Cyber Security via Minority Games with Epistatic Signaling

(Invited Paper)

William Casey
Software Engineering Institute, Carnegie Mellon University

Rhiannon Weaver
Software Engineering Institute, Carnegie Mellon University

Leigh Metcalf
Software Engineering Institute, Carnegie Mellon University

Jose Andre Morales
Software Engineering Institute, Carnegie Mellon University

Evan Wright
Software Engineering Institute, Carnegie Mellon University

Bud Mishra
Courant Institute, New York University

ABSTRACT

We present a game theoretic framework, modeling strategic interactions among humans and things, which are assumed to be interconnected by a social-technological network, as in Internet of Humans and Things (IOHT). Often a pair of agents in the network interacts in order for an informed sender-agent to signal an uninformed receiver-agent to take an action that benefits each of the players – the benefits to the pair of agents are modeled by two separate utility functions, both depending on the sender's private information, the signal exchanged, and the receiver's revealed (and unrevealed) action. In general, the two agents' utilities may not be aligned and may encourage *deceptive behavior*. For example, a sender, aware of his own private "state of ignorance," may seek useful information from a receiver who owns powerful computational resources to search a large corpora of webpages; the sender does so by sending a signal to the receiver in the-form of a keyword. Obvious examples of deceptiveness here range from attempts to hide one's intentions to auctioning the keywords on an ad exchange through real-time bidding. A rather troublesome situation occurs when deceptions are employed to breach the security of the system, thus making the entire social-technological network unreliable. Earlier, we proposed a signaling-game-theoretic framework to alleviate this problem. This paper further enhances it by reconfiguring signals to possess more complex structures (epistatic signals to represent attack and defense options over a given set of vulnerabilities). We explore two augmentations to the original evolutionary signaling game by *first* enhancing mutation bias toward strategies performing well in previous populations and *secondly* by allowing the parameters of the utility functions to dependent on population preferences giving rise to a minority game with epistatic signaling. The resulting game systems are empirically studied through extensive computer simulation.

BICT 2014, December 01-03, Boston, United States

DOI 10.4108/icst.bict.2014.257967

1. GAMES AND CYBER CONFLICTS

At the core of many dynamic online strategic interactions are simple *information-asymmetric games*, which do permit the agents to act *deceptively* to gain advantages. Take for example the flashlight app for smartphones which was also discovered to open a GPS-tracking backdoor to gain private information by logging the device's physical locations (discovery reported in [6]). Whereas the producer (i.e., sender) of the flashlight app may advertise (i.e., signal) that the application is designed to provide a flashlight feature for smartphones, the sender creates the deceptive impression of respecting the user's privacy by giving the app a benign sounding name: "Flashlight App." A typical user's expectations of privacy would proscribe the surveillance capabilities (physically tracking the user's device via GPS tracking) and not foresee encroachment by an app that is *prima facie* simple, benign, and desirable. In this case (and others like it) a typical consumer (receiver) may recognize that they had been deceived, and may label the producer (sender) as a miscreant and tarnish the producer's reputation with a negative ranking and comments labeling the app as "backdoor," "Trojan," or "malware." Such verification processes are aimed at protecting the future consumers. However, the encounter, concluded before the discovery of the attack, has its costs and benefits: the cost to the receiver is the loss of privacy, and the benefit to the sender is the ability to gain strategic informational advantages with unanticipated usages.

In considering signaling games for cyber security to model interactions, such as the one above, we envision that security properties such as *non-surveillance* can be checked, dynamically and efficiently, via two additional mechanisms: namely, (i) a social-technological recommendation-verification system involving additional players, and (ii) a currency system, represented by M-Coins certificates backing the proofs concerning the behavior of the agents.

We also extend the receiver's strategy space by providing it means to *challenge* the sender. Note that, without proof or certification that the app's behavior complies with reasonable security properties, the receiver is left with the options to either trust the sender or attempt to challenge them. Such challenges may seek their own or otherwise trusted proofs or certificates to let the receiver decide whether the sender is being deceptive.

Motivated by biological systems, we provide another extension to allow our recommendation-verification system to address the many distinct attacks that a producer (sender) could use to deceptively ensnare a consumer (receiver). Here, we describe this extension of signaling games which include diverse attack vectors and we term this extension *epistatic signaling games*. After defining epistatic signaling games we present experiments designed to understand their dynamics empirically and how such a system could operate in practice. We relegate a formal description of the system and proofs of its various properties to the full paper.

1.1 Signaling Games in Cyber Security

A *signaling game* is a dynamic game with two players, the sender (S) and the receiver (R). The sender is assumed to possess a certain type, $t \in \mathcal{T}$, which is selected by nature (we will think of it as sender's private information; thus, the sender observes his own type while the receiver does not know the type of the sender). The sender chooses a message α from a set of possible messages $\mathcal{M}$: in epistatic signaling games the message may contain attacks upon any subset of K distinct vulnerabilities[1], denoted $V = \{v_1, \ldots, v_K\}$, including the empty set which we term a *clean* or benign signal. The receiver observes the message but not the type of the sender or the attacks implicitly encoded in the message. Then the receiver chooses an action γ from a set of feasible actions $\mathcal{C}$, which include challenges to various attacks: Letting c_i be the check for an attack on the ith vulnerability v_i the sender's options are subsets of C (also denoted by the vulnerabilities, v_i's, $i = 1, \ldots, K$, being challenged), with the empty set representing the option of receiving messages with no challenge at all (either trusting or an insouciant option). The utility functions are given by $U_S : \mathcal{T} \times \mathcal{M} \times \mathcal{C} \to \mathbb{R}^+$ for the sender and $U_R : \mathcal{T} \times \mathcal{M} \times \mathcal{C} \to \mathbb{R}^+$ for the receiver. In the context of cyber security, we always consider the symmetric game with repetitions (as opposed to one-shot), in which both players play both roles. Basic signaling games were introduced by In-Koo Cho and David M. Kreps in a 1987 article[5].

By considering a singular attack option and singular checking action we explore the effects of deceptive agents in cyber security problems via simulations; these simulations reveal a range of outcomes for system behavior over the space of payoff parameters([3]).

Epistatic signaling games differ from signaling games originally introduced by us for cyber security in the following two ways. First, in signaling games the strategic options for sender and receiver are limited to a single attack and a single challenge option; a signaling game is a special case of the general epistatic signaling formulation when $K = 1$. Higher, but bounded, values of $K > 1$ add realism to the model by constructing the attack surface to be K vulnerable objects. The second way in which this approach differs from traditional signaling games is that we simplify the transitions in strategies for repeated games. In this approach we are limiting the agents to two transitions based on whether or not a detection event occurred. While this constraint may appear to be limiting, it is more realistic, since agents are primarily interested resolving an attack (i.e., detection event); note particularly that in the case of an undetected attack, the user will not have immediate access to what attack succeeded.

[1]Vulnerabilities may be considered code objects which can be exploited and attacked by a malicious user.

We briefly review the strategic options and payoff of signaling games for cyber security to fully demonstrate the relation between signaling games, signaling games introduced previously for cyber security and this approach of epistatic signaling games.

The strategic options: In signaling games (when $K = 1$) the sender may select to send cooperatively C or to send an attack D. Similarly the options for the receiver are to accept trusting C or to challenge D. We encode all options of the symmetric game using strings where the first letter denotes the sender's type and the second the receiver's action. Using this encoding the option space for a single round of symmetric signaling games is the set $\{CC, CD, DC, DD\}$.

Game Payoff: The payoff matrix for the symmetric signaling game (with $K = 1$) is then defined over the product of row-player options and column-player options. In this matrix, d is the benefit of an attack for the sender (assumed to be a zero-sum quantity), e is the cost of getting caught attacking as sender, f is the prize for catching an attacker, and g is the cost of challenging a sender as receiver. The payoff matrix is:

row / *col*	CC	CD	DC	DD
CC	0 0	0 $-g$	$-d$ d	$-d$ $d-g$
CD	$-g$ 0	$-g$ $-g$	$f-g$ $-e$	$f-g$ $-e-g$
DC	d $-d$	$-e$ $f-g$	0 0	$-d-e$ $d+f-g$
DD	$d-g$ $-d$	$-e-g$ $f-g$	$d+f-g$ $-d-e$	$-e+f-g$ $-e+f-g$

2. MINORITY GAMES WITH EPISTATIC SIGNALING

2.1 Epistatic Signaling

Central to epistatic signaling games are the outcomes of receiver challenges against sender attacks, which may result in detection events (or otherwise). In epistatic signaling games there are K vulnerabilities denoted as a set V. A sender may exploit some subset of vulnerabilities as an option. Likewise the receiver may check (if the sender's message exploits) any subset of vulnerabilities as an option. We denote the sender and receiver options as α and γ and note that they are both subsets of V. A sender who sends a benign signal is encoded by $\alpha = \emptyset$. A receiver who offers no challenges may also be encoded as $\gamma = \emptyset$. In principle both the sender and receiver have 2^K options.

Letting α, γ be the options employed by sender and receiver during an interaction (message), there are four possible outcomes concerning each vulnerability v:

Attack detection: When $v \in \alpha \cap \gamma$ the attacker attempts to exploit vulnerability v at a cost of H, the cost of attacking a vulnerability. The receiver incurs a challenge cost of G; however, in this case the result is a detection. The detection

imparts a heavy reputational cost on the sender which we term E, the cost of getting caught. Within a social network the detection may also confer to the receiver a reward of F, the benefit for catching an attacker. Furthermore, the benefit for catching an attacker is higher when the challenging receiver is in a *minority*, as she shares the benefit with few others.

Futile defense: When $v \in \gamma \setminus \alpha$ the receiver checks vulnerability v at a cost of G. The challenge is unnecessary as the sender employs no deceptive attacks on v.

Effective attack: When $v \in \alpha \setminus \gamma$ the attacker exploits vulnerability v which goes undetected by the receiver. In this case the receiver avoids the cost of checking G, but will incur the cost of being attacked D while the attacker is rewarded D at the cost of H.

Benign sender and trusting receiver: When $v \notin \alpha \cup \gamma$ the sender and receiver incur no costs nor receive any rewards stemming from either deceptions or challenges.

We organize the outcomes (across the entire vulnerability space) into two revealing outcomes to be used for the control mechanisms of agent strategies in repeated games. These outcomes are: *detection* when an attack detection occurs for any $v \in V$. (i.e., $|\alpha \cap \gamma| > 0$), and *otherwise* when there are no attacks detected for any vulnerability. (i.e., $|\alpha \cap \gamma| = 0$).

2.2 Epistatic Signaling Games

We next examine the utility functions in the epistatic signaling game, as an extension of signaling games. We begin by discussing strategic options and game payoffs; to assist in computing payoffs, we introduce auxiliary counting functions. We organize the symmetric game into two stages: the play is in offense when the agent is a sender facing a potentially challenging receiver and in defense when the agent is a receiver facing a possibly deceptive sender. Finally we present the payoff function.

The strategic options: By considering a larger signal space, the sender and receiver will have vastly more options for strategic selection. The options for sender (offense) include every element of the power set 2^V as do the options for receiver (defense).

In the symmetric game form the agent must select one option for offense and one for defense. We let $W = 2^V \times 2^V$ comprise the full set of strategic options for an agent in symmetric epistatic signaling games.

2.2.1 Payoff Structure for Epistatic Signaling Games

We consider a special form of the payoff matrix for the epistatic signaling game here, which may be formalized as an assignment of payoff (for the row-player i against column-player j) over the product space of signals $W \times W$. Letting $w_i = (\alpha_i, \gamma_i) \in W$ and $w_j = (\alpha_j, \gamma_j) \in W$ be the strategic option for the row-player and the column-player respectively, we may reference the offense options as α_i, α_j and the defensive options as γ_i, γ_j employed by the row-player and column-player respectively.

The payoff matrix for epistatic signaling games will be denoted as $M(w_i, w_j)$ to quantify the payoff (for the row-player only) when the row-player i employs option w_i and column-player j employs option w_j.

Payoff evaluation: To compute $M(u_i, u_j)$ we organize a single round of play into stages and introduce a few simple auxiliary functions; the stages are offense (when the row-player is a sender) and defense (when the row-player is a receiver).

Offense: In a stage of play the row-player i launches a total number of attacks against the column-player j counted as $\textsc{Attacks-Fielded}(i,j) = |\alpha_i|$, while the number of successful attacks by the row-player i against the column-player j is counted as $\textsc{Attacks-Achieved}(i,j) = |\alpha_i \setminus \gamma_j|$. For each attack launched by the sender a fixed cost H is added to the overall cost of the sender option. This fixed cost may be associated with the cost to develop/deploy an attack, identify a software vulnerability, develop an exploit, or apply resources to attack. For each attack achieved by the row-player i against the column-player j a fixed zero-sum equity of D is transferred to the row-player as a benefit at the expense of the column-player. This zero-sum equity is intended to model the value of a digital asset, authorization token, credential, personally identifiable information, digital currency (e.g., bitcoin or more specifically, M-Coin), etc.

Defense: In a stage of play the row-player i fields a total number of defenses (or checks) against the column-player j, denoted as $\textsc{Defenses-Fielded}(i,j) = |\gamma_i|$, while the number of effective defenses (or detection events) for the row-player i against column-player j is counted as $\textsc{Detects}(i,j) = |\gamma_i \cap \alpha_j|$; the futile challenges for player i against player j are counted as $\textsc{Futile-Challenge}(i,j) = |\gamma_i \setminus \alpha_j|$. For each defense fielded by the receiver a fixed cost G is applied to the strategic option; this cost can be treated as a cost to develop the detector algorithms and may be amortized and scaled to affordable quantities via a social-technical network where detection methods are deployed. Each detection event imposes a heavy cost of E on the sender and will also confer a benefit of F to the receiver. The cost associated with a detection event for the sender is designed to model the loss of reputation, loss of security certifications, M-Coin tokens, etc. As an example, a code project that imparts users with a large vulnerability surface[2] will naturally suffer a reputational loss as multiple receivers prove its deficiencies. Defenses that are fielded but do not result in detections may be considered futile (at least for that round) and will carry a cost burden for the receiver. This cost imposes a natural pressure on agents to be parsimonious with detection and establishes an incentive to measure the effectiveness of receiver options so that the most effective methods for detection can be selected and propagated in a population.

For row-player i selecting option $w_i = \alpha_i \times \gamma_i$ playing against column-player j who selects option $w_j = \alpha_j \times \gamma_j$, the row-

[2]Vulnerabilities may result from technical deficits such as sloppy code writing and leave a user exposed to an attacker.

player payoff is defined as follows:

$$\begin{aligned} M(w_i, w_j) \;=\; & D \cdot \textsc{Attacks-Achieved}(i,j) \\ & - D \cdot \textsc{Attacks-Achieved}(j,i) \\ & + F \cdot \textsc{Detects}(i,j) \\ & - E \cdot \textsc{Detects}(j,i) \\ & - H \cdot \textsc{Attacks-Fielded}(i,j) \\ & - G \cdot \textsc{Defenses-Fielded}(j,i). \end{aligned}$$

Note that the settings of parameters D, E, F, G, H are shown to be critically important for the behavior of a system for evolving populations in [3]. The important distinction for this model (epistatic) is that costs/benefits are allowed to scale to the expanded signal space, introducing a more realistic setting for cyber security.

These scale laws naturally place incentives to select effective options. They also provide a means to study various interesting system behavioral outcomes such as system effects for various rates of evolution in attacks vs. defenses. Our motivation for studying this problem is rooted in the following practical questions: whether a social-technological recommendation-verification system can be effective in providing defenses flexibly, and if so, what mechanisms can achieve this desideratum?

2.2.2 *Strategy for Repeated Games*

In the repeated game form the encountering agents will play a sequence of rounds, each round a single symmetric epistatic signaling game. Because interactions among agents in IOHT are strongly dependent on prior interactions, a strategy should address how agents may adapt prior revealed outcomes.

A strategy for repeated games may include a control mechanism which incorporates previously revealed outcomes as conditional information for an agents' updated strategy. To incorporate such a control mechanism (based on attack detection) into strategy for repeated games, we model each agent's strategy as a labeled deterministic finite state automata (DFA), just as we did in our earlier work. Labeled DFA provides a means to evolve complex strategic interactions spanning multiple plays of a repeated game. This technique (used in [2, 8]) enhances the dynamics that are possible, while mutation provides a means for exploration (of a vast strategic space), thus allowing an ensemble of agents to adapt strategies to population-dependent fitness-landscapes.

Below in Figure 1, we show how strategy structures can evolve, and in Figure 2, we illustrate how a mutational process for strategies generates diverse strategies over time.

2.3 Strategic vs Oblivious Mutation

In our original signaling games the mutation process places fixed probabilities on mutation types. While such a process can be shown to make all labeled DFAs reachable its fixed nature throughout the simulation of a population's history ignores a realistic aspect of the social-technological system we are hoping to construct. We foresee that within a social-technological system strategic agents would naturally bias the mutations of sender/receiver options toward a performance measure (rather than oblivious mutation via a fixed process). These considerations led us to incorporate performance biased mutation processes within the context of a social-technological system [4].

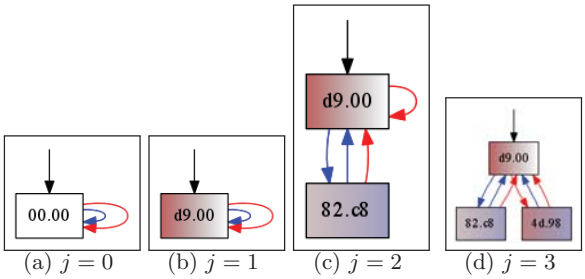

(a) $j = 0$ (b) $j = 1$ (c) $j = 2$ (d) $j = 3$

Figure 1: Agent strategies for epistatic signaling games are succinctly represented as a deterministic finite state automaton (DFA), which evolves over time via a mutation process. As an example we show a sequence of four mutations with $K = 8$. Each sending (attack) and receiver (defense) option is denoted by a number in hexadecimal notation which represents the subset of V employed. Starting with (a), the initial seed strategy employs no attacks and no defenses; the label 00.00 represents the selected attacks (two hexadecimal digits to the left) and selected defenses (two hexadecimal digits to the right). The transitions (edges) in the DFA instruct the agent on which state to transition to next based on revealed outcomes: when the strategy detects an attack the agent will use the red transition and blue otherwise. Next, in (b), the sending signal is modified from 00 to $d9$, which encodes the new attack option exploiting vulnerabilities $\{1, 4, 5, 7, 8\}$. To illustrate the number of attacks and checks we use a gradient coloration from red on the left to blue on the right which indicates the density of attacks (intensity of red on the left) and defenses (intensity of blue on the right). Next, (c) presents an added state labeled $82.c8$ which encodes the sender option exploiting vulnerabilities $\{2, 8\}$ and receiver option which checks vulnerabilities $\{4, 7, 8\}$. Last, in (d), a state labeled $4d.98$ is added that represents attack option exploiting vulnerabilities $\{1, 3, 4, 7\}$ and defense option checking vulnerabilities $\{4, 5, 8\}$. In this mutation sequence, options for a newly created state are selected uniformly random over the option space.

2.4 Minority Signaling Games

In signaling games played in social-technological systems, we may consider the possibility of variable costs/payoffs depending on bulk population behavior. In this context, there will be certain advantages (e.g., reputational gain) by being in the minority as a challenging receiver. These considerations led us to formulate minority signaling games. If early adapters (minorities) have slight preferential advantage there may also be incentives for the population to develop and maintain diverse challenging options. It may also be possible that a population that develops and sustains diversity in strategies may mitigate some of the most wild

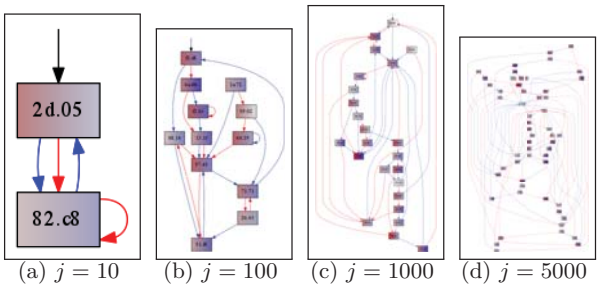

(a) $j = 10$ (b) $j = 100$ (c) $j = 1000$ (d) $j = 5000$

Figure 2: Mutation of strategy creates diverse strategies with complex transitions. Continuing the mutation sequence from Figure 1, we illustrate the strategies explored after (a) 10, (b) 100, (c) 1000, and (d) 5000 mutations. Transitions for repeated games (e.g., the edges) are also mutable and correspond to the two revealing outcomes: detection (colored red) and otherwise (colored blue).

dynamics observed in signaling games, which include oscillation between low to high levels of attacks and checking (all players deciding to challenge or all players deciding to be indifferent).

To study this problem we consider the effects of non-constant coefficients in the payoff structure, a mechanism that gives rise to dynamics similar to the El Farol bar problem [1]. To introduce El Farol bar dynamics into the epistatic signaling games, we consider allowing the cost parameter G to vary based on bulk population behavior; the simplest adjustment is a step function which increases the cost (by a multiple ζ) when the fraction of outcomes in a population exceeds a given fractional threshold τ. We define the set of agents as $U = \{u_1, u_2, \ldots, u_M\}$ and consider all the games occurring during encounters in a given generation. Summing over all encounters during a generation we let $\mathcal{X}$ be a monitor for the fractional amount of checks deployed among all defensive options compared to the total possible capacity for checking during the generation (i.e., if all receiver options employed every check).

In minority signaling games the general form of the payoff for row-player is a slight modification to the equation for $M(u_i, u_j)$ where the coefficient G is modified to be a step function depending on the population quantity $\mathcal{X}$ computed during the games of a generation:

$$G(U) = \begin{cases} G & \text{if } \mathcal{X}(U) \leq \tau \\ \zeta \times G & \text{otherwise.} \end{cases}$$

3. SIMULATIONS

Motivating our investigation are the questions:

- What are the effect of strong and transparent measures for the challenge options in a population vs. random selection? This experiment seeks to compare the system behavior in each of the following two cases:
  - Receiver challenge options are selected uniformly randomly over the receiver option space (when mutation events occur).
  - Receiver challenge options are selected based on performance measures proven in the previous generation of games (when mutation events occur). Some fraction ξ of mutations that will affect receiver options will be selected uniformly randomly over the entire receiver option space.
- What are the effects of minority games and El Farol dynamics when applied as a step function for sender costs? This experiment introduces the population-behavior-based step function $G(U)$ already defined with fractional behavior quantity $\mathcal{X}$ and threshold τ. It also explores if this mechanism can diversify sender options in a population and lead to effects on system dynamics.

The investigation of these questions are meaningful for prospective engineering of cyber security in social-technical networks. In the first experiment, we investigate the effects of strong and transparent measures for challenge options as a means to organize a distributed cyber response within a epistatic signaling game, but also related to other notions of cellular immune response systems ([7]).

In the first experiment the fraction ξ will have effects on fixation probabilities because it imposes differential rates of exploration for receiving and sending options. In addition ξ will have effects on the ability for receiver options to adapt defense to novel attack strategies as it makes previously high performing strategies more persistent. While the effects of mutation rates and ξ are of practical interest a more general theoretical framework remains to be constructed (see the full paper).

The second experiment addresses some of the wild dynamics observed in these systems, which include defection invasions and spontaneous cooperation as well as wild oscillation between them. The experiment is designed to investigate the possible effects of a mechanism, which may incentivize the parsimonious use of and the diversification of defense options and may increase stability in these complex dynamics. Such a mechanism may either be designed as part of a system or may be discovered as a natural factor.

Below the results obtained from the experimentation are reported in images as well as exposition of what this may mean for security in social-technical systems.

3.1 Simulation Outline

We outline the general simulation and provide descriptions of how we can augment or modify each step to achieve the analytic steps outlined above.

Shape Parameters: $\langle M, K, N \rangle$: population size, option set size, and number of generations.

System Parameters: $\langle D, E, F, G, H, \delta, \mu, \rangle$: payoff settings, continuation factor, mutation rate.

Initialize: A population U of M users initialized with random strategies.

For each generation:

- **Encounter:** Using the population of strategies (time n) we create pairwise encounters for game play.
- **Play:** For each encounter: repeated games are played using agent strategies. Number of rounds determined by continuation parameter δ. Each player aggregates a vector of outcomes.
- **Aggregate and Evaluate Scores**: Total performance measures are aggregated across strategies and unique options used during the encounters for generation n. Scores and measures are computed using epistatic signaling game payoff matrix, outcome vectors resulting from play, and system parameters.
- **Re-create**: A population of M strategies is recreated (for next generation n+1) by sampling the existing strategies with probability density proportional to performance scores.
- **Mutate**: Players are chosen with rate μ for mutation. Each mutation event may modify the strategic encoding of strategy.

The encounters may be created in a variety of ways, including random pairing, use of an underlying neighborhood graph to describe kinship or geographical relations, or various hybrid notions. In ([8]) population structure parameters α, δ were introduced to study a mixture of encounters ranging from random to structured encounters. In our experiments we use $\frac{M}{2}$ encounters selected as random encounters only. During the play the continuation parameter δ is used to determine number of rounds by generating a random geometric derivate with δ as continuation parameter. For pairwise agent encounters in repeated games each agent will use their strategy (described by a labeled DFA), which is used to compute options and outcomes for each round of play in during the repeated epistatic signaling games. The labeled DFAs are used in the following way (described for the row-player): Starting from the start-state the *sending* and *receiving* signals are determined. If the row-player detects an attack from the column-player then the red transition edge is used to determine the next strategic options for both sending and receiving. If an attack was not detected then the blue transition edge is used to determine the next strategic option for the row-player. In either case, in the next round, the option including both send and receive is determined. By following this sequence of steps in the strategic automata each agent may aggregate a vector of outcomes (e.g., number of attacks, number of defenses, number of detections, number of times opponent detects their attacks). These aggregate counts are stored for the next step where the strategies are scored.

Mutation of strategy is performed on the generation of M strategies with base rate μ with an expected number of mutants as μM per generation. Given that a strategy is selected for mutation, one of the five mutation types is selected according to a mutational type frequency vector, which throughout the experiments will be fixed at $\nu = [0.15, 0.15, 0.1, 0.3, 0.3]$. Next we describe the mutational types:

- *type-i*: mutate the sender option.
- *type-ii*: mutate the receiver option. The selection distribution is the subject of experiment titled: Effect of strong and transparent measures.
- *type-iii*: mutate an edge (selected uniformly randomly in all experiments).
- *type-iv*: create a new strategy state with randomly selected edges. (Through out these experiments we limit the size of automata to 256).
- *type-v*: remove a strategy state. (Throughout these experiments we limit the size of automata to be one or more states).

In the first experiment, where we investigate the effects of strong and transparent measures on receiver options, we also track the number of times each receiver option detects an attack. When a mutation event modifies a receiver option we replace the option with a random selection with probability $1 - \xi$, and with probability ξ we use a performance-scaled density over the defense options at play in generation n. The first outcome (with probability $1-\xi$) represents the oblivious choice of a fixed mutation process while the second outcome should allow the population to track existing attack vectors within the population more effectively.

In the second experiment, where we investigate the El Farol dynamics, we augment the aggregate-and-evaluate step to compute $\mathcal{X}$ and update cost per defense using function $G(U_n)$ for generation n. This allows us to draw some conclusions about the use of such a mechanism in epistatic signaling games.

3.2 Experimental Results

Using shape parameters $M = 320, K = 8, N = 80,000$ with system parameters $D = 10, E = -100, F = 4, G = -2, H = -2, \mu = 0.03, \delta = 0.5$ and letting our encounter mechanism be random pairs $\alpha = 0.0$, we conduct experiments by generating 100 histories of simulations of the following systems. Throughout the mutation rates remain fixed at $\nu = [0.15, 0.15, 0.1, 0.3, 0.3]$.

- **S1** : Epistatic signaling games with receiver options mutated uniformly randomly over the option space. (Single history illustrated in Figure 3)
- **S2** : Epistatic signaling games with receiver options scored as a strong and transparent measure in the population $\xi = 0.5$. (Single history illustrated in Figure 4)
- **S3** : Epistatic signaling games with minority step function $G(U)$ with $\tau = 0.4, \zeta = 4.5$.
- **S4** : Epistatic signaling games with receiver options scored as a strong and transparent measure in the population $\xi = 0.5$, and minority step function $G(U)$ with $\tau = 0.4, \zeta = 4.5$. (Single history illustrated in Figure 5)

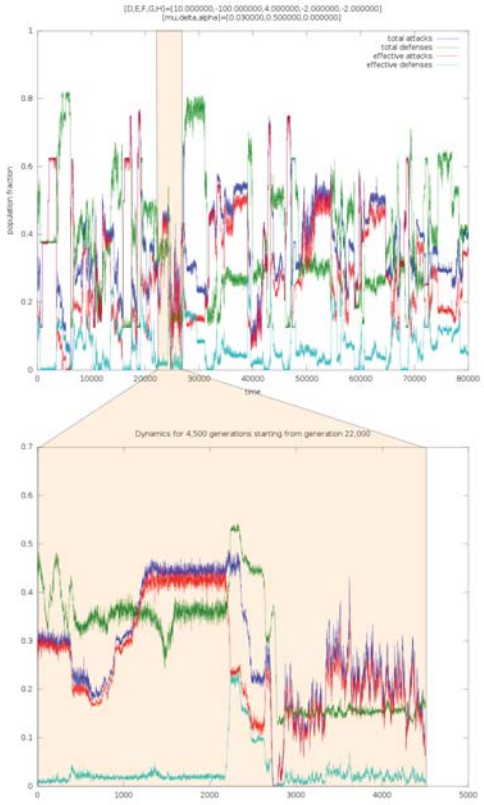

Figure 3: S1: Dynamics of epistatic signaling games. Fractional quantities of attacks (blue), effective attacks (red), defenses (green), and effective defenses (cyan) in 80,000 generations. Below the details are shown in higher resolution for 4,500 generations.

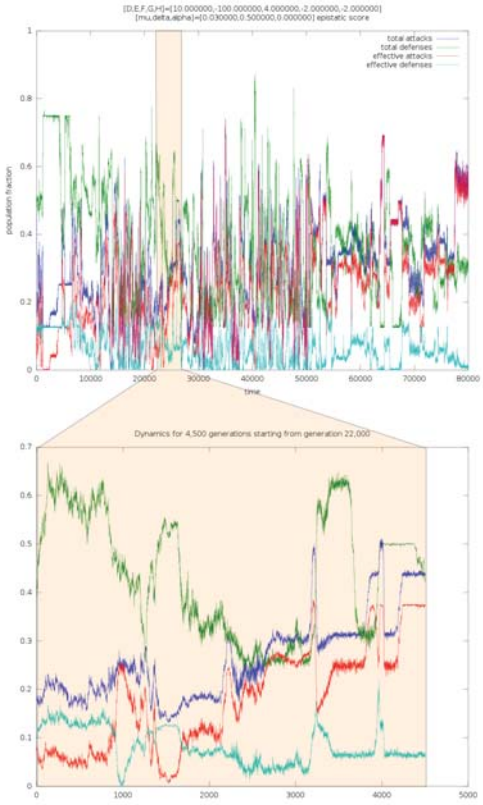

Figure 4: S2: Dynamics of epistatic signaling games when mutation for receiver options is biased toward strong and transparent performance measures proven in previous rounds against employed attacks. Fractional quantities of attacks (blue), effective attacks (red), defenses (green), and effective defenses (cyan) in 80,000 generations. Below the details are shown in higher resolution for 4,500 generations.

3.2.1 Insights from Experiments

In Figure 6 we compare the behavior of each system using the quantities which measure the fraction of all attacks sent as A, the fraction of attacks that are not detected as $[A]$, the fraction of defenses which detect attacks as $[D]$, and the fraction of defenses fielded as D. All fractions are made by comparing the number of observations to the total possible capacity of users if they all expressed every option at every play.

The effect of strong and transparent measures for challenge options does not appear to decrease the number of attacks but seems to reduce the number of defenses fielded, while maintaining an equivalent detection rate. The effect of minority games, which introduce a multiplier cost to G (the cost of fielding defenses), seems to also have an equivalent effect to that of imposing strong and transparent measures on the receiver options. This finding suggests that early adapters who are rewarded preferentially (minority game) and strong transparent measures that bias selection toward strong performance measures may have similar effects on a population. This is somewhat surprising and of interest for engineering cyber security goals as each path may have very different implementations. When strong and transparent measures (performance biased mutation) and early adapter advantage (minority games) act together they appear to have compounding effects. We also observe that as systems gain efficiency in detection, there is also an increase in attack efficacy suggesting that attackers may better exploit the higher dimensional signal space as detectors improve.

4. CONCLUSION

We have shown a natural role for signaling games in modeling various strategic interactions among agents in a social-technological network such as an internet of humans and things. In particular, we have studied the effect of augmenting the two-player sender-receiver games with a recommendation-verification system, further augmented by a newly-devised crypto-coin (e.g., M-Coin). In this paper, we have primarily advanced the design by introducing a complex form of signaling called epistatic signaling and explored the role of minority games in this context. Our simulations have identified some counter-intuitive behaviors (for instance, the behavior of the attackers in exploiting the signal complexity as the

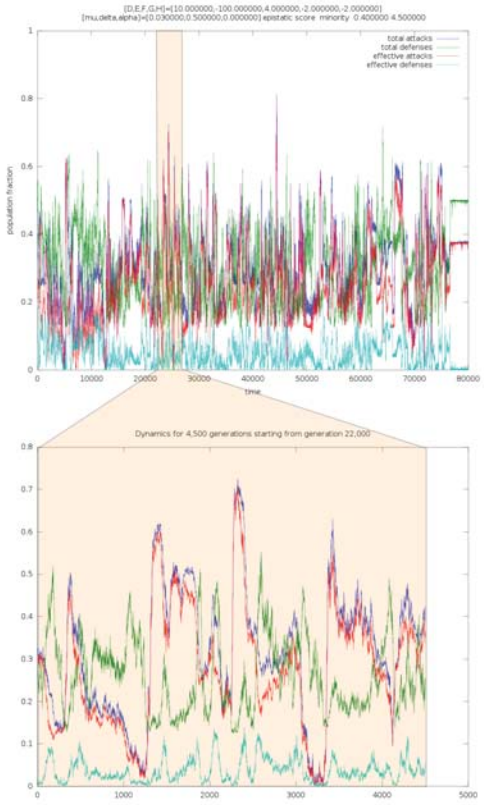

Figure 5: S4: Dynamics of epistatic signaling games when the mutation for receiver options is biased toward strong and transparent performance measures and a minority step function $G(U)$ is used to determine the cost of each defense. Fractional quantities of attacks (blue), effective attacks (red), defenses (green), and effective defenses (cyan) in 80,000 generations. Below the details are shown in higher resolution for 4,500 generations.

dimension of the attack and checking vectors grew). In our future work, we plan to explore the natural trade-offs that exist between complexity of signals and levels of deception.

Acknowledgments. This work is funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. We wish also to thank Thomson Nguyen and Jonathan Spring for engaging discussions.

References

[1] W. B. Arthur. Inductive reasoning and bounded rationality. *The American economic review*, pages 406–411, 1994.

[2] K. G. Binmore and L. Samuelson. Evolutionary Stability in Repeated Games Played by Finite Automata. *Journal of Economic Theory*, pages 278–305, 1992.

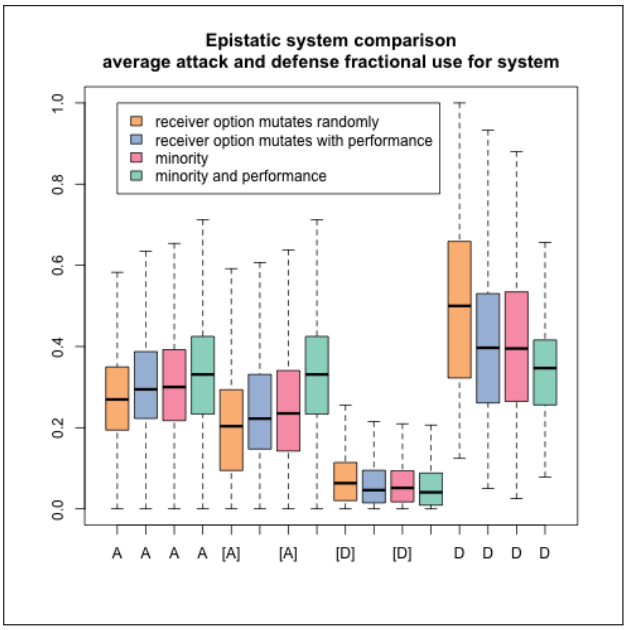

Figure 6: Experiment to uncover the dynamics of epistatic signaling games in differing systems: system S1 (orange), system S2 (blue), system S3 (pink), and system S4 (green). We show the quartile statistics (taken over all histories) of behavioral quantities: the fraction of attacks sent A, the fraction of effective attacks $[A]$, the fraction of detected attacks $[D]$, and the fraction of defenses fielded D. All fractional quantities are computed by dividing the observed over the total possible capacity of users to field attacks or defenses.

[3] W. Casey, J. A. Morales, T. Nguyen, J. Spring, R. Weaver, E. Wright, L. Metcalf, and B. Mishra. Cyber security via signaling games: Toward a science of cyber security. In *ICDCIT*, pages 34–42, 2014.

[4] W. Casey, E. Wright, J. A. Morales, M. Appel, J. Gennari, and B. Mishra. Agent-based trace learnng in a recommendation-verification system for cybersecurity. In *9th IEEE International Conference on Malicious and Unwanted Software, MALCON 2014*.

[5] I.-K. Cho and D. M. Kreps. Signaling games and stable equilibria. *The Quarterly Journal of Economics*, pages 179–221, 1987.

[6] M. Kassner. Android flashlight app tracks users via GPS, FTC says hold on, December 2013. [Online; posted December 11, 2013, 9:49 PM PST].

[7] H. A. van den Berg. Design principles of adaptive cellular immunity for artificial immune systems. *Soft Comput.*, 13(11):1073–1080, 2009.

[8] M. van Veelen, J. García, D. G. Rand, and M. A. Nowak. Direct reciprocity in structured populations. *Proceedings of the National Academy of Sciences*, 109(25):9929–9934, June 2012.