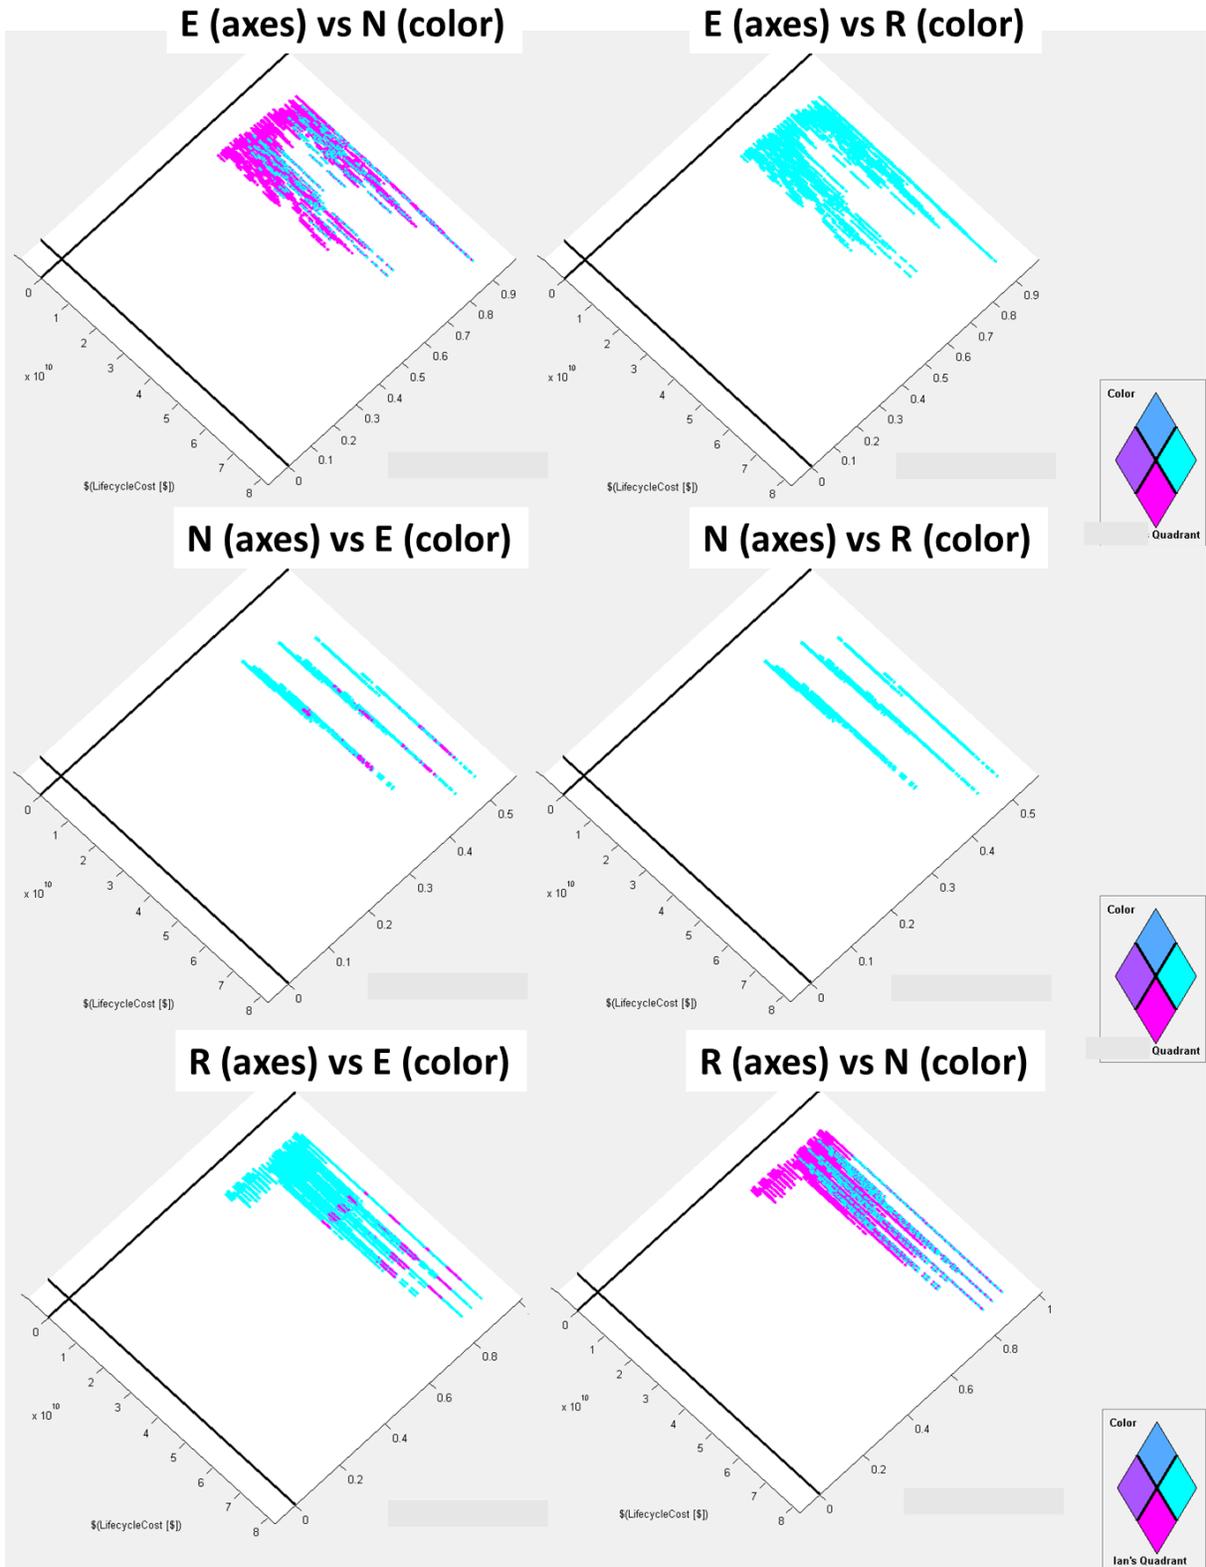


Figure 9-15: Negotiation tradespaces for context #72 in revised Satellite Radar

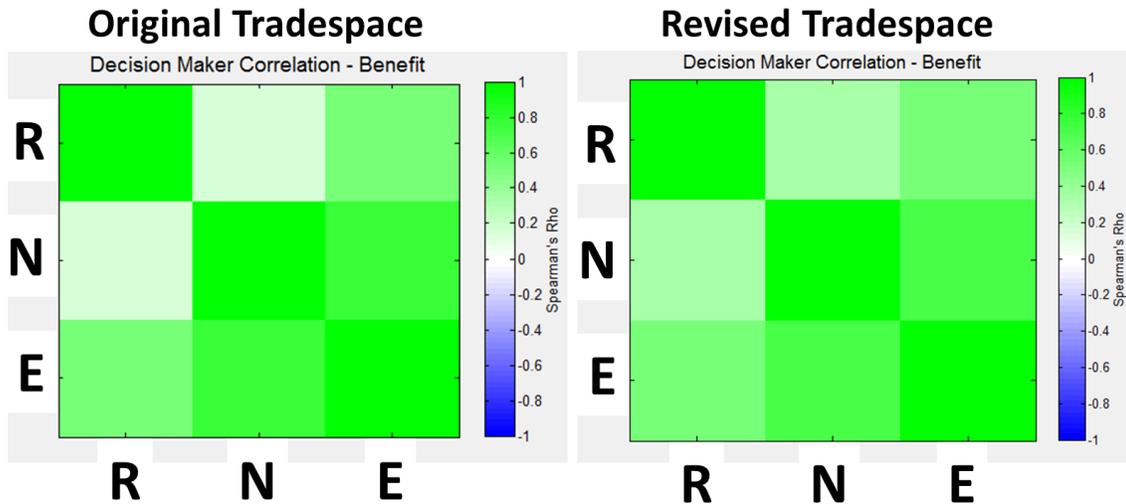


Figure 9-16: Stakeholder utility correlations for original and revised tradespaces

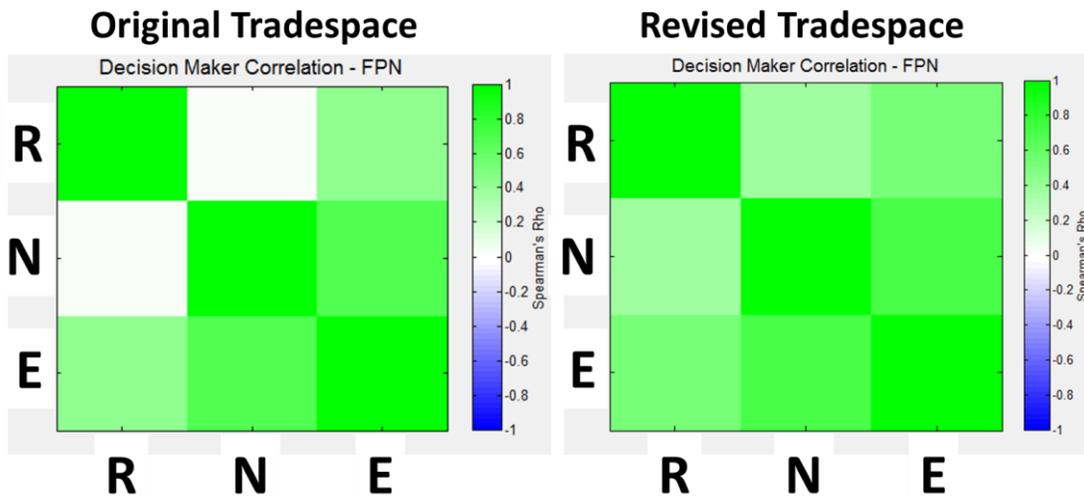
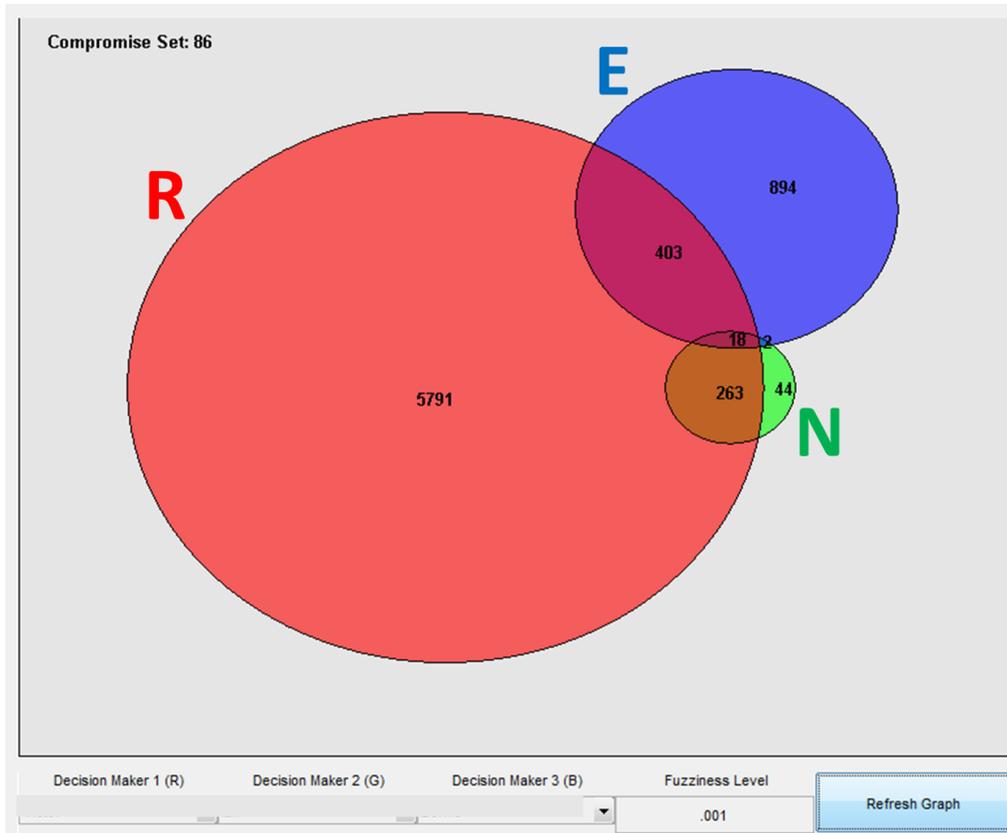


Figure 9-17: Stakeholder FPN correlations for original and revised tradespaces

As a final check on the revised tradespace, we can consider the Pareto Venn diagram again and compare the shared features of the efficient designs. Figure 9-18 shows the 0.1% fuzzy Venn diagram, which shows a similar pattern to the original but with many more designs in the sets, driven by the addition of the Mass Budget design variable adding more alternatives with equal utility but different cost along the top of the tradespace, especially for R. R again has the largest set of near-efficient alternatives, followed by E and then N.



**Figure 9-18: Pareto Venn diagram (0.1% fuzzy) for revised tradespace**

Again, the three-way joint intersection and the three pairwise intersections are all populated with alternatives. Table 9-14 shows the shared design variables of these sets, marking in bold those variables that are the same as under the original tradespace. N's new derived requirement for an 800 km altitude has had a trickle-down effect on the rest of the sets. The joint set is nearly identical to the previous joint set but now is at 800 km and, as a result, can use lower power but requires the use of the comm downlink (due to the reduced line-of-sight area on the ground). The two pairwise sets including N also must switch to 800 km altitude but again share largely similar features as the original analysis. The N-E set is still the most similar to the joint set, but is now actually smaller because the 10 satellite constellations it previously contained are not feasible at the lower altitude. The R-N set is still comprised mostly of more expensive variants of the joint set, with larger antennas, higher power, and little preference amongst the cheaper add-ons like comm downlink, maneuverability, and constellation option. Finally, the R-E set cleans up the remainder of the space not feasible for N, heavily featuring high altitudes and small antennas in the low-cost region of the tradespace.

Table 9-14: Shared design variables of revised Pareto set intersections (bold indicates same as original)

	Joint	N-E	R-N	R-E
# Planes	<b>5</b>	<b>5</b>	<b>5</b>	<b>5</b>
# Satellites	<b>20</b>	20	-	-
Inclination	<b>53 deg</b>	<b>53 deg</b>	<b>53 deg</b>	-
Altitude	800 km	800 km	800 km	-
Antenna Area	<b>40 m<sup>2</sup></b>	40 m <sup>2</sup>	-	-
Bandwidth	<b>2 GHz</b>	0.5 or 1 GHz	<b>2 GHz</b>	-
Peak Power	1.5 kW	<b>1.5 kW</b>	-	-
Comm Downlink	Yes	Yes	-	-
Tactical Downlink	<b>Yes</b>	<b>Yes</b>	<b>Yes</b>	-
Maneuverability	<i>any</i>	<b>Base</b>	-	-
Constellation Option	<i>any</i>	<b>None</b>	-	-
Mass Budget	1 or 0.75	0.75	-	-
Notes		Like Joint, but with fixed maneuverability/constellation option and variable bandwidth	40 m <sup>2</sup> + 10 sats OR 40 m <sup>2</sup> + 20 sats OR 100 m <sup>2</sup> + 20 sats all with various sources of “extra” cost driven by other features	No clear patterns, but features variable levels unsupported by N (high altitude, small antenna, etc.) and a low-cost cluster of 10 m <sup>2</sup> antenna and 1.5 kW power designs

Thus, we can be confident that our understanding of the relationships between the stakeholders is largely unchanged. Since the motivation for revising the Satellite Radar model was to look across different contexts in order to collect insights about the stakeholders and their preferred alternatives while acknowledging the uncertainty in the operational environment of Satellite Radar, we can now examine the impact of those contextual changes on their relationships. Waiting for improved radar technology or communications infrastructure may allow for lighter and cheaper satellites to dominate the tradespace, while more difficult targets may require different orbit parameters.

Looking at the tradespace yield – the fraction of alternatives that are feasible – for each stakeholder across the context space reveals the impact of changing mission environment. The yields are shown as a histogram, one for each stakeholder, in Figure 9-19. All three stakeholders show a large spike in their yield distribution in the zero-yield bin amounting to just over one third of the context space (108). Closer investigation reveals that no designs are feasible for *any* stakeholder when the target velocity is at its slowest setting (20 m/s) as the best minimum

detectable velocity (MDV) in the tradespace is near 23 m/s<sup>29</sup>. Other than the group of “impossible” missions, there are two main patterns in these histograms. E and R share a set of “easy” contexts with yield near 100%, characterized by large, fast targets and available communication infrastructure, while the bulk of contexts are distributed around 50% yield with a mix of “easy” and “hard” context variables. In contrast, N’s yield never exceeds 11%, showing that he remains the most difficult stakeholder to please. This limit can be traced back to strong derived preferences on specific design variables: in any given context, N will only accept alternatives with 800 km altitude (1/3 of the design space), 40 or 100 m<sup>2</sup> antenna area (2/3), and with tactical communication available (1/2).

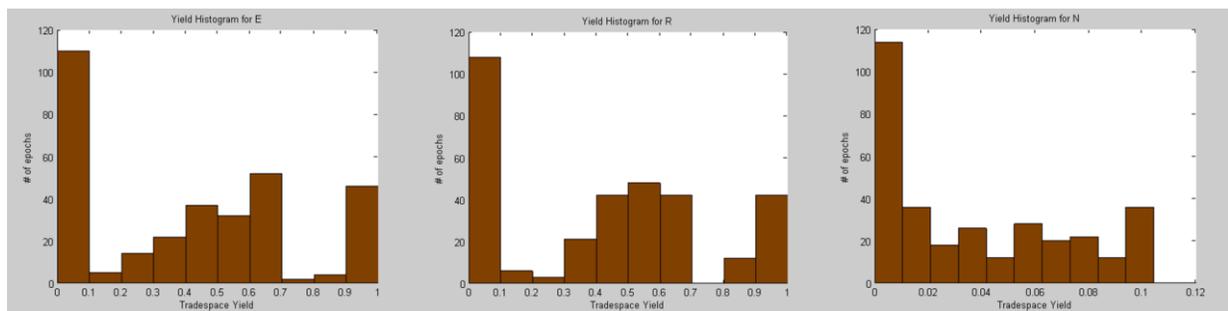


Figure 9-19: Yield histograms for E (left), R (center), and N (right) across the context space

## 9.8 Passively Robust Alternatives

In this research, we will use the word “robustness” to indicate the property of a design that is able to provide value while withstanding changes in environment or operational context that may occur during its lifetime (Fricke and Schulz, 2005). Robustness is often considered an “ility” as it describes a positive property of a system that does not contribute directly to its ability to perform its function. “Ilities” rarely have stable, consensus definitions; the word “robustness” is used interchangeably with other “ilities” such as survivability and resilience in various sources. To begin, we will identify passively robust solutions in the tradespace by calculating the Normalized Pareto Trace (NPT) of each design: the fraction of contexts in which the design is on the Pareto front for a given stakeholder (Ross, Rhodes, and Hastings, 2009). The single alternative, or set of alternatives, with the highest NPT for each stakeholder is shown in Table 9-15. Additionally, the table also shows the “maximin” NPT alternative, by grading each alternative by the stakeholder it is *worst* for; i.e., an alternative must be on the Pareto front for all

---

<sup>29</sup> This is not a particularly impressive MDV because the underlying model uses a very simplified and conservative estimate for the synthetic aperture. Additionally in this tradespace, MDV is driven mostly by antenna size and slightly by altitude. The important variable *not* used to influence MDV here is the beam frequency/wavelength, which was assumed to be in the X band (10 GHz / 3 cm) in the original tradespace enumeration and not changed in the revised tradespace. If the stakeholders were interested in exploring higher MDV designs, adding frequencies in the K<sub>U</sub> band or above to the tradespace may be desirable, though it would require additional modeling of the tradeoffs in signal attenuation in the atmosphere.

three stakeholders (joint) to count for the maximin NPT. The maximin metric appeals to the common fairness criterion of “make the worst-off person as well-off as possible”, but with stakeholder participation, any fairness criterion could be selected to target the best *group* designs across the context space.

**Table 9-15: Highest NPT alternative(s) for each stakeholder, and highest maximin NPT across stakeholders**

	<b>R</b>	<b>N</b>	<b>E</b>	<b>Maximin</b>
<b>NPT</b>	0.1667	0.3333	0.1667	0
<b># Alternatives</b>	1	1	15	-
<b># Planes</b>	5	5	5	-
<b># Satellites</b>	5	20	5 or 10	-
<b>Inclination</b>	53 deg	53 deg	53 deg	-
<b>Altitude</b>	1500 km	800 km	1500 km	-
<b>Antenna Area</b>	40 m <sup>2</sup>	100 m <sup>2</sup>	10 or 40 m <sup>2</sup>	-
<b>Bandwidth</b>	2 GHz	2 GHz	<i>any</i>	-
<b>Peak Power</b>	10 kW	1.5 kW	1.5 kW	-
<b>Comm Downlink</b>	No	No	No	-
<b>Tactical Downlink</b>	No	Yes	No	-
<b>Maneuverability</b>	Base	Base	Base	-
<b>Constellation Option</b>	None	None	None	-
<b>Mass Budget</b>	0.5	0.5	1 or 0.5	-
<b>Notes</b>			40 m <sup>2</sup> antenna only with 5 sats	No alternatives are Pareto efficient for all three stakeholders in any context

Table 9-15 shows the conflict in the *individual* preferences of each stakeholder quite clearly. Other than a common thread of minor cost-reducing features (no communications downlink or constellation options, base maneuverability), the preferred alternatives through NPT are different in the most impactful design variables: number of satellites, altitude, antenna area, and power. As a result of the significant differences in their utility functions, the stakeholders never share a joint Pareto efficient design in any context, hence the maximin NPT of zero. Additionally, the NPTs themselves are quite low – neither R nor E have any alternatives that are efficient for more than one-sixth of the context space. Note that the lowest setting for mass budget, 50% of the maximum required mass, is the most common setting in this table: when the highest NPT is so low, it is still possible to “cheat” on the mass budget and remain competitive.

In contrast, Table 9-16 shows the top 0.1% fuzzy NPT (fNPT) alternatives; incorporating nearly-efficient designs paints a much different picture of the relationship between stakeholders.

**Table 9-16: Highest 0.1% fNPT alternatives for each stakeholder, and highest maximin fNPT across stakeholders**

	<b>R</b>	<b>N</b>	<b>E</b>	<b>Maximin</b>
<b>0.1% fNPT</b>	0.6667	0.5556	0.3580	0.3580
<b># Alternatives</b>	432	2	54	9
<b># Planes</b>	5	5	5	5
<b># Satellites</b>	20	20	20	20
<b>Inclination</b>	<i>any</i>	53 deg	53 deg	53 deg
<b>Altitude</b>	<i>any</i>	800 km	800 km	800 km
<b>Antenna Area</b>	100 m <sup>2</sup>	100 m <sup>2</sup>	100 m <sup>2</sup>	100 m <sup>2</sup>
<b>Bandwidth</b>	2 GHz	1 or 2 GHz	<i>any</i>	2 GHz
<b>Peak Power</b>	10 or 20 kW	1.5 kW	1.5 kW	1.5 kW
<b>Comm Downlink</b>	<i>any</i>	No	Yes	Yes
<b>Tactical Downlink</b>	<i>any</i>	Yes	<i>any</i>	Yes
<b>Maneuverability</b>	<i>any</i>	Base	<i>any</i>	<i>any</i>
<b>Constellation Option</b>	<i>any</i>	None	<i>any</i>	<i>any</i>
<b>Mass Budget</b>	1	1	1	1
<b>Notes</b>	Major overlap between all sets Fuzziness allows expensive but consistent designs to dominate Maximin is a subset of E			

First, the fNPT scores are much higher than the NPT scores. R can find alternatives that are near-efficient across two-thirds of the context space, which is the maximum score possible when considering that the third of the contexts with slowest targets were discovered to be “impossible” missions in this tradespace. To achieve scores this high, the mass budgets for all of these alternatives is set to the full 100%: lowering the mass budget results in power systems (batteries / solar panels) that are undersized for the observation tasks of harder missions. Second, it is clear that this type of analysis lends itself to finding “gold plated” alternatives: very expensive designs that are consistently good for all stakeholders. Every design in Table 9-16 has the maximum number of satellites and the largest antenna size in addition to the full mass budget, and there are many similarities between the preferred alternatives of each stakeholder. The number of alternatives in each set is also quite high, reflecting the fact that these alternatives are in the top region of the tradespace where many expensive designs have equal utilities, especially for R.

The maximin set is the subset of E's set with design variables that he was indifferent towards (bandwidth / tactical downlink) set to values that the other stakeholders prefer (2 GHz / Yes), exactly mimicking the “easy compromises” performed in a single context during the original live negotiation. Now we have additional insight into those tradeoffs: they remain beneficial in the complete context space and can be used to allow the stakeholder most sensitive to changing contexts, E, to maximize his robustness while minimizing the impact on the value of R and N, who remain efficient in 48% and 44% of the context space, respectively.

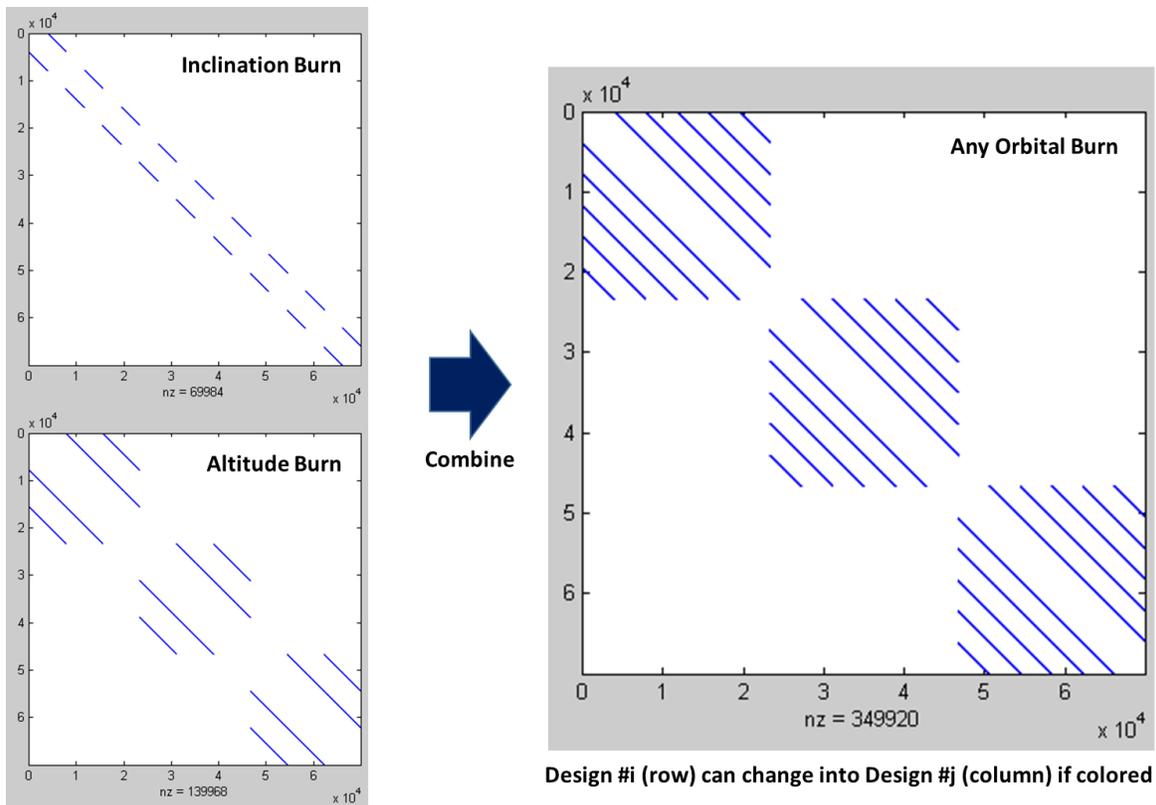
## **9.9 Changeability-enabled Robust Alternatives**

In addition to passively robust designs, some alternatives may be able to actively change in response to context shifts that negatively impact their value. This changeability therefore enables robustness and offers an alternative design paradigm for combating uncertainty. To find these changeability-enabled robust alternatives, we will calculate *effective* NPT and fuzzy NPT (eNPT / efNPT) using a changeability strategy according to the Valuation Approach for Strategic Changeability (VASC), an offshoot of Epoch-Era Analysis (Fitzgerald et al., 2012). In this case, we will use “maximize efficiency” – i.e. change to be as close to the Pareto front as possible – as a strategy, in order to provide a consistent analogy when comparing to the NPT results. These metrics grade alternatives not by their own performance, but their “effective” performance given the ability to change into other designs.

First, we must define the change mechanisms that allow designs in the tradespace to alter their design variables. We will consider the use of fuel burns for changing both altitude and inclination. Satellites carry on-board fuel that, though typically used for small corrections counteracting orbit decay, can be burned in order to change their defining orbit parameters. These orbit maneuvers are measured in  $\Delta V$ : the amount of velocity needed to change the orbit. For this case, the maneuverability design variable determines the amount of fuel carried by each satellite: the “base” setting for maneuverability corresponds to enough fuel for 545 m/s of  $\Delta V$ , and the higher settings for 1,090 and 2,180 m/s. More maneuverable designs are therefore able to make more orbit changes before needing to remain in a single orbit<sup>30</sup>. Figure 9-20 shows the accessibility matrices for the different change mechanisms, and the combined matrix for executing multiple fuel burns simultaneously. A given row of these plots corresponds to an initial design and the filled columns indicate the other designs it can change into by modifying only the targeted design variable. For example, there is one blue point on each row in the Inclination Burn matrix since the enumeration of the design space has only two inclinations. The combined matrix on the right indicates that each design can execute an orbital burn to five other designs.

---

<sup>30</sup> Or be tugged / refueled by an on-orbit servicing infrastructure. Such space maintenance missions have been proposed as use cases for future “space tug” satellites, including in the original Satellite Radar reports which considered these as additional change mechanisms (Rhodes et al., 2008, 2010).



**Figure 9-20: Combining the accessibility matrices of different types of orbital burns**

In principle, this analysis evaluates groups of designs that are connected via these change mechanisms, rather than individual designs. With these designs known, we can evaluate changeability-enabled robustness by using the *preferred* design in the set for any given calculation. Therefore, for Satellite Radar we calculate eNPT/efNPT on a satellite architecture that does not include a defined orbit altitude or inclination, instead using the “best” set of orbit parameters for each context with the understanding that the satellites will simply move into that orbit should the context arise. Calculating eNPT for Satellite Radar results in an underwhelming insight: the highest eNPT alternatives for each stakeholder are simply the highest NPT alternatives (again, without specifying an orbit type). Additionally, E is the only stakeholder who even uses the change mechanisms to achieve a higher eNPT than NPT, achieving a modest improvement from 0.1667 to 0.1914. Instead, we will focus on the 0.1% efNPT results, as shown in Table 9-17.

**Table 9-17: Highest 0.1% efNPT alternatives for each stakeholder, and highest maximin efNPT across stakeholders<sup>31</sup>**

	<b>R</b>	<b>N</b>	<b>E</b>	<b>Maximin</b>
<b>0.1% efNPT</b>	0.6667	0.6667	0.6049	0.5556
<b># Alternatives</b>	432	648	324	54
<b># Planes</b>	5	5	5	5
<b># Satellites</b>	20	20	20	20
<b>Antenna Area</b>	100 m <sup>2</sup>	100 m <sup>2</sup>	100 m <sup>2</sup>	100 m <sup>2</sup>
<b>Bandwidth</b>	2 GHz	1 or 2 GHz	<i>any</i>	2 GHz
<b>Peak Power</b>	10 or 20 kW	<i>any</i>	1.5 kW	1.5 kW
<b>Comm Downlink</b>	<i>any</i>	<i>any</i>	Yes	Yes
<b>Tactical Downlink</b>	<i>any</i>	Yes	<i>any</i>	Yes
<b>Maneuverability</b>	<i>any</i>	<i>any</i>	<i>any</i>	<i>any</i>
<b>Constellation Option</b>	<i>any</i>	<i>any</i>	<i>any</i>	<i>any</i>
<b>Mass Budget</b>	1	1	1	1
<b>Notes</b>	Same as fNPT: designs do not need to change	Less restrictive than fNPT set on power/comm/maneuverability	Same set as fNPT but without fixed altitude/inclination	Shared subset of N and E sets

The key variables of these designs are similar to those in the fNPT table: still very similar between the different stakeholders and featuring many satellites and large antennas. However, how each set responds to the potential use of changeability is different. R’s set of preferred alternatives is unchanged. Under his preferences, these designs gain nothing from changeability because they already had the maximum possible score (0.6667) without needing to change orbits. E has the largest performance increase, becoming efficient in ~20% more of the context space (0.3580 to 0.5679). Because his preferred alternatives are also the same as under fNPT, though now without a fixed altitude or inclination, we can conclude that changeability can effectively mitigate much of his sensitivity to context. Changeability drives a smaller improvement for N (0.5556 to 0.6667) but he hits the maximum efNPT while also diversifying the set by removing the fixed values for communication downlink, maneuverability, and constellation option.

<sup>31</sup> For readers wondering why efNPT is seemingly unaffected by many cost-driving design decisions (such as fuel mass, via the “maneuverability” variable): this is entirely a side-effect of the high-cost high-utility nature of all of these designs. As we saw earlier, the tradespaces for these stakeholders have “flat” tops – many designs (including those in this table) are tied with the highest achievable utility but vary in cost. In this circumstance, using a fuzzy Pareto set will include all of these tied alternatives, hence the indifference we see in this table. This is a weakness of fuzzy Pareto analysis when applied to tradespaces with this shape; however we will see in the next section that we can apply some additional screening to consider the impact of varying amounts of fuel.

Like E's set, the maximin set is the same as under fNPT but without the fixed altitude and inclination. Because of the increased diversity in N's preferred alternatives, the maximin set is now a subset of both N and E's sets – in fact, it is exactly the intersection of N and E's sets. This presents further evidence that N and E have very similar interests across the context space (and not just in the “baseline” context used in the correlation plots). Unlike with fNPT however, the efNPT maximin score does not match E's score. Though R's individual optimum efNPT is 0.6667, relying on high power (10 or 20 kW), he scores only 0.5556 in the shared E-N set, as the low power (1.5 kW) makes him vulnerable to failing his requirement on Number of Target Boxes when the context shifts to a harder mission. The 20% increase in worst-case efficiency across the context space for all stakeholder (0.3580 to 0.5556) is directly attributable to the availability of orbit-modifying change mechanisms and demonstrates the effectiveness of changeability at improving outcomes for *multiple* stakeholders just as it can for single stakeholders like E.

### **9.10 Leveraging Changeability to Satisfy Multiple Stakeholders at Low Cost**

The robustness analysis of the previous subsections has shown that, though no designs remain Pareto efficient across the entire context space, there are available alternatives that are near-efficient and satisfy all stakeholders for over one-third of the context space *without* changeability and half of the context space *with* changeability. However, the “best” Satellite Radar designs according to these TSE / EEA screening metrics are all very expensive: the cheapest alternative in the efNPT maximin set has an estimated lifecycle cost of over \$34B, driven by the selection of large antennas and many satellites<sup>32</sup>. Changeability has the potential to enable not only superior performance but lower cost, by moving away from “gold plated,” passively robust designs in favor of smaller designs that can change more frequently to cover their flaws. Since rising costs were cited as justification for the cancellation of real-world space radar programs, in addition to the US Department of Defense's institutional emphasis on minimizing costs, as referenced in both Analysis of Alternatives and the Better Buying Power initiative (Office of Aerospace Studies, 2013; Kendall, 2014), we can be confident that the search for low-cost Satellite Radar alternatives would be of primary interest to the stakeholders.

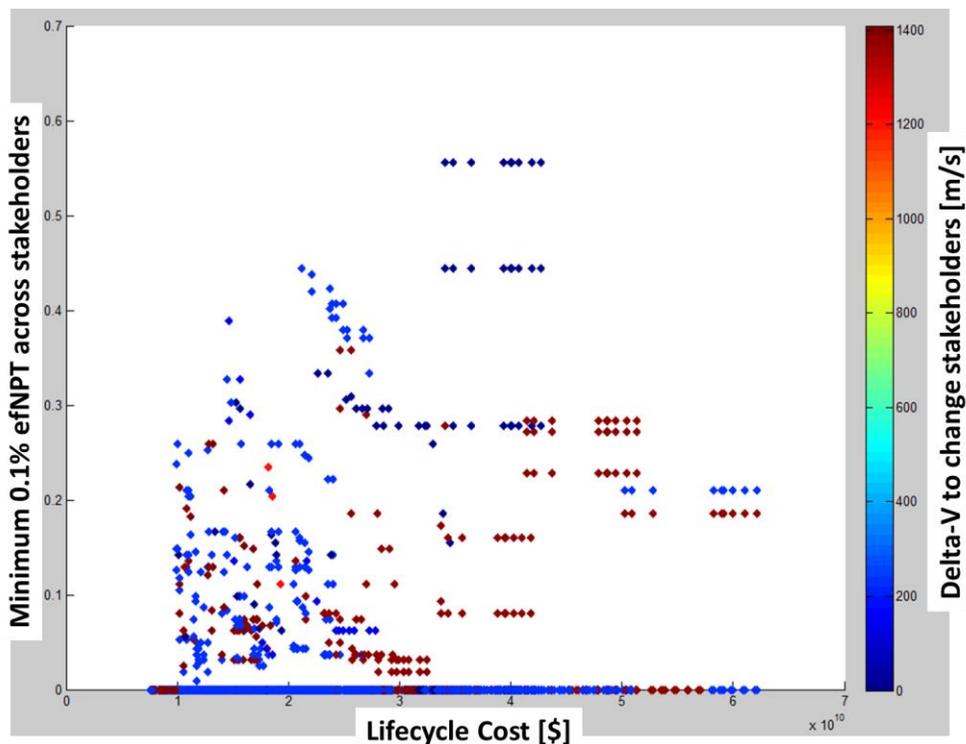
For Satellite Radar in particular, there is an opportunity to leverage changeability to temporarily optimize the system for a single stakeholder. Though multiple government agencies

---

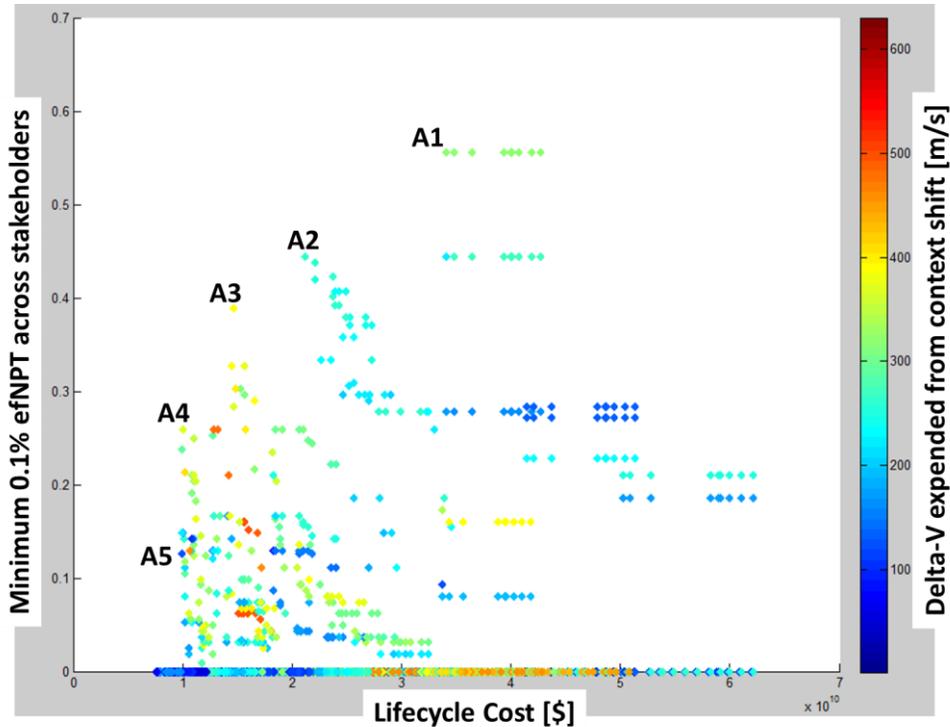
<sup>32</sup> This cost is clearly too large to ever be a feasible proposal: compare it against other very large unmanned space programs, such as GPS (estimated \$14B in 1995 dollars for operations through 2016; Pace et al., 1995) or the James Webb Space Telescope (\$8.8B, after considerable cost overruns, for a planned ten years of operation starting in 2018; National Aeronautics and Space Administration, 2015). This is a combination of the Satellite Radar program being undeniably large and ambitious, as well as this case's use of an unsophisticated parametric cost model, which does not accurately account for cost savings when replicating satellites for a constellation. Regardless of the source of the high costs, feasible lower-cost alternatives would be highly desirable solutions in the tradespace.

would be combining resources to build and share the resulting system, it not necessarily true that all the agencies will need to use the system *simultaneously*. If the stakeholders take turns using the system based on the changing priority of their targets of interest, it is possible that orbit burns could switch between local optima for each stakeholder where no single orbit would suffice.

To do this, we can deepen our search from just the *highest* maximin efNPT solutions to include the tradeoff between maximin efNPT and lifecycle cost. Figure 9-21 shows a scatterplot of these metrics, where each point is a physical satellite design (i.e., without a fixed orbit). Additionally, the points are colored by the average  $\Delta V$  expended to transition the satellite from the preferred orbit of one stakeholder to another. Figure 9-22 is the same, but colored by the average  $\Delta V$  expenditure to respond to a context shift. These two different types of fuel burn are not necessarily correlated: some satellite designs will be passively robust to context shifts but require different orbits for each stakeholder, while other designs may be highly sensitive to context but all stakeholders prefer the same orbits at the same time. To provide some perspective, the approximate cost of an altitude change in this tradespace is 150-350 m/s, while an inclination change is about 1,750 m/s. Five designs of interest are marked on the second plot, labeled as A1 through A5.



**Figure 9-21: Satellite Radar design groups, by cost and minimum 0.1% efNPT, colored by  $\Delta V$  cost to change stakeholders**



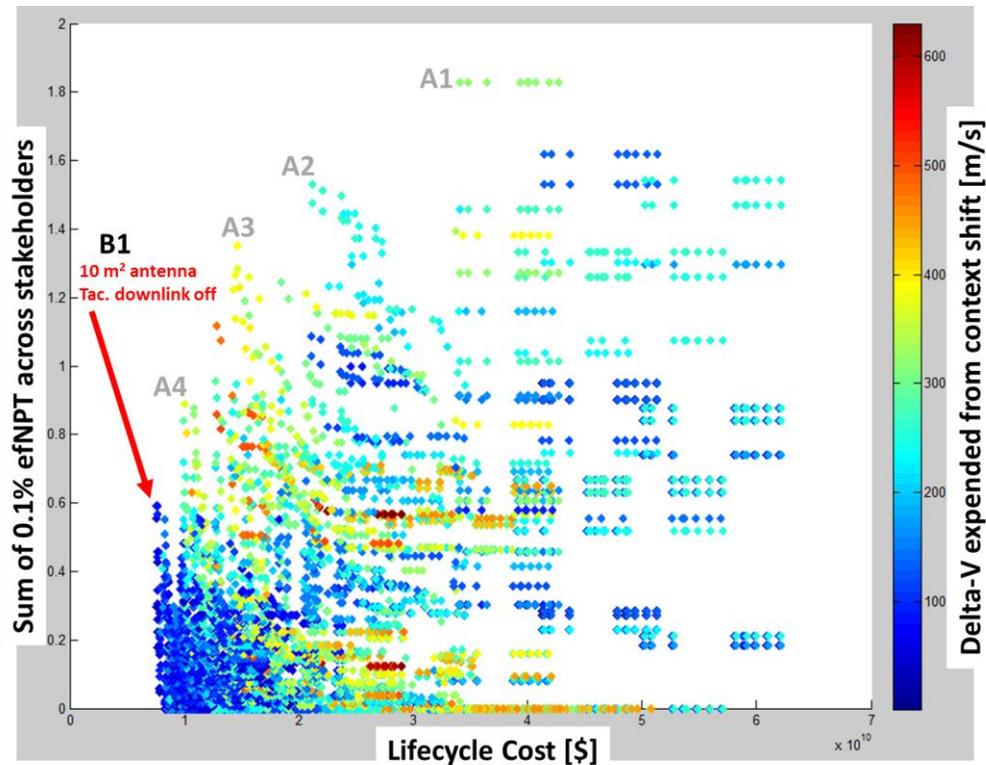
**Figure 9-22: Satellite Radar design groups, by cost and minimum 0.1% efNPT, colored by  $\Delta V$  cost expended in response to a context shift**

The details for designs A1-A5 (which, again, do not specify orbit parameters) are shown in Table 9-18. Design A1 and the points directly to its right are the maximin designs identified in the previous efNPT-only search. As cost is lowered, the corresponding worst efNPT between the three stakeholders is also lowered. The patterns in the design variables indicate a gradual trading-down of features, beginning with antenna size and gradually reducing the number of satellites and communications downlink. Additionally, some of these lower cost alternatives use mass budgets less than one: when attempting to lower costs, shaving weight is effective but naturally results in a lower efNPT due to some number of failed missions because of the undersized power system. The rate of  $\Delta V$  consumption amongst these alternatives is mostly stable around 225 m/s for a stakeholder change and 250 – 350 m/s for a context change. The notable exceptions are at the ends of the cost spectrum: all stakeholders share a preferred orbit for A1 (leading to zero  $\Delta V$  consumption on a change), and A5 changes the least in response to context shifts, requiring only 125 m/s on average.

Table 9-18: Designs of interest trading cost and minimum 0.1% efNPT

	A1	A2	A3	A4	A5
<b>Minimum 0.1% efNPT</b>	0.5556	0.4444	0.3889	0.2593	0.1265
<b>Lifecycle Cost</b>	\$34B	\$21B	\$15B	\$10B	\$9.9B
<b>Avg. Stakeholder change <math>\Delta V</math></b>	0	225 m/s	135 m/s	225 m/s	225 m/s
<b>Avg. Context change <math>\Delta V</math></b>	320 m/s	260 m/s	385 m/s	370 m/s	125 m/s
<b># Planes</b>	5	5	5	5	5
<b># Satellites</b>	20	20	10	5	5
<b>Antenna Area</b>	100 m <sup>2</sup>	40 m <sup>2</sup>	40 m <sup>2</sup>	40 m <sup>2</sup>	40 m <sup>2</sup>
<b>Bandwidth</b>	2 GHz	2 GHz	2 GHz	2 GHz	2 GHz
<b>Peak Power</b>	1.5 kW	10 kW	10 kW	1.5 kW	1.5 kW
<b>Comm Downlink</b>	Yes	Yes	Yes	Yes	No
<b>Tactical Downlink</b>	Yes	Yes	Yes	Yes	Yes
<b>Maneuverability</b>	Base	Base	Base	Base	Base
<b>Constellation Option</b>	None	None	None	None	None
<b>Mass Budget</b>	1	1	0.75	1	0.75
<b>Notes</b>	Part of the maximin set found previously	Trades down antenna size, up in power	Trades down number of satellites and mass budget	Trades down number of satellites and power, up in mass budget	Trades down comm downlink and mass budget

All alternatives with lifecycle cost under \$9.9B have a minimum efNPT of zero. As we previously saw, this was driven by N, who will not accept any designs with the smallest antenna area. It is possible that one of these cheap designs is a viable two-stakeholder solution. To quickly screen for a good choice in that area, we can adjust the plot to use the sum of each stakeholder’s efNPT rather than the minimum, as shown in Figure 9-23. The designs of interest from the previous plot remain dominant solutions despite appealing to a different, more utilitarian fairness criterion: a consequence of the net-positive value correlation between the stakeholders, who generally agree on what designs are “good.” However, now a new, lower cost alternative (B1) dominates A5. B1 sits at the bottom of the cost range near \$7.6B, similar to A5 but using a 10 m<sup>2</sup> antenna and no tactical downlink: two cost-reducing decisions not acceptable to N. Design B1 is therefore potentially of interest to stakeholders R and E only.



**Figure 9-23: Satellite Radar design groups, by cost and sum of 0.1% efNPTs, colored by  $\Delta V$  cost to change stakeholders**

All of these designs are potentially useful solutions for the three stakeholders. However, the rate of  $\Delta V$  consumption is noticeably high, particularly since the designs of interest all have only the base maneuver package. Design A5, the initial choice with the least fuel expenditures, will still consume its entire 545 m/s  $\Delta V$  budget in (on average) 4 to 5 context shifts, even without switching stakeholder control. Given the variety of observation targets of interest to the defense and intelligence communities, it is exceedingly unlikely that the mission context will change only five times over the system lifetime (10 years minimum). Even if the stakeholders accept the relatively minor additional cost of the larger maneuverability packages, increasing the  $\Delta V$  budget to 2,180 m/s, the system will still only be able to locally optimize for stakeholder and mission an average of 7 – 20 times depending on the expected transition costs of the chosen alternative. Once the  $\Delta V$  budget is consumed the system will exist in a static orbit until it is decommissioned<sup>33</sup>. If the reversion to a passive system occurs mid-life, it will pressure the initial design decision back towards the expensive, passively robust designs such as design A1. This raises an important question with regards to *rationing* the use of changeability in the system. Is it possible to use changeability to preserve the feasibility of the system across

<sup>33</sup> Which actually also costs a few hundred meters per second of  $\Delta V$  for deorbiting, but this is included in the base system design.

multiple stakeholder preferences for its entire lifetime, rather than using it to maximize its efficiency for a fraction of the lifetime?

The usage strategy for limited resources such as  $\Delta V$  is an additional dimension on which the participating stakeholders can craft an agreement. To evaluate the effects of limited fuel usage, we can adjust our changeability strategy from “maximize efficiency” to “prevent failure.” Under this strategy, a stakeholder will only burn fuel and change orbit if he controls the system and the current orbit fails to meet his mission requirements when a different orbit would. If so, the resulting fuel burn will use the least amount of  $\Delta V$  necessary to achieve mission success – defined as meeting all requirements and thus being “in” the tradespace. As a result, the system will execute fewer, smaller burns between context shifts while maintaining the same mission success rate as the previous strategy (though with worse utility). Figure 9-24 and Figure 9-25 show the resulting minimum 0.1% efNPT vs. cost scatterplots for the reduced burn strategy. As expected, the plots show significantly less  $\Delta V$  usage but worse performance, with efNPT only reaching as high as 35% instead of almost 60%. When only attempting to avoid failure,  $\Delta V$  consumption can easily be limited to 0 – 60 m/s per context shift, extending the duration of changeability in the lifecycle essentially indefinitely. Thus, a strategy that defaults to “prevent failure” will leave room for a limited number of discretionary decisions during operations to aggressively burn fuel to maximize performance for particularly important or difficult missions.

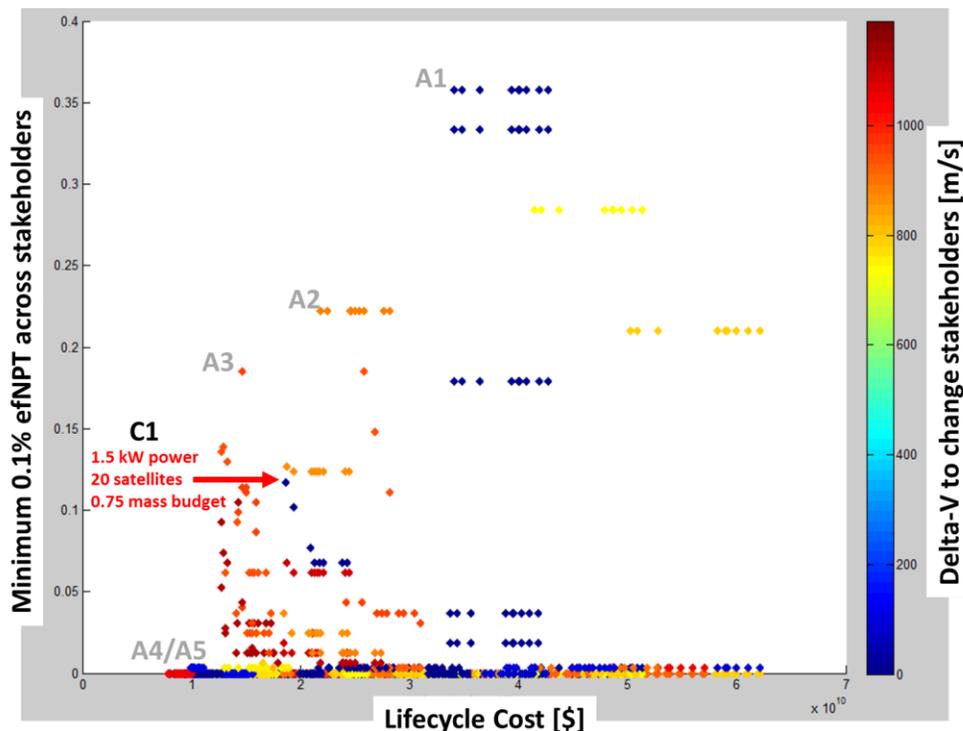
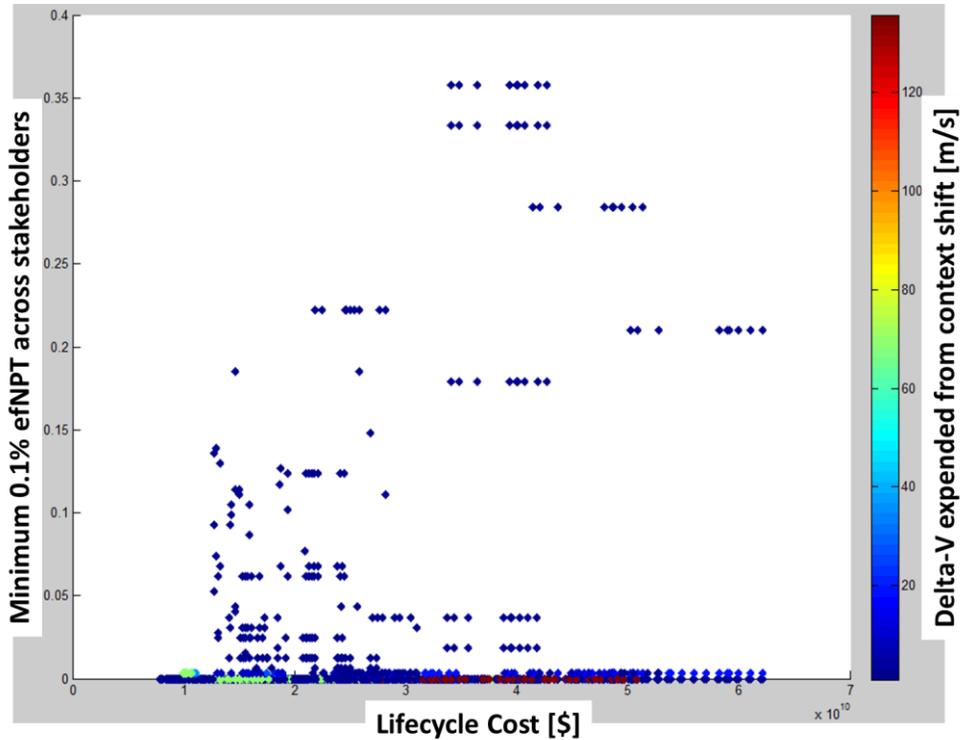


Figure 9-24: Satellite Radar design groups, by cost and minimum 0.1% efNPT, colored by  $\Delta V$  cost to change stakeholders, reduced burn strategy



**Figure 9-25: Satellite Radar design groups, by cost and minimum 0.1% efNPT, colored by  $\Delta V$  cost expended in response to a context shift, reduced burn strategy**

Under this new strategy, designs A1-A5 are still dominant, but A2 and A3 have noticeably higher  $\Delta V$  costs for stakeholder changes, implying that the orbits with fewest failed missions for each stakeholder occupy different inclinations (the more expensive burn type). An 800+ m/s stakeholder exchange cost is infeasible in the long term given the quantity of  $\Delta V$  on board these satellites. Though its efNPT scores are dominated by a handful of designs, design C1 appears to be the best intermediate cost solution for the “prevent failure” strategy, as it dominates among the solutions with low stakeholder change cost. However, C1 is a passively robust solution, relying on 20 satellites like A1 and A2 to maintain mission success without needing to burn fuel, while saving money with cutbacks on power and mass.

Table 9-19 and Table 9-20 collect the minimum 0.1% efNPT and average mission success rates for each of the designs of interest discussed above. Additionally, it includes two designs of interest, one passively robust and one flexible, identified from an earlier Satellite Radar TSE analysis, before the introduction of the three-stakeholder negotiation (Rhodes et al., 2008). The “2008 robust” design was selected based on its high performance in both imaging and tracking missions and its high NPT, while the “2008 flexible” design was selected for its high Filtered Outdegree (FOD), a quantification of changeability that measures the number of

available design changes under a specified feasible cost threshold<sup>34</sup> (Ross, 2006). The “2008 robust” alternative is Design #3435 from the original analysis (the shared Pareto front solution for R and E, with tactical downlink turned on for N), which is similar in design to A2 but with no communication downlink and a full mass budget. The “2008 flexible” alternative is similar to A3, but with maximum power, lower bandwidth, and more fuel.

**Table 9-19: Minimum 0.1% efNPT for designs of interest under different change strategies**

	<b>Cost</b>	<b>No Change (out of fuel)</b>	<b>Prevent Failure</b>	<b>Maximize Efficiency</b>
<b>A1</b>	\$34B	0.3580	0.3580	0.5556
<b>A2</b>	\$21B	0.2037	0.2222	0.4444
<b>A3</b>	\$15B	0.1049	0.1852	0.3889
<b>A4</b>	\$10B	0.0741	0.0031	0.2593
<b>A5</b>	\$9.9B	0.0463	0.0031	0.1265
<b>(R and E only) B1</b>	\$7.6B	0.2037	0.2037	0.2593
<b>C1</b>	\$19B	0.1358	0.1358	0.1667
<b>2008 Robust</b>	\$14B	0.0247	0.0988	0.1667
<b>2008 Flexible</b>	\$35B	0	0	0

---

<sup>34</sup> It is slightly unfair to compare these designs on equal footing with the new designs of interest, given that the revised problem formulation and models were not used for their original evaluation and selection. For example, the NPT for the “2008 robust” alternative was over 70% under the original model, and we can see in the table that it performs much worse with the new model, even when using the more forgiving efNPT metric. However, since we confirmed in section 9.7 that the revised model did not have a drastic impact on the relative ranking of different alternatives, we can at least compare these alternatives as long as we don’t judge the classic TSE approach too harshly for their inferior performance to the newly selected designs. Additionally, mission success rate is less sensitive to model changes than efNPT as it is a significantly broader set, and thus even less likely to present comparison issues.

**Table 9-20: Mission success rate (stakeholder average) for designs of interest under different change strategies**

	Cost	No Change (out of fuel)	Prevent Failure	Maximize Efficiency
<b>A1</b>	\$34B	0.6667	0.6667	0.6667
<b>A2</b>	\$21B	0.6214	0.6296	0.6296
<b>A3</b>	\$15B	0.4198	0.5679	0.5679
<b>A4</b>	\$10B	0.4321	0.5432	0.5432
<b>A5</b>	\$9.9B	0.2160	0.2716	0.2716
<b>(R and E only) B1</b>	\$7.6B	0.2963	0.2963	0.2963
<b>C1</b>	\$19B	0.4959	0.4959	0.4959
<b>2008 Robust</b>	\$14B	0.5103	0.5926	0.5926
<b>2008 Flexible</b>	\$35B	0.6049	0.6296	0.6296

Beyond the clear correlation between cost and performance in both metrics, there are two main patterns. First, executing fuel burns to “prevent failure” increases the mission success rate for most designs over the “no change” strategy – either by intent or from running out of fuel. Remembering that 0.6667 is the best score possible in this context space, only B1 and C1 have room for improvement that is not at all accessible via orbit changes. The mission success rate is the same for “prevent failure” and “maximize efficiency” for all designs as, naturally, moving towards the Pareto front will always require meeting all requirements. Second, the “maximize efficiency” strategy increases the efNPT performance of each alternative over “prevent failure” much more than “prevent failure” over “no change”, though we saw earlier that it comes at a much higher  $\Delta V$  cost<sup>35</sup>. These two trades – increasing mission success rate by burning (small) amounts of fuel, and increasing efficiency by burning (large) amounts of fuel – are shown in the bubble plots of Figure 9-26 and Figure 9-27. The size of the bubbles corresponds to the system cost<sup>36</sup>, and the color to the method used to identify that design (A/B/C/2008). The line with slope 1 corresponds to fuel burns not increasing mission success / efficiency. The distance of a point above that line indicates its potential value-added from additional fuel burns.

---

<sup>35</sup> Note that “prevent failure” actually *lowers* efNPT compared to “no change” for A4 and A5. This is counterintuitive at first, but simply indicates that there is a trade taking place between an orbit with a high success rate but low efficiency (preferred by “prevent failure”) and one with a lower success rate that lies closer to the Pareto front when it is valid (preferred by “no change”).

<sup>36</sup> The size of B1 has been scaled up to account for the fact that it is nominally a 2-stakeholder solution. This adjustment makes it slightly more expensive on a per-stakeholder basis than A4 or A5.

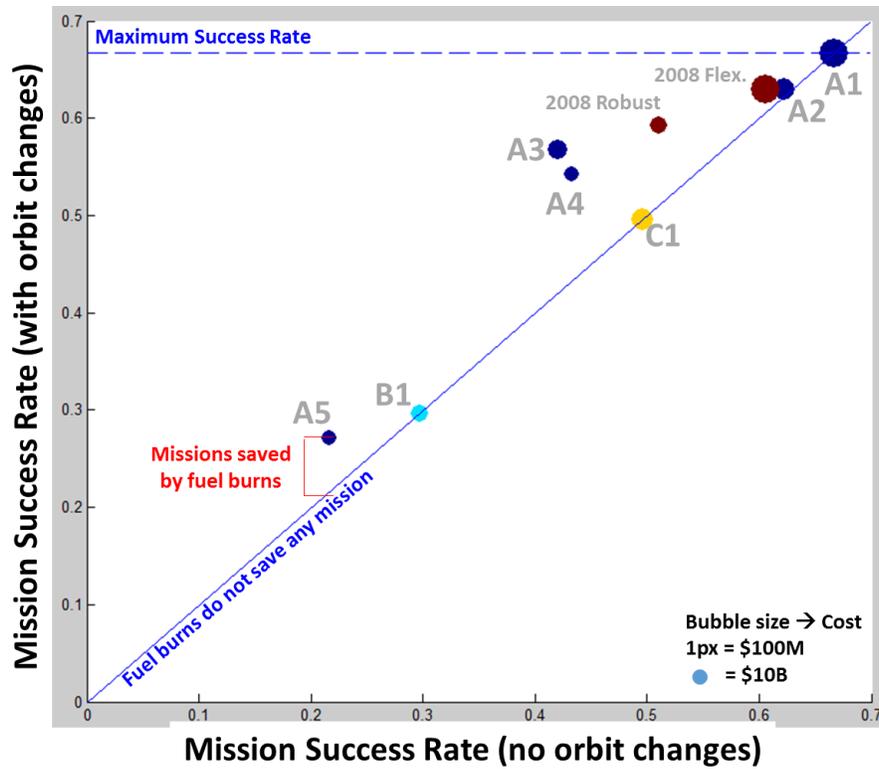


Figure 9-26: Mission success rate for designs of interest with and without orbit changes

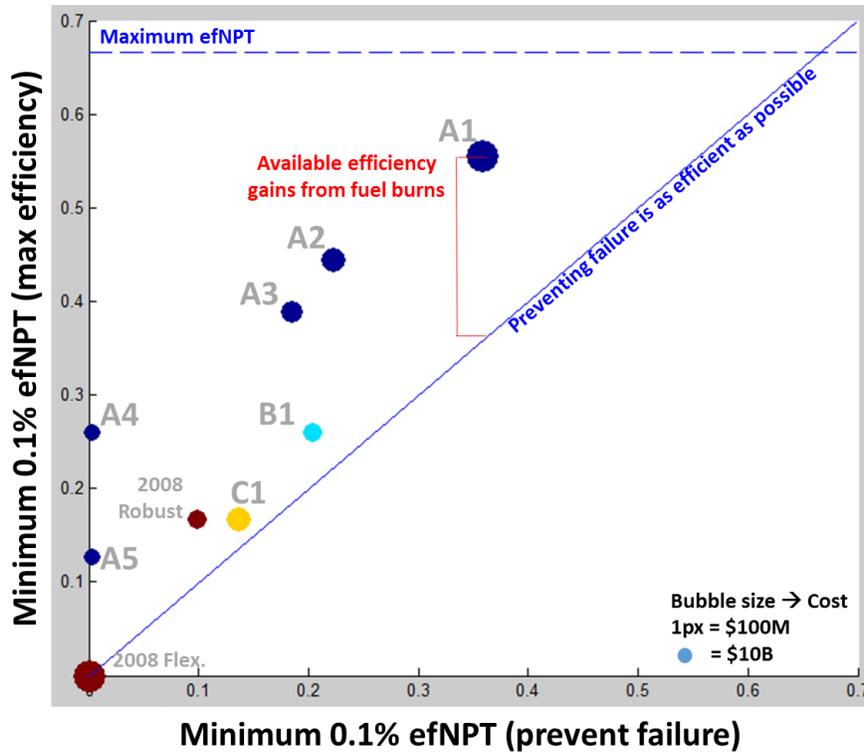


Figure 9-27: Minimum 0.1% efNPT for designs of interest, "prevent failure" vs. "maximize efficiency"

From these figures, we can gain some insight into the performance of these alternatives for a mixed strategy that emphasizes “prevent failure” but may occasionally need to “maximize efficiency”. In Figure 9-26, we can see that the handful of expensive, robust designs in the upper-right corner do not see significant improvements in mission success from orbit changes (A1, A2, 2008 Flexible), while the remaining designs above the line maintain a roughly constant 5-10% improvement in success rate with the manageable  $\Delta V$  consumption of the “prevent failure” mode. Notably, both B1 and C1 sit on the diagonal line, indicating that they do not complete any additional missions when allowed to change orbit despite their middling static performance. Expensive designs again perform better in Figure 9-27, with the exception of the 2008 flexible solution that is never efficient. In this plot, there is again a roughly constant 20-25% available improvement in efficiency across the context space, regardless of cost, that is captured fully by the A designs and less so by the others. We can conclude that the change mechanisms available to Satellite Radar offer a consistent but limited range of performance impact, such that inexpensive alternatives cannot “catch up” to the more expensive alternatives by changing orbit.

That said, the performance improvement derived from the freedom to change orbits is considerable despite the fact that it doesn’t increase as cost (and static performance) decrease. In a context space with many different targets, inclination can strongly impact coverage on targets in different locations, while altitude offers a tradeoff between revisit interval / field of regard and accuracy / resolution. Orbit changeability allows these tradeoffs to be made in response to an emergent target. For example, we know that stakeholder N needs the lowest altitude (800km) to meet his accuracy requirement, but fuel burns allow R and E to increase altitude in support of their preference on revisit interval when the target is large enough to see from a high orbit. The leverage provided by the change mechanisms makes the Pareto front “reachable” for a certain fraction of the context space: in this example, contexts with targets large enough to see from high orbit. That fraction is the maximum possible improvement in efNPT derivable from the change mechanisms, which is approached by designs A1-A4 based on our initial screening.

With cost as a major driver of value then, it appears that A4 may be the “best” solution amongst these designs of interest. It has the lowest per-stakeholder cost other than A5 but with a dramatic increase in mission success rate, driven in part by using its full mass budget. Our original concern over the rate of fuel consumption for A4 was mitigated by switching to the “prevent failure” strategy, which lowered the per-context-shift  $\Delta V$  cost to 67 m/s and the per-stakeholder-change  $\Delta V$  cost to 96 m/s. However, A4 maintains the option to reach maximum efficiency for just under half of the feasible targets. Though A4 never reaches the same level of performance as the “gold plated” solution (A1) with or without changeability, the flat improvements make it proportionally more attractive. Without orbit changes, A1 has a 3:2 ratio of mission success compared to A4 and a nearly 5:1 ratio for mission efficiency; those reduce to approximately 5:4 and 2:1, respectively, when fuel burns are included in the evaluation. A4 is also under one-third the cost of A1 at \$10B, putting it within the realm of large but successfully

funded and strategically valuable space programs such as GPS. This trade may be acceptable to all stakeholders, particularly if there is more emphasis on basic mission success than high-level performance. Though each stakeholder may individually prefer a different alternative (which would be a topic of discussion in a live negotiation, as we covered in previous sections), we can be confident that low-cost solutions that can meet the needs of all stakeholders do exist. This demonstrates the potential positive impact that evaluating changeability can have on a negotiation, by offering a creative source of mutual benefit beyond static design analysis.

## **9.11 Discussion**

Summarizing the results of the Satellite Radar case study using the key objectives outlines in the chapter introduction:

### **1. Demonstrate the use of MSTSE analysis applied to a technically complex system with many potential alternatives and relatively few stakeholders**

We used the recommendations for informal MSTSE to walk through an example analysis of Satellite Radar: a three-stakeholder system with 69,984 alternatives and 324 contexts (after our slight rescoping of the tradespace from the original analysis). The size of the tradespace, despite being much larger than those used in the experiment and interviews, did not present any notable difficulties for our MSTSE visualizations. Additionally, the nature of these structural elements was used to guide our subsequent analysis in useful directions, including an emphasis on understanding the pairwise relationships between stakeholders rather than repeating the favorite-alternatives driven analysis of the previous live negotiation session.

### **2. Revisit prior TSE studies and stakeholders in order to capture changes in insight or outlook prompted by the new techniques**

Arguably the most important result for the purposes of this research, this case study served to highlight some important features of the hypothetical Satellite Radar program specific to its multi-stakeholder nature that were not captured in the original Epoch-Era Analysis implementation of TSE. Specifically, we identified:

- N's low tradespace yield is impactful mostly due to its elimination of low cost alternatives, which the other stakeholders may still want
- N and E have a naturally high correlation in value, though they are both still positively aligned with R as well.
- Each pairwise combination of stakeholders has a set of shared preferred designs that share consistent design variables. These sets offer different trades away from the three-stakeholder joint set, with R and N favoring high-cost trades (mostly with larger antennas) and R and E favoring low-cost trades (mostly with fewer

satellites). N and E have the smallest pairwise set because their interests are the most correlated (and thus mostly contained within the joint set).

- R and N have a small set of nearly-efficient shared alternatives but all other alternatives have a large drop-off for one or the other, contributing to their lower utility correlation.
- Preferences on latency (imaging for E and tracking for N) drive the most differences between stakeholders. These “pain points” arise because key latency-reducing design choices, particularly the inclusion of communications and tactical downlinks, increase cost without supporting other performance attributes.
- After adjusting the model and verifying that the same writ-large relationships were present, we leveraged the cross-context comparison of EEA to find both passive and changeability-enabled robust solutions. Achieving near-optimal performance across the context space for all stakeholders requires an almost-certainly infeasible cost both in dollars and  $\Delta V$ . However, there are available and realistically-priced solutions that use small fuel burns to achieve nearly-equal rates of mission success (though with lower performance above those requirements).

In collecting these results, we have collected supporting evidence for the performance validity of MSTSE, in terms of its value-added over classic TSE. The new visualizations proved effective at quickly highlighting important relationships between the stakeholders’ value functions. Key areas of interest in the tradespace were also identifiable despite the large design space and the fact that all alternatives were located in Quadrant 1 as a result of the “do-nothing” BATNA for Satellite Radar. In the process, we learned more about the types of trades each stakeholder would prefer to make individually and in pairs, supporting the possible selection of alternatives for subgroups of stakeholders.

### **3. Identify the potential for -ilities to create opportunities for additional value by allowing transitions between designs that favor individual stakeholders**

The analysis of changeability, specifically targeting the use of fuel burns to support temporarily optimizing the system for a given in-control stakeholder, provided a good example of the potential for -ilities to support negotiations by increasing mutual benefit of less expensive (and less passively robust) alternatives. Without the use of changeability, the Satellite Radar system is unable to exceed a mission success rate of ~50% for less than ~\$20B. Utilizing fuel burns to avoid mission failure while exchanging control of the system between the three stakeholders can support a ~55% success rate effectively for the complete expected satellite lifetime while halving the necessary costs to nearly \$10B.

The exploration of additional change options represents a potential means of further improvement to the hypothetical Satellite Radar system. We identified that the limited amount

of on-board fuel presents a barrier to aggressively changing orbit in pursuit of optimal performance (as opposed to merely mission success) for every context shift and stakeholder change, limited to no more than 20 of such burns on average over the lifetime. Our method of mitigating this concern was to focus on a more conservative change strategy that only burned fuel to avoid mission failure. However, other potential change mechanisms have been theorized for such a purpose. The presence of on-orbit infrastructure could serve to enable more extensive orbit maneuvers. For-hire orbital transfer vehicles (OTVs) often called “space tugs”, designed to push other satellites to new orbits, could provide a means of changing orbits without the need for on-board fuel, extending the time period for which Satellite Radar could reap the benefits of changeability (Collins et al., 2001; McManus and Schuman, 2003). Instead, such maneuvers would cost money (for hiring the tug) and time (rather than an effectively instantaneous fuel burn, a given satellite may need to wait for a tug to be in the area). Alternatively, on-orbit refueling via OTV or orbiting fuel depots could allow Satellite Radar to execute fuel burns indefinitely by resupplying fuel over time (Whelan et al., 2000; Chandler et al., 2007). Should such enablers become available in the future, they could be evaluated by MSTSE analogously to the mechanisms discussed in this case. Satellite Radar could plausibly see a significant increase in *potential* value, if the cost of such services were affordable enough for regular use.

Changeability also can enter the Satellite Radar program in other ways. No consideration was given in this case study to the possibility for either staged deployment or heterogeneous constellations, as the original Satellite Radar models did not incorporate features to evaluate these possibilities. The diversity inherent in heterogeneous constellations, consisting of satellites with different components and/or orbit parameters, has the potential to mitigate target uncertainty by providing partial coverage of a wider range of targets. Starting with a homogenous constellation, fuel burns could be deployed on a fraction of the member satellites in order to temporarily capture the benefits of heterogeneity. For example, if a high-latitude target arises that receives insufficient coverage from the current 10 satellite and 53 degree inclination constellation, five satellites could make the (expensive) inclination burn to 67 degrees if they were enough to achieve the desired coverage, saving fuel compared to changing inclination for all ten satellites. Staged deployment could follow a similar process, launching a small number of satellites to a desired high-value orbit and reserving the option to launch additional satellites to potentially different orbits. Lowering the cost of staging by pre-purchasing additional parts or whole satellites is the nominal purpose behind the inclusion of the “constellation option” design variable, which did not see much use in either this case or the original analysis because the orbit module was not sophisticated enough to truly exploit it. Future improvements to the model could target these areas in order to measure other potential sources of value from changeability.

### **Concluding thoughts:**

One structural feature of the Satellite Radar problem that we did not cover during our exploration was the divisibility of cost and the potential for an unequal distribution of those costs

between the three stakeholders. Analysis was performed on this possibility by incorporating an additional “change mechanism” into the cross-context analysis, corresponding to an imaginary option to add or remove satellites from the constellation. Of course, a functional satellite would not be removed from orbit under any normal circumstances. However, it is possible that within one group of designs (with all design variables fixed except for orbit parameters and number of satellites), one or more stakeholders could prefer the cost savings of fewer satellites while others could prefer the increased performance of many satellites. In such a situation the larger constellation could be built to satisfy all stakeholders, but it would likely be considered “fair” to have the former stakeholder(s) pay a smaller fraction of the total cost. Unfortunately, this analysis yielded no notable results for Satellite Radar: none of these 3-change-mechanism design groups had superior efNPT to the 2-mechanism groups found in the previously discussed analysis, suggesting that all stakeholders in this case agree on the best number of satellites given a fixed antenna size, power, etc.

This case study also brought to light a number of potential challenges for MSTSE and opportunities for further research. For example, the frequent reoccurrence or prominence of “gold plated” designs in the various visualizations and analyses we performed demonstrates a challenge of anchoring and reference points that has not been addressed by MSTSE to this point. By their nature, “gold plated” designs are good at everything the system has been asked to do and thus they receive high marks across different value metrics. This includes across the different individual stakeholder value functions and also from different cross-contextual metrics (as we saw with design A1 having the highest value for efNPT and mission success under multiple changeability execution strategies). This type of persistent reinforcement may be partially responsible for the observed bias towards this type of solution in multi-stakeholder engineering problems. Future research could attempt to isolate this phenomenon and test its impact on anchoring MSTSE participants’ reference points onto those designs. Balancing the need to accurately convey the high performance of such solutions (given that they are, in the end, excellent but expensive alternatives) without “normalizing” the idea of paying such a high cost would support exploration of a wider range of the tradespace.

The MSTSE visualizations developed in this research target single-context tradespace analysis; when we began analysis across multiple contexts, we mostly returned to using basic scatterplots of cross-context data as well as tables for describing designs and performance trades. From a research perspective, it was understandable to “start with the basics,” but the importance of uncertainty in engineering design mandates the need for MSTSE to specifically support analysis of uncertainty with dedicated visualizations. Future MSTSE research should consider this a priority. The bubble plot used to visualize the impact of varying fuel burn strategies was an interesting idea that would benefit from the same testing and refinement with practitioner feedback that the other visualizations received. Additionally, the cross-context analysis we performed largely relied on screening alternatives for all stakeholders using a fairness criterion (using the worst-case stakeholder via a maximin solution, except when screening for two-

stakeholder solutions when we used the sum across stakeholders). This type of uncertainty analysis should also support single-stakeholder visualizations augmented with information on the other stakeholders, similar to the negotiation tradespace. The bubble plot can potentially be adjusted for this purpose, for example by changing the axes to single-stakeholder mission success (rather than average) and including vectors corresponding to the ranges of success for the other stakeholders. A mockup of such a plot for three stakeholders is shown in Figure 9-28. The red and green vectors point towards the location of that bubble on the other stakeholders' axes: small vectors indicate close consensus amongst stakeholders about the value of that design and vectors pointing down/left indicate designs that are better for Stakeholder 1 than the others (and vice-versa for up/right). However, a plot like this is only manageable for small numbers of alternatives before the vector lines clutter the graph to the point of unreadability.



**Figure 9-28: Mockup of a "Stakeholder 1" bubble plot, with vectors indicating relative position of Stakeholders 2 (red) and 3 (green)**

Overall, the application of MSTSE has contributed new knowledge to this case study in Satellite Radar, demonstrating the nature of the relationships between the stakeholders and the impacts of changeability on the value of the group. At some level, this can be attributed to macro framing: the recent history of Satellite Radar has been treated mostly as a design exercise that was unable to solve a strictly technical problem – balancing tracking and imaging requirements – within a reasonable budget. When we view the problem through a new framing, one that emphasizes the role of negotiation between stakeholders, we learn more about how the differences between the stakeholders manifest in the value space and how that affects the types of solutions that are desirable for the group.



## 10 Case Study – Northeast Corridor

This chapter presents a case study on the Northeast Corridor (NEC): the high-traffic transportation system amongst destinations roughly between Washington DC and Boston. The case focuses on the potential for investment in high-speed rail (HSR) infrastructure within the NEC, which would drive many potential benefits (such as reduced travel time and pollution) but with a significant up-front development cost. In contrast to Satellite Radar, the primary challenge for any HSR project in the NEC is the large number of involved stakeholders. Because the NEC crosses many jurisdictions, the ability of any HSR proposal to satisfy all potential vetoes and build considerable constituent support is paramount, making it a highly social multi-stakeholder problem despite its relatively straightforward technical implementation – HSR was first constructed in Japan in 1964 and is now in use around the world. The difficulty introduced by the presence of many stakeholders is compounded by the small number of feasible HSR alternatives, driven by the considerable number of physical and organizational constraints imposed by the system’s large footprint.

### **The NEC case will be approached with the following key objectives:**

1. Demonstrate the use of MSTSE analysis applied to a socially complex system with few feasible alternatives but many stakeholders
2. Compare and contrast the emergent coalitions resulting from MSTSE analysis with prior NEC stakeholder analysis
3. Identify key controllable design parameters that can be leveraged to adjust promising alternatives in order to create value and attract support from additional stakeholders

### **10.1 Northeast Corridor Overview**

The Northeast Corridor (NEC) is the transportation system servicing both passenger and freight traffic in the most densely populated area of the United States stretching from Boston to Washington DC. Given the long history of the area and its span across multiple state jurisdictions, the network of transportation modes is a hodgepodge of originally-unconnected and now-aging systems. In particular, the rail network has continued to suffer performance degradation under the strain of persistent under-maintenance and a constantly increasing local population. Establishing an improved, unified vision for the rail system, possibly including the addition of high-speed rail, has the potential to revitalize the business of intercity rail travel and drive considerable economic benefits to the region, including productivity and environmental improvements. However, the large cost of such a project and the broad resulting impacts result in a large set of affected stakeholders, creating a challenging problem of both system design and interest-based negotiation.

This case study was performed in collaboration with the MIT Regional Transportation Planning and High-Speed Rail Research Group<sup>37</sup>, which has previously performed extensive analysis of the NEC for their various sponsors, including the Institution for Transportation Policy Studies (ITPS) and the East Japan Railway Company. The supporting documents for those research partnerships (Sussman, 2012; Sussman et al., 2015) are considered the foundation of this case study and the Benefit-Cost Analysis spreadsheet detailed within is used here as the basis for the evaluative model of NEC. Readers are encouraged to refer to those sources for further detail on their original analysis, which utilized the CLIOS Process as a central framework and supplemented it with a combination of supporting techniques including stakeholder analysis. Additionally, much of the financial data in those reports is sourced from the UPenn School of Design’s studio projects on high-speed rail, which may also be of interest to the reader (PennDesign, 2011, 2012).

The following sections will detail the creation of an MSTSE-compatible problem formulation from the existing data of those reports and the analysis of the resulting tradespace.

## **10.2 Creating the NEC Tradespace**

In order to adapt the prior NEC research from the CLIOS Process into a problem formulation capable of directly supporting MSTSE, it was helpful to go through the existing models and data and categorize them according to a tradespace framework – thereby structuring the tradespace and identifying any gaps that must be filled. Figure 10-1 illustrates a model-centric decision making framework designed to support TSE (Ross, Rhodes, and Fitzgerald, 2015). The “Decision (Problem)”, e.g. what to do in the NEC, has already been discussed, but in order to proceed with analysis and presumably reach the “Decision (Solution)” conclusion of this framework there must be a clearly defined design space, epoch space, and model set. These will be described here.

---

<sup>37</sup> For more information: <http://web.mit.edu/hsr-group/>

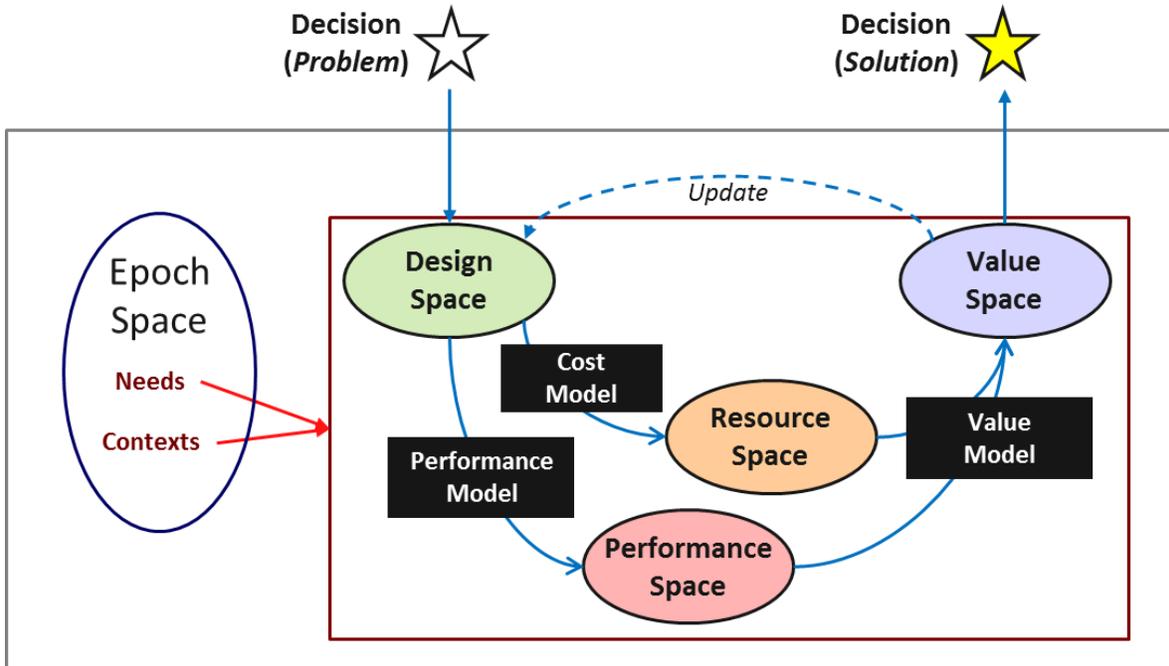
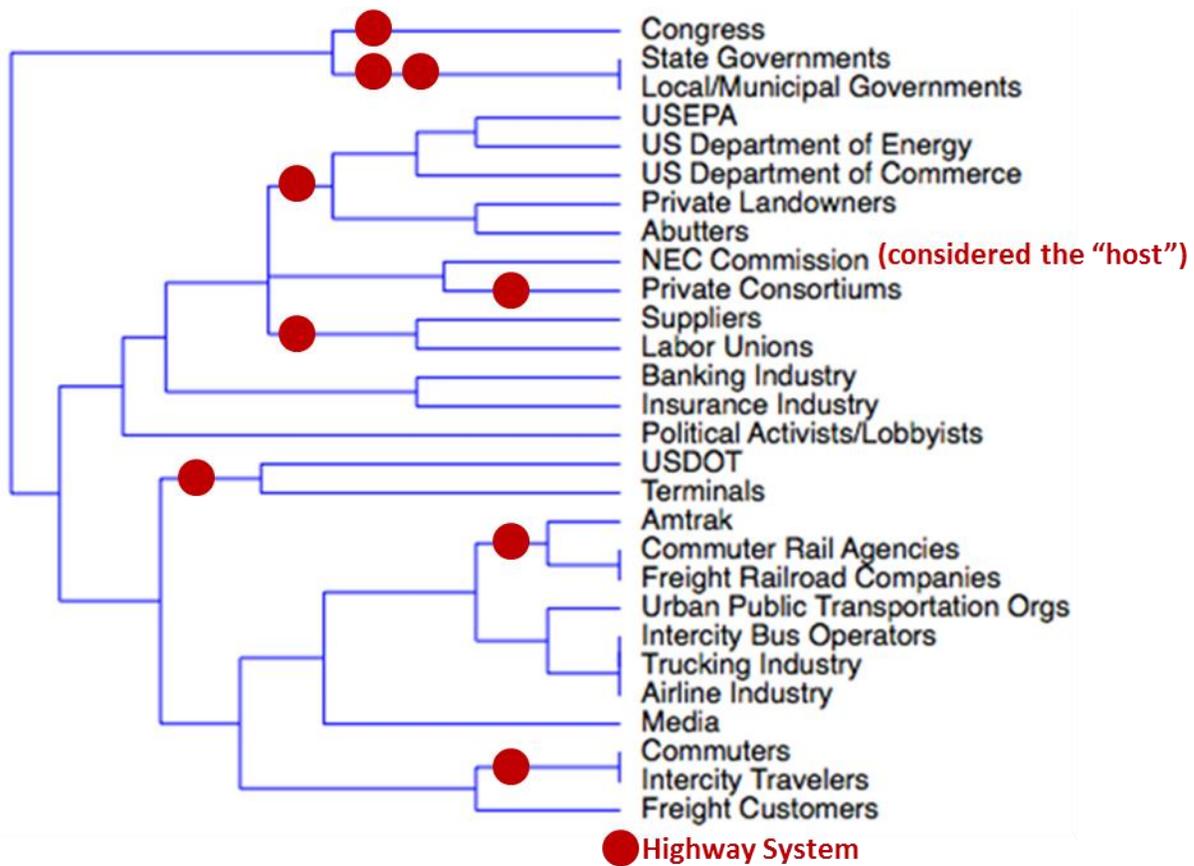


Figure 10-1: A generic model-centric decision making framework, capable of supporting TSE

### 10.2.1 Needs and Value Models

Our primary interest in this case is for its negotiation features; therefore the variety of stakeholders (and their needs) in the problem is of maximum importance. Because the inherited analysis was structured using the CLIOS Process [refer to section 8.1 for details], the stakeholders are documented in the institutional sphere. In this case, 28 stakeholders were identified. Figure 10-2 shows a version of an interest-based clustering analysis of the 28 stakeholders from the 2015 report, annotated with red dots to show the selection of 10 stakeholders for MSTSE.



**Figure 10-2: NEC Stakeholder clustering (Sussman et al., 2015; Moody, 2016), with additional annotations in red indicating the chosen MSTSE stakeholders**

The selection of which stakeholders to carry into MSTSE and which to discard was performed using the following criteria:

- *Aggregation.* Many of the red dots are located some distance up the clustering tree to the left, indicating that they consist of more than one of the original 28 stakeholders, but by virtue of sharing similar interests are easily grouped together. For example, Amtrak has been combined with the other rail agencies that use the tracks, both commuter and freight. These decisions were made with respect to the level of detail available in the models, which take a top-down perspective of the system. Without a higher level of detail, it becomes difficult to meaningfully quantify the differences in the needs for stakeholders who are clustered closely. Thus, we can choose to consider them as a single interest group when working in this low-detail, early-concept environment.
- *Degenerate.* Some stakeholders have interests that are functionally degenerate, which in this situation means that they do not present the opportunity for meaningful or interesting tradeoffs. For example, the banking industry will collect interest on any loans, which means they want to loan money with no specific reservations (again, at least for this level

of detail). They do not add depth to the negotiation and will not add depth to the goals of this case study; therefore they are not included.

- *Competitive.* The intercity bus operators, trucking industry, and airline industry are stakeholders in the NEC by virtue of being direct competitors to the services the rail network offers. Because MSTSE has been scoped to support cooperative negotiation and the Full, Open, and Truthful Exchange principle, it would be unwise to invite these stakeholders to participate in a private, closed-door meeting where they are incentivized to obstruct any potential agreements among the other stakeholders. By conducting MSTSE with only cooperative stakeholders, it is possible to build momentum behind a particular proposal before announcing it publicly, thereby weakening the bargaining position of the competitive stakeholders. Negotiations can include preparations for dealing with competitive stakeholders after they become involved in the next stage of project development. In order to capture the concept of a stakeholder with a direct interest in another mode of travel while still being cooperative, the Highway System was added to the list of NEC stakeholders.
- *Simplification.* To some extent, all of these decisions to reduce to 10 stakeholders are simplifications of the larger problem. More specifically, some stakeholders, such as political lobbyists, have motivations too complex to be captured at the available level of detail and were left out for that reason. This also includes the NEC commission: a group comprised of representatives of other stakeholders in the institutional sphere (such as Amtrak and the Department of Transportation) tasked with managing and planning the NEC infrastructure. As a collection of different interests already represented by other stakeholders, it is challenging to ascribe a single set of needs to the NEC Commission that avoids redundancy. In reality, the NEC Commission is more likely to be a driving force or “host” for negotiations and activities such as MSTSE as it already occupies a role at the interface of the various other stakeholders in the NEC.

Note that there are two stakeholders marked under the state and local governments branch of the tree. It was deemed necessary to divide the states into two categories: northern states (New York and north) and southern states (New Jersey and south). This is because, as will be discussed in the next section, the potential alternatives available for the design of the NEC rail system include choices with service level differences between these two areas. Thus their interests, in both benefits and costs, must be kept separate in order to capture the inter-stakeholder tension inherent in those alternatives

With the ten stakeholders chosen, the value model was created by assigning to each (1) a benefit function using a Keeney-Raiffa multi-attribute utility function (Keeney and Raiffa, 1993) and (2) a cost attribute representative of their interests. These value functions were created in

collaboration with the researchers of the MIT High-Speed Rail Group<sup>38</sup>. A list of the ten stakeholders and a summary of the attributes in their benefit and cost functions is shown in Table 10-1; a complete description of the attribute definitions, basis functions, and swing weights is available in Appendix B. They are not included in the main body of text because the analysis of this case study does not focus on the tradeoffs of individual attributes and thus they are not required knowledge to understand the remainder of the case. Note that “quality of service” is a superset of three attributes related to the effectiveness of the passenger transport on the NEC, including on-time performance, safety, and time savings (of HSR, compared to the current conventional rail system).

**Table 10-1: Summary of NEC stakeholder value models**

<b>Stakeholder</b>	<b>Benefit Function</b>	<b>Cost Function</b>
<b>Dept. of Transportation (USDOT)</b>	Quality of service Road congestion Emissions	Public funding
<b>Amtrak and rail agencies</b>	Discounted financial returns Quality of service	Private funding
<b>Congress</b>	Economic returns Discounted financial returns	Public funding
<b>Northern corridor states</b>	Economic returns Passengers (North) Quality of service (North)	North State funding
<b>Southern corridor states</b>	Economic returns Passengers (South) Quality of service (South)	South State funding
<b>EPA and landowners</b>	Emissions Environmental mitigation	Land use
<b>Private consortiums</b>	Private financial returns Payback period	Private funding
<b>Suppliers and labor unions</b>	Construction cost Duration of construction	<i>(none)</i>
<b>Highway System</b>	Road congestion	Diversion
<b>Travelers</b>	Quality of service	Fares

---

<sup>38</sup> Of course, ideally these would be elicited from the stakeholders themselves, but in the absence of access to those stakeholders (as is the case here, for this research) subject matter experts are the next best source for value estimates.

## 10.2.2 Design Space

The design space under consideration for this case consists of the six “bundles of strategic alternatives” identified as feasible potential solutions in the execution of the CLIOS Process. These bundles were the result of a combination of expert opinion and principles of factorial designs in order to provide a representative span of the complete realm of possible NEC implementation alternatives. Table 10-2 shows the six bundles with their associated design variable levels, grouped into the CLIOS categories of physical configuration, organizational structure, and funding mechanism.

**Table 10-2: NEC design space, comprised of six bundles (Sussman et al., 2015)**

<b>Bundle ID#</b>	<b>Physical Configuration</b>	<b>Organizational Structure</b>	<b>Funding Mechanism</b>
<b>1</b>	Incremental HSR Existing Alignment Shared Track	Vertically Integrated Amtrak	Public Infrastructure Public Operations
<b>2</b>	Incremental HSR Existing Alignment Shared Track	Vertically Separated Amtrak	Public Infrastructure PPP Operations
<b>3</b>	Piecewise International Quality HSR New Alignment Dedicated Track	Vertically Integrated Non-Amtrak Single Op	Public Infrastructure Private Operations
<b>4</b>	Piecewise International Quality HSR Existing Alignment Shared Track	Vertically Separated Competing Operators	PPP Infrastructure PPP Operations
<b>5</b>	All-over International Quality HSR New Alignment Shared Track	Vertically Integrated Amtrak	PPP Infrastructure Public Operations
<b>6</b>	All-over International Quality HSR New Alignment Dedicated Track	Vertically Separated Competing Operators	PPP Infrastructure Private Operations

Readers can look to the source references for exact detail on the meanings of each design variable. In terms of the direct impact on the tradespace, the two most important variables are the HSR type (incremental, piecewise international, or all-over international) and the funding mechanism. The HSR type determines if and where HSR is implemented in the NEC. *Incremental* HSR involves no new tracks or trains, but rather just State of Good Repair (SOGR) improvements to the existing infrastructure, allowing for faster travel and opening up potential future adoption of HSR. *Piecewise* HSR involves adding true international quality high-speed train service between Boston and New York, while *all-over* HSR offers that service throughout the NEC, from Boston to Washington DC. Note that, in accordance with factorial design of experiments principles, there is two of each type of HSR implementation. This variable has the

most impact on the total cost of the system and its resulting performance attributes, as the upgrade from SOGR to HSR is a costly but beneficial change. The funding mechanism can be set to public, private, or PPP (public private partnership) for both infrastructure and operations. These variables impact what fraction of the total cost of the NEC system each stakeholder has to pay and are the main differentiator in the tradespace between bundles of the same HSR type. The other variables have smaller impacts on various performance attributes of interest to subsets of the stakeholders.

The small design space described here is a defining feature of this tradespace. Six alternatives are vastly fewer than most applications of TSE, which presents challenges in the effectiveness of many common types of tradespace analysis. However, this is a realistic possibility entering an engineering negotiation, so it is important to evaluate MSTSE's ability to deliver useful insights in a problem of this scale.

### **10.2.3 Evaluative Models (Performance and Cost)**

The core evaluative model used for this case is the Benefit-Cost Analysis spreadsheet created by the MIT High-Speed Rail Group for the purpose of assessing the financial and economic impacts of the NEC bundles. This model, building on the work of PennDesign, Amtrak, and the NEC Commission, uses a 55-year time horizon (from 2010 to 2065) with associated cost and demand projections in order to quantify the performance of the system under the different combinations of physical, organizational, and financial design variables. It is from this model that most of the value-driving attributes are calculated for use in the tradespace.

Additionally, a few, very simple, supplementary models were created and used to calculate specific attributes identified by the stakeholder needs that were not covered by the BCA spreadsheet. These include:

- *Additional Post-Processing.* BCA typically aggregates benefit and cost for the purpose of analytical simplicity. Because a benefit-cost tradespace requires that these values be separated, additional post-processing of the BCA spreadsheet was necessary to separate the time-discounted benefits and costs of the different bundles. Additional financial metrics such as payback period (number of years before revenues recoup capital expenditures) and private financial returns (the fraction of returns assigned to the percentage of private investment) were also added.
- *Cost Division Model.* The BCA spreadsheet divided the costs of funding into only two sources: public (meaning federal) and "Private and State". In order to accurately assign costs to stakeholders in the tradespace, a further cost division model for the second pool of funds was necessary. This simple assignment model used the parameters defining each bundle, with particular emphasis on the designated funding mechanism, to directly set costs between the private, northern state, and southern state pools. Table 10-3 shows the division and reasoning for each bundle.

**Table 10-3: Cost division for Private+State pool**

<b>Bundle ID#</b>	<b>Division</b>	<b>Reasoning</b>
<b>1</b>	<i>(none)</i>	No funding in the pool
<b>2</b>	States split evenly, no private	Private funding will not likely extend to non-HSR options as there is lower growth potential
<b>3</b>	No states, all private	Bundle has an “operations only” assignment to private funding with a private operator, matching the cost of rolling stock
<b>4</b>	Northern states and private split evenly, no southern states	Even split for the PPP solution, but southern states will not pay for an HSR solution that does not serve their constituency
<b>5</b>	Private matches Bundle 3, states split remainder	Private follows same logic as Bundle 3, states cover the remainder evenly
<b>6</b>	Private matches Bundle 4, states split remainder	Private follows same logic as Bundle 4, states cover the remainder evenly

- *Safety Model.* Safety is an important metric for capturing the perception of rail travel by consumers. The BCA spreadsheet accounts for safety indirectly by considering “lives saved” due to diversion from road travel, which is more dangerous, and including them in the economic benefits. However, a more appropriate metric to capture quality of service would be a direct quantification of the safety of the resulting NEC system. To accomplish this, a linear model was used to quantify the passenger casualties (including fatalities and serious injuries) per 100,000,000 passenger-miles. The linear model uses safety rates for conventional and high-speed rail systems and weights them by the percentage of traffic on each mode according to the BCA spreadsheet<sup>39</sup>.
- *On-time Performance Model.* On-time performance is another important service quality metric that provides benefit for passengers, and is affected by access to HSR and thus must be calculated for the northern and southern states separately. This metric was assigned to be Amtrak’s stated NEC goal of 92.5% combining conventional and HSR (NEC Master Plan Working Group, 2010) for areas with HSR access, and the average of that value and the reported 72.4% for the 2014 fiscal year (Amtrak, 2014) for areas with just SOGR improvements.
- *Highway Congestion Model.* In order to capture the impact to highway service quality caused by diversion of passengers from the roads to the rails, it was necessary to estimate the congestion on the highways before and after diversion. This model begins with the

---

<sup>39</sup> The safety rates were based on data from the FRA safety statistics website for the year 2015 (Federal Railroad Administration, 2015) and a report of the TGV rail system in France (SNCF, 2013) for conventional and high-speed rail, respectively: varying between roughly 7 and 11 casualties per 100M passenger-miles. The HSR safety rate was adjusted between these values for bundles with shared tracks as opposed to dedicated HSR tracks.

data from a congestion relief study for Interstate 95 conducted for the Connecticut Department of Transportation (CDM Smith, 2014), which provides the current exit-by-exit vehicle density at rush hour on the highway. This traffic is then proportionally scaled according to a number of parameters, including diversion (from the northern states only, based on the location), natural population-driven and induced-demand-driven traffic increase, and average passengers per vehicle<sup>40</sup>. The resulting traffic is categorized as unstable, or operating above capacity, if it falls into the level of service grade of F, corresponding to a vehicle density of greater than 62 vehicles per lane, per mile. The congestion metric assigned to each bundle is the percentage of the highway designated as unstable due to excess traffic. This model captures basic nonlinearities in highway congestion as a result of natural, demand-driven traffic patterns at exits.

Appendix B includes a table showing, for each attribute of interest used in the value model of a stakeholder, which evaluative model or models it was calculated by.

#### 10.2.4 Contexts

Context variables are unknown parameters that impact the performance of the system, as measured by the attributes in the value model. Considering the impact of contextual shifts is important for the creation of robust systems and a cornerstone of Epoch-Era Analysis (Ross and Rhodes, 2008). The analysis of this case will not follow EEA exactly, but will consider the impact of context in a similar way by allowing the stakeholders to compare the alternatives in a variety of scenarios to illustrate the impact that it can have on reaching agreement. To do this, a full-factorial enumeration was performed on three context variables, totaling 18 contexts, as shown in Table 10-4.

**Table 10-4: NEC context variables**

<b>Context Variable</b>	<b>Steps</b>	<b>Levels</b>
Demand projection	3	Default, Optimistic, Pessimistic
DCF Discount rate	3	12%, 7%, 4%
Time horizon	2	2030, 2065

The default demand projection matches the one used in the NEC BCA spreadsheet, as assembled from other primary sources. The optimistic and pessimistic projections were developed as hypothetical deviations from the default, indicating the importance of acknowledging the uncertainty in projections – both in terms of their overall accuracy and the possibility of distinct impetuses affecting the assumptions used to build them. The optimistic

---

<sup>40</sup> Data for diversion came from the BCA spreadsheet, while the parameters were sourced from: population-driven traffic (Office of Highway Policy Information, 2014), induced demand (Litman, 2015), and passengers per car (Bureau of Transportation Statistics, 2015)

projection assumes the future possibility of a federal gas tax, raising traffic from road diversion by 33%, plus a generic additional 5% to all traffic. The pessimistic projection assumes that no passengers will be diverted from air travel (for example, if the airlines' response to the introduction of HSR is to dramatically reduce prices) and a 5% decrease of all other traffic. These different demand forecasts are used illustratively in this case study and not intended to be fully vetted worst-case or best-case scenarios. In a real-world application of MSTSE, each stakeholder could utilize their own demand projection, however they personally choose to set it. Differences in these projections can be explored later and potentially reconciled, as this case will demonstrate.

The Discounted Cash Flow discount rate and time horizon variables are parameters directly affecting aggregated DCF attributes and "snapshot" attributes of the NEC, respectively. The three levels of discount rate are sourced from different cited discount rates in other NEC and rail analyses (Public-Private Infrastructure Advisory Facility et al., 2011; PennDesign, 2012, and PennDesign, 2011, respectively), demonstrating that some evaluative parameters are heavily subject to interpretation, even among subject matter experts. Similarly, stakeholders may also use different time horizons for making decisions depending on their preferences for short-term or long-term planning and the urgency of their needs. 2030 is the original target year of the NEC BCA spreadsheet, while 2065 is the final year of its projections. These variables could alternatively be included as a part of the value models, as it can be argued that they represent subjective assessments of how and when to evaluate the system. The decision to model them as context variables was made in order to more closely match their implementation in the BCA spreadsheet.

Referring back to Table 10-4, we will consider the <Default, 12%, 2030> context triplet to be the "base" context for our analysis, as it matches the parameters used in the BCA spreadsheet.

### **10.3 Negotiation Structure**

Before identifying key components of the negotiation structure within the NEC, the best alternative to a negotiated agreement (BATNA) for each stakeholder needs to be identified. Fortunately, determining the BATNA is relatively straightforward for this case despite the large number of stakeholders. Because the NEC spans many jurisdictions, no one stakeholder has the authority to impose a comparable, region-spanning transportation system. Similarly, the large fixed cost of infrastructure makes it highly unlikely that anyone other than the federal government would be willing to pay for such a system alone. As a result, it is safe to say that, should no agreement among the stakeholders be reached, each stakeholder's best alternative will be to continue using the existing rail system. This status quo solution is similar to Bundle 1 in that there is no HSR and the federal government bears the costs. However, without an agreement, the costs come in the form of sustained remedial funding to continue operation as opposed to the concerted SOGR efforts of the bundle. These features all match the "No Action

Alternative” described in the recent NEC Environmental Impact Statement (US Department of Transportation and Federal Railroad Administration, 2015). “No Action Alternative” is a turn of phrase unmistakably kindred in spirit to the BATNA. Though the document itself due to its scoping mentions “negotiation” only once, it is as a “challenge” before the implementation of any alternative, presumably with the “No Action Alternative” as a fallback, matching the role of BATNA perfectly.

For this case, the BATNA is evaluated like Bundle 1 using the same models, but with changes designed to capture the following:

- funding consists of the FY2014 federal spending proposal of \$550M per year
- without SOGR, on-time reliability remains at its current level (72%)
- without SOGR, capacity remains at its current level (rather than +10%)

Without needing to look in detail, it is apparent that the BATNA will have a lower benefit than all of the bundles and will be strictly dominated in benefit and cost by *at least* Bundle 1 for any stakeholders with federal spending as their cost metric.

With the BATNA established, examining structural features of the NEC negotiations that are apparent from the problem formulation can serve to highlight useful directions with which to approach the analysis. Table 10-5 shows a list of some key features and the consequences they may have on analysis, based on the categorization scheme outlined in chapter 6. Similar to the previous case study, because this is an analysis-only implementation of MSTSE without live interaction between stakeholders, the logistics category is not relevant.

**Table 10-5: Structural features of NEC tradespace**

<b>Category</b>	<b>Key Features</b>	<b>Consequences</b>
Stakeholders	<i>Number</i> – 10 (large) <i>Representation</i> – Legislation	Search for possible agreements within coalitions; identify design variables that drive value; limit modification of value models
Preferences	<i>Alignment</i> – Some prior analysis <i>Divisible Attributes</i> – Funding	Solidify understanding of coalitions; leverage funding to create new alternatives
Alternatives	<i>BATNA Type</i> – Existing system <i>BATNA Quality</i> – Poor <i>Tradespace Size</i> – 6 (small) <i>Tradespace Completeness</i> – No	Low cost of withdrawal socially, high cost of withdrawal technically; use analyses that show all alternatives at once; emphasize creation of new alternatives
Uncertainty	<i>Contextual Uncertainty</i> – Scenarios	Analyze beyond “base” context
Logistics	Informal MSTSE	No logistic challenges

The emergent consequences of the problem structure indicate that coalitions are an important feature of the NEC tradespace. In addition to the large number of stakeholders diffusing individual influence over the solution, the passive nature of the BATNA lowers the social cost of withdrawal by being a “failure to act” rather than a distinct course of action for some or all of the stakeholders, allowing for perpetual delays or tabling of the discussion – a trend that is visible in the slow pace of real-world NEC planning. If breaking this pattern and reaching an agreement is the ultimate goal of this negotiation, using the high technical cost in the form of the BATNA’s relatively poor performance to spur the formation of interest-based coalitions is a potentially fruitful endeavor, as coalitions can exert more social pressure on stakeholders who might otherwise walk away from the negotiations. The existing stakeholder cluster analysis from the CLIOS Process execution is a useful point for comparison when analyzing this topic.

The other defining feature of the NEC tradespace is its size, with only 6 alternatives. With this few alternatives, it is feasible to consider the entire design space at once without undue cognitive burden but more difficult to assess the tradeoffs between the various performance attributes of interest due to the limited “density” of the tradespace. Correspondingly, the emphasis of the exploration should be on the comparison of all six bundles and how the design variables impact their value, rather than on marginal performance trades. At the same time, it is important to remember that this is far from a complete sampling of all potential NEC implementations, with or without high-speed rail. Therefore the testing of new alternatives, based on the analysis of value drivers, is a promising avenue for finding solutions agreeable to more stakeholders and creating additional value.

It is also worth noting that some of the stakeholders in the NEC are representatives of a larger whole. For example, were this a live negotiation with stakeholders, “Congress” would be a Congressman, representative of the larger Congress, which is in turn representative of an even larger population. This would limit the ability to update their value model upon exposure to new information or emergent insight, as the larger interest group will not be rapidly introduced to this knowledge. This issue is of less importance in analysis-only MSTSE, where value models can be edited for the purpose of analyzing hypothetical situations, but we will hold the value models fixed in this case to imitate that condition.

#### **10.4 Negotiation Success Criteria**

In order to determine what constitutes a “good” solution, success criteria for reaching a viable agreement were established before analysis was performed (so as to avoid shifting goals in response to the data). The first-order need for a plan to succeed in the Northeast Corridor negotiation is support from as many stakeholders as possible; thus a minimum of 8 out of 10 stakeholders should be in support of a given bundle before it is acceptable. There is no “correct” supermajority to use in a negotiation, as it is ultimately up to the stakeholders themselves to determine when a critical mass has been reached to proceed forward. Holden (2004)

demonstrates, on a two-vote case where stakeholders receive value based on a distance between the chosen solution and their single desired solution, that the preferred supermajority is positively correlated with the importance of the decision but negatively correlated with the number of stakeholders. Given that the NEC features an important decision with many stakeholders, a moderate supermajority seems appropriate (at least for the purposes of exploring the decision space). Searching for solutions that satisfy all 10 stakeholders is in the best interest of building momentum going forward from MSTSE but is not strictly required.

Without conducting interviews or full MSTSE with the real stakeholders, it is difficult to anticipate exactly what would constitute “support” from a stakeholder. In the absence of more detailed information on metapreferences, we will simply assume that each stakeholder is willing to support any bundle on their Pareto front, though they may have a distinct favorite bundle they prefer over the others. Additionally, in the absence of any other feasible agreements, stakeholders may accept designs within a small fuzziness margin of their Pareto front as long as they are still preferable to the BATNA. This type of tiered decision making is common in TSE, as stakeholders work to find the best possible solution that is still feasible – recall that the MSTSE visualizations were designed to support decision making of this type without establishing the Pareto front as a reference point over the BATNA.

Though unanimous agreement is not required to accept a given bundle, some individual stakeholders are required. These stakeholders have some critical purpose in one or more bundles, and the bundles are identified in Table 10-6. The US Department of Transportation and Amtrak stakeholders are always mandatory by virtue of being identified as definitive stakeholders in the NEC according to the implementation of the Mitchell stakeholder saliency framework (Mitchell, 1997) in the CLIOS stakeholder analysis. Congress is also always mandatory, as they control the ability to legislate and provide public funding. The states are mandatory only when a bundle would build high-speed rail in their jurisdiction, as it is unclear whether federal effort to compel such cooperation would even be possible. Finally, though private investment in small amounts will likely be available regardless of the chosen bundle, the large scale of private funding necessary for the all-over HSR bundles requires that the private consortiums be satisfied with the bundle.

Table 10-6: NEC stakeholders mandatory for some bundles

Stakeholder	Mandatory for Bundles	Reasoning
Dept. of Transportation (USDOT)	all	Definitive stakeholder
Amtrak and rail agencies	all	Definitive stakeholder
Congress	all	Required for legislative authority and public funding
Northern corridor states	3, 4, 5, 6	Required for northern HSR
Southern corridor states	5, 6	Required for southern HSR
EPA and landowners	-	
Private consortiums	5, 6	Private funding necessary for large-scale projects with all-over HSR
Suppliers and labor unions	-	
Highway System	-	
Travelers	-	

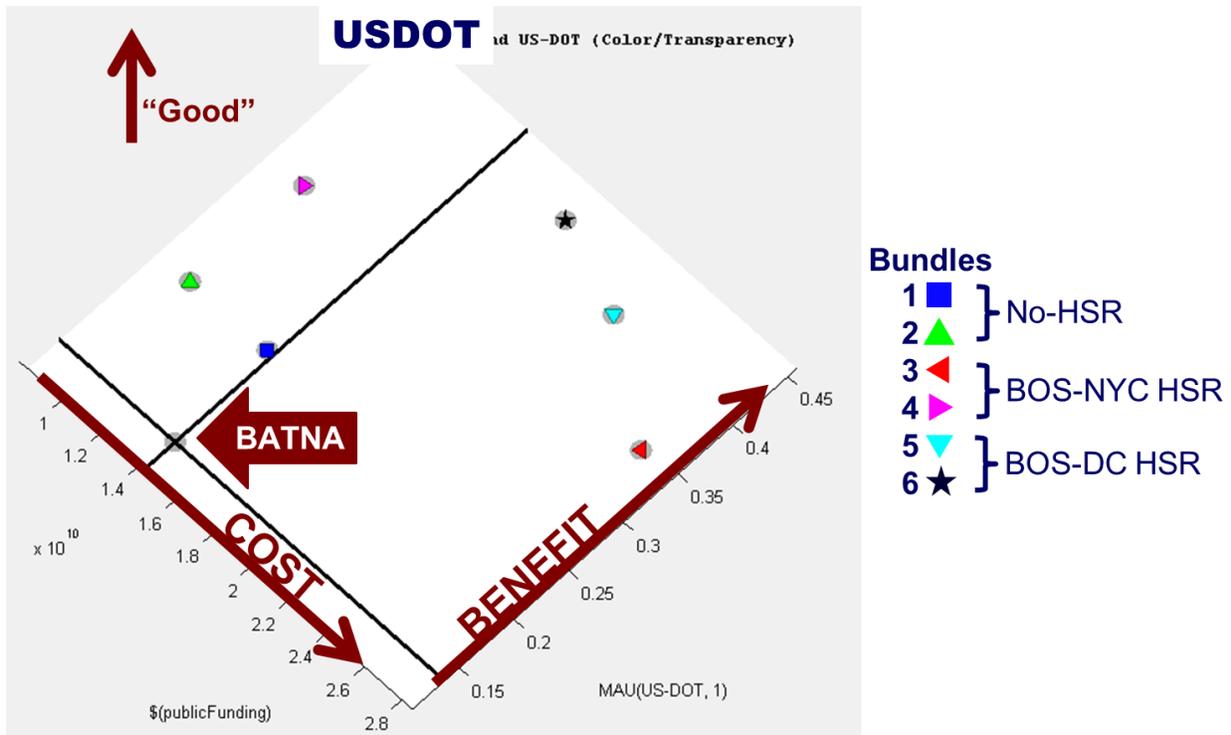
In summary, the assumptions we are using for stakeholder behavior in the NEC are:

- 8 out of 10 stakeholders must agree on a bundle for it to succeed
- All mandatory stakeholders for a given bundle must agree for it to succeed
- Stakeholders agree according to the following tiered process, preferring agreements higher up but willing to accept those lower down
  - “Favorite” or most preferred alternatives
  - Other Pareto efficient alternatives
  - Alternatives within 5% fuzziness of Pareto efficiency

### **10.5 Negotiation Tradespace and Favorite Alternatives**

With the tradespace now set up, analysis can begin. Figure 10-3 shows the negotiation tradespace in the base context for the Department of Transportation stakeholder. Each of the six bundles is marked with a custom marker, rather than utilizing the color and transparency features. This is because of the small number of alternatives, which (1) makes the color/transparency less informative due to the sparseness of the tradespace and (2) incentivizes engagement with all alternatives together in order to avoid positional bargaining driven by anchoring on a preferred solution. Additionally, because the color/transparency features of the negotiation tradespace require toggling between different stakeholders, they are less effective at

providing high-level summaries of the complete multi-stakeholder problem when there are a large number of stakeholders to toggle between.



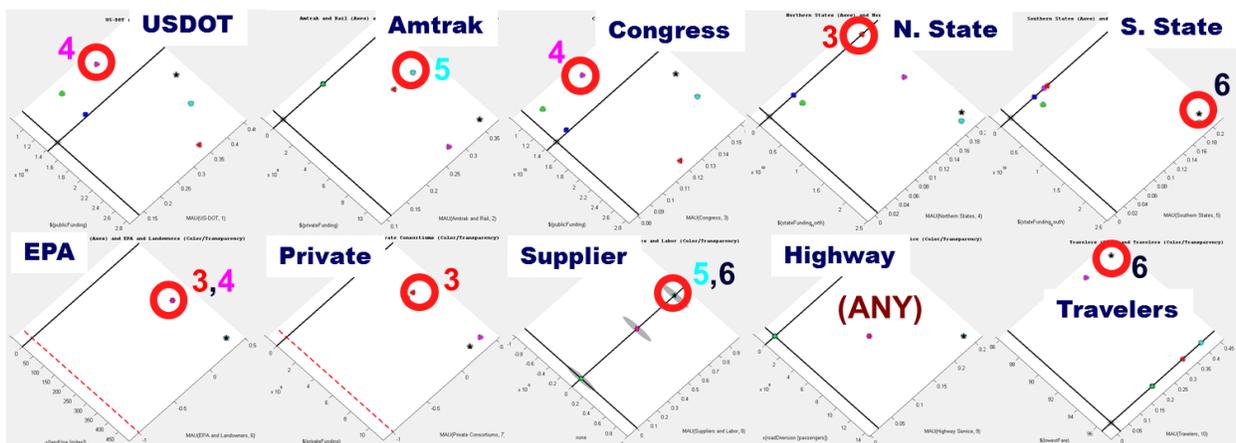
**Figure 10-3: USDOT negotiation tradespace in base context**

The negotiation tradespace for the USDOT is shown here to illustrate that the basic patterns that were expected from the problem formulation do in fact appear in the tradespace. The bundles are roughly tiered in terms of benefit received by the scope of the HSR implementation: the no-HSR bundles (1 and 2) have utilities near 0.2, the piecewise implementations (3 and 4) near 0.35, and the all-over bundles with at least 0.4. All bundles provide more benefit than the BATNA, but can be more or less expensive. In this case, because the USDOT uses federal funding as a cost metric, the bundles that save money over the BATNA are Bundles 1, 2, and 4. This leads to a tradespace populated entirely in quadrants I and II (which is also true for every other stakeholder), reflecting the lack of interest in exploring choices that could *reduce* service in order to save money: a scoping decision that is reasonable for the participants as long as it is intentional. If one or more stakeholders were interested in evaluating the potential savings of scaling back or shutting down rail operations in the NEC, additional bundles would need to be added in order to have a tradespace representative of the full range of solutions; for this analysis, we will assume that the original six bundles are an accurate depiction of the desire to explore expanding NEC rail service.

If the USDOT stakeholder were to look at Figure 10-3, they would likely look to the Pareto front to find their most preferred bundle. In this case, the Pareto efficient bundles are 2, 4,

and 6. Bundles 2 and 4 save money relative to the BATNA and provide some additional benefit and are therefore strictly superior to continuing operation of the existing rail system. Bundle 6 costs more than the BATNA but also provides the most benefit in the tradespace, making it a potentially interesting alternative. Upon cursory examination, Bundle 4 is likely to be the “favorite” design of the USDOT, because it has the best benefit-cost tradeoff and is located in quadrant II.

Figure 10-4 shows all ten stakeholders’ negotiation tradespaces in the base context, highlighting the bundle that is most likely to be their “favorite” in the same way that the USDOT prefers Bundle 4. Some stakeholders are indifferent between certain bundles (EPA, suppliers) and others may not have a clear favorite due to a linear Pareto front (Highway System). Regardless, it is clear that the stakeholders prefer a variety of different designs. If each one looked at the problem individually, with no regard for the need for group agreement or coordination driven by MSTSE or other negotiation assistance, it is highly likely that positional bargaining between these different “favorites” would result. A development plan for the NEC, as a CLIOS system, cannot feasibly be determined by any sort of simple majority-rule voting, but we can see that a vote would fail to resolve the debate anyway. As it stands, Bundles 3 (2 votes), 4 (2 votes), 5 (1 vote), and 6 (2 votes) would all receive support. Even if the three stakeholders with multiple potential favorites were to vote for the same bundle, none of the bundles would reach a simple majority, deadlocking well short of the necessary 8/10 supermajority for a successful agreement.



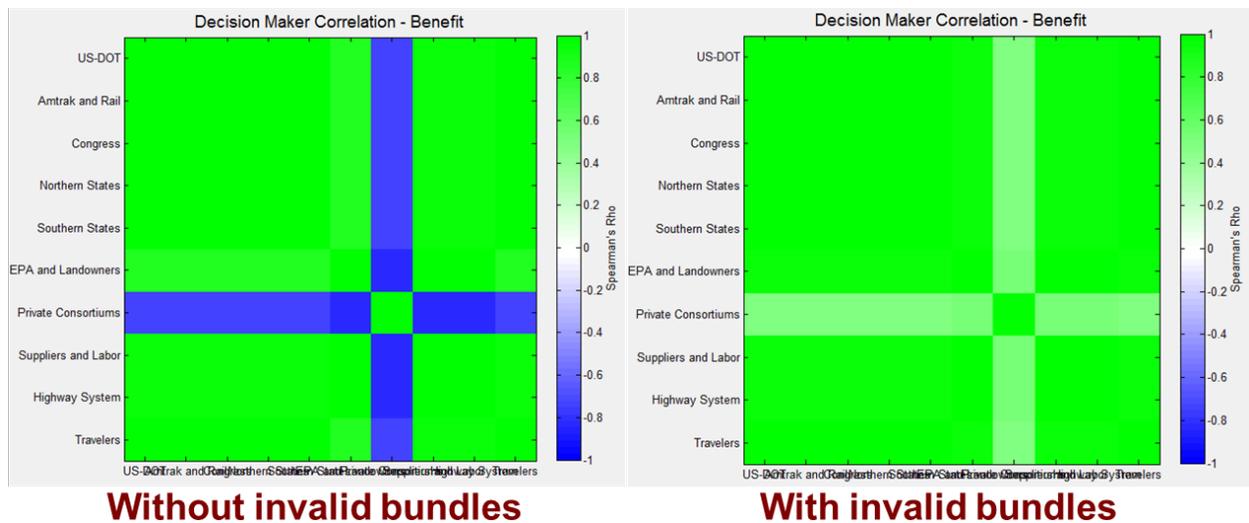
**Figure 10-4: Negotiation tradespaces for all stakeholders in the base context, with "favorite" bundles circled**

Despite the failure to identify an emergent shared “favorite”, one notable trend in this data is that neither Bundle 1 nor 2 (nor the BATNA) are any stakeholder’s preferred solution. This is true even though most stakeholders have at least one of those bundles on their Pareto front. The cost of building HSR appears to always “pay off” with a large amount of benefit for each stakeholder for at least one combination of other design variables.

## 10.6 Stakeholder Alignment and Coalitions

As predicted, due to the large number of stakeholders, a jointly preferred favorite design is not available in this tradespace. Accordingly, the development of coalitions is a likely outcome in the NEC and an important feature to analyze in MSTSE. Understanding the possible interest blocs and what brings them together (and separates them from the others) builds valuable intuition for the underlying problem – the main goal of TSE and MSTSE. To do this, we can look at the alignment of stakeholder interests as measured by their correlation.

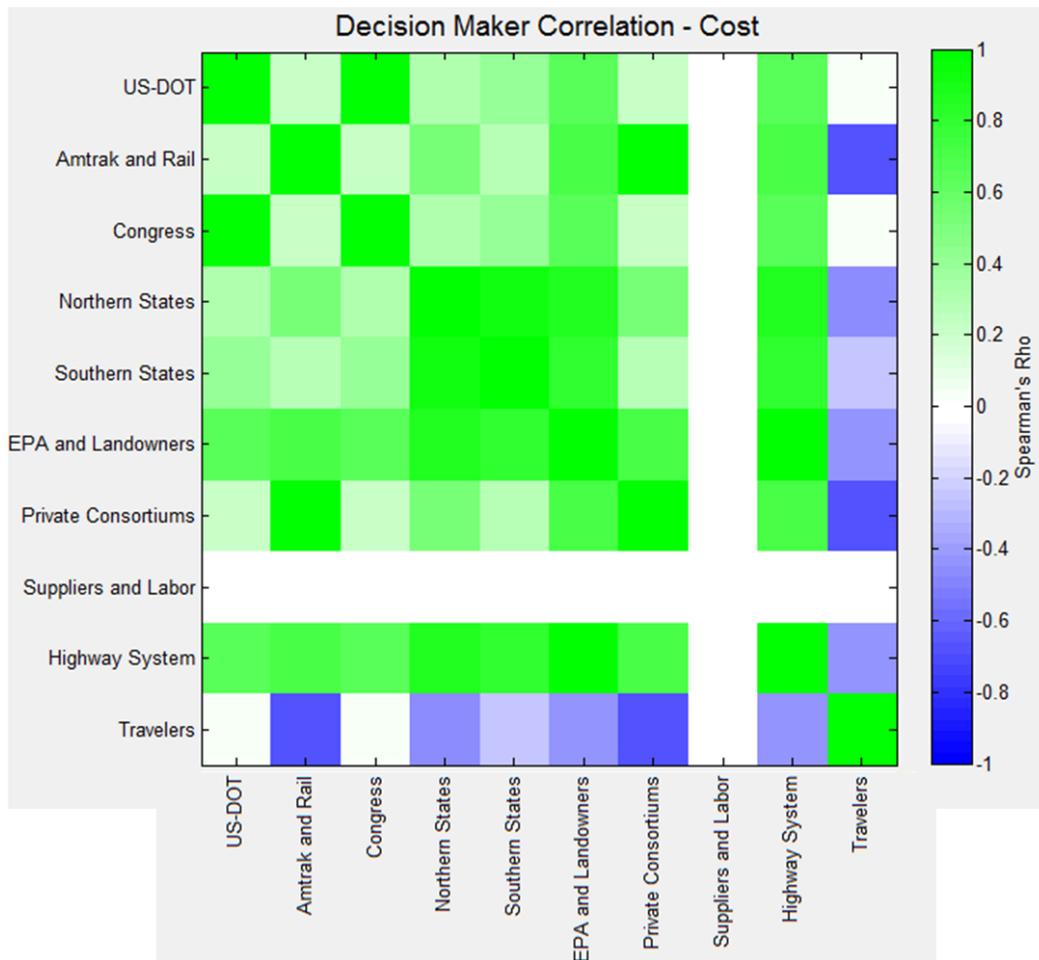
Figure 10-5 shows two correlation heatmaps for the utility functions of the ten stakeholders. Most of the heatmap is green, suggesting that the stakeholders are largely in agreement on what bundles provide the most benefit. This confirms the prior conclusion that the bundles with a larger HSR implementation always drive more benefit. The one exception to this rule, and the source of the blue stripe on the left plot, are the Private Consortiums. Because they have a utility-driving preference on payback period – the time it takes for the NEC to earn back the capital expenditures of the project – they derive more benefit from piecemeal-international than all-over HSR bundles, which cost more and take longer to reach full capacity. However, the blue stripe indicating opposed interests is also an artifact of invalid bundles, as designated by failing to meet the utility function minimum requirements in one or more attributes, exerting significant impact on the small design space. For the Private Consortiums, Bundles 1 and 2 (and the BATNA) are marked as invalid by the utility function, specifically for failing to reach enough private returns to merit lending support to the plan. When these bundles are included in the rank correlation with a “utility” of zero, the blue stripe becomes light green, indicating that the Private Consortiums are only slightly different from the other stakeholders in terms of benefits as a result of that single rank reversal.



**Figure 10-5: NEC stakeholder utility correlations without (left) and with (right) invalid bundles included**

With the benefits highly correlated, it is likely that costs are the main driver of the differences in “favorites” of each stakeholder. Figure 10-6 shows the cost correlation plot, which

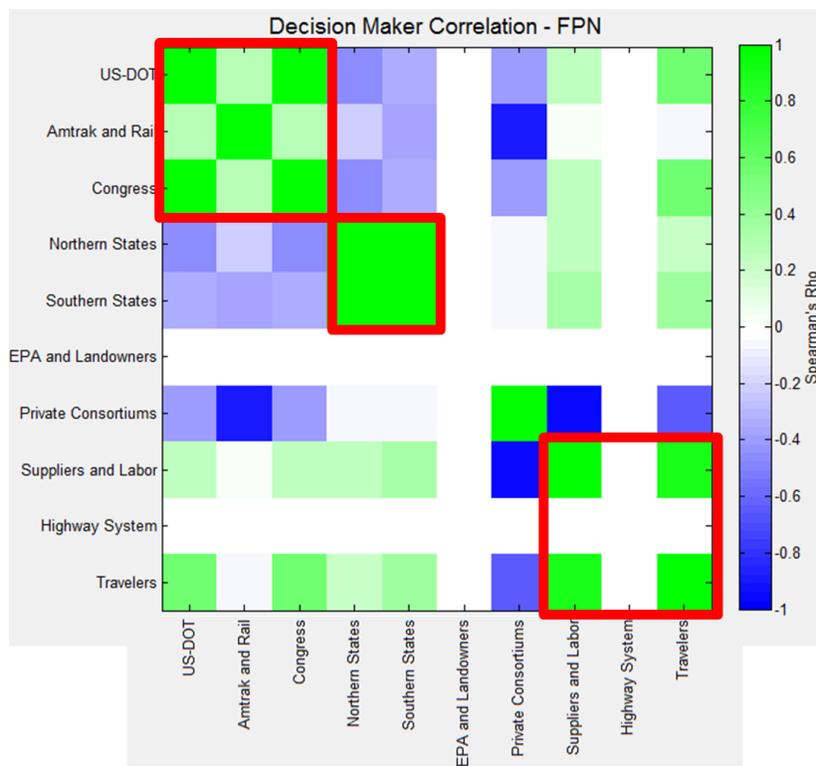
reveals more significant differences between stakeholders. The correlations are still mostly positive, but are closer to 0.5 than 1 for most stakeholder pairs. This is again driven by the direct correlation between HSR scope and cost establishing three distinct tiers of total system cost that are shared by most stakeholders. Within those tiers however, the cost-ordering of those bundles changes depending on the funding mechanism and how it allocates cost between public, private, and state sources. The only stakeholder with negative cost correlation to the others is the Travelers stakeholder, which uses Fares as a cost metric. Fares are decreased in the presence of competing operators (the organizational structure variable), which are attached to the more costly with-HSR bundles. Note that the Suppliers stakeholder shows all-white in this plot, as they are a benefit-only stakeholder and do not bear any costs of the system (and thus cannot be correlated).



**Figure 10-6: NEC stakeholder cost correlation**

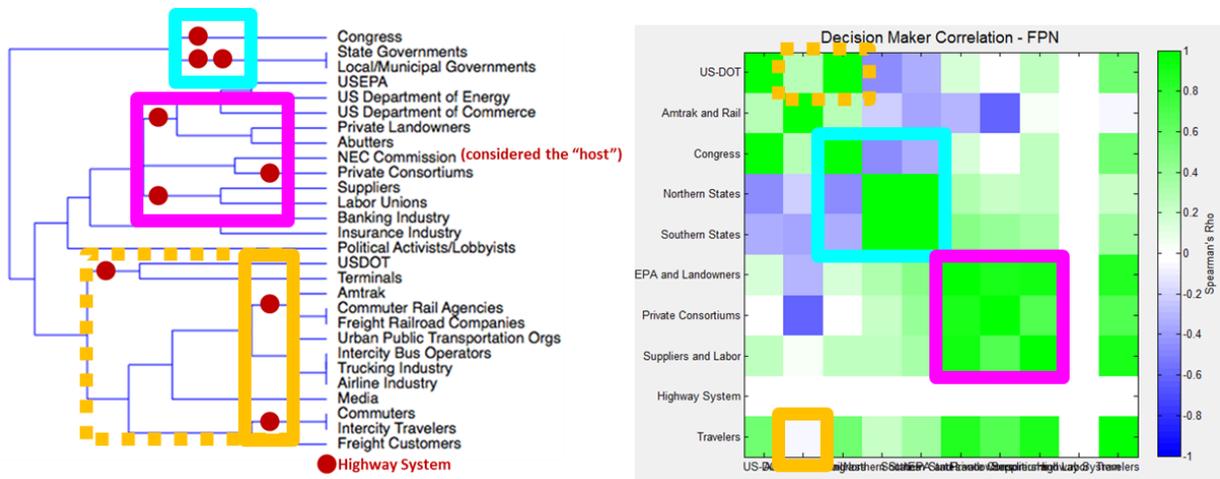
Clearly the issue of “who pays” for the NEC is more contentious than what constitutes a “good” system. Combining these two issues and correlating a value metric like cost-benefit efficiency can reveal which stakeholders agree about the “best value” designs: the alternatives on and closest to the Pareto front. Figure 10-7 shows the correlation of Fuzzy Pareto Number (FPN) between the ten stakeholders. Three main groups of stakeholders emerge as potential

coalitions and are highlighted with red squares: (1) USDOT, Amtrak, and Congress, (2) the Northern and Southern States, and (3) the Suppliers and Travelers. Coalition #3 is weakly correlated with both other coalitions, but interestingly Coalitions #1 and #2 are opposed, as indicated by the light blue rectangles adjacent to them. The outlier stakeholders include the Private Consortiums (driven by the already-identified difference in utility compared to the other stakeholders) and the EPA and Highway System (which have all valid designs on the Pareto front and thus no variability in FPN with which to perform a rank-order correlation and which could presumably ally with any coalition). As no single coalition has enough member stakeholders to force through an agreement, the resolution of this negotiation requires identification of a bundle that is acceptable to more than one coalition, specifically one that has the proper balance of funding pools.



**Figure 10-7: NEC stakeholder FPN correlation, with three main coalitions highlighted**

Before returning to the comparison of the different bundles, we will compare these predicted coalitions against those of the clustering analysis performed in (Sussman et al., 2015) and with additional detail in (Moody, 2016). Because the clustering analysis was performed on interest vectors with no connection to specific valid/invalid bundles, it is more appropriate to include the invalid bundles in the FPN correlation than to leave them out when making this comparison. Figure 10-8 shows both the clustering analysis and FPN correlation, with the coalitions identified by the clustering analysis highlighted in both.



**Figure 10-8: Coalitions identified by clustering analysis, compared to FPN correlation (with invalid bundles)**

There are three main coalitions formed by the branches of the clustering dendrogram, each with a different relation to the FPN correlations:

- *EPA, Private Consortiums, and Suppliers (magenta)*. This coalition of interests is also apparent in the FPN correlations, with all members well over correlation 0.5 with each other. The two methods are in agreement on this grouping.
- *Congress, Northern and Southern States (cyan)*. This coalition is not borne out by the tradespace analysis. As previously identified, the states are positively correlated with each other, but negatively with Congress. This is likely the result of two factors: (1) FPN correlation accounts for the different funding pools accessed by the states and Congress, while the clustering was performed only on benefits, and (2) the “state governments” are divided into separate northern and southern stakeholders. At an aggregate, all-states level of analysis, the states shared values are more similar to national interests than they are when broken into smaller regions. Thus, we can confidently assert that the fidelity improvement of MSTSE has revealed a weak coalition in the original analysis.
- *Amtrak and Travelers (orange) followed by USDOT (dotted)*. Interestingly, Amtrak and Travelers are essentially uncorrelated in the tradespace but both of them are positively correlated with the USDOT. This is probably caused by a combination of the small tradespace size, where small changes in ordering can have a large impact on overall correlation, and its particular form of expert-influenced enumeration. Amtrak and Travelers have a strong negative correlation in cost, driven by the enumeration of the tradespace combining the fare-lowering competing operators only with large private investment bundles. Based on the tradespace analysis, this is either (1) an artifact of the limited tradespace size, or (2) a reality of the problem uncaptured by the interest vectors used in the clustering analysis, where competing operators *require* a sizable private

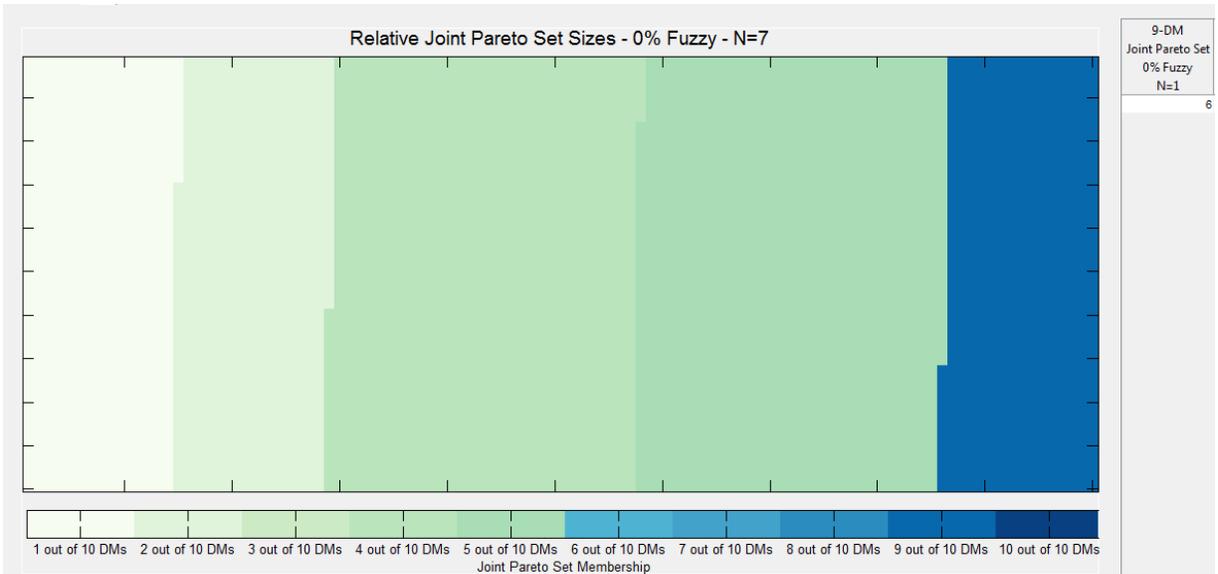
investment. This type of insight is valuable when performing analysis-only MSTSE as a discussion point to bring back to other stakeholders or subject matter experts, as their reactions will either inform iterative improvement of the tradespace or confirm a previously unconsidered aspect of the negotiation.

Overall, the interest vector clustering from the CLIOS Process execution and the FPN correlation methods of predicting coalitions display considerable convergence in the insights they generate. The differences are attributable to fidelity improvements in the tradespace by accounting for costs, and also provide useful discussion points to either bring into a negotiation after completing analysis-only MSTSE or to further refine and improve the tradespace.

As a final insight into interest-based coalitions in the NEC, it is worth pointing out that the “mandatory” stakeholders do not exercise complete control of the solution. There are four stakeholders who are mandatory for no bundles: the EPA, Suppliers, Highway System, and Travelers. Based on the low individual influence of these stakeholders, it might seem at first glance that the other six stakeholders are more important to consider and satisfy when picking a bundle. However, those four stakeholders have enough *combined* power to prevent a successful agreement, which we have assumed will require at least 8 out of the 10 stakeholders. In addition to that, all four of those stakeholders have Bundle 6 on their Pareto front; if they wanted to insist as a group on Bundle 6 they could at least force an impasse in the negotiation. This type of coalition analysis is similar to the follow-up to the cluster analysis in (Moody, 2016) in which the stakeholder saliency framework was used to predict the most likely coalitions from the dendrogram based on the fastest ways to assemble power, urgency, and legitimacy together on the graph, but MSTSE uses the extra layer of specific alternatives to potentially bridge between different branches. For this example, we already identified a likely coalition including the EPA, Suppliers, and Travelers based strictly on interests; however, based specifically on Bundle 6 we find crossover with the Highway System – a stakeholder previously unattached to any of the three major potential coalitions. Using specific bundles to consider expansions of coalitions across branches of interest alignment is another valuable aspect of analyzing the preferences and alternatives together with MSTSE rather than separately.

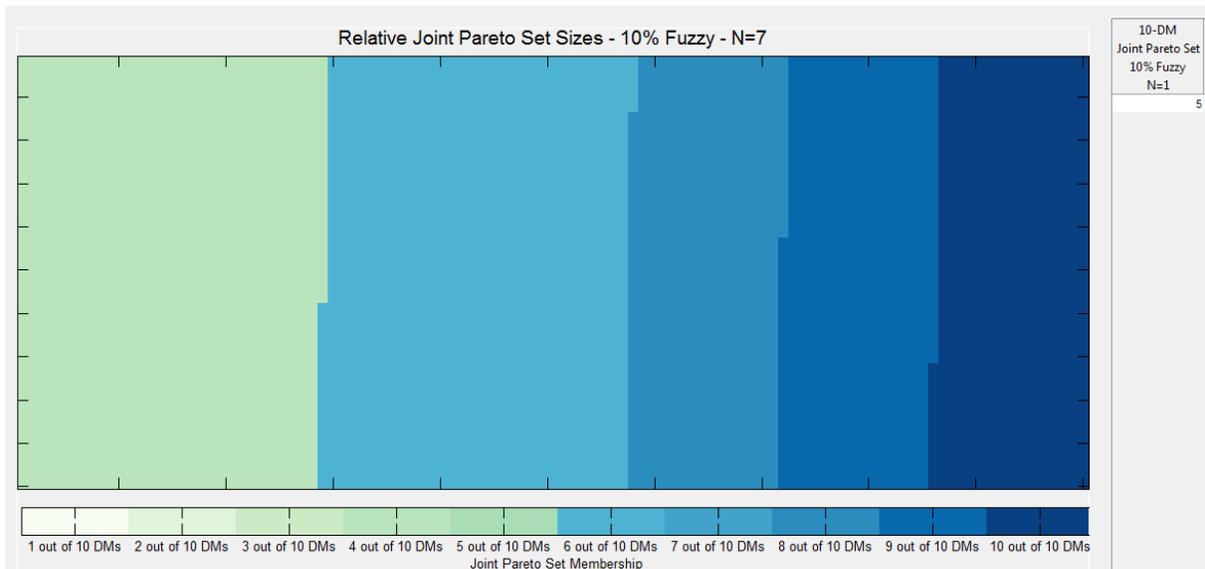
## **10.7 Finding Mutually Beneficial Alternatives**

Returning our attention to evaluating the six alternatives and searching for the most promising solution, we will use a gridmap to get an idea of how appealing each of the bundles are across all of the stakeholders. Figure 10-9 shows a gridmap of the base context, indicating that Bundle 6 is on the Pareto front for nine of the stakeholders while none of the other bundles get above five. The high correlation amongst all the stakeholders’ utility functions drives this, as Bundle 6 has the highest utility (and is therefore on the Pareto front, by definition) for nine of them, though based on our inspection of the benefit-cost tradespaces Bundle 6 is not the favorite choice for many of them.



**Figure 10-9: Gridmap of base context, showing Bundle 6 with 9/10 stakeholders and none other above 5**

Alternatively, adding a 10% fuzziness buffer to the gridmap, as in Figure 10-10, shows that Bundle 5 is the nearest to all Pareto fronts and the first bundle to reach all ten stakeholders: the minimax FPN solution. Bundle 6 remains at 9/10 stakeholders even with the 10% buffer. Either of those designs are interesting choices to build support behind, particularly Bundle 6 which we already identified as a potential bridge between coalitions.



**Figure 10-10: 10% fuzzy gridmap of base context, Bundle 5 is first to 10/10 stakeholders**

However, neither Bundle 5 nor 6 can be a successful agreement under the rules outlined at the start of the case. Table 10-7 lists the bundles passing the 8/10 stakeholders threshold at 5% fuzziness. Each of these three bundles is unacceptable for a mandatory stakeholder. Without some further information, there will be no valid agreement.

**Table 10-7: Bundles within 5% of the Pareto front in the base context**

<b>Bundle</b>	<b>5% Fuzzy Pareto Efficient for</b>	<b>Inefficient Stakeholders</b>
6	9/10	Private Consortiums ( <i>mandatory</i> )
5	9/10	USDOT ( <i>mandatory</i> )
4	8/10	Amtrak ( <i>mandatory</i> ) Suppliers and Labor

Bringing in analysis of different contexts is the type of information that could increase the attractiveness of one or more of these designs. With such a small tradespace of highly differentiated alternatives, it is unlikely that the context variables scoped into this analysis will significantly reorder the attractiveness of the bundles; they are too far apart to change the membership on the Pareto front, making a full Epoch-Era Analysis treatment excessive based on the limited insights that would result from multi-epoch or era analysis. However, the relative attractiveness of a bundle could go up or down depending on the underlying assumptions, and a change for even a single stakeholder can make the difference in securing a valid agreement.

For these reasons, we will consider a scenario planning approach in which each stakeholder has a preferred evaluative scenario, represented here as a choice of context variables. Given the subjective nature of the context variables for this case, the choice of scenario is a form of metapreference, indicating a preferred mode of making decisions with respect to optimism/pessimism in forecasting and the relative value of near- and long-term benefits. Each stakeholder can evaluate the alternatives in their chosen scenario independently from the other stakeholders. In this metapreference-driven space, it does not matter if they disagree on what will actually happen (the context) as long as they agree on what course of action (bundle) is the best way to move forward. If no agreement is able to be reached, addressing the forces behind these assumptions and determining if they can be reconciled between stakeholders is one means of pushing forwards.

To demonstrate this type of analysis, each stakeholder was assigned a preferred context. The context variables were assigned based on presumed metapreferences for each stakeholder: comprised of a desire for long term benefits (over short-term, as in the base), a more positive or negative demand outlook for rail travel, and the option to lower the discount rate. Again, in a full-participation MSTSE implementation, the stakeholders would be consulted directly to establish these preferences, but in their absence, subject matter experts are the next best source<sup>41</sup>.

---

<sup>41</sup> In most cases you would anticipate a bigger drop-off in quality when capturing metapreferences as opposed to utility-driving preferences when using subject-matter experts: inclinations such as choosing worst-case preparation over expected-value planning are by their nature more personal to the stakeholder and less likely to be known just from expertise in the technical area of interest.

Table 10-8 contains a list of each stakeholder’s preferred context and how it differs from the base in terms of the context variables.

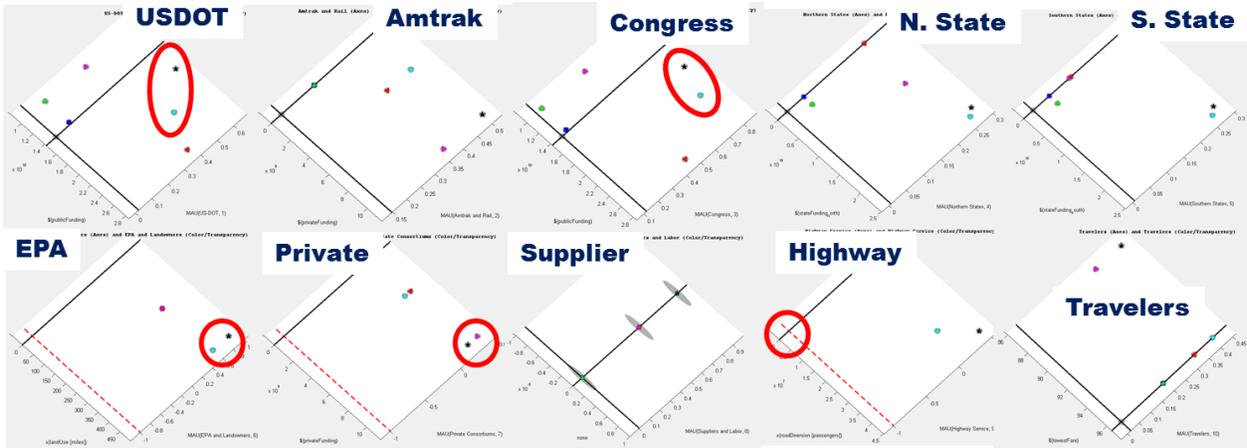
**Table 10-8: NEC stakeholders’ preferred contexts (changes from the base are bolded)**

Stakeholder	Metapreferences / Context Assumptions	Context ID#	Year	Demand	Discount Rate
<b>Dept. of Transportation (USDOT)</b>	Long term benefit, neutral outlook	10	<b>2065</b>	Default	12%
<b>Amtrak and rail agencies</b>	Positive outlook, reduced discount rate	8	2030	<b>Optimist</b>	<b>7%</b>
<b>Congress</b>	Very long term benefit (minimum discount rate)	3	2030	Default	<b>4%</b>
<b>Northern corridor states</b>	Reduced discount rate	2	2030	Default	<b>7%</b>
<b>Southern corridor states</b>	Reduced discount rate	2	2030	Default	<b>7%</b>
<b>EPA and landowners</b>	Long term benefit, Positive outlook	16	<b>2065</b>	<b>Optimist</b>	12%
<b>Private consortiums</b>	Negative outlook	4	2030	<b>Pessimist</b>	12%
<b>Suppliers and labor unions</b>	Base	1	2030	Default	12%
<b>Highway System</b>	Long term benefit, Negative outlook	13	<b>2065</b>	<b>Pessimist</b>	12%
<b>Travelers</b>	Base	1	2030	Default	12%

If each stakeholder chooses to use a specific context, the tradespace for each will be affected. Figure 10-11 shows the updated negotiation tradespaces for each stakeholder. As predicted, though the bundles change utility in each tradespace, there is little overall reordering of the designs due to their mostly large initial separation. Some notable changes are circled in red. These include:

- Bundle 6 improves relative to Bundle 5 for stakeholders with long-term decision windows (USDOT, Congress, EPA), as the full benefits of maximum HSR capacity are not experienced until later in the lifetime of the system after demand builds up.
- Bundle 6 improves slightly relative to Bundle 4 for the Private Consortiums due to a smaller difference in payback period under a pessimistic demand outlook.
- The BATNA and Bundles 1 through 4 are now all invalid for the Highway System. Under worst-case planning (long term with very few people riding the

trains), they require all-over international quality HSR in order to divert enough travelers to maintain feasible congestion on the roads.



**Figure 10-11: Negotiation tradespaces for all stakeholders in chosen contexts, with notable changes from base circled**

Reevaluating Table 10-7 in the new contexts results in Table 10-9. This clearly illustrates how allowing personal context selection, as a reflection of stakeholder metapreferences, can help or hurt potential agreements. Bundle 6 is now a viable agreement, as it has moved within 5% of the Pareto front for the Private Consortiums after adjusting their tradespace to the lower demand estimate. On the other hand, Bundles 4 and 5 have both dropped a stakeholder from their prior consensus. Bundle 4 is now under the 8/10 stakeholders threshold after the Highway System marks it as invalid. Bundle 5 lost another mandatory stakeholder in Congress due to the relative improvement of Bundle 6.

**Table 10-9: Bundles within 5% of the Pareto front in stakeholder-chosen contexts**

Bundle	5% Fuzzy Pareto Efficient for	Inefficient Stakeholders
6	10/10	
5	8/10	USDOT ( <i>mandatory</i> ) Congress ( <i>mandatory</i> )
4	7/10	Amtrak ( <i>mandatory</i> ) Suppliers and Labor Highway System

Overall, it appears that Bundle 6 is the alternative in the tradespace most likely to lead to a successful negotiation and implementation. After allowing for stakeholders to set their own contexts, it is the only alternative that meets the criteria for a successful agreement and is even

unanimously preferable to the BATNA (though not strictly cost-benefit efficient for each stakeholder)<sup>42</sup>.

## **10.8 Adding New Alternatives**

With such a small number of alternatives, we identified during problem formulation that the addition of new alternatives was an important value-creation activity for the NEC. Recall, the central goal of principled negotiation is to “create and explore many options”. With only six alternatives enumerated of a highly complex, multidimensional design space, it is certainly possible that other bundles exist that would be attractive to a larger subset of stakeholders. Creating new designs can be done traditionally, by returning to problem formulation with subject matter experts and establishing new design variables or levels of existing variables, or by modifying existing designs to target improvement in certain attributes. For the NEC, the divisibility of funding between the different sources provides easy leverage for modifying existing bundles in ways that will directly benefit stakeholders who are blocking agreement on the original bundle. We can explore this as a rapid means of testing new designs and their impact if introduced to the negotiation.

Table 10-10 shows the three bundles that were identified as closest to agreement for all stakeholders and describes a funding shift to support the stakeholders for whom each bundle was worst. Bundle 6, though a valid agreement based on the negotiation rules, is still clearly worst for the Private Consortiums, who in application may choose to fight for a more favorable design such as Bundle 3. The stakeholders who use public funding as a cost metric are generally more in favor of Bundle 6. Thus, a new design (Bundle 9) was created based on Bundle 6 but with the private funding reduced to a rolling-stock-only level, shifting the difference into the public and state funding pools so that the total funds remain constant. Conversely, Bundle 5 was worst for the USDOT and Congress, so a new design (Bundle 8) was created by reducing the public funding to the same level as Bundle 6, shifting the difference into the private funding pool. Bundle 4 cannot be modified effectively using this technique, as the Suppliers and Labor and Highway System stakeholders do not use the divisible funding pools as cost metrics. Creating a new design based on Bundle 4 to satisfy these stakeholders would require returning to problem formulation and considering adding new variables that can be adjusted to affect the attributes that those particular stakeholders are interested in<sup>43</sup>.

---

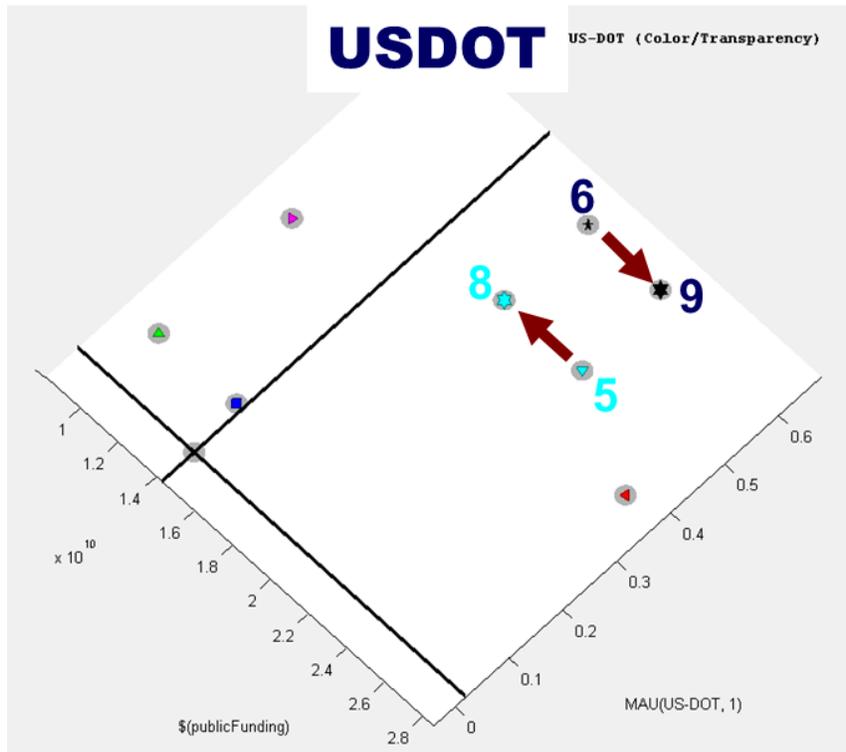
<sup>42</sup> Remember: we assumed in the problem formulation that each stakeholder would accept any bundles within 5% of their Pareto fronts over the BATNA. This of course may not be true in reality or for any different cases – each stakeholder would have to be consulted directly during a “real” application of MSTSE – but this captures the “close to the Pareto front” criteria that many stakeholders use when looking at a tradespace.

<sup>43</sup> For the curious: there is no “Bundle 7” because the tradespace software we used needed to assign a number to the BATNA – thus it was set as “7” *before* we attempted to create new alternatives in this way.

**Table 10-10: Creation of new alternatives by shifting funding sources**

<b>Original Bundle ID</b>	<b>Worst for</b>	<b>Funding Change</b>	<b>New Bundle ID</b>
6	Private Consortiums	Private reduced to rolling stock only, shifted to Public and States	9
5	USDOT Congress	Public reduced to match Bundle 6, shifted to Private	8
4	Amtrak Suppliers and Labor Highway System	N/A	N/A

With two new alternatives defined, the tradespace can be updated. Figure 10-12 shows the USDOT negotiation tradespace with Bundles 8 and 9 added. Clearly, they are cost-shifted versions of the existing Bundles 5 and 6; in this case, 8 has improved and 9 has worsened compared to their originals based on the USDOT’s preference on public funding. From this plot, we can see that Bundle 8 is not-quite on the Pareto front: it is weakly dominated by Bundle 6 by virtue of having the same cost but a lower utility. This is akin to being fuzzy Pareto efficient in the limit as fuzziness approaches zero. Designs that are weakly dominated, while not being individually preferable for a given stakeholder, are prime candidates for selection across multiple stakeholders when no designs are jointly efficient (as is the case for the NEC).



**Figure 10-12: USDOT negotiation tradespace in preferred context, with new bundles**

With the new bundles added to the tradespace, we can reevaluate the consensus around each bundle. Table 10-11 shows the three original interesting bundles plus the two new bundles, along with how many stakeholders find each efficient at three levels of fuzziness: on the Pareto front, weakly dominated, or within 5% (the agreement condition).

**Table 10-11: Number of stakeholders with each bundle at different levels of efficiency, in preferred contexts**

<b>Bundle ID</b>	<b>Pareto Front</b>	<b>Weakly Dominated</b>	<b>5% Fuzziness</b>
9	5/10	<b>10/10</b>	<b>10/10</b>
8	2/10	5/10	8/10
6	8/10	9/10	<b>10/10</b>
5	2/10	5/10	8/10
4	4/10	6/10	7/10

From this data we can draw a number of conclusions:

- Neither Bundle 8 nor 9 is particularly good at the Pareto front level. These bundles are not replacing the most preferred bundles for most individual stakeholders.
- However, the introduction of these bundles *does* impact some of the original bundles. For example, Bundle 6 is down to 8 out of 10 stakeholders on the Pareto front. This is because Bundle 9 dominates Bundle 6, previously Pareto efficient, for Amtrak.
- Bundle 9 joins Bundle 6 as an acceptable agreement, unanimous at the 5% fuzziness level.
- Most importantly, Bundle 9 is actually unanimous at the weakly dominated level: it is Pareto efficient for five stakeholders and weakly dominated for the other five. This makes it the “fastest” design to reach consensus by adding to the fuzziness buffer. Bundle 6 has only nine stakeholders at this level and the missing stakeholder is mandatory (Private Consortiums).

Thus, Bundle 9 has created some additional value, in terms of distance from the Pareto front, by redistributing the costs of Bundle 6 in favor of the private-funding stakeholders. Bundle 6 is good enough for public and state funding stakeholders that it remains efficient as Bundle 9 even after carrying the extra cost. This is not true for Bundle 8, as for each stakeholder gained by reducing public funding another is lost by increasing private funding – neither Amtrak nor the Private Consortiums like Bundle 5 enough to bear the extra cost. Bundle 9 appears to have a slight edge over Bundle 6 as the best solution when considering all ten stakeholders.

The ease with which new funding divisions can be added to the tradespace and tested allows this type of analysis to be performed without necessarily referring each potential new design back to subject matter experts. However, given that the original design space for the NEC

was heavily scoped with expert opinion, it is important to consult those experts again on the feasibility of any new alternatives developed *ad hoc* during MSTSE that are potentially interesting solutions. The final section of this case study will discuss the response of subject matter experts to the insights developed here, including the possible feasibility of Bundle 9.

## **10.9 Beyond FOTE MSTSE**

MSTSE has been used within this research as a technique for supporting negotiations engaged with the principle of Full, Open, and Truthful Exchange. This was mainly the result of two considerations: (1) that FOTE is associated with mutually beneficial and satisfactory negotiations and thus is a worthwhile ideal to pursue, and (2) that this early step towards integrating the insights of negotiation theory into TSE and engineering should focus on the basics without becoming preoccupied over concerns about strategic competitive behavior from the participants, especially given the many realistic examples of good-faith cooperation in engineering. However, these assertions do not prevent competition from influencing engineering negotiations – indeed, it may even be a reach to categorize the NEC as an “engineering” negotiation, given that it is highly politicized and, as we identified earlier, the technical engineering challenge for developing a HSR system is minimal at this point in time. The NEC provides two “real world” examples of incentivized competitive behavior that are prominent enough to consider directly.

First, there are stakeholders with a vested interest *against* the construction of HSR in the NEC. Some were able to be combined into larger interest groups in the transition to MSTSE. For example, landowners and abutters (who are cost-only stakeholders experiencing property devaluation due to noise or environmental pollution, identified as possessing only *urgency* in the stakeholder saliency framework) were predicted to willingly ally with the EPA who, despite potentially favoring HSR, share an interest in reducing pollution and have both *urgency* and *legitimacy*. This was not possible for all stakeholders, including the other commercial passenger and freight modes in the corridor: highway and air. Directly competitive stakeholders involved in these industries are incentivized to obstruct negotiation and prevent agreements even if they agree to follow the FOTE principle, as the “truth” of their interests will place them at odds with any functional agreement. In tradespace terminology, the BATNA will dominate all other possible agreements, since the NEC bundles were chosen to include only expansion of rail service. We recommend conducting private negotiation sessions *without* these stakeholders as a means of building momentum behind a particular solution; but how should they be approached after completing MSTSE?

Consider the airline industry as an example<sup>44</sup>. Why are they opposed to HSR? Clearly, any improvement to rail travel is likely to divert passengers from air to rail and HSR represents a considerable speed improvement over conventional rail. As a first order impact, fewer air customers will lead to less revenue for the airlines. Though passenger diversion would come with a positive side-effect of reducing air travel congestion, HSR is highly unlikely to receive support from airlines if it comes at the cost of reduced profits. Various unilateral ameliorative actions are available to the airlines to lessen the impact of HSR. Most directly, fares can be lowered to better compete with HSR on price and stem diversion, but this also leads to reduced revenue. Alternatively, airlines could reduce capacity offered within the NEC and shift it towards long-haul routes on which rail cannot directly compete. There has been some discussion (PennDesign, 2011; 2012) on the positive impacts of such an adjustment, citing the “cost-effectiveness” of long-haul flights over short-haul flights, but this is both (1) optimistic about the ability to spur increased long-haul demand to replace the lost short-haul demand, and (2) dependent on the idea that airlines are forced to offer unprofitable short-haul service within the NEC to feed the rest of their networks, despite the fact that the majority of air traffic in the NEC is point-to-point rather than hub-and-spoke (Anderson, 2007). A more realistic example of the impact of HSR on air travel is the introduction of London-Paris high-speed service upon the opening of the Chunnel in 1994. Point-to-point air travel between those cities decreased almost 50% between 1994 and 2007 with most of the losses in the first two years (French Air Transport Directorate, 2008). This decrease occurred despite overall air travel between the UK and France increasing over the same period by about 50%, driven by the rise of low-cost carriers: a mitigating shift that has already occurred in the United States.

Therefore, the most realistic airline response to proposed HSR is to obstruct the plan. This behavior has previously been noted amongst US carriers, particularly in Southwest Airlines’ successful attempt to block a potential HSR project in its hub of central Texas in the early 1990s and its current noncommittal but wary response to a possible revival (Wray and Nicholson, 2015; Campoy and Tsuneoka, 2015). Even if airlines are not invited to participate in private MSTSE negotiations, they will almost certainly insert themselves into the following planning and legislation phases and the momentum from a consensus derived from MSTSE may not be enough to brush them aside. How might the MSTSE participants who want to move forward with HSR choose to approach the airline challenge?

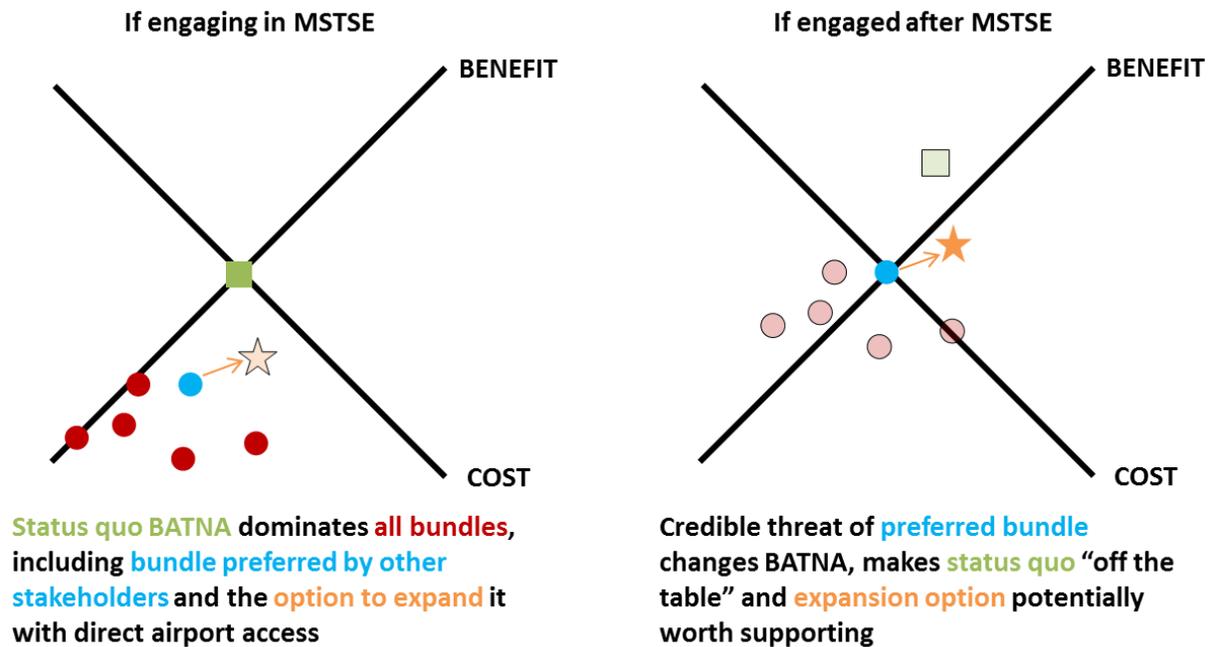
The –ilities are a potentially attractive means for preparing a response to challenges from competitive stakeholders post-MSTSE. In this case, the participating stakeholders could investigate changeability in the NEC bundles in the form of specific options designed to appeal

---

<sup>44</sup> As with most stakeholders at our level of analysis, the “air” stakeholder could be described in more detail as a collection of similar-interest stakeholders (in this case, passenger airlines, air freight, and airports), but their shared motivation to support air traffic is clear enough that focusing on airlines can be a representative example.

to airlines and airports. Prominent suggestions in this domain include expanding codeshare services between high-speed rail and airlines, and the construction of direct rail access to airports. Codeshares provide easy integration of ground and air travel services for customers, and allow for some outsourcing of ticket sales on the part of the airlines to rail operators. Direct intercity rail access to airports, currently limited to Newark in the NEC, can greatly increase the catchment area for airports and airlines by decreasing travel time necessary to reach the airport. High-speed rail stations at airports have developed large modal shares of departing travelers in Europe (up to 30%, see Kouwenhoven, 2008) compared to the typically under 5% share for conventional rail in the United States (US Government Accountability Office, 2013). The NEC is one of the few areas in the US where these are feasible: featuring both the frequency of rail service and intermodal oversight/leadership, in the form of the NEC Commission, that were deemed to be requirements for the successful funding, construction, and operation of airport rail terminals with codeshares (US Government Accountability Office, 2013).

Thus, the MSTSE stakeholders might choose to consider what bundles could feasibly include airport integration, which could depend on some or all of the physical, organizational, or funding variables. If the desired bundle *does* offer a feasible change option of this type and such an option is agreeable to the participating stakeholders, they could use it in a “carrot and stick” approach to dealing with the air industry. The “stick” is the momentum behind the original bundle, which is nominally bad for the airports and airlines; the “carrot” that can be offered up is the potential to change that bundle to one that is better for the air industry, if not as good as no rail improvements at all. If the momentum behind the bundle is strong enough to be a credible threat to the air industry, they may choose to align with the project in order to maximize their own benefit by improving their catchment area through the change option. The importance of that credible threat, for the purposes of potentially removing the obstacle of competitive stakeholders, increases the value of reaching unanimity during MSTSE. Figure 10-13 illustrates the impact a credible threat and expansion option can have on a notional “air industry” tradespace in the NEC. Note how the eliminating the original status quo and changing the BATNA to the MSTSE-stakeholders’ preferred bundle switches the expansion option from being dominated to being a potentially worthwhile tradeoff.



**Figure 10-13: Comparison of notional Air Industry tradespace if engaged during MSTSE vs. after agreement on preferred bundle**

In addition to exploring possible changeability in the bundles, it is also worth considering the ability of each bundle to circumvent post-MSTSE competitive responses during the negotiations. This information can be used as a tiebreaker between equally attractive designs, or even as a component in value functions of one or more stakeholders. In the NEC, the source of funding is an avenue for competitive stakeholders to delay or cancel any agreement. When Southwest Airlines successfully blocked the Texas HSR plan, they did so largely by exploiting the large quantity of public funding in the plan, tying it up in legislative red tape until the momentum was gone and the project was scuttled; this is a key reason that the current Texas HSR plan depends only on private investment (Wray and Nicholson, 2015). A similar move by airlines would be possible in the NEC, potentially increasing the attractiveness of bundles with lower amounts of public or state funding.

The second avenue for competitive behavior to influence MSTSE for the NEC is through the introduction of new alternatives. Recall that Bundles 8 and 9 were created by redistributing costs of existing bundles amongst the stakeholders. For example, Bundle 6 was favored by the stakeholders with public costs, who were asked to pay more in Bundle 9 to increase its attractiveness to other stakeholders. In our analysis-only example, we took the role of the NEC Commission as a party with a primary interest in reaching an agreement. But what if we took the role of the USDOT, performing informal MSTSE analysis to advise our later negotiation tactics? Bundle 6, identified as most likely to reach agreement of the original six bundles, dominates Bundle 9 for the USDOT. The USDOT has little incentive to offer Bundle 9 as an addition to the tradespace of bundles under consideration if they believe Bundle 6 can be

agreed upon; instead, they might choose to hold Bundle 9 “in their pocket,” only to be offered if negotiations are about to dissolve. This behavior is certainly realistic and illustrates how the application of the FOTE principle can often be difficult in a high-stakes negotiation.

There is no easy way to control for competitive behavior like this; without omniscient oversight of all stakeholders, it is unclear how to identify secret individual knowledge or deliberately unstated preferences/objectives. A simple response would be to accept it as a reality of negotiating between different self-interested actors: the stakeholders will likely behave this way and, ultimately, if Bundle 9 is better than the BATNA the given stakeholder will still offer it before negotiation ends and an agreement will occur. This may be acceptable but runs the risk of drawing out the negotiations unnecessarily, both delaying the implementation of the ultimate agreement (if one is reached) and possibly damaging inter-stakeholder relationships. The threat to relationships can also be exacerbated if the other stakeholders *do* choose to negotiate with FOTE and suspect that it is not reciprocated. Practically, attempts to avoid this predicament should begin with reinforcing the importance of joint modeling activities with designated time for generating and evaluating new alternatives *during* the negotiation session based on the system insights developed through MSTSE, rather than encouraging stakeholders to do such activities individually.

## **10.10 Discussion**

Summarizing the results of the Northeast Corridor case study using the key objectives outlines in the chapter introduction:

### **1. Demonstrate the use of MSTSE analysis applied to a socially complex system with few feasible alternatives but many stakeholders**

MSTSE proved to be a capable extension of an existing CLIOS Process study (as discussed in chapter 8.1), requiring minimal restructuring in order to support tradespace analysis. The NEC also presented a test case for MSTSE and its accompanying visualizations on a tradespace with very few alternatives and many stakeholders. Overall, MSTSE has been designed to target a more traditional tradespace size of hundreds to tens of thousands designs (or more). A tradespace with only six alternatives, or seven if including the shared BATNA, presents a unique challenge by limiting the amount of data available to analyze and draw conclusions from. For example, the NEC tradespace was too sparse to effectively use color or transparency in the negotiation tradespace; the large gaps between “adjacent” design points resulted in sharp jumps in transparency that reduced readability of the graph. Similarly, the attribute-by-attribute correlation plots were not even presented in this write-up, as the large discrete jumps resulted in unpredictable behavior in the correlations after isolating individual components of value. At the same time, the presence of ten stakeholders requires more time and effort spent analyzing each individual alternative of interest to understand its feasibility as a potential agreement.

Two general conclusions on this particular conflict were emergent from this case. First, having very few alternatives limits the advantages of advanced tradespace visualization techniques such as the negotiation tradespace and the gridmap. When every design is important enough to consider directly, analysis will ultimately return to specific comparisons of each alternative (such as the breakdowns of Pareto front membership amongst the stakeholders at different levels of fuzziness) that carry a higher level of detail impossible to use effectively on an entire normal-sized tradespace. These visualizations can still be used, but with only limited reward over more traditional trade study analysis techniques. Fortunately this does not present a conflict of interest, because both the TSE and negotiation aspects of MSTSE still encourage the creation of many alternatives for exploration; MSTSE will simply always be more naturally suited to problems that *are* able to create large evaluated design spaces. Second, on the other hand, the presence of many stakeholders mandates the assistance of visualizations specifically geared toward many-stakeholder analysis. Ten is far too many stakeholders to toggle through the supplementary color axes in the negotiation tradespace and maintain situational awareness of their relationships, further limiting its effectiveness in this case. The stakeholder correlation plot and gridmap, however, were very effective when scaled up to ten stakeholders: particularly the correlation plot because it does not directly visualize the challengingly small design space of the NEC.

## **2. Compare and contrast the emergent coalitions resulting from MSTSE analysis with prior NEC stakeholder analysis**

The stakeholder correlation data calculated from the tradespace was used to predict plausible coalitions in the NEC. Three coalitions naturally inclined to rally around similar bundles were identified: (1) the Department of Transportation, Amtrak, and Congress, (2) the northern and southern states, and (3) the suppliers and travelers. These coalitions were positively cross-validated with existing NEC analysis, with most stakeholders displaying similar relationships as those predicted via interest vector clustering. When consulting subject matter experts in the High-Speed Rail Research Group upon completing the analysis, the emergent coalitions based on FPN correlation were considered believable. Additionally, the FPN correlation data was able to identify a weakness in the original clustering analysis resulting in an unlikely clustering of Congress and states. This difference between the methods was attributed to a higher level of detail in the FPN correlation (via the inclusion of separate funding pools and specific alternatives), allowing MSTSE to highlight a tension in the relationship driven by divergent funding interests.

## **3. Identify key controllable design parameters that can be leveraged to adjust promising alternatives in order to create value and attract support from additional stakeholders**

The structural analysis of the problem formulation pointed to the importance of leveraging funding to create new alternatives, which proved to create additional value in the negotiation. For the NEC, we revealed considerable evidence that Bundle 6 is the most likely of the original bundles to reach agreement across the pool of stakeholders. However, the ability to leverage the assignment of funding led to the creation of Bundle 9 based on Bundle 6, which was a superior solution in terms of minimizing potential cost-benefit inefficiencies for all stakeholders. Bundles 6 and 9 are the only alternatives in the tradespace capable of uniting all three of the predicted coalitions.

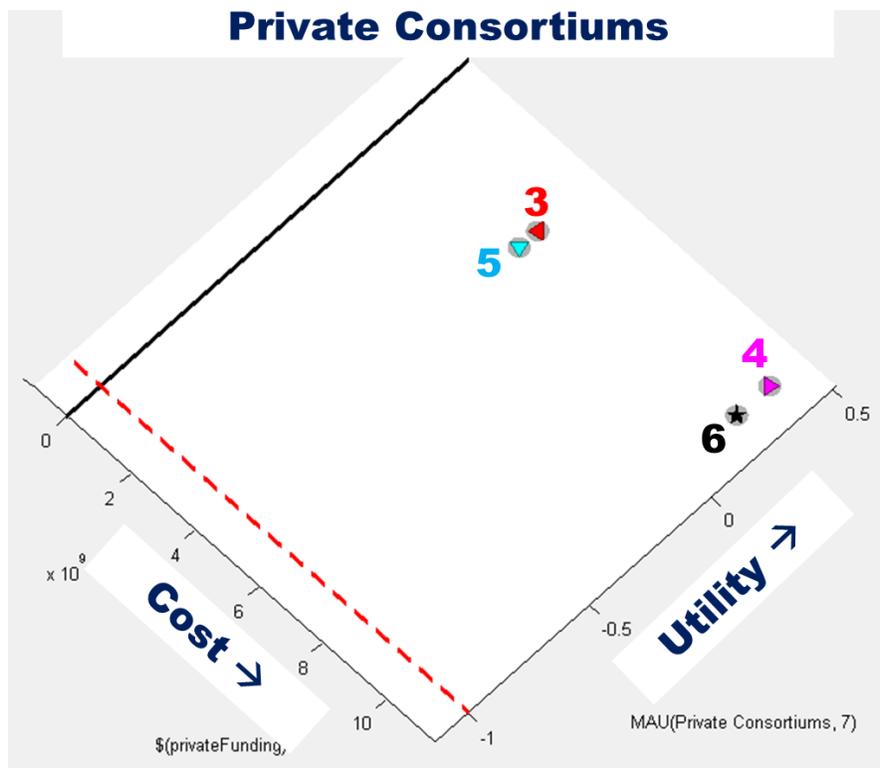
The results of this case were also brought back to the subject matter experts in the High-Speed Rail Research Group for discussion. Feedback from experts is a valuable check on the validity of the conclusions, particularly when conducting informal MSTSE without experts participating in the exploration. In addition to discussion of the validity of the conclusions, the discussion focused on the possibility of integration of MSTSE into their work. The main points of feedback were:

- Bundle 9 is potentially feasible in reality, despite its raised federal contributions to a system with competing private operators. This was attributed to a heavy emphasis on unanimity in federal transportation planning, to the point that the government will willingly pay more to attract more stakeholders. Competing private operators were also not considered to be a significant barrier to this bundle either, as there is a perceived domain bias towards privatization – a form-dependent preference described in Sclar (2001) that may be unhealthy overall but works to the advantage of negotiating an agreement here.
- Insights such as the identification of new, interesting bundles such as Bundle 9 are the most likely to impact the implementation of the CLIOS Process for the NEC or other systems. MSTSE was considered to have the most direct correspondence with the tasks of scoping bundle design and selection of alternatives for analysis, particularly between iterations of the process. Though this case study took the existing bundles and used them as the design space for the tradespace, future applications could consider applying MSTSE earlier in the process and using it to downselect into a smaller set of bundles for detailed analysis.

### **Concluding thoughts:**

In addition to the key objectives listed above, we elicited feedback from subject matter experts to identify a key area for improvement in future analysis or iterations of the CLIOS Process for the NEC. Concern was raised over the tradespace of the Private Consortiums, specifically about their apparent preference for piecewise HSR (Bundles 3 and 4) over all-over HSR (Bundles 5 and 6) as shown in Figure 10-14, which was deemed counterintuitive. The Private Consortiums' utility function used two attributes that were nominally intended to be in

tension: private returns, favoring large systems, and payback period, favoring smaller investments. The private returns were modeled as a fraction of the total financial returns corresponding to the private “stake” in the capital expenditures. This ultimately favored the piecewise HSR bundles, as the extra cost of all-over HSR diluted the private “stake” more than the additional revenues could compensate. After discussing this modeling assumption, the experts expressed interest in creating a more detailed model for private returns, accounting for organizational structure which would likely additionally favor private investors in bundles with competing private operators such as Bundle 6. Such a model could be incorporated back into the tradespace and/or BCA spreadsheet in future analysis to either verify the original insight or improve the quality of data.



**Figure 10-14: Private Consortiums tradespace: piecewise-HSR bundles (3/4) dominate all-over-HSR bundles (5/6) due to unanticipated model interaction**

Overall, the application of MSTSE to the Northeast Corridor was an effective “stress test” for the ideas of this research, where we were able to:

- Deploy our new visualizations and analysis on a dataset with fewer designs and more stakeholders than any of our previous examples.
- Demonstrate the ability to use MSTSE to supplement the insights of a structured method like the CLIOS Process with no explicit guidelines for tradespace

exploration, which is a promising sign for the breadth of potential applications for MSTSE.

- Leverage existing data to identify new alternatives of interest to the stakeholders, which is a powerful result that demonstrates the potential for tradespace exploration to improve not only the *analysis* of alternatives but also to iteratively refine the *creation* of alternatives: an underdeveloped phase of decision making in the current literature (Tang, 2006).

This chapter was the final application of MSTSE in this thesis. The remaining chapters will cover discussion points that were not applicable to a specific chapter to this point and conclude the work.



## 11 Discussion

This chapter will address final discussion points for this research on MSTSE that were not directly applicable to any of the previous chapters. The benefits of MSTSE over TSE and its applicability (as described here) to different types of stakeholders are considered. Different directions for future research are also identified, which is of particular importance given the relatively underdeveloped state of MSTSE literature as a whole.

### 11.1 *Benefits and Applicability*

As a part of the larger TSE literature, this research into MSTSE began with a set of assumptions about the benefits of exploring tradeoffs among design choices. In particular, the use of tradespaces as a means of involving stakeholders in the design process and building their intuition for complex, technical problems is viewed as a key benefit of TSE (Ross et al., 2010b). This led to the original conception of MSTSE as an interactive environment to support in-person stakeholder negotiation (Ross et al., 2010a). This research has continued with that assumption, while also acknowledging that the actual *explorers* of TSE are often not stakeholders but rather the systems engineers and analysts who create the tradespace models and datasets. When these workers do the exploration in classic TSE, they are responsible for deriving insights from the tradespace and passing them up to the stakeholders or decision makers in the form of reports and/or presentations. This format of TSE does not allow for the benefits of direct stakeholder-data interaction (intuition building, reduced iteration, etc.), but is nevertheless common.

In the shift to MSTSE, these drawbacks will typically be exacerbated as the internalized domain knowledge that makes stakeholders valuable participants in TSE is often related to social implications of the system and its relationship to the “big picture”, including the other stakeholders. However, given the logistic challenge of arranging face-to-face negotiation time amongst stakeholders, it is perhaps even *less* likely that MSTSE will see the desired level of stakeholder participation than standard TSE. This should not be taken to limit the applicability of MSTSE. The concept of deploying engineers as “proxy” stakeholders for a mock negotiation in order to capture some of the same bargaining dynamics was addressed directly in chapter 7.4, but arose in multiple places throughout the research, as noted here:

- Raised as a common topic in the practitioner interviews of chapter 5.6, with mixed opinions on the importance of stakeholder interaction
- Discussed in chapter 6.5.1 as it relates to low stakeholder participation as a structural feature of a given design effort
- Specifically noted as likely for military Analysis of Alternatives in chapter 8.2

What then are the benefits of MSTSE, particularly when stakeholders may not be available to negotiate with each other directly and thus we do not gain the benefits of stakeholders interacting with the tradespace? The case studies of chapters 9 and 10 are

instructive here. The cases were conducted on existing data sets that had previously been analyzed by other methods, with only minor additions necessary to accommodate MSTSE, and yet we were able to extract new insights about the relationships between stakeholders even without their direct participation. This outcome – the gathering of new information from the same underlying data – is a powerful result for MSTSE. It demonstrates that MSTSE allows designers to dive deeper into the social aspects of the tradespace and understand how different stakeholders relate to each other at a level more fundamental than shared favorite alternatives.

In addition to the increased detail, MSTSE can provide some needed structure to engineers wanting to model multi-stakeholder systems. The applications of MSTSE to the case studies in this research provide examples for how to incorporate the social features of a multi-stakeholder system into the tradespace formulation with relatively little overhead in addition to the standard TSE tasks of defining design variables, context variables, and other key components of any tradespace. If stakeholders *do* choose to conduct a negotiation with MSTSE as an aid, MSTSE combines the communication of technical analysis of TSE with an accessible means of comparing and contrasting stakeholders. If stakeholders *do not* participate in negotiation, as in our case studies, MSTSE can allow engineers to understand and explain at least some of the dynamics between stakeholders in the “big picture” and then communicate that information back to the stakeholders. This expands what engineers can contribute to solving multi-stakeholder problems and correspondingly reduces the amount of analysis that *must* be conducted by the decision makers who will negotiate the final solution, whether formally or informally.

So when should MSTSE be applied, particularly in contrast to other multi-stakeholder design techniques such as those discussed in the literature review? Similar to the criteria we used for selecting case studies, the problem at hand needs multiple stakeholders and should nominally be complex enough that TSE does not represent excessive effort. The main driver for MSTSE over other methods, however, is the desire to engage the stakeholders in the case *prescriptively* rather than *normatively*. Whether with game theory, dynamic programming, simulated negotiation, or another technique that “solves” the problem, algorithmic approaches to solving multi-stakeholder problems rely on the assumption that the stakeholders subscribe to the assumptions and implied objectives/values of the norm used to solve the designated question. MSTSE does not seek to replace optimization. Rather, optimization is a powerful tool in the engineer’s arsenal of analysis techniques and one that should be deployed when normative analysis is appropriate. For example, a single stakeholder with a rigorously elicited utility function could have that function optimized with considerable confidence that the resulting solution is the one he would most prefer. Or, as discussed before, a group of tightly connected stakeholders such as divisions within a single company may be able to leverage optimization to guide their strategic coordination during a design task by combining their preferences. Alternatively, engineers who are asked to perform analysis and provide *their own* recommendation for a design problem may choose to optimize to find a single, defensible solution.

It is the transition to multiple independent stakeholders that weakens the appeal of optimization. It can be challenging to get *one* stakeholder to trust the results of an optimization; getting many stakeholders to trust the *same* optimization is even more challenging. Thus, a prescriptive approach such as MSTSE is more likely to successfully engage a group of stakeholders by simply providing new and accessible information while allowing the stakeholders to still make the decision without being told the “correct” answer. This mirrors the feedback provided during our practitioner interviews, based on their experience attempting to bring technical analysis to stakeholders. At the very least, MSTSE should be considered as an alternative to optimization for approaching problems of this type, particularly when the primary challenge of the design task lies in brokering an agreement between stakeholders with highly disparate preferences who are unlikely to share a nearly-optimal solution.

## **11.2 Future Work**

This research has sought to develop the foundations of MSTSE as an extension of TSE concerned with accurately representing the needs of multiple stakeholders and enabling them to make educated decisions in a negotiation rather than as a stand-alone design task. Because this is a big first step for a technique with a very modest associated literature, there are many gaps in our knowledge of MSTSE that can be filled by future research. In the following sections we suggest a few of the many direct extensions of the current research, increasing detail in our understanding of how stakeholders interact with tradespaces. Additionally, we provide some suggestions that could help expand the applicability of MSTSE to a wider range of problem types.

### **11.2.1 Developing More Visualizations**

Certainly, one of the most straightforward ways to continue building MSTSE is to develop more tradespace visualizations specifically geared towards multi-stakeholder analysis. In this research, we identified and mitigated some framing flaws in existing visualizations as well as testing a handful of new visualizations. The new visualizations, which focused on the comparison of stakeholders rather than the comparison of designs, ended up being the most powerful additions to the MSTSE repertoire for individually analyzing the case studies but were not tested experimentally for their impact on face-to-face negotiations. Further research into developing the suite of visualizations available to prospective MSTSE users could expand the types of information available to stakeholders before reaching key decision points.

The most desirable new visualizations would target design under uncertainty. The visualizations described in chapter 5 were deliberately simple, starting only with the “basics” of tradespace exploration: single-context analysis, where each design has some fixed performance and gives each stakeholder some fixed value. We deployed those visualizations across multiple contexts in the case studies by simply replicating them when necessary, but this is an inherently limited approach. Visualizations designed to exploit data across many contexts could drastically improve negotiations by allowing stakeholders to find highly robust solutions (for one or all

parties) without the need to repeat single-context analysis. The bubble plot used in chapter 9.10 to summarize the relative impact of changeability on a small set of designs of interest is one such example, though one that is likely unable to scale up to the size of a full tradespace and remain readable. Current research into Interactive Epoch-Era Analysis, which expands tradespace uncertainty analysis using advanced insights from visual analytics, may drive improvement in this area (Curry and Ross, 2016). The results of that research thread will have positive impact for MSTSE applications that use Epoch-Era Analysis to model uncertainty, and may be extendable to other uncertainty frameworks as well.

## 11.2.2 Expanding Data Collection on Tradespace Framing

The quantity of data collected on framing for this research was modest; there remains an opportunity to improve our understanding of framing in tradespace exploration by dramatically expanding the testing of MSTSE visualizations. Increasing the sample size of mock negotiations such as those conducted in chapter 4 could build additional confidence in the use of the negotiation tradespace and its impacts on the interactions that take place during MSTSE. Testing on a wider range of variables would also increase the generalizability of the conclusions; some of the modifications identified as of particular interest were:

- Utilize real datasets with a higher level of technical complexity and a longer time available for exploration, to more closely approximate real applications
- Test the impact of training, both in TSE (a la Wolf et al., 2011) and negotiation (a la Herbst and Schwarz, 2011)
- Test thresholds for *just-noticeable differences* in TSE – how far apart do alternatives need to be in order to be different, and how much better than the BATNA must an alternative be in order to be “worth” working together?
- Test the effects of tradespace shape on perceived value
- Test if the observed bias favoring high-cost-high-benefit (“gold plated”) solutions is intrinsic or driven by framing

Widening the subject pool for experimental testing beyond student engineers and into professional engineers or even non-technical decision makers at organizations with potential MSTSE applications would also support generalizability. Conducting an experiment that isolates the different changes of the negotiation tradespace (e.g. being BATNA-centered *or* using transparency, rather than both) could identify which mechanisms are most impactful. Similar framing experiments could be designed to test any new MSTSE visualizations developed in the previous research thread. Experimental negotiation studies could also be structured to analyze specific biases, such as the perceived bias for “gold plated” solutions that are expensive but high-performing for all stakeholder, which has been noted both anecdotally (Ross et al., 2010a) and in the solution patterns for the experiment in this research. A carefully controlled study may be able to determine whether this bias is natural (e.g. an empirical human tendency to favor increasing benefits over reducing costs) or is introduced to the negotiation via framing (e.g. by

the lack of an accessible reference point or by the high frequency with which these designs are called out by screening metrics since they are always “good”).

Future research could also look to move beyond mock negotiations in order to address some of the key weaknesses of that approach. For example, data recording *during* the negotiations was limited to observational data in order to limit intrusiveness and mitigate any experiment-driven bias that could plausibly be caused or exacerbated by recording the sessions (such as acquiescence bias or friendliness bias between the participants). Individual interaction with the MSTSE visualizations could be quantified in greater detail using the methods of human-computer interface and GUI testing. In particular, mouse- or eye-tracking software could serve to measure the attention being paid to different regions of the tradespace and draw inferences about reference points and the use of the tradespace as an “attention-directing tool” in line with other multi-criteria decision making tools (Roman et al., 2004).

Another weakness of mock negotiations as a testing ground for MSTSE is the inability to replicate the *importance* of a real decision, in which stakeholders may be less likely to accept sub-optimal solutions or more likely to bargain aggressively for their own interests. A series of embedded case studies of multi-stakeholder engineering problems, some deploying MSTSE and some not, could address this weakness by offering a first demonstration of the impact of MSTSE when the consequences of the decision are “real” (Yin, 2009). Realistically, this goal lies beyond expanded testing: using MSTSE in its current form as a part of a live negotiation would require a particularly progressive set of stakeholders willing to take a risk on incorporating a new technique into their design process.

### **11.2.3 Culture of Negotiation in Engineering**

A broader study of the culture surrounding the practice of negotiation among engineers or stakeholders as a part of engineering design is a promising research path for the identification of other topics with possible macro framing impacts on MSTSE and multi-stakeholder engineering design in general. Though the commentary on macro framing within this research was taken implicitly or explicitly from primary sources (including practitioners, official SE process guidelines, and subject matter experts), the limited range of subjects falls far short of a comprehensive ethnographic study that could be generalized to all types of engineering with full confidence. The military/aerospace and infrastructure examples emphasized in this work have different relationships with negotiation: for example, “negotiation” (the word itself) is accepted as integral to infrastructure planning on the scale of the NEC, but the multi-stakeholder design task supporting Satellite Radar may not be considered a “negotiation” by some stakeholders. It is highly possible that other domains of engineering have very different relationships with negotiation than either of these examples. Some of these other domains may not find MSTSE to be a good fit. For example, since TSE is typically applied during conceptual design where leverage to explore many different solution types is highest, MSTSE would not be the best choice of technique for a domain where negotiation between stakeholders mostly occurs during

detailed design. However, it is not unrealistic to suggest that the lessons learned from studying macro framing in a wider set of domains could improve MSTSE by identifying common beliefs or practices that impact the perceived value of alternatives in a negotiation.

One macro framing topic discussed in this research that would be valuable to consider in other domains is the notion of a “compromise” in engineering design. The colloquial use of “compromise” in TSE literature to describe designs that are not optimal for a single stakeholder but rather only optimal between stakeholders is potentially a problematic influence on the reference point in the tradespace, especially given the already substantial emphasis on the individual Pareto front. Testing the impact of this terminology on negotiation outcomes with an experiment would be extremely difficult: as with all macro framing topics, it is nearly impossible to strictly control the previous training and exposure of the subjects to the idea of “compromise” in design. Instead, it would be instructive to research how pervasive the use of “compromise” (relative to hypothetical and possible infeasible design choices rather than concrete existing solutions or other BATNAs) has become in other design communities dealing with multi-stakeholder problems. It would be especially interesting to see if the presence of “compromise” in official documentation or reports of decisions was correlated with the frequency of reaching agreement or stakeholder satisfaction with the end solution.

The perceived divide between engineers and stakeholders/management is another potentially fruitful topic for further investigation. This was a common refrain in our interviews with practicing systems engineers, who expressed frustration with the challenge of communicating the results of their technical analysis to stakeholders through documents and presentations. On the surface, this has been one of the motivations for performing TSE and MSTSE with stakeholders in the loop, allowing them to develop insight into the decision interactively. However, for multi-stakeholder problems, the engineers said they generally moved *away* from interactive negotiation with the stakeholders and *towards* a system of generating a single recommendation that “balanced” the needs of all parties in order to avoid the possible interference of “personalities” in the decision. A wider sample of the systems engineering workforce would clarify how common this attitude is. Communities or organizations with a high level of cooperation between stakeholders and engineers could provide useful guidelines for structuring these interactions, and a tighter coupling between the two could eventually increase the feasibility of full stakeholder participation in MSTSE negotiations.

Finally, the culture surrounding interactive negotiation in engineering is also influenced by the infrastructure in place to support it. MSTSE has additional requirements for effective implementation beyond those of standard tradespace analysis. This includes the availability of physical space able to accommodate a negotiation between multiple stakeholders with shared access to interactive tradespace tools and visualizations. Organizations without such space will be limited in their ability to implement MSTSE as a negotiation aid. The additional modeling effort necessary to create models that are trusted by all participating stakeholders is another

barrier to applying MSTSE that could be mitigated through virtual infrastructure, specifically the creation and curation of shared, validated model libraries. These topics are both of interest to the current research thrusts of Interactive Model-Centric Systems Engineering (of the Systems Engineering Research Center; see IMCSE, 2014 and 2015) and Engineered Resilient Systems (of the Department of Defense; see Goerger et al., 2014). Further research into the ways organizations currently implement such modeling and exploring infrastructure could lead to insights that would support their broader adoption.

#### 11.2.4 Game Theory in MSTSE

Game theory has obvious applications to any negotiation-related fields and MSTSE is no exception. The assumption used by this research to skirt the application of game theoretic concepts was that stakeholders interested in applying MSTSE can and will choose to engage in Full, Open, and Truthful Exchange negotiations. This is arguably the most likely of our assumptions to be violated in practice. Though FOTE bargaining has been demonstrated to improve negotiation outcomes on integrative problems (Raiffa, 2002) and there is reason to believe that FOTE bargaining is at least plausible in some circumstances relevant to MSTSE (such as intra-government, inter-agency cooperation), there is no question that many negotiations *don't* or *can't* achieve that level of open collaboration. Therefore, research on the adaptation necessary for MSTSE to support stakeholders engaged in competitive or potentially zero-sum negotiation would enable MSTSE to be applied to a wider range of realistic problem types.

In this document, the only subjects related to competitive negotiations that were briefly discussed were (1) the use of black-boxed models to protect private information and (2) the potential use of informal MSTSE for a single stakeholder to *prepare* for a competitive negotiation by analyzing the anticipated interests, alternatives, and responses of the other stakeholders. Research emphasizing the application of MSTSE in competitive scenarios would need to address and/or utilize game theoretic principles. Mechanism design, the reverse-game-theory problem, could inform new protocols for organizing offer/counteroffer bargaining in as productive and fair a way as possible. Alternatively, algorithmic applications of game theory such as those mentioned in chapter 2.1.2 could be adapted for use within MSTSE, perhaps as a part of a new interactive visualization that suggests alternatives of interest based on an input pattern of offered alternatives by the participants.

Game theory may also have applications within MSTSE even when FOTE negotiation is feasible. For example, we assumed in this research that each stakeholder could define a BATNA by determining their intended course of action should an agreement not be reached. This assumption is challenged when the value of the alternatives outside the tradespace are dependent on the actions of the other stakeholders. For example, consider a slightly modified version of the “used car” case discussed in chapter 4 where now Vic, who wanted the car to drive home on the weekends, also has the option to move out of the apartment and back home should they not buy a car. Vic’s new individual alternative has a strongly negative impact on Nat’s prior BATNA of

riding the bus to work since Nat can't afford to live in the city (and on the bus route) without a roommate, which may make a different one of his individual options superior (perhaps sleeping on another friend's couch within walking distance of work). Suddenly, there is a traditional "game" taking place between Nat and Vic should they *not* agree. Game theory could be utilized to find equilibria in this game in order to reach an agreement on the *backup* plan should the main MSTSE negotiation not reach agreement, using this agreement as the BATNA. Additional complexity is added to this problem when non-unanimous agreements are allowed. If subgroups of stakeholders can implement designs in the tradespace and these agreements impact the choice of BATNA for the excluded stakeholders (or vice-versa, if individual BATNAs can impact the value of designs in the tradespace), then the BATNA for each stakeholder will not be determinate before the negotiations begin. Game theory may provide insights into how to quantify an indeterminate "BATNA" for each stakeholder to leverage within MSTSE.

### 11.2.5 Dividing Costs for Shared Systems

The topic of leveraging certain divisible attributes (particularly monetary costs) as additional dimensions on which to craft an agreement was raised multiple times in this research, particularly in chapter 6.2.6. We also demonstrated the use of this freedom to add new alternatives to the tradespace in the Northeast Corridor case by directly modifying existing alternatives. However, additional research efforts put towards more elaborate schemes for dividing costs could be integrated into MSTSE and potentially reveal more potentially attractive agreements in the tradespace. As previously mentioned, costs can easily be divided algorithmically to "optimize" the distribution (for a given "fairness" objective) between stakeholders in a fixed-context tradespace. Future research could seek to apply this idea across uncertainty, accounting for the possibility that some alternatives may favor different stakeholders in different contexts.

Systems that deploy changeability would strongly benefit from advanced cost-division schemes. Our analysis of changeability in Satellite Radar did not include the possibility of an unequal cost distribution between stakeholders, despite explicitly assuming that one stakeholder would "control" the system at a time. In addition to the relative *benefit* for each stakeholder, time-sharing systems such as this would likely include the anticipated relative *time in control* as a part of the objective function used to set the costs for each stakeholder. Final agreements could also expand their scope to explicitly include flexible cost divisions that would allow final costs for each stakeholder to be adjusted via side-payments in response to the actually experienced lifecycle (rather than an estimate based on an uncertainty model). Decisions of that type would be new to the TSE literature and would likely demand new visualization research to support effective exploration. When using Epoch-Era Analysis, simulation of eras (lifecycles) could provide estimates and statistics for the resulting cost distributions but would be reliant on the ability to accurately estimate the probability of occurrence for different future contexts.

## 11.2.6 MSTSE for Detailed Design

MSTSE has been scoped in this research as a technique for early concept design, selecting the high-level features of the system that feed into detailed design. Future research could seek to alter the framing of MSTSE to better apply in later design phases. In the acquisition timeline discussed in chapter 8.2, MSTSE was notionally co-located with Analysis of Alternatives as pre-Milestone-A activities. How could MSTSE contribute farther down that timeline? The challenge of increasing the level of detail contained within the tradespace is a problem for all of TSE, not just MSTSE. More detail typically requires higher fidelity models, increasing the time needed to evaluate each alternative and decreasing the range of context and/or design variables for which the models are valid, limiting the tradeoffs available to be explored. Certainly, there is interest within the TSE community for new research in the area of multi-fidelity exploration and optimization (Spero et al., 2014).

More directly relevant to MSTSE is the potential to support different types of multi-stakeholder problems that arise only later in the design process. For example, contractors are continually evaluating requests for proposal (RFPs) from a variety of different stakeholders. These RFPs roughly correspond to the *output* of early-concept design on the part of the stakeholders, who have determined what type of system they want and its approximate performance characteristics. Despite corresponding to different final products for each stakeholder, the RFPs may share many similarities. Consider a contractor who does detailed design and manufacturing for unmanned aerial vehicles (UAVs): some fraction of the open RFPs for UAVs may have approximately equal requirements for the propulsion system of the final product (e.g. turbojet, five thousand pounds of thrust). The contractor could potentially model this situation as a “negotiation” between the requirements (i.e. interests) of the different RFPs (i.e. stakeholders) and use MSTSE to search for designs that would be effective for multiple RFPs. Correlation plots could group the RFPs into “coalitions” that are appropriate to pursue as a group. These insights could be used to increase the number of RFPs the contractor can handle at one time by increasing productivity, leveraging the similarity between designs to maximize the value of the manpower spent on design. Correspondingly, this could allow for more competitive bidding on the contracts.

Contractors already perform analysis along these lines – attempting to find synergy between active contracts – but may benefit from the structure and associated visualizations of MSTSE. However, it will be necessary to adjust the framing of the MSTSE visualizations given the different type of multi-stakeholder problem in question. For example, the BATNA may be impossible to determine because there is no active negotiation in the contract-bidding system and no “individual” action for the contractor to take: the bids are submitted and the issuer makes a decision on which bid to accept by himself. Future research to develop this type of MSTSE would need to reassess the framing of the problem with respect to both the new type of multi-stakeholder decision *and* the features of high-fidelity TSE for detailed design. The research

would also benefit from collaboration with the contracting community in order to address their needs in the bidding phase and any perceived weaknesses of their current methods for bidding.

## 12 Conclusion

This final chapter concludes the research by revisiting the original research questions and assessing what has been accomplished in pursuit of their answers. Other key contributions of the research outside of the research questions are also highlighted.

### 12.1 Answering the Research Questions

We began this research with a set of four research questions targeting the role that tradespace exploration could play in the negotiations that are essential in multi-stakeholder engineering design problems. The questions are revisited here – presented with a concise answer as well as with a summary of the pertinent conclusions that have been borne out by this research.

**RQ1. Are the principles of tradespace exploration (TSE) fundamentally aligned with those of complex, sociotechnical negotiations?**

**Yes. TSE is a strong choice of analysis technique for negotiations, particularly due to its emphasis on creating and exploring many different alternatives, which is considered highly desirable in the negotiation literature.**

The first research question can be answered with a confident affirmative. A broad literature review touching on the fields of TSE, negotiation, framing, value modeling, and visual analytics found no significant conflicts between the underlying principles or active tasks of TSE and currently accepted methods for productive, mutually beneficial bargaining – i.e. for negotiations with potentially integrative or synergistic value between stakeholders. In fact, the main features of TSE have striking similarity to the cornerstones of *principled negotiation* (Fisher, Ury, and Patton, 1991) and *Full, Open, and Truthful Exchange* (Raiffa, 2002). These negotiation tenets rely on using objective measures of subjective interests, avoiding attachment to single positions, exploring many different solutions, and open sharing of information between stakeholders: all of which are needs shared or well met by TSE. Also, recent research in sociotechnical negotiation, focusing on sustainability, has provided additional evidence that the use of models can improve negotiation outcomes, particularly when they are used to test many alternatives throughout the negotiation period (Czaika, 2015). TSE is naturally positioned to capitalize on its strengths if deployed in negotiations, explaining its originally “coincidental” effectiveness when applied to multi-stakeholder problems (Ross et al., 2010a).

**RQ2. Has the evolution of multi-stakeholder tradespace exploration (MSTSE), as an offshoot of single-stakeholder TSE, resulted in unintentional framing effects impacting decision making, and can those effects be controlled?**

**Yes. The classic benefit-cost scatterplot encourages the use of the Pareto front as a reference point, which negatively frames the value of all other alternatives in the**

**tradespace. Slight modifications to the scatterplot were shown to improve the user's understanding of gains and losses as well as lead to more positive search patterns.**

The framing of classic, single-stakeholder TSE is unquestionably individualistic, which has influenced the early development of MSTSE. TSE activities are consistently geared toward defining, quantifying, visualizing, and (often eventually) optimizing the value of a specific stakeholder. At a *macro* framing level, familiarity with TSE could lead participants in MSTSE to treat the multi-stakeholder problem as multiple, individual explorations that could ultimately combine their insights into a workable single decision. Early research into the structure and effectiveness of MSTSE followed such a pattern (Ross et al., 2010a). This research has resulted in recommendations that conceptually differentiate TSE and MSTSE as early as during the problem formulation, emphasizing the role that social factors can play in the resulting negotiation and preparing stakeholders to engage with them productively. Due to resource limitations and the high difficulty of rigorously controlling macro framing (which is influenced by personal history extending far before problem formulation), we have left the study of the impact of such changes in framing to future research.

We are able to assert that *micro* framing has a noticeable impact (framing effect) on the decision making that takes place during a face-to-face MSTSE negotiation, and that this impact can be controlled by modifying the visualizations available to the participants. The traditional benefit-cost tradespace represents value information for only a single stakeholder, and we hypothesized that the Pareto front of this plot is the most accessible *reference point* for decision making. We used a controlled experiment, comprised of mock negotiations, to test the insertion of indicators for both the *best alternative to a negotiated agreement* (BATNA) as well as the value for another stakeholder. This experiment yielded significant improvements in the treatment-condition subjects' ability to indicate alternatives superior to the BATNA. It also caused a significantly different pattern of offers through the tradespace, shown in Figure 12-1, beginning near the BATNA and improving value over time rather than on the Pareto front and reducing value: a change that theoretically leads to reduced cognitive load and higher satisfaction (Gelfand et al., 2004; Gonzalez et al., 2005). Additional insights related to the preferred agreements and subject feedback were also discussed. These changes suggest that MSTSE can improve on classic TSE visualizations by deliberately controlling the way they frame the problem.

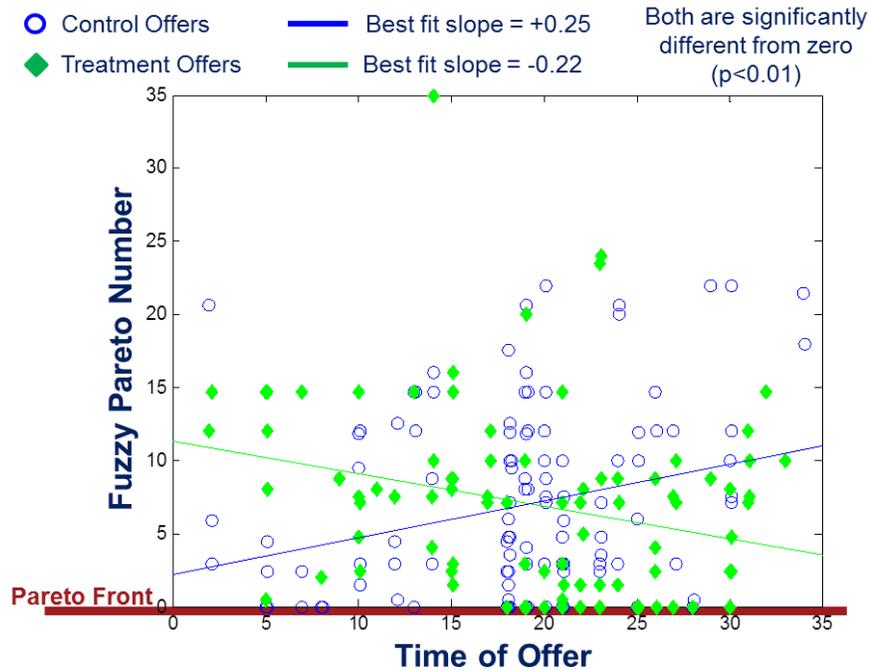


Figure 12-1: Negotiation experiment offers by FPN and time, showing change in initial reference point and search direction based on treatment

**RQ3. How can MSTSE be effectively incorporated into a design process, such that it best complements the tasks required by practicing engineers and the needs of decision makers?**

**MSTSE is relevant to the tasks of modern systems engineers and is capable of being deployed within larger systems engineering methods as a means of performing quantitative analysis of the relationships between stakeholders – even if the stakeholders themselves do not participate.**

This research has developed a set of preliminary recommendations for applying MSTSE “in the field”. Starting with the working theory of how framing can impact decision making in the tradespace and the results of our negotiation experiment, we added a variety of practical considerations that should position MSTSE to contribute effectively to a systems engineering design process. Primary sources were the most important component for answering this research question, as our visualization-aided group interviews with practicing systems engineers and critical analysis of the framing of instructions for larger systems engineering methods were enlightening of the expected role of trade studies in the design process. The interviews were integrated with the development of new visualizations that targeted insights needed by systems engineers yet not provided by their current favored techniques, focusing on the relationships between stakeholders rather than directly comparing alternatives. The SE handbooks for

applying end-to-end design methods such as Analysis of Alternatives provided some excellent examples of how the framing of the design process can either support or hinder multi-stakeholder problem solving, as well as clarifying the ways in which MSTSE could plausibly be deployed within these methods to help organizations that use them. Perhaps most critically, MSTSE enables quantitative comparison of different stakeholders' preferences, supplementing the qualitative analysis (for this type of multi-stakeholder problem) that is currently integrated with these methods.

We also used a conceptual framework for a multi-stakeholder problem to make deductions (based on the working theory) about how problem *structure* could logically affect the appropriateness of specific MSTSE tasks and visualizations. This discussion of the variation in problem structure was folded into the recommendations for conducting MSTSE as an explicit part of problem formulation: explicitly acknowledging the impact of framing and preparing to address it most effectively in the eventual negotiation. This in turn led to a discussion of *informal* MSTSE, conducted when access to stakeholders and/or decision makers is too limited for a live negotiation. This was posed as a way for engineers to extract information about multi-stakeholder problems without necessarily requiring a large time commitment from the relevant decision makers, which had been identified by practitioners as the most significant barrier to entry for applying MSTSE.

**RQ4. Can –ilities contribute to MSTSE as a potential avenue for creating mutual value and breaking impasses?**

**MSTSE can identify changeability-driven value that is experienced between stakeholders (rather than for an individual stakeholder), which potentially increases the attractiveness of less expensive solutions that are not able to simultaneously meet the demands of each stakeholder (but can change to satisfy any one).**

The application of MSTSE to the Satellite Radar case study demonstrated the potential value of changeability in a negotiation setting. Analysis across a wide uncertainty space showed that the ability to change orbit altitude and inclination was able to increase not only individual stakeholder value (as has been shown in many prior studies) but also the value for *the group* by allowing stakeholders to customize the design if they alternate individual control of the system. This did not, however, answer if changeability can break an impasse between stakeholders with divergent interests, as the stakeholders in the case had naturally aligned interests. We did demonstrate that changeability can be used to decrease the proportional performance gap between high- and low-cost solutions, as in Figure 12-2, which may prove effective at reconciling differences in more contentious negotiations where expensive “gold plated” designs were previously the only solutions to satisfy all stakeholders. We also used the Northeast Corridor case to discuss the use of change options *after* a cooperative negotiation to attract the support of competitive stakeholders with highly divergent interests. It is possible that other

–ilities could also contribute to negotiations in ways not captured in their traditional single-stakeholder analysis, but no other –ilities were tested for this research.

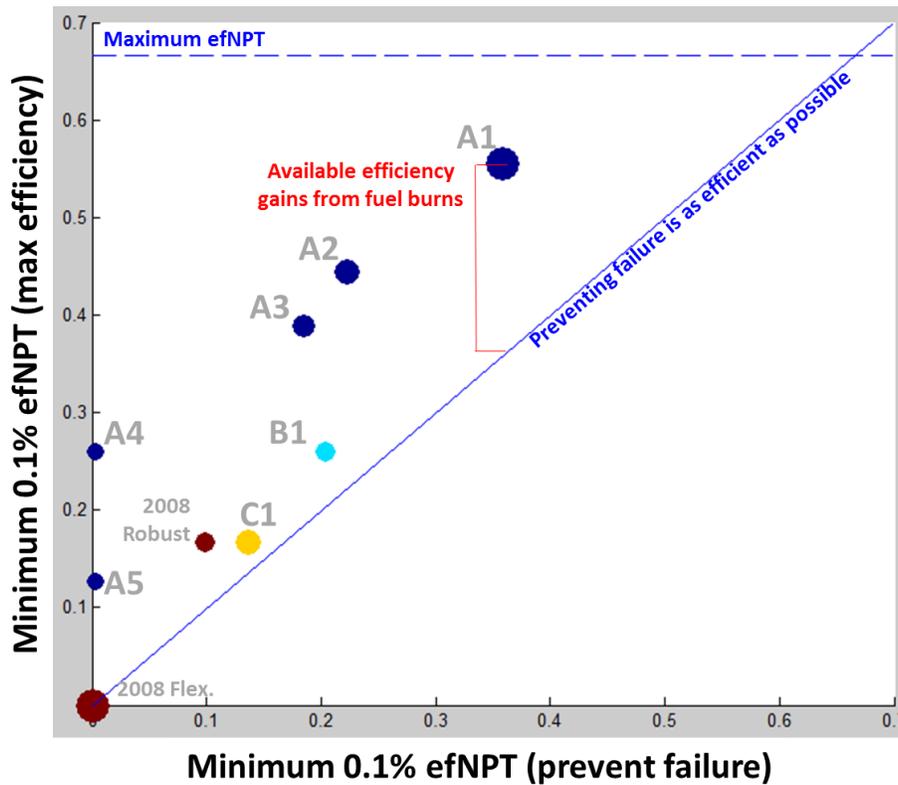


Figure 12-2: Fixed gains from changeability (height above diagonal) provide greater proportional improvement to low cost solutions (smaller circles)

## 12.2 Additional Contributions

This research is, at its core, TSE research and represents a number of first steps for the TSE community, including the first concerted efforts to incorporate fundamental insights of the negotiation and framing literatures into the activities and visualizations of multi-stakeholder TSE. In addition to these firsts and the conclusions drawn on the guiding research questions, there were a number of other outputs of this research that constitute contributions to the TSE/MSTSE literature. These contributions include:

### 1. A remodeled benefit-cost tradespace designed to support MSTSE negotiations.

The benefit-cost tradespace is the foremost visualization in TSE. The modified version used in the negotiation experiment and shown in Figure 12-3 has been shown to improve negotiators' grasp of gains and losses relative to the BATNA, resulting in more positive negotiation tactics. Though we do not assert that this visualization is "optimal" in any meaningful way (and it certainly could be improved by future research), it is demonstrably better

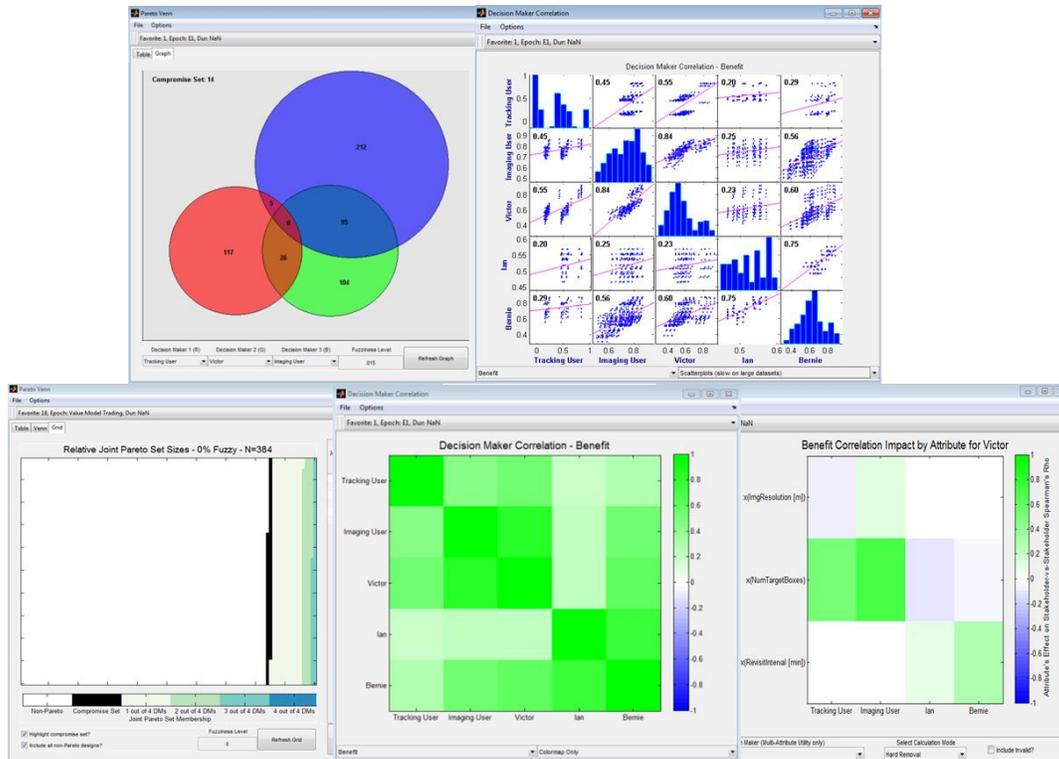
than the standard tradespace when deployed on a multi-stakeholder problem and can easily be recreated by any other MSTSE researchers or practitioners.



Figure 12-3: The "negotiation tradespace" and its key modifications from the standard tradespace

## 2. Visualization support for relationships between stakeholders.

The paradigm shift away from using TSE strictly to view alternatives and compare them should not be understated. Simultaneous visual communication of the pertinent value information for each stakeholder in a multi-stakeholder problem is extremely challenging when dealing with the thousands of designs in a normal tradespace: there is only so much information that a user can consider at once. Attempting to solve this problem with brute force may be akin to fitting a square peg in a round hole – if you want to learn about what each stakeholder wants, why not show that directly? The acknowledgement that the relationships amongst the stakeholders are as critical to understanding of the problem as the alternatives themselves is, to our knowledge, something that has received little more than lip-service in the TSE community. In support of this objective and with the cooperation of practicing systems engineers, we have created some simple visualizations that leverage a full tradespace dataset but avoid showing individual alternatives in favor of representing inter-stakeholder relationships, shown in Figure 12-4. As with the negotiation tradespace, these are also easily able to be replicated by other members of the TSE community.



**Figure 12-4: Relationship visualizations, deemphasizing the individual alternatives relative to most classic TSE visualizations**

**3. Demonstration of informal MSTSE on two case studies.**

- a. Problem structure used to guide following analysis.**
- b. Ability to extract new insights out of the same quantity of data previous used in TSE.**
- c. Ability to identify –ility contribution to group value.**
- d. Case-specific insights.**

The final step of this research was to demonstrate MSTSE using the visualizations and recommendations resulting from RQ2 and RQ3. Testing new analysis techniques supports validations and is especially important for subjective human-in-the-loop techniques such as tradespace exploration (which typically cannot be *mathematically* validated). Additionally, though full MSTSE with live negotiation was tested with success in our experiment, it was also necessary to test informal MSTSE *without* negotiation given the considerable importance attached to stakeholder-independent activities by the engineers who participated in our interview sessions. Grounding the admittedly abstract nature of most of this research using realistic applications of MSTSE has hopefully helped clarify the pros and cons of applying it. The description of our analysis was intended to be clear and detailed enough to provide a guideline for prospective adopters or researchers interested in further developing MSTSE.

The Satellite Radar and Northeast Corridor cases also offered four key contributions to this research beyond serving as instructional examples for informal MSTSE. First, they demonstrated the practical application of considering the multi-stakeholder structure of a problem during problem formulation, which was used to guide the following analysis by identifying appropriate visualizations and promising avenues for finding mutually beneficial solutions. Second, because the cases were based on existing datasets, we were able to compare the findings of MSTSE against the original analyses. This directly supported the validation of MSTSE by offering a “baseline” to compare against, while also demonstrating a key contribution of MSTSE to the TSE field: an effective way to uncover additional insights about a problem (mostly insights about the relationships between the interests of each stakeholder, as in Figure 12-5) without needing to engage in more modeling or add more data. Third, the ability to quantify the impact of –ilities at the intersection of multiple stakeholders and their different interests is a new concept for the TSE community, which to this point has considered –ilities either from a single value or value-neutral perspective. Finally, though not specifically relevant to the goals of this research, it is possible that the insights we gained about Satellite Radar and the Northeast Corridor, even with our relatively simple models, could be useful for researchers in those domains.

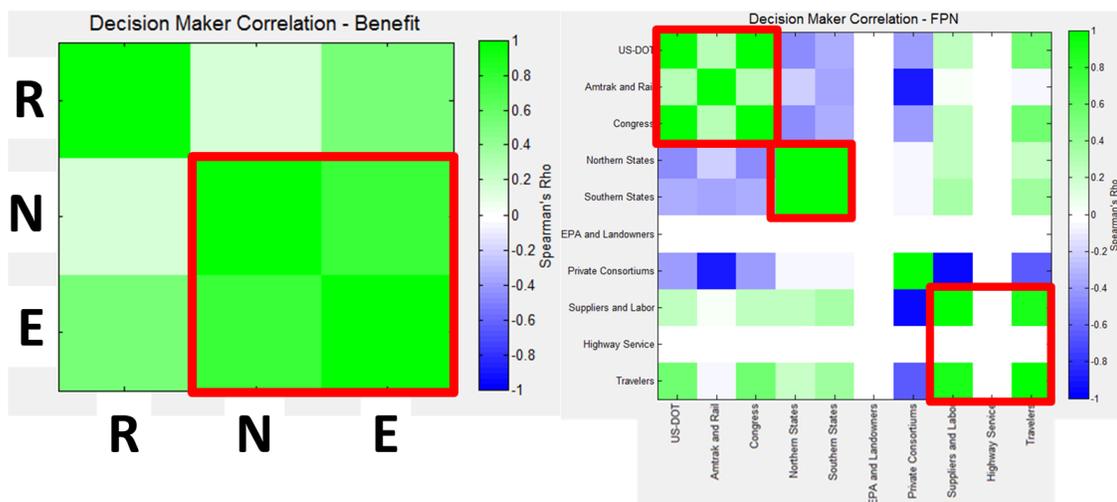


Figure 12-5: Correlation plots, identifying possible stakeholder coalitions for Satellite Radar (left) and Northeast Corridor (right)

### 12.3 Final Thoughts

This research has sought to take the ad-hoc applications of TSE to multi-stakeholder problems common to this point and supplement them with the fundamentals of the negotiation and framing literatures. Ultimately, it is our belief that MSTSE is positioned to offer a valuable service to systems engineers by guiding the incorporation of “softer” factors beyond technical design into the analysis process. The feedback that has been received both from within this research (as a part of the experiment and group interviews) and from without (from colleagues at conferences and workshops) has been positive to the point that we are confident that MSTSE is

addressing an unmet need and moving our ability to understand multi-stakeholder engineering systems in the right direction. MSTSE has many years of development ahead of it, particularly as a new technique incorporating ideas previously un- or underexplored by the systems engineering community: we hope that this research has pointed out a number of the fruitful research topics that remain and are excited to see what the future brings.



## 13 Glossary

The following pages will provide brief definitions for key terms used in this research, as well as reminders for other mentioned acronyms. These definitions are not intended to be complete descriptions nor necessarily representative of a larger research community consensus: they are simple reminders of this document's terminology for people who need one while reading.

### 13.1 Key Terms

#### **BATNA – Best Alternative to a Negotiated Agreement**

Each stakeholder in a negotiation has their own BATNA, corresponding to what they will do if no agreement is reached.

#### **Decision Maker**

A type of stakeholder characterized by the ability to directly determine the final decision of a design process (e.g. make an agreement, withdraw, etc.). In general, this research concerns stakeholders who are also decision makers unless otherwise noted.

#### **FOTE – Full, Open, and Truthful Exchange**

A negotiation philosophy characterized by telling the truth, e.g. not engaging in strategic deception. Improves outcomes for all parties in cooperative, integrative negotiations.

#### **FPN – Fuzzy Pareto Number**

A tradespace metric that measures the cost-benefit efficiency of an alternative as a normalized distance from the Pareto front.

#### **Framing**

The context in which a decision is made. Can influence decision outcomes without changing substantive issues (*framing effect*). Framing can influence decisions at both a *macro* (values, beliefs, etc.) and *micro* (interpretation, reference points, etc.) level.

#### **JFF – Joint Fact Finding**

A cooperative effort between multiple stakeholders to establish objective facts (or models, in conjunction with Collaborative Modeling) upon which to base a negotiation in order to preempt subjective arguments.

#### **MATE – Multi-Attribute Tradespace Exploration**

Tradespace exploration specifically considering the value derived from multiple system performance attributes, often with multi-attribute utility (MAU).

**MAU / MAUT – Multi-Attribute Utility / Multi-Attribute Utility Theory**

A methodology for combining the value of multiple attributes into an aggregate value (for a single stakeholder). Canonically uses expected-value lotteries to achieve a normative result, but is often elicited informally for use in prescriptive methods such as MATE.

**MSTSE – Multi-Stakeholder TradeSpace Exploration**

Tradespace exploration, with more than one stakeholder and thus more than one benefit / cost function to consider.

**NEC – NorthEast Corridor**

The transportation system servicing both passenger and freight traffic in the densely populated area of the United States stretching from Boston to Washington DC.

**Principled Negotiation**

A negotiation strategy popularized by Fisher, Ury and Patton (1991) that supports mutually beneficial agreements by evaluating choices on their objective merits and agreed-upon fairness criteria.

**Reference Point**

A concept from Prospect Theory, it is the neutral value against which gains and losses are judged. Asymmetry in the perceived impact of gains and losses makes setting an appropriate reference point important for decision making.

**Satellite Radar**

A potential system of earth-observing radar satellites, of strategic value to the defense and intelligence communities.

**SE – Systems Engineering**

A multidisciplinary field focusing on the design and management of successful systems, with its own methods and frameworks for doing so.

**Stakeholder**

A person or group with a vested interest in the outcome of a decision. This research focuses on stakeholders who are participating in the decision process, though it can also include stakeholders who are only affected.

**TSE – TradeSpace Exploration**

A modern engineering technique that explores a design space by enumerating and evaluating a large number of potential designs, including apparently sub-optimal designs

(however “optimality” may be defined), with the understanding that certain valuable behaviors may not be captured by a particular stated value metric

### **13.2 Other Acronyms**

**AFROC** – Air Force Requirements Oversight Committee

**AFRRG** – Air Force Requirements Review Group

**AISR** – Airborne Intelligence, Surveillance, and Reconnaissance

**AoA** – Analysis of Alternatives

**CLIOS** – Complex, Large-scale, Interconnected, Open, Sociotechnical

**DCF** – Discounted Cash Flow analysis

**DoD** – Department of Defense

**DOTmLPF-P** – Doctrine, Operations, Training, (existing) materiel, Leadership/education, Personnel, and Facilities Analysis

**DSM** – Design Structure Matrix

**DVM** – Design Value Mapping

**EEA** – Epoch-Era Analysis

**ELM** – Elaboration Likelihood Model

**JROC** – Joint Requirements Oversight Committee

**KPP** – Key Performance Parameter

**GMTI** - Ground-Moving Target Identification

**HSR** – High-Speed Rail

**INCOSE** – International Council on Systems Engineering

**IVTea** – Interactive Value-driven Tradespace Exploration and Analysis suite

**MDD** – Materiel Development Decision

**MODA** – Multi-Objective Decision Analysis

**MoE** – Measure of Effectiveness

**MSD** – Multi Stakeholder Dialogue

**NPV** – Net Present Value

**OAS** – Office of Aerospace Studies

**PPP** – Public Private Partnership

**QFD** – Quality Functional Deployment

**RAF** – Risk Assessment Framework

**RFP** – Request For Proposal

**ROA** – Real Options Analysis

**RSC** – Responsive Systems Comparison method

**SAI** – Systems Architecting forilities process

**SAM-PD** – Stakeholder Assisted Modeling and Policy Design

**SAR** – Synthetic Aperture Radar

**SEArI** – Systems Engineering Advancement Research Initiative

**SOGR** – State Of Good Repair

**SoS** – System of Systems

**SPO** – System Program Office

**TRL** – Technology Readiness Level

**TSAT** – Transformational Satellite System

**UAV** – Unmanned Aerial Vehicle

**USAF** – United States Air Force

**VASC** – Valuation Approach for Strategic Changeability

**VCG** – Vickrey-Clark-Groves mechanism

**VBTS** – Value Based Theory of Systems Engineering

**WGS** – Wideband Global SATCOM (satellite communications)

## 14 References

- 16.89, Space System Engineering, "X-TOS: Final Design Report." Cambridge, MA, Massachusetts Institute of Technology: 105, 2002.
- F. Ackerman and L. Heinzerling, "Pricing the Priceless: Cost-Benefit Analysis of Environmental Protection," *Univ. PA. Law Rev.*, vol. 150, no. 5, p. 1553, May 2002.
- R. Adner and D. A. Levinthal, "What is not a real option: Considering boundaries for the application of real options to business strategy." *Academy of management review*, vol. 29, no. 1, pp. 74-85, 2004.
- Air Force Space Command, "Wideband Global SATCOM Satellite," 2015. [Online]. Available: <http://www.afspc.af.mil/library/factsheets/factsheet.asp?id=5582>. [Accessed: 21-Jan-2016].
- Amtrak, "Monthly Performance Report for September 2014", 2014.
- W. P. Anderson, "Air Transportation in the Northeast Corridor: Challenges for the Future," prepared for the Northeast Corridor Transportation Project, Institute for Public Administration, University of Delaware, 2007.
- P. Antunes, R. Santos, and N. Videira, "Participatory decision making for sustainable development—the use of mediated modelling techniques", *Land Use Policy*, vol. 23, pp. 44–52, 2006.
- K. J. Arrow, "Social Choice and Individual Values" 2nd Edition, Yale University Press, 1963.
- E. Babbie, "The Practice of Social Research", Twelfth Edition, Ch. 4-5, Wadsworth publishing company, Cengage Learning, Belmont CA, 2009.
- D. Bahler, C. Dupont, and J. Bowen, "A Mixed Quantitative/Qualitative Method for Evaluating Compromise Solutions to Conflicts in Collaborative Design," *Artif. Intell. Eng. Des. Anal. Manuf.*, vol. 9, no. 04, p. 325, Feb. 1995.
- R. Balling, "Design by Shopping: A New Paradigm?," Proceedings of the Third World Congress of Structural and Multidisciplinary Optimization ?WCSMO-3?, Buffalo, NY, University at Buffalo, pp. 295–297, 1999.
- R. Bartle, "Hearts, clubs, diamonds, spades: Players who suit MUDs", *Journal of MUD research*, vol. 1, no. 1, pp. 19, 1996.
- M. H. Bazerman, J. R. Curhan, D. A. Moore, and K. L. Valley, "Negotiation," *Annu. Rev. Psychol.*, vol. 51, pp. 279–314, Jan. 2000.
- J. C. Beesemyer, "Empirically Characterizing Evolvability and Changeability in Engineering Systems," Master of Science thesis, Aeronautics and Astronautics, MIT, June 2012.
- H. R. Bernard, "Research Methods in Anthropology – Qualitative and Quantitative Approaches", Fourth Edition, AltaMira Press, Lanham MD, 2006.

- B. Boehm and A. Jain, "The Value-Based Theory of Systems Engineering: Identifying and Explaining Dependencies," in *INCOSE International Symposium 2007*, 2007.
- M. A. Brach, *Real options in practice*. Vol. 198. John Wiley & Sons, 2003.
- M. Braganca, "The History of Space Radar - A Brief Survey of Radar Capabilities," Center for Strategic and International Studies, 2011. [Online]. Available: <http://csis.org/blog/history-space-radar-brief-survey-radar-capabilities>.
- J. M. Brett, "Negotiating group decisions," *Negot. J.*, vol. 7, no. 3, pp. 291–310, 1991.
- S. Bryson, "Direct Manipulation in Virtual Reality", *The Visualization Handbook* Chapter 21, edited by C. D. Handsen and C. R. Johnson, Academic Press, 2011.
- Bureau of Transportation Statistics, "Table 4-22: Energy Intensity of Light Duty Vehicles and Motorcycles," [http://www.rita.dot.gov/bts/sites/rita.dot.gov/bts/files/publications/national\\_transportation\\_statistics/html/table\\_04\\_22.html](http://www.rita.dot.gov/bts/sites/rita.dot.gov/bts/files/publications/national_transportation_statistics/html/table_04_22.html), accessed on October 30, 2015.
- B. G. Cameron, "Value Network Modeling: A Quantitative Method For Comparing Benefit Across Exploration Architectures," Master of Science Thesis, Aeronautics and Astronautics and Engineering Systems Division, Massachusetts Institute of Technology, 2007.
- A. Campoy and C. Tsuneoka, "Texas Group Races for High-Speed Rail; Line Would Link Houston and Dallas in 90 Minutes, Take No Taxpayer Funds," *Wall Street Journal*, available at <http://www.wsj.com/articles/texas-group-races-for-high-speed-rail-1423592965>, February 10 2015.
- P. R. Carlile, "A Pragmatic View of Knowledge and Boundaries: Boundary Objects in New Product Development," *Organ. Sci.*, vol. 13, no. 4, pp. 442–455, 2002.
- P. R. Carlile and C. M. Christensen, "The Cycles of Theory Building in Management Research," Working Paper, 2004.
- A. B. Carter, "Better buying power: Guidance for obtaining greater efficiency and productivity in defense spending" [Memorandum] 2010a. Available: [http://www.acq.osd.mil/docs/USD%28AT&L%29\\_Memo\\_to\\_Acquisition\\_Professionals\\_June\\_28\\_2010.pdf](http://www.acq.osd.mil/docs/USD%28AT&L%29_Memo_to_Acquisition_Professionals_June_28_2010.pdf)
- A. B. Carter, "Better buying power: Mandate for restoring affordability and productivity in defense spending" [Memorandum] 2010b. Available: [http://www.acq.osd.mil/dpap/cpf/docs/USD\\_ATL\\_Guidance\\_Memo\\_September\\_14\\_2010\\_FINAL.PDF](http://www.acq.osd.mil/dpap/cpf/docs/USD_ATL_Guidance_Memo_September_14_2010_FINAL.PDF)
- CDM Smith, "I-95 Corridor Congestion Relief Study – Data Collection / Existing Conditions," technical memorandum for the Connecticut Department of Transportation, February 26, 2014.
- S. Chaiken, A. Liberman, and A. H. Eagly, "Heuristic and Systematic Information Processing within and beyond the Persuasion Context," in *Unintended Thought*, ed. J. S. Uleman and J. A. Bargh, Guilford Press, 1989.
- F. Chandler, D. Bienhoff, J. Cronick, and G. Grayson, "Propellant Depots for Earth Orbit and Lunar Exploration," AIAA Space 2007, Long Beach, CA, 2007.

- R. Chang, C. Ziemkiewicz, R. Pyzh, J. Kielman, and W. Ribarsky, "Learning-based evaluation of visual analytic systems," *ACM SIGCHI BELIV 2010 Work. (Beyond time errors Nov. Eval. methods Inf. Vis.*, pp. 29–34, 2010.
- J. H. Chen, K. M. Chao, N. Godwin, and V. W. Soo, "A multiple-stage cooperative negotiation," in *IEEE International Conference on eTechnology eCommerce and eService 2004*, 2004, pp. 131–138.
- M. V Cilli and G. S. Parnell, "Systems Engineering Tradeoff Study Process Framework," in *INCOSE International Symposium 2014*, 2014.
- E. Clarke, "Multipart Pricing of Public Goods," *Public Choice* vol. 11 no. 1, pp. 17–33, 1971.
- J. T. Collins, H. F. Meissinger, and R. Bell, "Small Orbit Transfer Vehicle (OTV) for On-Orbit Satellite Servicing and Resupply," 15th Annual AIAA/USU Conference on Small Satellites, 2001.
- P. D. Collopy, C. L. Bloebaum, and B. L. Mesmer, "The Distinct and Interrelated Roles of Value-Driven Design, Multidisciplinary Design Optimization, and Decision Analysis," in 12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and 14th AIAA/ISSM, 2012, no. September, pp. 1–9.
- Congressional Budget Office, "Alternatives for Military Space Radar." Congressional Budget Office report, 2007.
- K. Corcoran, "Higher Eyes in the Sky: The Feasibility of Moving AWACS and JSTARS Functions into Space." Masters Thesis, School of Advanced Airpower Studies, Air University, Maxwell Air Force Base, AL, 2000.
- L. Crump and A. I. Glendon, "Towards a Paradigm of Multiparty Negotiation," *Int. Negot.*, vol. 8, no. 2, pp. 197–234, 2003.
- J. R. Curhan, M. A. Neale, and L. Ross, "Dynamic valuation: Preference changes in the context of face-to-face negotiation," *J. Exp. Soc. Psychol.*, vol. 40, no. 2, pp. 142–151, 2004.
- M. Curry and A. M. Ross, "Designing System Value Sustainment using Interactive Epoch Era Analysis: A Case Study for On-orbit Servicing Vehicles," 14th Conference on Systems Engineering Research, Huntsville, AL, March 2016.
- E. Czaika, "Model Use in Sustainability Negotiations and Decisions," PhD Thesis, Institute for Data, Systems, and Society, Massachusetts Institute of Technology, 2015.
- N. C. Dalkey, "The Delphi method: An experimental study of group opinion," The RAND Corporation, Santa Monica, CA, 1969.
- M. J. Daskilewicz and B. J. German, "Aspects of Effective Visualization of Multidimensional Design Spaces," in *AIAA Aviation Technology, Integration, and Operations Conference*, September 2009.
- R. DeLap and S. Suhr, "Transition of Airborne Surveillance/Reconnaissance to Space." *AIAA Space Programs and Technologies Conference*, Huntsville AL, 1996.
- A. Deshmukh, M. Wortman, D. Jacques, T. Housel, K. Sullivan, and P. Collopy, "RT-18: Value of Flexibility," Phase I Progress Report, Systems Engineering Research Center, 2010.
- B. DiCicco-Bloom and B. F. Crabtree, "The qualitative research interview," *Med. Educ.*, vol. 40, no. 4, pp. 314–321, 2006.

- M. Dwyer, Z. Szajnarfarber, D. Selva, and E. F. Crawley, "The Cost of Jointness and How to Manage It," in *AIAA Space 2014*, August 2014.
- M. Dwyer and Z. Szajnarfarber, "A Framework to Assess the Impacts of Jointness," in *4th International Engineering Systems Symposium*, 2014.
- J. R. Ehrmann and B. L. Stinson, "Joint Fact-Finding and the Use of Technical Experts", in *The Consensus Building Handbook: A Comprehensive Guide to Reaching Agreement*, ed. L. E. Susskind, S. McKernan, and J. Thomas-Lamar, Sage Publications, 1999.
- M. R. Endsley, R. Hoffman, D. Kaber, and E. Roth, "Cognitive Engineering and Decision Making: An Overview and Future Course," *J. Cogn. Eng. Decis. Mak.*, vol. 1, no. 1, pp. 1–21, Mar. 2007.
- Federal Railroad Administration, "3.01 – Accident Trends – Summary Statistics", <http://safetydata.fra.dot.gov/officeofsafety/publicsite/summary.aspx>, accessed on October 30, 2015.
- Federation of American Scientists, "Discoverer II (DII) – STARLITE," maintained by Robert Sherman, <http://fas.org/spp/military/program/imint/starlight.htm>, accessed on January 17, 2016, last updated January 24, 2000.
- Z. Fang and D. DeLaurentis, "Multi-stakeholder dynamic planning of system of systems development and evolution," *Procedia Comput. Sci.*, vol. 44, Conference on Systems Engineering Research, pp. 95–104, 2015.
- W. Feng, E. F. Crawley, O. De Weck, R. Keller, and B. Robinson, "Dependency Structure Matrix Modelling for Stakeholder Value Networks," 12th Int. Depend. Struct. Model. Conf., no. July, pp. 3–16, 2010.
- W. Feng, E. F. Crawley, O. L. De Weck, D. R. Lessard, and B. G. Cameron, "Understanding the Impacts of Indirect Stakeholder Relationships - Stakeholder Value Network Analysis and Its Application to Large Engineering Projects," in Strategic Management Society (SMS) 32nd Annual International Conference, 2012, pp. 1–37.
- B. Fischhoff, "Value Elicitation: Is There Anything in There?," *Am. Psychol.*, vol. 46, no. 8, pp. 835–847, 1991.
- R. Fisher, W. Ury, and B. Patton, "Getting to Yes: Negotiating an Agreement without Giving In", Penguin Books, New York NY 1991.
- M. E. Fitzgerald and A. M. Ross, "Mitigating Contextual Uncertainties with Changeability Analysis in the Multi-Epoch Domain," 6th Annual IEEE Systems Conference, Vancouver Canada, 2012.
- M. E. Fitzgerald, A. M. Ross, and D. H. Rhodes, "Assessing Uncertain Benefits: a Valuation Approach for Strategic Changeability (VASC)," INCOSE International Symposium 2012, Rome, Italy, July 2012.
- M. E. Fitzgerald and A. M. Ross, "Guiding Cooperative Stakeholders to Compromise Solutions Using an Interactive Tradespace Exploration Process," 11th Conference on Systems Engineering Research, Atlanta, GA, March 2013.
- M. E. Fitzgerald and A. M. Ross, "Controlling for Framing Effects in Multi-Stakeholder Tradespace Exploration," Conference on Systems Engineering Research, Redondo Beach CA, 2014.
- French Air Transport Directorate, "Thematic Report – Air traffic between France and the United Kingdom: air transport between the Regions, train between capitals," December 2008.

- E. Fricke and A. P. Schulz, "Design for changeability (DfC): Principles to enable changes in systems throughout their entire lifecycle," *Systems Engineering*, vol. 8, no. 4, pp. 342-359, 2005.
- M. Garber, S. Sarkani, and T. A. Mazzuchi, "Multi-Stakeholder Trade Space Exploration Using Group Decision Making Methodologies," in *25th Annual INCOSE International Symposium*, 2015.
- R. M. Gates, "Defense Budget Recommendation Statement," Delivered April 6, 2009. [Online]. Available: <http://www.globalsecurity.org/military/library/news/2009/04/dod-speech-090406.htm>. [Accessed: 21-Jan-2016].
- N. Gatti and F. Amigoni, "An approximate Pareto optimal cooperative negotiation model for multiple continuous dependent issues," in *IEEEWICACM International Conference on Intelligent Agent Technology*, 2005, pp. 565-571.
- M. J. Gelfand, A. C. Fulmer, and L. Severance, "The Psychology of Negotiation and Mediation," in in *APA Handbook of Industrial and Organized Psychology*, Vol 3, 2004, pp. 495-554.
- B. J. German, K. M. Feigh, and M. J. Daskilewicz, "An Experimental Study of Continuous and Discrete Visualization Paradigms for Interactive Trade Space Exploration," *J. Comput. Inf. Sci. Eng.*, vol. 13, no. 2, Apr. 2013.
- J. J. Gibson, "The Ecological Approach to Visual Perception." Boston: Houghton Mifflin, 1979.
- S. R. Goerger, A. M. Madni, and O. J. Eslinger, "Engineered resilient systems: A DoD perspective," *Procedia Comput. Sci.*, vol. 28, pp. 865-872, 2014.
- A. Golkar and E. F. Crawley, "A Framework for Space Systems Architecture under Stakeholder Objectives Ambiguity," *Syst. Eng.*, vol. 17, no. 4, pp. 479-502, 2014.
- C. Gonzalez, J. Dana, H. Koshino, and M. Just, "The framing effect and risky decisions: Examining cognitive functions with fMRI," *J. Econ. Psychol.*, vol. 26, no. 1, 2005.
- M. Grace, B. Norvell, K. Higgins, M. Gilbert, and H. Kazemi, "Active lens: a mass, volume, and energy efficient antenna for space-based radar," *Proc. 2004 IEEE Radar Conf. (IEEE Cat. No.04CH37509)*, vol. 0048, pp. 1-5, 2004.
- F. L. Greitzer, C. F. Noonan, and L. R. Franklin, "Cognitive foundations for visual analytics," Pacific Northwest National Laboratory, PNNL-20207, Prepared for the U.S. Department of Energy, 2011.
- P. T. Grogan, "Interoperable Simulation Gaming for Strategic Infrastructure Systems Design", PhD Thesis, Engineering Systems Division, Massachusetts Institute of Technology, 2014.
- T. Groves, "Incentives in Teams," *Econometrica*, vol. 41 no. 4, pp. 617-631, 1973.
- L. Harrison, F. Yang, S. Franconeri, and R. Chang. "Ranking visualizations of correlation using weber's law." *IEEE Transactions on Visualization and Computer Graphics* (2014); 20(12): 1943-1952.
- A. Hassidim and Y. Singer, "Submodular Optimization under Noise," working paper, 2015.

- G. A. Hazelrigg, "A framework for decision-based engineering design," *J. Mech. Des.*, vol. 120, no. 4, pp. 653–658, 1998.
- G. A. Hazelrigg, "Validation of Engineering Design Alternative Selection," *Eng. Optim.*, vol. 35, no. 2, pp. 103–120, 2003.
- U. Herbst and S. Schwarz, "How Valid Is Negotiation Research Based on Student Sample Groups? New Insights into a Long-Standing Controversy," *Negotiation Journal* vol. 27, no. 2, pp. 147–170, 2011.
- R. Holden, "Supermajority Voting Rules," SSRN Electron. J., pp. 1–33, 2004.
- E. Horowitz, J. Lee, and J. Lee, "WinWin: a system for negotiating requirements," *Softw. Eng.*, pp. 646–649, 1999.
- IMCSE, Phase One Technical Report SERC-2014-TR-048-1, September 2014.
- IMCSE, Phase Two Technical Report SERC-2015-TR-048-2, February 2015.
- INCOSE, "Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities", version 3.2.2, San Diego, CA, USA: International Council on Systems Engineering (INCOSE), INCOSE-TP-2003-002-03.2.2, 2012.
- INCOSE, "Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities", version 4.0. Hoboken, NJ, USA: John Wiley and Sons, Inc, ISBN: 978-1-118-99940-0, 2015.
- S. Islam and L. E. Susskind, "Water Diplomacy: A Negotiated Approach to Managing Complex Water Networks," RFF Press, 2013.
- D. Kahneman and A. Tversky, "Choices, Values, and Frames." Cambridge University Press, 2000.
- D. Kahneman, "Thinking Fast and Slow," Farrar Straus and Giroux, 2011.
- H. A. Karl, L. E. Susskind, and K. H. Wallace, K.H., "A Dialogue, Not a Diatribe: Effective Integration of Science and Policy through Joint Fact Finding", *Environment: Science and Policy for Sustainable Development*, vol. 49, no. 1, pp. 20–34, 2007.
- T. Y. Katz-Navon and C. Goldschmidt, "Goal orientations in negotiations: the influence of goal orientations on fixed-pie perceptions and bargaining outcomes.," *Int. J. Psychol.*, vol. 44, no. 1, pp. 60–70, Mar. 2009.
- R. L. Keeney, "Value Focused Thinking: A Path to Creative Decision-making," Harvard University Press, Cambridge, Massachusetts, 1992.
- R. L. Keeney and H. Raiffa, "Decisions with multiple objectives—preferences and value tradeoffs," Cambridge University Press, Cambridge & New York, 1993.
- D. A. Keim, F. Mansmann, J. Schneidewind, J. Thomas, H. Ziegler, "Visual Analytics: Scope and Challenges," in *Visual Data Mining*, ed. S. Simoff, M. H. Bohlen, and A. Mazeika, Springer, 2008.
- F. Kendall, "Value Engineering (VE) and Obtaining Greater Efficiency and Productivity in Defense Spending." Memorandum from the Office of the Under Secretary of Defense, 2011.

- F. Kendall, "Better Buying Power 3.0 White Paper," Office of the Under Secretary of Defense, Acquisition, Technology and Logistics, 2014.
- M. Klein, P. Faratin, H. Sayama, and Y. Bar-Yam, "Negotiating Complex Contracts," *Gr. Decis. Negot.*, vol. 12, no. 2, 2003.
- M. Kouwenhoven, "The Role of Accessibility in Passengers' Choice of Airports," Joint Transport Research Centre, Discussion Paper 2008-14, prepared for the round table on Airline Competition, Systems of Airports, and Intermodal Connections, August 2008.
- T. S. Kuhn, "The Structure of Scientific Revolutions," University of Chicago Press, 1962.
- A. Kusiak and J. Wang, "Negotiation in engineering design," *Group Decision and Negotiation*, vol. 3, iss. 1, pp 69-91, April 1994.
- S. Langsdale, A. Beall, E. Bourget, E. Hagen, S. Kudlas, R. Palmer, D. Tate, and W. Werick, W., "Collaborative Modeling for Decision Support in Water Resources: Principles and Best Practices", *Journal of the American Water Resources Association*, vol. 49, no. 3, pp. 629-638, 2013.
- K. N. Lee, "Compass and Gyroscope: Integrating Science and Politics for the Environment," Island Press, 1994.
- I. Levin, S. Schneider, and G. Gaeth, "All Frames Are Not Created Equal: A Typology and Critical Analysis of Framing Effects," *Organ. Behav. Hum. Decis. Process.*, vol. 76, no. 2, pp. 149-188, Nov. 1998.
- H. A. Linstone and M. Turoff, "Delphi Method: Techniques and Applications", Addison-Wesley Publishing, Boston, MA, 1975.
- T. Litman, "Generated Traffic and Induced Travel – Implications for Transport Planning," Victoria Transport Policy Institute white paper, January 27, 2015.
- Z. Liu and J. Heer, "The Effects of Interactive Latency on Exploratory Visual Analysis," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 12, 2014.
- Los Angeles Air Force Base, "Factsheets: Space Radar," *Los Angeles Air Force Base Website*, 2011. [Online]. Available: <http://www.losangeles.af.mil/library/factsheets/factsheet.asp?id=5308>.
- A.V. Lotov, V.A. Bushenkov, G.K. Kamenev, "Interactive Decision Maps: Approximation and Visualization of Pareto Frontier," Springer, 2004.
- S. Lu, W. Elmaraghy, G. Schuh, and R. Wilhelm, "A Scientific Foundation of Collaborative Engineering," *CIRP Ann. - Manuf. Technol.*, vol. 56, no. 2, pp. 605-634, Jan. 2007.
- D. N. Mavris, O. J. Pinon, and D. Fullmer Jr., "Systems Design and Modeling: A Visual Analytics Approach," *27th Congr. Int. Council. Aeronaut. Sci. ICAS*, pp. 1-28, 2010.
- I. S. Mayer, I.S., "The Gaming of Policy and the Politics of Gaming: A Review", *Simulation & Gaming*, vol. 40, no. 6, pp. 825-862, 2009.
- H. McManus and T. Schuman, "Understanding the Orbital Transfer Vehicle Trade Space," AIAA Space 2003, Long Beach, CA, September 2003.

- H. L. McManus and D. E. Hastings, "A Framework for Understanding Uncertainty and Its Mitigation and Exploitation in Complex Systems," *IEEE Engineering Management Review*, Vol. 34, No. 3, pp. 81-94, Third Quarter 2006.
- H. L. McManus, M. G. Richards, A. M. Ross, and D. E. Hastings, "A Framework for Incorporating "ilities" in Tradespace Studies," *AIAA Space 2007*, Long Beach, CA, September 2007.
- B. L. Mesmer, C. L. Bloebaum, and H. Kannan, "Incorporation of Value-Driven Design in Multidisciplinary Design Optimization," in *10th World Congress on Structural and Multidisciplinary Optimization*, 2013, pp. 1–19.
- R. K. Mitchell, B. R. Agle, and D. J. Wood, "Toward a Theory of Stakeholder Identification and Salience: Defining the Principle of Who and What Really Counts," *Acad. Manag. Rev.*, vol. 22, no. 4, pp. 853–886, 1997.
- J. C. Moody, "Development of a Predictive Coalition Building Analysis for Stakeholders of Sociotechnical Systems," Master of Science Thesis, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, 2016.
- A. Mostashari and J. M. Sussman, "Engaging Stakeholders in Engineering Systems Representation and Modeling," in *Engineering Systems Symposium*, 2004, pp. 1–15.
- A. Mostashari, "Stakeholder-Assisted Modeling and Policy Design Process for Engineering Systems," PhD Thesis, Engineering Systems Division, Massachusetts Institute of Technology, 2005.
- G. Muller, "Reflective Practice to Connect Theory and Practice; Working and Studying Concurrently," *Procedia Comput. Sci.*, vol. 44, no. 1877, pp. 679–688, 2015.
- National Aeronautics and Space Administration, *Systems Engineering Handbook*, Rev. 1, Washington DC, 2007.
- National Aeronautics and Space Administration, "Standard For Models And Simulations," NASA-STD-7009, Washington DC, 2008.
- National Aeronautics and Space Administration, "James Webb Space Telescope – Technical Frequently Asked Questions (FAQ)," 2015. [Online] Available: [http://jwst.nasa.gov/faq\\_scientists.html#cost](http://jwst.nasa.gov/faq_scientists.html#cost).
- NEC Master Plan Working Group, "The Northeast Corridor Infrastructure Master Plan," 2010.
- D. A. Norman, *The psychology of everyday things*. New York: Basic Books, 1988.
- P. C. Nutt, "Surprising but true: Half the decisions in organizations fail.," *Acad. Manag. Perspect.*, vol. 13, no. 4, pp. 75–90, 1999.
- P. C. Nutt, "Investigating the success of decision making processes," *J. Manag. Stud.*, vol. 45, no. March, pp. 425–455, 2008.
- Office of Aerospace Studies, "Analysis of Alternatives (AoA) Handbook," 2013.

- Office of Highway Policy Information, “Travel Trends – December 2014 – Travel Monitoring and Traffic Volume,” [http://www.fhwa.dot.gov/policyinformation/travel\\_monitoring/14dectvt/page2.cfm](http://www.fhwa.dot.gov/policyinformation/travel_monitoring/14dectvt/page2.cfm), accessed on October 30, 2015.
- M. G. O’Neill, “An Approach to Analyze Tradeoffs for Aerospace System Design and Operation,” PhD Thesis, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 2013.
- S. Osterweil, “Real Games for Real STEM Education,” Presentation to the DOEMSP Regional Meeting, Jan 13 2010.
- C. P. Ozawa, “Recasting Science: Consensual Procedures in Public Policy Making,” Westview Press, 1991.
- O. C. Ozcanli, Y. Dong, J. L. Mundy, H. Webb, R. Hammoud, and T. Victor, “Automatic geo-location correction of satellite imagery,” *IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. Work.*, pp. 307–314, 2014.
- S. Pace, G. P. Frost, I. Lachow, D. R. Frelinger, D. Fossum, D. Wassem, M. M. Pinto, “The Global Positioning System – Assessing National Policies,” Appendix B, GPS History, Chronology, and Budgets, Rand Corporation, 1995.
- G. S. Parnell, M. V Cilli, and D. Buede, “Tradeoff Study Cascading Mistakes of Omission and Commission,” in *INCOSE International Symposium*, Las Vegas NV, 2014.
- K. Pedersen, J. Emblemsvåg, R. Bailey, J. K. Allen, and F. Mistree, “Validating Design Methods & Research: The Validation Square,” in *ASME Design Engineering Technical Conferences*, 2000.
- PennDesign, “High-Speed Rail in the Northeast Megaregion: From Vision to Reality,” University of Pennsylvania School of Design, Philadelphia PA, 2011.
- PennDesign, “Early Actions for High Speed Rail,” University of Pennsylvania School of Design, Philadelphia PA, 2012.
- R. Petty and J. T. Cacioppo, “Communication and Persuasion: Central and Peripheral Routs to Attitude Change,” Springer, 1986.
- S. Postrel, “Islands of shared knowledge: Specialization and mutual understanding in problem-solving teams,” *Organ. Sci.*, vol. 13, no. 3, pp. 303–320, 2002.
- Public-Private Infrastructure Advisory Facility, World Bank, and Transportation Research Support, “Railway Reform Toolkit,” 2011.
- W. E. Purdy, P. W. Gaiser, G. A. Poe, E. A. Uliana, T. Meissner, and F. J. Wentz, “Geolocation and pointing accuracy analysis for the windSat sensor,” *IEEE Trans. Geosci. Remote Sens.*, vol. 44, no. 3, pp. 496–504, 2006.
- H. Raiffa, “The Art and Science of Negotiation,” Harvard University Press, 1982.
- H. Raiffa, “Negotiation Analysis: The Science and Art of Collaborative Decision Making,” Belknap Press of Harvard University Press, 2002.

- J. Rawls, "A Theory of Justice," Harvard University Press, Cambridge MA, 1971.
- D. H. Rhodes, A. M. Ross, and D. E. Hastings, "Responsive Systems Comparison Method for Performance at Less Cost". *DII Year One Final Report*, 2008.
- D. H. Rhodes, A. M. Ross, and D. E. Hastings, "Responsive Systems Comparison Method for Performance at Less Cost". *DII Year Two Final Report*, 2010.
- N. Ricci, M. E. Fitzgerald, A. M. Ross, and D. H. Rhodes, "Architecting Systems of Systems with Ilities: An Overview of the SAI Method," 12th Conference on Systems Engineering Research, Redondo Beach, CA, March 2014.
- M. G. Richards, L. Viscito, A. M. Ross, and D. E. Hastings, "Distinguishing Attributes for the Operationally Responsive Space Paradigm," in Responsive Space 6 Conference, 2008.
- M. G. Richards, A. M. Ross, D. E. Hastings, and D. H. Rhodes, "Survivability Design Principles for Enhanced Concept Generation and Evaluation," INCOSE International Symposium 2009, Singapore, July 2009.
- M. G. Richards, A. M. Ross, D. Stein, and D. E. Hastings, "Multi-Attribute Tradespace Exploration for Survivability: Application to Satellite Radar," AIAA Space 2009, Pasadena, CA, September 2009.
- M. G. Richards, "Multi-Attribute Tradespace Exploration for Survivability," Doctor of Philosophy Dissertation, Engineering Systems Division, MIT, June 2009.
- F. Roman, N. Rolander, M. G. Fernández, B. Bras, J. K. Allen, F. Mistree, P. Chastang, P. Depince, and F. Bennis, "Selection without Reflection is a Risky Business..." in AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, 2004.
- E. Romanhuki, M. Fuckner, F. Enembreck, B. Avila, and E. E. Scalabrin, "Improving bilateral negotiation with evolutionary learning," in *2008 12th International Conference on Computer Supported Cooperative Work in Design*, 2008, pp. 343–348.
- A. M. Ross, D. E. Hastings, J. M. Warmkessel, and N. P. Diller, "Multi-Attribute Tradespace Exploration as Front End for Effective Space System Design," *Journal of Spacecraft and Rockets*, 41 (1), 2004.
- A. M. Ross and D. E. Hastings, "The Tradespace Exploration Paradigm," in *INCOSE International Symposium 2005*, Rochester, NY, 2005.
- A. M. Ross, "Managing Unarticulated Value: Changeability in Multi-Attribute Tradespace Exploration," PhD Thesis, Engineering Systems Division, Massachusetts Institute of Technology, 2006.
- A. M. Ross, and D. H. Rhodes, "Using Natural Value-centric Time Scales for Conceptualizing System Timelines through Epoch-Era Analysis," INCOSE International Symposium 2008, Utrecht, the Netherlands, June 2008.
- A. M. Ross, H. L. McManus, A. Long, M. G. Richards, D. H. Rhodes, and D. E. Hastings, "Responsive Systems Comparison Method: Case Study in Assessing Future Designs in the Presence of Change," AIAA Space 2008, San Diego, CA, September 2008.

- A. M. Ross, D. H. Rhodes, and D. E. Hastings, "Using Pareto Trace to Determine System Passive Value Robustness," *3rd Annual IEEE Systems Conference*, pp. 285-290, Vancouver, Canada, March 2009.
- A. M. Ross, H. L. McManus, D. H. Rhodes, D. E. Hastings, and A. M. Long, "Responsive Systems Comparison Method: Dynamic Insights into Designing a Satellite Radar System," *AIAA Space 2009*, Pasadena, CA, 2009.
- A. M. Ross, H. L. McManus, D. H. Rhodes, and D. E. Hastings, "A Role for Interactive Tradespace Exploration in Multi-Stakeholder Negotiations," *AIAA Space 2010*, Anaheim, CA, September 2010a.
- A. M. Ross, H. L. McManus, D. H. Rhodes, and D. E. Hastings, "Revisiting the Tradespace Exploration Paradigm: Structuring the Exploration Process," *AIAA Space 2010*, Anaheim, CA, September 2010b.
- A. M. Ross, M. G. O'Neill, D. E. Hastings, and D. H. Rhodes, "Aligning Perspectives and Methods for Value-Driven Design," *AIAA Space 2010*, Anaheim, CA, September 2010c.
- A. M. Ross, D. H. Rhodes, and M. E. Fitzgerald, "Interactive Value Model Trading for Resilient Systems Decisions," *13th Conference on Systems Engineering Research*, Hoboken, NJ, March 2015.
- J. Saleh, G. Mark, and N. Jordan, "Flexibility: a multi-disciplinary literature review and a research agenda for designing flexible engineering systems," *J. Eng. Des.*, vol. 20, no. 3, pp. 307–323, Jun. 2009.
- Satellite Today* Staff, "Space Radar Program Cancelled," *Satellite Today*, 2008. [Online]. Available: <http://www.satellitetoday.com/telecom/2008/03/07/space-radar-program-cancelled/>.
- M. A. Schaffner, M. S. Wu, A. M. Ross, and D. H. Rhodes, "Enabling Design for Affordability: An Epoch-Era Analysis Approach," *Proceedings of the 10th Annual Acquisition Research Symposium- Acquisition Management*, April 2013.
- M. A. Schaffner, A. M. Ross, and D. H. Rhodes, "A Method for Selecting Affordable System Concepts: A Case Application to Naval Ship Design," *12th Conference on Systems Engineering Research*, Redondo Beach, CA, March 2014.
- M. A. Schaffner, "Designing Systems for Many Possible Futures: The RSC-based Method for Affordable Concept Selection (RMACS), with Multi-Era Analysis", Master of Science Thesis, Aeronautics and Astronautics, MIT, June 2014.
- K. Schmeer, "Guidelines for Conducting a Stakeholder Analysis," *Report for Partnerships for Health Reform*, Abt Associates Inc., Bethesda, MD, 1999.
- D. Schon, "The Reflective Practitioner: How Professionals Think in Action," Basic Books, 1983.
- D. Schon and M. Rein, "Frame Reflection: Toward the Resolution of Intractable Policy Controversies," Basic Books, 1994.
- E. Sclar, "You Don't Always Get What You Pay For: The Economics of Privatization," Cornell University Press, 2001.
- M. J. Scott and E. K. Antonsson, "Formalisms for Negotiation in Engineering Design," in *8th International Conference on Design Theory and Methodology*, 1996, pp. 165–182.

- M. J. Scott and E. K. Antonsson, "Arrow's Theorem and Engineering Design Decision Making," *Res. Eng. Des.*, vol. 11, no. 4, pp. 218–228, 2000.
- SEARI, Feedback and action items from sponsor visit, Unpublished manuscript, January 2010.
- N. B. Shah, J. Wilds, L. Viscito, A. M. Ross, and D. E. Hastings, "Quantifying Flexibility for Architecting Changeable Systems," in *Conference on Systems Engineering Research*, 2008.
- D. Shahan and C. C. Seepersad, "Bayesian Networks for Set-Based Collaborative Design," in *ASME 2009 International Design Engineering Technical Conferences*, 2009.
- J. Shanteau, "Competence in experts: The role of task characteristics," *Organ. Behav. Hum. Decis. Process.*, vol. 53, no. 2, 1992.
- H. A. Simon, "Models of Man," Wiley, 1957.
- B. Smiley, "The absolute and relative geolocation accuracies of QB02 and WV01," *ASPRS 2009 Annu. Conf. Balt. Maryl.*, pp. 9–13, 2009.
- SNCF, "Sécurité de l'exploitation sur le réseau ferré national Français – Rapport annuel 2013", available from [http://medias.sncf.com/sncfcom/pdf/transparence/RAPPORT\\_SECURITE\\_2013.pdf](http://medias.sncf.com/sncfcom/pdf/transparence/RAPPORT_SECURITE_2013.pdf), 2013.
- T. J. Spaulding, "Tools for Evolutionary Acquisition: A Study of Multi-Attribute Tradespace Exploration (MATE) Applied to the Space Based Radar (SBR)", Master of Science Thesis, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, June 2003.
- E. Spero, C. L. Bloebaum, B. J. German, A. Pyster, and A. M. Ross, "A Research Agenda for Tradespace Exploration and Analysis of Engineered Resilient Systems," 12th Conference on Systems Engineering Research, Redondo Beach, CA, March 2014.
- H. I. J. Spoelstra and W. D. Pienaar, *Negotiation: Theories, Strategies and Skills*. Juta and Co. Ltd, 1999.
- G. Stump, S. Lego, M. Yukish, T. W. Simpson, and J. A. Donndelinger, "Visual Steering Commands for Trade Space Exploration: User-Guided Sampling With Example," *J. Comput. Inf. Sci. Eng.*, vol. 9, no. 4, 2009.
- N. P. Suh, "Axiomatic design theory for systems," *Res. Eng. Des. - Theory, Appl. Concurr. Eng.*, vol. 10, no. 4, pp. 189–209, 1998.
- L. Susskind, D. Fairman, M. Ferenz, and B. Fuller, "Multistakeholder Dialogue at the Global Scale," *Int. Negot.*, vol. 8, no. 2, pp. 235–266, 2003.
- J. M. Sussman, J. B. McConnell, A. Mostashari, and S. Sgouridis, "CLIOS User's Guide." 2009.
- J. M. Sussman, "Transportation in the Northeast Corridor of the U.S.: A Multimodal and Intermodal Conceptual Framework," Research report for the Institution for Transportation Policy Studies, 2012.
- J. M. Sussman, J. Moody, M. Pena-Alcaraz, and P. Doyle, "The Evaluation of International HSR Markets for JR East: A Systems Approach. Volume 3 - The CLIOS Process Applied to the Northeast Corridor (NEC)," Research report for the East Japan Railway Company, 2015.

- K. Sycara, "Utility Theory in Conflict Resolution," *Ann. Oper. Res.*, vol. 12, pp. 65–83, 1988.
- V. Tang, "Corporate Decision Analysis: An Engineering Approach," PhD Thesis, Engineering Systems Division, Massachusetts Institute of Technology, 2006.
- J. J. Thomas and K. A. Cook, "Visualization Viewpoints: A Visual Analytics Agenda," *IEEE Comput. Graph. Appl.*, vol. 26, no. February, pp. 10–13, 2006.
- J. A. Tirpak, "The Space Based Radar Plan," *Air Force Magazine*, August, 2002.
- A. Tversky and D. Kahneman, "Judgment under Uncertainty: Heuristics and Biases," *Science*, vol. 185, no. 4157, pp. 1124–31, Sep. 1974.
- A. Tversky and D. Kahneman, "The Framing of Decisions and the Psychology of Choice," *Science*, vol. 211, no. 4481, pp. 453–458, 1981.
- US Air Force, "Analysis of Alternatives (AoA) Handbook: A Practical Guide to Analyses of Alternatives," 2010.
- US Department of Transportation and Federal Railroad Administration, "NEC Future Tier 1 Draft Environmental Impact Statement," November 2015.
- US Government Accountability Office, "Space-Based Radar Effort Needs Additional Knowledge before Starting Development," Report to Congressional Committees GAO-04-759, July 2004.
- US Government Accountability Office, "Intermodal Transportation: A Variety of Factors Influence Airport-Intercity Passenger Rail Connectivity," Report to Congressional Committees GAO-13-691, August 2013.
- US Naval Research Laboratory, "WindSat Payload description." [Online]. Available: <http://www.nrl.navy.mil/WindSat/Description.php>. [Accessed: 06-Feb-2016].
- A. van Lamsweerde, "Requirements Engineering: From System Goals to UML Models to Software Specifications." New York, NY, USA: Wiley, 2009.
- W. Vickrey, "Counterspeculation, Auctions, and Competitive Sealed Tenders," *The Journal of Finance*, vol. 16 no. 1, pp. 8–37, 1961.
- J. Von Neumann and O. Morgenstern, "Theory of Games and Economic Behavior," 3<sup>rd</sup> Edition, Princeton University Press, 1953.
- W. K. Viscusi, J. E. Harrington, and J. M. Vernon, "Economics of Regulation and Antitrust", 4<sup>th</sup> Ed., Ch. 20, MIT Press, Cambridge MA, 2005.
- M. Watkins, "Building Momentum in Negotiations: Time-Related Costs and Action-Forcing Events," *Negot. J.*, vol. 14, no. 3, pp. 241–256, 1998.
- J. R. Wertz and W. J. Larson (eds.), "Space Mission Analysis and Design," 3<sup>rd</sup> ed., Microcosm Press, El Segundo 1999.
- D. A. Whelan, E. A. Adler, S. B. Wilson III, and G. M. Roesler Jr., "DARPA Orbital Express program: effecting a revolution in space-based systems", Proceedings of SPIE Vol 4136 - Small Payloads in Space, 2000.

- T. D. Wilson and J. W. Schooler, "Thinking Too Much: Introspection Can Reduce the Quality of Preferences and Decisions," *J. Pers. Soc. Psychol.*, vol. 60, no. 2, pp. 1–23, 1991.
- D. Wolf, J. Hyland, T. W. Simpson, and X. (Luke) Zhang, "The Importance of Training for Interactive Trade Space Exploration: A Study of Novice and Expert Users," *J. Comput. Inf. Sci. Eng.*, vol. 11, no. 3, 2011.
- D. Wray and E. Nicholson, "On the Line: Will the Houston/Dallas Bullet Train Revolutionize Texas or Divide it Forever?" *Houston Press*, available at <http://www.houstonpress.com/news/on-the-line-will-the-houston-dallas-bullet-train-revolutionize-texas-or-divide-it-forever-7679365>, August 18 2015.
- R. K. Yin, "Case Study Research: Design and Methods," Fourth Edition, Sage Publishing, 2009.
- F. Zwicky, "Discovery, Invention, Research: Through the Morphological Approach," Macmillan, 1969.

## 15 Appendices

### Appendix A Description of Tradespace Redesign Experiment Questionnaire

The questionnaire began with short answer questions:

1. *Out of the 100 different available cars, how many potential choices do you think you would have preferred over your specified no-agreement alternative (i.e. ride the bus/train)?*
2. *How many of those choices do you think would also have been preferred by your partner over their alternative?*
3. *Using the axes to the right as a guide, indicate the regions of the tradespace where you would have accepted a car over your no-agreement alternative. (circle or shade in the chosen area or areas)*
4. *Describe the general process your group used to find a solution. (1-2 sentences)*

These four questions ask the subject to describe their impressions of the problem space and solution method. Questions 1 and 2 provide numerical estimations of the number of feasible solutions, for use in scaling (subjects with low reported answers to question 1 are somewhat more difficult to please, which could correlate with other indicators) and in judging communication (groups with similar answers for question 2 are likely to have had clearer discussion of needs). Question 3 was accompanied with a miniature cost-benefit tradespace, and the subject response can be used to categorize them by what region of the design space was most interesting to them. Question 4 allows for an open response, the information from which can be used with an exploratory coding to find unexpected commonalities in the responses of different subjects.

The bulk of the questionnaire asked subjects to respond using a 7-point Likert-type scale from Strongly Disagree to Strongly Agree:

5. *I felt that I understood my benefits and costs.*
6. *I was able to judge whether or not a car was valuable according to my needs.*

Questions 5 and 6 target the subject's perception of their own grasp of the problem and their ability to make value judgements.

7. *I felt that I understood my partner's benefits and costs.*
8. *It was difficult to find choices that were fair: approximately equally good for both my partner and me.*
9. *I found the design task to be stressful.*

Questions 7 and 8 target the evaluation of the group problem, particularly whether or not it was clear what each subject's partner wanted and the ease with which both participants' needs could be balanced. Question 9 expands on the idea of difficulty to see if the subject was stressed by the task.

*10. There were cars that were superior to my no-agreement alternative in benefit and cost.*

*11. There were cars that were superior to the no-agreement alternatives for both my partner and me in benefit and cost.*

*12. My partner's needs were opposite of my own.*

Questions 10 through 12 are a miniature "quiz" about the tradespace. Questions 10 and 11 ask the subject to make a judgement on whether there were options that were strictly superior to the BATNA for one or both subjects. Question 10 should identify some subjects that fail to even grasp what their BATNA is, which is hypothesized to be an effect of the "classic" tradespace views. The data set is set up such that, though there are many choices that are strictly superior for *one* person, there are only three choices that are superior for both. That car is not necessarily the "best" choice, as it is only barely better than each subject's BATNA and a tradeoff for more benefit or less cost could be preferable. However, question 11 *could* identify underconfidence in finding a mutually beneficial solution if respondents disagree, because it would imply that a dually-dominant design was not found. Similarly, respondents who agree with question 12 are expressing a common assumption of self-interested negotiators: that their partners are "against" them even when (as is the case here) their needs are mostly uncorrelated.

*13. We found and discussed the car that was best for both of our needs.*

*14. I am satisfied with our final decision.*

*15. Our final decision was better for my partner than it was for me.*

Questions 13 through 15 ask for the subjects to comment on the effectiveness of the process, with regards to the outcome. The first asks if the best alternative in the tradespace was found, and the second asks if, whether or not that choice was what was agreed on, the final result was satisfactory. Subjects who disagree with these sentiments are expressing either dissatisfaction with the exploration task or corresponding negotiation, respectively. Question 15 asks whether or not the subject believes they compromised more than their partner.

*16. The computer software helped me understand the problem.*

*17. Access to simpler tools (tables, line graphs, etc.) would have helped us find good solutions more easily.*

*18. Access to more customization (colors, sizes, etc.) would have helped us find good solutions more easily.*

Questions 16 through 18 ask for impressions on the computer interface, including whether it was helpful and asking if alternatives would be preferable. Low agreement on the first two questions will help to identify subjects who may not have been comfortable with the task or had difficulty applying the brief TSE training document. Because the experiment is necessarily constrained, it was expected that many subjects would agree with question 18, but a lower rating would imply that the visualizations are pushing on the boundary of “information overload” and that further complication should be avoided.

*19. At least once, I tried to convince my partner that a particular car was a good choice even though they were hesitant to accept.*

*20. At least once, I tried to misrepresent my needs to my partner in order to get a more favorable outcome.*

Questions 19 and 20 ask for self-reflection on the techniques used by the subject during the task. Question 19 describes one-sided positional bargaining, while question 20 describes strategic deception. Both of these are considered “acceptable” in a negotiation setting, suggesting a low risk of non-reporting, but are typically indicative of a mindset that is considered counterproductive to integrative (mutually beneficial) negotiation.

*21. At least once, I think my partner tried to misrepresent their needs to me in order to get a more favorable outcome.*

*22. My partner was cooperative.*

*23. My partner was willing to compromise.*

*24. I would be willing to work with my partner again.*

Questions 21 through 24 allow the subjects to rate their partner in a variety of dimensions, including suspicion of deception, cooperation, and willingness to compromise. According to the principles of “full, open, truthful exchange” (Raiffa, 2002), these scores should improve on average when more group-centric information is presented to the participants in the experimental condition, although this may be obscured due to high variance between subjects due to different personalities.

*25. I would be willing to use a similar method to solve a real-life group problem in the future.*

Question 25 asks for an overall satisfaction with the efficacy of MSTSE, separated from the car-choice problem and the subject’s current partner.

## Appendix B NEC Value Model

The following tables describe the utility function used for each stakeholder in the Northeast Corridor case study, as developed in collaboration with the MIT Rail / High-Speed Rail Research Group. Each stakeholder has their own table, and each attribute in the function is described in a row with the following categories:

- Unit – the unit of measurement
- Swing Weight – the k-weight assigned to the attribute
- Requirement – the worst acceptable performance in that attribute, corresponding to a utility of 0. Requirements for this case were set low specifically to avoid preemptively marking any bundles as invalid thus allowing for their inclusion in the tradespace, with few exceptions.
- Curvature – the marginal value of the attribute, set to either linear (constant), increasing, or decreasing. Increasing and decreasing attributes also include one inflection point in the utility curve, written as [X;Y] meaning  $U(X)=Y$ .
- Objective – the performance which results in maximum value for the attribute, corresponding to a utility of 1. In this case, the objectives were set to be “better than the best” value in the tradespace so that all differences in performance result in differences in value
- Source Model – the source of the data. The different models are described within the case study. Most attributes are taken directly from the “BCA spreadsheet” of Sussman et al., 2015, or are based on “BCA postprocess” of that spreadsheet with basic arithmetic functions.

The values for swing weight, requirement, curvature, and objective are all estimates based on subject matter expertise and are not suggested to be perfect substitutes for well-structured utility elicitation interviews with the actual stakeholders. For the purposes of this case study, however, this assumption is particularly palatable given the small number of highly-differentiated feasible alternatives in the NEC tradespace. The large differences between each possible bundle in the NEC tradespace make the resulting organization / relative ordering of the tradespace insensitive to small changes in these parameters.

The utility function used to aggregate the attributes is the Keeney-Raiffa multi-attribute utility function, which is as follows:

$$U(\hat{X}) = \frac{[\prod_{i=1}^n (K \cdot k_i \cdot U_i(X_i) + 1)] - 1}{K}, \text{ where } K = -1 + \prod_{i=1}^n (K \cdot k_i + 1)$$

Here  $K$  is the normalization constant,  $U(\hat{X})$  is the aggregate MAU value across the multiple single attributes  $X_i$  and their respective single attribute utilities  $U_i(X_i)$ ;  $k_i$  is the elicited swing weighting factor for attribute  $X_i$ ;  $n$  is the number of attributes.

The US Department of Transportation is modeled as having an interest in quality of service metrics – including for alternative modes of travel such as the roads – and environmental emissions. Their primary concern is safety, marked by a slightly higher swing weight, but they are balanced overall.

**Table 15-1: US Dept. of Transportation Utility function**

<b>Attribute</b>	<b>Unit</b>	<b>Swing Weight</b>	<b>Requirement (Utility=0)</b>	<b>Curvature</b>	<b>Objective (Utility=1)</b>	<b>Source Model</b>
Casualty Rate	Passenger casualties per 100K passenger-miles	0.3	8	Increasing [3;0.25]	0	Safety Model
Time Saved	Million minutes	0.2	0	Increasing [4000;0.25]	8,000	BCA Spreadsheet
On-time Performance	-	0.2	0.72	Decreasing [0.9;0.8]	1	On-Time Model
Road Congestion	Proportion jammed	0.2	0.2	Decreasing [0.15;0.75]	0.1	Road Congestion Model
Net Emissions \$	\$	0.1	0	Linear	\$160M	BCA Spreadsheet

Amtrak and the other rail agencies that use the NEC are modeled as having a primary interest in the financial returns of the rail system, as a business. The remainder of their utility function is spread over the quality of service metrics, with more emphasis on on-time performance and casualty rate as the more “visible” features for customers.

**Table 15-2: Amtrak and Rail Agencies Utility function**

<b>Attribute</b>	<b>Unit</b>	<b>Swing Weight</b>	<b>Requirement (Utility=0)</b>	<b>Curvature</b>	<b>Objective (Utility=1)</b>	<b>Source Model</b>
Financial Returns	\$	0.5	0	Linear	\$80B	BCA Postprocess
On-time Performance	-	0.2	0.72	Decreasing [0.9;0.8]	1	On-Time Model
Casualty Rate	Passenger casualties per 100K passenger-miles	0.2	8	Increasing [3;0.25]	0	Safety Model
Time Saved	Million minutes	0.1	0	Increasing [4000;0.25]	8,000	BCA Spreadsheet

Congress defines benefit with the highest-level BCA attributes of economic and financial return. Their primary interest is on the economic benefits of the system to the surrounding population, but they also want to ensure that the resulting rail system is financially solvent in order to avoid a situation like the present in which they must continually provide remedial funding to support the system.

**Table 15-3: Congress Utility function**

<b>Attribute</b>	<b>Unit</b>	<b>Swing Weight</b>	<b>Requirement (Utility=0)</b>	<b>Curvature</b>	<b>Objective (Utility=1)</b>	<b>Source Model</b>
Economic Returns	\$	0.6	0	Linear	\$120B	BCA Postprocess
Financial Returns	\$	0.4	0	Linear	\$80B	BCA Postprocess

Like Congress, the northern states are primarily interested in the economic effects of the rail system. However, the states also care about the local traffic and quality of service on the rails – benefit is tied to the experience of the local constituency.

**Table 15-4: Northern States Utility function**

<b>Attribute</b>	<b>Unit</b>	<b>Swing Weight</b>	<b>Requirement (Utility=0)</b>	<b>Curvature</b>	<b>Objective (Utility=1)</b>	<b>Source Model</b>
Economic Returns	\$	0.35	0	Linear	\$120B	BCA Postprocess
Rail Passengers (north)	Passengers	0.3	5.9M	Linear	60M	BCA Spreadsheet
Time Saved (north)	Million minutes	0.15	0	Increasing [4000;0.25]	8,000	BCA Spreadsheet
On-time Performance (north)	-	0.1	0.72	Decreasing [0.9;0.8]	1	On-Time Model
Casualty Rate (north)	Passenger casualties per 100K passenger-miles	0.1	8	Increasing [3;0.25]	0	Safety Model

The southern states use the same utility function as the northern states, but for their own local region. The two state stakeholders are differentiated by the different levels of service between the regions in some bundles.

**Table 15-5: Southern States Utility function**

Attribute	Unit	Swing Weight	Requirement (Utility=0)	Curvature	Objective (Utility=1)	Source Model
Economic Returns	\$	0.35	0	Linear	\$120B	BCA Postprocess
Rail Passengers (south)	Passengers	0.3	5.9M	Linear	60M	BCA Spreadsheet
Time Saved (south)	Million minutes	0.15	0	Increasing [4000;0.25]	8,000	BCA Spreadsheet
On-time Performance (south)	-	0.1	0.72	Decreasing [0.9;0.8]	1	On-Time Model
Casualty Rate (south)	Passenger casualties per 100K passenger-miles	0.1	8	Increasing [3;0.25]	0	Safety Model

The EPA and landowners stakeholder is modeled as the primary environmental interest in the negotiation, with a large emphasis on the net emissions impact of the NEC and a minor interest in additional money earmarked for environmental impact mitigation spending. One important feature of this utility function is the raised requirement on net emissions, which unilaterally eliminates all non-HSR bundles from this stakeholder’s tradespace. Additionally, note that net emissions are calculated annually, thus explaining the low magnitude relative to the one-time environmental mitigation payment.

**Table 15-6: EPA and Landowners Utility function**

Attribute	Unit	Swing Weight	Requirement (Utility=0)	Curvature	Objective (Utility=1)	Source Model
Net Emissions \$	\$	0.7	\$1M	Linear	\$160M	BCA Spreadsheet
Env. Mitigation \$	\$	0.3	0	Linear	\$3B	BCA Spreadsheet

The private consortiums are the non-public investors in the NEC, and are interested in maximizing financial returns (scaled by their own contributions to the funding pool) and minimizing the payback period over which the system will recoup its capital expenditures. The \$1B requirement on private returns was included to prevent this stakeholder from ever supporting bundles in which they do not contribute *and* make money.

**Table 15-7: Private Consortiums Utility function**

Attribute	Unit	Swing Weight	Requirement (Utility=0)	Curvature	Objective (Utility=1)	Source Model
Financial Returns (private)	\$	0.5	\$1B	Linear	\$17B	BCA Postprocess
Payback Period	years	0.5	55	Linear	0	BCA Postprocess

Suppliers and Labor derive benefit from the construction attributes of the NEC plan, primarily from the total cost but with an additional component derived from construction duration (as longer construction leads to more job security).

**Table 15-8: Suppliers and Labor Utility function**

Attribute	Unit	Swing Weight	Requirement (Utility=0)	Curvature	Objective (Utility=1)	Source Model
Construction Cost	\$	0.8	0	Linear	\$90B	BCA Spreadsheet
Construction Duration	years	0.2	0	Linear	30	BCA Spreadsheet

The Highway System stakeholder has a preference on the same road congestion attribute as the US-DOT, but with a stricter requirement. This higher requirement represents an unwillingness to drop below current service levels and in some scenarios will eliminate non-HSR bundles which do not drive enough diversion to mitigate population-driven increases in road traffic.

**Table 15-9: Highway System Utility function**

Attribute	Unit	Swing Weight	Requirement (Utility=0)	Curvature	Objective (Utility=1)	Source Model
Road Congestion	Proportion jammed	1	0.16	Decreasing [0.12;0.75]	0.1	Road Congestion Model

Finally, the rail travelers derive benefit only from the quality of service of the rail system. In contrast to other stakeholders who share quality of service attributes, the travelers have priority on time saved and on-time performance, as these have the most immediate impact on the experienced benefit of high speed rail for all riders (as opposed to casualty rate, which is experienced by only a few).

**Table 15-10: Travelers Utility function**

<b>Attribute</b>	<b>Unit</b>	<b>Swing Weight</b>	<b>Requirement (Utility=0)</b>	<b>Curvature</b>	<b>Objective (Utility=1)</b>	<b>Source Model</b>
Time Saved	Million minutes	0.4	0	Increasing [4000;0.25]	8,000	BCA Spreadsheet
On-time Performance	-	0.4	0.72	Decreasing [0.9;0.8]	1	On-Time Model
Casualty Rate	Passenger casualties per 100K passenger-miles	0.2	8	Increasing [3;0.25]	0	Safety Model