



HAL
open science

Security under Uncertainty: Adaptive Attackers Are More Challenging to Human Defenders than Random Attackers

Frederic Moisan, Cleotilde Gonzalez

► **To cite this version:**

Frederic Moisan, Cleotilde Gonzalez. Security under Uncertainty: Adaptive Attackers Are More Challenging to Human Defenders than Random Attackers. *Frontiers in Psychology*, 2017, 8, 10 p. hal-03188217

HAL Id: hal-03188217

<https://hal.science/hal-03188217>

Submitted on 1 Apr 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Security under Uncertainty: Adaptive Attackers Are More Challenging to Human Defenders than Random Attackers

Frédéric Moisan¹ and Cleotilde Gonzalez^{2*}

¹ Economics Department, University of Cambridge, Cambridge, United Kingdom, ² Dynamic Decision Making Laboratory, Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA, United States

OPEN ACCESS

Edited by:

V. S. Chandrasekhar Pammi,
Allahabad University, India

Reviewed by:

Stefan Sütterlin,
Inland Norway University of Applied
Sciences, Norway
Ido Erev,
Technion—Israel Institute of
Technology, Israel

*Correspondence:

Cleotilde Gonzalez
coty@cmu.edu

Specialty section:

This article was submitted to
Cognitive Science,
a section of the journal
Frontiers in Psychology

Received: 01 April 2017

Accepted: 29 May 2017

Published: 22 June 2017

Citation:

Moisan F and Gonzalez C (2017)
Security under Uncertainty: Adaptive
Attackers Are More Challenging to
Human Defenders than Random
Attackers. *Front. Psychol.* 8:982.
doi: 10.3389/fpsyg.2017.00982

Game Theory is a common approach used to understand attacker and defender motives, strategies, and allocation of limited security resources. For example, many defense algorithms are based on game-theoretic solutions that conclude that randomization of defense actions assures unpredictability, creating difficulties for a human attacker. However, many game-theoretic solutions often rely on idealized assumptions of decision making that underplay the role of human cognition and information uncertainty. The consequence is that we know little about how effective these algorithms are against human players. Using a simplified security game, we study the type of attack strategy and the uncertainty about an attacker's strategy in a laboratory experiment where participants play the role of defenders against a simulated attacker. Our goal is to compare a human defender's behavior in three levels of uncertainty (Information Level: Certain, Risky, Uncertain) and three types of attacker's strategy (Attacker's strategy: Minimax, Random, Adaptive) in a between-subjects experimental design. Best defense performance is achieved when defenders play against a minimax and a random attack strategy compared to an adaptive strategy. Furthermore, when payoffs are certain, defenders are as efficient against random attack strategy as they are against an adaptive strategy, but when payoffs are uncertain, defenders have most difficulties defending against an adaptive attacker compared to a random attacker. We conclude that given conditions of uncertainty in many security problems, defense algorithms would be more efficient if they are adaptive to the attacker actions, taking advantage of the attacker's human inefficiencies.

Keywords: defense strategies, security games, game theory, human behavior, learning

INTRODUCTION

Security problems involve offensive and defensive actions across nations, institutions, and individuals. Attackers aim at stealing and getting access to assets, information, and goods while defenders allocate their limited security resources to prevent attackers from stealing their goods. In home security for example, a home owner may assign alarm systems to strategic locations of the house. Yet, attackers foreseeing the way home owners behave, may be able to find simple and unexpected ways to break into the house (i.e., through the front door). Defending against

intelligent unauthorized intrusions in the cyber world can be even more challenging, given the hyper-dimensionality of the environment, the type of digital weapons used, the speed of operations and large number of nodes to protect against a relative high number of potential attackers (Gonzalez et al., 2015). In this research we address a basic question of how do human defenders behave under several levels of uncertainty and various types of attack strategies.

Game theory is a common formalized way to inspire the development of defense algorithms in several security problems (e.g., Roy et al., 2010; Tambe, 2011; Shieh et al., 2012). For example, Stackelberg games (i.e., strategic sequential games in which one's own strategy is reactive to one's rival's actions), are common in the design of algorithms that help allocating limited security resources, and results have been successfully applied to multiple naturalistic settings (e.g., Jain et al., 2010; Fang et al., 2013). Although these strategies have been mostly demonstrated in the physical world, a parallel situation occurs in the cyber world, where there is a need of protecting from electronic criminal activities. Researchers have also turned to using game theory to understand security and defense strategies in the cyber world (Alpcan and Başar, 2010; Dutt et al., 2013; Gonzalez, 2013). Attacks in the cyber world (i.e., cyberattacks) use digital weapons that are often imperceptible to the human senses; they are not limited by geography and political boundaries; they require of highly sophisticated technical knowledge, and they may be highly dynamic and distributed. Thus, a defender in the cyber world may need strategies that are dynamic and adaptive to sophisticated attackers, in contrast to currently common static and non-adaptive defense algorithms (Bier and Azaiez, 2008; Abbasi et al., 2015a,b).

To build effective dynamic and adaptive defense algorithms we need to address at least two strong assumptions in the science of security games (Nguyen et al., 2016) and behavioral game theory more generally (Gonzalez et al., 2015): information certainty and human rationality. Current defense algorithms inspired by game theory assume that a defender has perfect information about the payoff matrix and the attacker's behaviors. They also often assume that players are perfectly rational and are able to account and process all information accurately. In a review of learning models from behavioral game theory, Camerer (2003) concludes that most models rely on a full information assumption, and they would not be able to predict behavior in conditions where this information is not available. Similarly, most of these models make cognitively implausible assumptions, such as players being able to remember large amounts of information (Gonzalez et al., 2015).

This paper contributes to understanding how to possibly address these two challenges by learning from a laboratory experiment how human defenders deal with information uncertainty in games in which the attacker algorithm is random, conservative (i.e., minimizes losses) or adaptive (i.e., adapts to the defender's behavior), using an adversarial security game with payoff asymmetry that mimics real life interactions between a defender and an attacker. The conclusions that we draw from our study are general in nature and have applications to several security problems. We discuss the implications our results have

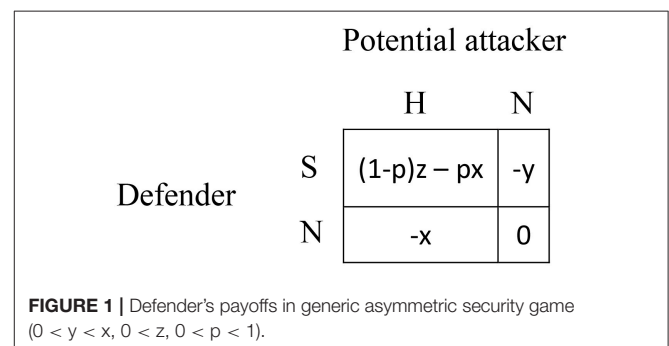
for the design of defense algorithms against adaptive attackers under conditions of uncertainty.

An Asymmetric Cybersecurity Game, Attack Algorithms, and Information Uncertainty

Figure 1 illustrates a generic non-cooperative game with two-players, each of them being able to take two possible actions. In the context of cybersecurity, attacker, and defender interact in non-cooperative ways: The players have conflicting interests (one's gains correspond to the other's losses) characterized by a zero-sum property. In other words, there is no value in cooperating in such interactions because no player can win without making the other one lose. Furthermore, the game is fully strategic since any player's best move strictly depends on the other player's move.

Consider a simple strategic game where a potential attacker may either hack the defender's network (H) or do nothing (N) and a defender that may either surveil the network (S) or do nothing (N). In such a scenario, there exists some probability p ($0 < p < 1$) that the hack is successful despite surveillance. This means that there is a probability $1 - p$ that the hack has been stopped. The Appendix provides a game-theoretic analysis of this scenario and discusses the various assumptions of the outcomes from the strategic actions (x , y , and z), as well as the assumptions regarding the probability of the hacker's success (p).

Experimentally controlling for different attack strategies while investigating human defenders, is a common research approach (Lieberman, 1962; Messick, 1967; Fox, 1972; Coricelli, 2004; Shachat and Swarthout, 2004, 2012; Dürsch et al., 2006; Spiliopoulos, 2008). In competitive scenarios, there often exist well-known simple strategies that vary in the level of rationality and level of adaptation to the defender's behavior. A *Random* strategy (Fox, 1972) is a fixed and static strategy consisting of choosing every option with an equal probability. Such rule is independent of the payoff matrix and the defender's behavior. A *Minimax* strategy implies the minimization of the possible loss of a worst case scenario (Lieberman, 1962; Messick, 1967; Fox, 1972; Shachat and Swarthout, 2004). When defined as a mixed strategy in a two-player zero-sum game, this principle guarantees the same expected payoff independently of the defender's choice (see Appendix). An *Adaptive* strategy may imply that an attacker accounts for the history of the defender's actions, estimates



the defender’s most likely move in the current round and choose a best response accordingly (Brown, 1951; Robinson, 1951; Messick, 1967; Dürsch et al., 2006). Unlike the other strategies, an Adaptive strategy takes into account the other player’s past behavior to decide on the own agent’s current behavior. The specific adaptive strategy that we consider in this study corresponds to the well-known fictitious play (Brown, 1951; Robinson, 1951).

Uncertainty regarding the attacker’s actions and outcomes in repeated strategic interactions is of particular interest in cybersecurity and the lack of explicitly described information naturally leads people to rely more on their own past experience. Risky behaviors may be contrastingly different in situations where people rely on descriptions or experience to make decisions (Hertwig et al., 2004). This body of research suggests fundamental psychological differences in how humans make decisions under uncertainty (Hau et al., 2008; Hertwig and Erev, 2009). In behavioral game theory, uncertainty effects are also of great interest (Martin et al., 2014; Gonzalez et al., 2015). For example, Bereby-Meyer and Roth (2006) have studied players’ abilities to learn cooperation in a repeated PD game where the payoff descriptions were more or less noisy: in each outcome, the payoffs were presented either deterministically or probabilistically. The speed of individual learning was substantially diminished when the payoffs were noisy (non-deterministic), even though players could monitor one another’s past actions perfectly. More generally, these studies show that even small changes in payoff environment can have a large effect on collective behavior. Our main general hypothesis is that introducing uncertainty in the attacker’s payoffs (through the probability p) can affect the human defender’s behavior in different ways depending on the attacker’s strategy: uncertainty may lead defenders to follow noisier and more unpredictable behavior unless the attacker’s strategy is sufficiently naive and exploitable.

METHODS

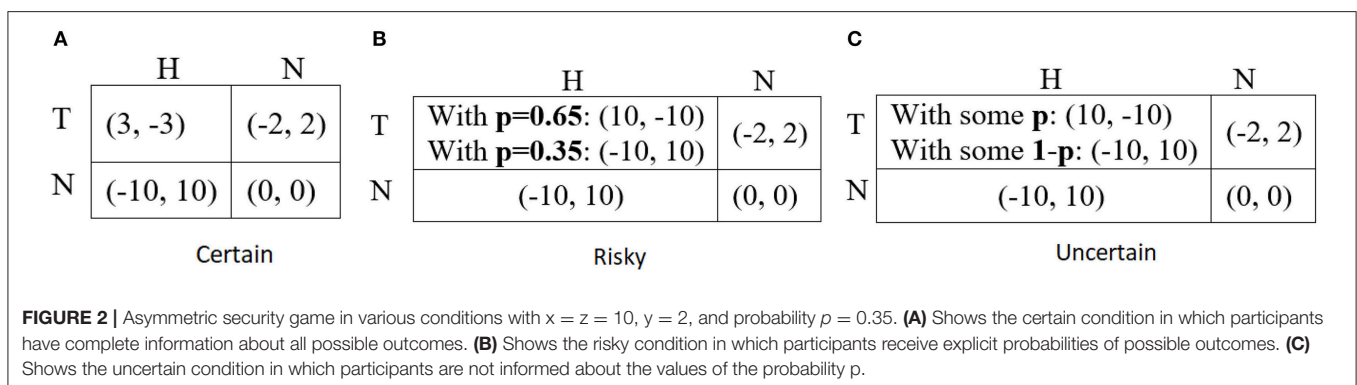
We examine the combined effects of (1) the attacker’s preprogrammed strategy and (2) the level of uncertainty in the outcome description, on a defender’s behavior in the asymmetric security game (Figure 1), using a 3 (Information Level: Certain, Risky, Uncertain) × 3 (Attacker’s strategy: Minimax, Random, Adaptive) between-subjects experimental design.

Under Certain information, participants have complete information about all possible outcomes; they receive the game as in Figure 2A, corresponding to a certainty case of the game presented in Figure 1. For example, when the defender chose S and the potential attacker chose N, the outcome (S,N) indicates a payoff of −2 points to the defender and 2 points to the attacker.

Under Risky information (Figure 2B), participants only have partial information about possible payoffs they can obtain as the defender chooses S and the attacker selects H. Under Uncertain information (Figure 2C), participants are never informed about the actual value of probability p that determines the (S, H) consequence. In this case, participants only know that one player will win while the other will lose depending on the unknown probability p . Importantly, the games in all information conditions are theoretically equivalent in the sense that the same cells of the game carry the same expected payoffs: the value of p in Figures 2B,C is 0.35; the deterministic payoffs in the case of the (S,H) outcome from the Certain conditions correspond to the expected payoffs in the same outcome (where $p = 0.35$) in the two other conditions (Risky and Uncertain descriptions).

The attacker’s strategy is manipulated into three algorithms: Minimax, Random, and Adaptive. In Minimax, the attacker always follows the Minimax principle by selecting H with $\frac{2}{15}$ probability (and therefore N with $\frac{13}{15}$ probability) in every round (see Appendix). These values derive from the payoffs in the baseline setting of Figure 2A. In this strategic game, the defender’s Minimax strategy is to select S with a $\frac{2}{3}$ probability and N with a $\frac{1}{3}$ probability (in each round, the defender’s corresponding expected payoff is $-\frac{4}{3}$ and potential attacker’s corresponding expected payoff is $\frac{4}{3}$). Note that this particular game allows for a clear distinction between both players’ optimal and the Random strategy, which consists of a purely random choice where the attacker selects H or N with equal probability (0.5) at every round. The Adaptive algorithm is as follows: in the very first round of the game, the attacker selects H or N with equal probability (as with the Random strategy), and in all subsequent rounds, its choice is based on the defender’s history of past moves (see Appendix).

In all conditions, participants are asked to repeatedly play against the same attacker strategy for 100 rounds (the participants are not informed about the exact number of rounds). At every round $r > 1$, participants receive feedback indicating the actual outcome in the previous round.



Participants

Nine Hundred and Twenty Seven American individuals (61% male; $M_{age} = 30.8$, $SD_{age} = 9.3$) were recruited through Amazon Mechanical Turk. Participants were randomly assigned to one of the 9 different conditions previously described: $N(\text{Minimax-Certain}) = 107$, $N(\text{Minimax-Risky}) = 99$, $N(\text{Minimax-Uncertain}) = 101$, $N(\text{Random-Certain}) = 103$, $N(\text{Random-Risky}) = 102$, $N(\text{Random-Uncertain}) = 105$, $N(\text{Adaptive-Certain}) = 110$, $N(\text{Adaptive-Risky}) = 100$, and $N(\text{Adaptive-Uncertain}) = 100$. Upon completion of the experiment, each person was paid based on their performance in the task. The average time spent to complete the task was 5 min 09 s, and the average amount of total payment was \$0.88, including a fixed participation fee of \$0.3. This research complied with the American Psychological Association Code of Ethics and the tenets of the Declaration of Helsinki and it was approved by the Institutional Review Board at Carnegie Mellon University. Informed consent was obtained from each participant.

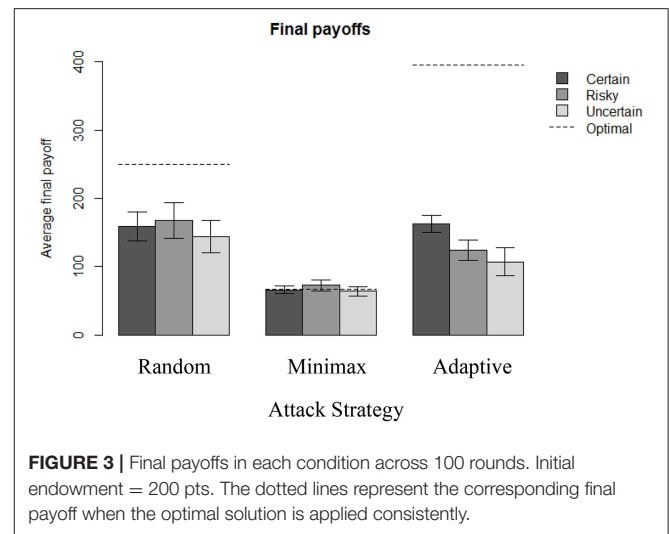
Procedure

Starting with an initial capital of 200 points, participants were instructed to repeatedly play a game with an unknown preprogrammed attacker strategy, where each round would determine either a gain of extra credits or a loss of points from their capital. They were informed that the immediate outcome of each round depended on their own choice between two options (A and B), as well as on the attacker's decision. Note that the labeling of the players and their actions were different from those shown in **Figure 2**, in order to avoid uncovering any focal point that may bias people's behavior ("opponent" instead of "attacker," option "A" instead of "S" or "H," and option "B" instead of "N"). Participants were told that their income by the end of the experiment (in US dollars) would be calculated based on the value of their remaining total capital (if any) with a conversion rate of US\$0.01 for every point.

Participants were not provided with any information regarding the attacker's strategy, except that the attacker was motivated to steal as many points as possible from their capital. They were provided with the payoff matrix corresponding to the experimental condition (as in **Figure 2**). The display of the payoff matrix was determined randomly from four combinations of actions [i.e., (A,A), (A,B), (B,A), or (B,B)] to control for possible effects of the game's display and the buttons' labels. In every round, feedback was provided about what choice each player made in the previous round together with the resulting payoffs.

RESULTS

The maximum expected final payoff for the defender was obtained when playing optimally against the Adaptive strategy (395.5 pts, see Appendix), and playing against Random (250 pts) or Minimax (66.7 pts) strategies was less beneficial. **Figure 3** provides a summary of the corresponding final payoffs that participants obtained in each condition. Note that although participants did not lose money in this experiment (any final negative final payoff was simply reduced to a zero gain), the



values used in **Figure 3** are based on raw data that potentially include negative final payoffs.

All payoffs can be compared to a reference payoff of 66.7 pts, which can be guaranteed to the defender against the Minimax strategy in every round, because the game is zero-sum. We observe in **Figure 3** that this reference payoff is reached under the Minimax strategy in all uncertainty conditions. Average final payoffs from all conditions are at least as good as this reference payoff. Also, defenders' behavior appear better when the attacker's strategy is Adaptive or Random, compared to Minimax. However, the defenders' behavior is suboptimal compared to the maximum expected payoff they could have obtained against a Random and Adaptive strategy over 100 rounds.

Looking at the interaction effect between the attacker's strategy type and the level of uncertainty on the participant's final payoff, **Table 1** provides evidence that those factors significantly influence people's payoffs in the game (after playing 100 rounds).

Figure 3 and **Table 1** reveal an effect of the attacker's strategy: the best final payoffs are obtained when playing against Random ($N = 310$, $M = 156$ pts, $SD = 125$), which are significantly better than when playing against Adaptive [$N = 310$, $M = 132$ pts, $SD = 87$; two sample t -test: $t_{(552)} = 2.8$, $p < 0.01$, $d = 0.22$]. Playing against Adaptive is, however, still more significantly beneficial than playing against Minimax [$N = 307$, $M = 67$ pts, $SD = 35$; two sample t -test: $t_{(408)} = 12$, $p < 0.001$, $d = 0.98$]. However, outcome uncertainty is also significant: playing in the Certain condition ($N = 320$, $M = 129$ pts, $SD = 89$) is significantly more profitable than playing in the Uncertain condition [$N = 306$, $M = 105$ pts, $SD = 101$; two sample t -test: $t_{(606)} = 3.1$, $p < 0.01$, $d = 0.25$]. The interaction effects presented in **Figure 3** suggest that uncertainty only has a significant effect on the overall payoff when playing against Adaptive strategy. Increasing uncertainty significantly decreases efficiency: final payoffs in the Adaptive-Certain condition ($N = 110$, $M = 163$ pts, $SD = 68$) are significantly better than Adaptive-Risky [$N = 100$, $M = 124$ pts, $SD = 78$; two sample t -test: $t_{(198)} = 3.8$, $p < 0.001$, $d = 0.53$] and Adaptive-Uncertain [$N = 100$, $M = 107$ pts, $SD =$

TABLE 1 | Factorial analysis of variance for final payoffs.

Source	ANOVA tests		
	Df	F	Eta-Squared
(A) Attacker's strategy	2	82.39**	0.15
(B) Outcome uncertainty	2	6.11*	0.01
A × B (interaction)	4	3.44*	0.01
Error (within groups)	918		

** $p < 0.001$; * $p < 0.01$.

104; two sample t -test: $t_{(169)} = 4.5, p < 0.001, d = 0.64$]. On the other hand, increasing uncertainty has no significant effect when playing against Random and Optimal.

Moreover, when the payoffs are all known and deterministic, people are as efficient against Random as against Adaptive ($M = 159$ pts in Random-Certain condition, and $M = 163$ pts in Adaptive-Certain condition). One could even argue that people play better against Adaptive because of the slightly higher final payoff on average and the smaller standard deviation ($SD = 113$ in Random-Certain condition, and $SD = 68$ in Adaptive-Certain condition). This observation is particularly surprising because of the very different complexity in computing the best response in both of these conditions: learning to play optimally against Random (by always playing S) is easier as it is shown when risky or uncertain information is introduced.

Choice Behavior

Figure 4 shows that playing the game in the Certain condition leads to an average behavior that closely approaches the prediction of the Optimal principle (selecting S with a $\frac{2}{3}$ probability, as shown earlier). The rate of selecting S in the Certain condition is 65% ($N = 320$), which is not significantly different from the theoretical Minimax solution of 66.66% [$t_{(319)} = -0.74, p = 0.46, d = 0.08$].

However, participants' behavior differed significantly in the first round [$F_{(2, 616)} = 4.22, p = 0.015$]. Introducing some uncertainty in the (S,H) outcome leads to a significant increase in the rate of selecting S: 75% of participants ($N = 301$) chose S during the first round in the Risky condition, which is significantly different from the Certain condition at the 1% level [$t_{(618)} = -2.84, p = 0.005, d = 0.23$]. Similarly, 72% of participants ($N = 306$) chose S during the first round in the Uncertain condition, which is significantly different from the Certain condition at the 5% level [$t_{(624)} = -2.03, p = 0.04, d = 0.16$].

Table 2 shows the effect of the attacker's strategy. Uncertainty level has no significant effect. Results indicate a significant effect of rounds on average behavior, but the interaction between uncertainty and attacker's strategy is not significant.

Figure 5 illustrates the differences in overall average choices depending on the attacker's strategy over the 100 rounds, and also reveals different levels of heterogeneity in individual behavior across conditions. The highest proportion of selecting S is reached when playing against Random ($N = 310, M = 82\%, SD = 19$), which is significantly higher than when playing against the

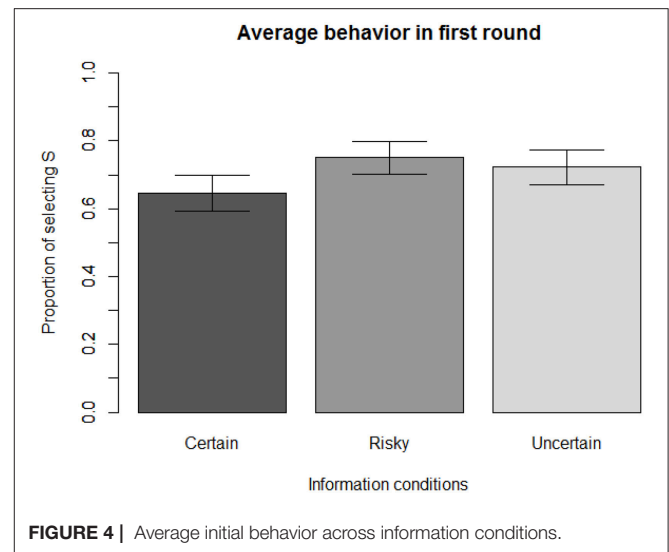


FIGURE 4 | Average initial behavior across information conditions.

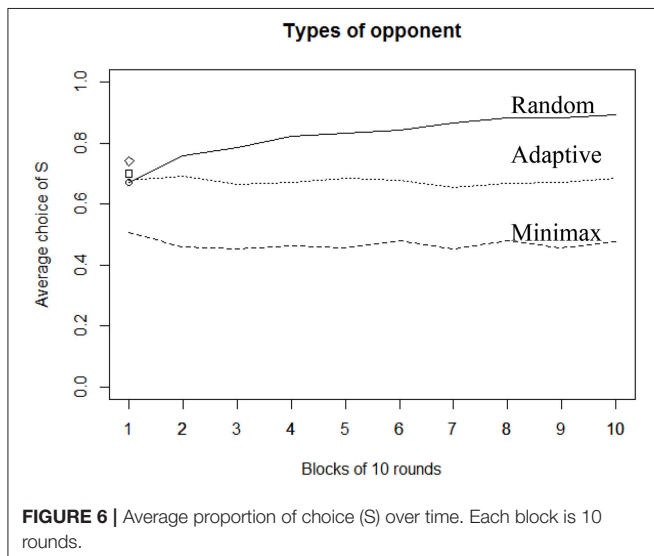
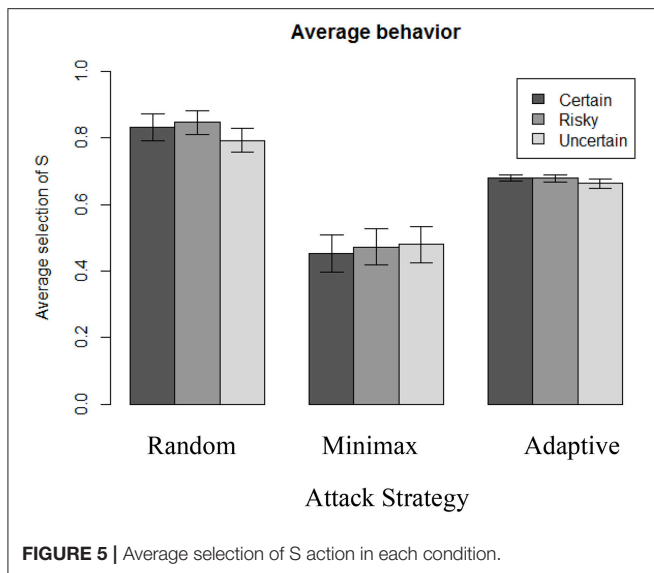
TABLE 2 | Factorial analysis of variance for average choice over blocks of 10 rounds.

Source	ANOVA tests			
	Df	Df _{error}	F	Eta-Squared
(A) Attacker's strategy	2	918	245.24**	0.21
(B) Outcome uncertainty	2	918	0.89	<0.01
(C) Blocks of 10 rounds	7.43	6,820.28	9.45**†	<0.01
A × B (interaction)	4	918	0.93	<0.01
A × C (interaction)	14.86	6,820.28	14.04**†	0.01
B × C (interaction)	18	8,262	1.42	<0.01
A × B × C (interaction)	36	8,262	1.00	<0.01

** $p < 0.001$; †Greenhouse, Geisser corrected.

Adaptive [$N = 310, M = 67\%, SD = 6; t_{(364)} = 12.9, p < 0.001, d = 1.06$]. Playing against Adaptive, however, leads to a more frequent selection of S than when playing against the Minimax strategy [$N = 307, M = 47\%, SD = 28; t_{(332)} = 12.6, p < 0.001, d = 0.99$].

Figure 6 depicts the overall effect of the number of rounds over the proportion of S choices. The main observation is that only the Random opponent's strategy leads defenders to adapt and increase their proportion of selecting S over time. In this case, participants learn to play the best response against the Random strategy (which corresponds to always selecting S). Also, average behavior drifts away from the theoretical Minimax solution (playing S with probability 0.67). Instead, when playing against the Minimax opponent strategy, we observe rather stable behavior across time. Average selection of S is above 65% in the first round and drops to 44% within the first 20 rounds before stabilizing for the remaining rounds. As a consequence, during the very first rounds, people's behavior again drifts away from the theoretical Minimax solution, but in a different direction as compared to the previous condition (playing against Random). This observation indicates that people do not learn to minimize the variance of their payoff (which consists in always playing T).



Instead, they quickly become more indifferent between their options, regardless of the level of uncertainty.

Finally, when playing against Adaptive strategy, there is no learning effect. The standard deviation of the attacker’s average behavior is, however, larger than in the Random condition (SD in Adaptive = 19; SD in Random = 5), suggesting that react to Adaptive strategy, but do not follow an adaptive strategy themselves (if they did, the attacker’s average behavior would then converge toward the Nash equilibrium).

Switching Behavior

We analyzed the number of times that participants switch decisions from one option to the other in the course of 100 rounds. This analysis is relevant as it is possible that two individuals who share the same choice proportion of S reach that stage through different exploration strategies. As shown in

TABLE 3 | Factorial analysis of variance for switching behavior over blocks of 10 rounds.

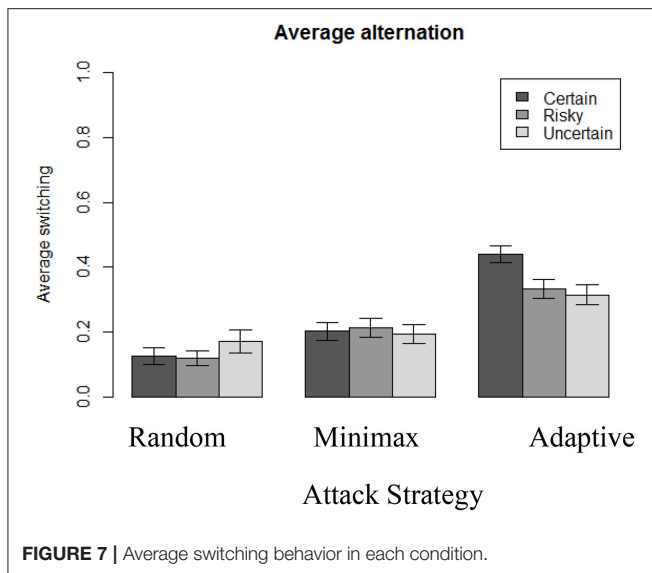
Source	ANOVA tests			
	Df	Df _{error}	F	Eta-Squared
(A) Attacker’s strategy	2	918	191.20**	0.17
(B) Outcome uncertainty	2	918	5.06**	<0.01
(C) Blocks of 10 rounds	7.40	6,791.28	20.09**†	0.01
A × B (interaction)	4	918	10.52**	0.02
A × C (interaction)	14.80	6,791.28	16.96**†	0.02
B × C (interaction)	18	8,262	0.76	<0.01
A × B × C (interaction)	29.59	6,791.28	2.01**†	<0.01

** $p < 0.001$; †Greenhouse, Geisser corrected.

in **Table 3**, there is a clear effect of both outcome uncertainty and the attacker’s strategy on switching behavior. Furthermore, the interaction effect between the two factors is significant, as illustrated in **Figure 7**.

Switching occurs more often when playing against Adaptive ($N = 310$, $M = 36\%$, $SD = 16$) than when playing against Minimax [$N = 307$, $M = 20\%$, $SD = 15$; $t_{(614)} = 13.06$, $p < 0.001$, $d = 1.03$], which is also significantly more than when playing against Random [$N = 310$, $M = 14\%$, $SD = 15$; $t_{(615)} = 5.38$, $p < 0.001$, $d = 1.42$]. We also observe a main effect of uncertainty: switching in the Certain condition ($N = 320$, $M = 26\%$, $SD = 19$) is higher than in the Risky condition [$N = 301$, $M = 22\%$, $SD = 17$; $t_{(613)} = 2.64$, $p < 0.01$, $d = 0.22$] and the Uncertain condition ($N = 306$, $M = 22\%$, $SD = 17$; $t_{(622)} = 2.64$, $p = 0.02$, $d = 0.22$). These results suggest an interaction effect of the two variables (type of attacker strategy and outcome uncertainty). Indeed, when playing against Random, participants alternate significantly more (at the 5% level) under the Uncertain condition ($N = 103$, $M = 17\%$, $SD = 18$) than under the Certain condition [$N = 105$, $M = 12\%$, $SD = 13$; $t_{(190.6)} = 2.08$, $p = 0.039$, $d = 0.32$]. On the other hand, when playing against Adaptive, participants alternate significantly less under the Uncertain condition ($N = 100$, $M = 31\%$, $SD = 15$) than under the Certain condition [$N = 110$, $M = 44\%$, $SD = 14$; $t_{(201.3)} = 6.12$, $p < 0.001$, $d = 0.90$].

Moreover, **Table 3** shows the existence of a significant temporal effect: participants tend to switch less over time, and there is also a significant interaction between outcome uncertainty, the attacker’s strategy, and time. Participants switch more over time when they play against Adaptive, whereas they switch less often over time when they play against Random or Minimax opponent’s strategy. Furthermore, the switching rate over time depends on outcome uncertainty: people learn to switch more in Adaptive-Certain condition than in the Adaptive-Risky and Adaptive-Uncertain conditions, where the switching rate remains constant over time. Similarly, when playing against Random, the decreased switching rate is also different depending on outcome uncertainty: people learn to alternate slightly less in both the Random-Risky and Random-Uncertain conditions than in the Random-Certain condition. No difference is found across conditions when playing against Minimax strategy.



Subjective Randomization

A possible explanation for the increased rate of exploration over time against an Adaptive strategy and with full information may be a reflection of players' attempt to become unpredictable. Such explanation suggests that people may deliberately try to behave more randomly themselves when they play against an Adaptive strategy than when they play against other types of attackers. Intentional random behavior is difficult for humans to detect and perceive (e.g., Rapoport and Budescu, 1997), and many current defense strategies rely on randomization of defense resources based on game-theoretic results (Nguyen et al., 2016). To test for this explanation we used a common non-parametric test for randomness to measure independent and identically distributed (i.i.d.) behavior: the Wald-Wolfowitz runs test. This test relies on the number of runs found in a time series. It compares the observed number of runs of a dichotomous variable (e.g., participant's choice between S and N in the above asymmetric security game) against the expected number of runs. More (less) runs than expected indicate the existence of over (under) alternation in the time series. **Figure 8** presents the proportion of participants for whom the subjective randomization is not supported in the first (left panel) and the last (right panel) 50 rounds.

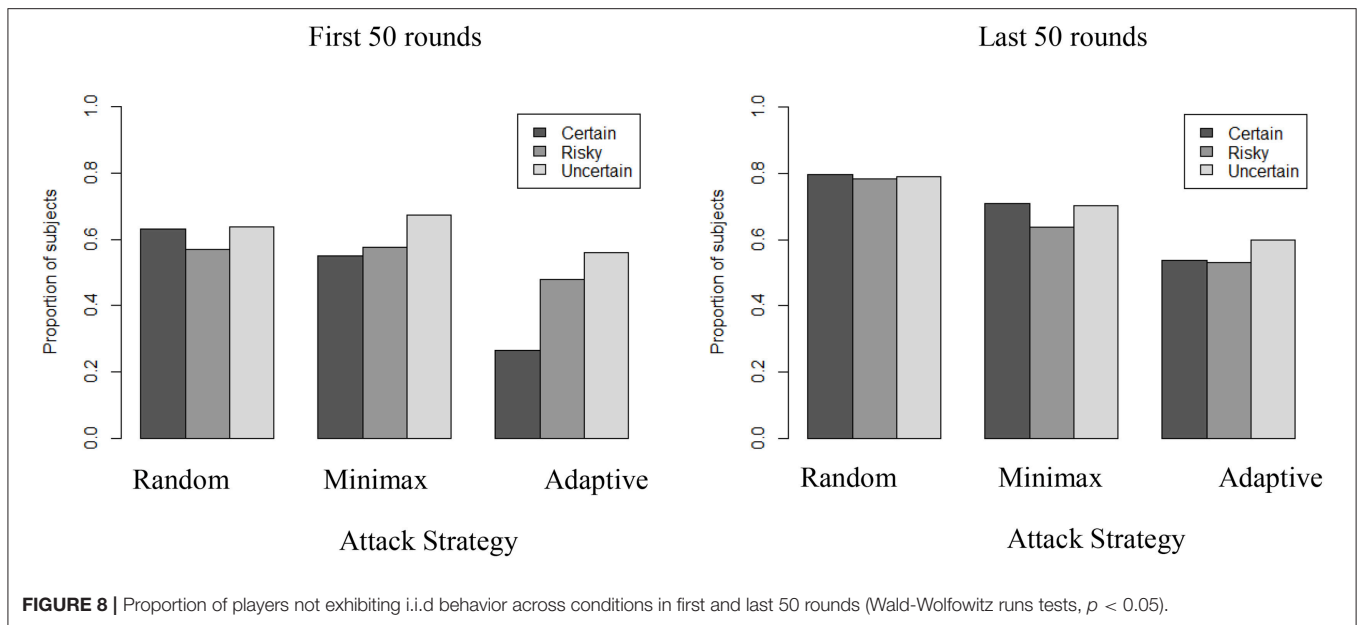
In general, a majority of participants (>50%) do not exhibit subjective random behavior and the level of uncertainty has a clear influence when playing against an Adaptive strategy, particularly in the initial rounds. Higher subjective randomness is observed in Adaptive-Certain than in either Adaptive-Risky or Adaptive-Uncertain in the first 50 rounds ($p < 0.001$), but not in the last 50 rounds. Also, in the Certain condition, subjective randomization is more common in the first 50 rounds if playing against Adaptive (26% of participants do not exhibit i.i.d.) than when playing against Minimax (55% of participants do not exhibit i.i.d.) or Random opponent strategy (63% of participants do not exhibit i.i.d.). Together with the results in the previous section, these results suggest that when playing

against an adaptive strategy, more certainty leads people to behave intentionally more randomly at first (first 50 rounds), before they uncover the benefits of regular alternations (last 50 rounds).

DISCUSSION

Perhaps the most surprising finding is that defenders behave more optimally and with less effort (switching) when confronting attackers represented by random strategies. Defenders seem to focus their attention on how to maximize their payoffs in the long run, even if it implies losses once in a while. They seem to gradually learn that the attacker's strategy is unchanging over time. These observations are interesting as they suggest that unpredictable or random strategies often used as defense mechanisms such as "moving target defense" (Evans et al., 2011; Zhuang et al., 2013) may not be as effective against human attackers. We also find that stochastic conservative strategies such as the Minimax attack strategy are easy to exploit by human defenders, and this is relevant because such attack strategies are commonly used to generate defense schedules in real life scenarios in Stackelberg security games (Pita et al., 2010, 2012; Yang et al., 2011). Our results suggest however, that humans are able to learn to exploit this opponent's strategy, and that an adaptive strategy, which accounts for the opponent's dynamics of behavior, would be more efficient.

Another surprising finding suggests that a reduction of outcome uncertainty did not necessarily lead defenders to execute more optimal actions. Unlike Bereby-Meyer and Roth (2006), we found no significant effect of uncertainty on the speed of learning. Furthermore, this observation distinguishes asymmetric security games from existing games against nature (individual non-strategic decision problems) that reveal significant behavioral differences depending on whether information is descriptive or experiential (Hau et al., 2008; Hertwig and Erev, 2009). Defenders learn to reach good performance when playing against a random attacker where the level of outcome uncertainty seems to have only little effect. In other words, people are no more or less efficient against a random attacker when there is no uncertainty. In contrast, under uncertainty, defenders have difficulty behaving more accurately against an adaptive attacker, and they exert more effort in their attempts. Because playing optimally against a random strategy (selecting the same fixed action at every trial) can be, in principle, more accessible than playing optimally against an adaptive strategy (regularly switching actions in a precise way), one would reasonably expect the former to be more profitable. Instead, we observe that playing against these two different attack strategies can be similarly profitable under certainty. A possible interpretation is that people are naturally more reluctant to play a very basic strategy, which can be easily exploited (e.g., always playing the same action could be easily learned by the opponent), rather than a more complex one, which appears to be more robust to protect them from any type of exploitation (e.g., dynamic behavior makes it more difficult for the opponent to anticipate one's future behavior).



Finally, we find that when defenders interact with adaptive strategies and are given full information, they exert more effort to behave more intentionally in a “random manner,” as a possible attempt to be unpredictable to the attacker. This effect is particularly relevant in early interactions. Initial intentional random behavior and the following transition toward more predictable behavior may also be seen as a smooth learning exploratory strategy that helps people detect some behavioral pattern approaching actual optimal play. This may be related to a known transition from exploration to exploitation in decisions under risk (Abbasi et al., 2016).

In conclusion, our study provides helpful insights regarding the effects of uncertainty and attack behavior on a human defender behavior. Currently, most defense mechanisms inspired on game-theory assume complete information, and make unreasonable assumptions regarding rationality and cognitive plausibility of the defense mechanisms (Abbasi et al., 2015a,b). Our results suggest that humans’ main defense vulnerability lies in their performance against adaptive attackers in conditions of uncertainty. Surprisingly, humans are able to handle random attacking behavior better than they are at handling adaptive attackers, suggesting that common randomized security algorithms might be less effective than adaptive human-inspired strategies for defense.

On a more practical ground, this study suggests the need to provide human defenders with more unambiguous information about possible outcomes. However, it is clear that such precise information can hardly be obtained in many real life security scenarios that strongly rely on uncertainty. To illustrate this, consider again the above situation involving a security analyst in charge of protecting a firm’s network infrastructure that may be hacked by some unknown individual. In this case, the analyst has no way to know beforehand how likely a potential attacker is to counteract a defense action. The analyst simply

ignores the payoffs that can result from tracking an individual that hacks the network: the hacker may indeed be more or less prepared to protect himself/herself from being identified. In order to improve the analyst’s behavior in this case, our study suggests the help of a decision support system, which would provide an estimated deterministic value of the various payoffs for each possible outcome (such expected payoffs could simply be determined based on statistical data). This way, the human analyst would perceive the situation as if its description were fully known and would be more efficient at defending the firm’s network. This study therefore suggests the need for increasing efforts in designing more and more efficient decision support systems that take into account the complexity of human behavior in such complex situations.

This work is only a first step toward a more general cybersecurity science. An obvious direction for future research relates to investigating behavior of trained analysts and hackers (e.g., experts) in more concrete security scenarios. Such analyses would undoubtedly help us uncover more sophisticated ways to help people protect themselves in the highly uncertain cyber world.

AUTHOR CONTRIBUTIONS

FM contributed to the design of the study, the implementation of experimental protocols, data collection, data analyses, and write up of the manuscript. CG developed the idea of the study, contributed to the design and implementation of the study, and contributed to the writing of the manuscript.

FUNDING

This research is supported by the Army Research Laboratory under Cooperative Agreement Number W911NF-13-2-0045

(ARL Cyber Security CRA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <http://journal.frontiersin.org/article/10.3389/fpsyg.2017.00982/full#supplementary-material>

REFERENCES

- Abbasi, Y. D., Short, M., Sinha, A., Sintov, N., Zhang, C., and Tambe, M. (2015a). "Human adversaries in opportunistic crime security games: evaluating competing bounded rationality models" in *Proceedings of the Third Annual Conference on Advances in Cognitive Systems ACS* (Atlanta, GA).
- Abbasi, Y. D., Sinha, A., Short, M., Sintov, N., Zhang, C. H., and Tambe, M. (2015b). "Human adversaries in opportunistic crime security games: how past success (or failure) affects future behavior," in *Workshop on Behavioral, Economic and Computational Intelligence for Security Held at the 24th International Joint Conference on Artificial Intelligence (IJCAI)* (Buenos Aires).
- Abbasi, Y. D., Ben-Asher, N., Gonzalez, C., Kar, D., Morrison, D., Sintov, N., et al. (2016). "Know your adversary: insights for a better adversarial behavioral model," in *38th Annual Meeting of the Cognitive Science Society (CogSci 2016)* (Philadelphia, PA).
- Alpcan, T., and Başar, T. (2010). *Network Security: A Decision and Game-Theoretic Approach*. New York, NY: Cambridge University Press.
- Bereby-Meyer, Y., and Roth, A. E. (2006). The speed of learning in noisy games: partial reinforcement and the sustainability of cooperation. *Am. Econ. Rev.* 96, 1029–1042. doi: 10.1257/000282806779468562
- Bier, V. M., and Azaiez, M. N. (eds.). (2008). *Game Theoretic Risk Analysis of Security Threats*. New York, NY: Springer Science & Business Media.
- Brown, G. W. (1951). Iterative solution of games by fictitious play. *Act. Anal. Prod. Alloc.* 13, 374–376.
- Camerer, C. F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Russell Sage Foundation.
- Coricelli, G. (2004). *Strategic Interaction in Iterated Zero-Sum Games*. University of Arizona. Available online at: <http://economics.eller.arizona.edu/downloads/coricelli.pdf>
- Dürsch, P., Kolb, A., Oechsler, J., and Schipper, B. C. (2006). Rage against the machines: how subjects learn to play against computers. *Econ. Theory* 43, 407–430. doi: 10.2139/ssrn.839084
- Dutt, V., Ahn, Y. S., and Gonzalez, C. (2013). Cyber situation awareness modeling detection of cyber attacks with instance-based learning theory. *Hum. Factors* 55, 605–618. doi: 10.1177/0018720812464045
- Evans, D., Nguyen-Tuong, A., and Knight, J. (2011). "Effectiveness of moving target defenses," in *Moving Target Defense: Creating Asymmetric Uncertainty for Cyber Threats, Advances in Information Security 54*, eds S. Jajodia, A. K. Ghosh, V. Swarup, C. Wang, and X. S. Wang (New York, NY: Springer), 29–48. doi: 10.1007/978-1-4614-0977-9_2
- Fang, F., Jiang, A. X., and Tambe, M. (2013). "Optimal patrol strategy for protecting moving targets with multiple mobile resources," in *Proceedings of The International Conference on Autonomous Agents and Multi-Agent Systems* (Saint Paul, MN: International Foundation for Autonomous Agents and Multiagent Systems), 957–964.
- Fox, J. (1972). The learning of strategies in a simple, two-person zero-sum game without saddlepoint. *Behav. Sci.* 17, 300–308. doi: 10.1002/bs.3830170305
- Gonzalez, C. (2013). "From individual decisions from experience to behavioral game theory: lessons for cyber security," in *Moving Target Defense II: Applications of Game Theory and Adversarial Modeling*, eds S. Jajodia, A. K. Ghosh, V. S. Subrahmanian, V. Swarup, C. Wang, and X. S. Wang (New York, NY: Springer), 73–86.
- Gonzalez, C., Ben-Asher, N., Martin, J. M., and Dutt, V. (2015). A cognitive model of dynamic cooperation with varied interdependency information. *Cogn. Sci.* 39, 457–495. doi: 10.1111/cogs.12170
- Hau, R., Pleskac, T. J., Kiefer, J., and Hertwig, R. (2008). The description–experience gap in risky choice: the role of sample size and experienced probabilities. *J. Behav. Decis. Mak.* 21, 493–518. doi: 10.1002/bdm.598
- Hertwig, R., Barron, G., Weber, E. U., and Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychol. Sci.* 15, 534–539. doi: 10.1111/j.0956-7976.2004.00715.x
- Hertwig, R., and Erev, I. (2009). The description–experience gap in risky choice. *Trends Cogn. Sci.* 13, 517–523. doi: 10.1016/j.tics.2009.09.004
- Jain, M., Tsai, J., Pita, J., Kiekintveld, C., Rathi, S., Tambe, M., et al. (2010). Software assistants for randomized patrol planning for the LAX airport police and the federal air marshal service. *Interfaces* 40, 267–290. doi: 10.1287/inte.1100.0505
- Lieberman, B. (1962). "Experimental studies of conflict in some two-person and three-person games," in *Mathematical Methods in Small Group Processes*, eds J. Criswell, H. Solomon, and P. Suppes (Stanford, CA: Stanford University Press), 203–220.
- Martin, J. M., Gonzalez, C., Juvina, I., and Lebiere, C. (2014). A description–experience gap in social interactions: information about interdependence and its effects on cooperation. *J. Behav. Decis. Making* 27, 349–362. doi: 10.1002/bdm.1810
- Messick, D. M. (1967). Interdependent decision strategies in zero-sum games: a computer-controlled study. *Behav. Sci.* 12, 33–48. doi: 10.1002/bs.3830120105
- Nguyen, T. H., Kar, D., Brown, M., Sinha, A., Jiang, A. X., and Tambe, M. (2016). "Towards a science of security games," in *Mathematical Sciences with Multidisciplinary Applications: In Honor of Professor Christiane Rousseau. And In Recognition of the Mathematics for Planet Earth Initiative*, ed B. Toni (Cham: Springer International Publishing), 347–381.
- Pita, J., Jain, M., Tambe, M., Ordóñez, F., and Kraus, S. (2010). Robust solutions to Stackelberg games: addressing bounded rationality and limited observations in human cognition. *Artif. Intell.* 174, 1142–1171. doi: 10.1016/j.artint.2010.07.002
- Pita, J., John, R., Maheswaran, R., Tambe, M., Yang, R., and Kraus, S. (2012). "A robust approach to addressing human adversaries in security games," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems* (Valencia: International Foundation for Autonomous Agents and Multiagent Systems), 1297–1298.
- Rapoport, A., and Budescu, D. V. (1997). Randomization in individual choice behavior. *Psychol. Rev.* 104, 603–617. doi: 10.1037/0033-295X.104.3.603
- Robinson, J. (1951). An iterative method of solving a game. *Ann. Math.* 54, 296–301. doi: 10.2307/1969530
- Roy, S., Ellis, C., Shiva, S., Dasgupta, D., Shandilya, V., and Wu, Q. (2010). "A survey of game theory as applied to network security," in *2010 43rd Hawaii International Conference on System Sciences (HICSS)* (Honolulu, HI).
- Shachat, J., and Swarthout, J. T. (2004). Do we detect and exploit mixed strategy play by opponents? *Math. Methods Oper. Res.* 59, 359–373. doi: 10.1007/s001860400354
- Shachat, J., and Swarthout, J. T. (2012). Learning about learning in games through experimental control of strategic interdependence. *J. Econ. Dyn. Control* 36, 383–402. doi: 10.1016/j.jedc.2011.09.007
- Shieh, E. A., An, B., Yang, R., Tambe, M., Baldwin, C., DiRenzo, J., et al. (2012). "PROTECT: an application of computational game theory for the security of the ports of the United States," in *AAAI* (Toronto, ON).
- Spiliopoulos, L. (2008). *Humans versus Computer Algorithms in Repeated Mixed Strategy Games*. Technical Report 6672, Munich Personal RePEc Archive.

- Tambe, M. (2011). *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. New York, NY: Cambridge University Press.
- Yang, R., Kiekintveld, C., Ordóñez, F., Tambe, M., and John, R. (2011). "Improving resource allocation strategy against human adversaries in security games," in *IJCAI Proceedings-International Joint Conference on Artificial Intelligence* (Barcelona).
- Zhuang, R., Zhang, S., Bardas, A., DeLoach, S. A., Ou, X., and Singhal, A. (2013). "Investigating the application of moving target defenses to network security," in *2013 6th International Symposium on Proceeding of the Resilient Control Systems (ISRC)* (San Francisco, CA).

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2017 Moisan and Gonzalez. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.