Deciphering Cyber Operations

The use of methods and simulations for studying military strategic concepts in cyberspace

Seth T. Hamman
Jelena Vićić
Jack Mewhirter
Philip White
Richard J. Harknett

ABSTRACT

The academic research community faces a significant hurdle when it comes to the study of nation-state cyber operations dynamics. For national security and commercial reasons, little to no cyber operations data is disclosed to the public. Without access to operational data, academic contributions will remain inhibited and the academy will be underutilized in the study of this important strategic domain. We claim that researchers can begin to overcome this information gap by designing experiments that take place within simulation environments. Such approaches are beneficial in that they allow researchers to generate data not easily collected or observed in real-world settings and increase the capacity of researchers to isolate causal effects. In this paper, we describe a simulation environment specifically designed to study cyber operations dynamics below the threshold of armed attack—the competitive space where nearly all nation-state cyber operations activity takes place today. We discuss the simulation environment and then, to illustrate how it can be leveraged to generate tests of research hypotheses, detail our pilot experiment which examines the escalatory dynamics of defend forward activities.

I. INTRODUCTION

The National Security Strategy of the United States (NSS) acknowledges, "today cyberspace offers state and non-state actors the ability to wage campaigns against American political, economic, and security interests without ever physically crossing the border."[1] The Department of Defense National Defense Strategy (NDS) designates cyber as a foundational capability under its Global Operating Model and sets the expectation that cyber operations will be "designed to help us compete more effectively below the level of armed conflict."[2]

Both of these central national documents, along with U.S. Cyber Command's (USCYBERCOM) Vision document and Department of Homeland Security Cyber Strategy expand the focus of cyber operations beyond contributing technically to wartime missions (the focus of previous policy) to include both defensive and offensive capabilities and actions that can delay, degrade, and deny an adversary's ability to produce cumulative effects below the threshold of war that can create strategic advantage.[3][4] According to Lt. Gen. Bruce T. Crawford, U.S. Army Chief Information Officer/G-6, "The bottom line, when it comes to the threat, is that never again will we have the luxury of operating in uncontested space. That's become a part of who we are now."[5]

The academic community faces a significant challenge to contribute empirically-based research findings relevant to this new, more active approach in cyberspace. For national security reasons, most government cyber operations remain classified. Operations involving the private sector are rarely discussed in detail due to commercial concerns. Thus, the empirical record rests primarily on after-action reporting through third-party commercial entities, general news reporting references, or broad attribution of prior attacks. The operations themselves remain opaque and generally unobservable to the academic community in real-time.[6]

The lack of available data constrains academic contributions to the field of cyber operations and conflict. A significant portion of research under the heading of "cyber operations," therefore, tends to focus on technical innovations that are developed under capabilities requirements. While this contribution is essential, it remains too narrow to improve the US position regarding the ongoing cyber campaigns waged against US national sources of power to which the NSS calls attention. Strategic considerations must set and shape technological advances. Without a robust mechanism to generate unclassified, unbiased data about the dynamics associated with cyber operations, academic research will remain constrained by the limits of deductive analysis and limited heuristic case study. In other words, security studies scholars, in particular, have to rely on the security firms' research and reporting by journalists, as well as on heuristics, to make sense of opaque case studies in the absence of unclassified reports.

In this paper we discuss how researchers can utilize experimental research designs and simulation environments to produce unique datasets that allow them to better examine cyber operation dynamics, and potentially, evaluate the efficacy of various cyber strategies. Such approaches may be potentially advantageous even if more observational data were made available, given the high potential of bias introduced through non-random case selection, measurement issues/errors, as well as spurious and/or endogenous relationships between outcome and predictor variables. We then illustrate how such a research design could be employed. We describe a computer environment created to simulate an iterated exchange between two cyber actors (one human player, one computer player). The computer is programmed to employ predefined actions which mimic different cyber strategies: by randomizing computer actions and observing actor response, we are able to evaluate how and to what extent various strategies elicit different types of responses.

Experimental approaches similar to the one described in this paper have long been utilized across various natural and social science disciplines to generate data not easily collected or observed in a natural setting and increase the capacity of researchers' ability to isolate causal effects. Leveraging this established approach and bringing it into the subfield of cyber security studies will help usher in a new range of academic contributions. This project is the first critical step to create that foundational method and we hope to contribute directly to the concerns of the NSS and NDS on how adversary behavior is threatening "the safety of the American people and the Nation's economic vitality."[7][1]

In this paper, we describe approaches to the study of cyberspace and provide background on experimental methodology and wargaming. We then describe our simulation environment, which incorporates some core elements of wargaming. Lastly, we illustrate how experiments can be embedded within the environment to study meaningful questions by detailing our pilot experiment which examines the escalatory dynamics of defend forward activities.

2. RELATED WORK

2.1. Limitations to the Study of Cyber Conflict

The lack of cyber-phenomena-related data inhibits the academic study of cyber conflict. The cyber conflict literature has been stalled in its formative stages due to researchers' inability to create datasets with the full methods and research design toolkit (otherwise available to political scientists and international relations scholars).[8] A study examining the methodological state of the field finds that out of the total number of 70 articles on cyber conflict published between 1990 and 2018, only 9 explicitly discuss their methodology.[9] In almost three decades, not only have there been very few articles focusing on cyber conflict, but the articles that were published suffered from limited or no data. Moreover, there has been a strong focus on theory-building versus theory-testing articles, and both qualitative and quantitative approaches to the study of cyber conflict remain limited.[10] As stated elsewhere, "[t]he need for a more scientifically structured approach is readily apparent to the [academic] community."[11] Overall, the history of cyber studies is plagued with challenges, including a lack of comprehensive case studies and a small number of large n-datasets, as well as high levels of secrecy that inhibit research across different topics.[12] Specifically, we were only able to identify two articles to-date on cyber conflict relying on quantitative data.[13][14][2] The challenge in studying events that are veiled in secrecy is that datasets are not complete and do not represent the whole universe of cases/events of cyber conflict. As such, biases are likely introduced by the limited number of cases that are represented in the sample. For example, only certain classes of cases may be reported by the media or uncovered by security researchers studying cyber conflict.

2.2. Experimental Methods

Capturing the impact that cyber operations strategies (e.g., persistent engagement) have on

[1] It is important to note that while our research questions in terms of strategy were informed by the American policy shift of 2018, the experimental model developed here can be applied to any state actor's strategy.

[2] While quantitative in approach, these studies still rely primarily on the opaque access to real world cases discussed above.

specified outcomes (the rate at which an entity is subject to cyber-attacks, one's resilience to cyber-attacks, the risk of escalation, etc.) is a complicated task. Notably, cyber operations come "with high degrees of anonymity and with plausible deniability; [...] more uncertain in the outcomes they produce; [...][and] involve a much larger range of options and possible outcomes, and may operate on time scales ranging from tenths of seconds to years."[15] As evidenced through the Rubin causal model, measuring the impact that a given strategy or policy has on a specified outcome requires that one observe outcome measures for a given unit (e.g., a country, a cyber-operator) across instances when that policy is enacted and when the policy is not enacted at the same time.[16] For example, if a researcher is interested in evaluating how the defend forward operational concept may impact the rate by which a country is subject to cyber-attacks, one must simultaneously observe the rate at which a country is attacked when that operation is underway and when it is not. The difference between the outcome measures across these two conditions is defined as the "treatment effect" or the "causal effect" of a given policy. Our inability to observe a person/entity at more than one state at a time is known as the "fundamental problem of causal inference," which limits our ability to quantify causal effects.[17]

For this reason, experimental methods have played a significant role in advancing the natural sciences and various social science disciplines, including, but not limited to: psychology, behavioral economics, and political science.[18][19][20][21][22] Such approaches are beneficial in that they allow researchers to generate data not easily collected or observed in a real-world setting and increase the capacity of researchers to isolate causal effects. In essence, social science experiments are meant to replicate—and analyze variation between—how individuals behave when assigned to one of two or more conditions. Each condition is meant to simulate an alternate real-world environment; for example, when the defend forward operational concept is in effect and when it is not. Because individuals are randomly assigned to groups, on average groups share the same characteristics (e.g., ethnicity, sex, abilities, etc.): the only difference is the experimental condition to which they are subject.[23] Because of this, the researcher can be assured that any difference observed in outcomes is *caused* by the treatment effect.

While experimental approaches are advantageous in that they allow the researcher to avoid the problems associated with unobserved counterfactuals and lack of access to real-world data, the extent to which findings can be generalized to the real-world is a function of how well the conditions replicate the real-world environment, and the extent to which treatment effects (again, variation in outcomes between treatment and untreated groups) observed in the sample are consistent with the population of interest (e.g., cyber operators, nation-states, etc.). As the simulation environment deviates from real-world conditions, and as samples differ from the population of interest, the likelihood that patterns of behavior accurately capture reflect real world processes decreases.[24][25]

2.3. Wargaming

In the words of Thomas Allen, "Wargaming is a simulation of war, a horror show without props."[26] It is an exercise envisioned to test, practice, and improve military strategies. An important aspect of wargaming is the realism of the exercise. Modern war games stress realism, as opposed to focusing on theory, and are useful in producing military contingency plans. According to Allen, "Real war games are blueprints of real war."[27]

The first "game center" in the US was operated by the Department of State in 1948, while wargaming was used as a military exercise all throughout the Cold War. From being played only by military personnel, wargames spread to universities such as MIT, and transformed into civilian military-politico exercises involving game theory.[28] Historically, variables of interest in wargaming were related to institutional, cognitive, tactical, and strategic conditions and their effects on decision-making.[29][30] More recently, wargames have been integrated into inexpensive digital game frameworks and new cloud computing architectures.[31] This has paved the way for science-based experimental methods to be combined with wargaming models to produce large datasets for quantitative analysis. In its original form, wargames "examine the processes of warfare and do not present quantitative analysis of military effectiveness."[32]

In the context of cybersecurity, wargaming is mostly used by security firms in anticipation of cyber threats and incident response planning. The first government-led cybersecurity exercise was operated by the Department of Homeland Security in 2006. Named Cyber Storm, it was designed to test "response, coordination, and recovery mechanisms in reaction to simulated cyber events," between international, federal, and state governments.[33] In academia, the Naval War College focuses on testing decision-making processes involving "students, academics, and military thinkers," rather than dynamics of conflict.[34] Academic studies which analyze cyber-related wargaming data are rare. The only existing academic study involving cyber-related wargames is a longitudinal analysis of the escalatory nature of cyber operations using data from wargames played at Naval War College between 2011 and 2016.[35] However, combining wargaming with replicable experimental design can be a useful tool for studying the dynamics of interactions in cyberspace that go beyond testing the decision-making processes in crisis situations.

3. THE SIMULATION ENVIRONMENT

3.1. Our Approach

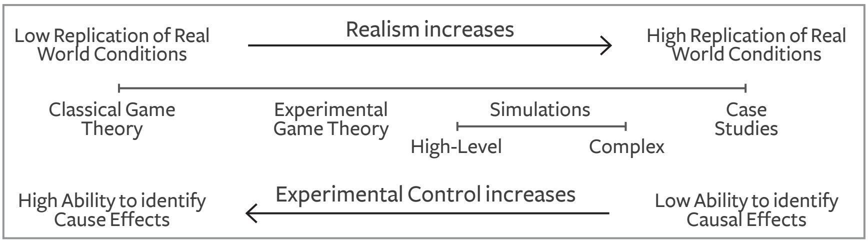
Low Replication of Real World Conditions
Realism increases
High Replication of Real World Conditions
Classical Game Theory
Experimental Game Theory
Simulations
High-Level
Complex
Case Studies
High Ability to identify Cause Effects
Experimental Control increases
Low Ability to identify Causal Effects

Fig. 3-1: The spectrum of scholarly research types for proposed national defense strategies.

The spectrum of potential scholarly support for proposed national defense strategies ranges from purely theoretical to case study-based (see Fig. 3-1). Theoretical approaches, as in classical game theory, are based on presuppositions such as player-perfect rationality and the presumed utility values of the players. This idealized environment enables mathematically-based deductions to be made and can forecast how players will respond in novel situations, assuming that all presuppositions are valid. Case studies, on the other hand, work inductively from real world examples, trying to identify salient characteristics and variables so that broadly applicable generalizations can be made and applied going forward.

In the case of cyber operations, both of these scholarly approaches are problematic. The presuppositions inherent in game theoretical models are likely to be questioned and (perhaps, rightly) distrusted because the domain is novel. As for case studies, the necessary numbers are not readily available for study both because the domain is new and access to existing cyber operations data is unavailable.

Our approach lies in the middle of this spectrum (see the "High-Level" marker in Fig. 3-1). We created a simulated cyber operations environment where data can be generated from diverse populations of research subjects on demand. Generating data from human subjects allows us to limit the number of assumptions that must be made regarding the players and their utility values, and we retain some of the dynamics of real-world cyber operations while still being able to identify causal effects by embedding experiments within the simulation environment.

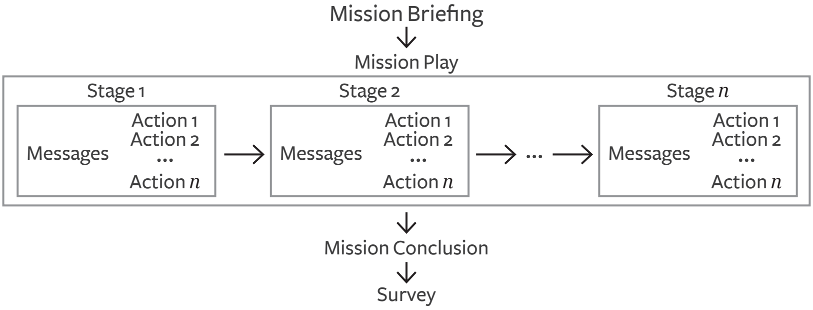

Fig. 3-2: The design of scenarios in the simulation environment.

The environment is a computer game that places research subjects in the role of a nation-state cyber actor. The research subjects participate as monads—each research subject plays against the computer. The research subjects log into the simulation via a web browser from any network-connected computer, and are presented with a *scenario* (see Fig. 3-2). Scenarios start with a *mission briefing*. In the mission briefing, the research subjects are charged with conducting a cyber operation against their adversary which may be modeled after a documented nation-state cyber operation. Details are changed to minimize the risks of research subjects playing out a known script. The mission briefing contains realistic but fictitious background information to contextualize the scenario with the goal of creating intrigue and buy-in from the research subjects. For example, they may be provided with information regarding their country's economy and military strength as well as that of their adversary's, and the scenario may include graphics such as a national seal or flag. These details may be modeled after actual nation-states, but fictitious names are used to avoid any potential psychological associations which could bias the research subjects' actions.

After the mission briefing, the *mission play* unfolds in *stages*. Stages have two elements: informational *messages* to provide the research subjects with situational awareness, and a set of *actions* related to the mission and its progress. In each stage the research subjects are tasked with selecting an action from the choices provided. When their choice is submitted, the next stage begins. The number of stages is variable and is determined by the study designer. The messages allow the study designer to inform the research subjects of the results of their previous actions as well as to update them about adversarial actions against their own network. Missions end with a *mission conclusion*—a final message that can be used to inform the research subjects of the overall results of the mission. At the conclusion of the simulation, the research subject can be presented with an optional *survey* to gather additional data.

The environment is a high-level simulation. The research subjects are not required to actually know how to technically conduct any of the steps of a cyber operation. For example, instead of having to manually run a Nmap network scan, the research subjects can just select the "Conduct Nmap network scan" action from the list of choices provided. In the next stage, they can be informed of the results of previous choices at whatever level of detail the study designer desires. This type of simulation sacrifices realism, but it accelerates the collection of data, and it lowers barriers to participation since trained expertise is not required. Therefore, in addition to technically trained and skilled cyber operators, non-technical managers and strategy-level personnel can participate in the experiments.

To help achieve realism, the stages of the operation are graphically mapped to the "cyber kill chain." First developed by Lockheed Martin in an effort to improve network defense, the Lockheed cyber kill chain "describes phases of intrusions, mapping adversary kill chain indicators to defender courses of action, identifying patterns that link individual intrusions into broader campaigns, and understanding the iterative nature of intelligence gathering."[36]

Because the environment is a high-level simulation, it cannot be used to capture operational nuances such as inadvertent effects that may occur in a real-world mission. However, player reactions to unforeseen events can still be studied through the use of messages. For example, a message could be used to inform research subjects of an unintended consequence of one of their actions (e.g., "Your Nmap scan created a denial of service condition on the enemy network").

The primary benefit of the high-level simulation design is that it allows the environment to be constrained to the exact specifications of the study designer. This allows all of the variables to be held constant except the treatment effect under study. This is vital for making causal inferences because as the level of realism increases in a simulation, confounding variables provide alternate explanations of the behavior observed, and this undercuts the study's ability to identify statistically significant causal effects. Of course, high-level simulations have limitations as well (see Section 4.3), but they can still be an important component of the study of complex phenomena.

3.2. Technical Details

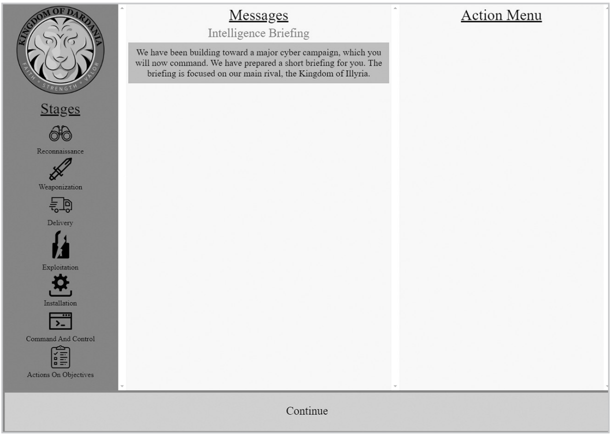

Intelligence Briefing

We have been building toward a major cyber campaign, which you will now command. We have prepared a short briefing for you. The briefing is focused on our main rival, the Kingdom of Illyria.

Fig. 3 3: The mission briefing screen in the simulation environment.

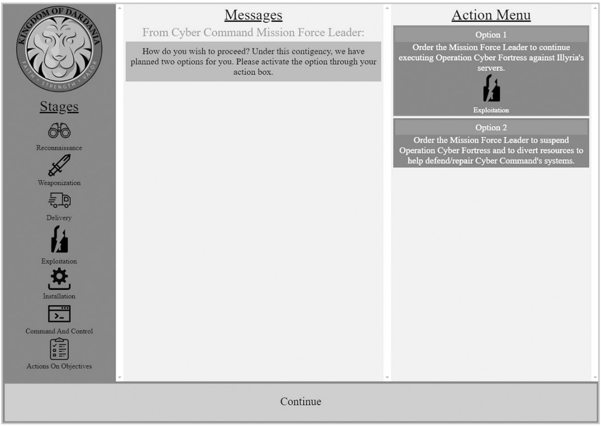

From Cyber Command Mission Force Leader:

How do you wish to proceed? Under this contigency, we have planned two options for you. Please activate the option through your action box.

Option 1

Order the Mission Force Leader to continue executing Operation Cyber Fortress against Illyria's servers.

Exploitation

Option 2

Order the Mission Force Leader to suspend Operation Cyber Fortress and to divert resources to help defend/repair Cyber Command's systems.

Fig. 3 4: The mission play screen in the simulation environment.

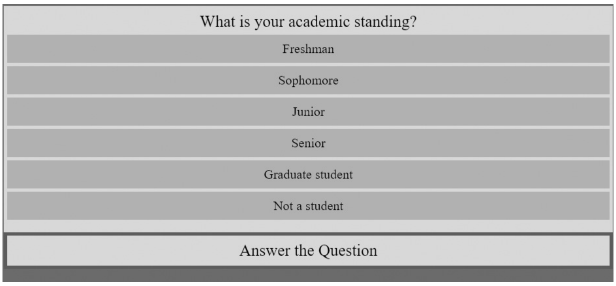

What is your academic standing?

- Freshman
- Sophomore
- Junior
- Senior
- Graduate student
- Not a student

Fig. 3 5: The survey screen in the simulation environment.

The simulation environment is a client-server web application (see Figs. 3-3, 3-4, and 3-5). The application is highly customizable within a framework of scenarios as outlined in the previous section. The research subjects access the scenarios via any network-connected computer by browsing to a specified IP address and port number. The environment supports all standard web browsers. The server application is designed to be hosted on a research administrator's laptop and has a user-friendly graphical user interface (GUI). Client load demands are minimal so hundreds of research subjects can be served from a commodity laptop.

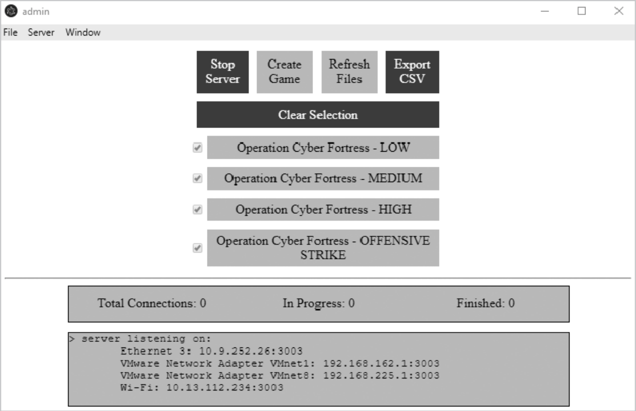
admin
File Server Window
Stop Server
Create Game
Refresh Files
Export CSV
Clear Selection
☑ Operation Cyber Fortress - LOW
☑ Operation Cyber Fortress - MEDIUM
☑ Operation Cyber Fortress - HIGH
☑ Operation Cyber Fortress - OFFENSIVE STRIKE

Total Connections: 0 In Progress: 0 Finished: 0

```
> server listening on:
      Ethernet 3: 10.9.252.26:3003
      VMware Network Adapter VMnet1: 192.168.162.1:3003
      VMware Network Adapter VMnet8: 192.168.225.1:3003
      Wi-Fi: 10.13.112.234:3003
```

Fig. 3 6: Scenario administration: selecting scenarios and starting the server.

We designed the experimental environment to be highly portable. A typical data collection scenario may look like the following: a research administrator walks into a room with a laptop, connects to the Wi-Fi access point, starts the scenario server on the laptop with the click of a button, and announces to the research subjects the IP and port number where they can access the environment (see Fig. 3-6). The research subjects then connect to the server from the web browser of their choice. The research administrator can monitor which scenarios are being served and the status of the clients. The data clients generate is saved to the research administrator's computer for later analysis. Alternatively, the simulation environment can be hosted in the cloud like a typical web application which would allow access from any Internet connected computer.

The environment is built on Electron. "Electron is an open source library developed by GitHub for building cross-platform desktop applications with HTML, CSS, and JavaScript. Electron accomplishes this by combining Chromium and Node.js into a single runtime and apps can be packaged for Mac, Windows, and Linux."[37] We chose Electron because of its versatility and portability, and because all coding is done in JavaScript which is a very popular and well-documented programming language.

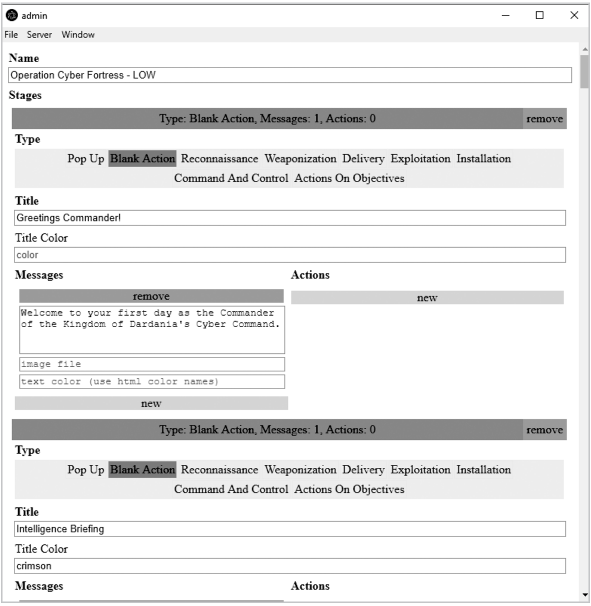
admin
File Server Window
remove
Pop Up Blank Action Reconnaissance Weaponization Delivery Exploitation Installation
Command And Control Actions On Objectives
remove
new
new
remove
Pop Up Blank Action Reconnaissance Weaponization Delivery Exploitation Installation
Command And Control Actions On Objectives

Name
Operation Cyber Fortress - LOW

Stages

Type: Blank Action, Messages: 1, Actions: 0

Type
Blank Action

Title
Greetings Commander!

Title Color
color

Messages

```
Welcome to your first day as the Commander
of the Kingdom of Dardania's Cyber Command.
```

image file

text color (use html color names)

Actions

Type: Blank Action, Messages: 1, Actions: 0

Type
Blank Action

Title
Intelligence Briefing

Title Color
crimson

Messages

Actions

Fig. 3 7: Scenario administration: the scenario builder.

Our primary goal was to make the application as user-friendly as possible for the research administrators. The application package is distributed via a zip file, and can be unpacked on Windows, Linux, and Macintosh computers. Once unpacked, research administrators can use the GUI to create, edit, and delete scenarios (see Fig. 3-7), serve scenarios to research subjects as outlined above, and export the collected experimental data to a comma-separated values (CSV) file for analysis.

4. EMBEDDING EXPERIMENTS WITHIN THE ENVIRONMENT: OUR PILOT STUDY

In this section we provide an example of how experiments can be embedded within the environment. Our study, detailed below, seeks to examine the potential escalatory or non-escalatory dynamics associated with possible defend forward activities.

4.1. The Interaction versus Escalatory Dynamics of Defend Forward Operations

The careful study of the interaction dynamics of nation-state cyber operations is critical because the stakes are high—no country wants to risk inadvertently provoking another into a costly conflict. However, in this new technological paradigm, much uncertainty exists because there are few precedents. Real-world cyber operations data is not widely available, and even if it was, the relative scarcity of data points as well as the unique dynamics of any given data point would make it difficult to generalize the results.

The USCYBERCOM operationalized approach of persistent engagement, which positions cyber operators to defend forward and to actively contest cyber adversaries, is a significant pivot in US cyber strategy and operations. It moves the emphasis from reacting to cyber intrusions to a more anticipatory footing with an emphasis on shifting the balance of initiative so as to keep US networks and data safer. [38][39] However, a defend forward posture involves continually conducting cyber operations on the computer networks of adversaries and, according to its critics, has the potential of provoking adversaries into an escalatory conflict.[40][41] At this early stage in the history of nation-state cyber competition, there exists little open-source evidence in either support or refutation of this claim. Therefore, the initial research questions we are investigating for our pilot experiment is:

> Under what conditions does persistent engagement, which positions cyber operators to defend forward and actively contest cyber adversaries, lead to escalation to armed conflict or intensification of cyber interactions?

In order to study this dynamic, we created a set of scenarios based on published reports of state interaction in cyberspace, while integrating experimental design and war gaming into a simulation environment.

4.2. Experimental Details

Upon logging in to the simulation environment, the research subjects are informed that they are to engage in a cyber operations campaign against an adversary that their initial mission is to launch a spearfishing attack. The subjects are randomly assigned to one of four groups. All four groups are presented with the same mission. The initial stages, messages, and actions are also the same. The groups differ only in the final stage of the scenario as the messages reveal different degrees of defend forward activity that has taken place against the research subject's home network. This variation is the treatment effect that is used to determine how defend forward activities affect the final action selected by the research subjects. We designated the four groups as Defend Forward Low, Defend Forward Medium, Defend Forward High, and Preemptory Strike. The fourth variant is not a defend forward activity but is included as an experimental control to differentiate reactions to defending forward from reactions to a clear (non-inadvertent) adversarial escalatory behavior.

Experimentally, we designate each of the actions presented in the final stage as either escalatory or non-escalatory. We define activities directly related to carrying out the mission and to the defense of the home network as non-escalatory, and offensive activities not directly associated with the carrying out of the mission as escalatory. The last action the research subjects choose is the data point we are studying for this experiment.

This game set-up allows us to examine whether and at what level of intensity defending forward leads to escalatory or non-escalatory responses. Specifically, we can analyze whether increases in the intensity of the defending forward strategy (moving from Low to High) impacts the probability that a respondent chooses an escalatory response or remains in an interactions mode. It also allows us to disentangle at what level of intensity this increase might occur.

Given that subjects are randomly assigned to a given condition (meaning that on average, actors share similar qualities across groups), we can be assured that different propensities to escalate conflict across groups are solely due to treatment effects. We assure balance across groups by comparing demographics across groups. Various pre-estimation balancing techniques (e.g., coarsened exact matching) can be employed, if needed.

4.3 Pilot Limitations

The extent to which our study will provide generalizable information is contingent on the extent to which i) the behavioral patterns observed by study participants parallel those of the population of interest (for example, Russian cyber actors), and ii) our experimental conditions mimic real-world conditions. Here, we address each of these issues.

With regard to the first issue, it is important to note that we are concerned with treatment effects, not the underlying probability of escalation in any treatment or control group. While the sample population (in our case, initially students and eventually US cyber professionals), may differ in many ways that impact the underlying likelihood of choosing an escalatory response (risk aversion, training, culture, etc.) in any given scenario, it is less clear as to why the treatment effects would vary significantly if played by students or professional military personnel. A key difference could be that the sample and target population vary in terms of restraint: when subjected to more aggressive tactics, those with less restraint should be more likely to choose an escalatory option. If cyber operatives, on average, possess more or less restraint than those in our sample, our findings may under or overstate the escalatory effects of defending forward. Researchers should be careful to detail how differences between sample and target populations may impact between group variation when interpreting results.

With regard to realism, things become a bit trickier. In the real-world, cyber operators *generally* work within the context of an organization (in our case, the military). Organizations have institutionalized practices and norms that guide the way the actors respond. We believe that this is a valid subject question to tackle. Organizations are composed of individuals that have unique incentive structures and who possess some level of discretion when making a given choice. When engaging with an adversary, operators have preferences that are then modified by the organizational incentives.[3] The strength by which preferences are modified vary across operators. Actors that prioritize their personal preferences may deviate from what is seen as organizationally correct, which in turn can have a profound impact on the organization as a whole. In military environments, rules of engagement can be expected to channel the behavior of individual operators. In our initial design, we capture a conservative rule of engagement in that the cyber commander making the decisions cannot escalate to conventional war directly, but rather can advise the central decision-maker (a king) to declare war and move to kinetic military operations. A looser rule of engagement can be modelled easily by changing this choice to a direct decision to respond to a cyber-attack with a kinetic attack. For this first design, we stayed with-in public reporting expectations about US rules of engagement. In this instance,

[3] That said, organizational culture and norms have been found to shape the goals and preferences of individual members. This effect tends to vary across organizations and across members in the same organization.

if specific types of defending forward strategies elicit particularly high rates of escalation they perhaps should be avoided (given the propensity to elicit strong personal reactions) knowing well that said reactions will generally be modified. In future iterations of our model we will seek to capture the more fluid real-world environment of interactions across a range of cyber operations to test the potential of inadvertence under more complex testing environments.

5. CONCLUSION

Cyber insecurity is not going away anytime soon. We need to marshal a much broader range of academic expertise and empower it with new research tools so that the technical advances we make are set against real and expected adversarial behavior. This research project is a step in that direction.

At the end of the Second World War, the academic community was harnessed with providing technical, tactical, operational, and strategic guidance to deal with the security implications of the nuclear revolution. We need to once again harness a broad spectrum of academic researchers from across the technical computing and social sciences to pull back the veil of the unexplored security dynamics that have emerged and will continue to emerge in and through cyberspace. The dynamics of cyber operations pose a daunting challenge for states as they face a congested and contested strategic cyberspace domain. Our simulation environment provides a much-needed foundation to leverage unclassified, unbiased, empirically driven academic research to advance interests in an open, global, secure, and stable cyberspace.

ACKNOWLEDGEMENT OF GOVERNMENT RIGHTS AND SPONSORSHIP AND DISCLAIMER

This project is sponsored by the National Security Agency under Grant Number H98230-18-1-0336. The United States Government is authorized to reproduce and distribute reprints notwithstanding any copyright notation herein. This manuscript is submitted for publication with the understanding that the United States Government is authorized to reproduce and distribute reprints. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Security Agency.

Author Bios

Seth T. Hamman

Seth T. Hamman, Ph.D., CISSP, GPEN, CEH, GLEG, is the Director of the Center for the Advancement of Cybersecurity at Cedarville University and an Associate Professor of Computer Science. He earned his B.A. degree in Religion from Duke University, his M.S. degree in Computer Science from Yale University, and his Ph.D. in Computer Science with a specialization in Cybersecurity from the Air Force Institute of Technology. His research interests include helping to shape the growing academic discipline of cybersecurity and collaborating across the multidisciplinary spectrum of cybersecurity.

Jelena Vićić

Jelena Vićić is a Ph.D. candidate in the Department of Political Science at the University of Cincinnati, focusing on international relations, methodology, and cyber strategy studies. She holds the Master of Advanced International Studies (MAIS) degree from the Diplomatic Academy of Vienna. Jelena's research examines the patterns of state strategic behavior in cyberspace.

Jack Mewhirter

Jack Mewhirter is an Assistant Professor in the Department of Political Science at the University of Cincinnati. His research focuses on cooperation problems and influence in polycentric governance systems. Most of his work is done in the context of regional water governance systems.

Philip White

Philip White is an undergraduate Student Fellow in the Center for the Advancement of Cybersecurity at Cedarville University. He is pursuing a specialization in cyber operations within the B.S. in Computer Science major. At Cedarville, he is a member of the Cyber Competition Team and the ACM Programming Contest Team. After graduation, he hopes to pursue a Ph.D. in computer science.

Richard J. Harknett

Richard J. Harknett is Professor and Head of the Department of Political Science and Chair of the Center for Cyber Strategy and Policy at the University of Cincinnati. He co-directs the Ohio Cyber Range Institute. He served as Scholar-in-Residence at US Cyber Command and National Security Agency and in Fulbright Scholar appointments in Cyber Studies (Oxford University, UK) and International Relations (Diplomatic Academy, Vienna, Austria). He has authored over 50 publications in the areas of international relations theory, international security, and cyber security studies.

NOTES

1. White House, "National Security Strategy of the United States," 2018.
2. Department of Defense, "Summary of the National Defense Strategy," 2018.
3. Department of Defense, "Achieve and Maintain Cyberspace Superiority: Command Vision for US Cyber," 2018.
4. Department of Homeland Security,"Cybersecurity Strategy," 2018.
5. Geoffrey Dobson, Anushul Rege and Kathleen Carley, "Virtual Cyber Warfare Experiments Based on Empirically Adversarial Intrusion Chain Behavior," Ed. Leenen, Louise. *ICCWS 2018 13th International Conference on Cyber Warfare and Security* (2018).
6. Brandon Valeriano, Benjamin M. Jensen, and Ryan C. Maness, Cyber Strategy: The Evolving Character of Power and Coercion, New York: Oxford University Press, 2018.
7. White House, "National Security Strategy of the United States," 2018.
8. Max Smeets and Robert Gorwa, "Cyber Conflict in Political Science: A Review of Methods and Literature," Working Paper Prepared for 2019 Annual Convention.
9. Ibid.
10. Ibid.
11. Thomas E. Carroll, David Manz, Thomas Edgar, and Frank L. Greitzer, "Realizing Scientific Methods for Cyber Security," *Proceedings of the 2012 Workshop on Learning from Authoritative Security Experiment Results*, New York: ACM, 2012, 19-24, DOI=http://dx.doi.org/10.1145/2379616.2379619.
12. Max Smeets and Jason Healey, "Cyber Conflict History." C*yber Conflict State of the Field Series: Economic Dimensions of Cyber Conflict*, Cyber Conflict Studies Association, 2017.
13. Brandon Valeriano and Ryan C. Maness, "The Dynamics of Cyber Conflict between Rival Antagonists, 2001–11." *Journal of Peace Research* 51, no. 3, May 2014, 347–360. doi:10.1177/0022343313518940.
14. Nadiya Kostyuk and Yuri M. Zhukov, "Invisible Digital Front: Can Cyber Attacks Shape Battlefield Events?" *Journal of Conflict Resolution* 63, no. 2 (Feb 2019), 317–347.
15. William A. Owens, Kenneth W. Dam, and Herbert S. Lin (eds.) "Excerpts from Technology, Policy, Law and Ethics Regarding U.S. Acquisition and Use of Cyberattack Capabilities," National Research Council, 2009.
16. James N. Druckman, "The power of television images: The first Kennedy-Nixon debate revisited." T*he Journal of Politics* 65, no. 2 (2003), 559-571.
17. Paul W. Holland, Clark Glymour, and Clive Granger, "Statistics and Causal Inference," ETS Research Report Series, 1985.
18. Hal Pashler (Ed.), *Stevens' Handbook of Experimental Psychology, Volume 2, Memory and Cognitive Processes*, Hoboken: John Wiley & Sons, 2004.
19. John Duffy, "Macroeconomics: A Survey of Laboratory Research." *Handbook of Experimental Economics Volume 2*, by J.H. Kagel and A.E. Roth (Eds.), Princeton: Princeton University Press, 2016, 1-90.
20. A. E. Roth, "Introduction." *Handbook of Experimental Economics*, by J.H. Kagel and A.E. Roth (Eds.), Princeton: Prince ton University Press, 1995, 1-110.
21. Rose McDermott, "Experimental Methods in Political Science." *Annual Review of Political Science* 5, no. 1 (2002), 31-61.
22. Samuel J. Eldersveld, "Experimental Propaganda Techniques and Voting Behavior." A*merican Political Science Review* 50 no. 1 (1956), 154-165.
23. Jasjeet S. Sekhon, "The Neyman-Rubin Model of Causal Inference and Estimation Via Matching." *The Oxford Handbook of Political Methodology* (2008).
24. David O. Sears, "College Sophomores in the Laboratory: Influences of a Narrow Data Base on Social Psychology's View of Human Nature." *Journal of Personality and Social Psychology* (1986), 515-530.
25. Alex Mintz, Steven B. Redd, and Arnold Vedlitz, "Can we Generalize from Student Experiments to the Real World in Political Science, Military Affairs, and International Relations?" *Journal of Conflict Resolution* 50 no. 5 (2006), 757–776.
26. Thomas B. Allen, *War Games*, McGraw-Hill Book Company: 1989.
27. Ibid.
28. Ibid.
29. Ibid.

NOTES

30. Jacquelyn Schneider, "The Information Revolution and International Stability: A Multi-Article Exploration of Computing, Cyber, and Incentives for Conflict." Thesis, The George Washington University: 2017.

31. Andrew W. Reddie, Bethany L. Goldblum, Kiran Lakkaraju, Jason Reinhardt, Michael Nacht, and Laura Epifanovskaya. "Next-Generation Wargames: Technology Enables New Research Designs, and More Data." *Science* 362, no. 6421 (2018).

32. Schneider, "The Information Revolution and International Stability: A Multi-Article Exploration of Computing, Cyber, and Incentives for Conflict."

33. Department of Homeland Security, "Cyber Storm: Exercise Report," 2006.

34. US Naval War College, "Cyber and Innovation Policy Institute Holds 1st War Game with Unique Twist. (November 2018), accessed July 2019. https://usnwc.edu/News-and-Events/News/Cyber-and-Innovation-Policy-Institute-holds-1st-war-game-with-unique-twist.

35. Schneider, "The Information Revolution and International Stability: A Multi-Article Exploration of Computing, Cyber, and Incentives for Conflict."

36. Eric M. Hutchins, Michael J. Cloppert, and Rohan M. Amin. "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains." *Leading Issues in Information Warfare and Security Research, Volume 1*, by Julie Ryan (Ed.) (Academic Conferences Limited: 2011).

37. Electron Community, "About Electron" (2019), accessed July 2019. https://electronjs.org/docs/tutorial/about.

38. Michael Fischerkeller and Richard Harknett, "Through Persistent Engagement, the U.S. Can Influence 'Agreed Competition," Lawfare: 2019, accessed October 2019. https://www.lawfareblog.com/through-persistent-engagement-us-can-influence-agreed-competition

39. Richard Harknett, "United States Cyber Command's New Vision: What it Entails and Why it is Important," Lawfare: 2018, accessed October 2018, https://lawfareblog.com/united-states-cyber-commands-new-vision-what-it-entails-and-why-it-matters.

40. James Miller and Neal Pollard, "Persistent Engagement, Agreed Competition, and Deterrence," Lawfare: 2019, accessed October 2019, https://www.lawfareblog.com/persistent-engagement-agreed-competition-and-deterrence-cyberspace.

41. Michael Fischerkeller and Richard Harknett, "A Response on Persistent Engagement and Agreed Competition," Lawfare: 2019, accessed October 2019, https://www.lawfareblog.com/response-persistent-engagement-and-agreed-competition.