



Conference on Systems Engineering Research (CSER 2014)

Eds.: Azad M. Madni, University of Southern California; Barry Boehm, University of Southern California;
Michael Sievers, Jet Propulsion Laboratory; Marilee Wheaton, The Aerospace Corporation
Redondo Beach, CA, March 21-22, 2014

Controlling for Framing Effects in Multi-Stakeholder Tradespace Exploration

Matthew E Fitzgerald*, Adam M Ross

Systems Engineering Advancement Research Initiative (SEArI), Massachusetts Institute of Technology, Cambridge MA, 02139

Abstract

Framing effects have been shown to have dramatic impact on human decision making in many domains, in certain circumstances even driving self-detrimental behavior. Multi-stakeholder tradespace exploration (MSTSE), an emerging technique for advanced multiparty decision making for engineering systems, has displayed many benefits with regards to insight-generation and identification of mutually beneficial solutions. However, for complex problems with no solutions that are individually optimal for each stakeholder, stakeholders may still resist “compromising” from their individual preferred solutions. This occasionally drives a failure to reach agreement, despite a design space with a considerable number of feasible designs with value for all parties. This paper hypothesizes that this result may be caused in part by an unintentional framing of the initial stages of MSTSE as an individual problem, establishing an unrealistically high reference point. Theoretically, this locks stakeholders into a mindset that forces them to “compromise” down, rather than more appropriately building up mutual value from the “no agreement” alternative. This paper addresses the current literature of multi-stakeholder system design, the ramifications of framing on MSTSE, considerations for establishing a more appropriate reference point, and example techniques and visual representations for doing so. A preliminary set of experiments is described to confirm the hypothesized framing effect and to validate visual representations for mitigating its impact.

© 2014 The Authors. Published by Elsevier B.V.

Selection and peer-review under responsibility of the University of Southern California.

Keywords: tradespace exploration, multi-stakeholder, multiple stakeholders, framing effects, negotiation

* Corresponding author. mattfitz@mit.edu

Nomenclature

TSE	tradespace exploration
MSTSE	multi-stakeholder tradespace exploration
Stakeholder	a person/group with a defined interest in a given system
Decision maker	a stakeholder with non-negligible control over the system design process and outcome

1. Introduction

As a paradigm for solving complex design problems, the majority of research in tradespace exploration (TSE) has focused on the analysis of the space of alternatives with the goal of uncovering design choices that are optimal or near-optimal^{1,2}. These designs feature desirable combinations of attributes for a given system stakeholder, including technical attributes, cost, and, more recently, -ilities³. Less tradespace research has been devoted to the multi-stakeholder problem, in which there are multiple parties with different desired attributes of performance, 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 groups⁴. Often, practitioners have taken to ignoring these results, optimizing some function of each stakeholder's utility, or simply selecting a baseline and allowing iterative design improvements only if they benefit every stakeholder individually^{5,6}. These techniques, though fast and sometimes effective, run counter to the goal of *exploration* by looking at only a subset of possible designs, limiting the potential for improving stakeholders' understanding of the problem, and increasing the possibility that a superior mutually desirable design is overlooked or not evaluated. 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 parallel⁷. This parallel exploration has the goal of uncovering as many 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 order to further improve the likelihood of discovering superior compromise solutions, MSTSE should move beyond parallel individual exploration and incorporate lessons from the behavioral economics and negotiation literature.

Framing effects have been shown to have dramatic impact on human decision making, with even subtle differences in framing resulting in surprising variation of outcome⁸. The current practice of MSTSE as a combination of individual exercises frames the problem for each participant *first* as an individual and *second* as a group. According to Prospect Theory, humans make decisions with regard to a reference point (often influenced by framing) and are strongly negative about decisions that result in outcomes below that reference ("losses"), more so than they are positive about results above the reference ("gains")⁸. By encouraging stakeholders to spend time on the problem individually, their reference point is reinforced at a high level, essentially centering the mind of each stakeholder on the question: "how much value could I get if everyone else had to agree with me?" For most complex problems, realistic compromises will be significantly below this reference point and thus appear as unappealing "losses," possibly blocking the acceptance of a mutually beneficial solution that could otherwise be agreed upon. Framing the multi-stakeholder problem as a *group* exercise from the start has the potential to establish a more appropriate reference point at the beginning of the exploration, and make a larger set of potential designs within the tradespace visible as "gains" that all parties might agree to develop.

This paper conducts a review of common practices for system design with multiple stakeholders and discusses their relevance and appropriateness for use with tradespace exploration. Then the ramifications of standard, individual-focused tradespace exploration techniques on the framing of the group decision problem are described. Discussion focuses on the creation of a more productive framing for MSTSE. Additionally, example techniques and visual representations designed to reinforce a group-focused reference point over the individual are presented.

2. Decisions with multiple stakeholders

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. Single stakeholder, multiple goals

Before diving into multiple-stakeholder analysis, we should 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. Aggregation allows the ranking of alternatives, 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 paper is not intended to be an exhaustive accounting of multi-criteria decision making methods and unfortunately some popular techniques (e.g. AHP, QFD, Pugh charts) stretch or defy the following categorization, but for narrative simplicity, we will cast this as a debate between cost-benefit analysis and utility analysis (representing multiple popular non-monetizing valuation techniques).

Utility theory, originated by von Neumann and Morgenstern and developed by many others^{4,9}, 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, though a broad technique with different varieties, is 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 called "utility" if the accounting procedure follows utility theory axioms)¹⁰. 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. 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 take¹¹. 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 people¹².

2.2. Subsystem teams and connected stakeholders

Large engineering design efforts are often organized into subsystem design groups, with each responsible for advocating for their own interests. These interactions have received considerable attention in the literature, mostly from the utility camp, including the modeling of the problem as a negotiation. A key assumption of much of this research is that the different stakeholders are closely tied, characterized by extensive information exchange and iterated communication. The field of collaborative engineering¹³ would refer to this degree of connection as either *collaborative* or *cooperative* (as opposed to the less-connected *coordinated*), and indeed collaborative negotiation has received significant attention in recent years¹⁴.

Many utility-oriented researchers have approached this problem with the goal of some sort of aggregate utility maximization between groups^{6,15}. This research thread has extended to the point of complete automation, using algorithms to generate “optimal” or Pareto-efficient solutions^{5,16,17}. 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 which can be compared across individuals. Aggregations of multiple people’s utilities sacrifice the normative benefits of utility theory, so the results of such optimizations should be taken with caution.

2.3. Multiple empowered decision makers

The methods of the previous section were 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 negotiations if they so desire. These types of negotiations are more complicated than those between connected stakeholders, as attention to interpersonal conflict and resolution is required to ensure completion. Automated, algorithmic negotiations are typically not appropriate for these problems, as decision makers are less likely to abdicate responsibility to a “black box.” This problem type is the main concern of this research, so the remainder of the paper will use “stakeholder” as shorthand for “decision maker.”

Utility theory has been applied to resolve conflicts between decision makers, modeling preferences and offering up potential compromises one at a time using a combination of joint maximization and equality to address interpersonal friction¹⁸. This application of utility theory is prescriptive, demonstrating effectiveness at *persuading* stakeholders to accept compromises. The explicit inclusions of insights from conflict resolution and negotiation theory are most prevalent in research at the decision maker level, including methods such as Joint Fact Finding, layering them into process-oriented systems engineering frameworks¹⁹. Note that there have *not* been serious efforts at aggregating and maximizing utility for multiple stakeholders in this problem type, because the higher degree of social choice characteristics in this domain places it firmly within the grip of Arrow’s Theorem³⁰ unlike the more structured problem of the prior section which can plausibly escape that result on axiomatic grounds³¹.

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)^{20,21} 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. 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.

2.4. Tradespace exploration and multiple stakeholders

The field of complex sociotechnical systems benefits dramatically from performing tradespace exploration with utility theory over cost-benefit analysis. 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. utilities). Though the analysis of some complex systems will be comfortable adopting 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 representative²². Additionally, negotiations are often accompanied with changes in stakeholder preferences as new information is revealed, which would necessitate frequent reassessment of willingness-to-pay²³. 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 gives the potential to account for modified preferences on the fly during a negotiation. For these reasons, this paper will focus on using utility theory to represent stakeholder preferences.

The uses of utility theory in negotiations, however, have been severely limited by the inability to mathematically aggregate the utilities of different stakeholders. Previous attempts have either accepted the risks of employing such an aggregation, including the premature and unnecessary compromising of compatible preferences, or resorted to simply “horse trading” the utility of the stakeholders as a measure of benefit. Neither of these solutions is conducive to exploring the problem and developing understanding, so previous attempts at conducting MSTSE have called for stakeholders to conduct individual TSE and then combine insights to search for mutual agreement⁷. The following section considers the implications of this procedure with regards to the possibility that framing effects can worsen the outcome of the negotiations.

3. MSTSE and Framing Effects

In spite of the above challenges related to the adoption of extant methods for aggregation or reconciliation of multi-stakeholder value statements, TSE remains extremely relevant to the negotiation problem. The main features of utility-driven TSE are aligned with the tenets of good negotiation, including information sharing, use of objective data, and a focus on interests (i.e. preferences) over positions (i.e. specific design points)²⁴. It seems that combining the tools and ideas of TSE with those of negotiation analysis will allow MSTSE to promote stakeholder understanding of the problem space and improve subsequent decisions. However, the individual, self-interested tools and tasks of classic tradespace exploration may inappropriately frame the group design problem of MSTSE, potentially driving a wedge between stakeholders that can derail attempts to find a mutually beneficial solution.

3.1. MSTSE using individual TSE techniques

As previously mentioned, early attempts at conducting structured MSTSE involved performing individual TSE tasks in parallel, then bringing the results together to discuss and search for mutually agreeable solutions via analysis of Pareto efficiency and updating utility functions in light of new information⁷. The individual TSE was performed with the expressed goal of finding specific design points to use as focusing points for the later discussion and was conducted using a series of standard tradespace views and analyses, including individual utility-cost tradespace plots, design variable comparisons, variable sensitivity plots, and animated views of context uncertainty³. These techniques were originally designed to help stakeholders find alternatives with the most value *for them*. In a single stakeholder problem, this makes sense: exploring the tradespace to find the design that best meets the stakeholder’s

perception of value is the ultimate goal. However, the ultimate goal of MSTSE requires mutual agreement. Searching for individually optimal or near-optimal solutions does not necessarily support this goal.

3.2. Negative framing effects in MSTSE

In a decision-making context, *framing* refers to the way in which a problem is posed. Framing *effects* are a phenomenon associated with changes in behavior between situations that are identical except for their framing, and have been documented extensively in many problem-solving domains. Often, the most powerful framing effects have been demonstrated when switching between a positive and negative frame, explained in part by the asymmetric nature of perceived value around a *reference point*, as described by Prospect Theory⁸. 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 relevance²⁵. Changing a reference point, once established, usually requires a deliberate effort.

One way to think about framing in MSTSE is to look at the influence of particular data visualization and exploration techniques with regards to the focus of the stakeholder. For TSE, the stakeholder’s utility function is usually elicited before the design points in the tradespace are evaluated but analysis of these stated preferences is held until after the stakeholder is shown the cost-utility tradespace plot. 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 design points), it is natural for the stakeholder to fixate on reaching maximum utility, or the complete satisfaction of all of their stated needs. However, it is extremely common for no design points to achieve maximum utility and it is important to prevent the stakeholder from establishing that as their reference point as it makes the other feasible designs in the tradespace look significantly worse by comparison. Instead, the cost-utility tradespace plot is typically the first view shown, illustrating feasible solutions for varying combinations of cost and utility. People with familiarity with TSE are comfortable with this view and the available reference point is the Pareto front: the set of designs with efficient cost-utility tradeoffs. This is a good reference for individual TSE, as the stakeholder is empowered to select any of those Pareto efficient designs if he desires.

In MSTSE, however, no single stakeholder has the ability to dictate decisions based on only their own value. Establishing the individual Pareto front as the reference point is overly *optimistic*, since the actual solution cannot be better than that reference (on cost-utility grounds) and will almost certainly be worse if other stakeholders have divergent preferences. If the individual Pareto front becomes the reference point, the 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 negotiations³². Thus, 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 in Figure 1.

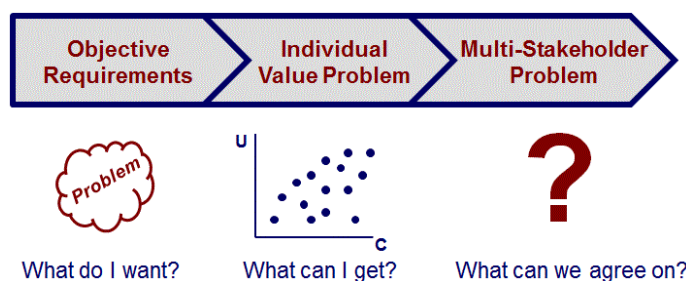


Fig. 1. Increasing complexity of decision from problem definition to physical constraints to relational concerns

4. Creating a better frame

4.1. Productive framing

Avoiding counterproductive framing effects is desirable but doing so requires determining what constitutes a *productive* frame. The establishment of the appropriate reference point is necessary to make the framing of the problem one that allows stakeholders to make the best possible decisions. In the negotiation literature, stakeholders are encouraged to always refer back to their best alternative to a negotiated agreement (BATNA)²⁶, as it represents a boundary of value that defines when leaving the negotiation is the best decision. The BATNA is therefore 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 so²⁷, 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. Ideally then, the exploration would be designed to frame the problem in such a way that the BATNA is reinforced as the reference point.

In addition to emphasizing a proper reference point, MSTSE can also benefit from increasing the availability of information relevant to the group problem that is absent in traditional TSE. In particular, it is important to show each stakeholder some indicator of their counterpart’s value. 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 beneficial²⁶. By making each stakeholder aware of the others’ interests and estimates of their value up front, MSTSE can reduce the time wasted and negative feeling associated with “offering” a potential solution that is clearly unacceptable to one or more parties.

4.2. Modifying tradespace visualizations

As discussed previously, the default view of traditional TSE is the individual cost-utility tradespace, as depicted in Figure 2a. Where is the BATNA in this visualization? With no additional information, it is 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 2b. With a more immediate reference point in the origin than the Pareto front (because it is directly indicated on the plot), designs 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 from leaving the negotiation with no solution.

Part of the challenge in redesigning the tradespace plot for MSTSE is breaking the bad habits of individual “value claiming” from stakeholders who are confident engaging with tradespaces. Presenting the same information in a new way can prevent people from leaping to comfortable, predefined conclusions and instead consider the available evidence more carefully. The heuristic-systematic model of information processing³³ would describe this as a switch from ‘heuristic’ processing (such as a prior commitment, based on experience, to choosing a design on the Pareto front) to ‘systematic’ processing, driven by an increased concern over reliability of old heuristics in the new framing. Even a simple change such as the rotation of the tradespace axes could force stakeholders to more

critically evaluate 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.

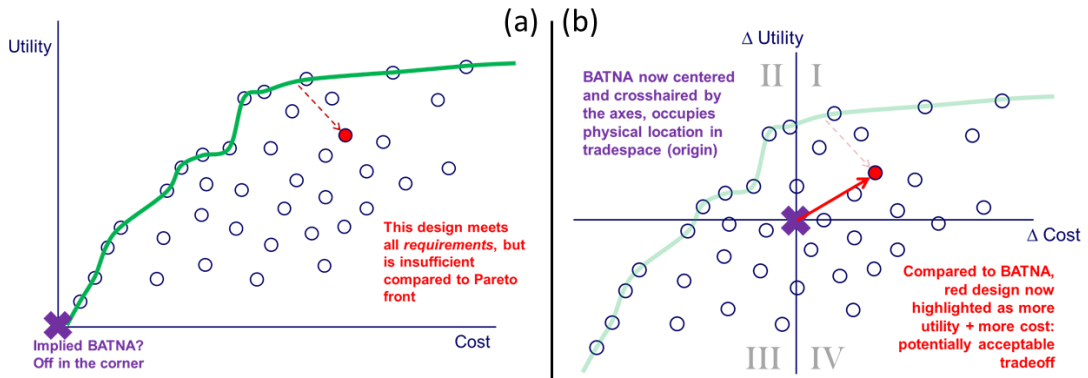


Fig. 2. A cost-utility tradespace before (a) and after (b) being centered on the BATNA

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 (note that this allows only one extra stakeholder to be viewed at once, though MSTSE is in general not restricted to only two participants). This makes 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). 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)²⁸. 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 example plot incorporating color and transparency is shown in Figure 3. This image is for conceptual demonstration only, as the deployment of color and transparency must be used carefully in order to properly communicate variables of different relationships (e.g. categorical, interval, etc.) and further investigation is required to find the best visualization.




Fig. 3. Example re-centered, rotated, and color/transparency MSTSE visualization

4.3. Experimental validation and further research

The conclusions of this paper, though attempting to be as thorough as possible, are still in need of experimental validation. We believe that the analysis of the concepts of framing effects, negotiation, and tradespace exploration contained within are sufficient for the creation of content validity, but experiments with human practitioners are necessary to confirm the impact that framing effects can have on MSTSE. A key challenge for these experiments will be the difficulty in replicating “real world” applications of MSTSE with regards to both the TSE expertise of the participants and the level of aggressive value-claiming behavior when the value is real and jobs or reputations are on the line. We anticipate that engineering students will be successful test subjects despite these concerns, as a short introduction to the basics of TSE combined with an engineering background should be sufficient to approximate the approaches of technical decision makers in a negotiation context²⁹. The behavior of non-technical decision makers will require additional research.

Two experiments are being organized in order to validate the theory presented in this paper. The first will focus on quantifying the effects of framing on stakeholders in a simplified MSTSE scenario. In addition to comparing the outcomes between a control group (guided through current MSTSE practice) and an experimental group (using the above visualization modifications) to determine effects on the likelihood of reaching a mutual solution and the quality of that solution, questionnaires will be administered to quantify the stakeholder response to the MSTSE process in terms of situational awareness and satisfaction. If significant differences are observed between the two groups, we can tentatively conclude that framing effects do impact the practice of MSTSE. The second experiment will attempt to get real engineering decision makers to participate in both controlled and experimental MSTSE on a dataset replicating an actual space architecture case study. Analysis of this experiment will focus on qualitative interviews with and observations of the participants rather than quantification, exploring their insights and feedback about both conditions. This should result in a more complete understanding of the ways in which practitioners choose to engage with the data presented in MSTSE, and allow for further refinement of updated visualizations to best support decision making and the identification of productive, mutually beneficial solutions.

Acknowledgements

This material is based upon work supported, in whole or in part, by the U.S. Department of Defense through the Systems Engineering Research Center (SERC) under Contract H98230-08-D-0171. SERC is a federally funded University Affiliated Research Center managed by Stevens Institute of Technology. The authors also gratefully acknowledge support for this research from the MIT Systems Engineering Advancement Research Initiative (SEARI).

References

1. Daskilewicz, M. J. and German B. J., “Aspects of Effective Visualization of Multidimensional Design Spaces,” in *AIAA Aviation Technology, Integration, and Operations Conference*, 2009
2. Ross, A. M., Hastings, D. E., Warmkessel, J. M., and Diller, N. P., “Multi-Attribute Tradespace Exploration as Front End for Effective Space System Design,” *J. Spacecr. Rockets*, vol. 41, no. 1, 2004..
3. Ross, A.M., McManus, H.L., Rhodes, D.H., and Hastings, D.E., "Revisiting the Tradespace Exploration Paradigm: Structuring the Exploration Process," *AIAA Space 2010*, Anaheim, CA, September 2010.
4. Keeney, R.L. and Raiffa, H.: *Decisions with multiple objectives—preferences and value tradeoffs*, Cambridge University Press, Cambridge & New York, 1993.
5. Gatti, N., & Amigoni, F. An approximate Pareto optimal cooperative negotiation model for multiple continuous dependent issues, *IEEE/WIC/ACM International Conference on Intelligent Agent Technology* 565–571, 2005.
6. Bahler, D., Dupont, C., & Bowen, J. A Mixed Quantitative/Qualitative Method for Evaluating Compromise Solutions to Conflicts in Collaborative Design. *Artificial Intelligence for Engineering, Design, Analysis and Manufacturing*, 9(04), pg 325, 1995.
7. Ross, A.M., McManus, H.L., Rhodes, D.H., and Hastings, D.E., "A Role for Interactive Tradespace Exploration in Multi-Stakeholder Negotiations," *AIAA Space 2010*, Anaheim, CA, September 2010.
8. Kahneman, D. and Tversky, A. “Choices, Values, and Frames.” Cambridge University Press, 2000.
9. Von Neumann, J. and Morgenstern, O., *Theory of Games and Economic Behavior* 3rd Edition, Princeton University Press, 1953.

10. Hazelrigg, G. A., “A framework for decision-based engineering design,” *J. Mech. Des.*, vol. 120, no. 4, pp. 653–658, 1998.
11. Viscusi, W.K., Harrington, J.E., Vernon, J.M., “Economics of Regulation and Antitrust”, 4th Ed., Ch. 20, MIT Press, Cambridge MA, 2005.
12. Ackerman, F. and Heinzerling L., “Pricing the Priceless: Cost-Benefit Analysis of Environmental Protection,” *Univ. PA. Law Rev.*, vol. 150, no. 5, p. 1553, May 2002.
13. Lu, S., Elmaraghy, W., Schuh, G., and Wilhelm, R., “A Scientific Foundation of Collaborative Engineering,” *CIRP Ann. - Manuf. Technol.*, vol. 56, no. 2, pp. 605–634, Jan. 2007.
14. Shahan, D. and Seepersad, C., “Bayesian Networks for Set-Based Collaborative Design,” in *ASME 2009 International Design Engineering Technical Conferences*, 2009.
15. Scott, M.J. and Antonsson, E.K., “Formalisms for Negotiation in Engineering Design,” in *8th International Conference on Design Theory and Methodology*, 1996, pp. 165–182.
16. Chen, J.H., Chao, K.M., Godwin, N., and Soo, V.W., “A multiple-stage cooperative negotiation,” *IEEE International Conference on eTechnology eCommerce and eService 2004*, pp. 131–138.
17. Romanhuki, E., Fuckner, M., Enembreck, F., Avila, B., and Scalabrin, E., “Improving bilateral negotiation with evolutionary learning,” 12th International Conference on Computer Supported Cooperative Work in Design, 2008, pp. 343–348.
18. Sycara, K., “Utility Theory in Conflict Resolution,” *Ann. Oper. Res.*, vol. 12, pp. 65–83, 1988.
19. Mostashari, A., “Stakeholder-Assisted Modeling and Policy Design Process for Engineering Systems,” MIT Engineering Systems Division, PhD Thesis, 2005.
20. B. Boehm and A. Jain, “The Value-Based Theory of Systems Engineering: Identifying and Explaining Dependencies,” in *INCOSE International Symposium 2007*, 2007.
21. Horowitz, E., Lee, J.H., and Lee, J.S., “WinWin: a system for negotiating requirements,” *Softw. Eng. 1999. ...*, pp. 646–649, 1999.
22. Ross, A. M., O’Neill, M.G., Hastings, D.E., and Rhodes, D.H., “Aligning Perspectives and Methods for Value-Driven Design,” in *AIAA Space 2010*, 2010.
23. Curhan, J.R., Neale, M.A., and Ross, L., “Dynamic valuation: Preference changes in the context of face-to-face negotiation,” *J. Exp. Soc. Psychol.*, vol. 40, no. 2, pp. 142–151, 2004.
24. Fitzgerald, M.E., and Ross, A.M., “Guiding Cooperative Stakeholders to Compromise Solutions Using an Interactive Tradespace Exploration Process,” 11th Conference on Systems Engineering Research, Atlanta, GA, March 2013.
25. Tversky, A. and Kahneman, D., “Judgment under Uncertainty: Heuristics and Biases,” *Science*, vol. 185, no. 4157, pp. 1124–31, Sep. 1974.
26. Fisher, R., Ury, W., and Patton, B., “Getting to Yes: Negotiating an Agreement without Giving In”, Penguin Books, New York NY 1991.
27. Richards, M.G., Viscito, L., Ross, A.M., and Hastings, D.E., “Distinguishing Attributes for the Operationally Responsive Space Paradigm,” in *Responsive Space 6 Conference*, Los Angeles CA, 2008.
28. Fitzgerald, M.E. and Ross, A.M., “Mitigating Contextual Uncertainties with Changeability Analysis in the Multi-Epoch Domain,” 6th Annual IEEE Systems Conference, Vancouver Canada, 2012.
29. Wolf, D., Hyland, J., Simpson, T.W., and Zhang, X., “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.
30. Arrow, K.J., “Social Choice and Individual Values” 2nd Edition, Yale University Press, 1963.
31. Scott, M.J. and Antonsson, E.K., “Arrow’s Theorem and Engineering Design Decision Making,” *Res. Eng. Des.*, vol. 11, no. 4, pp. 218–228, 2000.
32. Gelfand, M.J., Fulmer, A.C., and Severance, L., “The Psychology of Negotiation and Mediation,” in *APA Handbook of Industrial and Organized Psychology*, Vol 3, 2004, pp. 495–554.
33. Chaiken, S., Liberman, A., and Eagly, A.H., “Heuristic and systematic processing within and beyond the persuasion context,” in J.S. Uleman and J.A. Bargh (Eds.), *Unintended Thought*, pp. 212-252, New York, Guilford Press, 1989.