Available online at www.sciencedirect.com

SciVerse ScienceDirect

Procedia Computer Science 00 (2014) 000–000

Procedia
Computer Science

www.elsevier.com/locate/procedia

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

Exploring Stakeholder Value Models via Interactive Visualization

Nicola Ricci*, Michael A. Schaffner,
Adam M. Ross, Donna H. Rhodes, Matthew E. Fitzgerald

Systems Engineering Advancement Research Initiative (SEAri), Massachusetts Institute of Technology, Building E38-574, Cambridge, MA 02139

Abstract

In the design of complex systems, models can be used to predict a future system's performance as well as approximate stakeholder preferences on performance. This paper examines the problem of model truthfulness and the challenge of trusting models, with a focus on value models and how they are used to predict stakeholder preferences. A framework is proposed for the analysis of these issues (truthfulness and trust), which is used to discuss the relationship between models and decision outcomes. Interactive visualization is proposed as an efficient and effective method for increasing model truthfulness and model trust, and hence making better decisions. An interactive visualization tool is also presented, and an application of the tool to a complex decision case is discussed.

Keywords: Values; Model Representation; Performance Model; Value Model; Stakeholder Preferences; Interactive Visualization.

* Corresponding author. Tel.: +1-617-901-7791; fax: +1-617-253-3641.
E-mail address: nricci@mit.edu.

"Two things fill the mind with ever-increasing wonder… the starry heavens above me and the moral law within me."
– Immanuel Kant, Critique of Practical Reason (1788)

1. Introduction

Models are essential tools in system design and are used by analysts and engineers throughout the design process. A key function of models is to generate data when empirical sources are not available; this data is then used by stakeholders to discern among potential alternatives. This paper explores the challenge of building trust and the problem of building truthfulness in the models used to represent the performance of engineering systems. It also examines the challenge of building trust and the problem of building truthfulness in models that represent the value that stakeholders attribute to such performance. The design-performance-value loop (hereafter referred to as the "design loop") is discussed, and a framework is delineated for the comparison of important model types involved at various stages in the design loop (performance v. value, mental v. constructed). The problem of building trust and truthfulness in models is then defined within this proposed framework. Interactive visualization is proposed as a viable tool to help improve the model representations, thereby increasing model usefulness and stakeholder trust in the models.

1.1. Motivation

In the conceptual design phase of modern engineering systems, analysts are confronted with exploring and representing several layers of complex information about the design alternatives they are considering. Rhodes and Ross (2010) outline five types of complexity in engineering systems[1]: structural (related to systems' forms), behavioral (systems' operations), contextual (environment in which the systems operate), temporal (systems changing over time), and perceptual (stakeholder preferences on systems' performance outcomes). Given these complexities, not only is it difficult to model the performance of a system that does not yet exist, but also it is difficult to represent stakeholders' views on what the system should do (i.e., their needs). This brings about the problem of *trust* (do I "believe"?) in the models that are used to evaluate both the performance and the preferences on that performance (i.e., value models). Another important problem is related to how *truthful* (is it "correct"?) the models actually are – independent of trust (since it is possible to trust a very poor model).

In this context, effective *interactive visualization* techniques can become indispensible in building trust in models, as well as bringing them closer to a more truthful representation of performance and stakeholder values – as will be explored in §3.

1.2. The Design Loop

Conceptual design of complex systems involves a design-performance-value loop (see Fig. 1). First, a design space of alternatives is generated based on a preliminary understanding of stakeholder needs. Each design alternative is mapped onto a performance space, spanned by attributes of interest quantified by a constructed performance

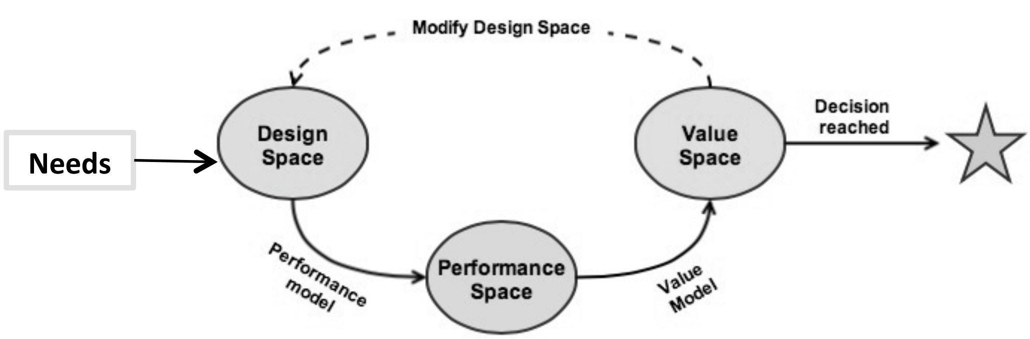

Fig. 1. Illustration of the design-performance-value loop.

model. The performance space is then mapped onto a value space via a value model (e.g., functional requirements, utility functions), through which the stakeholder evaluates the attractiveness of the alternatives. At this point, with an understanding of how each design alternative scores according to stakeholders' values, it is possible to either make a decision on what design (or set of designs) to focus on, or to go back and change the initial design space and repeat the loop.

1.3. Artificial Data in the Design Loop

Engineers designing complex systems of the future are often forced to use complicated models and simulations in order "to explore…system performance without actually producing and testing each candidate system"[2]. These models and simulations have embedded in them causal and functional relationships, as well as empirical data from the past, which enable the synthesis of new data describing how the system is going to perform. In these cases, the data generated is *artificial* (synthetic) – i.e., it is not obtained by direct measurement of system properties, since no identical (or even similar) system may yet exist. As a result, the artificial data (as well as the model) "cannot be classified as accurate or inaccurate in any absolute sense"[2]. Thus, artificial data stands in stark contrast to empirical data (e.g., temperature readings, historical stock prices), which is directly measured and thereby holds a potentially higher degree of validity (i.e., reliability with respect to the relevant components). Law and Kelton (2000, p. 265) express the difficulties of model validation for complex future systems[3]:

> The ease or difficulty of the validation process depends on the complexity of the system being modelled and on *whether a version of the system currently exists*… For example, a model of a neighbourhood bank would be relatively easy to validate since it could be closely observed. However, a model of the effectiveness of a naval weapons system in the year 2025 would be impossible to validate completely….

As Law and Kelton point out, however, even after a simulation is validated, it "can only be an *approximation* to the actual system"[3]. There must remain some departures from reality in the artificial data – as Herbert Simon (1969) notes: "artificiality connotes perceptual similarity but essential difference."[4] Regardless of the amount of "essential difference", though, it is the "perceived similarity" to reality that leads to model *credibility* (i.e., trustworthiness), which is when a "manager and other key project personnel accept [it] as 'correct.'"[3]

The issues related to the artificial data from performance models (of "unobservable" systems) also apply to the value models in the design loop (see Fig. 1). For the purposes of this paper, it is assumed that a stakeholder's values cannot be observed by direct measurement – meaning any data provided by a mental or constructed model of values (e.g., utility function, set of requirements) is artificial data (discussed further in §2.2). Therefore, when dealing with artificial data from the performance models as well as the value models, stakeholders and design engineers are inherently concerned with the following questions: Are the models truthful? Do I trust the models?

2. Model Purpose and Model Type

For the remainder of the paper – and in the context of complex engineering systems design – two meta-dimensions of models are explored: model purpose and model type. Model purpose can be to predict the *performance* of the system (i.e., performance model) or the *value* stakeholders assign to that performance (i.e., value model), as reflected in Fig. 1. Model type differentiates between where the abstraction of reality (or of values) resides: a *mental* model "represents entities and persons, events and processes, and the operations of complex systems"[5] and resides in the mind; a *constructed* model is a formalization of (one or more) mental model(s), and it can reside, for example, in a computer simulation or on a piece of paper in the form of a diagram. Mental models about reality are created automatically through perception and cognition[5], whereas constructed models are intentionally created for widespread applications, such as system dynamics models of climate change, a discrete-event simulation of a surveillance system, or a system of differential equations describing the diffusion of heat in a homogeneous body. These meta-dimensions of model type as well as model purpose are illustrated in Table 1.

Table 1: Matrix illustrating different kinds of models.

		Purpose: Performance	Purpose: Value
Type	Mental	Mental performance model	Mental value model
Type	Constructed	Constructed performance model	Constructed value model

2.1. The Performance Models

As discussed in §1.3, the performance model for a system is an attempt to represent some relevant portion of reality and predict the outcomes of interest that will occur in that portion of reality. As noted in the matrix above, two distinct types of performance models are present throughout the design process: *mental* and *constructed* models. The first type, the mental model, begins with observation of existing physical systems and their properties. The mental model of even the least complicated system is severely limited, however, by the inherent bounded rationality of human cognition[4]. These general intuitions regarding a system's properties lead to the construction of models of all kinds (e.g., abstract, physical, computer simulations) to predict, capture, and study the specific interactions and behavior of complex systems. The constructed model for a system thus becomes a "store" for all related models of reality (e.g., stochastic models, laws of physics, cost models), effectively formalizing human thought and extending the bounds of human rationality.

2.2. The Value Models

The value model discussed in §1.2 is an attempt to capture stakeholder preferences on performance (i.e., their values). Fischoff (1991) discusses a spectrum of philosophies with regard to the nature (and elicitation) of values[6]. On the one end of the spectrum is the philosophy of articulated values, which assumes that values are self-evident in people's choices. This philosophy is intrinsic, for example, in some valuation methods used by economists (e.g., empirical estimates of a demand curve, or of the value of life[7]). On the other end of the spectrum is the philosophy of basic values, which derives value models from some core set of values through an inferential process (e.g., interviews). This paper assumes a basic value philosophy, as it best approximates a "requirements-based" approach common in systems engineering. Additionally, the risks of such an assumption are potentially less impactful to a design outcome than the assumption of articulated values[6]. From this standpoint, the paper explores the resulting role of the mental models and the constructed models in accurately representing stakeholders' values. Other intermediate positions exist on the spectrum, such as that of partial perspectives[6].

With this understanding and assumption, a value model can be used to predict the preferential order of design alternatives. Such a model assumes that some approximation can be created for a stakeholder's basic values. The *mental* value model can be described as the stakeholder's own *grasp* of what is important to them and how important it is (i.e., the decision-making criteria and their respective weights they carry for the decision). Similar to the case of performance models, a *constructed* value model can store one or more mental models of values.

Within the context of complex system design, there are usually too many alternatives for a stakeholder to evaluate. In addition, the many dimensions of value (and their interactions) are often beyond the capability of a stakeholder's mental value model, due to *bounded valuation* – i.e., bounded rationality applied to the description/prediction of preferences. For the problem of deciding among too many alternatives, a constructed value model can serve as a "stand-in" for the stakeholder's mental value models. Similarly, for the problem of bounded valuation, a constructed value model can serve as a reliable predictor for a potentially overwhelmed or confused stakeholder. Examples of constructed value models are customer value models in economics (i.e., a representation of the worth – in monetary terms – of what a company does for its customers), or multi-attribute utility functions from classical decision theory[8,9]. The constructed value model's reliability is critical to success in system design; for as Hall points out, "to design the wrong value system is to design the wrong system."[10]

3. The Three Body Problem

The two types of models shown in Table 1 (mental and constructed) lead to three key relationships between the elements used in the performance model and in the value model of the design loop. The three elements involved in these key relationships include: a mental model, a constructed model, and a reference (reality or basic values).

3.1. Vis-à-vis Performance and Value Models

For the performance model used in the design loop (see Fig. 1), the three elements (i.e., "bodies") involved are constructed model, mental model, and a reference point being *reality*. The three-body problem thus formed is depicted below in Fig. 2a, where Δ_1^P and Δ_3^P represent Simon's "essential difference" of the mental model and constructed model from reality, respectively[4]. As discussed in §1.3 and §2.1, constructed performance models are formed to extend the bounds of human rationality. In the context of Fig. 2a, classic model validation is then the process of reducing Δ_3^P (i.e., increasing truthfulness). Since the reference of reality is objective – albeit changing (e.g., geopolitical shifts, technology advancement) – the constructed model can be modified in order to "better" approximate reality (with respect to those aspects deemed relevant). Once the difference between the constructed model and reality is small enough (i.e., the constructed model is truthful), the constructed model can be confidently used to update the mental model of reality (reducing Δ_2^P between constructed model and mental model). As a result, this implies a reduction in Δ_1^P between the updated mental model and reality.

In the case of value models, the reference point is basic values, with which the mental model and the constructed model interact. As mentioned in §2.2, this paper assumes a basic values philosophy on value elicitation[†], which leads to the three-body problem illustrated in the diagram of Fig. 2b. The decision maker holds a set of core basic values, and, when a decision needs to be made, he or she attempts to formulate a mental model of his or her values that will be used to make such decision. This mental model can vary in complexity and sophistication. Kahneman (2011) describes two different ways in which a mental model can be formed[11]: fast, automatic, emotional and

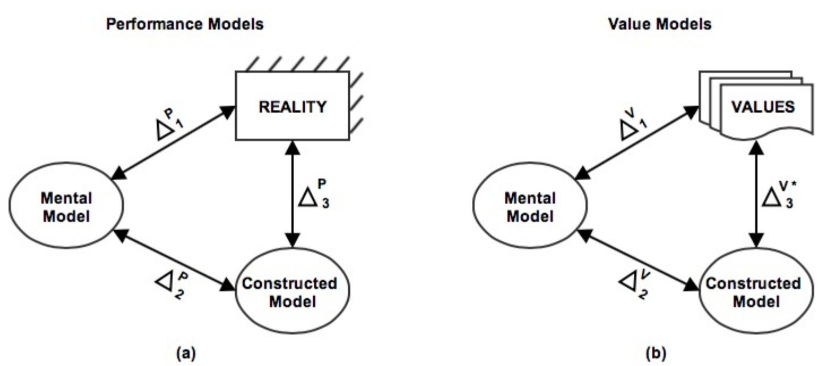

Fig. 2. Three-body problem for (a) performance models and (b) value models. The deltas indicate the amount of "essential difference" between the various elements, with the arrows representing the distance. The starred delta indicates the impossibility of direct comparison and control, since *Values* are only accessible through stakeholders' *Mental Models*.

[†] Although a basic values philosophy is assumed in the context of this paper, the larger research effort attempts to frame the same model-related decision-making problems when an articulated value or partial perspective philosophy are adopted as well.

subconscious (System I) or slow, effortful, logical and conscious (System II). Furthermore, it may also be the case that the decision maker is unable to express his or her values on a given decision and neither System I nor System II can help.

For complex decisions employing multiple objectives and a variety of decision criteria, as well as a multitude of alternatives to choose from, the decision maker can decide to build a constructed model of values (constructed with the help of System II-type mental models). A constructed model extends decision makers' bounded valuation (similar to the constructed performance model extending bounded rationality). It captures decision makers' values by augmenting the mental model's capabilities – whether by saving time in numerous evaluations, or by extending the bounds of valuation. In the case of the value models, since values are not directly observable, Δ_3^{V*} in Fig. 2b is also unobservable. Therefore, unlike the constructed performance model and reality, there is no possibility of validating the constructed value model directly with respect to values. Rather, one always accesses the actual values through the stakeholder's mental model of values. The ultimate goal of the construction of a value model is to allow the stakeholder to better understand his/her own values (i.e., low Δ_1^V), which can be enabled by a low Δ_2^V and Δ_3^V.

Difficulties arise in any attempt to reduce the Δ's present in both the performance and value models. In the performance model, barriers to reducing Δ_3^P (i.e., constructed model validation) include lack of experience or existence of appropriate abstractions (e.g., modeling fluid dynamics before the inception of calculus), which leads to untruthful constructed performance models. Barriers to reducing Δ_2^P include a lack of trust in the constructed performance model. Barriers to reducing Δ_1^P include the bounded rationality of stakeholders to grasp complex systems. In the value model, barriers to reducing Δ_1^V are *cognitive biases* – such as anchoring or availability biases[12] – as well as *bounded valuation*. Barriers to reducing Δ_2^V are *explanation*-related *biases*, which arise, for example, when trying to elicit utility functions from stakeholders through interviews, or if trying to infer the price value of a good through empirical observation of the market. They can also arise due to biases present in choosing a particular value model over another[13]: choosing utility functions to represent a stakeholder's values may lead to different results than if using an AHP model[14]. Finally, since Δ_3^V can only be indirectly observed through Δ_2^V and Δ_1^V, the barriers to reducing Δ_3^V include the barriers inherent to reducing Δ_2^V and Δ_1^V.

3.2. Trust and Truthfulness of Performance and Value Models

In the context of system design, a good decision is one based on a trusted, truthful representation of both reality and values. In performance models, the concept of trust is related to how small Δ_3^P is perceived to be. The concept of truthfulness is related to how close the constructed model actually is to reality (i.e., Δ_3^P). The three-body problem described in §3.1 leads to the matrix shown in Table 2 for performance models, which reflects the comparison of both the perceived and the actual constructed model to reality.

Table 2: Trust and mistrust in the performance model.

		Constructed Model v. Reality	
		Low **Δ₃** (truthful)	High **Δ₃**
Perceived Constructed Model v. Reality	Low **Δ₃** (trust)	**Correct Trust:** Trust present. Correct representation of reality.	**Type I error:** Trust present. Incorrect representation of reality.
	High **Δ₃**	**Type II error:** Trust not present. Correct representation of reality.	**Justified Mistrust:** Trust not present. Incorrect representation of reality.

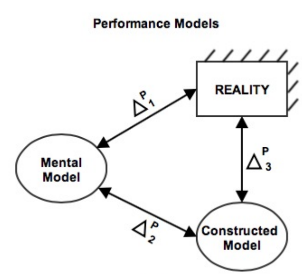

The traditional early system design process consists of validating constructed performance models, and the matrix describes the possible outcomes of this process. The top left quadrant of the matrix is the desirable outcome of the model construction/validation process. There exist three other potential outcomes to such activities, however, ranging from justified mistrust of an untruthful model (bottom right), to an untrusted but truthful model (bottom left), to a trusted and untruthful model (upper right). These outcomes are analogous to null hypothesis tests in statistics, with associated errors of conclusion[15].

In the context of systems design, once the mental and constructed performance models have converged to the upper row of Table 2 (i.e. the performance models are trusted), the outputs of the constructed performance model can be processed by the value models. While the performance models (both mental and constructed) receive great attention in the traditional design process, the same cannot be said of value models. However, decisions ultimately rest on value models – whether explicitly (with constructed value models) or implicitly (with mental models that are assumed truthful). Similar to the performance model, the three-body problem of value model described in §3.1 leads to the matrix shown in Table 3. Unlike the performance model, it is not possible to compare the constructed value model directly to unobservable values (cf. §3.1). The key difference is that the constructed value model is validated through comparison to the mental value model (the only access to values), the same entity responsible also for trust in the constructed model. In other words, determining trust and truthfulness both require the same entity – i.e., the mental value model. Therefore, decreasing the distance Δ_2^V can decrease the distance between the constructed value model and the values themselves, simultaneously increasing the truthfulness of and trust in the constructed model

In practice, it is often assumed that stakeholders' values are self-evident (i.e., articulated, see Fischoff[6]), or that some preliminary constructed model of values (e.g., utility curve) is valid throughout the design process. Constructed models of values could include sets of requirements, if a stakeholder believes that they need a certain level of performance of the system. When dealing with constructed value models, analysts and decision makers tend to implicitly trust the value models they have created (whether utility or requirements-based), risking a Type I error. Type I errors in this case would occur when cognitive biases are present in the mental model of values, and the trusted constructed model reflects those biases, or when a stakeholder sets requirements inappropriately. An acceptable situation is again the lower right quadrant of the matrix, where not making any decision is justifiable. Type II errors can occur when the mental value model is a truthful representation of values, but the constructed model and mental model are divergent (e.g., explanation biases, value model selection biases[13]). The desirable situation is again the top left quadrant, where the mental value model is not biased, and the constructed model is a truthful representation of the mental model of value. Since Δ_1^V and Δ_2^V are reduced, Δ_3^{V*} reduction is also achieved.

Table 3: Trust and Truthfulness in value models.

		Mental Model v. Values	
		Low Δ_1 (truthful)	High Δ_1
Constructed Model v. Mental Model	Low Δ_2 (trust)	**Correct Trust:** Trust present. Correct representation of values.	**Type I error:** Trust present. Incorrect representation of values.
	High Δ_2	**Type II error:** Trust not present. Correct representation of values.	**Justified Mistrust:** Trust not present. Incorrect representation of values.

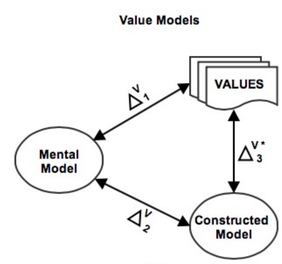

3.3. Interactive Visualization for Delta Reductions in the Value Models

Considering the matrix illustrated in Table 3, the goal is to end up in the upper left quadrant, which implies the reduction of both Δ_2^V and Δ_1^V, and the making of a good decision. Several methods exist for reducing Δ_1^V (e.g., reducing cognitive biases), when a basic values philosophy is assumed. One proposed method is to allow a person to simulate "actually" making a decision between alternatives and then note any regrets that they may have about their choice. Another method is to allow them to consider the ramifications of their expressed (i.e., mental model of) values for many different situations and outcomes (e.g., multi-criteria decision analysis methods such as MATE[16]). A third method might be to directly discuss and uncover biases that skew a person's understanding of his/her actual values (e.g., anchoring, availability[12]). Furthermore, it is possible to introduce decision makers to new attributes that they had not yet considered.

Likewise, methods exist for reducing Δ_2^V (i.e., building trust), when a basic values philosophy is assumed. One method is a rigorous interview process for the construction of utility functions, which reduce framing and other biases that may be present in standard requirements-as-value-statements approaches. Another method is direct comparison of a person's ranked list of alternatives with a ranked list produced by constructed value models (e.g., utility-based rankings). Furthermore, it is possible to compare the ranked lists of alternatives produced by two or more different constructed value models.

In this paper, the authors hypothesize that interactive visualization can facilitate these methods for the reduction of Δ_2^V and Δ_1^V through an iterative process. This process would consist of 1) displaying a constructed value model alongside its outcomes (e.g., ranked list of design alternatives); 2) allowing real-time modifications to the constructed model (e.g., simulating different sets of preferences); and 3) providing instant feedback on the newly changed constructed model, as well as providing comparison of the new results with previous results. Each of these steps potentially allows a user to simulate making a decision, observe their mental model of values propagating throughout the design space to each alternative, discover biased regions of their mental model, and compare their expectations of preferred designs with the constructed model's prediction. As a result of efficiently enabling all of these behaviors, interactive visualization could provide an effective means for reliably representing a stakeholder's actual values through both mental and constructed models.

In the context of the matrix shown in Table 3, this iterative process (through interactive visualization) of delta reductions just described traces a clockwise trajectory, regardless of the starting location. For example, if the initial constructed value model is truthful but untrusted, the stakeholder risks a Type II error in the decision outcome. Interactive visualization allows a stakeholder to observe the constructed model's predictions and compare them with his mental model, thus building trust in the constructed model and pushing the outcome toward the top left – a constructed value model that is both trusted and truthful. Alternatively, if the initial constructed value model is untruthful and untrusted, interactive visualization will alternately allow changing the constructed model and immediately testing those changes so to update the mental model, thereby building the truthfulness first and the trust second. The trajectory will first move toward the Type II error, which as already discussed will quickly converge to the desirable outcome of a constructed model that is both trusted and truthful. Finally, if the initial constructed value model is untruthful but trusted, the stakeholder risks a Type I error in the decision outcome. Through interactive visualization, the errors will quickly become clear as already discussed, immediately causing (proper) mistrust of the constructed model, pushing the outcome toward the bottom right quadrant. From there, the iterative process will follow the clockwise trajectory, ending in a constructed model that is both trusted and truthful.

3.4. Case Demonstration

In this section, a decision about the design of a Space Tug system is used as a case demonstration of part of the interactive process described in §3.3. The Space Tug system is a hypothetical satellite servicer/tender proposed to perform various missions in orbit around the Earth[17]. The visualization tool used for the case demonstration was created with web technologies (D3/Javascript/HTML/CSS), and it runs in web browsers such as Google Chrome and Mozilla Firefox. The tool presents a tradespace of designs, with cost on the x-axis and multi-attribute utility (MAU[8]) on the y-axis. Furthermore, it features a modifiable set of stakeholder preferences (single attribute utility curves, SAU) over design performance – i.e., a constructed model of value. Moving anchor points on the individual preference curves triggers a recalculation of the MAU, which is immediately reflected on the tradespace through a smooth D3 transition[18]. Individual designs on the tradespace can be selected for further inquiry, and their performance levels are also displayed as vertical lines on each attribute's single attribute utility plot.

In order to perform the Space Tug case demonstration, a previously constructed (physics-based) performance model of the system has been used[19]. Using this model, over 900 different Space Tug designs are evaluated. Furthermore, a person familiar with the Space Tug study has been used as a proxy stakeholder for the purposes of this representational demonstration. For the sake of simplicity, only two performance attributes have been chosen over which the stakeholder could state his preferences: Speed and available Delta V. After a first-order value elicitation and Single Attribute Utility (SAU) curves construction, an initial tradespace was derived. Fig. 3a shows this initial tradespace of different Space Tug designs (each point is a distinct design alternative), as well as the stakeholder preferences (SAU curves) associated with this tradespace.

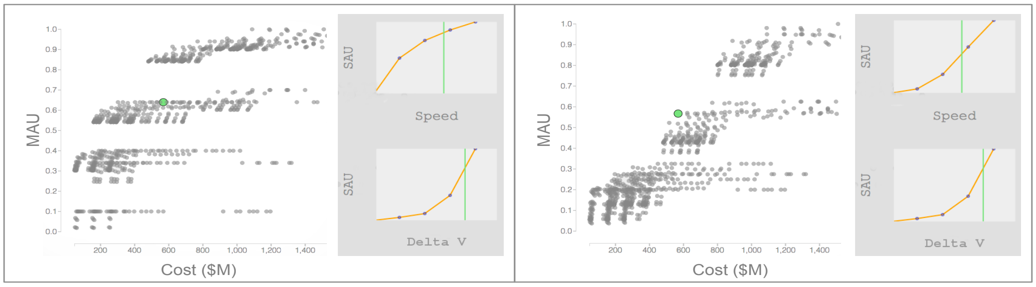

(a) (b)

Fig. 3: (a) Initial tradespace and preferences over performance attributes. (b) Revised preferences and final tradespace.

The stakeholder was allowed to interact with the tradespace: the instant feedback of information about the designs allowed him to efficiently compare and contrast them, thereby gaining insights about the trades in value. After some time, the stakeholder identified a design with favourable Delta V and more than adequate Speed (shown as a larger point in Fig. 3a); however, this design was suboptimal (i.e., far from the Pareto frontier of best MAU for cost). Further investigation made him realize that he would prefer this design to many others that were closer to the Pareto frontier. This event caused him to start doubting the previously derived value model (i.e., SAU curves), hence decreasing his trust in it. In terms of the matrix in Table 3, this process prevented the stakeholder from committing a Type I error, placing him in the “justified mistrust” quadrant of the matrix (movement illustrated in Fig. 4a to Fig. 4b). At this point, the stakeholder started to interact with the initial preferences over performance (shown in Fig. 3a). He soon realized that his stated preferences over Speed were not accurate (due to an availability bias toward previous systems): he cared more about designs with higher Speed levels, and the current preferences gave too high utility to lower Speed designs. By interactively modifying the SAU curve for Speed, the stakeholder arrived at a new SAU curve, shown in Fig. 3b. This SAU curve gives more marginal value to higher speed designs, more accurately reflecting the stakeholder’s preferences (pushing him to the bottom-left quadrant of the matrix in Table 3 – i.e., Fig. 4c: a more truthful model, but not yet trusted). The updated tradespace is shown in Fig. 3b. The initial design is now part of the Pareto efficient frontier, which increases stakeholder trust in the new value model. Moreover, the stakeholder continued to interact with the new tradespace, further building trust in the current value model. This process finally led the stakeholder to trusting the revised utility curves as a more truthful representation of his values, hence landing in the top-left corner of the matrix in Table 3 – i.e., movement from Fig. 4c to Fig. 4d – and leaving him poised to make a “good decision.” The stakeholder was then able to more confidently continue to analyze designs along the Pareto frontier, and ultimately chose a cheaper design alternative for further development.

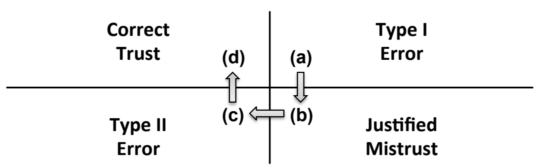

Fig. 4: Stakeholder’s clockwise trajectory through the matrix in Table 3, enabled by visually interactive exploration of preferences.

4. Discussion

While other value philosophies exist, this paper assumes a basic values philosophy, for which stakeholders' basic values can be elicited and used in a value model. This has enabled the creation of the value three-body diagram in Fig. 2b, analogous to the performance three-body diagram of Fig. 2a. The basic values in Fig. 2b are assumed to be a nebulous entity that does not change, but is nonetheless difficult to comprehend due to the obscure nature of its internal structure. The described framework posits that interacting with mental and constructed models of value can help in elucidating internal values, so as to facilitate better decision making. Moreover, this framework, based on the basic values philosophy, encourages attention to cognitive biases existing at the deltas between the various elements of the three-body problem in Fig. 2.

In spite of its relevance for requirements-based approaches, other value elicitation philosophies exist beyond that of basic values, although there is no generally accepted philosophy[6]. One of the goals of ongoing research is to understand the implications of other value philosophies to the described framework. For example, assuming the partial perspective philosophy on value elicitation would automatically imply variations to the three-body diagram shown in Fig. 2b, as the values would be allowed to change in different situations (i.e., different frames[12]).

The representational case presented (Space Tug) serves to provide an initial demonstration of how the process of interactive visualization influences the deltas shown in Fig. 2b, and it provides preliminary face validity of the hypothesis put forth in §3.3. Through visually interacting with his mental and constructed value models, the stakeholder was able to form more truthful and trusted value models. Fig. 4 illustrates the clockwise trajectory taken by the stakeholder through interactive visualization, which ultimately allowed him to make a "good decision" – i.e., one based on a trusted, more truthful representation of his values. Further research will refine the presented framework and use it to formulate experiments, methods, processes, and tools for effectively improving systems engineering analysis and decision making through interactive models, both for performance and value.

5. Conclusion

This paper serves to frame the challenge of building trust and truthfulness in performance and value models used in the design of complex systems. The relationships between the models present in the design loop (mental v. constructed; performance v. value) are described, and the resulting three-body problems of performance models and value models are introduced. Cognitive theories of value and mental models are also discussed within the proposed framework, with special focus given to stakeholder preferences in complex system design. Interactive visualization is proposed as an efficient means for increasing trust and truthfulness of constructed and mental value models. Finally, a custom interactive visualization tool was used to illustrate the framework through a demonstration case, showing how a stakeholder can build correct trust in models for engineering system design decision making.

Acknowledgements

The authors gratefully acknowledge funding for this research provided through MIT Systems Engineering Advancement Research Initiative (SEAri, http://seari.mit.edu) and its sponsors.

References

1. Rhodes DH, Ross AM. Five aspects of engineering complex systems: emerging constructs and methods. *4th Annual IEEE Systems Conference*, San Diego, CA, April 2010.
2. Blanchard BS, Fabrycky WJ. *Systems engineering and analysis.* International Series in Systems Engineering. (2006)
3. Kelton WD, Law AM. *Simulation modeling and analysis*. Boston, MA: McGraw Hill. (2000)
4. Simon HA. *The sciences of the artificial*. Cambridge, MA (1969).
5. Johnson-Laird PN. Mental models and thought. *The Cambridge handbook of thinking and reasoning* (2005): 185-208.
6. Fischhoff B. Value elicitation: is there anything in there? *American Psychologist* 46.8 (1991): 835.
7. Viscusi WK, Harrington Jr JE, Vernon JM. *Economics of regulation and antitrust*. Fourth Edition, The MIT Press; Cambridge, MA. Ch. 20, pp. 730-731.
8. Keeney RL, Raiffa H. *Decisions with multiple objectives: preference and value tradeoffs*. Published by John Wiley & Sons, Inc., Ch. 2, (1976)
9. Von Neumann J, Morgenstern O. Theory of games and economic behavior. *Bull. Amer. Math. Soc* 51 (1945): 498-504.
10. Hall AD. *Metasystems methodology: a new synthesis and unification*. New York: Pergamon Press, 1989.
11. Kahneman D. *Thinking, fast and slow*. Macmillan, 2011.
12. Kahneman D, Tversky A. Choices, values, and frames. *American psychologist* 39.4 (1984): 341.
13. Slovic P. The construction of preference. *American psychologist* 50.5 (1995): 364.
14 Saaty TL. Decision making – the Analytic Hierarchy and Network Process (AHP/ANP). *Journal of Systems Science and Systems Engineering*, Vol. 13, No.1, 2004, pp.1-35.
15. Neyman J, Pearson ES. The testing of statistical hypotheses in relation to probabilities a priori. *Joint Statistical Papers*. Cambridge University Press. pp. 186–202.
16. Chattopadhyay D, Ross AM, Rhodes DH. Combining attributes for systems of systems in multi-attribute tradespace exploration. *7th Conference on Systems Engineering Research*, Loughborough University, UK, April 2009.
17. Fitzgerald ME. *Managing uncertainty in systems with a valuation approach for strategic changeability*. Master of Science Thesis, Aeronautics and Astronautics, MIT, June 2012.
18. D3.js. Retrieved October 1st, 2013, from http://d3js.org.
19. McManus H, Schuman T. Understanding the Orbital Transfer Vehicle Trade Space. *AIAA Space 2003*, Long Beach, CA, September 2003.