Empirical Comparisons of Descriptive Multi-objective Adversary Models in Stackelberg Security Games

Jinshu Cui and Richard S. John

University of Southern California, Department of Psychology
Los Angeles, CA, USA

Abstract. Stackelberg Security Games (SSG) have been used to model defender-attacker relationships for analyzing real-world security resource allocation problems. Research has focused on generating algorithms that are optimal and efficient for defenders, based on a presumed model of adversary choices. However, relatively less has been done descriptively to investigate how well those models capture adversary choices and psychological assumptions about adversary decision making. Using data from three experiments, including over 1000 human subjects playing over 25000 games, this study evaluates adversary choices by comparing 9 adversary models both nomothetically and ideographically in a SSG setting. We found that participants tended to be consistent with utility maximization and avoid a target with high probability of being protected even if the reward or expected value of that target is high. It was also found in two experiments that adversary choices were dependent on the defender's payoffs, even after accounting for attacker's own payoffs.

Keywords: adversary modeling, Stackelberg Security Game, utility function.

1 Introduction

Relationships between attackers and defenders have been modeled as Stackelberg Security Games (SSG). In SSG, a defender moves first as a leader, an attacker then observes the defender's strategy and choose a target to attack. Security resource allocation research has focused on identifying defenders' optimal strategy. One approach is to generate a robust method that is independent of adversaries' strategies[1]. Another approach to determine a defender's optimal strategy is to model adversaries' strategies and construct an optimal defense in response[2, 3]. The approach that considerably models adversaries' choices has been proved to be more effective.

However, relatively less has been done descriptively to investigate how well the adversary-based defenders' algorithms capture adversary decision making and the psychological assumptions of adversaries' choice behavior. This study aims to explore adversaries' choices by comparing different adversary models in a SSG setting. Using data from three experiments, including over 1000 human subjects playing over 25000 games, nine models were evaluated nomothetically and ideographically.

The models compared in this paper all measure adversaries' choices as probabilistic choices, that is, if the probability of choosing one target is higher than that of

R. Poovendran and W. Saad (Eds.): GameSec 2014, LNCS 8840, pp. 309–318, 2014.

choosing an alternative, the adversary will choose that target to attack. In decision making research, Luce's Choice Axiom (LCA)[4] assumes that choice behavior is probabilistic instead of deterministic. McFadden[5, 6] applied LCA to preferential choice in economic analysis. His model was able to exaggerate the differences between different alternatives by exponentiating utilities and the optimal choice is consistent with utility maximization. McKelvey and Palfrey[7] later developed Quantal Response Equilibrium (QRE) in economics, which assumes that the chance of selecting a non-optimal strategy increases as the level of error increases, in which λ captures the rational level (absence of errors) of a player. Since expected utility maximizing is the baseline of a rational decision maker and it is easier to measure a parameter close to 0, we adjusted the quantal response model by reversing the parameter and let λ represent the level of error (softmax): a player chooses randomly when $\lambda \to \infty$ and maximizes expected utility when $\lambda \to 0$. Let $q_i \in [0, 1]$ represent the probability that target t_i will be attacked:

$$q_i(x_i) = \frac{e^{U_{A_i}(x_i)/\lambda}}{\sum_{t_k \epsilon T} e^{U_{A_k}(x_k)/\lambda}},\ \lambda \geq 0 \tag{1}$$

Using the softmax function, we evaluated adversary decision making by assessing four different aspects of the proposed choice models:

(1) Consistency level with utility function maximizing. As suggested by bounded rationality[8], inconsistency with utility function maximization could result when an attacker has limited time and resources to contemplate the optimal choice. The actual choice could deviate from optimal choice and magnitude of deviation is represented by the inconsistency level (λ).
(2) Attention to probability of success. It has been assumed that adversaries pay more attention to probability rather than consequences such that they tend to choose targets with higher probability of success ("soft targets")[9]. We hypothesized that an attacker would pay extra attention to the probability of sucess.
(3) Dependence on defender's utility. Given that adversaries may be driven by emotion, it is reasonable to assume an attacker could "sacrifice" part of their own reward to "hurt" the enemy. We anticipated that terrorists would choose a target that could create more damage to the targeted population, even though that choice could have a lower expected return.
(4) Risk attitude. Past research indicates that emotions can influence risk attitude, such that fear can lead to risk-aversion and anger can lead to risk-seeking[10]. There is little basis to assume adversaries are risk neutral or risk averse[11], especially for an attacker who could experience strong emotions.

2 Models

For each of the proposed models, we aim to capture an adversary's consistency with utility maximization, attention to probability of success, dependence on defender's utility, and risk attitude. The various proposed utility models all utilize the softmax function to calculate the attacker's probability of choosing a particular target. The nine utility functions can be partitioned into five categories: (1) attacker's expected value (EV), (2) attacker's expected utility (EU) accounting for risk attitude, (3) lens model[12, 13] with a weighted average of p(success), attacker's reward and penalty and defender's reward and penalty, (4) lens model accounting for risk attitude, and (5) multi-attribute utility (MAU) model with a weighted average of p(success), attacker's EV and defender's EV[14].

A summary of the nine models grouped in five categories is presented in the Table 1. All nine models capture the inconsistency level (λ). EV is the baseline model. The five lens models and the MAU model capture an attacker's trade-offs among competing cues (or objectives). The EU–α, lens–3–α and lens–5–α allow risk attitude to be accounted for; lens–4, lens–5, lens–5–α, and MAU model take defender's utility into account for attacker's utility function.

2.1 Attacker's Expected Value

The basic utility function of an adversary only captured the expected utility of an attacker who is risk neutral (expected value). The model was first introduced by Yang and colleagues[2] in the name of Quantal Response model. If the attacked target t_i (i = 1,2,..,8) is covered by the defender, the attacker receives penalty P_{A_i} and the defender receives reward R_{D_i}; if the attacked target is not covered by the defender, the attacker received reward R_{A_i} and the defender receives penalty P_{D_i}. Let x_i denotes the probability of a guard at t_i, attacker's expected utility at t_i is

$$U_{A_i}(x_i) = x_i P_{A_i} + (1 - x_i) R_{A_i} \tag{2}$$

Yang et al. [2] further modified the model by adding an extra weight (λ_s, $\lambda_s \geq 0$) to the target that is least protected by the defender, that is, the least defended target is given a bonus in the SOFTMAX calculation. This assumption is consistent with "soft target" hypothesis. Let $S_i(x_i)$ denote whether a target is covered by the least resource:

$$S_i(x_i) = \begin{cases} 1, if x_i \leq x_{i'} \\ 0, otherwise \end{cases} \tag{3}$$

2.2 Attacker's Expected Utility Accounting for Risk Attitude

A simple power utility function was constructed by adding a parameter α to capture risk attitude where α>1 indicates risk seeking and 0<α<1 indicates risk aversion. Assuming the same risk attitude for gain and loss, expected utility of target t_i is:

$$U_{A_i}(x) = x_i \mathrm{P}_{A_i}^{\alpha} + (1 - x_i) \mathrm{R}_{A_i}^{\alpha} \tag{4}$$

2.3 Lens Model

The lens model suggests that attacker judgments depend on a linear combination of multiple observable cues. Therefore, the expected utility function of an attacker can

Table 1. A Summary of the Nine Models Grouped in Five Categories

Category	Model	Abbreviation	Equation
Attacker's expected utility models	Attacker's expected utility model	EU	$q_i(x_i) = \frac{e^{[x_i P_{A_i} + (1-x_i) R_{A_i}]\lambda}}{\sum_{t_k \epsilon T} e^{[x_k P_{A_k} + (1-x_k) R_{A_k}]/\lambda}}$
	Attacker's expected utility model accounting for soft target	EU–soft target	$q_i(x_i) = \frac{e^{[x_i P_{A_i} + (1-x_i) R_{A_i}]/\lambda + \lambda_s S_i(x_i)}}{\sum_{t_k \epsilon T} e^{[x_k P_{A_k} + (1-x_k) R_{A_k}]/\lambda + \lambda_s S_k(x_i)}}$
Attacker's expected utility model accounting for risk attitude	Attacker's expected utility model accounting for risk attitude	EU–α	$q_i(x_i) = \frac{e^{[x_i \mathrm{P}_{A_i}^{\alpha} + (1-x_i) \mathrm{R}_{A_i}^{\alpha}]/\lambda}}{\sum_{t_k \epsilon T} e^{[x_i \mathrm{P}_{A_k}^{\alpha} + (1-x_i) \mathrm{R}_{A_k}^{\alpha}]/\lambda}}$
Lens models	Lens model – three parameters	Lens–3	$q_i(x_i) = \frac{e^{(w_1 x_i + w_2 P_{A_i} + w_3 R_{A_i})/\lambda}}{\sum_{t_k \epsilon T} e^{(w_1 x_k + w_2 P_{A_k} + w_3 R_{A_k})/\lambda}}$
	Lens model – four parameters	Lens–4	$q_i(x_i) = \frac{e^{(w_1 x_i + w_2 P_{A_i} + w_3 R_{A_i} + w_4 (P_{D_i} + R_{D_i}))/\lambda}}{\sum_{t_k \epsilon T} e^{(w_1 x_k + w_2 P_{A_k} + w_3 R_{A_k} + w_4 (P_{D_k} + R_{D_k}))/\lambda}}$
	Lens model – five parameters	Lens–5	$q_i(x_i) = \frac{e^{(w_1 x_i + w_2 P_{A_i} + w_3 R_{A_i} + w_4 P_{D_i} + w_5 R_{D_i})/\lambda}}{\sum_{t_k \epsilon T} e^{(w_1 x_k + w_2 P_{A_k} + w_3 R_{A_k} + w_4 P_{D_k} + w_5 R_{D_k})/\lambda}}$
Lens models accounting for risk attitude	Lens model – three attributes accounting for risk attitude	Lens–3–α	$q_i(x_i) = \frac{e^{(w_1 x_i + w_2 P_{A_i}^{\alpha} + w_3 R_{A_i}^{\alpha})/\lambda}}{\sum_{t_k \epsilon T} e^{(w_1 x_k + w_2 P_{A_k}^{\alpha} + w_3 R_{A_k}^{\alpha})/\lambda}}$
	Lens model – five attributes accounting for risk attitude	Lens–5–α	$q_i(x_i) = \frac{e^{[w_1 x_i + w_2 (P_{A_i}^{\alpha} + R_{A_i}^{\alpha}) + w_3 (P_{D_i} + R_{D_i})]/\lambda}}{\sum_{t_k \epsilon T} e^{[w_1 x_k + w_2 (P_{A_k}^{\alpha} + R_{A_k}^{\alpha}) + w_3 (P_{D_k} + R_{D_k})]/\lambda}}$
Multi-attribute utility model	Multi-attribute utility model	MAU	$q_i(x_i) = \frac{e^{w_1 x_i + w_2 [x_i P_{A_i} + (1-x_i) R_{A_i}] + w_3 [x_i P_{D_i} + (1-x_i) R_{D_i}]}}{\sum_{t_k \epsilon T} e^{w_1 x_k + w_2 [x_i P_{A_k} + (1-x_i) R_{A_k}] + w_3 [x_i P_{D_k} + (1-x_i) R_{D_k}]}}$

be a linear combination of three attributes that are important to the decision (x_i, R_{A_i}, and P_{A_i}). The model, labeled the Subjective Utility Quantal Response (SUQR), was first proposed by Nguyen and colleagues [3]. The utility function was defined as:

$$U_{A_i}(x_i) = w_1 x_i + w_2 P_{A_i} + w_3 R_{A_i} \tag{5}$$

We then extended this utility function to a linear combination of five cues with four weighting parameters (x_i, R_{A_i}, P_{A_i},R_{D_i}, and P_{D_i}) with a common weight for the sum of defender's penalty and reward. We also extended this model to a linear combination of all five cues with separate weighting parameter for each cue:

$$U_{A_i}(x_i) = w_1 x_i + w_2 P_{A_i} + w_3 R_{A_i} + w_4 (P_{D_i} + R_{D_i}) \tag{6}$$

$$U_{A_i}(x_i) = w_1 x_i + w_2 P_{A_i} + w_3 R_{A_i} + w_4 P_{D_i} + w_5 R_{D_i} \tag{7}$$

2.4 Lens Model Accounting for Risk Attitude (lens-α)

Risk attitude can be captured by introducing the parameter α to the lens model:

$$U_{A_i}(x_i) = w_1 x_i + w_2 P_{A_i}^{\alpha} + w_3 R_{A_i}^{\alpha} \tag{8}$$

Risk attitude was also captured in the lens model with five cues. To reduce the number of parameters, we assumed a common weight on attacker's reward and penalty and another common weight on defender's reward and penalty. The evaluation of choosing target t_i then is:

$$U_{A_i}(x_i) = w_1 x_i + w_2 (P_{A_i}^{\alpha} + R_{A_i}^{\alpha}) + w_3 (P_{D_i} + R_{D_i}) \tag{9}$$

2.5 Multi-Attribute Utility Model

Inspired from the lens model which assumed expected utility as a linear combination of different attributes, we developed a new model of multi-attribute utility assuming that the adversary had multiple objectives. We assumed adversaries had three objectives: (1) maximize the probability of success, (2) maximize their expected utility and (3) minimize defender's expected utility. The probability of choosing target t_i is:

$$U_{A_i}(x_i) = w_1 x_i + w_2 EU_{A_i} + w_3 EU_{D_i}$$

$$= w_1 x_i + w_2 [x_i P_{A_i} + (1 - x_i) R_{A_i}] + w_3 [x_i P_{D_i} + (1 - x_i) R_{D_i}] \tag{10}$$

3 Experiment

3.1 Method

The three experiments used the same game paradigm called "The Guards and The Treasure" written in PHP. Each participant was asked to play as an attacker and choose one out of eight gates to attack given x_i, R_{A_i}, P_{A_i},R_{D_i}, and P_{D_i} for each alternative. The three experiments differ in attacker and defender payoff matrixes, defender's guarding strategies and experiment procedures[1]. The published work [1-3] focused on evaluating algorithms for defender strategy in terms of defender EV. This paper reports new analyses of data from the three experiments, focusing on evaluating attackers' choices.

Amazon Mechanical Turk (AMT) was used to collect data. In experiment I, 102 participants, each played 40 rounds, and completed 4080 rounds in total. Forty of the 102 participants were from the US and 48 were from India. Thirty-six (35%) were female. In experiment II, a total of 653 US participants, each played 25 rounds and completed 16325 rounds in total. Two-hundred and seventy-two (42%) were female. In experiment III, a total of 294 US participants, each played 25 to 33 rounds and completed 8538 rounds in total. Eighty-nine (30%) were female.

[1] Please refer to the published worksfor the game procedures, payoff matrices and algorithms.

Table 2. Estimates of Parameters and AIC for Experiments I, II and III

Model	Experiment I		Experiment II		Experiment III	
	AIC	Parameters estimation	AIC	Parameters estimation	AIC	Parameters estimation
EU	15036	λ=.09	60674	λ=.08	33334	λ=.20
EU–soft target	14820	λ=(.09,.59)	50548	λ=(.07,1.89)	27802	λ=(.41, 1.79)
EU–α	15012	λ=.08, α=.86	59169	λ=.06, α =.7	31065	λ=.08, α=.33
Lens–3	14670	λ=.05, w=(-.32,.44,.24)	52014	λ=.04, w=(-.42,.35,.23)	25445	λ=.07, w=(-.16,.18,.67)
Lens–4	14656	λ=.02, w=(-.31,.44,.23,.02)	48218	λ=.01, w=(-.36,.30,.20,.14)	22937	λ=.02, w=(-.47,.03,.19,.31)
Lens–5	14658	λ=.05, w=(-.30,.43,.23,.02,.02)	43265	λ=.02, w=(-.31,.26,.17,.04,.20)	22592	λ=.04, w=(-.44,-.01,.04,.30,.21)
Lens–3–α	14645	λ=.05, w=(-.32,.35,.34), α=1.47	51929	λ=.04, w=(-.42,.30,.28), α =1.25	25159	λ=.07, w=(-.09,.11,.80), α =1.86
Lens–5–α	14624	λ=.08, w=(-.46,.5,.04), α=1.51	48121	λ=.04, w=(-.45,.32,.18), α =1.32	23228	λ=.05, w=(-.58,.07,.35), α =.47
MAU	14973	λ=.08, w=(-.06,.84,.10)	45335	λ=.03, w=(-.32,.39,.29)	26540	λ=.07, w=(-.66,-.01,.33)

3.2 Results

Nomothetic Analysis. Maximum Likelihood Estimation (MLE) [15] was employed to fit the data over all the games played in each of the three experiments and estimate parameters for all nine models. The likelihood function for each model is:

$$\mathrm{L} = \prod_{i=1,2,..,N} q_i(x_i) \tag{11}$$

The Akaike Information Criterion (AIC) [16] was calculated using equation 12 for each model in the three experiments, where k is the number of parameters of a model. AIC is an estimate of the expected, relative distance between the fitted model and the unknown true mechanism that generated the observed data [17]. The model with the minimum AIC is the best among the alternatives.

$$\mathrm{AIC} = -2 \ln \mathrm{L} + 2\mathrm{k} \tag{12}$$

The estimates of the parameters and AICs for the nine models tested in experiments I, II and III are summarized in Table 2. In experiment I, AIC results indicate that models EV and EU–αwere similar in terms of fit; model EU–soft target was slightly better than EU and EU–α. The lens models fit better than model EU–soft target; among lens models, lens–5–α was the best. The MAU model did not fit as well as the linear utility models. Parameter estimates indicate that participants were consistent with maximization of the various evaluation functions ($\lambda < 0.1$) for all nine models. Both the lens models and the MAU model resulted in a negative weight on the probability of being caught, which suggests that participants tended to give a bonus to targets that are less likely to be guarded. Parameter estimates for the four models that captured the weight participants put on defender's rewards and penalties (Lens–4, lens–5, lens–5–α and MAU) suggest that the weight on defender's side was much lower than that put on attacker's rewards and penalties (about 1/10). Finally, model EU–α indicated that participants were risk-averse, while lens–3–$\alpha$and lens–5–α indicated that attackers were risk-seeking.

In experiment II, the AIC fit indices indicated consistency with experiment I; model EV was the worst model among the nine. Model EU–α was slightly better than EV but was worse than EU–soft target. The lens models and MAU were again better than EU–soft target. The MAU model was not as good as the lens models. Among the five lens models, lens–4 was better than lens–3 and lens–5 was better than lens–4. Adding a parameter for risk attitude on lens–3 (lens–3–α) improved the model slightly. Adding a parameter for risk attitude on lens–5 and combining attacker's side and defender's side (lens–5–α) did not improve the model. Parameter estimates indicated that participants were rational ($\lambda<0.1$ for EU–α, lens models, and MAU while $\lambda<0.5$ for EV and EU–soft target). Again, lens models and the MAU model indicated that a negative weight was put on the probability of being caught. Results of the four models that capture the weight attackers place on the defender's rewards and penalties (lens–4, lens–5, lens–5–α and MAU) suggested that the weight on defender's rewards and penalties was as high as the weight on attacker's side. Finally, EU–α and EU–5–α indicated that participants were risk-averse while EU–3–α indicated that participants were risk-seeking.

In experiment III, AIC results were consistent with those from experiments I and II in that model EV was the worst model among the nine. EU–α was slightly better than EV but was worse than EU–soft target. EU–soft target was better than lens–3 and lens–3–α, and was worse than lens–4, lens–5, lens–5–$\alpha$and MAU (all models accounted for defender's rewards and penalties). Among the five lens models, lens–4 was better than lens–3 and lens–5 was better than lens–4. Adding a parameter for risk attitude on lens–3 (lens–3–α) improved the model slightly. Adding a parameter for risk attitude on lens–5and combining attacker's rewards and penalties and defender's rewards and penalties (lens–5–α) did not improve the model. The MAU model did not fit as well as lens–5 but was better than the other seven models. Parameter estimates indicated that participants were rational ($\lambda<0.1$) for all models. Again, lens models and MAU indicated that a negative weight was put on the probability of being caught. Results of the four models that captured the weight participants put on defender's rewards and penalties (lens–4, lens–5, lens–5–α and MAU), suggested that the weight put on defender's rewards and penalties was as high as that put on attacker's rewards and penalties. Finally, EU–α indicated that attackers were risk-averse while lens–3–$\alpha$and lens–5–α indicatedthat attackers were risk-seeking.

Ideographical Analysis. We expected there were individual differences in utility function parameters. For instance, some attackers may have multiple objectives of maximizing expected utility, minimizing the chance of being caught and minimizing their enemies' (defenders) expected utility at the same time (captured in MAU). Some attackers may only maximize their own expected value (captured in EV). It is impossible to differentiate different types of "adversaries" with the nomothetic analysis alone. An ideographical analysis allows us to evaluate how each individual attacker made the decision and how that person is different from others. Again, parameters were estimated using MLE. Since the sample size (N) is small with respect to the number of parameters (k) (N/k < 40 using the k from the most complex model), AICc was calculated for comparisons over different models[17]:

$$\text{AICc} = -2 \ln \text{L} + 2\text{k} \left(\frac{N}{N-k-1}\right) \tag{13}$$

The number of times each model has a minimum AICc is summarized in Table 3. Out of 102 attackers (each playing 40 games) in Experiment I, results indicated that lens–3 scored the minimum AICc most often. MAU model, EU–α and lens–4 also scored the minAICc more often than the other models. In Experiment II, out of 653 attackers (playing 25 games each), results indicated that lens–5 scored the minimum AICc most often; MAU and lens–4 also scored the minimum AICc more often than other models. In Experiment III, out of 294 attackers (each playing 25-33 games), results indicated that lens–5 scored the minimum AICc most often, and lens–4 more often scored the minimum AICc compared to other models. EU and EU–α never scored the minimum AICc across all 294 attackers.

Table 3. Number of Times Model i has Minimum AICc for Experiments I, II and III

	EU	EU – soft target	EU–α	Lens – 3	Lens – 4	Lens – 5	Lens–3–α	Lens–5–α	MAU	Total
Experiment I	7	3	14	28	11	3	7	8	21	102
Experiment II	7	13	6	32	105	282	14	26	168	653
Experiment III	0	29	0	34	56	96	11	33	35	294

4 Discussion and Conclusion

We found that attackers in all three experiments tended to behave consistently with the proposed evaluation functions ($\lambda \rightarrow 0^{+}$). This suggests that in general attackers select targets based on maximizing one of the proposed evaluation functions. The EV model never provided as good a fit as the other eight models, suggesting that the traditional expected value model for an attacker cannot account for adversary choice. Moreover, while model EU–α was superior to model EV, it did not perform as well as the other seven models, suggesting that risk attitude alone does not fully explain adversaries' deviations from EV.

In addition to maximizing attackers' own expected utility, it was found that another predictor of adversaries' choices is defender's payoffs and rewards. In the nomothetic analysis, Experiment I demonstrated that evaluation functions with more parameters (e.g., lens–5) did not fit any better than evaluation functions with fewer parameters (e.g., lens–3). However, in both experiments II and III the evaluation functions with more parameters were better. Model lens–5–α was the best model in experiment I, and lens–5 was the best model in experiments II and III. Both models indicate that attackers take defender's rewards and penalties into account when selecting a target. Additionally, results from experiments II and III indicated a comparable weight of defender's payoffs with the weight of attacker's own payoffs, which implied that attackers gave as much weight to the defenders' rewards and penalties as they did to their own payoffs. The idiographic analysis revealed substantial variability among attackers; however, model lens–5 was found to provide the best fits for the most attackers, consistent with findings from the nomothetic analysis.

We also found that another determinant of adversaries' target selection is the likelihood of success. Participants tended to overvalue the target that was less likely to be guarded. For instance, in the MAU model, which double-counts the probability of success (or probability of being caught) both directly and in the EU calculation, was found to be a competitive model in both the nomothetic analysis and the idiographic analysis. We also found consistently in all three experiments that models accounting directly for success probability (lens models and MAU model) are better than models that account for success probability only in the calculation of EV or EU.

Results from the idiographic analysis indicated that there is no best model among the nine that generally accounts for most of the attackers' choices. Our results suggest that attackers used different evaluation functions to compute the "best" choice in a game. Therefore, individual differences in adversaries (diversity) should be taken into consideration when predicting attacker behavior. It is necessary to identify different types of adversaries in order to predict their choices and to compute optimal strategies for defenders.

References

1. Pita, J., John, R., Maheswaran, R., Tambe, M., Yang, R., Kraus, S.: A robust approach to addressing human adversaries in security games. In: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems, vol. 3 (2012)
2. Yang, R., Kiekintveld, C., Ordóñez, F., Tambe, M., John, R.: Improving resource allocation strategies against human adversaries in security games: An extended study. Artificial Intelligence 195, 440–469 (2012)
3. Nguyen, T.H., Yang, R., Azaria, A., Kraus, S., Tambe, M.: Analyzing the effectiveness of adversary modeling in security games. In: Conference on Artificial Intelligence (2013)
4. Luce, R.D.: Individual choice behavior. John Wiley & Sons, Inc., New York (1959)
5. McFadden, D.L.: Quantal choice analaysis: A survey. Annals of Economic and Social Measurement 5(4), 363–390 (1976)
6. McFadden, D.: Economic choices. American Economic Review, 351–378 (2001)
7. McKelvey, R.D., Palfrey, T.R.: Quantal response equilibria for normal form games. Games and Economic Behavior 10(1), 6–38 (1995)
8. Simon, H.A.: A behavioral model of rational choice. The Quarterly Journal of Economics 69(1), 99–118 (1955)
9. Asal, V.H., Rethemeyer, R.K., Anderson, I., Stein, A., Rizzo, J., Rozea, M.: The softest of targets: A study on terrorist target selection. Journal of Applied Security Research 4(3), 258–278 (2009)
10. Lerner, J.S., Keltner, D.: Beyond valence: Toward a model of emotion-specific influences on judgement and choice. Cognition & Emotion 14(4), 473–493 (2000)
11. Stott, H.P.: Cumulative prospect theory's functional menagerie. Journal of Risk and Uncertainty 32(2), 101–130 (2006)
12. Brunswik, E.: The conceptual framework of psychology, vol. 1. University of Chicago Press (1952)
13. Hammond, K.R.: Probabilistic functioning and the clinical method. Psychological Review 62(4), 255 (1955)

14. Keeney, R.L., Raiffa, H.: Decisions with multiple objectives: Preferences and value trade-offs (1976)
15. Scholz, F.: Maximum likelihood estimation. Encyclopedia of Statistical Sciences (1985)
16. Akaike, H.: A new look at the statistical model identification. IEEE Transactions on Automatic Control 19(6), 716–723 (1974)
17. Burnham, K.P., Anderson, D.R.: Model selection and multimodel inference: A practical information-theoretic approach. Springer (2002)