



Decision Analysis

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Optimal Off-line Experimentation for Games

Theodore T. Allen, Olivia K. Hernandez, Abdullah Alomair

To cite this article:

Theodore T. Allen, Olivia K. Hernandez, Abdullah Alomair (2020) Optimal Off-line Experimentation for Games. Decision Analysis 17(4):277-298. <https://doi.org/10.1287/deca.2020.0412>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2020, The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Optimal Off-line Experimentation for Games

Theodore T. Allen,^a Olivia K. Hernandez,^a Abdullah Alomair^a

^aIntegrated Systems Engineering, The Ohio State University, Columbus, Ohio 43210

Contact: allen.515@osu.edu,  <https://orcid.org/0000-0002-9522-3252> (TTA); hernandez.127@osu.edu (OKH); alomair.3@osu.edu (AA)

Received: July 2, 2019

Revised: February 7, 2020

Accepted: February 11, 2020

Published Online in Articles in Advance:
May 22, 2020

<https://doi.org/10.1287/deca.2020.0412>

Copyright: © 2020 The Author(s)

Abstract. Many business situations can be called “games” because outcomes depend on multiple decision makers with differing objectives. Yet, in many cases, the payoffs for all combinations of player options are not available, but the ability to experiment off-line is available. For example, war-gaming exercises, test marketing, cyber-range activities, and many types of simulations can all be viewed as off-line gaming-related experimentation. We address the decision problem of planning and analyzing off-line experimentation for games with an initial procedure seeking to minimize the errors in payoff estimates. Then, we provide a sequential algorithm with reduced selections from option combinations that are irrelevant to evaluating candidate Nash, correlated, cumulative prospect theory or other equilibria. We also provide an efficient formula to estimate the chance that given Nash equilibria exists, provide convergence guarantees relating to general equilibria, and provide a stopping criterion called the estimated expected value of perfect off-line information (EEVPOI). The EEVPOI is based on bounded gains in expected utility from further off-line experimentation. An example of using a simulation model to illustrate all the proposed methods is provided based on a cyber security capture-the-flag game. The example demonstrates that the proposed methods enable substantial reductions in both the number of test runs (half) compared with a full factorial and the computational time for the stopping criterion.



Open Access Statement: This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “Decision Analysis. Copyright © 2020 The Author(s). <https://doi.org/10.1287/deca.2020.0412>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Funding: This work was supported by the Division of Social and Economic Sciences [Grant 1409214] and Directorate for Computer and Information Science and Engineering [Grant 1912166].

Keywords: applications • security • experimental • game theory • statistics

1. Introduction

In many realistic situations, the individual decision maker is not in complete control of all factor settings that influence outcomes. Instead, multiple decision makers select options and receive rewards that depend on the selections made by all players (Nash 1951). In many of these situations, it can be helpful to estimate the rewards for all possible player action combinations, perhaps focusing on the combinations most likely to be played. Here, we seek efficient experimental methods and stopping criteria to estimate mean rewards or utilities to support decision making in which the starting point is access to low-consequence experimentation, for example, engagement simulations or experimental war-games.

One well-studied set of action combinations of potential interest is Nash equilibria, which involves setting options such that no player could benefit through individual adjustments. The relevance of Nash equilibria is rationalized by expected utility theory (EUT, von Neumann and Morgenstern 2007). Many other explanations for the relevance of Nash equilibria have been provided in the literature. A common view is the “self-enforcing” agreement relating to possible communications between players before play (Brandenberger and Dekkel 1987). In part because of the possibility of these agreements, war-game or other game designers can use the structure of the equilibria to suggest system improvements or other incentives to make the equilibria more desirable, that is, “mechanism”

design (Conitzer and Sandholm 2002, De Clippel et al. 2018).

Correlated equilibria (Aumann 1987) and cumulative prospect theory (CPT) equilibria are generalizations of Nash equilibria (Tversky and Kahneman 1992, Keskin 2016, Phade and Anantharam 2019). More general equilibria have motivations that include the subjective nature of probabilities and the irrationality of decision makers. Selten and Chmura (2008) study multiple types of equilibria and demonstrate that some make more accurate predictions of human behavior than Nash equilibria. The purpose of this article is to provide algorithms for pre-experiments “off-line” to support a variety of equilibria estimates and related mechanism design objectives.

Player rewards are often not known with certainty by all the players. This uncertainty may be an intrinsic property of the game requiring strategies for mitigation (e.g., see Harsanyi 1967). More commonly perhaps, it may be possible to learn the rewards and treat them afterward as known constants. It is possible that some apparent violations of EUT motivating relevant generalizations such as CPT (Keskin 2016, Phade and Anantharam 2019) might relate to parametric uncertainty rather than irrationality. Also, much research addresses how players can learn to reduce the uncertainty by repeatedly playing the real game (e.g., see Chapman et al. 2013, Foster et al. 2016). For these problems, Nash equilibria are sometimes not considered to be relevant. Instead, learning the so-called no-regret decision options (analogous to Nash equilibria) is an important objective. Yet, what if the parameters are unknown and the game is not repeated?

To overview, we start pregame preparations with unknown payoff matrices but also with an ability to experiment off-line; for example, we have a simulation model. We play the game off-line many times, choosing actions for each player following our experimental planning and analysis methods. In each run, we observe the payoffs for all players. Then, we use meta-models to predict all the mean payoffs and determine whether off-line experimental stopping conditions are met. Our goal is to predict the information needed to support decision making, for example, the Nash or cumulative prospect theory equilibria. With off-line experimentation complete (at least temporarily), the decision maker can either redesign the related systems

(e.g., mechanism design of a combat aircraft) and/or play the actual game (i.e., play the real game).

Therefore, this paper is focused on a new problem. The decision maker can reduce parametric uncertainty using so-called off-line experimentation for variance reduction similar to risk reduction (Delquíé 2012). For example, off-line testing might involve a simulation or a low-consequence pregame period. Even though the experiments are off-line, they are not free. For example, creating test ranges and performing sets of war-games might cost billions of dollars. Consider that one player may have 10 or more options, as may the opponent. Then, off-line experimentation would need to support the estimation of literally hundreds of payoff parameters from simulation or real-world tests, which may have replication or other errors.

The original examples of games in the management science literature include the design of advertising strategies and military tactics (Shubik 1955, 2002). In the context of modern online advertising, the internal working of the ad-placement algorithms and the decisions of potential customers create an opaque game system. Through relatively inexpensive experimentation on test markets, decision makers can develop analytical models, predict estimated rewards, and enumerate equilibria to facilitate large-scale campaigns. Similarly, in military (or cybersecurity) contexts, simulation models with inputs from red and blue teams can be tested with replication to produce the inputs for game-theoretic studies, leading to further insights into system vulnerabilities and strategic policy selections.

This paper makes several important contributions:

1. A new class of experimental design problems is introduced that supports the enumeration of equilibria relevant to predicting behavior or mechanism design. The empirical exploration of off-line systems can leverage the need to explore in detail option combinations that relate to settings that decision makers are likely to select.
2. A tractable and scalable sequential algorithm for off-line experimentation to support single-period games is provided. Empirical models are developed and used to predict the expected rewards, and the initial batch of experiments minimizes the prediction variance over the player option combinations of interest. The algorithm converges to the true equilibria

(Nash, correlated, cumulative prospect theory, or other) with probability one.

3. To clarify the convergence properties, an analytical formula for the probability that specific Nash equilibria exist as a function of the available inputs is included. Monte Carlo can be applied for other equilibria. The probability relates to an integral over a multivariate normal distribution with inputs that depend on the data and often realistic assumptions.

4. Demonstration of the novel methods for a simulation-based, cybersecurity capture-the-flag (CTF) red team/blue team game is given. By planning and executing experiments involving inputs from more than one decision maker, advice is provided for participants in upcoming games as well as insights for game or system designers.

The rest of this paper is organized as follows. Section 2 introduces the notation and reviews studies of deterministic games, gaming under uncertainty, and relevant empirical modeling and experimental design methods. Section 3 combines the previous results to provide an initial heuristic procedure and an augmentation algorithm for equilibria estimation. Section 4 characterizes the experimental methods with respect to chances for identifying the true equilibria both for finite samples and asymptotically. Section 5 describes an application in cybersecurity. Section 6 offers conclusions and opportunities for future work. Note that some of the details about the cybersecurity case study are omitted because of space limitation, but these appear in the appendix.

2. Literature Review

This research combines methods and results from the game theory and experimental design literatures. We begin by introducing the notation. Then, we describe the literature and concepts relating to the basic game formulation, empirical modeling methods, and relevant experimental design results.

2.1. Notation

Our notation contains elements from both the game theory and experimental design literatures. Let m and n be the number of real game options for players A and B, respectively. We use A_{ij} and B_{ij} as the reward if player A selects option i and player B selects option j for

players A and B, respectively. Also, \hat{A}_{ij} and \hat{B}_{ij} are the associated estimated quantities. Vectors μ_A and μ_B with dimensions mn are the vectorized mean values of the matrices A and B , respectively, with joint $(2mn) \times (2mn)$ covariance matrix $\tilde{\Sigma}_{AB}$. Let q be the number of Nash equilibria (either true or estimated, depending on the context). The number and location of the equilibria are uncertain because of our uncertainty about the rewards or, equivalently, the payoffs. The real game decision variables are w_A and w_B , which represent probability over the m and n options for players A and B, respectively. The decision variable w_C would apply to a third player. The scalars α^i and β^i represent optimal payoff values that players A and B achieve at Nash equilibria i , and all the candidate equilibria are (w_A^i, w_B^i) for $i = 1, \dots, q$ or simply (w_A^0, w_B^0) for a specific candidate under consideration. Equilibria are “pure” if the vector has probability one on a single action or “mixed” otherwise. The vectors e and l have all entries equal to one and dimensions m and n , respectively (and o is for player 3). For three players, the tensors A, B, C are payoff cubes.

Here, each player action represents a combination of factor-level settings. Also, we use regression models to predict the mean rewards for all combinations of player actions. Let K represent the number of regression model terms and N denote the number of experimental runs. The initial number is N_0 , and M_A and M_B are the number of decision factors for players A and B, respectively. The assumption parameters β_A and β_B relate model coefficients for predicting the player A and player B reward matrices, respectively. The corresponding estimated quantities are $\hat{\beta}_A$ and $\hat{\beta}_B$. The standard deviations of repeated experimental outputs are σ_A and σ_B under the simplest equal variance assumption considered.

In the context of either linear regression or Gaussian stochastic regression (GSR) models, the random errors are N dimensional vectors ε_A and ε_B for players A and B, respectively. For GSR, the correlation function between points is ϕ and the covariance matrix is C . Because of experimental uncertainty, the existence of a candidate equilibrium (as defined in Section 3) is uncertain with probability $p_N(x_0, y_0)$.

Decision factors represent dimensions along which specific player options are available. For example, in a cybersecurity capture-the-flag example, player 1

has choices relating to whether to try to exploit the firewall or pivot immediately to the internal machine. We say that the factor is “firewall pivot” and the levels are “first” and “never,” which means that options or player policies are referred to as a combination of factor levels. In this example, we are implying that the player policies are designed off-line before the real game begins and then followed. Our analysis activity is intended to help the players design these policies. We believe that the decision factor decomposition of the strategy space is relevant for many real-world situations in which key policies are effectively set in a single round, for example, the combinations of chess openings and team preset strategies for cyber CTF.

The experimental design decision variables include $x_{A,k,l}$ and $x_{B,k,r}$ for the setting selection for an experimental run indexed by k , decision factors l and r , and players A and B, respectively. The vectors $x_{A,k}$ and $x_{B,k}$ are M_A and M_B dimensional vectors of settings for run k for players A and B, respectively, with $(x'_{A,k}|x'_{B,k})$ in the region of interest set S_e . Corresponding values relate to the decision options in the game, which may be assumed to represent a discretization of the factor levels: $\tilde{x}_{A,i,l}$ and $\tilde{x}_{B,j,r}$ are the game settings for player A, option i , and decision factor l and for Player B, option j , and decision factor r . These are $\tilde{x}_{A,i}$ and $\tilde{x}_{B,i}$ indexed to run i in vector form. The parameters $r_{i,j}$ weight the option combinations by subjective importance, all set to one by default. The vector $f(x_A, x_B)$ is K dimensional and includes the model terms (e.g., one and $x_{A,2}x_{B,1}$). The design matrix X is $N \times K$ corresponding to the model terms and experimental runs (for coefficient estimation). The vectors Y_A and Y_B are N dimensional response values of players A and B at the experimental points. These responses could be simulation game scores or the income from test markets.

The design matrix \tilde{X} is $mn \times K$ based on the real game available options (for reward matrix estimation). Intermediate matrices for calculating equilibrium probabilities are W_1 and W_2 , which are $m \times (mn)$ and $n \times (mn)$, respectively. Also, $T(m, n)$ is an $(mn) \times (mn)$ permutation matrix. A key intermediate random vector, Z , has dimension $(m + n)$.

In the context of sequential augmentation experimentation, the set of irrelevant option combinations for establishing whether all candidates are equilibria is $S_{irrelevant}$, and the associated random search parameter is $p_{irrelevant}$.

2.2. Bimatrix and Multiplayer Games

In the standard single-period (bimatrix) game, player A sets the probability vector w_A , and player B sets the probability vector w_B . The standard formulation assumes that the reward or payoff matrices A and B are known. We preliminarily entertain this (often unrealistic) assumption for the sake of reviewing a seminal contribution of Nash (1951). With known A and B , the rewards received for the players are derived using the joint formulation

$$\begin{aligned} & \max_{w_A} w'_A A w_B \\ & \text{s.t. } |e' w_A - 1| = 0; w_A \geq 0, \\ & \max_{w_B} w'_A B w_B \\ & \text{s.t. } |e' w_B - 1| = 0; w_B \geq 0. \end{aligned} \quad (1)$$

The payoff values may, in general, represent mean profits or mean utilities. Here, we propose regression-based prediction of payoffs A and B from off-line experiments, and the assumption that these matrices represent mean utilities.

The joint formulation in Equation (1) leads, without loss of generality, to Nash equilibria (w_A^i, w_B^i) for $i = 1, \dots, q$. Selections not among these equilibria indicate (potentially) irrationality. Each Nash equilibrium satisfies the well-known property that player A cannot do better in the first optimization than w_A^i if player B does w_B^i , and player B cannot do better in the second optimization than w_B^i if player A does w_A^i . Players can benefit by knowing the equilibria because they can select among them to maximize their game rewards. Game designers or system owners can benefit from knowing them because they may want to design incentives for players to change their behaviors.

Generalizations to multiplayer games have been explored extensively, including generalizations of Nash equilibria (e.g., Phade and Anantharam 2019). Yet the numbers of rewards need to be estimated and

the notational complexity grow with the number of players. For example, consider the extension of Nash equilibria from a bimatrix game to a three-player game. The payoff cubes are A, B, C and “ \otimes ” is the Kronecker product. Nash equilibria satisfy (Lee and Baldick 2003)

$$\begin{aligned} A \otimes [w_A^i, w_B^i, w_C^i] &\geq A \otimes [w_A, w_B^i, w_C^i] \text{ for all } w_A \in R^{N_1}, \\ &w_A \geq 0, w_A^i e = 1, \\ B \otimes [w_A^i, w_B^i, w_C^i] &\geq B \otimes [w_A^i, w_B, w_C^i] \text{ for all } w_B \in R^{N_2}, \\ &w_B \geq 0, w_B^i l = 1, \text{ and} \\ C \otimes [w_A^i, w_B^i, w_C^i] &\geq C \otimes [w_A^i, w_B^i, w_C] \text{ for all } w_C \in R^{N_3}, \\ &w_C \geq 0, w_C^i o = 1. \end{aligned} \quad (2)$$

2.3. Equilibrium Conditions

Even with known A and B , the general problems of finding the number of equilibria q and the actual equilibria (w_A^i, w_B^i) are NP-hard in terms of the numbers of options m and n (Chen and Deng 2006, Conitzer and Sandholm 2008, Daskalakis et al. 2009). However, state-of-the-art solution methods can practically enumerate equilibria for problems in which both players have hundreds of options (Savani and von Stengel 2015). Also, necessary and sufficient conditions for the equilibria (Mangasarian and Stone 1964) relate to the existence of scalar α^0 and β^0 satisfying

$$\begin{aligned} w_A^{0'} A w_B^0 - \alpha^0 &= 0, \\ w_A^{0'} B w_B^0 - \beta^0 &= 0, \\ A w_B^0 - \alpha^0 e &\leq 0, \\ B' w_A^0 - \beta^0 l &\leq 0 \\ e' w_A^0 - 1 &= 0; w_A^0 \geq 0, \text{ and} \\ l' w_B^0 - 1 &= 0; w_B^0 \geq 0. \end{aligned} \quad (3)$$

More general multiplayer conditions such as those in Equation (3) are available (Phade and Anantharam 2019). The key features of all correlated equilibria only involve the rows and columns associated with nonzero values of w_A and w_B .

2.4. Empirical Prediction of Payoff Matrices

Key to our approach is the use of planned experiments and empirical regression models to predict simultaneously all

the mean parameters in both payoff matrices A and B . Whereas the decision variables are weights or probabilities (i.e., w_A, w_B) for real games, the empirical model building decision variables are the factor-level settings (i.e., $x_{A,i}, x_{B,i}$) for off-line experimental games. The experiments can be off-line or not “real” in the sense that they do not require playing the game; for example, one can experiment on a simulation model of the game as we illustrate for our cybersecurity-planning example in Section 5. The experiments could also be relatively low-consequence pre-experiments, for example, involving test markets.

Consider that system options are potentially combinations of factor levels; that is, option combination or run i is represented by the settings $(x_{A,i}, x_{B,i})$. The standard linear model functional form is $f'(x_{A,1}, x_{B,1})$. The “design” matrix (e.g., see Goos and Jones 2011) is

$$X = \begin{pmatrix} f'(x_{A,1}, x_{B,1}) \\ \vdots \\ f'(x_{A,N}, x_{B,N}) \end{pmatrix}. \quad (4)$$

Gaussian process regression is a generalization of ordinary linear models (e.g., see Gorodetsky and Marzouk 2016). The multivariate expressions of the rewards, Y_A and Y_B , derive from model coefficients, β_A and β_B , and random errors, ϵ_A and ϵ_B :

$$\begin{pmatrix} Y_A \\ Y_B \end{pmatrix} = \begin{pmatrix} X\beta_A \\ X\beta_B \end{pmatrix} + \begin{pmatrix} \epsilon_A \\ \epsilon_B \end{pmatrix}, \quad (5)$$

where the random errors could derive from simulation lack of repeatability, for example, Monte Carlo random errors in cyberattack simulations. A common assumption is that the random errors follow a multivariate normal (MN) distribution with variance covariance matrix, Σ_{AB} :

$$\begin{pmatrix} \epsilon_A \\ \epsilon_B \end{pmatrix} \sim MN[0, \Sigma_{AB}]. \quad (6)$$

Here, we consider both the standard $\Sigma_{AB} = \sigma^2 I$ linear model regression assumption and a more general GSR assumption in terms of scalar variance parameter, τ :

variance, σ ; directional parameters, θ_k ; and variance-covariance matrix, C . This gives

$$\Sigma_{AB} = I\tau^2 + \begin{bmatrix} C & 0 \\ 0 & C \end{bmatrix}$$

$$\text{where } C = \begin{bmatrix} \phi[(x_{A,1}, x_{B,1}), (x_{A,1}, x_{B,1})] & \dots & \phi[(x_{A,1}, x_{B,1}), (x_{A,N}, x_{B,N})] \\ \vdots & \ddots & \vdots \\ \phi[(x_{A,1}, x_{B,1}), (x_{A,N}, x_{B,N})] & \dots & \phi[(x_{A,N}, x_{B,N}), (x_{A,N}, x_{B,N})] \end{bmatrix} \quad (7)$$

$$\text{and } \phi[(x_{A,i}, x_{B,i}), (x_{A,j}, x_{B,j})]$$

$$= \sigma^2 \exp \left[\sum_{k=1}^n \left(\frac{(x_{A,i}, x_{B,i}) - (x_{A,j}, x_{B,j})}{\theta_k} \right)^2 \right].$$

In our case study, we consider only categorical factors, so continuous variable GSR is not relevant. Yet we believe that the flexibility of GSR is critical for problems involving continuous factors, and specific results apply to the more general GSR assumptions as we clarify. Using the standard linear model case assumption ($C = 0$), the least squares coefficient estimates are $\hat{\beta}_A = (X'X)^{-1}X'Y_A$ and $\hat{\beta}_B = (X'X)^{-1}X'Y_B$.

Then, the payoff matrix estimates, \hat{A} and \hat{B} , can be predicted using the regression in their vectorization forms, $vec(\hat{A})$ and $vec(\hat{B})$, and a full factorial design matrix \tilde{X} :

$$vec(\hat{A}) = \begin{pmatrix} \hat{A}_{1,1} \\ \vdots \\ \hat{A}_{n,1} \\ \vdots \\ \hat{A}_{i,j} \\ \vdots \\ \hat{A}_{n,m} \end{pmatrix} = \tilde{X}\hat{\beta}_A, \quad vec(\hat{B}) = \tilde{X}\hat{\beta}_B, \quad \text{and where}$$

$$\tilde{X} = \begin{pmatrix} f'(\tilde{x}_{A,1}, \tilde{x}_{B,1}) \\ \vdots \\ f'(\tilde{x}_{A,n}, \tilde{x}_{B,1}) \\ \vdots \\ f'(\tilde{x}_{A,i}, \tilde{x}_{B,j}) \\ \vdots \\ f'(\tilde{x}_{A,n}, \tilde{x}_{B,m}) \end{pmatrix}. \quad (8)$$

2.5. Experimental Design

It is well known that the accuracy of the empirical model greatly depends on the experimental design points used in its construction. In our game context, the experimental runs are pairings of level settings chosen by both players: $(x_{A,1}, x_{B,1}), \dots, (x_{A,N}, x_{B,N})$. The accuracy also depends on the discrete points that form the options for the game: $(\tilde{x}_{A,1}, \tilde{x}_{B,1}), \dots, (\tilde{x}_{A,m}, \tilde{x}_{B,m})$. Typically, perhaps, there may be many more combinations of player options than experimental budgets can afford, that is, $mn \gg N$. This makes the use of optimization particularly critical to permit prediction of the payoff matrices and, thus, accurate estimation of the Nash equilibria.

The accuracy of a linear model also depends on the terms included or the so-called functional form implied by the vector functions $f_A'(x_{A,1}, x_{B,1})$ and $f_B'(x_{A,1}, x_{B,1})$ for deriving estimated payoff matrices \hat{A} and \hat{B} , respectively. A concise and relevant functional form includes only the first order terms and players A and B interactions:

$$f_A'(x_{A,1}, x_{B,1}) = f_B'(x_{A,1}, x_{B,1})$$

$$= \begin{pmatrix} 1 & x_{A,1} & \dots & x_{A,M_A} & x_{B,1} & \dots & x_{B,M_B} \\ x_{A,1}x_{B,1} & \dots & x_{A,M_A}x_{B,M_B} \end{pmatrix}. \quad (9)$$

For continuous variables, more detailed and accurate models may also be of interest, including adding quadratic terms, for example, $x_{A,1}^2$. The standard regression model with parameters $\hat{\beta}_A$ and $\hat{\beta}_B$ to predict a generic mean reward is

$$\hat{y}_A = f'(\tilde{x}_A, \tilde{x}_B)\hat{\beta}_A \quad \text{and} \quad \hat{y}_B = f'(\tilde{x}_A, \tilde{x}_B)\hat{\beta}_B$$

$$\hat{y}_B = f'(\tilde{x}_A, \tilde{x}_B)\hat{\beta}_B. \quad (10)$$

These models have prediction variances of the form

$$\text{var}[\hat{y}_A(x_A, x_B)] = \sigma_A^2 f'(x_A, x_B)(X'X)^{-1}f(x_A, x_B) \quad \text{and}$$

$$\text{var}[\hat{y}_B(x_A, x_B)] = \sigma_B^2 f'(x_A, x_B)(X'X)^{-1}f(x_A, x_B). \quad (11)$$

A natural objective to generate the initial experimental points is to minimize the average prediction errors over the lattice of player options. The standard prediction variance formula yields the following experimental design formulation. Gorodetsky and Marzouk (2016) provide a formulation relevant to Gaussian process regression. Here, we focus on an initial design with the linear model terms because our application has only categorical variables.

Assume that the relevant variance is σ_A^2 (could be σ_B^2) and $r_{i,j}$ is the weight for player A, game option i and player B, game option j . In our example, we assume that all option combinations are equally of interest ($r_{i,j} = 1$) at the start of experimentation, but more general assumptions could be important. Therefore, weighted prediction variance (WPV) reduces to the average prediction variance (APV), and the optimization is over points in the experimental set S_e . Then, the WPV formulation, including the relevant option combinations for a fixed number of runs, N , is

$$\text{Minimize : } \frac{\sigma_A^2}{(\mathbf{x}'_{A,1}, \mathbf{x}'_{B,1}), \dots \in S_e} \sum_{j=1}^n \sum_{i=1}^m r_{i,j} \mathbf{f}'(\tilde{\mathbf{x}}_{A,i}, \tilde{\mathbf{x}}_{B,j}) (\mathbf{X}'\mathbf{X})^{-1} \mathbf{f}(\tilde{\mathbf{x}}_{A,i}, \tilde{\mathbf{x}}_{B,j}).$$

For the $r_{i,j} = 1$ case, the APV can be simplified as

$$\text{Minimize : } \frac{\sigma_A^2}{(\mathbf{x}'_{A,1}, \mathbf{x}'_{B,1}), \dots \in S_e} \text{Tr}[\tilde{\mathbf{X}}' \tilde{\mathbf{X}} (\mathbf{X}'\mathbf{X})^{-1}], \quad (12)$$

where mn is the number of decision points; the number of runs, N , is fixed by the dimensions of \mathbf{X} ; and “Tr” is the trace or sum of the diagonal elements. Even with only linear regression modeling, the APV formulation in Equation (12) is NP-hard (Ko et al. 1995). However, using the Meyer and Nachtsheim (1995) coordinate exchange algorithm with 1,000 random starting points, as suggested by Goos and Jones (2011), all the problems considered here were approximately solved in 10 seconds to within 0.1% of optimality with JMP[®] software.

3. Experimental Procedures

In this section, we combine the previous results to create a one shot and sequential empirical equilibrium enumeration procedures. Also, we describe decision making about the initial number of experimental runs (N_0).

3.1. Initial Equilibria Estimation Procedure

Procedure 1 begins by optimally planning and executing off-line test runs, for example, game simulations. Then, the payoff matrix inputs to the bimatrix game formulation are predicted. Finally, the estimates can be used to enumerate the Nash or other equilibria with standard equilibrium enumeration methods (e.g., Savani and von Stengel 2015) based on the approximate assumption that the bimatrix inputs in Equation (1) are known.

Procedure 1 (Initial Experimentation and Equilibria Estimation).

1. Identify the factor levels for experimentation and mn game combinations of interest.
2. Solve the APV formulation in Equation (12) with N_0 runs (see Section 3.2 for choosing).
3. Collect experimental data following the optimal plan to derive the vectors \mathbf{Y}_A and \mathbf{Y}_B .
4. Estimate the empirical model parameters, for example, using least squares estimation.
5. Estimate the payoff matrices $\hat{\mathbf{A}}$ and $\hat{\mathbf{B}}$ using coefficients $\hat{\beta}_A$ and $\hat{\beta}_B$ with Equation (8).
6. Derive the candidate equilibria, (w_A^i, w_B^i) for $i = 1, \dots, q$, for example, by solving Equation (3) assuming $\mathbf{A} = \hat{\mathbf{A}}$ and $\mathbf{B} = \hat{\mathbf{B}}$.

In general, candidate equilibria from empirical procedures (such as procedure 1) may not be true equilibria of interest. This follows because experimental random errors (and model bias) can make it so that $\mathbf{A} \neq \hat{\mathbf{A}}$ and $\mathbf{B} \neq \hat{\mathbf{B}}$. Yet, with sufficiently large experiments, that is, $N \gg 0$, the empirically derived equilibria can be expected to converge to the desired equilibria as we describe in Section 4.

3.2. Average Prediction Variance Designs

As mentioned previously, the APV objective offers a natural formulation for initial experimentation. Intuitively, if the variances of parameters $\hat{\mathbf{A}}$ and $\hat{\mathbf{B}}$ are minimized, the chances of identifying key insights and the correct equilibria improve. Also, excluding interactions from the model that are not critical, for example, third order or higher interactions, is a technique commonly used in many types of experiments to limit costs (Goos and Jones 2011). In bimatrix games, the interactions between player one and player two variables are intuitively critical for identifying equilibria. Related results can likely be established rigorously because, without interactions, players could optimize separately, making the games uninteresting to play and study. Higher order interactions are likely critical for games with multiple players. Therefore, regression model forms need to be changed to consider multi-player games.

Selecting the initial number of runs, N_0 , unavoidably involves some amount of subjectivity. In general, N_0 , values greater than or equal to the number of model terms, K , are desirable so that the estimation problem is

fully determined. Yet setting N_0 at much less than the full number of option combinations, mn , is also desirable to reduce costs.

Table 1 shows the $N_0 = 24$ experimental plan used in the cybersecurity CTF case study. Additional information about the five factors involved is provided in Section 5. Here, we simply note that the model form in Equation (9) includes the three-level categorical factor firewall pivot ($x_{A,i,1}$) and $K = 16$ terms.

Using JMP software, it is possible to solve the APV formulation with equal weights ($r_{ij} = 1$ for $i = 1, \dots, m$ and $j = 1, \dots, n$) for varying numbers of runs. Figure 1 shows the plot of the average prediction variance for these optimal designs. The plot shows the typical “elbow” curve such that incurring additional expense beyond a certain point carries diminishing returns for prediction accuracy. This explains the choice of $N_0 = 24$ for the case study shown in Table 1. Note that the full factorial experimental design is APV optimal or near optimal with $N = 48$ for this problem. Yet $N = 48$ offers only an incremental benefit for prediction accuracy as evidenced in Figure 1 on the right-hand side.

It is true, however, that using a design with $N > 16$ (the minimum number for estimating the model we employ) could permit a more complete model form than the one in Equation (9) to be fitted. Yet, in the context of the case study, the resulting model would (with high probability, see Section 5) result in the same Nash equilibrium being identified with increased experimental cost.

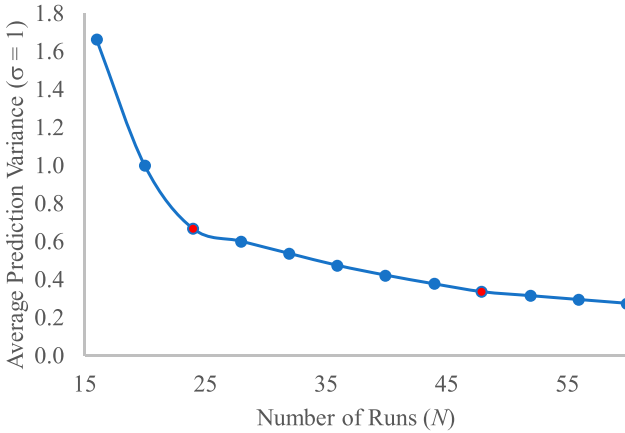
3.3. An Augmentation Procedure

After initial experimentation and equilibria estimation, significant uncertainties can remain; that is, the variances in the \hat{A} and \hat{B} estimates in Equation (12) may not be negligible, leading to uncertainty about the relevant equilibria. This depends on the variances of the experimental random errors σ_A^2 and σ_B^2 and, possibly, on the bias from the model form approximation in Equation (9). The initial procedure unavoidably generates at least a single candidate equilibrium because of the fundamental equilibrium existence theorem and the enumeration algorithms (Nash 1951). Intuitively, some experimental options are irrelevant to establishing

Table 1. Average Variance Experimental Design for Cybersecurity Case Study with Five Factors and the Indexed Option Selections

Run	Firewall pivot ($x_{A,i,1}$)	Tenacity ($x_{A,i,2}$)	Start ($x_{B,i,1}$)	Forensic firewall ($x_{B,i,2}$)	Tenacity ($x_{B,i,3}$)	Player 1	Player 2	$Y_{A,i}$	$Y_{B,i}$
1	First	Persist	Firewall	Never	Persist	1	3	40.7	60.8
2	Never	Persist	PC	Inspect firewall	Move on	3	5	63.2	50.5
3	Never	Persist	PC	Never	Persist	3	4	52.6	30.3
4	Wait	Persist	Firewall	Never	Move on	2	7	50.8	64.6
5	Never	Move on	Firewall	Inspect firewall	Persist	6	1	64.1	58.0
6	Never	Move on	PC	Inspect firewall	Move on	6	6	47.7	64.0
7	Never	Persist	Firewall	Inspect firewall	Persist	3	1	63.9	56.8
8	Never	Move on	Firewall	Never	Move on	6	7	41.5	70.7
9	Wait	Persist	Firewall	Inspect firewall	Persist	2	1	54.1	61.4
10	Wait	Persist	PC	Inspect firewall	Move on	2	6	51.6	47.8
11	Never	Persist	Firewall	Never	Move on	3	7	52.4	67.6
12	Wait	Persist	PC	Never	Persist	2	4	49.8	28.9
13	Wait	Move on	Firewall	Never	Move on	5	7	40.7	63.6
14	First	Persist	Firewall	Inspect firewall	Move on	1	5	37.7	60.7
15	First	Persist	PC	Inspect firewall	Persist	1	2	43.3	50.9
16	First	Move on	PC	Inspect firewall	Persist	4	2	35.9	56.2
17	First	Move on	PC	Never	Move on	4	8	38.8	31.6
18	First	Move on	Firewall	Inspect firewall	Move on	4	5	35.2	49.9
19	First	Move on	Firewall	Never	Persist	4	3	37.6	63.0
20	Never	Move on	PC	Never	Persist	3	4	50.4	26.1
21	First	Persist	PC	Never	Move on	1	8	53.5	28.0
22	Wait	Move on	Firewall	Inspect firewall	Persist	5	1	52.1	65.9
23	Wait	Move on	PC	Never	Persist	5	4	47.8	27.2
24	Wait	Move on	PC	Inspect firewall	Move on	5	6	50.5	55.1

Figure 1. (Color online) Minimum Average Variances for Designs with Variable Run Numbers (N)



whether candidate equilibria are the true equilibria. Here, the equilibria considered could be any type. Yet the structure of the irrelevant set may vary. Our proposed procedure applies to Nash and correlated equilibria because related conditions involve only specific row and column values (Phade and Anantharam 2019).

Definition 1. True equilibria are those that would be derived after all parametric uncertainty is removed (e.g., from suitable infinite experimentation).

Avoiding irrelevant experiments for finding true equilibria is the driving objective of the proposed random search method (procedure 2). A key parameter of this method is the probability of selecting from among the irrelevant options, $p_{irrelevant}$. We suggest 0.05 as a default value to minimally evaluate likely unhelpful options.

Procedure 2 (Augmentation Algorithm)

1. Update model (e.g., using sample mean estimates for \hat{A} and \hat{B} or, alternatively, regression prediction) and the associated candidate equilibria (w_A^i, w_B^i) for $i = 1, \dots, q$.

2. (Optional) If stopping conditions are met, stop.

3. Update the irrelevant set:

$$S_{irrelevant} = \{i, j | w_{A,i}^k = 0, w_{B,i}^k = 0 \forall k = 1, \dots, q\}. \quad (13)$$

4. Sample $U \sim [0,1]$

If $U \leq p_{irrelevant}$, then random sample $i, j \in S_{irrelevant}$.

Else, random sample $i, j \notin S_{irrelevant}$.

5. Perform experiment at $(\tilde{x}_{A,i}, \tilde{x}_{B,j})$.

6. Go to step 1.

In the next section, the rigorous properties of the irrelevant set, $S_{irrelevant}$, and the convergence properties of procedure 2 are investigated. A stopping criterion relating to the probability identifying true equilibria is proposed.

4. Properties of Experimental Procedures

In this section, the finite sample and asymptotic convergence properties of the procedures in Section 3 are characterized. We start with the probability that candidates are correctly identified. Then, we apply the results in the context of the initial experimentation and analysis (procedure 1) and provide convergence results of the augmentation methods (procedure 2).

4.1. The Probability That a Candidate Is a Nash Equilibrium

Consider that A and B are uncertain in the sense that the decision maker does not know fully what they are; that is, there is “parametric uncertainty” that can be reduced through experimentation. Given a current parametric uncertainty level in the payoff values, there is uncertainty about whether a candidate equilibrium (w_A^0, w_B^0) would be discovered to be an equilibrium if all the parametric uncertainty were removed through experimentation.

Definition 2. The probability that a candidate equilibrium (w_A^0, w_B^0) is a true equilibrium, p_N , is the chance that the equilibrium is a Nash equilibrium for a random realization of the payoff matrices, A and B .

Theorem 1 provides a method to calculate this probability without the need for time-consuming Monte Carlo simulation of entire enumeration procedures. It also provides insights relating to the data sufficiency of many types of empirical methods for Nash equilibria estimation. The theorem starts with a given candidate equilibrium (w_A^0, w_B^0) , which is a pair of probability vectors. The theorem is based on assumed known values for the means (μ_A, μ_B) and variances (V_A, V_B) of the vectorized payoff matrices, that is, $vec(A)$ and $vec(B)$. A key permutation matrix is $T(m, n)$, which relates $vec(B')$ to $vec(B)$. These payoff matrices are

assumed to be multivariate normally distributed, MN , which is an often-relevant assumption for estimates derived from regression models. This theorem is relevant for general GSR, which is also based on a multivariate normal distribution, that is, not merely least squares regression.

Theorem 1. Assume that the payoff matrices A and B are multivariate normally distributed: $\begin{bmatrix} \text{vec}(A) \\ \text{vec}(B) \end{bmatrix} \sim MN\left(\begin{bmatrix} \mu_A \\ \mu_B \end{bmatrix}, \tilde{\Sigma}_{AB}\right)$. Then, the probability that the candidate feasible solution (w_A^0, w_B^0) is a Nash equilibrium for a random instance of A and B , p_N , is

$$p_N(w_A^0, w_B^0) = \Pr(\{Z_1 \leq 0\} \cap \{Z_2 \leq 0\} \cap \dots \cap \{Z_{m+n} \leq 0\}), \quad (14)$$

where Z is a random $m + n$ dimensional vector. The distribution of Z is defined in terms of $T(m, n)$, which is a matrix that converts the vectorization of an $m \times n$ matrix into its transpose vectorization as

$$Z \sim MN\left(\begin{bmatrix} W_1 \mu_A \\ W_2 \mu_B \end{bmatrix}, \begin{bmatrix} W_1 & 0 \\ 0 & W_2 \end{bmatrix} \tilde{\Sigma}_{AB} \begin{bmatrix} W_1 & 0 \\ 0 & W_2 \end{bmatrix}'\right), \quad (15)$$

and where $W_1 = (w_A^0 \otimes (I - e w_A^0))$ and $W_2 = (w_B^0 \otimes (I - l w_B^0)) T(m, n)$. (16)

Proof. We seek to show that the event in Equation (14) is equivalent to the necessary and sufficient conditions in Equation (3) (Mangasarian and Stone 1964). If this is demonstrated, the probability in Equation (14) is the probability that (w_A^0, w_B^0) is a Nash equilibrium. The last sets of constraints in Equation (3) are satisfied automatically because (w_A^0, w_B^0) is a feasible solution. The first two constraints in Equation (3) are $\alpha^0 = w_A^0 A w_B^0$ and $\beta^0 = w_A^0 B w_B^0$. Plugging these into the following inequalities in Equation (3) and rearranging using scalar properties gives

$$\begin{aligned} A w_B^0 - w_A^0 A w_B^0 e &= I A w_B^0 - e [w_A^0 A w_B^0] \leq 0 \\ (I - e w_A^0) A w_B^0 &\leq 0. \end{aligned} \quad (17)$$

Using the Kronecker product (\otimes) and a standard vectorization (vec) identity (Searle 1982) gives

$$\begin{aligned} (I - e w_A^0) A w_B^0 &= \text{vec}\left(\left(I - e w_A^0\right) A w_B^0\right) \\ &= \left(w_B^0 \otimes (I - e w_A^0)\right) \text{vec}(A) \\ &= W_1 \text{vec}(A) \leq 0. \end{aligned} \quad (18)$$

Similarly, we have

$$\begin{aligned} B' w_A^0 (w_A^0 B y^0) l &= I B' w_A^0 - l (w_B^0 B w_A^0) \leq 0 \\ (I - l w_B^0) B' w_A^0 &= (w_A^0 \otimes (I - l w_B^0)) \text{vec}(B') \\ &= (w_A^0 \otimes (I - l w_B^0)) T(m, n) \text{vec}(B) = W_2 \text{vec}(B) \leq 0. \end{aligned} \quad (19)$$

Together, Equations (18) and (19) give $m + n$ inequalities. Introducing the random vector, Z , Equations (18) and (19) become

$$Z = \begin{bmatrix} W_1 & 0 \\ 0 & W_2 \end{bmatrix} \begin{bmatrix} \text{vec}(A) \\ \text{vec}(B) \end{bmatrix} \leq \begin{bmatrix} 0 \\ 0 \end{bmatrix}. \quad (20)$$

Substituting the assumption for the multivariate payoff matrices, $\text{vec}(A)$ and $\text{vec}(B)$, and using standard matrix identities (Searle 1982), Z is multivariate normal:

$$Z \sim MN\left(\begin{bmatrix} A_1 \mu_A \\ A_2 \mu_B \end{bmatrix}, \begin{bmatrix} W_1 & 0 \\ 0 & W_2 \end{bmatrix} \tilde{\Sigma}_{AB} \begin{bmatrix} A_1 & 0 \\ 0 & A_2 \end{bmatrix}'\right). \quad (21)$$

Therefore, the necessary and sufficient conditions for an equilibrium are satisfied if and only if each element of the vector, Z , is negative. Then, $p_N(w_A^0, w_B^0)$ is equal to the probability that this condition occurs assuming Equation (21). ∴

The central result in Theorem 1 is intuitive. Once there is a candidate equilibrium, there is no need to enumerate all the equilibria for thousands of model scenarios for probability estimation. Instead, simulations need only study whether the small number of normally distributed parameters exceeds the other normally distributed parameters involved in the equilibrium. Equation (14) supports efficient estimation while addressing correlated prediction estimates and mixed equilibria, that is, fractional probabilities. From the estimated equilibria, the player may be able to see which equilibrium an opponent would likely prefer (if there are more than one). If there is only one equilibrium, as appears to be the case in our cyber game example, it may be highly desirable to play related settings.

Note that Theorem 1 could apply to uncertainty of types other than parametric, that is, uncertainty not caused from limited experimental data. For example, it could relate to games with intrinsically random payoffs. Also, Theorem 1 permits the computationally efficient estimation of Nash equilibria probabilities.

The probabilities for more general equilibria, such as correlated or cumulative prospect theory equilibria, can, of course, be estimated using Monte Carlo simulation. Further, regardless of the specific stopping criterion applied, the information about the rewards can support many types of decision support activities. Next, we relate Theorem 1 to empirical uncertainties.

4.2. Application to Procedure 1

In real empirical investigations, the analyst does not have the mean values of the payoff estimates, μ_A and μ_B , and the true covariances, Σ_{AB} . Instead, the analyst has response data, Y_A and Y_B , which can be derived from procedure 1. From this data, estimates can be generated using regression, for example, linear regression in Equation (8). Many types of regression models are multivariate normal distributions. The standard linear regression multivariate normal assumption inspires the following corollary to Theorem 1, which characterizes the finite sample properties of procedure 1.

Corollary 1. Assume the experimental outputs derive from a standard linear model with variances σ_A^2 and σ_B^2 , that is, $\begin{pmatrix} Y_A \\ Y_B \end{pmatrix} \sim MN\left(\begin{pmatrix} X\beta_A \\ X\beta_B \end{pmatrix}, \begin{bmatrix} \sigma_A^2 I & 0 \\ 0 & \sigma_B^2 I \end{bmatrix}\right)$. Also, the functional form, $\mathbf{f}(x_A, x_B)$, for the design matrix, X , is the same as for the prediction design matrix \tilde{X} . Further, procedure 1 is applied. Then, the probabilities that any given candidate equilibrium (indexed by i) is a true equilibrium, $p_N(w_A^i, w_B^i)$, is given by Equation (14) substituting $\mu_A = \tilde{X}\hat{\beta}_A$, $\mu_B = \tilde{X}\hat{\beta}_B$, $(w_A^i, w_B^i) = (w_A^0, w_B^0)$, and

$$\tilde{\Sigma}_{AB} = \begin{bmatrix} \sigma_A^2 \tilde{X}(X'X)^{-1} \tilde{X}' & 0 \\ 0 & \sigma_B^2 \tilde{X}(X'X)^{-1} \tilde{X}' \end{bmatrix}. \quad (22)$$

Proof. Establishing the conditions for Theorem 1 proves that the probability applies. Applying procedure 1 involves using the least squares estimators and gives

$$\text{vec}(\hat{A}) = \tilde{X}\hat{\beta}_A = \tilde{X}(X'X)^{-1}X'Y_A \quad (23)$$

such that each entry is a linear combination of normally distributed random variables and, thus, is also multivariate normal. Similarly, $\text{vec}(\hat{B})$ is multivariate normal. The mean vector, μ_A , is (for mean zero random ε_A):

$$\begin{aligned} \mu_A &= E[\text{vec}(\hat{A})] = E[\tilde{X}\hat{\beta}_A] = E[\tilde{X}(X'X)^{-1}X'Y_A] \\ &= E[\tilde{X}(X'X)^{-1}X'(X\beta_A + \varepsilon_A)] = \tilde{X}\beta_A. \end{aligned} \quad (24)$$

Similarly, $\mu_B = \tilde{X}\beta_B$. The variance-covariance matrix is derived using $\text{var}[Tx] = T\text{var}[x]T'$, where T is an arbitrary matrix:

$$\begin{aligned} \text{var}[\text{vec}(\hat{A})] &= \text{var}[\tilde{X}\hat{\beta}_A] = \text{Var}[\tilde{X}(X'X)^{-1}X'Y_A] \\ &= \tilde{X}(X'X)^{-1}X' \text{var}[Y_A]X(X'X)^{-1} \tilde{X}' \\ &= \tilde{X}(X'X)^{-1}X' \begin{bmatrix} \sigma_A^2 I & 0 \\ 0 & \sigma_B^2 I \end{bmatrix} X(X'X)^{-1} \tilde{X}' \\ &= \begin{bmatrix} \sigma_A^2 \tilde{X}(X'X)^{-1} \tilde{X}' & 0 \\ 0 & \sigma_B^2 \tilde{X}(X'X)^{-1} \tilde{X}' \end{bmatrix}. \end{aligned} \quad (25)$$

Therefore, the payoff matrix estimates are multivariate normal with the prescribed mean and covariance, and Theorem 1 applies. ∴

Remark 1. Corollary 1 inspires an approximate method to estimate the probability that a given candidate equilibrium, (w_A^0, w_B^0) , is a true equilibrium. Procedure 1 generates the regression estimates: $\hat{\beta}_A$, $\hat{\beta}_B$, $\hat{\sigma}_A^2$, and $\hat{\sigma}_B^2$. Then, one can assume that $\hat{\beta}_A = \beta_A$, $\hat{\beta}_B = \beta_B$, $\hat{\sigma}_A^2 = \sigma_A^2$, and $\hat{\sigma}_B^2 = \sigma_B^2$ and use Equation (22) for inputs to estimate the probability the equilibrium is a Nash equilibrium from Theorem 1 in Equation (14). This approach can function as a stopping criterion in procedure 2, that is, stop when the estimated probabilities for all equilibria exceed a threshold such as 95%.

4.3. Irrelevant Experimental Options

Consider a list of candidate Nash equilibria from procedure 1. Some of these candidates may not be true equilibria. In determining the accuracy of the given candidate equilibria, some experimental options are irrelevant. Lemma 1 clarifies conditions under which option combinations are known to be irrelevant to equilibrium probability calculations.

Lemma 1. Assume that the payoff matrices A and B are random. The probability that (w_A^0, w_B^0) is a Nash equilibrium depends neither on the values $A_{i,j}$ with $w_{B,j}^0 = 0$, nor on the values of $B_{i,j}$ with $w_{A,i}^0 = 0$.

Proof. If we can show that the values in question are irrelevant to establishing the necessary and sufficient conditions to be an equilibrium in Equation (3), the result is proven. From Equation (17), we have $(I - \varepsilon w_A^0)Aw_B^0 \leq 0$, and from Equation (18), we have

$(I - Iw_B^0)B'w_A^0 \leq 0$. From the proof of Theorem 1 (and intuition), these are the only dependence that the necessary and sufficient conditions have on the values of the payoff matrices A and B . Therefore, columns of $A_{i,j}$ with $w_{B,j}^0 = 0$ are irrelevant, and so are rows of $B_{i,j}$ with $w_{A,i}^0 = 0$. \therefore

Extending this result to all correlated equilibria using the generalized equilibria conditions (Phade and Anantharam 2019) is proposed for future work. In a problem with numbers of equilibria smaller than the number of option combinations ($q \ll mn$), few payoff values are relevant. Intuitively, knowing that some option combinations are irrelevant should down-weight their importance in empirical data augmentation procedures. This suggests the application of the random search procedure 2.

4.4. Convergence Results

For finite amounts of off-line data, there is some general uncertainty about the Nash, correlated, cumulative prospect theory, or other equilibria. Lemma 1 sheds light solely on which experimental option combinations are relevant with respect to a given list of equilibria. Focusing only on relevant experimental options, therefore, could conceivably miss some of the equilibria, even in the limit of infinite experimentation. Theorem 2 clarifies the implications for long-run applications of procedure 2 augmentation.

Theorem 2. Assume the experimental outputs derive from a standard linear model with variances σ_A^2 and σ_B^2 , that is, $(Y_A, Y_B) \sim MN[(\mu_A, \mu_B), \tilde{\Sigma}_{AB}]$ for finite $\tilde{\Sigma}_{AB}$. Further, the augmentation procedure is applied with sample mean-based estimation and no stopping rule. Both

- i. If $0 < p_{irrelevant} < 1$ in the limit $N \rightarrow \infty$, then all true equilibria are identified with probability one.
- ii. If $p_{irrelevant} = 0$ in the limit $N \rightarrow \infty$, then all identified candidate equilibria are true (correlated or Nash) equilibria with probability one.

Proof. If $0 < p_{irrelevant} < 1$, then, in the limit, $N \rightarrow \infty$, we have $sample\ mean(Y_A) \rightarrow \mu_A$ and $sample\ mean(Y_B) \rightarrow \mu_B$ because of the central limit theorem with finite variances. This implies that (w_A^i, w_B^i) for $i = 1, \dots, q$ are the equilibria. Assume that $p_{irrelevant} = 0$ and the limiting candidate set (w_A^i, w_B^i) for $i = 1, \dots, q$. Then, the relevant sample means converge, and the others are irrelevant to the candidate list. Without loss of generality,

we assume that $sample\ mean(Y_A) \rightarrow \mu_A$ and $sample\ mean(Y_B) \rightarrow \mu_B$. This gives the result for the candidate correlated or Nash equilibria in both cases. \therefore

Theorem 2 implies that focusing on the relevant set of experimental options can increase the probability that points on a given list of candidates are true equilibria. Also, some amount of focus on (apparently) irrelevant options could conceivably aid in the identification of all the true Nash equilibria. This follows because lists of equilibria from finite samples may be incomplete. The implications of data from irrelevant combinations for related regression model-based, sequential procedures are a topic for future research.

4.5. A Stopping Criterion Based on the Expected Value of Information

Even expanded notions of equilibria may not predict behavior accurately. Therefore, the decision maker may desire to entertain specific assumptions about the policies of an opponent (or opponents), and then decision making can be based on the expected value of perfect information; for example, see Delquíe (2012). Assume that a specific w_B^0 is known or assumed, perhaps from studying the off-line simulations or because there is a single policy of interest. The estimated expected value of perfect off-line information (EEVPOI) is

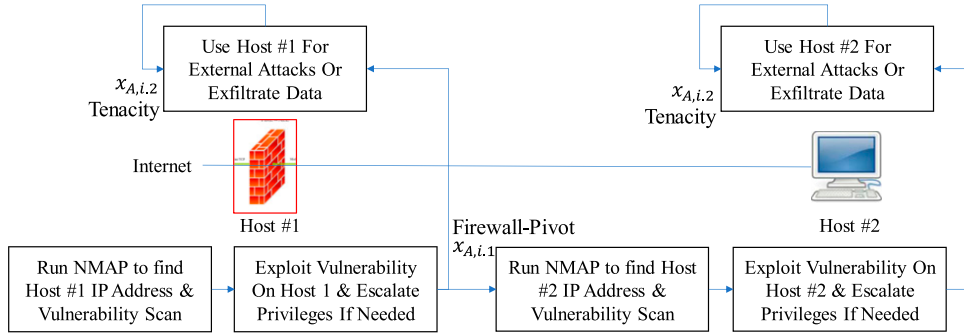
$$\begin{aligned} EEVPOI &= E \left[\max_{A, \epsilon} \left(w_A' (Aw_B^0 + \epsilon_A) \right) \right] \\ &\quad - \max_{w_A} E \left[\left(w_A' (Aw_B^0) + \epsilon_A \right) \right] \\ &= E \left[\max_{A, \epsilon} \left(w_A' (Aw_B^0) \right) \right] - \max_{w_A} \left(w_A' (\hat{A}w_B^0) \right), \quad (26) \end{aligned}$$

where $vec(A) \sim MN[\hat{\mu}_A, \hat{\Sigma}_A]$ with $\hat{\mu}_A$ are regression mean prediction covariance $\hat{\Sigma}_A$ matrix estimates. Instead of stopping off-line experimentation based on the Nash or other equilibrium probability estimates, stopping can be based on threshold values of the EEVPOI. In other words, stop when the expected gain in the utility (bound) is sufficiently small.

5. An Application to Red Team/Blue Team Capture the Flag

For our application, we consider a simulated cybersecurity CTF red team/blue team game. This game seeks to help train many types of students to learn both cybersecurity basics and related policy decision

Figure 2. (Color online) CTF System Diagram, Red Team Tasks, and Factors: Firewall Pivot and Tenacity



making. Additional details about the game and the discrete event simulation model are described in the appendix. Briefly, Figures 2 and 3 show key tasks for red and blue team players, respectively, with a port scanner (e.g., NMAP or Unicornscan) being a program to find IP addresses and scanning and exploiting activities supported by many commercial software.

Next, we describe the application of experimental methods described in Section 3. Zero-sum games have ($A = -B$). By default, we assume that the game is not zero sum to mirror real-life cybersecurity but converting to a zero-sum game is not difficult.

In applying procedure 1, we identify the factor levels for experimentation and $m \times n$ game combinations of interest. Figure 2 shows two player A variables: firewall pivot ($x_{A,i,1}$) and tenacity ($x_{A,i,2}$). A major policy decision that the red team needs to consider is what to do when it successfully exploits the firewall. Should it pivot to attempting to use the firewall as a general-purpose host, exfiltrating any data and launching external attacks or, alternatively, pivot to attacking the internal host? Possible levels include using the firewall

first, waiting until after the internal host is exploited, or never attempting to use the firewall other than to attack the internal host (first, wait, or never). Intuitively, firewalls rarely contain helpful data for exfiltration, and the access of firewall hosts is generally inferior to that of internal hosts. Another major policy choice for attackers is when to give up attempting to use the hosts that they compromise, that is, their tenacity. Levels include giving up immediately upon failure (move on) and never (persist).

Similarly, defenders have policy options as indicated in Figure 3. Unlike attackers, defenders can start their activities on either the firewall or the internal host because of their insider access (firewall or PC). Also, defenders can choose to do forensics on the firewall (inspect firewall) or to ignore its state of compromise (never inspect). Intuitively, forensic activities are time-consuming, and the compromise state of the firewall is not as important as the state of the internal host. Finally, the tenacity of defenders’ in patching attempts can be set. They can give up immediately upon failing to find a patch (move on), or they can persist in

Figure 3. (Color online) CTF System Diagram, Blue Team Tasks, and Factors: Start, Tenacity, and Forensics

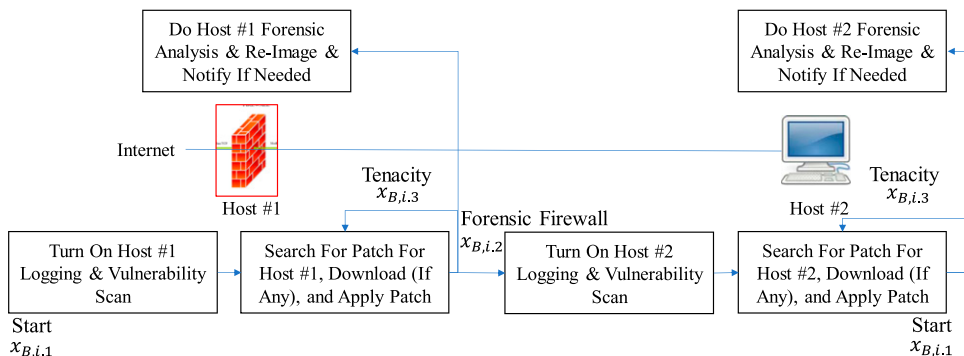


Table 2. Options as Factor-Level Combinations for the (a) Red Team and (b) Blue Team

Panel A: Red Team			
Option	Firewall pivot	Tenacity	
1	First	Persist	
2	Wait	Persist	
3	Never	Persist	
4	First	Move on	
5	Wait	Move on	
6	Never	Move on	

Panel B: Blue Team			
Option	Start	Forensic firewall	Tenacity
1	Firewall	Inspect firewall	Persist
2	PC	Inspect firewall	Persist
3	Firewall	Never	Persist
4	PC	Never	Persist
5	Firewall	Inspect firewall	Move on
6	PC	Inspect firewall	Move on
7	Firewall	Never	Move on
8	PC	Never	Move on

patching attempts (persist). Note that, in our analysis, the game options are the same as the experimental level combinations.

In step 2 of procedure 1, we formulate and solve the APV experimental design problem in Equation (12) with $N_0 = 24$ runs. The solution is in Table 1 based on the factor levels in Table 2. The choice of the number of runs, $N_0 = 24$, reflects a balance between experimental

economy and prediction model accuracy. It also represents the elbow point in Figure 1 as described in Section 3. In step 3, we collect experimental data using the SIMIO model indicated in part in Figure 4 with the 10 replicates following the optimal plan to derive the vectors Y_A and Y_B . In step 4, we estimate the empirical model parameters using least squares estimation with the results shown in Table 3.

Figure 5 plots the prediction model showing the interactions of player-level selections on the scores for both teams (e.g., $XA1 * XB1$ for the first Player A–Player B interaction). The most interesting interaction relates to the choice of the red team never to pivot to use the firewall. This choice benefits both teams unless the blue team is persistent in its patching attempts. Intuitively, this occurs because exploitation of the firewall permits the blue team to successfully patch the internal host, making its later exploitation by the attackers significantly more difficult. Also, starting at the firewall is generally more beneficial for the blue team regardless of the red team selections.

Then, in step 5, we estimate the bimatrix game parameters \hat{A} and \hat{B} , for example, $\hat{\beta}_A$ and $\hat{\beta}_B$ using Equation (8) as shown in Table 4. Enumerating the equilibria that are solutions to Equation (3) by inspection or using standard methods (Savani and von Stengel 2015) results in a single equilibrium as indicated in Table 4 (bolded). It is also apparently the only correlated equilibrium. The equilibrium is option 3

Figure 4. (Color online) SIMIO Model of the Red Team Tasks that Change the System States (Along with Blue Tasks Not Shown)

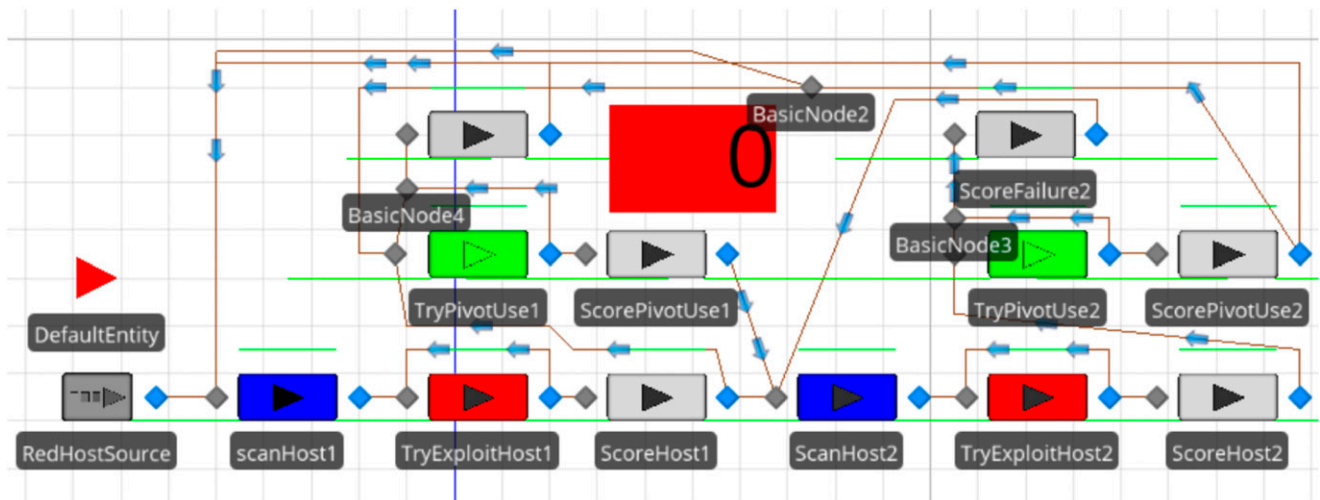


Table 3. Regression Model Summaries for (a) Red Team Score and (b) Blue Team Score

Panel A					Panel B				
Term	Coef	SE coef	T-value	P-value	Term	Coef	SE coef	T-value	P-value
Constant	33.1	2.74	12.07	0	Constant	56.38	3.65	15.43	0
XA1-firewall-pivot					XA1-firewall-pivot				
Never	21	3.36	6.25	0	Never	22.5	4.48	5.03	0.001
Wait	14.1	3.36	4.2	0.003	Wait	20.28	4.48	4.53	0.002
XA2-tenacity					XA2-tenacity				
Persist	6.7	2.74	2.44	0.04	Persist	-2.15	3.65	-0.59	0.572
XB1-start					XB1-start				
PC	3.87	2.74	1.41	0.196	PC	-14.88	3.65	-4.07	0.004
XB2-forensic firewall					XB2-forensic firewall				
Never	3.4	2.74	1.24	0.25	Never	-10.16	3.65	-2.78	0.024
XB3-tenacity					XB3-tenacity				
Persist	1.27	2.74	0.46	0.656	Persist	14.48	3.65	3.96	0.004
XA1-firewall-pivot * XB1-start					XA1-firewall-pivot * XB1-start				
Never PC	-7.07	3.36	-2.11	0.068	Never PC	-3.62	4.48	-0.81	0.441
Wait PC	-4.58	3.36	-1.36	0.21	Wait PC	-7.2	4.48	-1.61	0.146
XA1-firewall-pivot * XB2-forensic firewall					XA1-firewall-pivot * XB2-forensic firewall				
Never-never	-15.13	3.36	-4.5	0.002	Never-never	-0.08	4.48	-0.02	0.987
Wait never	-9.43	3.36	-2.81	0.023	Wait never	-2.9	4.48	-0.65	0.535
XA1-firewall-pivot * XB3-tenacity					XA1-firewall-pivot * XB3-tenacity				
Never persist	8.47	3.36	2.52	0.036	Never persist	-35.58	4.48	-7.95	0
Wait persist	4.48	3.36	1.33	0.219	Wait persist	-27.1	4.48	-6.06	0
XA2-tenacity * XB1-start					XA2-tenacity * XB1-start				
Persist PC	2.42	2.74	0.88	0.404	Persist PC	-4.1	3.65	-1.12	0.294
XA2-tenacity * XB2-forensic firewall					XA2-tenacity * XB2-forensic firewall				
Persist never	2.45	2.74	0.89	0.398	Persist never	3.17	3.65	0.87	0.411
XA2-tenacity * XB3-tenacity					XA2-tenacity * XB3-tenacity				
Persist-persist	-6.38	2.74	-2.33	0.048	Persist-persist	1.4	3.65	0.38	0.712

Note. Coef, coefficients; SE coef, standard error coefficients.

(never use the firewall and persist in using the internal host) for the red team and option 5 (start at the firewall, inspect the firewall, and move on if searching for a patch fails) for the blue team. Notice the Nash condition applies. The highest value in column (5) in Table 3, Panel A corresponds to option 3, and the highest value in row (3) corresponds to option 5 in Table 3, Panel B.

This candidate equilibrium is then evaluated using 10,000 simulations based on Equation (14). The result predicts an estimated (approximate) 96.5% chance of being a true equilibrium based on the right-hand side brute force python code solution of each scenario. For this brute force approach, the computation time is approximately 8.9 hours on an i5-3475 3.2 GHz CPU and python code. This estimate is approximate because correlations between predictions are ignored.

Using the right-hand side of Equation (14), Monte Carlo estimated the exact probability of 95.0% using 0.49 seconds on the MATLAB cloud. This demonstrates

the potentially critical computational advantage afforded by Theorem 1 in computational efficiency and accuracy. Therefore, also, procedure 2 immediately stops using only 24 runs, which is half of a full factorial ($3 \times 2 \times 2 \times 2 \times 2 = 48$ runs).

Similarly, if one assumes that player 2 will play action 5 with probability one (w_B^0), the EEVPOI can be estimated to a good approximation using column (5) in Table 4, Panel A. The table values give estimated means. A constant variance of 7.51 utility units squared is based on the regression results. Then, the first term in Equation (26) is 60.93, and the second term is 60.80 so that the EEVPOI is 0.13 or 0.21%. This may be regarded as negligible such that off-line experimentation can terminate.

The benefits of procedures 1 and 2 for hypothetical players are clear. There are higher point values associated with activities for the inner hosts, but determining sequence and when to give up is difficult. Yet, by selecting the levels indicated by the candidate Nash

Table 4. The Predicted (Mean) Payoff Matrices for Scores: (a) Red Team and (b) Blue Team

Panel A: Red Team								
	1	2	3	4	5	6	7	8
1	34.68	40.97	40.53	46.82	39.80	46.08	45.65	51.93
2	53.26	54.97	49.68	51.39	53.90	55.61	50.33	52.03
3	64.16	63.37	54.88	54.09	60.80	60.01	51.53	50.73
4	34.37	38.23	37.77	41.63	33.10	36.97	36.50	40.37
5	52.94	52.23	46.92	46.21	47.20	46.49	41.18	40.47
6	63.84	60.63	52.12	48.91	54.10	50.89	42.38	39.17

Panel B: Blue Team								
	1	2	3	4	5	6	7	8
1	70.10	51.13	63.11	44.13	54.23	35.25	47.23	28.26
2	63.28	37.10	53.38	27.21	74.50	48.33	64.61	38.43
3	57.03	34.43	49.96	27.36	76.73	54.13	69.66	47.06
4	70.85	55.98	60.69	45.82	56.38	41.50	46.22	31.34
5	64.03	41.95	50.97	28.89	76.65	54.58	63.59	41.52
6	57.78	39.28	47.54	29.04	78.88	60.38	68.64	50.14

equilibria, the player likely maximizes its payoffs in point scores accounting for other players' selections.

The benefits for game designers are also clear. Point selections may be adjusted if the goal is to make deception and decision making important aspects of the game. After each iteration of point value changes by the game designer, procedures 1 and 2 can be applied to generate the equilibria and balance the game (all equilibria have equal payoffs for both players). With the proposed experimental methods, the simulation

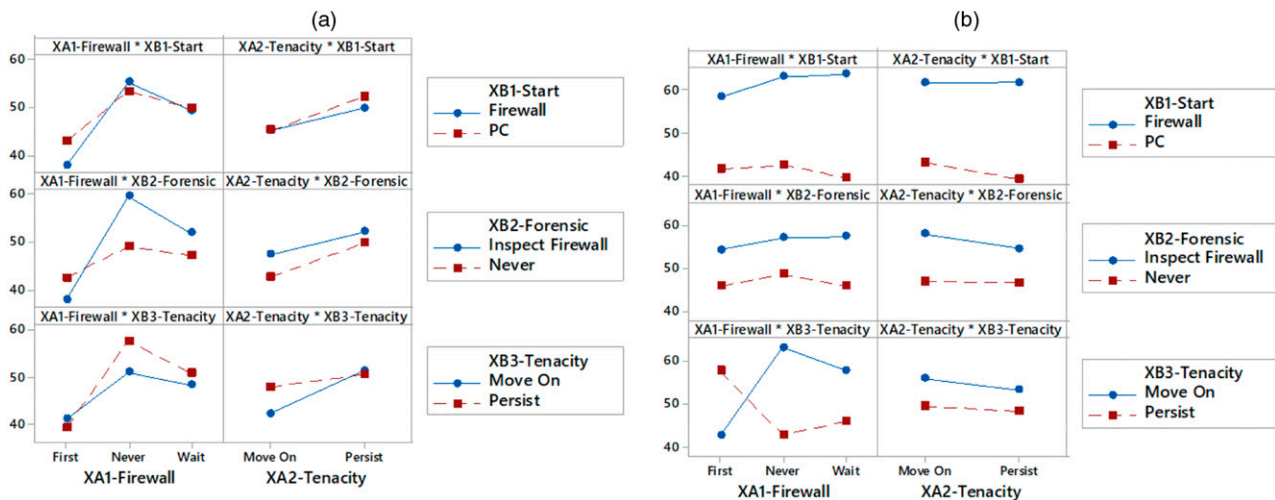
times are reduced by a factor of two from the costs of a full factorial, and the stopping condition times are greatly reduced (8.9 hours to 0.49 seconds). Therefore, approximate assurance is efficiently achieved such that the derived equilibria are actual Nash equilibria of the game irrespective of simulation replication errors.

6. Conclusions and Future Work

In many real management situations, payoff matrices are not readily available, but the ability to experiment off-line is. For example, the manager might have a simulation model with inputs from multiple decision makers or players. Also, there might be an ability to conduct relatively inexpensive test marketing experiments or sets of gaming exercises with a variety of stakeholders. These considerations have motivated a new class of experimental planning and analysis problems. We analyzed these problems and provided experimental plans for initial data collection, sequential methods for efficient follow-up experiments, and stopping rules, for example, stop when all candidate Nash equilibria are likely true equilibria. Additionally, we characterized the finite sample and convergence properties of the proposed experimental procedures.

In our case study game application, we demonstrated the practical benefits of the proposed experimentation and analysis procedures. These procedures permitted the estimation of the Nash equilibrium enumeration procedures at a fraction (half) of the off-line experimental costs of full factorials. Reduction

Figure 5. (Color online) Interaction Plots Showing Predictions for Scores: (a) Red Team and (b) Blue Team



such as this could be a critical enabler as simulations can take days to run and war-gaming exercises with key stakeholders can be difficult to arrange, making test runs extremely expensive. Also, the provided formulas greatly reduce the computational burden of the associated probability estimation process (from 8.9 hours to less than one second). At the same time, the results here focus on single-period games with deterministic payoffs. Limited results apply to cases in which the payoffs may be intrinsically random or estimated using Gaussian stochastic regression (Theorem 1).

Therefore, many opportunities for further research exist. First, efficient stopping rules such as the one in Theorem 1 can be developed for equilibria that more accurately predict human behavior than Nash equilibria, for example, cumulative prospect theory. Second, the subject of off-line experimental planning and analysis can be extended to address many other types of games, for example, repeated, learning, and distributed games. Third, generalizing to more than two players can be explored together with the associated three-factor or higher order interactions.

Fourth, the use of equilibrium probability models for improving the efficiency of sequential experimentation procedures can be investigated. Fifth, issues about approximate and mixed equilibria (e.g., see Feder et al. 2007) and related support points (pure strategy points with positive probability) and empirical estimation-related off-line supporting runs can be studied. Sixth, many applications additional to cyber CTF game design can be explored, including test market design and efficient methods for testing military systems building on previous experimental design results (Johnson et al. 2012). Seventh, results can be generalized to address cumulative prospect theory. Experiments can measure irrationality (in addition to reducing parametric uncertainty) as explored in Phade and Anantharam (2019). Seventh, extensions to multifidelity experimental design are possible building on research in Huang and Allen (2005). Finally, more advanced empirical modeling methods than linear models can be considered, including multifidelity modeling (possibly addressing real and off-line experiments), multiresponse, and Gaussian stochastic regression (Kleijnen and Mehdad 2014) methods can reduce total costs and further extend the practical relevance of game-theoretic

analyses and be related to relevant and irrelevant option combinations.

Acknowledgments

The authors thank Rakesh Vohra, Éva Tardos, and Chen Chen for helpful early discussions.

Appendix A. The Proposed Cyber CTF Game

CTF games divide into two types: jeopardy style in which all participants attack a static network and red team/blue team exercises in which some participants attack and others defend (Antonova et al. 2018). In this article, we propose a red team/blue team CTF game and a simulation model of that game. From our literature search, we believe that red team/blue team game designs are relatively rare. We differentiate in our terminology between actions (or equivalently tasks) and policy options (or factor-level combinations). Policy options are determined in an initial player meeting and govern action selection sequences within the game. Our expectation is that the game period is too short for policy changes as only a small number of actions are time feasible.

We believe that our proposed game offers multiple benefits including the following:

1. It supports relatively rapid training: Both the red and blue teams learn three actions, each of which requires only approximately one hour to study. Students can train for six hours (we estimate) and then play for three hours. At the end, they will have an understanding of scanning for IP address and *vulnerabilities* (e.g., bugs, weak passwords, or out-of-date encryption), *exploiting* vulnerabilities (e.g., applying a software script to gain access or cause mischief), *patching* vulnerabilities (i.e., applying code from vendors of the software to remove the vulnerability), *pivoting* to launch additional attacks (i.e., using the status gained from an exploit to score points, such as launching more attacks), *exfiltration* (i.e., stealing data), escalating privileges (i.e., moving up to system administrator), and performing simple *forensic analysis* (i.e., trying to find evidence of intrusions). Of course, the student experiences are limited, and many attack options in the MITRE framework (Strom et al. 2017) are omitted.

2. Actions are relevant to real-world cybersecurity professionals: The activities in the game are like those conducted currently by cybersecurity professionals and relate to multiple certifications.

3. Decision problems are relevant: Problems faced include pivoting options for the red team and starting options for the blue team. These choices can greatly affect the expected outcomes.

Appendix B. Game Description

Some CTF games include the exploration of an extensive network over multiple days. For example, the MERIT game

Figure B.1. (Color online) Network Model for Game: (a) PC Version, (b) Equivalent Manufacturing Version



covers a virtual small town. Our game focuses on a tiny network model in part so that all activities can reasonably occur within three hours. Figure B.1 shows the network involving two hosts. We can imagine that one host is a firewall (host #1) that is visible to the internet, and the other (host #2) is either (a) a PC or (b) an advanced manufacturing equipment device, such as a networked 3-D printer. The specifics do matter in relation to which vulnerabilities and patches are relevant. Yet, for the purposes of our simulation model, the game activities are simply modeled by the associated mean service times.

There are many ways to define the cybersecurity state of a computer host, which could be a PC, server, printer, exercise machine, 3-D printer, car, or cellphone. In the game system, there are four levels relating to the severity of the worst vulnerability on the host: low, medium, high, and critical. Recently, we have considered adding another state-based scheme including the presence of so-called “celebrity” vulnerabilities, such as “Heartbleed,” a bug famous enough to have its own logo (♥). Here, we consider only two types of hosts, that is, hosts whose worst vulnerability achieves a medium on the common vulnerability scoring system scale and those whose worst vulnerability achieves a critical score for simplicity. Critical vulnerabilities are often so problematic that they can be seen externally to the organization and exploits are widely published. Then, hackers may gain full or near full access to the host almost as easily as by logging in with a known password.

Hosts can also be compromised in the sense that unauthorized personnel can have partial or full access. Therefore, we consider hosts in four states as indicated in Table B.1. If a host has a critical vulnerability on it, it is easier to quickly and completely compromise it. If it is already compromised, it can be of use to hackers who can “pivot” to attack other hosts or exfiltrate data. Blue team personnel naturally seek to identify whether hosts are compromised

and transform them into not-compromised hosts, which have as many of their vulnerabilities patched as possible. Yet, of course, patching and forensic analyses take time as does compromising hosts through manually applied exploits.

We imagine that the multiple members of both red and blue teams follow the same workflow by agreement rather than branching out individually, a choice that might seem sensible given the limited amount of training. More complicated networks and independent team members can be considered in future work.

B.1. Red Team Actions

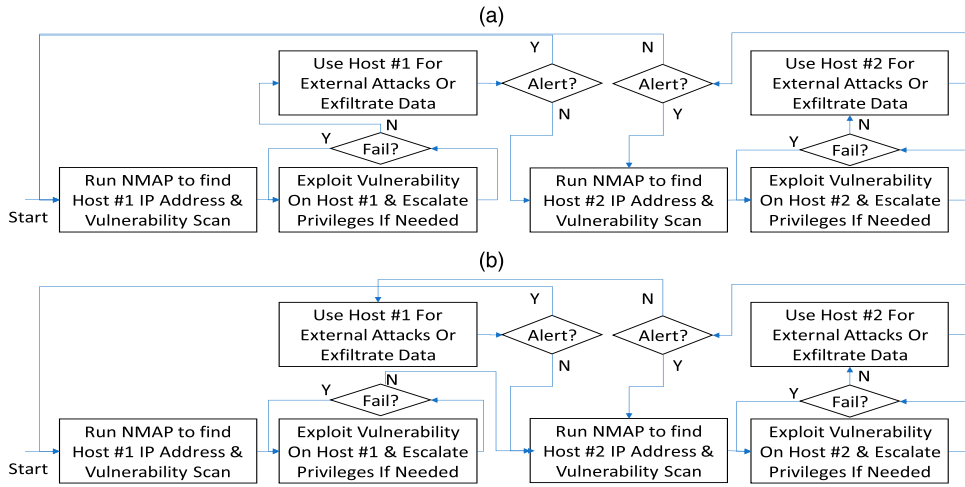
In our proposed game under development, we consider three red team actions:

1. Run script to find visible host IP address(es) and vulnerability scan: On virtually any computer, you can scan to find the visible IP addresses, for example, using a port scan program (e.g., NMAP). With an address, you can scan the host for vulnerabilities, for example, using the Nessus scanner from OpenVAS, Tenable Security, or the Rapid7 scanner.
2. Exploit vulnerability on host and escalate privileges if needed: After identifying the vulnerabilities present on the target host—generally these are so-called “network” vulnerabilities because they are, in part, visible without full access—you can look for associated available exploits. A product such as Metasploit facilitates finding the exploits and launching them. Depending on the level of the vulnerability and the quality of the exploit, one might not gain sufficient access to use the host. A privilege escalation activity may then be possible to gain additional access.
3. Use host for external attacks or exfiltrate data: Once the host is compromised and privileges have been achieved, the host is ready for use, that is, to pivot. Pivoting to external attacks (actually they are internal hosts in our game) is

Table B.1. Four Host States Relating to Compromise and Vulnerability Status

Host state	Compromise state	Vulnerability status
1	Not compromised	Critical and Medium Vulnerabilities
2	Not compromised	Medium vulnerabilities only
3	Compromised	Critical and medium vulnerabilities
4	Compromised	Medium vulnerabilities only

Figure B.2. (Color online) Red Team Workflow: (a) Greedy Version, (b) Patient Version



possible but risky in that the intrusion detection system or firewall rules might identify the compromise and block access. In fact, once access is blocked, personnel can easily isolate that host. A less risky step might be to exfiltrate or steal the information already on the host. Of course, most hosts do not contain monetizable data (e.g., medical records or possibly credit card data).

Figure B.2 shows a workflow that connects the red team actions. In the greedy version (Figure B.2(a)), the red team immediately attempts to use the firewall for gain. In Figure B.2(b), the red team is patient. Note that the workflow implies that the red team can use a host with either external attack or exfiltration but not both. Also, game rules dictate

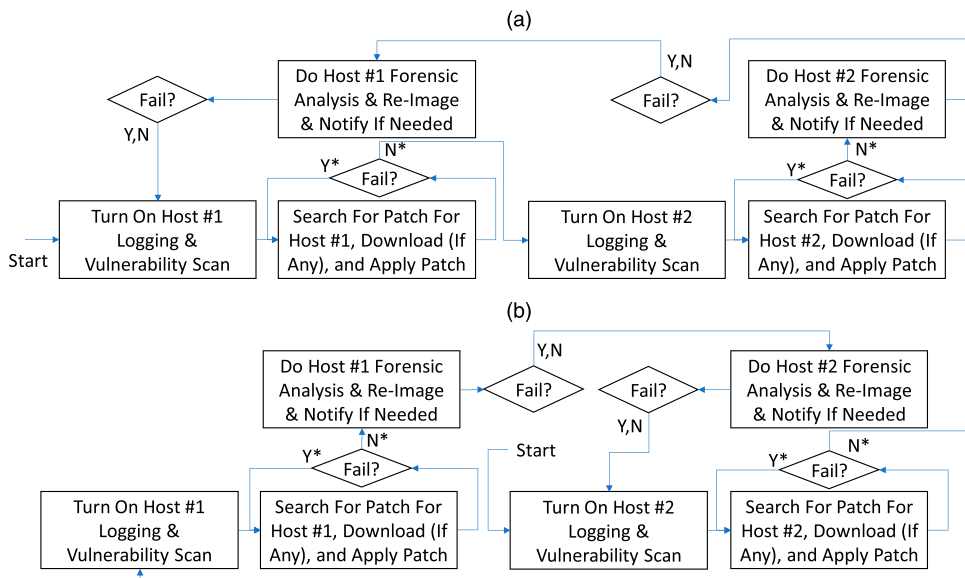
that the red team must attempt a major activity before returning to reuse a compromised host. By “alerts,” we mean declarations that hosts are compromised that limit direct access to attackers.

B.2. Blue Team Actions

Similarly, we consider three blue team (compound) actions:

1. Turn on host logging and vulnerability scan: There are many logging options to record which hosts authorized or unauthorized are doing during their sessions. Enabling basic logging, for example, through the Windows menus, can reasonably preserve privacy (sometimes) while facilitating effective forensic analyses of compromise. Also, the

Figure B.3. (Color online) Blue Team Workflow: (a) Start on the Left and Move Counterclockwise, (b) Opposite



blue team needs to scan for vulnerabilities in a manner similar to that used by the red team.

2. Search, download (if any), and apply patch (if any): Once vulnerabilities have been identified, there are often recommended patching or remediation actions provided by the scanner. Still, sometimes the security personnel must search the internet for patches and related information. Sometimes patches must be downloaded manually and applied, often only after successfully demonstrating that they do not interfere with needed software and services.

3. Do forensic analysis and reimage and notify if needed: Even during a CTF game with only two hosts, it is not clear at any given time whether a given host is compromised. Also, log analyses are supported by many software programs, but the process can be time-consuming. It might also fail to find compromised hosts. If a host is found to be compromised, there are often legal implications. Therefore, notification of affected individuals is likely to be legally required.

Figure B.3 shows two workflows for blue team members. One option is to start with the firewall like the red team, hoping to patch it before the red team exploits it. Alternatively, they can start on the PC or 3-D printer, which is

associated with many more points or payoffs in the game. They can try to patch the worst vulnerability on that host to make it more difficult and time-consuming for the red team. Note that the red team must start on the left-hand side because the firewall is the only host visible to the internet.

Our scoring system assigns points to red and blue teams based on the achievements of host compromises (minor) and host uses (major) with an emphasis on success with the internal host (#2). The blue team scores mirror the red team scores with emphasis on successful forensic work on the internal host.

Appendix C. Input Analysis

For time estimates for the tasks, we use YouTube videos illustrating the actions in applications. By selecting vulnerabilities carefully, we believe that we can control the service time distributions to some extent. Also, we believe that cybersecurity activity times are associated with a high coefficient of variation such that the exponential distribution might reasonably apply.

For one vulnerability for which codes are already loaded and available, exploitation might be quick. For another,

Table C.1. Mean Time Estimates (in Minutes) and Probability Estimates with Supporting YouTube Videos Used to Ballpark Initial Values for the Simulation

Parameter (property)	Description	Mean	Supporting video or notes (if any)
HRserv1	Exploiting host with critical vulnerability	15	https://www.youtube.com/watch?v=ZT7VYsJvh2Q
HRserv2	Exploiting with medium & escalation	30	https://www.youtube.com/watch?v=RdnVC0kNxN4
HRserv3	Entering compromised host	2	Similar to a usual login
HRserv4	Entering compromised host	2	Similar to a usual login
HRharvC	Pivoting to third party attack	20	https://www.youtube.com/watch?v=qIEHUUt2Wfc
HscanTime	Mapping and vulnerability scanning	10	https://www.youtube.com/watch?v=hMKIIRhfk74 , https://www.youtube.com/watch?v=9LA3iQfGGLY
HRprob1&2	Chance exploit works	0.5	Exploits can fail
HRharvNoC	Discovering access is lost	5	Attempted logins and failure
lowCompScore	Game score parameter	0	Chosen by the game designer
lowHarvScore	Game score parameter	10	Chosen by the game designer
highCompScore	Game score parameter	5	Chosen by the game designer
highHarvScore	Game score parameter	25	Chosen by the game designer
HBserv1	Enabling logs & vulnerability scanning	15	https://www.youtube.com/watch?v=hTK0pywfmDE
HBserv2	Enabling logs & vulnerability scanning	5	Fewer vulnerabilities and prescanned
HBserv3	Enabling logs & vulnerability scanning	15	https://www.youtube.com/watch?v=hTK0pywfmDE
HBserv4	Enabling logs & vulnerability scanning	5	Fewer vulnerabilities and prescanned
HBprob1	Patching critical vulnerabilities	0.9	Likely patches are available because of rating
HBprob2	Patching noncritical vulnerabilities	0.5	Likely patches are not available because of rating
HBprob3	Patching critical vulnerabilities	0.9	Likely patches are available because of rating
HBprob4	Patching noncritical vulnerabilities	0.5	Likely patches are not available because of rating
HBlogC	Forensic inspection, reimage, & notify	45	https://www.youtube.com/results?search_query=inspect+host+logs+for+cyber+security+compromise
HBlogNoC	Forensic inspection	30	See similar
HBlogPC	Chance inspection finds compromise	0.9	Chosen by the game designer
HBlogPnoC	Chance inspection finds compromise	0.9	Chosen by the game designer
LimitBRight		0	Chosen by the game designer
lowIndicentRepScore	Game score parameter	20	Chosen by the game designer
highIncidentRepScore	Game score parameter	30	Chosen by the game designer

searching, research, downloads, and testing might require considerable time in at least some instances. Also, players might not be aware that the exploits that they need are preloaded in their Metasploit software, for example. Table C.1 shows the parameters in our simulation model and their descriptions. It also shows the estimated mean values in minutes and the videos used for these estimates. Note that these videos describe the actions by team members with realistic example illustrations.

C.1. SIMIO Model

SIMIO is a commercially available software product with general simulation capabilities. It supports three-dimensional animation, including objects from the Google library. It also offers many experimental design options and associated visualizations. In our simulation model (see Figure 4), we use multiple features, including some for convenience (e.g., “dummy” hosts in which no task is performed):

Paths: Paths and user-defined properties are used to regulate the chances that activities are successful. With the user-defined variables, the model parameters or properties are accessible in the spreadsheet associated with the simulation experiments as well as the following:

States: States are used to store the values associated with the hosts. Then, using the Math.If() formula construction, the properties and the states can set the service times and the success probabilities.

Dummy servers: By using dummy servers, such as Score-Host1, the server features permit assignment conditions for the properties and states. This allows the scores of both teams to be updated as well as the states of the hosts. If actions fail, paths route the attention away from the dummy servers such that the scores and system states remain unchanged.

Duplication for red and blue teams: Even though the two teams work on the same two hosts, they are likely in different rooms and not aware of each other. Also, their service times, success probabilities, and attention paths differ greatly. Therefore, we developed three copies of the network that mirror the red team and two distinct blue team workflows. The blue team has an additional factor and flow because its operations can start either on the firewall (host #1) or on the internal host (host #2). This relates to having full internal access.

References

Antonova LV, Klyuchnikov MM, Loktionov AA, Samsonovich AV (2018) Model of communication and coordination in a capture-the-flag paradigm. *Procedia Comput. Sci.* 145:72–76.
Aumann RJ (1987) Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55(1):1–18.
Brandenburger A, Dekel E (1987) Rationalizability and correlated equilibria. *Econometrica* 55(6): 1391–402.
Chapman AC, Leslie DS, Rogers A, Jennings NR (2013) Convergent learning algorithms for unknown reward games. *SIAM J. Control Optim.* 51(4):3154–3180.

Chen X, Deng X (2006) Settling the complexity of two-player Nash equilibrium. Arora S, ed. *Proc. 47th Annual IEEE Sympos. Foundations Comput. Sci.* (IEEE, Piscataway, NJ).
Conitzer V, Sandholm T (2002) Complexity of mechanism design. Adnan D, ed. *Proc. 18th Conf. Uncertainty Artificial Intelligence* (Morgan Kaufmann Publishers Inc., San Francisco), 103–110.
Conitzer V, Sandholm T (2008) New complexity results about Nash equilibria. *Games Econ. Behav.* 63(2):621–641.
Daskalakis C, Goldberg PW, Papadimitriou CH (2009) The complexity of computing a Nash equilibrium. *SIAM J. Comput.* 39(1):195–259.
De Clippel G, Saran R, Serrano R (2018) Level-mechanism design. *Rev. Econom. Stud.* 86(3):1207–1227.
Delquí P (2012) Risk measures from risk-reducing experiments. *Decision Anal.* 9(2):96–102.
Feder T, Nazerzadeh H, Saberi A (2007) Approximating Nash equilibria using small-support strategies. Mackie-Mason J, ed. *Proc. ACM Conf. Econom. Comput.*, vol. 7 (ACM, New York), 352–354.
Foster DJ, Li Z, Lykouris T, Sridharan K, Tardos E. (2016) Learning in games: Robustness of fast convergence. Lee DD, Sugiyama M, Luxburg UV, Guyon I, Garnett R, eds. *Advances in Neural Information Processing Systems*, vol. 20 (Curran Associates, Red Hook, NY), 4734–4742.
Goos P, Jones B (2011) *Optimal Design of Experiments: A Case Study Approach* (John Wiley & Sons, Hoboken, NJ).
Gorodetsky A, Marzouk Y (2016) Mercer kernels and integrated variance experimental design: Connections between Gaussian process regression and polynomial approximation. *SIAM/ASA J. Uncertainty Quantification* 4(1):796–828.
Harsanyi JC (1967) Games with incomplete information played by “Bayesian” players, I–III Part I. The basic model. *Management Sci.* 14(3):159–182.
Huang D, Allen TT (2005) Design and analysis of variable fidelity experimentation applied to engine valve heat treatment process design. *J. Roy. Statist. Soc.: Ser. C (Appl. Statist.)* 54(2):443–463.
Johnson RT, Hutto GT, Simpson JR, Montgomery DC (2012) Designed experiments for the defense community. *Qualitative Engrg.* 24(1):60–79.
Keskin K (2016) Equilibrium notions for agents with cumulative prospect theory preferences. *Decision Anal.* 13(3):192–208.
Kleijnen JP, Mehdad E (2014) Multivariate vs. univariate Kriging metamodels for multi-response simulation models. *Euro. J. Oper. Res.* 236(2):573–582.
Ko CW, Lee J, Queyranne M (1995) An exact algorithm for maximum entropy sampling. *Oper. Res.* 43(4):684–691.
Lee KH, Baldick R (2003) Solving three-player games by the matrix approach with application to an electric power market. *IEEE Trans. Power Systems* 18(4):1573–1580.
Mangasarian OL, Stone H (1964) Two-person nonzero-sum games and quadratic programming. *J. Math. Anal. Appl.* 9(3):348–355.
Meyer RK, Nachtshheim CJ (1995) The coordinate-exchange algorithm for constructing exact optimal experimental designs. *Technometrics* 37(1):60–69.
Nash J (1951) Non-cooperative games. *Ann. Math. (2)*. 54(2):286–295.
Phade SR, Anantharam V (2019) On the geometry of nash and correlated equilibria with cumulative prospect theoretic preferences. *Decision Anal.* 16(2):142–156.
Savani R, von Stengel B (2015) Game Theory Explorer – Software for the applied game theorist. *Comput. Management Sci.* 12:5–33.

- Searle SR (1982) *Matrix Algebra Useful for Statistics*, Wiley Series in Probability and Statistics, 1st ed. (John Wiley & Sons, New York).
- Selten R, Chmura T (2008) Stationary concepts for experimental 2×2 games. *Amer. Econom. Rev.* 98(3):938–966.
- Shubik M (1955) The uses of game theory in management science. *Management Sci.* 2(1):40–54.
- Shubik M (2002) Game theory and operations research: Some musings 50 years later. *Oper. Res.* 50(1):192–196.
- Strom BE, Battaglia JA, Kemmerer MS, Kupersanin W, Miller DP, Wampler C, Whitley SM, Wolf RD (2017) Finding cyber threats with ATT&CK™-based analytics. Technical Report MTR170202, The MITRE Corporation, Bedford, MA.
- Tversky A, Kahneman D (1992) Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncertainty* 5(4):297–323.
- von Neumann J, Morgenstern O (2007) *Theory of Games and Economic Behavior* (Princeton University Press, Princeton, NJ).

Theodore T. Allen is an associate professor of integrated systems engineering at The Ohio State University. He is the

chief content officer of the Institute for Cybersecurity & Digital Trust and a member of the university security advisory board. His research focuses on the intersection of simulation, experimentation, and optimization. Application domains include cybersecurity, manufacturing, election systems, and supply chain engineering.

Olivia Hernandez is a PhD student in the integrated systems engineering department at The Ohio State University. She received her MS from The Ohio State University and her BS from Dennison. Olivia is a certified usability engineer and held various positions at JPMorgan Chase. Her interests are at the boundary of statistically designed experiments and human factors engineering.

Abdullah Alomair is a PhD student in the integrated systems engineering department at The Ohio State University. He received his MS from the University of Pittsburgh (2015) and his BS from King Fahd University of Petroleum and Minerals (2012). His interests are related to cybersecurity, operations research, and simulation.