SYM-AM-20-089

Proceedings of the Seventeenth Annual Acquisition Research Symposium

Acquisition Research: Creating Synergy for Informed Change

May 13–14, 2020

Published: April 23, 2020

Approved for public release; distribution is unlimited.

Prepared for the Naval Postgraduate School, Monterey, CA 93943.

The research presented in this report was supported by the Acquisition Research Program of the Graduate School of Defense Management at the Naval Postgraduate School.

A Game Theory Approach and Application for Government Acquisition

Scott Rosen, The MITRE Corporation, 7515 Colshire Drive, McLean, VA 22102, srosen@mitre.org

Kelly R. Horinek, The MITRE Corporation, 7515 Colshire Drive, McLean, VA 22102, khorinek@mitre.org

Andreas Tolk, The MITRE Corporation, 903 Enterprise Pkwy #200, Hampton, VA 23666, atolk@mitre.org

Abstract

This research focuses on the development of a game-theoretic approach to drive strategies for contract negotiations in government acquisition. The main objective of the approach is to provide the government the ability to move the bidding vendor into a preferred negotiation point and to understand potential trade-offs around these negotiation points. The approach also expedites the decision-making process in acquisition and provides more transparency of each side during the negotiation. Results from live experiments involving players with acquisition experience from government and industry using this approach will also be presented.

Introduction

The U.S. government has an untapped potential to leverage more quantitative approaches in the federal acquisition system in order to increase the agility of the acquisition process and to procure more preferable products. Game theory has been a research paradigm for studying conflict, bargaining, and negotiations for over 50 years and is widely applied throughout the business domain to develop strategies that reflect priorities and trade-offs. Bierman and Fernandez (1998) provide examples of successful game theory applications from business and industry. As programs become more technical and complex, game theory can help decision-makers identify strategies and leverage information to make data-driven decisions that reflect government priorities and trade-offs. This paper considers the application game theory to the federal acquisition process to facilitate negotiation strategies and proposes an automated framework and interface for applying game theory to bid negotiation in acquisition.

The focus of the game theory application in this paper is in a competitive source selection acquisition where multiple vendors are bidding to provide a product that is a non-commercial product that is still in the stages of development. Within a competitive source selection, there are multiple potential phases in the acquisition life cycle that game theory has the potential to provide support. First, it can steer the decision-makers toward the identification of the most meaningful program-specific evaluation criteria. Second, it enables a more objective and quantitative way to proceed through a source selection by better illuminating key attribute trade-offs for the government. For developmental items, the government can use game theory to take control and drive the key features of the product during the competitive source selection. Figure 1 depicts game theory's main potential intersection points with the acquisition process, using the DAU acquisition subway map to help with the illustration.

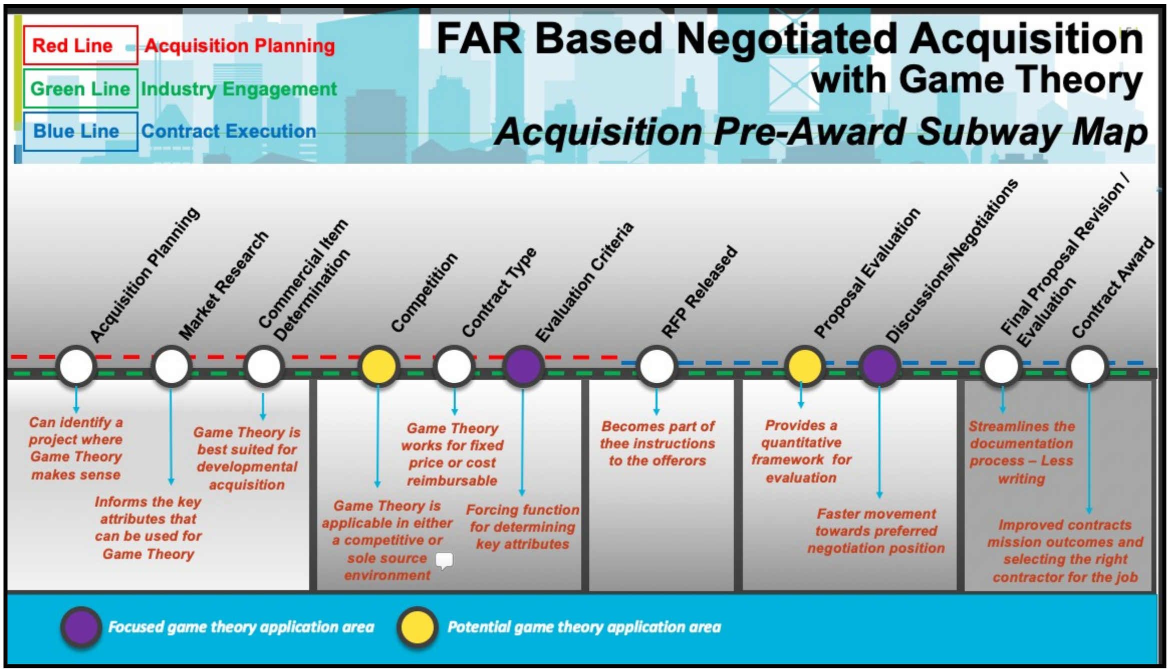

Figure 1. DAU Subway Map with Game Theory Intersection Points

(Currier, 2019)

The rules of the federal acquisition process must be considered before deriving a game theory framework for acquisition. The federal acquisition process is governed by a system of clearly defined rules and regulations codified in the Federal Acquisition Regulation (FAR) (GPO, 2017). However, federal source selections already adhere to a game theory configuration by the mandatory disclosure of evaluation criteria as key discriminators or attributes. By advertising its source selection criteria and relative order of importance, the government signals its trade-off considerations. The current federal acquisition system follows a structured process outlined in the FAR to guide the government in navigating the complexities and cumbersome nature of the source selection process. The FAR presents no obstacles to adopting a game theory framework for acquisition support. Moreover, FAR Part 1.102 “Statement of Guiding Principles for the Federal Acquisition System,” outlines an opportunity to introduce the agility and efficiencies of game theory by allowing strategies, practices, or procedures that are in the best interests of the government that are not specifically limited or prohibited by the FAR, Executive Order, or regulation.

A game theory approach can also be consistent with FAR Part 15.305, “Proposal Evaluation,” as evaluations may be conducted using any rating method or combination of methods, including color or adjectival ratings, numerical weights, and ordinal rankings. Additionally, this approach is consistent with the April 1, 2016, memorandum, *Department of Defense Source Selection Procedures* (Section 2.3 “Develop the Request for Proposals”), where evaluation criteria may be quantitative, qualitative, or a combination of both. Although numerical or percentage weighting of the relative importance of evaluation criteria may not be used in the DoD, assigning quantifiable or value trade-offs in evaluating an offeror’s proposal (found in the Subjective Trade-off and Value Adjusted Total Evaluated Price (VATEP) source selection trade-off approach) is allowable and harmonious with game theory.

This paper presents an overview of an approach to apply game theory to government acquisition as well as a web-based application to institute this approach. The next section continues with an overview of game theory and utility theory approaches. Following that, the Methodology section presents the methodology proposed for applying game theory to acquisition. A web-based prototype for automating the game theory framework is presented in the User Interface Design and Testing Scenario section, along with user testing and feedback. The final section presents the main conclusions and next steps with this research.

Overview of Game Theory and Utility Theory Approaches

Understanding of game theory begins with the proper definition of a game. A game can be viewed as an event that consists of at least two players, has some type of payoff received by the players in the end, and has a set of rules for which the players must abide. Naturally, in government acquisition, one player is the government acquisition office and the other players in the game are the vendors who are bidding on the proposal. The payoff for the government is the quality of the product or the degree that it meets threshold and objective values with respect to their attributes of interest. The payoff for the vendors is money, but it also extends into some implicit attributes like new business relationships and growth. As noted previously, the rules of the game are set by the FAR. Defining the acquisition in this context makes it the right environment for applying game theory. Given the definition of a game, game theory can be viewed as the application of mathematical models to analyze the potential outcomes of a game where players maximize their payoff given the strategies that the other agents will choose. There have been many different theoretical formulations for game theory explored by mathematicians and economists over the last 70 years, with many inspired by the concept of Nash equilibrium. This has allowed for solutions for zero-sum games, Stackelberg games, mixed strategy games, cooperative games, and dominant strategy games. This paper, though, considers acquisition to be a game of continuous payoffs and not a small discrete payoff matrix. This leads to a more generalized utility function approach, which is explained in detail in the Methodology section .

When surveying the literature with respect to game theory and acquisition, there are two primary papers that deserve special attention, as they investigate game theory with respect to acquisition decision-making. Levenson (2014) provides an overview of the constraints of DoD procurement, showing why typical solutions from commercial markets are often not applicable and lead to undesired and unforeseen results. He describes the effects of fixed price and competitive price contracts and concludes that "only when one or more competitors offer innovations that truly reduce the costs of development and production does the government substantially benefit from competition over sole-source procurement without the adverse side effects of cost overruns. Distinguishing between true innovation and optimistic cost estimating, however, can pose a challenge for DoD acquisition officials" (Levenson, 2014, p. 437).

Blott et al. (2015) compiled a set of auction and game theory-based recommendations for DoD acquisitions by synthesizing literature into specific military acquisition categories: procurement with unknown cost and no risk, items with known costs and existent but understood stochastic risk, and items with unknown costs and/or unknown stochastic risk. Some examples further evaluate if multiple competing vendors participate, and if the lot is to be procured from several bidders, potentially at different stages of the project. In summary, the literature survey provided some opportunities for immediate applications, but also revealed the need for continued research. The most pressing needs pertain to how to elicit preferences from decision-makers and apply these utility function methods in an environment under the special constraints of government acquisition. The utility model approaches selected as having the most potential, and both are described in detail along with recommendations for preference elicitation to support each of these methods. These three utility function methods with their corresponding preference elicitation procedures are then tested, and those results are then presented later in the paper.

As mentioned, the game theory framework is a utility function-based approach, so surveying utility models that can be used to best support this proposed framework was the next step in this research. The survey first focuses on multi-criteria decision-making methods, which encompass a utility function approach and on literature which focused specifically on utility function fitting methods. A highlight of the multi-criteria decision-making methods is listed here:

- Wallenius et al. (2008) provide a good overview of the various multiple criteria decision-making and multi-attribute utility theory methods currently in use, such as the analytical hierarchy process (AHP), evolutionary multi-objective optimization (EMO), or multi-objective linear programming (MOLP), placing them into a historical perspective, as the same author group conducted a similar review in 1992 as well.
- A more engineering-leaning perspective is given, among others, by Parnell, Driscoll, and Henderson (2011). They focus on methods and tools to cope with challenges resulting from numerous stakeholders, multiple competing objectives, substantial uncertainty, and significant consequences.
- Velasquez and Hester (2013) conducted a literature review and analysis of multi-criteria methods. They observe that outranking methods, which were prevalent in early approaches, were overtaken by value measurement approaches.
- Agarwal, Sahai, Mishra, Bag, and Singh (2011) provide an alternative viewpoint on the selection of the best multi-criteria decision-making method by focusing on the proper evaluation and selection of suppliers, which is highly relevant in acquisition as well. An additional insight provided by them is the need to evaluate the suppliers based on the inputs of the strategic, functional, and operational levels.

The literature highlighted on utility theory pertains to the formulation of utility functions to reflect the preferences of decision-makers. Slantchev (2012) defines preferences and utilities to support decision-making, including those to be made under uncertainty. As he is writing for political scientists, explanations and examples are easy to follow and do not require an in-depth education in game theoretic mathematical foundations. If data is available that reflects preferences of earlier decision-making processes for either side of the negotiating partners, the methods and algorithms described by Afriat (1967) are still relevant. The application of big data methods supporting utility function definitions is a topic of ongoing research with no predominant methods emerging thus far, although it was initially believed that more literature would be found.

An interesting variant for multi-issue closed negotiations addressing multi-time as well as multi-lateral negotiation strategies is described by Matsune and Fujita (2017), who developed not only the concept but demonstrated them in an agent-based simulation environment. Theoretically, nothing speaks against applying these ideas for acquisition specific challenges as well, but no applications in this domain within the survey were found. What makes the application described in this paper so interesting is the ability to learn the opponents' utility information from observing bidding choices within a strategy.

While the mathematics behind utility theory and utility functions is well understood, how to elicit the knowledge about their preferences from decision-makers is still a challenge in itself. Our survey of the literature did not reveal any predominant strategy. And this is a challenge the government must overcome in order to successfully apply game theory. This research has focused heavily on elicited preferences effectively from decision-makers. The Methodology section discusses that process, how it integrates into the overall game theory framework, and how it is automated within our web-based application.

Methodology

A game theory framework with a utility theory foundation provides an objective and mathematical framework to provide the government with insight into potential attribute trade-offs, as well as improve the transparency of the source selection process. Moreover, through integrating the government's utility function with utility functions of cost established by each vendor, a process for generating better initial bids can be established that gets the government and bidding vendor at a much better starting point for negotiation. This high-level concept for this framework was originally suggested in Simon and Melese (2011) and is extended further in this research. It is provided in Figure 2. The focus of the extension in this paper is a further delineation of a methodology for better calibrating and applying the government and vendor utility function. Other extensions are integrated into the high-level framework covered in this section and are also discussed further at the end of the paper, along with future areas of research.

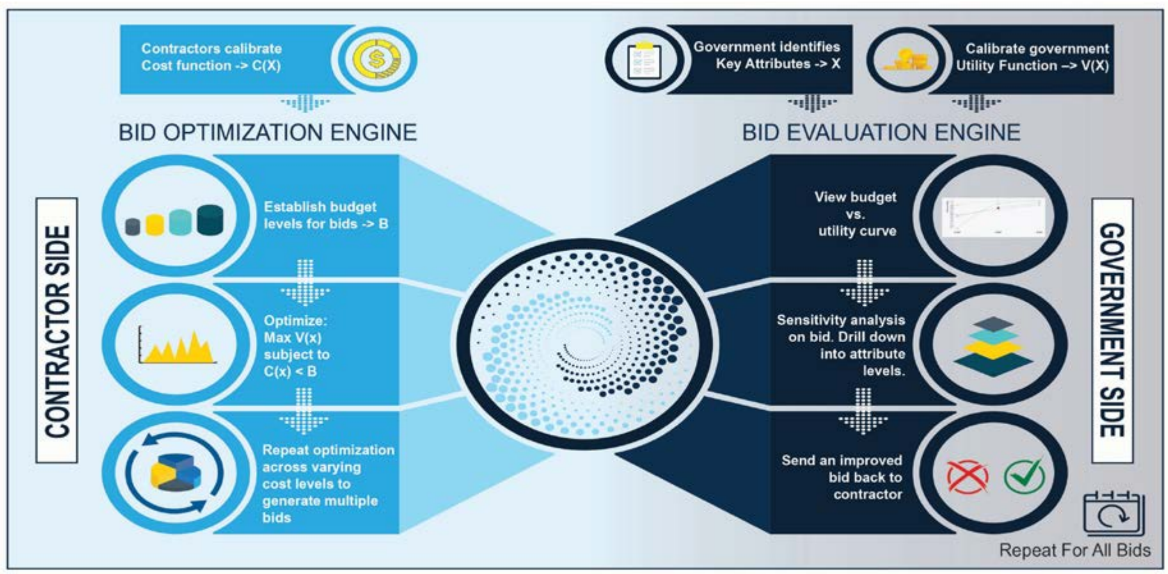

Figure 2. Process Model of Game Theory Framework

The main steps of the framework captured in Figure 2 are listed here:

- Formulate the government's preferences into a utility function parameterized by critical non-cost attributes, criteria, or discriminators.
- Enable industry to work with the utility function to generate bids that maximize the government's utility function.
    - Each vendor calibrates their own cost function parameterized by the same attributes as the government's utility function.
    - Software applications with embedded mathematical programming solvers automate the process of maximizing the utility function and generating the optimal bid for the vendor.
- The government can view bids from all vendors simultaneously with respect to utility and cost. The government can drill into preferred utility/cost values and then examine possible attribute trade-offs that can be used in the next phase of negotiation.

Utility for the government can be modeled through a utility function that represents the government's overall preference for a product. The utility function for the government is parametrized by the key attributes of the product being procured and maps levels of those attributes to a value of overall preference for the product. Utility functions for the bidding vendor can then focus primarily on the cost for achieving the product attribute levels in the government utility function solution. While generating their bid, vendors can then search for solutions that have attribute values that achieve a high value for the government's utility function and result in low cost for their own utility function. This automated optimization framework enables vendors to provide bids quickly and across a range of budget levels.

Moreover, incoming bids from vendors can be viewed as a Pareto curve capturing utility as a function of cost. This leaves the government the ability to do trade-offs between attribute values as well as between cost and attribute values to help provide a counteroffer to the bids that are submitted.

The approach for generating this utility function is presented in the Utility Modeling Approach section, which is parameterized by the key attributes affecting the acquisition decision. The output to this utility function is a level of preference for the government concerning a bid coming from a competing vendor. A bidding vendor uses this utility function as well in helping prepare their bid and integrates that information with their utility model or customized cost model. The procedure for developing this cost model is provided in the Vendor Cost Model Calibration section.

Through an envisioned web-based application, each potential vendor can attempt to maximize the government's utility function subject to their own cost function. The formulation for this optimization problem is presented in the section Bid Optimization. Each vendor's cost function is also parameterized by the same attributes or key criteria that parameterize the government's utility function. This enables the government to perform sensitivity analysis on each bid to evaluate what-if scenarios, and that concept is presented in the Sensitivity Analysis for Bid Evaluation section.

Utility Modeling Approach: The Best-Worst Method

The Best-Worst method originates from Rezaei (2015), and this research has extended that approach to work more smoothly for cases where there are a large number of attributes at hand and when the attributes are binary in nature (result in either a 0/1 or yes/no value). One of the Best-Worst method's features is its ability to perform calibration in a short series of questions. Moreover, these questions have the ability to be phrased to not be overly burdensome to the decision-maker(s). From our observations, having simple and clear acquisition questions to identify key discriminators facilitates the acquisition and conforms to best practices.

Consistent with source selection practices, the procedure for the Best-Worst method starts with selecting the attributes or discriminators that effect the decision. Then feasible ranges are assigned for each of these attributes. The next step is the assignment of weights for each attribute reflecting the preferences and importance. This applies specifically to the attribute to identify key discriminators and does not apply numerical weights to proposals in the source evaluation process. This step begins with selecting the most important attribute as well as selecting the least important attribute. From there, comparisons are made to understand the relative importance of the most important attribute to each of the other attributes. In a similar manner, comparisons are then made to assess the relative importance of the least important attribute to each of the other attributes.

The question phrasing to the decision-maker is the key to getting this approach to work effectively. The decision-maker needs to be directly asked how much more important the most important attribute is with respect to each of the other attributes individually. Mapping qualitative scales to numerical scales was shown to work well in our studies for preserving rank order. For instance, levels, such as, “just as important,” “slightly more important,” “more important,” “significantly more important,” and “extremely more important” were applied with good success while being mapped on a scale of 1–5.

The completion of the Best-Worst assessment procedure is to obtain a preference function in the form: $V(x) = w_1 v_1(x_1) + w_2 v_2(x_2) + \ldots \ldots w_n v_n(x_n)$. This is consistent with the general form of the utility function presented in Equation 1 with x_i = is the level for attribute i, w_i = the weight for attribute i, and v_i = the single attribute value function for the attribute i. The Best-Worst procedure primarily focuses on the weights. Suggestions in this paper for extending to the assessment of the single attribute utility functions $v_1(x_1)$ focus on fitting a function across sample points for each individual attribute. Sampling can be effective with just four points on the utility curve. When doing a qualitative mapping, those points can be referenced as the min, midpoint, target, and max. On a 0–1 scale those reference points were mapped to values of 0, 0.5, 0.75, and 1 respectively. The qualitative assessment questions can first focus on the target. Here the question is asked as to “what is the value of this attribute that you would really want to have.” Then the level representing satisfactory for the attribute is assessed: “What level for this attribute is acceptable and would not hinder my use. It can be considered being like a minimum requirement that is not ideal but gets the job done.” Then the max level for the attribute can be assessed: “What is the level for the most functionality that you could possible handle or need—any more wouldn’t make life any better.” Finally, the minimum level for the attribute is assessed: “What is the maximum attribute level where there is zero utility, or where you would have absolutely no use for this product if this attribute was at this level.”

The Best-Worst method was extended to be more applicable to acquisitions consisting of a large number of attributes (>20). For large number of attributes, the procedure was updated in the following manner:

1. Perform pairwise comparisons across adjacent pairs of attributes starting at attribute #1 and then work down the attribute list.
2. Bin the attributes based on whether the attributes were more important than two attributes, one attributes, or no attributes. End up with three bins: Prime, Mid, Low.
3. Reassess attributes in each bin to make sure they are in the right place.
    a. Ask for best and worst for each bin.
    b. Do pairwise comparison of best in mid and low bin with worst in the higher-level bin.
    c. Repeat 3A and 3B until no more changes are made.
4. Identify the attributes for inclusion into the best-worst method.
    a. Take all attributes in prime bin.
    b. Take best and worst in mid bin.
    c. Take best and worst in low bin.
5. Best-Worst method is then implemented on attributes in the prime bin.
6. Best-Worst method is then implemented on all other attributes mentioned previously.
7. Ask the level of difference between the worst attribute in prime and the best in mid. This level of difference then becomes the difference level for the weights in prime bin and the weights in the remaining bins and the weights are then scaled accordingly.

After these assessment procedures are made, the weights for the preference function can be solved through the optimization formulation outlined in Rezaei (2015). The pairwise comparisons given at the beginning of the assessment procedure can also be used to solve for the weights more effectively as well as for validation of the results.

Vendor Cost Model Calibration

The vendor's utility function is based strictly on cost in this framework and serves as a constraint when performing the utility model optimization. A cost function is required from each bidding vendor and is parametrized by the same key attributes that the government defines and that the government includes in their utility function. The following procedure is followed in this framework in order to calibrate cost functions for each vendor. The procedure essentially generates a sample of cost points across ranges of different attribute values and solves for the resulting equation through the method of least squares. This requires the vendor to generate estimates of cost based on meeting threshold and objective values of different attributes.

- Cost is first based on assigning threshold and objectives for each attribute of interest.
- A total cost for the product is assigned that meets all threshold values.
- For each attribute, the marginal cost for meeting the objective value for that attribute is assigned.
- Extract inflection points between the threshold and objective values for each attribute (if any).
- Interaction effects: If decision-maker prefers to incorporate interaction effects, proceed with a series of questions to measure such effects.
- Resulting sample points follow a fractional factorial design procedure allowing for the training of a regression model to represent cost.

Bid Optimization

The bid optimization can be performed with mathematical programming techniques to enable bids to be automatically generated. Below is a general utility optimization formulation for each vendor specific to their own individual cost constraints.

$$\text{Max } V(x) = \sum w_i v_i(x_i) \tag{1}$$

subject to: $\sum c_i x_i \leq B$

where: x_i = is the level for attribute i

w_i = the weight for attribute i

v_i = the single attribute value function for the attribute i

c_i = the offeror cost function for attribute i

B = budget constraint for maximizing utility i

Sensitivity Analysis for Bid Evaluation

The optimization formulation repeatedly solves for optimal points for each vendor across increasing budget levels that are predefined by the government user. These optimal points maximize the government utility function, subject to a specific budget level. After these solutions have been generated, the government analyzes Pareto curves in order to do comparisons and to select the best offer.

The government has the capability with these solutions and with the Pareto curve to drill down into points to understand the attributes values for the specific attributes of interest. In addition, the government can perform cost trade-offs with particular attributes as well as no cost trade-offs between two pairs of attributes. This allows the government to tweak the most preferred optimal solution coming out of the game theory model. The utility model is an approximation of preference and not a perfect representation of preferences, so this system allows the user to do a final tweaking with respect to the solution of choice. This updated solution can then be returned to the vendor to begin final negotiations on the terms of the contract. The section User Interface Design and Testing Scenario provides screenshots into the user interface that allows the user to perform this sensitivity analysis. In addition, screenshots for cost model and utility model calibration are also provided in this section.

User Interface Design and Testing Scenario

A web-based decision support prototype tool was developed to make the game theory approach an automated process that can be performed interactively between just the government or industry user and the computer. The user interface is shown in this section along with a hypothetical Future Combat Vehicle (FCV) testing scenario that incorporated hypothetical attributes and attribute values but demonstrates the process for using the game theory framework and web-based front end for acquisition.

The hypothetical testing scenario background information giving to the experiment role players consists of explaining that the FCV is critical to moving soldiers to positions of advantage on the battlefield. It is the capability of the future Combat Vehicle to transport and deploy cohesive squads significant distances and across different terrains that sets it apart from the current vehicle. It also has a weaponry package that allows for attacking buildings and targets. There are two combatant scenarios for a new DoD acquisition of FCVs. Each scenario has a different likelihood, but each one must be planned for FCVs not used for one of the missions to have the potential to be used in the other mission.

The acquisition will consist of evaluating the FCV at five budget levels. The base package is $500,000, and extra features that could be necessary to your mission can reach a total cost of upwards of $5,000,000. The government player is interested in receiving proposals at budget levels of $2,000,000, $3,000,000, $4,000,000, and $5,000,000. Proposals at these higher budget levels will contain features that will enable higher levels for the key combat vehicle attributes. The government player is interested in purchasing multiple tanks all with the same add-on packages, but the costs involved during the exercise will pertain to a single tank.

Table 1. Key Attributes Defining Acquisition Decision

Category	Attributes	Units	Definition
Survivability	Ballistic Protection	percent	Likelihood that your vehicle remains mission critical if hit by ballistic weapon
Survivability	IED Protection	percent	Likelihood that your vehicle remains mission critical if hit by an IED
Lethality	Building Target Likelihood	percent	Likelihood of inflicting catastrophic damage to a building target
Lethality	Vehicle Target Likelihood	percent	Likelihood of inflicting catastrophic damage to a vehicle target
Mobility	Road speed	mph	Maximum speed that the FCV can travel
Mobility	Gradient	slope (percent)	Slope in road that vehicle can handle
Mobility	Road range	miles	Related to MPG. How far FCV can go from deployment vehicle

The bidding vendor will be offering vehicles containing a combination of one or more packages for survivability, lethality, and mobility. Approximate prices for packages along with the attribute levels they provide are given here as a guide for the vendor player to provide a cost estimate. Table 2 pertains to price packages relating to het survivability attributes. The FCV shall provide survivability characteristics against (1) Ballistic and (2) Improvised Explosive Devices (IEDs) while meeting transportability requirements. To increase effectiveness, the solution can include a mix of Options 1–4 with multiple packages being able to be provided.

Table 2. Hypothetical Price Packages for Survivability Attributes

Notional Survivability Requirements for Future Combat Vehicle (FCV)

	Option	Packages	Cost ($, unit + Integ)	Weight (lbs)	Ballistic Protection	IED Protection
		Threshold Requirement		45,000	50%	40%
		Objective Requirement		38,000	95%	85%
Survivability	Base	FCV		33,000	25%	15%
	1	Ballistic Armor + Explosive Reactive Armor	$ 300,000.00	6,000	60%	35%
	2	Ballistic Armor + Active Protection System	$ 750,000.00	3,500	80%	85%
	3	Armor Package	$ 900,000.00	10,000	98%	85%
	3A	Ballistic sub-package	$ 400,000.00	3,000	98%	0%
	3B	IED sub-package	$ 500,000.00	4,000	0%	85%
	4	Ballistic + Bar Armor	$ 350,000.00	1,000	75%	15%

The FCV will need the lethality that is required for the mission and it will be another category of the government's key attributes. The lethality attributes each pertain to a given target type and are given as a probability of effectiveness. To increase effectiveness, the solution can include a mix of Options 1–5 with multiple packages being able to be provided. Approximate prices for packages are given here as a guide to the vendor player for a providing a cost estimate along with the attribute levels they provide.

Table 3. Hypothetical Price Packages for Lethality Attributes

Notional Lethality Requirements for Future Combat Vehicle

	Option	Packages	Cost ($, unit + Integ)	Building	armored vehicles
		Threshold Requirement		40%	35%
		Objective Requirement		80%	65%
	1	TOW Missile	$ 150,000.00	10%	40%
Lethality	2	25mm cannon	$ 250,000.00	70%	20%
	3	Javelin Missile	$ 250,000.00	10%	80%
	4	35mm cannon	$ 350,000.00	80%	45%
	5	40mm cannon	$ 650,000.00	90%	65%

The FCV will also need different mobility attributes dependent on what is required for the mission. Multiple packages are provided to guide the bidding vendor. These packages relate to different engine types that could be built with the vehicle.

Table 4. Hypothetical Price Packages for Mobility Attributes

Notional Mobility Requirements for Future Combat Vehicle

	Option	Packages	Cost ($, unit + Integ)	Road-Speed (mph)	Gradient (%)	Road-Range (miles)
		Threshold Requirement		48	27%	200
		Objective Requirement		65	38%	330
Mobility	1	**1,200 HP Diesel**	$ 400,000.00	50	25%	265
	3	**1,000 HP Hybrid**	$ 500,000.00	48	29%	300
	2	**1,500 HP Diesel**	$ 550,000.00	60	35%	200
	4	**1,500 HP Hybrid**	$ 650,000.00	56	33%	365
	5	**1,750 HP Turbine**	$ 900,000.00	70	49%	160

Login and Initial Calibration

The FCV scenario discussed above provides the background information for the government and industry player to role play and participate in the game. The web-based application is used to execute the game.

The remainder of this section provides screenshots and a discussion of the different components of the application. The first step involves creating an account and logging in as a government or vendor user. Government users and vendor users will have different interfaces that guide them through the process. Figure 3 provides a screenshot of the initial login screen of the application.

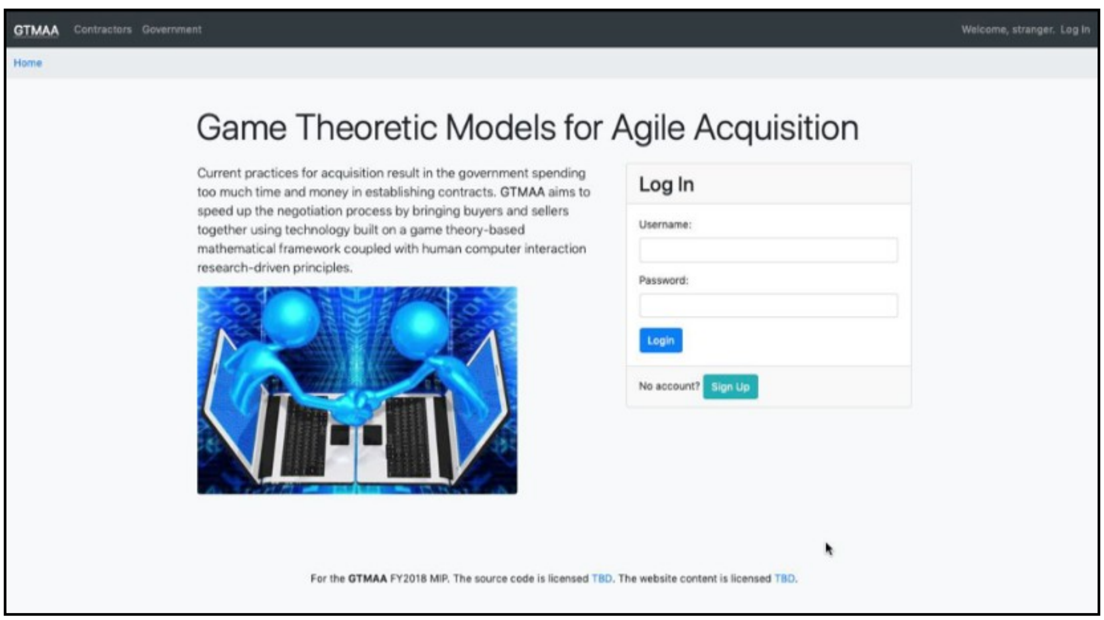

Figure 3. Login Screen for Game Theory Web-Based Application

After logging in, the government user is asked to create a project and specify specific budget levels for which they would like to receive bids at from the vendor. For this example, the budget levels were assigned to $2,000,000, $3,000,000, $4,000,000, and $5,000,000 by the government player. The remainder of this section provides the government screenshots for the utility model calibration in the Utility Model Calibration section. The Cost Model Calibration section provides the main user screen for the vendor cost function calibration. The Bid

Evaluation section then provides the main screen for the government user's bid evaluation.

Utility Model Calibration

First step is for calibrating weights using the Best-Worst method. Figure 4 shows a sample input screen for the Best-Worst method that was used for the FCV scenario. Here, the user has the ability to first select the attributes that they believe are the most important and least important, respectively. After that selection is made, the user then makes pair-wise comparisons between each attribute and the attribute determined to be the most important. There is a drop-down box with a qualitative scale of five options ranging from equally important to extremely less important. Then the user makes these same pair-wise comparisons in the adjacent column with respect to the attribute which was determined to be the least important. There is also a drop-down box for these inputs with options for qualitative responses ranging from equally important to extremely more important. After these selections have been made, the user can submit their inputs. This is all the information that is needed for the weights of the utility function to be calibrated.

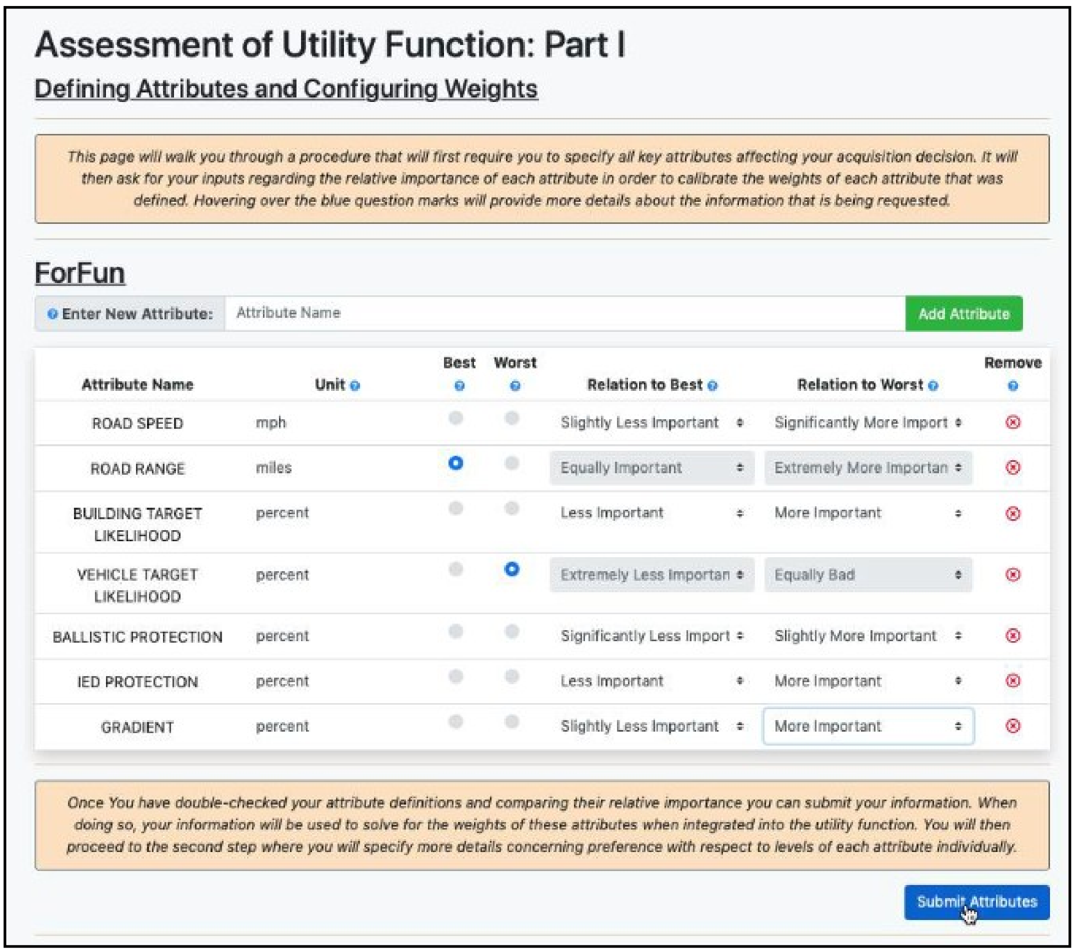

Figure 4. Automated Utility Function Weight Calibration Through Best-Worst Method

The second part of the assessment procedure involves building component utility functions for each attribute. This component utility functions $v_i(x_i)$ are basis utility functions in the aggregate utility function equation: $V(x) = w_1 v_1(x_1) + w_2 v_2(x_2) + \ldots \ldots w_n v_n(x_n)$. The weight terms were computed in the previous step using the Best-Worst method. These component utility functions $v_i(x_i)$ map an attribute value x_i to a level of preference, which is scaled between 0 and 1. A simple procedure is executed in Figure 5 to obtain the information from the user to build the component utility functions for each attribute. This procedure consists of the user specifying the attribute value which represents a minimum value for utility, a threshold value for utility, an objective value for utility, and a maximum value for utility. The specific wordings for minimum, threshold, objective, and maximum are as follows:

- Minimum—The minimum value which equates to zero utility and where the user would have absolutely no use for the product.
- Threshold—The value for the attribute where it becomes acceptable and would not hinder the government's use.
- Objective—The value for the attribute that the user would aspire to have.
- Maximum—The value for the attribute where it equates to as much functionality as the user could possibly handle or need.

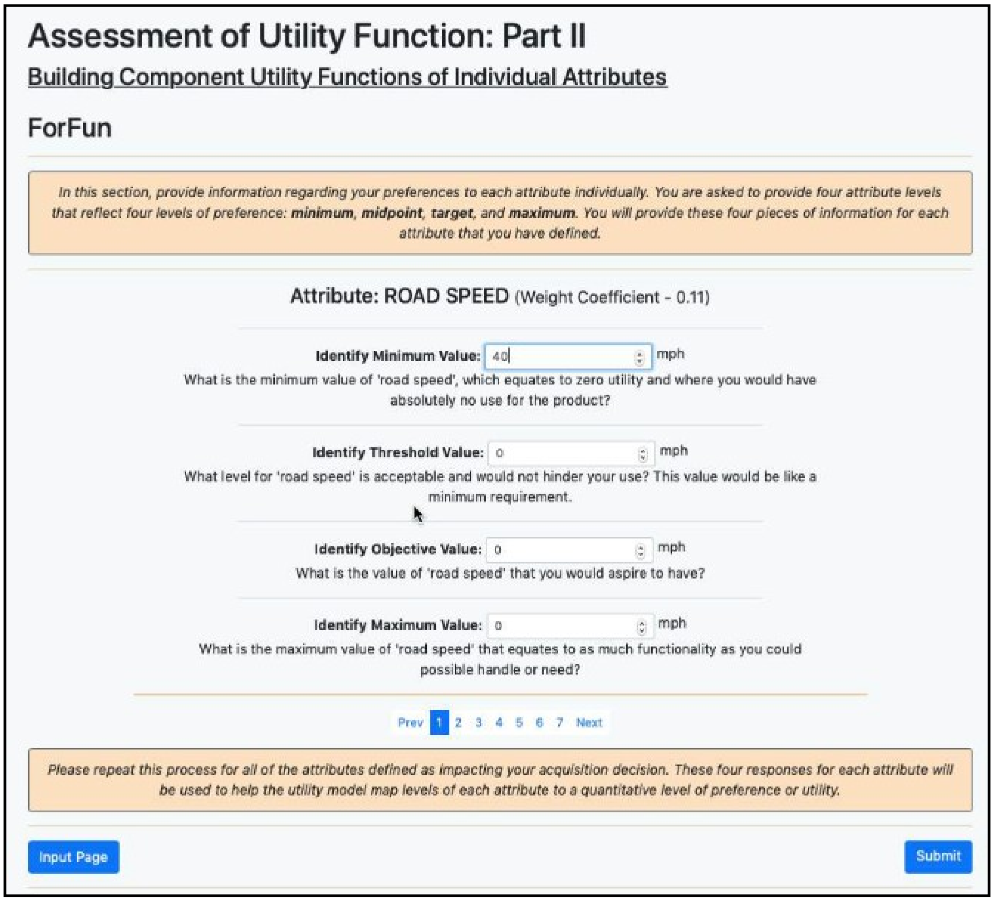

Figure 5. Calibrating Component Utility Functions for each Attribute

The minimum value, threshold value, objective, and maximum value for utility are assigned numerical values of 0, 0.5, 0.75, and 1, respectively. Then given these four sample points of attribute and preference values, a function can be fit through these points to generate an equation representing the component utility function for an attribute.

Cost Model Calibration

This section describes the vendor's main screen for calibrating the cost model. A screenshot from the web-based interface is provided in Figure 6. In this screen, the vendor can view the names of the key attributes defined by the government as well as the threshold and objective values defined for each attribute. The first step for the vendor is to specify the cost needed to produce a product meeting all of the threshold values for the attributes. The vendor then provides a cost for meeting all of the attributes at their objective values. After these to cost points have been assigned, the vendor then assigns marginal costs for each attribute that involves achieving the given attribute at its objective value while all other costs are held at their threshold values. The vendor also has the opportunity to provide cost points to represent interaction effects that may occur between two attributes. That is all of the information that is required from the vendor in order for the game theory framework to start generating recommended bids across the budget levels the government has interested in reviewing bids.

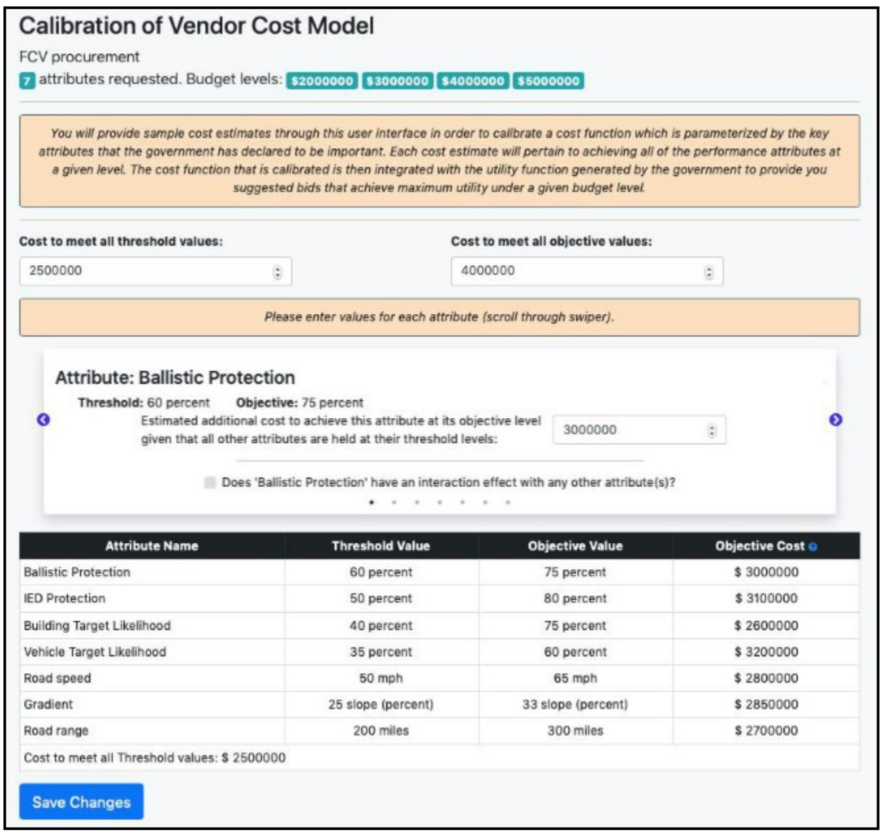

Figure 6. Interface for Vendor Cost Model Calibration

After the vendor user submits their cost estimates, the web-based application is able to solve for their cost function through least squares regression. When that cost function is obtained, it presents it to the user for review. Then the application pairs the vendor's cost function with the government's utility function to formulate a mathematical program. The government utility function becomes the objective function and the vendor cost function becomes the main constraint function. The optimization is then solved at a cost constraint of $2,000,000 to obtain an optimal solution for the vendor at that cost point. The optimization is repeated at $3,000,000, $4,000,000, and $5,000,000 (the budget levels initially set by the government user at the beginning of the FCV example) for the cost bounds to generate optimal solutions for the vendor at the other cost points. After a review of the solution generated by the optimization, the vendor user is then able to submit their bid.

Bid Evaluation

After each vendor user has calibrated their cost model and run the optimization and submitted their bid, the government user can log in to assess how each vendor meets their utility across the range of their pre-specified budget levels. The government user does so by viewing each bid as a Pareto curve, which plots cost versus utility. Each curve on the bid evaluation screen represents a vendor, and each point represents an optimal solution generated by the solver that maximizes the government's utility with respect to cost. In addition, each point contains information as far as the attribute levels that the vendor would provide at that cost point. A screenshot of the bid evaluation screen is provided in Figure 7.

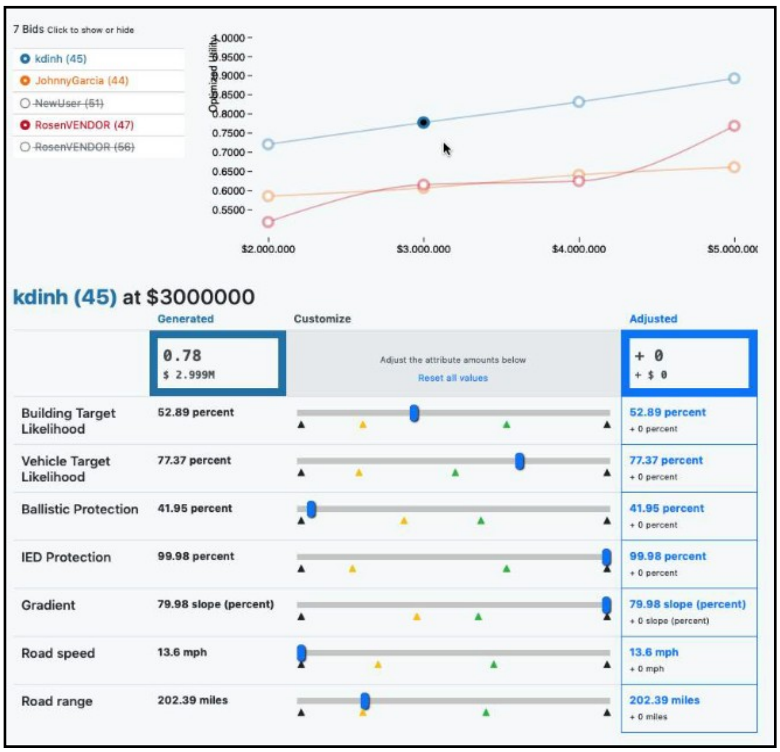

Figure 7. Bid Evaluation Screen for Trade-Off Analysis

The government first has the option of screening out bidders who are not in the ballpark in order to evaluate only the top contenders. In this screenshot, two vendors have already been removed. Then clicking into optimal point on a Pareto curve for a vendor opens up the sensitivity analysis slide bar for the attributes that follow. This first allows the government user to view all of the values associated with each attribute for the bid that the selected vendor is able to provide. The government user can accept this bid as is or can perform what-if sensitivity analysis to explore different possibilities. This can be done by using the sensitivity analysis slide bar to increase or decrease attribute values. When the user does so, they can view the resulting increase or decrease in utility as well as the resulting increase or decrease in cost. This allows the government user to fine-tune their bid to either try to lower cost or to increase an attribute of interest to potentially meet a threshold or objective value. After the sensitivity analysis is performed, the government user can reach out to the selected vendor and begin the final negotiation process with an agreed upon solution that is very close to the final agreed upon solution.

Summary of Feedback from Experimentation

The government and industry players who participated in this experiment each provided feedback from their experience in using the tool for a hypothetical acquisition. The government user's main feedback points are summarized as follows:

- There is so much potential for improved efficiencies due to less proposal writing and no proposal review.
- The competitive range is small; only one or two get there, so being able to quickly remove bids from the competitive range is a plus.
- Having the utility function equations as visual plots really helps with validation
- Additional sensitivity analysis screen for government with ability to better examine trade-offs across vendors would make this even more powerful
- Past performance integration into game theory framework would be a useful compliment to the current framework

The industry participants' main feedback points are summarized below:

- The best benefit to this is that it "decrypted the government's aspirations for an acquisition."
- This is a great tool for examining the best you can give when you can't provide everything.
- Vendors spend tons of money preparing bids/proposals, so this is a better structure for getting closer and requires far fewer resources.
- Vendors noted they were comfortable with a larger number of attributes and could have handled more in this framework.
- Pricing for interaction effects was noted as being very important because there are situations where it was easy to achieve objectives other attributes after achieving and single objective.
- For this to be adopted, website security is key.

Validation of utility function and visualizations was a recommendation and a continued research area. As noted in the feedback, this protype tool can also plot component utility functions 8 by attribute level vs. utility, as shown in Figure 8, to help provide the government user a means for validation. In addition, the weight coefficients are summarized for the user as well as the minimum, threshold, objective, and maximum values that were provided for each attribute.

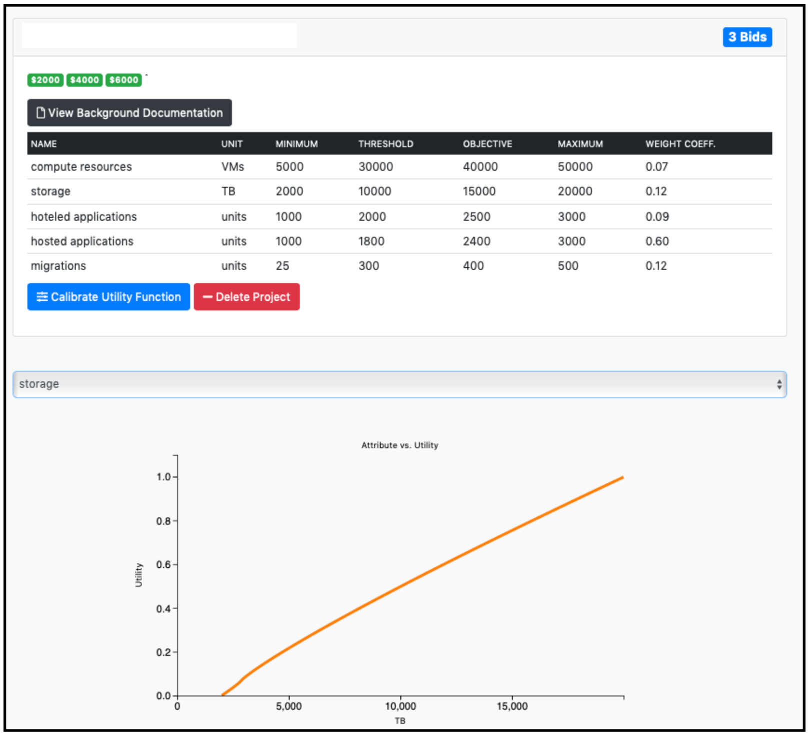

Figure 8. Government Validation Screen for Fitted Utility Function

Conclusions and Future Research

Current practices for acquisition result in the government spending too much time and money in establishing contracts. The idea presented in this paper is to enable negotiation more at the speed of technology by using technology to bring buyers and sellers together. It is shown in this paper that a game theory mathematical framework coupled with Human Computer Interaction principles provides this supporting negotiation technology. This game-theoretic framework is intended to provide a means to illuminate better contracting trade-offs for the government, provide insight into strategies that move the contractor into the government's preferred negotiation point, and expedite the decision-making process in acquisition.

In summary, the game theory framework in support of government acquisition is implemented through a web application consisting of two front end GUIs: One for the government and one for the bidding vendor. The GUI to be used by the vendors is envisioned to help prepare bids by generating automated solutions through the maximization of the government's calibrated utility function. In this GUI, the vendor submits information to calibrate a cost function parameterized by the same key attributes as the utility function and then the optimization is performed The GUI on the government side enables the user to visualize all incoming bids simultaneously and perform sensitivity analysis on their highest valued bids. Lastly, all outputs from the government GUI are traceable, well documented, unambiguous, and repeatable, which can support the government as well when protests occur.

Future research is intended to center around further experimentation with acquisition specialists from government and industry. One open challenge is to better represent uncertainties in choices and consequences. While some sophisticated statistical models exist, they often are not favored by practitioners due to the needs to provide a multitude of data required to apply them. Instead, heuristics that keep the amount of required additional information small seem more likely to be applied, such as fuzzification methods, or the use of certainty factors (Mousavi & Gigerenzer, 2017; Soni et al., 2028).

In addition, research will also involve incorporating vendor past performance into the game theory framework to enable the government user access to that information during the bid evaluation process. Research on how to leverage the quantitative information captured during the game theory process to support future bid protests will also be performed. Preparation for these protests can be a time-consuming process, which can be made more efficient through information captured using this approach.

References

Afriat, S. N. (1967). The construction of utility functions from expenditure data. *International Economic Review, 8*(1), 67–77.

Agarwal, P., Sahai, M., Mishra, V., Bag, M., & Singh, V. (2011). A review of multi-criteria decision making techniques for supplier evaluation and selection. *International Journal of Industrial Engineering Computations, 2*(4), 801–810.

Bierman, H. S., & Fernandez, L. F. (1998). *Game theory with economic applications*. Addison Wesley.

Blott, J., Boardman, N., Cady, A., Elliott, J., Griffin, W., Mastronardi, N., & Quinn, P. (2015). *Auction and game theory based recommendations for DoD acquisitions* (No. USAFA-AM-14-189). Colorado Springs, CO: Air Force Academy.

Currier, M. (2015). DAU contracting subway map. Retrieved from https://www.dau.edu/tools/Documents/Contracting%20Subway%20Map/index.html

Garvey, P. R., & Pinto, C. A. (2009, June). Introduction to functional dependency network analysis. *The MITRE Corporation and Old Dominion, Second International Symposium on Engineering Systems, 5.* Cambridge, MA: MIT.

Garvey, P. R., Pinto, C. A., & Santos, J. R. (2014). Modelling and measuring the operability of interdependent systems and systems of systems: Advances in methods and applications. *International Journal of System of Systems Engineering, 5*(1), 1–24.

Levenson, W. J. (2014). DoD acquisition–To compete or not compete: The placebo of competition. *Defense ARJ, 21*(1), 416–440.

Matsune, T., & Fujita, K. (2017, July). Weighting estimation methods for opponents' utility functions using boosting in multi-time negotiations. In *Agents (ICA), 2017 IEEE International Conference* (pp. 27–32). IEEE.

Mousavi, S., & Gigerenzer, G. (2017). Heuristics are tools for uncertainty. *Homo Oecon, 34*, 361–379.

Office of the Secretary of Defense. (2016, April 1). *Department of Defense source selection procedures* (Memorandum). Washington, DC: Author.

Parnell, G. S., Driscoll, P. J., & Henderson, D. L. (Eds.). (2011). *Decision making in systems engineering and management* (Vol. 81). John Wiley & Sons.

Rezaei, J. (2015). Best-worst multi-criteria decision-making method. *Omega, 53*, 49–57.

Schmidt, B. K. (2017*). Generalized Addition Tallying Operation (GATO)*. Bedford, MA: MITRE Corporation.

Servi, L. D., & Garvey, P. R. (2017). Deriving global criticality conditions from local dependencies using Functional Dependency Network Analysis (FDNA). Systems *Engineering, 20*(4), 297–306.

Simon, J., & Melese, F. (2011). A multiattribute sealed-bid procurement auction with multiple budgets for government vendor selection. *Decision Analysis, 8*(3), 170–179.

Slantchev, B. L. (2012). *Game theory: Preferences and expected utility* (2nd ed.). Technical Report, Political Science Courses. San Diego, CA: University of California.

Soni, M., Mene, S., & Singh, M. M. (2018). Fuzzy mathematical approach to game theory. *International Journal of Statistics and Applied Mathematics, 3*(1), 272–277.

U.S. Government Printing Office (GPO). (2017). *Code of Federal Regulations; Title 48 - Federal Acquisition Regulations System.* Retrieved from https://www.gpo.gov/fdsys/pkg/CFR-2017-title48- vol1/xml/CFR-2017-title48-vol1-sec1-101.xml

Velasquez, M., & Hester, P. T. (2013). An analysis of multi-criteria decision making methods. *International Journal of Operations Research, 10*(2), 56–66.

Wallenius, J., Dyer, J. S., Fishburn, P. C., Steuer, R. E., Zionts, S., & Deb, K. (2008). Multiple criteria decision making, multiattribute utility theory: Recent accomplishments and what lies ahead. *Management Science, 54*(7), 1336–1349.

Acquisition Research Program
Graduate School of Defense Management
Naval Postgraduate School
555 Dyer Road, Ingersoll Hall
Monterey, CA 93943

www.acquisitionresearch.net