Illegal Roaming and File Manipulation on Target Computers

Assessing the Effect of Sanction Threats on System Trespassers' Online Behaviors

Alexander Testa

David Maimon

Bertrand Sobesto

Michel Cukier

University of Maryland—College Park

Research Summary

The results of previous research indicate that the presentation of deterring situational stimuli in an attacked computing environment shapes system trespassers' avoiding online behaviors during the progression of a system trespassing event. Nevertheless, none of these studies comprised an investigation of whether the effect of deterring cues influence system trespassers' activities on the system. Moreover, no prior research has been aimed at exploring whether the effect of deterring cues is consistent across different types of system trespassers. We examine whether the effect of situational deterring cues in an attacked computer system influenced the likelihood of system trespassers engaging in active online behaviors on an attacked system, and whether this effect varies based on different levels of administrative privileges taken by system trespassers.

This research was conducted with the support of the National Science Foundation Award 1223634 and National Security Agency Lablet Award. We thank Adam Bossler, Jean McGloin, Ray Paternoster, and the anonymous reviewers for their helpful comments. We also wish to thank Gerry Sneeringer and the Security Team of the Office of Information Technology at the University of Maryland for their insight on this research. Direct correspondence to Alexander Testa, Department of Criminology and Criminal Justice, University of Maryland, 2220 LeFrak Hall, College Park, MD 20742 (e-mail: atesta@umd.edu).

DOI:10.1111/1745-9133.12312

By using data from a randomized experiment, we find that a situational deterring cue reduced the probability of system trespassers with fewer privileges on the attacked computer system (nonadministrative users) to enter activity commands. In contrast, the presence of these cues in the attacked system did not affect the probability of system trespassers with the highest level of privileges (administrative users) to enter these commands.

Policy Implications

In developing policies to curtail malicious online behavior committed by system trespassers, a "one-policy-fits-all" approach is often employed by information technology (IT) teams to protect their organizations. Our results suggest that although the use of a warning banner is effective in reducing the amount of harmful commands entered into a computer system by nonadministrative users, such a policy is ineffective in deterring trespassers who take over a network with administrative privileges. Accordingly, it is important to recognize that the effectiveness of deterring stimuli in cyberspace is largely dependent on the level of administrative privileges taken by the system trespasser when breaking into the system. These findings present the need for the development and implementation of flexible policies in deterring system trespassers.

Keywords

cybercrime, deterrence, restrictive deterrence, honeypots, randomized field experiment

System trespassing—the illegal access of a computer or computer network (Berthier and Cukier, 2009; Brenner, 2010)—is one of the fastest growing, yet least understood forms of cybercriminal activity. According to recent surveys, nearly half of U.S. firms faced a data breach incident in 2014 (Ponemon Institute, 2014) and most Americans are more fearful of falling victim to cybercrime than to all other serious violent and property crimes (Riffkin, 2014). The concerns related to the potential harm caused by system trespassers pose such a large threat that in April 2015, President Obama issued an executive order authorizing a series of new legal sanctions intended to prevent individuals from "engaging in significant malicious cyber-enabled activities."[1] Still, despite growing concerns regarding the threats posed by system trespassers, there is limited understanding and empirical investigation into the use of sanction threats in deterring malicious online behaviors initiated by system trespassers.

In addressing this issue, we explore the effectiveness of sanction threats presented to system trespassers during the progression of a system trespassing event in dissuading trespassers from navigating and manipulating files on the attacked computer. By drawing on restrictive

[1] For more information, see whitehouse.gov/the-press-office/2015/04/01/executive-order-blocking-property-certain-persons-engaging-significant-m.

deterrence theory (Gibbs, 1975; Jacobs, 2010) and prior interdisciplinary research (Maimon, Alper, Sobesto, and Cukier, 2014; Wilson, Maimon, Sobesto, and Cukier, 2015), the aim of this work is to expand our understanding of system trespassers' response to deterring stimuli in two key ways. First, whereas the aim of prior research has been on the effects of a sanction threat on system trespassers' passive attempts to *avoid* detection on the attacked system (Maimon et al., 2014; Wilson et al., 2015), we examine the effect of a sanction threat on the likelihood of system trespassers to engage in *active online behaviors, including* "roaming" the attacked system and manipulating files permission on the attacked system. Second, although the results of previous research have demonstrated the effectiveness of deterring cues in influencing all types of system trespassers (Maimon et al., 2014; Wilson et al., 2015), we first examine whether the effect of such cues is conditional on the level of administrative privileges imposed by the system trespasser in the attacked computer. To examine these questions, this work employs data collected from a randomized field experiment, in which target computers designed for the purposes of being infiltrated by system trespassers were deployed on the Internet infrastructure of a large American university. Before describing our research design and presenting the results, we begin by describing the nature of system trespassing and by reviewing prior research examining restrictive deterrence in both cyberspace and the physical world.

System Trespassing

System trespassing, alternatively referred to as *computer hacking*, *computer cracking* (Yar, 2006), or *cyber-trespassing* (Wall, 2001), refers to the illegitimate access into a computer or computer network and to the redesign of the hardware or software configuration of these systems in an effort to alter their intended function (Bachmann, 2010; Berthier and Cukier, 2009; Brenner, 2010; the Computer Fraud and Abuse Act of 1986 [§ 1030 (a)(5)(c)]).[2] System trespassers gain illegitimate access to a computer or computer network by seeking out computer vulnerabilities in either random (i.e., any computer system with vulnerabilities) or specific targets (i.e., specific computers to which system trespassers are trying to infiltrate) through a variety of techniques used to search for accessible entry ports (Maimon, Wilson, Ren, and Berenblum, 2015). Once vulnerabilities are discovered, trespassers exploit these security gaps and infiltrate the target computer system, often using credentials of legitimate users or system administrators (Lee, Roedel, and Silenok, 2003). After obtaining unauthorized access, and depending on their motivation (e.g., monetary gain, revenge, exploration, risk seeking, obsession, and ideology [McQuade, 2006; Xu, Hu, and Zhang, 2013]), personality traits (e.g., rationality level and self-control [Bachmann,

[2] Although there are definitional debates on what is considered *system trespassing* and *computer hacking* (Bachmann, 2010; Holt, 2007; Jordan, 2016; Jordan and Taylor, 1998; Schell and Dodge, 2002), our use of the term *system trespassing* is consistent with the interpretation of the term put forth in the Computer Fraud and Abuse Act of 1986 [§ 1030 (a)(5)(c)] as we believe it more accurately describes the scope of illegal behavior that is the focus of the current study.

2010; Bossler and Burruss, 2011]), and skill level (e.g., low or highly skilled [Giboney, Goel, Proudfoot, and Valacich, 2015; Zhang, Tsang, Yue, and Chau, 2015]), system trespassers can set up access, corrupt, and remove any private files hosted on the attacked system. Additionally, system trespassers can use the network to launch subsequent cyberattacks on other computers and victims (Bacher, Holz, Kotter, and Wicherski, 2005; Keizer, 2009; Liu, Xiao, Ghaboosi, Deng, and Zhang, 2009; Provos and Holz, 2007), and once obtaining information from the attacked system, they can sell it or use it for their own needs (Hutchings and Holt, 2015).

Based on the login credentials used to break into the computer network, system trespassers can infiltrate the attacked computer as either an *administrative user* or a *nonadministrative user* (Downing, 2004). Administrative accounts are usually held by information technology (IT) officers and provide a fuller range of access to components of the computer network. System trespassers who choose to infiltrate the attacked system as administrative users (commonly referred to as *root users* or as *super users*) gain admission into the network with the highest level of access (Downing, 2004). Such access guarantees trespassers a distinct set of extended privileges such as access to all files and components of the computer system; the ability to install, remove, or change any program or file on the computer; as well as the ability to cause serious damage to the attacked system. The benefits of gaining administrative privileges during a system trespassing event have been recently demonstrated in a large-scale attack on the U.S. federal government. The system trespassers who infiltrated the system "executed a sophisticated attack that gave them *administrative privileges* into computer networks at the Office of Personnel Management, mimicking the credentials of people who run the agency's system" (Sanger, Perlroth, and Shear, 2015, emphasis added), and they stole millions of personnel records of U.S. government employees.

Alternatively, system trespassers can access a target computer through nonadministrative accounts that are used by legitimate users of the system and that provide a more limited degree of privileges on the attacked computer (Thakare, Chandurkar, and Deshmukh, 2013). Once infiltrating a computer system as a nonadministrative user, system trespassers can still access, alter, remove, or install files on the attacked computer system, yet in a more restricted manner (depending on the account setting). As such, although the privacy of a legitimate computer user is still violated, the potential damage (both intentional and not intentional) that a system trespasser with nonadministrative privileges on the system can cause is minimal. The distinction in the potential for harm caused between trespassers with and without administrative privileges remains so pronounced that cybersecurity experts have recently emphasized that preventing the exploitation of privileged user accounts is essential to enhancing cybersecurity, noting that protecting against attacks on privileged user credentials "has become a critical element of our national defense as recent attacks on government systems reveal an escalation in attacks from cybercrime to cyberespionage. Stealing and exploiting privileged accounts is a central element of the kill chain in cyberattacks of all kinds, regardless of attacker origin" (Taft, 2015).[3] Still, researchers have not investigated whether system trespassers' access to varying levels of administrative privileges on the attacked system influences their online behaviors during the progression of a system trespassing event. Because system trespassers' privileges on an attacked computer system may indicate intruders' intentions and expected utility throughout a system trespassing event, we suspect that the effectiveness of deterring security policies and practices in mitigating the consequences of a system trespassing event will depend on system trespassers' imposed level of privileges on the system. The assumption that deterring cues in the environment may influence system trespassers' online behaviors during a system trespassing event draws on recent extension of the classical deterrence model.

Theoretical Background

Deterrence and Restrictive Deterrence

Deterrence theory traces its roots back more than two centuries to the work of Beccaria (1963 [1764]) and later to Bentham (1970 [1789]) who posited that individuals are rational actors who choose to engage in criminal conduct based on a hedonistic weighting of the potential costs and benefits of a criminal act. In particular, Beccaria (1963 [1764]) claimed that the severity, celerity, and certainty of sanctions would be influential in deterring would-be offenders from criminal activity. Since these early writings, a voluminous amount of literature has followed with researchers seeking to test and expand on the principles and propositions outlined by early deterrence philosophers (Nagin, 1998, 2013; Paternoster, 2010).

Notably, Gibbs (1975) elaborated on traditional deterrence theory and proposed a distinction between absolute and restrictive deterrence. Absolute deterrence, according to Gibbs, is the complete inhibition of criminal activity in response to a threat of sanction or the actual imposition of sanctions. Restrictive deterrence, by contrast, is the partial reduction of criminal conduct in light of deterring cues. According to Gibbs, restrictive deterrence occurs when an offender is not wholly deterred from engaging in criminal conduct, but instead, the offender adapts his or her behavior during the progression of the criminal event in a specific manner that will reduce the probability of detection and punishment. In the original theorizing of the concept, Gibbs focused primarily on the reduction in the frequency of criminal conduct as a unique restrictive deterrence strategy. Later extensions to Gibbs's concept by other researchers proposed that deterring cues could influence offenders' behavior in ways other than reduction in the frequency of criminal behavior (Jacobs, 1993, 2010). Specifically, Jacobs (2010) noted that in the presence of deterring cues (e.g., CCTV cameras and police officers), offenders will be more likely to (a) reduce the seriousness of their criminal activity (believing that punishment will not be as severe for more minor

[3] "Cyber Kill Chain" is a model developed by Lockheed Martin to describe the various stages of a cyberattack (see Hutchins, Cloppert, and Amin, 2011).

violations of the law), (b) take measures to decrease the likelihood of detection, and (c) change the location of the criminal event.

Interestingly, although the aim of extensive criminological research has been to examine the effect of absolute deterrence on crime (Nagin, 1998, 2013; Paternoster, 2010), limited applications of restrictive deterrence still exist. Nevertheless, some empirical evidence provides support for the Gibbs's (1975) and Jacobs's (2010) theoretical propositions in application to street-level drug dealers (Jacobs, 1993, 1996a, 1996b), auto-thieves (Jacobs and Cherbonneau, 2014), and marijuana cultivation (Nguyen, Malm, and Bouchard, 2015). For instance, the Jacobs (1996a) interviews with street-level crack dealers demonstrate that dealers are responsive to threats from law enforcement, but rather than completely curtail their offending behavior after the threat of sanctions, offenders actively change locations within neighborhood contexts to avoid apprehension.

Online Restrictive Deterrence

After acknowledging the extensive criminological literature that comprises studies aimed at theorizing and investigating the effect of deterrence practices in preventing and restricting individuals' involvement in off-line crimes (Gibbs, 1975), contemporary cybercrime scholars have debated the relevance of deterrence in preventing and reducing the scope of individuals' involvement in online crime. On the one hand, some scholars have suggested that traditional deterrence mechanisms will be ineffective in preventing and dissuading cybercriminals' illegitimate online behaviors (Andress and Winterfeld, 2011; Blank, 2001; Brenner, 2010; Harknett, 1996; Libicki, 2009). According to these scholars, as a result of the anonymous nature of cyberspace, and the difficulty involved in identifying and detecting online criminals, the certainty of an official punishment for online crimes (and system trespassers in particular) is minute. Consequently, offenders' perceived risk of detection is low (Geerken and Gove, 1975), and the perceived threat of sanctions is greatly diminished (Denning and Baugh, 1999; Harknett, 1996; Harknett, Callaghan, and Kauffman, 2010).

On the other hand, other scholars have suggested that it is not necessary to identify specific cyberoffenders for deterring cues to be effective in cyberspace (Goodman, 2010). Accordingly, the implementation of deterring cues in a computing environment may alter the risk perceptions of a system trespasser, and they may result in an online behavioral change that increases trespassers' exposure on the system. Similarly, the small probability of official punishment for online crime does not necessarily mean that the probability of unofficial punishment (e.g., hack-back or losing access to the target computers) is small as well (see Holzer and Lerums, 2016; Riofrio, 2013). Recent empirical evidence tends to show support for the view espousing potential restrictive influences of deterrence-based strategies on online criminals' illegitimate operations (Kigerl, 2015; Maimon et al., 2014; Stockman, Heile, and Rein, 2015; Wilson et al., 2015). Maimon et al. (2014) and Stockman and colleagues (2015) reported that system trespassers are responsive to sanction threats, showing that in line with restrictive deterrence theory, displaying a message issuing a warning to system trespassers can serve as a credible threat that does not lead to an immediate termination of a trespassing incident but does significantly reduce the duration of trespassing incidents. Similarly, Wilson et al., (2015) found that the presence of a surveillance banner reduces the probability of system trespassers to enter computer commands on the network during a first system trespassing event. Moreover, these scholars found that the probability of commands being entered on the attacked system in future system trespassing events is conditional on whether commands were typed in previous system trespassing events.

Although these works are important to our understanding of system trespassers' attempts to *avoid* detection in the presence of deterring cues, they still do not result in providing a link between the presence of deterring cues in the environment and offenders' willingness to escalate the severity of crimes they commit. By drawing on the results of these previous studies, we suspect that deterring situational stimuli in an attacked computer system could influence system trespassers' online behaviors on an attacked computer system during the progression of a system trespassing incident. Nevertheless, because the effect of deterrence tends to vary across offenders (Piquero, Paternoster, Pogarsky, and Loughran, 2011; Pogarsky, 2002; Thomas, Loughran, and Piquero, 2013) and situations (Maimon and Browning, 2012), we also examine whether system trespassers' differential access to criminal opportunities in the attacked computer system conditions their response to deterring cues.

Deterrence, Criminal Opportunities, and Criminal Offending

All in all, criminological theory often anticipates that varying degrees of access to resources and criminal opportunities may play an important role in conditioning the effect of deterrence on criminal behavior (Nagin, Solow, and Lum, 2015; Piliavin, Gartner, Thornton, and Matsueda, 1986). For instance, Piliavin and colleagues (1986: 116) noted that perceptions of the opportunities and returns of crime are unstable over time and vary across situations such that "assessments of risk are to some extent situationally-induced, transitory, and unstable." More recently, Nagin et al. (2015) contended that police can deter criminal activity by affecting the distribution of criminal opportunities available to would-be offenders and by reducing offenders' perceptions that criminal events can be successfully completed. Nevertheless, only scant criminological research has been aimed at investigating these relationships in either noncyber- or cyberenvironments. Moreover, no researchers to date have investigated the intersection of access to privileges, criminal opportunities, and situational deterring cues on the progression of a criminal event. Given a growing body of research that has produced results demonstrating rationality in decisions to offend in cyberspace (Maimon et al., 2014; Png and Wang, 2009; Rege, Ferrese, Biswas, and Bai, 2014; Wilson et al., 2015), we explore whether different administrative privileges taken by intruders during a system trespassing event result in consistent responses to sanction threats and dissuade system trespassers' navigation and manipulation of file permission on the attacked computer system.

Current Study

The use of sanction threats as a means to raise the perceived costs of criminal behavior has for centuries been the primary policy of the criminal justice system's efforts to deter would-be offenders (Nagin, 1998). Although questions remain regarding the effectiveness of criminal sanctions as a means to curtail adverse behavior (Paternoster, 2010; Pratt, Cullen, Blevins, Daigle, and Madensen, 2006), the use of sanction threats serves as a flexible policy applied to several behaviors ranging from traditional criminal offending (Nagin, 1998; Paternoster, 2010), to terrorism (Dugan and Chenoweth, 2012; LaFree, Dugan, and Korte, 2009), and recently, to malicious activity in cyberspace (Maimon et al., 2014).

Within the United States, recommendations regarding the applications of sanction threats as a means to prevent and mitigate criminal behavior in cyberspace are outlined in policy guidelines written by the National Institute for Standards and Technology (NIST, 2009). In part, NIST recommendations provide minimum-security guidelines for governmental, industrial, and private agencies, and they recommend that IT managers display an approved system-use notification to users in the form of a banner or message before granting access to the network. A banner is recommended to include the following information: (a) organizational policies regarding unauthorized access and use of the system and (b) the unauthorized use of the information system is prohibited and subject to criminal and civil penalties.[4]

We examined the effect of sanction threat in the form of a warning message on system trespassers' (a) navigation in an attacked computer system and (b) manipulation of file permission during the progression of a system trespassing event.[5] In drawing on the theoretical claims of Gibbs (1975) and Jacobs (2010), we tested whether the presence of a sanction threat installed in line with NIST's (2009) policy recommendations influences the potential harm caused by system trespassers after unlawfully accessing the computer network. Consistent with the results of prior research indicating that the use of a banner message can successfully reduce system trespassers' time (Maimon et al., 2014; Stockman et al., 2015) and the likelihood to enter computer commands in the system (Wilson et al., 2015), we hypothesized that *the presence of a sanction threat in an attacked computer system would dissuade system trespassers' navigation in the attacked system and discourage their manipulation of file permission during the progression of a system trespassing event.* Specifically, we suspected that in the presence of a warning banner, system trespassers would be more focused and task oriented during the progression of the criminal event and would be less likely to explore the attacked system extensively and to navigate back and forth between the

[4] There is no indication how prevalent the practice of deploying a warning banner is among universities, government agencies, or private companies as organizations typically do not advertise security practices.

[5] Progression of a system trespassing event is the outcome of the series of commands entered by a system trespasser in the attacked computer over time during a unique system trespassing event.

attacked system directories. Moreover, we suspected that in the presence of a warning banner, system trespassers would be reluctant to change file permissions to avoid an escalation of the system trespassing event (Jacobs, 2010).

Our second goal was to determine whether the effect of sanction threats on system trespassers' navigation in the system and manipulation of file permission varied based on system trespassers' level of administrative privileges on the system. Given the benefits associated with greater access and privileges on the network and the ability to remove traces easily from the attacked computer, there are two reasons to believe that administrative users will not be deterred by the presence of a warning banner on the system. First, accessing the system with administrative credentials that provide extra privileges on the network may be associated with high expected utility for carrying out a criminal act (e.g., monetary gain, revenge, and espionage). Therefore, it may be difficult to dissuade trespassers from extensively exploring (by navigating on the system) and changing file permissions on the attacked system throughout the progression of the system trespassing event when the potential benefits from a criminal event are heightened. Second, it is possible that those who access the system with administrative credentials may have high criminal efficacy. Brezina and Topalli (2012: 1059) noted that criminal self-efficacy refers to one's belief that he or she can perform a criminal task well enough that many offenders continue to engage in crime even when faced with the threat of sanctions "bolstered by confidence in their ability to succeed at crime and motivated by the possibility of greater success in the future." Related to this, the results of prior research suggest that offenders' optimism regarding their ability to be successful at crime may lead to an increase in offending behavior (Cherbonneau and Copes, 2006). In the case of system trespassing, it is possible that trespassers who infiltrate the system with the highest level of administrative privileges believe they can successfully complete their criminal task and avoid detection even when faced with the threat of sanctions. Indeed, system trespassers who obtain administrative credentials have many opportunities to engage in detection avoidance strategies such as removing files from the attacked system and cleaning computer event logs (i.e., special files that record a significant event on the computer such as when a user logs in) because they have the right privileges to do so (Provos, Friedl, and Honeyman, 2003; Shen, Chen, Li, and Liu, 2013; Snowberger and Thain, 2005). Opportunities to engage in detection avoidance practices may in turn reduce trespassers' perceived risk of detection and punishment (Jacobs, 1993, 1996a). Therefore, we hypothesized that *a warning banner may be ineffective in restricting the scope of navigation and file manipulation behaviors of system trespassers with administrative privileges on an attacked computer system.*

In contrast, system trespassers who infiltrate the attacked system with no administrative privileges face limited opportunities and few potential benefits on an attacked computer system. The motivation of these system trespassers might be different than the motivation of system trespassers with administrative privileges and include exploration and risk seeking only. Thus, we propose that for system trespassers with few criminal opportunities and less potential for substantial rewards, the marginal increase in costs resulting from the presence of a sanction threat should have a stronger effect on dissuading navigation back and forth between directories and changing file permissions during the progression of the system trespassing event. Moreover, it is also plausible that those who do not obtain access with administrative credentials may have low criminal efficacy because they are not able to cover their traces on the attacked computer in an effective manner (system trespassers with nonadministrative privileges cannot change the attacked computer event log files). Therefore, our third research hypothesis suggests that *system trespassers with nonadministrative privileges were likely to restrict the scope of their navigation on the attacked system and avoid manipulation of file access after being exposed to a banner.*

Data and Method

Design

To test our research hypotheses, we used data collected from a randomized experiment conducted at a large American university. The experiment used 300 public Internet Protocol (IP) addresses provided by the university information technology team.[6] These IP addresses were used for deployment on high-interaction honeypot computers. Overall, honeypot computers are target computers that are used as a decoy for security or research purposes. These target computers are set to allow system trespassers with access to these systems and to collect information about intruders' behavior during a real system trespassing event (Spitzner, 2002). In this study, we deployed "high-interaction honeypots" that allowed system trespassers to interact with the application or services of the network in real time. Low-interaction honeypots, on the other hand, allow only limited interaction between the system trespasser and the computer system (see Riden and Seifert, 2008; Zhuge et al., 2007). The deployed honeypots (which will be identified as target computers henceforth) were set up using the Linux Ubuntu 10.04 operating system (Canonical Group Limited, London, U.K.), and they were designed to imitate the activity of the regular university computer.

Procedure

In this experiment, we did not actively recruit subjects; rather, target computers (i.e., targets of opportunities) were deployed on the university network for a period of 6 months (from October 2011 to April 2012).[7] During this time, the research team waited for system

[6] An IP address is a unique identification number (similar to a telephone number), which a computer system uses to connect to the Internet and to correspond with other systems. Each IP address is a unique combination of four numbers arranged from 1 to 256 (e.g., 111.11.111.11).

[7] For security reasons, the university technology team prohibits the disclosure of specific details regarding our server capacity and capabilities, the number of nodes in the network, and data size and transmissions. In general, the servers we used in the current study were deployed on a large public university network. The network provides services to more than 35,000 students, greater than 9,000 faculty and staff, and an alumni network of more than 300,000 living individuals.

trespassers to identify the target computers, scan and identify vulnerable entry points through which they can infiltrate the system, and use specialized toolkits that enable the intruder to gain access to the network by guessing the correct combination of legitimate users' passwords and usernames (McQuade, 2006).[8] All in all, when accessing a target system, prospective trespassers can attempt to crack the combination of username and password that is associated with either *administrative* or *nonadministrative* users on the system. Consistent with these two possible ways to gain an illegitimate access to a Linux computer, we allowed system trespassers to infiltrate our systems either as administrative or as nonadministrative users. To simulate a genuine computer network, the target computers were designed to reject login attempts by system trespassers until reaching a predefined threshold between 150 and 200 attempts.[9] A login attempt on the *n*th attempt was treated as correctly cracking legitimate login credentials, and the trespasser was allowed into the system (either as administrative or nonadministrative user) and was able to access the system at a later time using the same credentials.[10] Each time a user logged back into the system with the same credentials, it was recorded as a new system trespassing event.

After successfully cracking into the system either as administrative or as nonadministrative users, trespassers were randomly assigned to 1 of 16 conditions.[11] To account for technical attributes of the computing environment, we employed a 2 [warning banner, no banner (control)] × 2 [low (512 MB) RAM, high (2.25 GB) RAM] × 2 [low (128 kb/s) bandwidth, high (512 kb/s) bandwidth] × 2 [low (5 GB) disk space, high (30 GB) disk

[8] Toolkits are software applications that permit system trespassers to select a username to be used to attempt to gain access to the system.

[9] The number of attempts limits the ability of human users to break into the system by arbitrarily typing login credentials by hand. Therefore, it prioritizes the use of brute force toolkits by system trespassers. Most brute force toolkits do not report the number of trials needed to gain access to the system. Accordingly, whether the threshold was set at 150 or 1,500 attempts should not have an impact on the sample in the current study.

[10] It is important to note that there are several ways in which a system trespasser can gain access to administrative account privileges by using brute force methods. For instance, trespassers may use technical or sophisticated strategies such as phishing attacks, malware, or automated tools to exploit vulnerable systems. Additionally, a less common approach in the current experimental study was for system trespassers to obtain administrative credentials to log on as a nonadministrative user by using a username and password combination associated with a nonadministrative account and attempting to gain access as an administrative user by using Linux commands such as *su*, *kdesu*, and *sudo*. Achieving administrative status through this path was rare in our experiment (observed in less than 2% of the target computers). Accordingly, in the current study, most trespassers gained administrative privileges by cracking username and password credentials associated with these accounts. Nevertheless, this is only one of several approaches that could be employed by system trespassers to gain access to such accounts.

[11] Like many computers on the university network intended for public use, the target computers deployed in this experiment did not have personal files contained on them. Rather, the only files on the target computer were system and configuration files. Although no additional files were added to the target computers, system trespassers could have downloaded or installed files onto our systems.

space] factorial design.[12] Those assigned to the warning group were exposed to the following message appearing on the computer screen immediately after successfully breaking into the system:

> *The actual or attempted unauthorized access, use, or modification of this system is strictly prohibited. Unauthorized users are subject to institutional disciplinary proceedings and/or criminal and civil penalties under state, federal, or other applicable domestic and foreign laws. The use of this system is monitored and recorded for administrative and security reasons. Anyone accessing this system expressly consents to such monitoring and is advised that if monitoring reveals possible evidence of criminal activity, the Institution may provide the evidence of such activity to law enforcement officials.*

Those assigned to the control group did not receive any message on their computer screen when logging onto the computer.

After successfully accessing the target computer, system trespassers were given access to the system for 30 days; at which point, they could freely log in and log out of the system, share access to the target computers with others, and use it to attack other computer systems. During the time in which a trespasser used the system, various information about the user's activity was collected with special software (Sebek keylogger). The breadth of data collection provided a detailed record of system trespassers' online behavior through an investigation of the specific commands associated with the Linux operating system entered by system trespassers on the target computer. To ensure that trespassers did not cause serious harm to other computer networks or systems, we employed network flow tools (NFDUMP) that provided constant surveillance of the activity transferred from the target computers to the Internet. After 30 days, the trespassers were removed from the system; target computers were cleaned and redeployed onto the network. Thus, to regain access and initiate subsequent trespassing events on our target computers, the system trespasser had to break into the system again.

During the 6 months of the experimental period, 502 target computers were infiltrated by system trespassers; of those, 221 target computers were used by system trespassers to enter commands. (See the Appendix for a distribution of the number of computers infiltrated and used by system trespassers per the 16 experimental conditions.) As we have focused our work on the progression of a system trespassing event and on analysis of computer commands entered into the attacked system, we examined data collected by target computers that

[12] We used this factorial design to understand whether certain components of the computing system (i.e., high RAM or high bandwidth) would be more attractive to certain system trespassers when this experimental study was being designed. These technical components did not carry direct or indirect effects on our dependent variables.

had computer commands recorded.[13] Moreover, because an examination of all commands entered on the target computers indicated that none of the trespassers entered the specific commands necessary to unveil details of the three technical components (i.e., RAM size, disk space, and bandwidth), we only tested for the effect of the warning banner in this work.[14]

Overall, most target computers (79.1%) were infiltrated by system trespassers who employed administrative credentials to the system.[15] Because random assignment to groups was done after cracking into the system, the proportions of target computers with warning banners that were infiltrated by administrative and nonadministrative system trespassers were similar. Specifically, although 49.7% of the target computers that were infiltrated by system trespassers with administrative credentials had a banner installed (87 computers), 52.2% of the target computers that were infiltrated by system trespassers with nonadministrative credentials had a warning banner installed (24 computers). A *t* test for the differences in the proportion of attacked computers with the banner condition across administrative and nonadministrative system trespassers revealed no statistically significant difference between the groups. Similarly, no statistical difference was found between the proportions of system trespassing events that were launched from target computers with a warning banner installed.

Overall, the 221 target computers yielded 553 unique system trespassing events with at least one computer command. A total of 415 of the trespassing events were initiated by system trespassers who employed administrative credentials on the system, whereas the other 138 trespassing events were launched by nonadministrative system trespassers. Approximately half of the system trespassing events initiated by administrative and non-administrative system trespassers were launched from target computers that had a warning banner installed (i.e., 48.7% and 49.3% of all system trespassing events initiated by these groups, respectively).

Outcome Measures

In this study, we focused on two key online behaviors commonly exhibited by system trespassers during the progression of a system trespassing event: navigation on the attacked computer system and changing file permission. Navigation on Linux computers is usually

[13] In analysis not shown, we reran our analyses and included attacked computers that did not record a keystroke as well ($n = 502$). The results from these analyses yielded similar results to those reported in the article.

[14] Additional analyses including the technical component (i.e., RAM size, disk space, and bandwidth) indicated no significant effects of these measures on our list of dependent measures. Furthermore, none of the interaction terms between the technical configuration of the system and warning on our list of dependent variables were statistically significant.

[15] A supplemental analysis comprising a sample including attacked computers that did not record a keystroke ($n = 502$) similarly revealed that most target computers (74.98%) were infiltrated by system trespassers who employed administrative credentials to the system.

performed through the typing of three key navigation commands: change directory, list files, and print working directory.[16] The change directory command (cd) is used to change the directory the user is currently working on to navigate through the target computer system and reach a desired point during the progression of a system trespassing event. Computer users apply the change directory command followed by a pathname to specify the specific directory in which they seek to access. For example, if a user seeks to advance to a file contained in the "/user/folder1" directory, a user can access the desired directory using the "cd/user/folder1" syntax. Similarly, the list files (ls) command is used to list information visually about the files and contents that are available on the directory that a user is currently working on. Accordingly, to generate information about the files and contents available on the current directory, a user can type the "ls" command to call forth a tabular listing of all files available and use this information to decide how to navigate through the system. Finally, the print working directory (pwd) command is used to report the full pathname of a working directory, which is used for storing files. For instance, if a user has advanced to a file contained in the "/user/folder1" directory, a trespasser can confirm his or her location by calling forth information on the pathname of his or her current location with the "pwd" command.

The change file permission command (chmod) is used to change the permissions of files, which define the way in which a particular file can be accessed. File permissions are ways to protect files by designating which users may access particular files and what level of access a given user has to the files. For instance, a particular file may be designated for access only by the owner of the file or by multiple members of a designated group. Similarly, files are designated with particular levels of access such as "read-only," where users can read but not manipulate a file. The change file permission command can change the classifications of permissions and access to a given file to give unauthorized users access to a file or change file permissions in such a way to allow the file to be altered by a user. Both administrative and nonadministrative users can change file permission on the system. Nevertheless, although administrative users can change file permissions of all files on the system, nonadministrative users can change the file permission of files they created and own only.

To capture the effect of a warning banner on the presence of navigation and change file access commands in the attacked computer system, we estimated the proportion of target computers and system trespassing events in which these commands were recorded. For the navigation commands, we also estimated the proportion of target computers recording at least one of the three navigation commands on the system. To tap the volume of navigation and change file access commands on the target computers, we constructed a command rate measure by counting the number of times each command was entered on the

[16] We restrict the analysis to the commands "cd," "ls," and "pwd" as these are the most commonly used navigation commands. Any other commands that fit our definition of "navigation" commands, such as "less," appeared too infrequently to be included in the analysis.

system, dividing by the overall number of system trespassing events initiated from the target computer, and then multiplying the outcome by three.[17] For the navigation commands, we also calculated an overall navigation commands rate including the three navigation commands. Finally, to capture the frequency in which navigation and change access files commands were entered during a system trespassing event, we counted the number of times each command was recorded during the event. For the navigation commands, we also calculated the overall frequency of the navigation by adding the overall number of the three navigation commands recorded during a system trespassing event.

Table 1 presents descriptive statistics of both the proportions of target computers and the system trespassing events with "navigation" and "change file permission" commands recorded, as well as the frequency in which these commands were recorded by system trespassers. As indicated in the table, navigation commands were commonly entered in most target computers (85%) and system trespassing events (77%). Moreover, the average rate (13.35 navigation commands per three system trespassing events) and frequency (4.53 commands per three system trespassing events) of navigation commands entered on the target computers were high. Nevertheless, although target computers and system trespassing events with "change directory" and "list files" commands were common, the "print working directory" command was recorded on few of the target computers and system trespassing events, and in a lower frequency. In contrast to the pattern revealed for the navigation commands, 43% of the target computers and 23% of the system trespassing events recorded the change file permission command. Both the average rate of the "change file permission" command on an attacked computer system (1.30 per three system trespassing events) and frequency of appearance during the progression of system trespassing event (0.37) were consistently low.

Importantly, Table 1 also presents the proportions and frequencies of navigation and change file permission commands recorded on computers infiltrated and system trespassing events initiated by administrative and nonadministrative system trespassers. Note that with the exception of one command (list files), there were no significant differences between the proportions and frequencies of navigation and change file permission commands recorded on computer systems attacked and system trespassing initiated by system trespassers with administrative privileges and system trespassers with nonadministrative privileges.

Results

Target Computer

We begin by exploring our first research hypothesis using data at the target computer level, and we estimate the differences between the proportion of warning and no warning target

[17] Although the average number of system trespassing events with commands recorded on our system is 2.5 system trespassing events, we opt to construct a command rate per 3.0 system trespassing events for ease of interpretation.

TABLE 1

Descriptive Statistics of Target Computers and System Trespassing Events With Commands Entered On by System Trespassers' Administrative Privileges

	Target Computers			System Trespassing Events		
	Proportion of Computers With Commands			Proportion of Events With Commands		
Command	**Full Sample ($N = 221$)**	**Infiltrated by Administrative Trespassers ($n = 175$)**	**Infiltrated by Nonadministrative Trespassers ($n = 46$)**	**Full Sample ($N = 553$)**	**Initiated by Administrative Trespassers ($n = 415$)**	**Initiated by Nonadministrative Trespassers ($n = 138$)**
Navigation Commands						
Change directory	0.80	0.82	0.72	0.67	0.65	0.70
	(0.40)	(0.39)	(0.46)	(0.47)	(0.48)	(0.46)
List files	0.73	0.72	0.76	0.61	0.58	0.72^{***}
	(0.44)	(0.45)	(0.43)	(0.49)	(0.49)	(0.45)
Print working directory	0.11	0.13	0.04	0.05	0.06	0.02
	(0.32)	(0.34)	(0.21)	(0.22)	(0.24)	(0.15)
Navigation composite	0.85	0.86	0.80	0.77	0.75	0.82
	(0.36)	(0.34)	(0.40)	(0.42)	(0.43)	(0.39)
Manipulate File Access Command						
Change file permission	0.43	0.45	0.37	0.23	0.24	0.21
	(0.50)	(0.50)	(0.49)	(0.42)	(0.43)	(0.41)

(Continued)

TABLE 1

Continued

Command	Average Rate of Commands Entered on Target Computers per 3 System Trespassing Events			Average Number of Commands Entered During A System Trespassing Event		
	Full Sample ($N = 221$)	Infiltrated by Administrative Trespassers ($n = 175$)	Infiltrated by Non-Administrative Trespassers ($n = 46$)	Full Sample ($N = 553$)	Initiated by Administrative Trespassers ($n = 415$)	Imitated by Nonadministrative Trespassers ($n = 138$)
Navigation Commands						
Change directory	6.87	6.86	6.91	2.20	2.24	2.08
	(7.33)	(6.80)	(9.15)	(2.74)	(2.81)	(2.54)
List files	6.14	5.92	6.94	2.24	2.29	2.07
	(7.16)	(7.17)	(7.11)	(3.56)	(3.87)	(2.40)
Print working directory	0.34	0.39	0.14	0.09	0.10	0.06
	(1.48)	(1.60)	(0.88)	(0.48)	(0.48)	(0.46)
Navigation composite	13.35	13.17	13.99	4.53	4.63	4.21
	(13.93)	(13.39)	(15.99)	(5.92)	(6.28)	(4.71)
Manipulate File Access Command						
Change file permission	1.30	1.25	1.48	0.37	0.38	0.36
	(2.62)	(2.28)	(3.68)	(0.90)	(0.90)	(0.93)

Note. Standard deviation in parentheses.
***p < .01.

computers with navigation and change file permission commands recorded. Table 2, Panel A presents our findings from a *t* test for the difference between two group proportions. Contrary to our first hypothesis, the presence of a warning banner on an attacked computer system had no statistically significant effect on the probability of either navigation or change file permission commands to be entered on the system. Specifically, although 85% of the no-warning target computers recorded at least one navigation command, 86% of the warning computers recorded at least one navigation command as well ($p > .10$). Moreover, although the proportion of warning target computers recording "change file permission" commands (47%) appeared to be slightly higher than the no-warning computers recording the same command (40%), this difference was statistically nonsignificant ($p > .10$). Investigation of the differences between the average rate of navigation and change file permission commands per three system trespassing events recorded on the warning and no-warning computers revealed similar findings. Specifically, results from a *t* test for the difference between two groups' means indicated that the differences between the rates of either navigation or change file permission commands entered on warning and no-warning computers were statistically nonsignificant ($p > .10$).

In turning to our second research hypothesis, we next examined whether a warning banner in an attacked computer system influenced the presence and volume of navigation and change file permission commands on target computers attacked by system trespassers who took over administrative privileges on the system. The findings from these analyses are presented in the right columns of Table 2, Panel A. As indicated in the table, the presence of a warning banner did not affect either the probability or the average rate of navigation commands to be entered on computers infiltrated by system trespassers with administrative privileges. Nevertheless, the presence of a warning banner significantly increased the proportion of target computers with "change file permission" commands recorded. Specifically, we found that the proportion of warning target computers attacked by an administrative system trespasser who entered a "change file permission" command on the attacked system (52%) was significantly higher than the proportion of no-warning target computers attacked by an administrative system trespasser who entered a "change file permission" command (39%; $p < .05$). Moreover, the average rate of "change file permission" command entered on warning target computers (1.69 commands per three system trespassing events) was significantly higher than the average rate of "change file permission" command on no-warning computers (0.83 commands per three system trespassing events; $p < .05$). Thus, in support of our second hypothesis, the results suggested that system trespassers with full administrative privileges were less likely to restrict the scope of navigation and change file permission commands after being exposed to a warning banner.

Finally, we investigated the effect of the warning banner on the probability and volume of navigation and change file permission commands to be entered on computer systems attacked by system trespassers with nonadministrative privileges on the system. The findings from these analyses are presented in Table 2, Panel B. In support

TABLE 2

Proportion of Target Computers With Commands and Average Rate of Commands Per *X* System Trespassing Events Recorded on Target Computers by System Trespassers' Administrative Privileges and Experimental Conditions

Panel A. Target Computers Infiltrated by Any System Trespassers and by Administrative System Trespassers

	Full Sample ($N = 221$ computers)				**Infiltrated by Administrative Trespassers ($N = 175$ computers)**			
	Proportion of Target Computers		**Average Rate of Commands per 3 System Trespassing Events**		**Proportion of Target Computers**		**Average Rate of Commands per 3 System Trespassing Events**	
Command	**No Banner ($n = 110$)**	**Banner ($n = 111$)**	**No Banner ($n = 110$)**	**Banner ($n = 111$)**	**No Banner ($n = 88$)**	**Banner ($n = 87$)**	**No Banner ($n = 88$)**	**Banner ($n = 87$)**
Navigation Commands								
Change directory	0.81	0.78	7.09	6.65	0.80	0.84	6.47	7.26
	(0.04)	(0.04)	(0.76)	(0.64)	(0.04)	(0.04)	(0.72)	(0.73)
List files	0.72	0.75	6.38	5.88	0.70	0.75	5.80	6.05
	(0.04)	(0.04)	(0.75)	(0.60)	(0.05)	(0.05)	(0.82)	(0.71)
Print working directory	0.13	0.10	0.48	0.19*	0.16	0.10	0.60	0.17*
	(0.03)	(0.03)	(0.18)	(0.07)	(0.04)	(0.03)	(0.23)	(0.07)
Navigation composite	0.85	0.86	13.97	12.73	0.84	0.89	12.88	13.48
	(0.03)	(0.03)	(1.48)	(1.15)	(0.04)	(0.03)	(1.53)	(1.32)
Manipulate File Access Command								
Change file permission	0.40	0.47	1.13	1.47	0.39	0.52**	0.83	1.69**
	(0.05)	(0.05)	(0.24)	(0.26)	(0.05)	(0.05)	(0.15)	(0.30)

(Continued)

TABLE 2

Continued

Panel B. Target Computers Infiltrated by Nonadministrative System Trespassers ($N = 46$ Computers)

	Proportion of Target Computers		Average Rate of Commands per 3 System Trespassing Events	
Command	**No Banner ($n = 22$)**	**Banner ($n = 24$)**	**No Banner ($n = 22$)**	**Banner ($n = 24$)**
Navigation Commands				
Change directory	0.86	0.58^{**}	9.58	4.47^{**}
	(0.07)	(0.10)	(2.42)	(1.18)
List files	0.77	0.75	8.74	5.29^{*}
	(0.09)	(0.09)	(1.83)	(1.04)
Print working directory	0.00	0.08	0.00	0.26
	(0.00)	(0.06)	(0.00)	(0.25)
Navigation composite	0.86	0.75	18.31	10.03^{**}
	(0.07)	(0.09)	(4.15)	(2.24)
Manipulate File Access Command				
Change file permission	0.46	0.29	2.35	0.69^{*}
	(0.11)	(0.09)	(1.03)	(0.39)

Note. Standard deviation in parentheses.
$^{*}p < .05.$ $^{**}p < .01.$ $^{***}p < .001.$

of our third hypothesis, the findings suggested that a warning banner substantially reduced the use of both navigation and change file permission commands on computers attacked by system trespassers with nonadministrative privileges. By beginning with navigation commands, we found that although 86% of the no-warning computers attacked by system trespassers with no administrative privileges recorded change directory commands, 58% of the warning computers recorded that command ($p < .05$). Moreover, although 86% of the no-warning computers attacked by system trespassers with no administrative privileges recorded at least one navigate command on the system, 75% of the warning computers recorded that command. Indeed, although the difference for the composite navigation command was statistically nonsignificant, it still demonstrated an important trend ($p > .10$). Investigation of the differences between the average rates of navigation commands recorded on warning and no-warning computers revealed further that the average rate of "change directory" command entered on warning target computers (4.47 commands per three system trespassing events) was significantly lower than the average rate of "change directory" command on no-warning computers (9.58 commands per three system trespassing events; $p < .05$). Furthermore, the average rate of the unique "list files" command was significantly lower on warning computers (5.29 commands per three system trespassing events) relative to that on no-warning computers (8.74 commands per three system trespassing events; $p < .10$). Finally, the overall rate of navigation commands was significantly lower on warning (10.03 commands per three system trespassing events) than on no-warning target computers (18.31 commands per three system trespassing events) attacked by nonadministrative system trespassers ($p < .05$).

Consistent with the lower proportions and rates of navigation commands entered in the systems attacked by nonadministrative system trespassers, we also found that a warning banner significantly reduced the presence and rate of the "change file permission" command to be entered on computers attacked by nonadministrative system trespassers. Specifically, we found that the proportion of warning target computers that recorded a change file permission command (29%) was lower than the proportion of target computers with no warning that recorded that command (46%; $p > .10$). Although the difference in the proportion of warning and no-warning target computers that recorded the change file permission command was statistically nonsignificant, the rate in which this command was entered in the two types of computer was. Specifically, although the average rate of "change file permission" command recorded on warning computer attacked by nonadministrator system trespassers was 0.69 commands per three system trespassing events, the average rate of this command on no-warning computers attacked by nonadministrator system trespassers was 2.35 commands per three system trespassing events ($p < .10$).

System Trespassing Event

To assess the robustness of our findings, as well as to investigate whether these findings held at the single system-trespassing event level, we reran our analyses and tested for significant differences between the probability and volume of navigation and change file permission commands entered by system trespassers during a system trespassing event. By starting again with examination of our first research hypothesis, we estimated the differences between the proportions of system trespassing events initiated on warning and no-warning target computers and during which navigation and change file permission commands were recorded. Table 3, Panel A presents findings from a *t* test for the difference between the two groups' proportions. As indicated in the left columns of Panel A, the presence of a warning banner on an attacked computer system had no statistically significant effect on the probability of navigation commands being entered during a system trespassing event. Specifically, although 76% of the system trespassing events initiated on no-warning target computers recorded at least one navigation command, 77% of the system trespassing events initiated on warning target computers recorded at least one navigation command as well ($p > .10$). Investigation of the differences between the average frequency of navigation and change file permission commands recorded during system trespassing events initiated on warning and no-warning computers revealed similar findings. Specifically, the difference between the frequency of navigation commands entered during a system trespassing event initiated from warning and no-warning computers was statistically nonsignificant. In contrast, the proportion of system trespassing events recorded on warning target computers and entering "change file permission" command (26%) was significantly higher than the proportion of system trespassing events recorded on no-warning target computers and entering "change file permission" command (20%; $p < .10$). Similarly, the differences between the average frequency of "change file permission" command recorded during system trespassing events initiated on warning computers (0.44 commands) was significantly higher than the frequency recorded on no-warning computers (0.31 commands; $p < .10$).

Next, we examined whether a warning banner in an attacked computer system influenced the presence and volume of navigation and change file permission commands to be entered by administrative system trespassers during a system trespassing event. The findings from these analyses are presented in the right columns of Table 3, Panel A. As indicated in the table, the presence of a warning banner did not affect either the probability or the frequency of navigation commands to be entered during a system trespassing event initiated by system trespassers with administrative privileges on the attacked system. Nevertheless, the presence of a warning banner significantly increased the proportion of system trespassing events in which the "change file permission" command was recorded and which initiated by administrative system trespassers. Specifically, we found that the proportion of system trespassing events in which the "change file permission" command was recorded and which was initiated by administrative system trespassers over warning target computers (29%) was significantly higher than the proportion of system trespassing events in which the "change file permission" command was recorded and which was initiated by administrative system trespassers over no-warning target computers (19%; $p < .01$). Moreover, the frequency in which the "change file permission" command was entered during a system trespassing event initiated by administrative system trespassers over warning

TABLE 3

Proportion of System Trespassing Events With Commands and Average Number of Commands Entered During a System Trespassing Event By System Trespassers Administrative Privileges and Experimental Conditions

Panel A. System Trespassing Event Initiated by Any System Trespassers and by Administrative System Trespassers

	Full Sample ($N = 553$ Events)				Initiated by Administrative System Trespassers ($N = 415$ Events)			
	Proportion of Trespassing Events with Commands		Average Number of Commands Entered During an Event		Proportion of Trespassing Events with Commands		Average Number of Commands Entered During an Event	
Command	**No Banner ($n = 283$)**	**Banner ($n = 270$)**	**No Banner ($n = 283$)**	**Banner ($n = 270$)**	**No Banner ($n = 213$)**	**Banner ($n = 202$)**	**No Banner ($n = 213$)**	**Banner ($n = 202$)**
Navigation Commands								
Change directory	0.68	0.65	2.22	2.18	0.64	0.66	2.08	2.42
	(0.03)	(0.03)	(0.17)	(0.16)	(0.03)	(0.03)	(0.19)	(0.20)
List files	0.59	0.64	2.16	2.31	0.54	0.61	2.11	2.48
	(0.03)	(0.03)	(0.23)	(0.20)	(0.03)	(0.03)	(0.28)	(0.25)
Print working directory	0.05	(0.05)	0.11	0.07	0.07	0.05	0.12	0.07
	(0.01)	(0.01)	(0.03)	(0.02)	(0.02)	(0.02)	(0.04)	(0.03)
Navigation composite	0.76	0.77	4.49	4.55	0.73	0.77	4.31	4.97
	(0.03)	(0.03)	(0.37)	(0.34)	(0.03)	(0.03)	(0.44)	(0.43)
Manipulate File Access Command								
Change file permission	0.20	0.26^{*}	0.31	0.44^{*}	0.19	0.29^{***}	0.25	0.51^{***}
	(0.02)	(0.03)	(0.05)	(0.06)	(0.03)	(0.03)	(0.05)	(0.08)

(Continued)

TABLE 3

Continued

Panel B. System Trespassing Events Initiated by Nonadministrative System Trespassers ($N = 138$ Events)

	Proportion of System Trespassing Events with Commands		Average Number of Commands Entered During an Event	
Command	**No Banner ($n = 70$)**	**Banner ($n = 68$)**	**No Banner ($n = 70$)**	**Banner ($n = 68$)**
Navigation Commands				
Change directory	0.80	0.60^{***}	2.66	1.49^{***}
	(0.05)	(0.06)	(0.36)	(0.21)
List files	0.73	0.72	2.31	1.82
	(0.05)	(0.05)	(0.34)	(0.21)
Print working directory	0.01	0.03	0.07	0.04
	(0.01)	(0.02)	(0.07)	(0.03)
Navigation composite	0.87	0.76	5.04	3.35^{*}
	(0.04)	(0.05)	(0.67)	(0.41)
Manipulate File Access Command				
Change file permission	0.24	0.17	0.49	0.24^{*}
	(0.05)	(0.05)	(0.14)	(0.07)

Note. Standard deviation in parentheses.
$^{*}p < .05$. $^{**}p < .01$. $p^{***} < .001$.

computers (0.51 commands) was significantly higher than the average frequency in which the "change file permission" command was entered during a system trespassing event initiated by administrative system trespassers over no-warning computers (0.25 commands; $p < .01$).

Finally, we investigated the effect of the warning banner in influencing the probability and frequency in which system trespassers with nonadministrative privileges on the system entered navigation and change file permission commands during a system trespassing event. The findings from these analyses are presented in Table 3, Panel B. Consistent with the pattern observed at the target computer level, our analyses indicated that a warning banner substantially reduced the use of both navigation and change file permission commands by system trespassers with nonadministrative privileges during the progression of a system trespassing event. By beginning with navigation commands, we found that although 80% of the system trespassing events launched from no-warning computers by system trespassers with no administrative privileges recorded "change directory" commands, 60% of the system trespassing events launched from warning computers recorded that command ($p < .01$). Moreover, although 87% of the system trespassing events were launched from no-warning computers by system trespassers with no administrative privileges, 76% of the system trespassing events launched from warning computers recorded that command ($p > .10$). Investigation of the differences between the frequency of navigation commands recorded during system trespassing events on warning and no-warning computers revealed further that the average frequency of "change directory" command entered during a system trespassing event launched from a warning target computer (1.49 commands) was significantly lower than the average frequency of "change directory" command entered during a system trespassing event launched from no-warning computers (2.66 commands; $p < .01$). Furthermore, the average frequency of all navigation commands entered during a system trespassing event launched from a warning target computer (3.35 commands per event) was significantly lower than the average frequency of all navigation commands entered during a system trespassing event launched from no-warning target computers attacked by nonadministrative system trespassers (5.04 commands per event on average; $p < .10$).

Similarly, we also found that a warning banner significantly reduced the presence and frequency of the "change file permission" command on computers attacked by nonadministrative system trespassers. Specifically, we found that the proportion of system trespassing events that recorded a change file permission command and that were launched from warning target computers (17%) was lower than the proportion of system trespassing events that recorded a change file permission command and that were launched from no-warning target computers (24%; $p > .10$). Although this difference was statistically nonsignificant, the frequency in which this command was entered in system trespassing events launched from the two types of computer was. Specifically, although the average frequency of "change file permission" command recorded during a system trespassing event launched from a warning computer attacked by nonadministrator system trespassers was 0.24 commands per event, the average frequency in which this command was entered on system trespassing events launched from no-warning computers attacked by nonadministrator system trespassers was 0.49 commands per event ($p < .10$).

Discussion

The emergence of cybercrime as a preeminent social issue has resulted in mounting attention being devoted toward developing feasible policies that mitigate the potential damage done by cyberattacks. Although researchers have begun to assess the effectiveness of cyberdeterrence policies in reducing the harm done by system trespassers, no previous study has been conducted to investigate whether such policies influence system trespassers' specific online behaviors during the progression of a system trespassing event, and whether these policies have a consistent effect across the entire population of system trespassers. By drawing on the restrictive deterrence perspective (Gibbs, 1975; Jacobs, 2010), through the current study, we asked whether sanction threats in the form of a warning banner installed on an attacked computer could restrict system trespassers' use of navigation and change file permission commands during the progression of a system trespassing event. Moreover, we explored whether the effect of a warning banner varied by level of administrative privileges imposed on the attacked computers by system trespasser during a system trespassing event. The findings from a randomized field experiment yielded several important insights regarding the effectiveness of a warning banner in restricting the scope and reducing the seriousness of a system trespassing event.

First, given the benefits associated with infiltrating computer systems with administrative credentials, we found that most system trespassers who attacked our computer system gained access to the system with administrative privileges. This finding supports Taft's (2015) observations regarding the attractiveness of privileged account credentials to system trespassers. Still, approximately one fourth of our target computers were infiltrated by system trespassers who took nonadministrative privileges in the system. It is possible that the minority of trespassers who accessed the system without administrative credentials were less skilled trespassers or had a different set of objectives that did not require the access to administrative credentials. Indeed the findings from past research have suggested that the range of skills evident in hackers and hacker communities varies substantially and is connected to the attack capability of individual hackers and hacker groups (see Brenner, 2010; Denning, 2012; Holt and Kilger, 2012; Jordan and Taylor, 1998). Based on the techniques of system trespassers observed in the current study, we can only speculate regarding the skill levels of individual system trespassers and the variation in trespassers' skill levels across the sample (Andress and Winterfeld, 2011; Denning, 2012; Holt and Bossler, 2016; Holt and Kilger, 2012). Still, investigating the relationship between system trespassers' skill level and the techniques used and responsiveness to deterring stimuli in cyberspace is a crucial area of future research.

Second, we found that the presence of a sanction threat in an attacked computer system did not restrict the overall scope of system trespassers' navigation in the attacked computer system. In contrast to our first research hypothesis and predictions from deterrence theory, we found preliminary evidence that the presence of a banner had no consistent effect on the probability, frequency, or rate of navigation commands entered into the attacked computer during the progression of a system trespassing event. Moreover, our analysis of the impact of a warning banner on the online behavior of trespassers with administrative access found that *within* administrative users, the use of the change file permission command increased after the presence of a warning banner. These results ran counter to the expectations from deterrence theory and suggested that the presence of sanction threats in an attacked computer system escalated the manipulation of file permission during the progression of a system trespassing event. Given the different levels of administrative privileges and opportunities available to trespassers who break into the computer system, it was possible that there was heterogeneity in the response to sanction threats and the deterrent effect of a warning banner may be consequential for a subgroup of the system trespasser population (Na, Loughran, and Paternoster, 2015; Piquero et al., 2011; Pogarsky, 2002; Thomas et al., 2013).

In light of this possibility, our second hypothesis suggested that a warning banner would have a muted effect on trespassers with access to full administrative privileges on the network. In line with this hypothesis, we observed that the presence of a sanction threat in an attacked computer system did not reduce the probability and rate in which system trespassers with full administrative privileges entered navigation commands on the attacked system and during a system trespassing event. Moreover, we found that the probability and frequency of "change file permission" commands were higher on warning target computers attacked by administrative system trespassers. Although it remains unclear exactly why system trespassers with full administrative privileges increased their use of the "change file permission" command on target computers that displayed a warning banner, we offer two possible explanations. First, trespassers who successfully obtain administrative credentials may have high criminal self-efficacy and are confident in their ability to avoid detection and successfully progress through a criminal event even after exposure to a sanction threat (Brezina and Topalli, 2012). After the imposition of a sanction threat, this group may be inclined to escalate rather than to restrict the scope of their attack as a result of being overly optimistic in their ability to avoid detection (Cherbonneau and Copes, 2006). Alternatively, by drawing on defiance theory (Sherman, 1993), it is possible that after obtaining administrative credentials, this group of system trespassers may be more likely to act defiantly in response to sanction threats that threaten to remove the highest level of access and privileges from them, and in response to such a threat, administrative trespassers may escalate their offending in response to a sanction threat perceived as illegitimate.

Although further research is needed to uncover the exact mechanisms that drive administrative system trespassers to increase their malicious activity in the presence of sanction threats, the finding that a warning banner did not deter system trespassers who obtained administrative privileges on the system carries important policy implications given that most trespassers accessed the network with administrative privileges. Although the outcomes of previous research focusing on deterrence in the physical world found evidence for a group of offenders who are not susceptible to sanction threats, this group is typically a small minority of the population. For instance, Pogarsky (2002) found a small group of potential offenders who are incorrigibles, characterized as being strongly committed to offending and unresponsive to the threat of legal sanctions. In contrast, our findings suggest that system trespassers who are unresponsive to simple sanction threats may compose a much larger proportion of trespassers. This is an important theoretical implication as our finding that most system trespassers may be unresponsive or even escalate the use of activity commands after the presence of a sanction threat in the form of a banner message appearing on the screen of attacked computer undermines the effectiveness of such a policy as a broad strategy to mitigate harm committed by system trespassers. Given that an estimated three quarters of the sample in the current study accessed target computers with administrative privileges, these findings suggest that NIST and other security agencies may need to revise the content of the warning messages. Although it is unclear at this time which alternative tactics would be more effective at generating deterrent effects across a larger proportion of system trespassers, future research should consider developing approaches targeted specifically at users with administrative privileges. For instance, policies such as a repeated series of visual warnings instead of a single warning banner or changing the content of the warning to issue more severe sanction threats to raise risk perceptions associated with system trespassing may be more effective policies than the current NIST warning banner. Indeed, the development of targeted policies aimed at deterring administrative users is a critical area of future research given that cybersecurity experts have noted the potential harm caused by system trespassers with administrative credentials (Taft, 2015) and other scholars have noted that once a system trespasser "attains the highest level of privilege, such as root or super user in UNIX, there is no reliable remedy" (Kumar, 2014: 119).

Finally, and in line with our third research hypothesis, we found evidence that nonadministrative system trespassers behaved in line with the rationale of the restrictive deterrence perspective. Specifically, we found that users without administrative privileges reduced the use of "navigation" and "change file permission" commands when on a target computer that displayed a warning banner. These findings were consistent across both the target computer level and the system trespassing event level, as well as across measures of proportion, frequency, and rate of "navigation" and "change file permission" commands entered on the system. At a practical level, these results suggest that for nonadministrative users, a sanction threat may be an effective policy in reducing trespasser activity on the system. On a theoretical level, this finding is in line with the results of previous research that have found that a credible sanction threat is effective at curtailing adverse behavior in cyberspace and that have suggested further evidence for the relevance of the application of restrictive deterrence to cyberspace (Maimon et al., 2014; Wilson et al., 2015).

The findings from the current study continue to advance knowledge regarding the effects of cyberdeterrence policies aimed at curtailing harmful behavior committed by system trespassers in cyberspace. Although the results of previous research have suggested that policies such as the presence of banner messages appearing on the screen of target computers may offer some promise for shortening system trespassers time and restricting their engagement with the system (Maimon et al., 2014; Wilson et al., 2015), this topic is in need of further research. The findings from the current study advance prior scholarship by developing theoretically based measures, which predict system trespassers' susceptibility to sanction threats. Moreover, through our study, we differentiated between system trespassers' imposed privileges on the attacked system and, in doing so, presented evidence that access to privileged information and opportunities matter in the decisions to carry out criminal activity in cyberspace and in restricting the scope of adverse behavior when faced with sanction threats.

Although the current study makes an important contribution regarding cyberdeterrence, it comprises several limitations worth noting. First, the generalizability is limited as the experiment took place using a computer network set up at one large academic institution. Specifically, as the intentions for attacking a university network are not clear, we believe that the system trespassers who attacked our computers infiltrated our systems because they were looking for targets of opportunity. In this sense, some scholars may argue that the system trespassers that participated in our experiment could be different than the system trespassers that are looking for specific targets to attack (Thomas and Stoddard, 2012). According to these scholars, system trespassers attempt to limit their engagement in large-scale attacks in an effort to reduce the likelihood of detection (Andress and Winterfeld, 2011; Rid, 2013), and consequently, they select their targets carefully based on their skill levels and motivations (Denning, 2012; Holt and Kilger, 2012). In contrast, other scholars suggest that most system trespassers are looking to infiltrate as many computers as possible, independent of the systems' owners (Spitzner, 2002). Based on this rationale, our findings could still be generalized to the broad population of system trespassers. Still, future research should aim to investigate whether system trespassers who target financial institutions or critical infrastructure that may be perceived to serve as target of choice differ meaningfully in their tactics and responses to deterring stimuli versus system trespassers who target academic institutions or other sectors. Related to this, we could not discern whether some trespassers were affiliated with the university and attacked the system as "insiders" versus whether others trespassers were unaffiliated with the university and attacked our computers from external locations (McQuade, 2006; Willison and Warkentin, 2009). Nevertheless, whether attacks came from within or outside the organization should not substantially impact the findings of the current study. We encourage future research to replicate the findings by deploying computers in various Internet infrastructures, such as private businesses or government institutions, and to uncover detailed information about the origins of the attackers' locations.

Second, from the point of entry to the system, we could not fully distinguish between human-driven and bot-based attacks.[18] Importantly, this issue would result in a downward bias on any observed effects reported in this study because if most trespassers were in fact bots, it would be difficult to observe any significant differences across the experimental conditions. Future research should be focused on deciphering trespassing events that are driven by bots from human driven attacks, and it should be aimed at testing the effectiveness of current policies in influencing both type of attacks. Similarly, in the current study, we assumed that the system trespassers were English literate. Although this limitation would similarly serve as a downward bias to the results, future research that is aimed at identifying trespassers who are English literate would be beneficial.

Third, it is possible that login credentials were shared with other system trespassers. In this case, the overall detection risk may have been diffused across several parties and may have potentially altered perceptions of risk if sharing access information systematically varied across administrative users and nonadministrative users. Similarly, it is also possible that different people may have trespassed into the same target computer by using toolkits to crack username and passwords. In this case, different system trespassing events were not showing the progression of an attack by the same trespasser on a given honeypot but different trespassing events by different trespassers. Thus, the effects in the target computer-level analysis may be smaller than what was hypothesized because of several people using the target computer. Nevertheless, even if multiple trespassers were using the same login credentials, this would not impact the findings of the system trespassing level analysis as each separate login (even with the same credentials) was recorded as an individual system trespassing event. Accordingly, this behavior should not impact the results of the system trespassing level analysis that analyzed the progression of an attack *within* a unique session.

Fourth, we could not discern differences in skill level or experience among system trespassers in the current study. Indeed, it may be possible that many of those included in the current study were unsophisticated system trespassers (i.e., script-kiddies), whereas highly skilled system trespassers or those employed by a nation-state may target government agencies, financial institutions, or critical infrastructure rather than academic institutions. Still, the results of prior research have suggested that most hackers are low skilled (Holt and Kilger, 2012). Accordingly, even if in the current study we are only capturing this subgroup of the system trespasser population, then the findings are still generalizable to *most* system trespassers and are relevant for generating deterrence-based polices aimed at this group. Nevertheless, research that is focused on identifying various skill levels of system trespassers and on examining whether the effectiveness of deterring stimuli

[18] A bot is a computer infected with malware without the user's knowledge and is controlled by cybercriminals.

such as a warning banner varies conditionally on the skill level and experience of system trespassers is a critical area for future research (Giboney et al., 2015; Zhang et al., 2015).

Finally, we have no way of determining whether system trespassers realized that they were using a target computer set up for the purposes of being attacked (Hayatle, Otrok, and Youssf, 2012, 2013; Holz and Raynal, 2005; Krawetz, 2004; Provos and Holz, 2007). Although system trespassers' probability of discovering that the target computers they attack is a fake computer is still unknown and depends on their skill level, we acknowledge that some highly skilled system trespassers may have been suspicious about our target computers. Although the fact that we used high-interaction honeypots, which allowed system trespassers to use the target computer as if these computers were real, mitigates this concern, this issue remains a potential limitation.

Conclusions

Although system trespassing has been an issue of interest to computer scientists and security experts since the late 1970s (Hollinger and Lanza-Kaduce, 1988; Parker, 1979), it has only become of interest to criminologists within the last two decades. In aiming to mitigate the consequences of system trespassing, policy makers and cybersecurity experts have turned to a broad range of deterrence-based programs and policies aimed at discouraging criminal activity in cyberspace. We expanded on research on the relevance of deterrence in cyberspace by suggesting that the effectiveness of sanction threats in influencing system trespassers' navigation and change file permission in an attacked computer system may vary on the basis of system trespassers' access to privileges on the attacked system. Our results suggest that the answer to whether the presence of a banner message reduces potentially harmful commands entered during a system trespassing event largely depends on the level of privileged access to information and opportunities afforded to trespassers. In line with the findings from decades of perceptual deterrence research in the physical world, our results demonstrate "the effect of sanctions on compliance is not one size fits all, and it is important to recognize the differential deterrability that exists across different people with respect to sanction threats" (Piquero et al., 2011: 338). Thus, our findings continue to demonstrate the relevance for applying deterrence principles and criminological theories to the realm of cyberspace and provide evidence for IT administrators and policy makers to develop flexible policies such as repeated visual and verbal cues that can be responsive to a diverse group of offenders and situations in cyberspace to reduce the harm caused by system trespassers.

Appendix: Target Computers Deployed by Experimental Conditions (Full Sample)

					Target Computers Deployed		Target Computers Deployed and Used to Enter Commands		System Trespassing with Commands Recorded	
Configuration	Memory	Disk Space	Bandwidth	Warning	Number	System Trespassing Events Recorded	Number	System Trespassing Events Recorded	Mean	SD
1	High	High	High	No	23	117	12	30	2.50	2.61
2	Low	High	High	No	30	170	13	30	2.31	1.93
3	High	Low	High	No	24	186	10	27	2.70	1.42
4	Low	Low	High	No	41	406	24	58	2.42	2.36
5	High	High	Low	No	34	213	14	34	2.43	1.40
6	Low	High	Low	No	27	179	13	19	1.46	0.88
7	High	Low	Low	No	32	192	13	28	2.15	1.95
8	Low	Low	Low	No	32	264	11	57	5.18	4.73
9	High	High	High	Yes	34	380	12	24	2.00	1.04
10	Low	High	High	Yes	29	305	14	40	2.86	2.41
11	High	Low	High	Yes	34	266	15	61	4.07	4.07
12	Low	Low	High	Yes	33	216	15	33	2.20	1.37
13	High	High	Low	Yes	37	243	14	27	1.93	1.86
14	Low	High	Low	Yes	34	248	9	19	2.11	1.17
15	High	Low	Low	Yes	29	186	16	27	1.69	1.45
16	Low	Low	Low	Yes	29	197	16	39	2.44	1.63
Total					502	3,768	221	553	2.50	2.33

Notes. ANOVA tests found no statistically significant differences in mean session between the experimental conditions. SD = standard deviation.

References

Andress, Jason and Steve Winterfeld. 2011. *Cyber Warfare: Techniques. Tactics and Tools for Security Practitioners*. Amsterdam, the Netherlands: Syngress Media.

Bacher, Paul, Thorsten Holz, Markus Kotter, and Georg Wicherski. 2005. *Know Your Enemy: Tracking Botnets*. Ann Arbor, MI: The Honeynet Project.

Bachmann, Michael. 2010. The risk propensity and rationality of computer hackers. *International Journal of Cyber Criminology*, 4: 643–656.

Beccaria, Cesare. 1963 [1764]. *On Crimes and Punishments*. New York: Macmillan.

Bentham, Jeremy. 1970 [1789]. *An Introduction to the Principles of Morals and Legislation*. Oxford, U.K.: Oxford University Press.

Berthier, Robin and Michel Cukier. 2009. An evaluation of connection characteristics for separating network attacks. *International Journal of Security and Networks*, 4: 110–124.

Blank, Stephen. 2001. Can information warfare be deterred? In (David S. Alberts and Daniel S. Papp, eds.), *Information Age Anthology, Volume III: The Information Age Military*. Washington, DC: Command and Control Research Program.

Bossler, Adam M. and George W. Burruss. 2011. The general theory of crime and computer hacking: Low self-control hackers. In (Thomas J. Holt and Bernadette H. Schell, eds.), *Corporate Hacking and Technology-Driven Crime: Social Dynamics and Implications*. Hershey, PA: IGI Global.

Brenner, Susan W. 2010. *Cybercrime: Criminal Threats from Cyberspace*. Santa Barbara, CA: ABC-CLIO.

Brezina, Timothy and Volkan Topalli. 2012. Criminal self-efficacy exploring the correlates and consequences of a "successful criminal" identity. *Criminal Justice and Behavior*, 39: 1042–1062.

Cherbonneau, Michael and Heith Copes. 2006. "Drive it like you stole it" auto theft and the illusion of normalcy. *British Journal of Criminology*, 46: 193–211.

Denning, Dorothy E. 2012. Stuxnet: What has changed? *Future Internet*, 4: 672–687.

Denning, Dorothy E. and William E. Baugh. 1999. Hiding crimes in cyberspace. *Information, Communication & Society*, 2: 251–276.

Downing, Richard W. 2004. Shoring up the weakest link: What lawmakers around the world need to consider in developing comprehensive laws to combat cybercrime. *Columbia Journal of Transnational Law*, 43: 705–763.

Dugan, Laura and Erica Chenoweth. 2012. Moving beyond deterrence: The effectiveness of raising the expected utility of abstaining from terrorism in Israel. *American Sociological Review*, 77: 597–624.

Geerken, Michael R. and Walter R. Gove. 1975. Deterrence: Some theoretical considerations. *Law & Society Review*, 9: 497–513.

Gibbs, Jack P. 1975. *Crime, Punishment, and Deterrence*. New York: Elsevier Scientific.

Giboney, Justin S., Sanjay Goel, Jeffrey Gainer Proudfoot, and Joseph S. Valacich. 2015. Measuring hacking ability using a conceptual expertise task. *Proceedings of the Conference on Digital Forensics, Security and Law*, 123–134.

Goodman, Will. 2010. Cyber deterrence: Tougher in theory than in practice? *Strategic Studies Quarterly*, 4: 102–135.

Harknett, Richard J. 1996. Information warfare and deterrence. *Parameters*, 26: 93–107.

Harknett, Richard J., John P. Callaghan, and Rudi Kauffman. 2010. Leaving deterrence behind: War-fighting and national cybersecurity. *Journal of Homeland Security and Emergency Management*, 7: 1–24.

Hayatle, Osama, Hadi Otrok, and Amr Youssef. 2012. A game theoretic investigation for high interaction honeypots. *IEEE International Conference on Communications*, 6662–6667.

Hayatle, Osama, Hadi Otrok, and Amr Youssef. 2013. A Markov decision process model for high interaction honeypots. *Information Security Journal: A Global Perspective*, 22: 159–170.

Hollinger, Richard C. and Lonn Lanza-Kaduce. 1988. The process of criminalization: The case of computer crime laws. *Criminology*, 26: 101–126.

Holt, Thomas J. 2007. Subcultural evolution? Examining the influence of on-and off-line experiences on deviant subcultures. *Deviant Behavior*, 28: 171–198.

Holt, Thomas J. and Adam M. Bossler. 2016. Technology and violence. In (Carlos A. Cuevas and Callie M. Rennison, eds.), *The Wiley-Blackwell Handbook on the Psychology of Violence*. Chichester, U.K.: Wiley Blackwell.

Holt, Thomas J. and Max Kilger. 2012. Know your enemy: The social dynamics of hacking. *The Honeynet Project*. Retrieved from honeynet.org/papers/socialdynamics

Holz, Thorsten and Frederic Raynal. 2005. Detecting Honeypots and Other Suspicious Environments. Paper presented at the Sixth Annual IEEE SMC Information Assurance Workshop.

Holzer, Corey T. and James E. Lerums. 2016. *The Ethics of Hacking Back*. West Lafayette, IN: CERIAS Tech Report.

Hutchings, Alice and Thomas J. Holt. 2015. A crime script analysis of the online stolen data market. *British Journal of Criminology*, 55: 596–614.

Hutchins, Eric M., Michael J. Cloppert, and Rohan M. Amin. 2011. Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. *Leading Issues in Information Warfare and Security Research*, 1: 80.

Jacobs, Bruce A. 1993. Undercover deception clues: A case of restrictive deterrence. *Criminology*, 31: 281–299.

Jacobs, Bruce A.1996a. Crack dealers and restrictive deterrence: Identifying narcs. *Criminology*, 34: 409–431.

Jacobs, Bruce A. 1996b. Crack dealers' apprehension avoidance techniques: A case of restrictive deterrence. *Justice Quarterly*, 13: 359–381.

Jacobs, Bruce A. 2010. Deterrence and deterrability. *Criminology*, 48: 417–441.

Jacobs, Bruce A. and Michael Cherbonneau. 2014. Auto theft and restrictive deterrence. *Justice Quarterly*, 31: 344–367.

Jordan, Tim. 2016. A genealogy of hacking. *Convergence: The International Journal of Research into New Media Technologies*. https://doi.org/1354856516640710.

Jordan, Tim and Paul Taylor. 1998. A sociology of hackers. *The Sociological Review*, 46: 757–780.

Keizer, Gregg. 2009. Russian cyber militia knocks Kyrgyzstan offline. *Computerworld*, 1: 28.

Kigerl, Alex C. 2015. Evaluation of the CAN SPAM Act testing deterrence and other influences of e-mail spammer legal compliance over time. *Social Science Computer Review*, 33: 440–458.

Krawetz, Neal. 2004. Anti-honeypot technology. *IEEE Security and Privacy*, 2: 76–79.

Kumar, Sandeep. 2014. Advances in intrusion detection systems with applications to data mining. *International Journal of Applied Science and Engineering*, 2: 117–125.

LaFree, Gary, Laura Dugan, and Raven Korte. 2009. The impact of British counterterrorist strategies on political violence in Northern Ireland: Comparing deterrence and backlash models. *Criminology*, 47: 17–45.

Lee, Cynthia B., Chris Roedel, and Elena Silenok. 2003. *Detection and Characterization of Port Scan Attacks* (Unpublished doctoral dissertation). Los Angeles: Department of Computer Science and Engineering, University of California.

Libicki, Martin C. 2009. *Cyberdeterrence and Cyberwar*. Santa Monica, CA: RAND.

Liu, Jing, Yang Xiao, Kaveh Ghaboosi, Hongmei Deng, and Jingyuan Zhang. 2009. Botnet: Classification, attacks, detection, tracing, and preventive measures. *EURASIP Journal on Wireless Communications and Networking*, 1184–1187.

Maimon, David, Mariel Alper, Bertrand Sobesto, and Michel Cukier. 2014. Restrictive deterrent effects of a warning banner in an attacked computer system. *Criminology*, 52: 33–59.

Maimon, David and Christopher R. Browning. 2012. Underage drinking, alcohol sales and collective efficacy: Informal control and opportunity in the study of alcohol use. *Social Science Research*, 41: 977–990.

Maimon, David, Theodore Wilson, Wuiling Ren, and Tamar Berenblum. 2015. On the relevance of spatial and temporal dimensions in assessing computer susceptibility to system trespassing incidents. *British Journal of Criminology*, 55: 615–634.

McQuade, Samuel C. 2006. *Understanding and Managing Cybercrime*. Boston, MA: Pearson Education.

Na, Chongmin, Thomas A. Loughran, and Raymond Paternoster. 2015. On the importance of treatment effect heterogeneity in experimentally-evaluated criminal justice interventions. *Journal of Quantitative Criminology*, 31: 289–310.

Nagin, Daniel S. 1998. Criminal deterrence research at the outset of the twenty-first century. In (Michael Tonry, ed.), *Crime and Justice: A Review of Research*, vol. 23. Chicago, IL: University of Chicago Press.

Nagin, Daniel S. 2013. Deterrence in the twenty-first century. In (Michael Tonry, ed.), *Crime and Justice: A Review of Research*, vol. 42. Chicago, IL: University of Chicago Press.

Nagin, Daniel S., Robert M. Solow, and Cynthia Lum. 2015. Deterrence, criminal opportunities, and the police. *Criminology*, 53: 74–100.

National Institute for Standards and Technology (NIST). 2009. *Recommended Security Controls for Federal Information Systems and Organization*. Washington, DC: U.S. Department of Commerce.

Nguyen, Holly, Aili Malm, and Martin Bouchard. 2015. Production, perceptions, and punishment: Restrictive deterrence in the context of cannabis cultivation. *International Journal of Drug Policy*, 26: 267–276.

Parker, Donn B. 1979. *Computer Crime: Criminal Justice Resource Manual.* Washington, DC: National Institute of Justice.

Paternoster, Raymond. 2010. How much do we really know about criminal deterrence? *Journal of Criminal Law and Criminology*, 100: 765–824.

Piliavin, Irving, Rosemary Gartner, Craig Thornton, and Ross L. Matsueda. 1986. Crime, deterrence, and rational choice. *American Sociological Review*, 51: 101–119.

Piquero, Alex R., Raymond Paternoster, Greg Pogarsky, and Thomas A. Loughran. 2011. Elaborating the individual difference component in deterrence theory. *Annual Review of Law and Social Science*, 7: 335–360.

Png, Ivan P. L. and Qui-Hong Wang. 2009. Information security: Facilitating user precautions vis-à-vis enforcement against attackers. *Journal of Management Information Systems*, 26: 97–121.

Pogarsky, Greg. 2002. Identifying "deterrable" offenders: Implications for research on deterrence. *Justice Quarterly*, 19: 431–452.

Ponemon Institute. 2014. The Second Annual Study on Data Breach Preparedness (Research report). Retrieved from ponemon.org/blog/is-your-company-ready-for-a-big-data-breach-the-second-annual-study-on-data-breach-preparedness.

Pratt, Travis C., Francis T. Cullen, Kristie R. Blevins, Leah E. Daigle, and Tamara D. Madensen. 2006. The empirical status of deterrence theory: A meta-analysis. In (Francis T. Cullen, John Paul Wright, and Kristie R. Blevins, eds.), *Taking Stock: The Status of Criminological Theory—Advances in Criminological Theory*, vol. 15. New Brunswick, NJ: Transaction.

Provos, Niels, Markus Friedl, and Peter Honeyman. 2003. Preventing Privilege Escalation. Paper presented at the USENIX Security Symposium. Retrieved from usenix.org/conference/11th-usenix-security-symposium/preventing-privilege-escalation.

Provos, Niels and Thorsten Holz. 2007. *Virtual Honeypots: From Botnet Tracking to Intrusion Detection*. Boston, MA: Pearson Education.

Rege, Aunshul, Frank Ferrese, Saroj Biswas, and Li Bai. 2014. *Adversary Dynamics and Smart Grid Security: A Multiagent System Approach*. Paper presented at the IEEE Resilient Control Systems (ISRCS), 2014 7th International Symposium.

Rid, Thomas. 2013. *Cyber War Will Not Take Place*. New York: Oxford University Press.

Riden, Jamie and Christian Seifert. 2008. *A Guide to Different Kinds of Honeypots* (Symantec Research Report). Feb. 13. Retrieved from symantec.com/connect/articles/guide-different-kinds-honeypots.

Riffkin, Rebecca. 2014. Hacking tops list of crimes Americas worry about most. *Gallup Poll News Survey*. Oct. 27. Retrieved from gallup.com/poll/178856/hacking-tops-list-crimes-americans-worry.aspx.

Riofrio, Melissa. 2013. Hacking back: Digital revenge is sweet but risky. *PC World*. May 9. Retrieved from pcworld.com/article/2038226/hacking-back-digital-revengeis-sweet-but-risky.html.

Sanger, David E., Nicole Perlroth, and Michael D. Shear. 2015. Attack gave Chinese hackers privileged access to U.S. systems. *The New York Times*. June 20. Retrieved from nytimes.com/2015/06/21/us/attack-gave-chinese-hackers-privileged-access-to-us-systems.html?_r=0.

Schell, Bernadette H. and John L. Dodge. 2002. *The Hacking of America: Who's Doing It, Why, and How*. Westport, CT: Greenwood.

Shen, Mou, Mengdong Chen, Min Li, and Lianzhong Liu. 2013. Research of least privilege for database administrators. *International Journal of Database Theory and Application*, 6: 39–50.

Sherman, Lawrence W. 1993. Defiance, deterrence and irrelevance: A theory of the criminal sanction. *Journal of Research in Crime and Delinquency*, 30: 445–473.

Snowberger, Phil and Douglas Thain. 2005. *Sub-identities: Towards Operating System Support for Distributed System Security* (Technical Report 2005-18). South Bend, IN: Department of Computer Science and Engineering, University of Notre Dame.

Spitzner, Lance. 2002. *Honeypots: Tracking Hackers*. Reading, MA: Addison-Wesley.

Stockman, Mark, Robert Heile, and Anthony Rein. 2015. *An Open-Source Honeynet System to Study System Banner Message Effects on Hackers*. Paper presented at the 4th Annual ACM Conference on Research in Information Technology.

Taft, Darryl K. 2015. CA to acquire Xceedium for privileged identity management. *Eweek*. Aug. 4. Retrieved from eweek.com/it-management/ca-to-acquire-xceedium-for-privileged-identity-management.html.

Thakare, Shailesh P., Prerana Chandurkar, and Maithili S. Deshmukh. 2013. Computer attacks and intrusion detection system: A need review. *International Journal of Computer Science and Applications*, 6: 425–436.

Thomas, Kyle J., Thomas A. Loughran, and Alex R. Piquero. 2013. Do individual characteristics explain variation in sanction risk updating among serious juvenile offenders? Advancing the logic of differential deterrence. *Law and Human Behavior*, 37: 10–21.

Thomas, Tom and Donald Stoddard. 2012. *Network Security First-Step*. Indianapolis, IN: Cisco Press.

Wall, David. 2001. *Crime and the Internet*. New York: Routledge.

Willison, Robert and Merrill Warkentin. 2009. *Motivations for Employee Computer Crime: Understanding and Addressing Workplace Disgruntlement Through the Application of Organisational Justice*. Paper presented at the IFIP TC8 International Workshop on Information Systems Security Research. International Federation for Information Processing.

Wilson, Theodore, David Maimon, Bertrand Sobesto, and Michel Cukier. 2015. The effect of a surveillance banner in an attacked computer system additional evidence for the relevance of restrictive deterrence in cyberspace. *Journal of Research in Crime and Delinquency*, 52: 829–855.

Xu, Zhengchuan, Qing Hu, and Chenghong Zhang. 2013. Why computer talents become computer hackers. *Communications of the ACM*, 56: 64–74.

Yar, Majid. 2006. *Cybercrime and Society*. Thousand Oaks, CA: Sage.

Zhang, Xiong, Alex Tsang, Wei T. Yue, and Michael Chau. 2015. The classification of hackers by knowledge exchange behaviors. *Information Systems Frontiers*, 17: 1239–1251.

Zhuge, Jian-wei, Xin-hui Han, Yong-lin Zhou, Cheng-yu Song, Jing-peng Guo, and Wei Zou. 2007. HoneyBow: An automated malware collection tool based on the high-interaction honeypot principle. *Journal—China Institute of Communications*, 28: 8.

Statute Cited

Computer Fraud and Abuse Act, 18 U.S.C. § 1030 (a)(5)(c) (1986).

Alexander Testa is a Ph.D. candidate in the Department of Criminology and Criminal Justice at the University of Maryland—College Park. His research interests include the consequences of criminal justice contact across the life course, criminal justice decision making, cross-national causes and consequences of violence, and criminal justice policy.

David Maimon is an associate professor of criminology and criminal justice at the University of Maryland—College Park. His research interests include cybercrime, experimental methods, and community and crime.

Bertrand Sobesto is senior IT security engineer in the Division of Information Technology at the University of Maryland—College Park. He operates a network of about 1,500 honeypots to support the Cybercrime research initiative, a project aiming at studying the attackers' behavior on compromised systems.

Michel Cukier is the director for advanced cybersecurity experience for students (ACES) and the associate director for education for the Maryland Cybersecurity Center (MC2). Michel is an associate professor of reliability engineering with a joint appointment in the Department of Mechanical Engineering at the University of Maryland—College Park. His research covers dependability and security issues. He has published more than 80 papers in journals and refereed conference proceedings in those areas.