Abstract

Title of Thesis:	RESTRICTIVE DETERRENCE AND THE SEVERITY OF HACKERS' ATTACKS ON COMPROMISED COMPUTER SYSTEMS
	Theodore Henry Wilson II, Master of Arts, 2014
Thesis directed by:	Assistant Professor David Maimon Department of Criminology and Criminal Justice

There is a lack of consensus within the literature assessing whether surveillance is effective in reducing the seriousness of criminal events, with almost no prior study investigating its operation in cyberspace. This thesis seeks to address both of these domains while drawing on the deterrence perspective. Data were obtained from an experiment conducted over seven months at a large, public university within the United States. Specifically, a series of virtual computers with known vulnerabilities were deployed into the university's computer network as part of a randomized controlled trial. This thesis seeks to examine 1) whether a surveillance banner reduces the severity of offending through inhibiting hackers from escalating to active engagement with the system upon gaining access on the first session and 2) whether the deterrent effect of a surveillance banner persists beyond the first session. This surveillance banner produced a restrictive deterrent effect for the first and second sessions.

RESTRICTIVE DETERRENCE AND THE SEVERITY OF HACKERS' ATTACKS ON COMPROMISED COMPUTER SYSTEMS

by

Theodore Henry Wilson II

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Master of Arts
2014

Advisory Committee:

Assistant Professor David Maimon, Chair
Associate Professor Thomas A. Loughran
Professor Raymond Paternoster

Copyright Statement

Table of Contents

1. Introduction

The incidence of computer trespassing in particular has been growing rapidly, and is one of the fastest growing branches of criminal conduct within the cyber domain (Furnell 2002). McQuade defines computer trespassing as, "illegally gaining access to one or more computer systems after exploiting security vulnerabilities or defeating a security barrier" (2006: 83). Indeed, the true prevalence of computer trespassing incidents throughout society is difficult to estimate due to the inability to adequately detect instances of computer trespassing at large. Nevertheless, a recent survey of IT managers conducted by the Ponemon institute (2011) found that over 90% of the represented US corporations detected multiple instances of computer trespassing incidents. These trespassing incidents are more than simple violations of privacy, as they may put sensitive data, financial sources, and proprietary records at risk of deletion, exploitation, or manipulation (Whitman 2003). According to Whitman (2003), data breaches to private and public U.S. corporations result in financial losses of billions of dollars annually in the U.S. These computer trespassing incidents were criminal acts that serve as the platform for a wide range of further criminal conduct ranging from fraud to the outright destruction or corruption of industrial devices and software (McQuade 2006).

Extensive research has been conducted by scientists toward improving the detection of technical characteristics of computer trespassing incidents, and the development of defensive technologies that will prevent individuals from being able to trespass on a given computer system (Zhang et al 2003; Mackey 2003; Berthier and Cukier 2009). However, this branch of research has not been driven by criminological models, and has typically been confined to means of target hardening at the border of a network of computers connection with the internet through firewalls and other forms of intrusion prevention software (see Salles-Lousto et al 2011). Addressing this issue, the present study draws on draws on the deterrence perspective (Gibbs 1975; Jacobs 2010) and assesses the impact of a surveillance banner and a surveillance process upon the behavior of intruders during system trespassing incidents. This analysis seeks to address several research questions related to an extension of deterrence theory to incorporate restrictive deterrence principles in identifying reductions in severity of offending, as constructed by Gibbs (1975) and Jacobs (2010). Specifically, this thesis seeks to address questions pertaining to absolute and restrictive deterrence. First, I seek to answer whether a surveillance banner displayed upon entry to a computer system reduces the likelihood of recidivism in terms of subsequent system trespassing incidents. Second, I ask whether a surveillance banner displayed upon entry to a computer system reduces the probability of an intruder actively engaging the system with a keystroke (presence of any commands typed into the console) during the first system trespassing incident? Third, I explore if the effect of a surveillance banner on an intruder's decision to actively engage with the system during a system trespassing incident persists beyond the first system trespassing session to subsequent sessions? Lastly, I investigate whether a surveillance banner reduces the ratio of system trespassing incidents with keystrokes to all observed system trespassing incidents on a target system?

2. Theoretical Background

2.1 System Trespassing

System trespassing necessarily involves the unauthorized use and violation of the private use of property (Mcquade 2006). This can be seen to be analogous to home invasions and automotive theft, which also involve the unauthorized use and violation of private property rights. Unlike these somewhat analogous actions within the physical realm however, system trespassers can access a computer remotely through an internet connection in addition to locally logging into the computer through a medium physically and directly connected to the computer (Anderson 1980). The damage resulting from a system trespassing session necessarily depends upon the motivation(s) of the intruder, and attacks can range in seriousness from benign queries of the internet to active manipulation of essential files on the target system and even launching denial of service attacks (Stallings 2005; McQuade 2006).

The past three decades have seen numerous efforts from technical researchers and politicians toward reducing the potential and incidence of system trespassing incidents. Technical researchers have been investigating techniques, software, and related solutions to detect and prevent unauthorized access of computer systems (Mackey 2003). Congress enacted the “Computer Fraud and Abuse Act” in 1986 in an effort to establish formalized sanctions for acts of computer violation and misuse; the sanctions inherent to violating this act extend to up to ten years of imprisonment. The intended impact of this act and the belief that providing sanctions for a type of behavior will reduce the incidence of that conduct is in line with the deterrence-based assumption that individual behavior can be modified due to a fear of and the imposition of pending sanctions for criminal conduct (Paternoster 1987).

2.2 Deterrence Theory and Restrictive Deterrence

Deterrence theory can be traced back to the utility-based principles presented through the works of Bentham (1785) and Beccaria (1764). Both Bentham and Beccaria conceived of individuals as rational actors who were susceptible to the influence of sanctions as they weighed the potential costs and benefits of committing a potential criminal act in what is termed a hedonistic calculus (Bentham 1785). Deterrence theory extends this tradition to focus more exclusively on the individual influence of the severity, celerity, and certainty of sanctions in leading an individual to refrain from criminal conduct (Paternoster 1987). The belief within this framework is that severe, certain, and swift punishments will be effective in increasing the costs associated with crime to the point of outweighing the benefits of crime, which leads to deterrence. The theory has evolved and grown to incorporate additional factors and concepts, as theorists and researchers continue to investigate deterrence principles. Branches of deterrence theory now attempt to account for the impact of punishment avoidance, whereby an offender commits a criminal act and goes without punishment in updating the offender's perception of the certainty of punishment (Stafford and Warr 1993). The work of Stafford and Warr is intrinsically related to the work of Becker (1968) as it distinguishes between objective and subjective sanctions in raising the distinction between the actual risks of sanctions and the individual's perception of the risk of sanctions within a given context. Anwar and Loughran (2011) provide evidence that the relationship between an individual's perception of the risk of sanctions is dynamic and that individuals update their perception of the risk of sanctions in response to new information on the "true" risk of sanctions in a Bayesian learning model. Importantly, theorists have also incorporated alternative outcomes to this deterrence model in order to differentiate between absolute and restrictive deterrence (Gibbs 1975).

According to Gibbs (1975), absolute deterrence is whereby an offender is wholly deterred from engaging in any criminal conduct due to fear of the potential for receiving sanctions. This abstention of criminal behavior within absolute deterrence is “absolute” in that it involves a complete refraining of engaging in this particular type of criminality on the part of the individual. In contrast, restrictive deterrence is a process whereby an offender is not wholly deterred from engaging in criminal conduct, but the offender’s behavior is modified in a specific manner to reduce the probability of receiving sanctions (Gibbs 1975; Jacobs 2010). Restrictive deterrence primarily focuses upon reductions in the frequency (Gibbs 1975) and the severity (Jacobs 1993) of criminal offending in an effort to reduce the probability of being detected and receiving sanctions in response to deterring cues in the environment. Criminologists have devoted considerable attention to the aforementioned contributions to the deterrence framework (Pratt et al 2006) with the notable exception of the restrictive deterrence principles. Prior to the current research, researchers have provided minimal treatment to these alternative outcome measures that should also be incorporated into deterrence-based research (Paternoster 1987; Jacobs 1993).

Gibbs (1975) describes restrictive deterrence as, “the curtailment of a certain type of criminal activity by an individual during some period because in whole or in part the curtailment is perceived by the individual as reducing the risk that someone will be punished as a response to the activity” (1975: 33). Gibbs’ initial conception of restrictive deterrence focused primarily upon the reduction in the frequency of criminal conduct, which necessarily limits the applicability of the constructs to those offenders who are already engaged in criminal conduct. This platform was later extended by Jacobs to enable for both the analysis of an individual modifying the initial act of crime (Jacobs 1993), and further to include reducing the seriousness of the criminal act, engaging in modification behavior or measures that reduce the risk of detection, and altering the spatial-temporal component of the criminal event (Jacobs 2010).

Empirical analyses of restrictive deterrence theory as presented by Gibbs and Jacobs are still preliminary within the field with only a handful of authors attempting to test the theory against nature (Jacobs 1993, 1996a, 1996b; Jacobs and Cherbonneau 2012; Jacobs and Miller 1998; Paternoster 1989). Jacobs (1996a) reports frequent location changes employed by crack dealers to avoid detection by both formal and informal agents within a neighborhood. Jacobs (1996b) interviewed crack dealers and reported the nonverbal and verbal cues crack dealers were primed for to ascertain the likelihood of conducting a sale with a prospective buyer that was actually an undercover police officer; one such “test” was to check who a buyer knew in the neighborhood to establish credibility. Beauregard and Bouchard (2010) identify concealment measures and the use of gloves by rapists to reduce both the evidence left at the scene and the eventual likelihood of apprehension. Jacobs and Cherbonneau (2012) observe auto thieves changing the license plate(s) of the car(s) they steal to avoid the cars coming up on the “hot” list of license plate numbers employed by police officers to identify stolen cars. Each of these measures is employed by the respective criminal(s) for the purpose of avoiding apprehension, but these findings do not touch upon the offender reducing the seriousness of their offense as a means to avoid apprehension. Most of these analyses were largely qualitative (Jacobs 1993, 1996a, 1996b; Jacobs and Cherbonneau 2012) and based on relatively small sample sizes (Beauregard and Bouchard 2010; Jacobs 1996b; Jacobs and Cherbonneau 2012). These studies also largely focused upon the modification behaviors employed by criminals to reduce the probability of receiving sanctions. This thesis seeks to contribute to this ongoing literature by investigating whether the seriousness of offending may also be impacted through this process and to see whether the principles of restrictive deterrence are evident in the realm of cyberspace.

The literature regarding the theoretical applicability of deterrence principles to cyberspace is mixed. Some scholars have discounted the applicability and utility of deterrence to cyberspace on the grounds that the internet affords a high level of anonymity and this anonymity affords a great deal of difficulty in tying offenders to the criminal incidents with formal sanctions (Blank 2011). Goodman (2010) disagrees with this assertion to claim that the immediate link between cyber-attacks and the respective criminal is not necessary to activate deterrence. While this literature as to the applicability of deterrence to cyberspace remains inconclusive in the abstract, Maimon and colleagues (2014) have produced empirical results toward revealing a deterrence effect upon hacker behavior by virtue of a warning banner.

Maimon and colleagues (2014) assessed the impact of a warning banner displayed upon entry to the computer system on subsequent intruder behavior; they found that such a warning banner reduced the duration of system trespassing incidents on both the first session and repeated sessions on a target system. This result provides empirical evidence for cyber attackers' susceptibility to deterrence cues and the applicability of deterrence within cyberspace. The experiment discussed by Maimon and colleagues (2014) is part of a broader collection of ongoing experiments presently under study at a large, public university in the United States toward providing further empirical evidence regarding the applicability of deterrence and social science-derived measures in efforts at reducing computer-focused crimes.

2.3 Surveillance and Crime

Geerken and Grove (1975) define a system of deterrence as a “communication mechanism” that’s purpose is to inform offenders of three components: (1) the probability that a criminal act will be detected by authorities is high, (2) the probability of receiving sanctions once detected is high, and (3) the severity, or cost, associated with the sanction is great enough to offset the potential reward of the criminal act. As such, a deterrent element or process’ degree of success is determined by the adequacy with which this deterring message is conveyed to the target actors, with those messages targeting the total population of potential offenders having the greatest success. Geerken and Grove also distinguish between formal deterrence systems that operate through legal agencies and informal deterrence systems that operate at the interpersonal level. I propose that, in keeping with this rationale afforded by Geerken and Grove (1975), *displaying a targeted and specific surveillance warning or banner should serve as an effective platform for the purpose of imparting deterring cues.* Such a message should directly impart information pertaining to the risk of detection for a given offender. An increase in the perception of the risk of detection is an immediate result of an effective deterrence system according to Geerken and Grove (1975) and it should activate the restrictive deterrent principles acknowledged by Jacobs (2010), resulting in an immediate reduction in the seriousness of offending in order to reduce the perceived risk of detection.

In addition to a notification of a sanction threat, I suspect that the presence of surveillance means in the environment should also serve as an effective deterrent against the occurrence of a criminal incident (Nagin and Pogarsky 2003). One means of surveillance within the physical realm is the use of closed circuit television (CCTV) cameras to record video images of a target area and send a signal containing that footage to a secondary location. While there has been extensive research conducted regarding the efficacy of CCTV cameras, results have been mixed with many studies finding a reduction in criminal conduct and others finding no effect for CCTV cameras (Welsh and Farrington 2008; Ratcliffe et al 2009). Additionally, there has been scant attention on the interplay between the surveillance process and the individual's perception of the risk of detection and perception of sanctions (Ratcliffe et al 2009). This surveillance-based research has been confined to physical landscapes with regard to monitoring of real world spaces; this work seeks to extend this scope with respect to the potential for surveillance processes within cyber environments.

2.4 The Present Study

This thesis seeks to directly address the aforementioned empirical gaps through assessing the impact of a surveillance banner upon the seriousness of offending by intruders on a compromised computer system. In accordance with restrictive deterrence principles, the seriousness of the system trespassing incident is assessed as an outcome measure related to the presence of a surveillance banner on the target computer system. Further, an increase in seriousness within a cyber-environment will result in a more extensive log file for and from the target system. This can contribute to a would-be offender abstaining from escalating his or her offending if they believe this abstention will reduce the likelihood of detection, which could occur through the collection of log files in this instance.[1] Importantly, absolute deterrence is not the appropriate mechanism

[1] Log file collectors are not necessarily identifiable by interfacing with a target computer directly as the collector could be on an external device to the target computer. This results in a necessary ambiguity as to whether such a collector is actively collecting data on activity on the target computer.

in the context of this work, as I am interested in studying how contextual information intended to alter the intruder's perceptions regarding the risk of detection and the corresponding probability of receiving sanctions operate within an environment with a low likelihood of receiving sanctions. Nevertheless, the potential for an absolute deterrent effect is directly assessed as a primary hypothesis of this analysis.

The active engagement with a target system by typing directly into the console through inputting keystrokes of any kind is conceptualized as an increase in the seriousness of offending for the present analysis. Each system trespassing incident is a criminal act and involves unauthorized access to a computer system. However, any and all of the actions that can be engaged in during a system trespassing incident necessarily involve the execution of keystroke commands on the target system. This thesis distinguishes between those system trespassing incidents during which the intruder typed into the system directly from those system trespassing incidents during which the attacker refrained from engaging with the system directly by typing in any commands or keystrokes. Regardless of the intent of the intruder, any actions performed on the system by the intruder beyond the initial unauthorized access of the system is a form of active manipulation of the system through either the use of resources to run the entered command or the purposive deletion of essential files. It is following this rationale that this thesis advocates a keystroke session as a more serious form of a system trespassing incident. This conception also fits within a restrictive deterrence framework as opposed to an absolute deterrence framework as the system trespassing incidents are still criminal events, but can be of a less serious nature due to a potential lack of keystrokes. This lack of keystrokes can also be observed, albeit without complementary anecdotal evidence, as a means to reduce the potential for sanctions, as there would not be as much information inserted into the aforementioned log file for the system if active monitoring and data recording were believed to be present by the intruder.

The five hypotheses for this work are as follows:

(1) The presence of a surveillance banner will not have an effect upon the likelihood of an attacker recidivating with repeated system trespassing incidents.

A surveillance banner on a target system would only be displayed to a prospective offender upon successful entry to the computer system, which would constitute a system trespassing incident. This form of banner is otherwise not visible to the offender while not connected to the target system. As such, the effect of the banner upon an offender's decision to return to the system, to recidivate, is expected to be marginal or lacking. The presence of the surveillance banner on a target system is expected to alter the behavior during the given system trespassing incident, but not the offender's decision to connect to the target system during subsequent instances.

(2) The presence of a surveillance banner upon entry will reduce the probability of an attacker typing on the system on the first system trespassing incident.

A surveillance banner should impart information to a would-be intruder that there is a heightened level of detection and/or monitoring on the target system. As such, this should increase the perceived risk of detection on the part of the offender. One means by which the offender can reduce this potential for detection is through abstaining from actively engaging within the system through typing on the system through the input of keystrokes. This rationale follows that offered by Jacobs (2010) that offenders will reduce the seriousness of their offending if that reduction will reduce the risk of detection.

(3a) A surveillance banner will reduce the probability of an intruder typing on the system on the second system trespassing incident if the intruder did not type on the system during the first system trespassing incident for a given intruder.

The restrictive deterrent effect hypothesized for the first system trespassing incident should extend to the second system trespassing incident for those intruders who do not actively engage with the system during the first system trespassing incident. This class of offenders who choose not to actively engage with the system during the first system trespassing incident would include those individuals who otherwise would not have engaged with the system in terms of entering keystrokes as well as those individuals who were restrictively deterred by the presence of this banner on the system. It is possible that this group of individuals retained perceptions of the certainty of detection that crossed the tipping point identified by Loughran and colleagues (2012) beyond which deterrence is most likely to occur. While this tipping point is not observed directly due to an inability to query the offenders directly on their perceptions of the certainty of detection and receiving sanctions, it is a potential mechanism by which the surveillance banner may alter the behavior of the offenders. In line with this rationale, this group of offenders should still be capable of being deterred on their respective second system trespassing incidents; these individuals' perception of the certainty of punishment should theoretically remain unchanged as they did not actively engage with the system during the second system trespassing incident. These individuals do not receive additional information with which to update their perceptions of punishment, as they did not escalate the offense and did not receive sanctions. As such, these intruders should still be prone to restrictive deterrence during the second system trespassing incident as they abstain from actively engaging with the target computer in the hope of reducing the risk of detection.

(3b) A surveillance banner will not have an effect on the probability of an intruder typing on the system on the second system trespassing incident if the intruder previously typed into the system during the first system trespassing incident for a given intruder.

The previous argument pertains to the sample of individuals who choose to not actively engage with the system during the first system trespassing incident, and are subsequently reasoned to be subject to deterrence principles during the second system trespassing incident. However, the other sample of individuals, who chose to actively engage with the system, are conceived as containing those individuals who would engage in the criminal conduct regardless of the deterrence mechanism at play and those individuals whose perceptions of the certainty of detection was not sufficiently altered by the banner to prompt them to abstain from actively engaging with the system. As this group of individuals was not previously deterred from actively engaging with the system, it is expected that they will likewise not be deterred during their second system trespassing incident. Additionally, if there were no sanctions or consequences levied as a result of their active engagement with the system during the first system trespassing incident, this should result in a lower updated perception of the certainty of detection in line with the findings of Loughran and colleagues (2012). As such, a null effect is expected with regard to the banner's potential to reduce the probability of a keystroke on the second session for this group of individuals that actively engaged with the system previously during the first system trespassing incident.

(4) A surveillance banner displayed upon entry during repeated system trespassing incidents will reduce the ratio of system trespassing incidents with commands typed into the computer to the total number of system trespassing incidents for a given intruder.

The aforementioned potential effects of the surveillance banner upon reducing the seriousness of the first system trespassing incident and a conditional impact upon the second system trespassing incident should also be observable within the ratio of system trespassing incidents with keystrokes to the total number of system trespassing incidents on a target system. The restrictive deterrent effect should be in lesser magnitude for this calculated ratio, as once an intruder actively engages with the system and does not receive sanctions or consequences, then the perception of the certainty of detection and punishment should update to a reduced degree that is more reflective of reality (Loughran et al 2012). The greatest impact should be most readily apparent on the first session during which the offender will not have experienced punishment avoidance (Stafford and Warr 1993) or have a reduced perception of the certainty of punishment due to repeatedly accessing the target computer without receiving sanctions. Even so, this thesis predicts that there will still be a global reduction in the seriousness of the system trespassing incidents that occur in the presence of a surveillance banner.

3. Experimental Design

To test the above research hypotheses, the present study employed secondary data collected from a previously administered randomized controlled trial. This experiment was conducted by a team of criminologists and engineers, and featured a two (surveillance banner vs no surveillance banner) x two (surveillance process vs no surveillance process) factor design. These four conditions can be described as (1) control, (2) surveillance banner only, (3) surveillance process only, and (4) both surveillance process and surveillance banner.

3.1 Design

The experiment in question utilized 300 public IP addresses from a large, public university to deploy experimentally controlled and standardized target computers into the university's network.[2] The 300 IP addresses were from a single, continuous subnet devoted to this experiment.[3] The target computers were virtually deployed and maintained from a set of physical servers housed within the respective university. While there was not a physical box for each of the deployed target computers utilized in this experiment, the intruder cannot ascertain that the machine is virtually deployed and it does not affect the behavior of the system. Each of these target computers was designed to be the equivalent of a real computer system running a version of the Linux-based operating system (CentOS). Intruders who found the vulnerable computers sitting behind the 300 IP addresses will attempt to brute force entry onto the devices – typically with toolkits[4] designed expressly for the purpose of cracking the legitimate login credentials to user accounts (McQuade 2006). To aid in the simulation of the intruder acquiring the legitimate credentials to the system, a predefined number of attempts for each intruder

[2] An IP address is the public identification number that computer systems use to connect to the internet and communicate with one another. IP addresses are allocated to individual machines attempting to join an existing network. IP addresses can be either dynamic (changing while the machine is actively connected to the network) or static. A single IP address consists of four numbers in sequence ranging from 1 to 256 (e.g., 111.11.111.11).

[3] A subnet is a range of IP addresses (e.g., 111.11.111.1-111.11.111.256).

[4] These toolkits enable the prospective intruder to select a username to attempt to gain access to the system with. Many users will initially target the user, root, as it is the standard admin user for linux-based systems that will retain heightened privileges within the computing environment over standard users. This is supported by the findings presented in Tables 3 and 4 with the distribution of the user names used by intruders during the deployment.

was randomly set between 150 and 200 attempts. This number of attempts limits the ability of human users to arbitrarily type in login credentials by hand, and stimulates the inclusion of those individuals who use brute force tool kits. Most brute force toolkits do not report the number of trials needed to gain legitimate access to the system, so whether the threshold was set at 100 or 1000 attempts should not have an immediate bearing upon the resulting sample included in the study. When this threshold is met, the login credentials used during this nth attempt are treated as the legitimate credentials for the target system. These are the credentials that must be used to conduct system trespassing sessions following the initial deployment.

Randomization to a condition occurs at the point of deployment. As soon as the intruder reached the predetermined nth attempt and gained access to the system, a random number was generated between 1 and 4 inclusive. This determined the experimental condition the offender's target computer will receive for the full 30-day duration of the respective deployment. This randomization procedure took place before the intruders could access the system, and can thus be considered to be exogenous to any and all differences observed on the target system over the course of a deployment.

3.2 Overview of the Present Study

The surveillance banner employed in the experiment takes advantage of a default welcome prompt to display a custom message. This banner appears as follows on the target computer system:

```
    This system is under continuous surveillance.
      All user activity is being monitored and
                     recorded.
```

This banner is the first set of text seen upon each entry to a target computer with the surveillance banner condition. The surveillance process employed in the experiment is a zabbix process inserted onto the system's list of active processes called "Monitor." This process does not impact system performance, and is only visible if the intruder calls forth a display depicting all of the processes running on the system.[5]

3.3 Sampling Procedure

This experiment did not actively solicit or attempt to encourage intruders to launch brute force attacks and receive the assignment of a target computer. In contrast with most experiments conducted within criminology, an active sampling methodology was not employed in favor of a passive sampling methodology that placed computers with active vulnerabilities on specific port numbers into the university's computer network. Target computers were deployed onto the network for a period of seven months from April 4, 2013 until November 3, 2013 with data collection pertaining to system trespassing incidents continuing for an additional 30 days beyond this period. These vulnerable computers, or honeypots, naturally attracted hackers that scan networks in search of computers with these known vulnerabilities (McQuade 2006). As such, this experiment is largely passive with regard to sample matriculation. It should be noted that while passive, this sampling procedure necessarily limits generalizable statements to those hackers who actively scan the subnets of this particular American university and conduct brute force attacks upon finding vulnerable systems. This is not an innocuous limitation, as this necessarily implies that the sample imposed should theoretically retain a low perceived certainty of detection and/or receiving sanctions due to computer-

[5] The linux command to call this information is "ps." This is analogous to bringing up the task manager and looking at the processes tab within a Windows OS environment.

focused crimes. This is not far from reality as the actual likelihood of receiving official sanctions for scanning activity and potentially brute forcing entry to a single target system is believed to be low (McQuade 2006). Further, this sample may not generalize directly to the entire population of hackers, as these computers are targets of opportunity as opposed to targets of choice. Hackers that are more sophisticated and looking for specialized documents or information would not necessarily attempt to enter a linux-based device that retains no inherent sensitive information or utility beyond its computing power. However, there is reason to believe that the majority of hackers are looking for targets of opportunity as opposed to targets of choice, which should limit the impact of this sampling bias (McQuade 2006).

3.4 Data

The collected data from this experiment were stored using three separate tables within a SQL[6] database: “deployment,” “session,” and “keystroke.” The deployment table retains information on the experimental condition assigned to each target computer, the IP address of the intruder, the country of origin of the intruder, the time of the deployment, the login credentials utilized to gain access, and a unique identifier for each deployment to join the deployment table with the session table. The session table retains information on all of the system trespassing incidents recorded for the target computers following successful deployments. This table has information on the start and end times for each session, the IP address of the intruder during the session, the country of origin of the intruder during the session, a unique identifier for each session to join the session table with the keystroke table, the login credentials utilized to gain access to the system,

[6] SQL is a database programming language

and a key indicative of the deployment to which a session belongs. The keystroke table retains information on all of the commands entered into the system during system trespassing incidents. This table contains the unique id associated with each session and every command typed into the system during the respective session as well as indicators for the screen that the command is attached to.[7]

These tables were joined using SQL queries and manipulated by hand to enable the proposed analysis. Sessions had to be sequenced by hand to properly ascertain the ordering of sessions for a given deployment. This operation was beyond the scope of the capabilities of SQL. Numerous checks were employed to ensure and verify that all data managed by hand was authentic and accurate. The detailed nature of this process revealed three observations for which there were technical errors in the data present within this database due to random factors resulting in negative session duration and multiple keystroke session entries; these observations have been excluded from the present analysis. Information on the specific commands employed by the intruders during these sessions is beyond the scope of this analysis and not present within this resulting dataset. This keystroke data in its raw form is not conducive to statistical analysis and will require extensive programming support from computer scientists to rework the data into a usable format.

[7] It is possible for an intruder to open and type into additional screens during a session on a linux system. One example pertains to the command "vim," which opens up a text editor window. The intruder can then script and insert commands into a file, which is then inserted into system memory. However, the commands inserted into this file are not directly executed until the file in which they are entered is run. Information indicating whether commands are directly entered into the system or into a text editor is available within the keystroke table.

3.5 Outcome Measures

Several terms must be explicitly defined for the purposes of the present analysis. A deployment occurs when an attacker launches a target computer following a successful brute force entry. Each time the intruder returns and logs into the system is considered to be a system trespassing incident or a session. A keystroke session is a system trespassing incident during which the hacker typed commands into the system. A keystroke session is indicative of active engagement with the system that goes beyond simply connecting to the system via SSH.[8]

The relevant outcome measure for the present analysis pertains to whether any keystrokes are present and recorded for a given system trespassing session. The presence of any keystrokes in a system trespassing incident is coded as a binary variable with (1) indicating the presence of at least one keystroke and (0) else. This dichotomization is assessed within a sequence analysis of the sessions for a given deployment conditioning on whether a keystroke session was present on prior sessions for the respective deployment. This analysis is kept at this level of abstraction as to simply whether any keystrokes are recorded for the given session.

4. Methods

4.1 Descriptive Statistics

[INSERT TABLE 1 HERE]

[INSERT TABLE 2 HERE]

Descriptive information pertaining to the intruder's actions following the deployment are presented in Tables 1 and 2. As depicted in both tables, 1059 honeypots

[8] SSH refers to secure shell, which is one of several means by which remote connections can be established.

were deployed over the seven-month period of deployments. The total number of deployments across conditions ranges from 250 deployments in the process only condition to 275 deployments in the control condition. Over 60 percent of the deployments retained at least one system trespassing incident during the respective thirty-day data collection period for each deployment. The average number of sessions per deployment with sessions is 4.458 with minimal differences between those honeypots that retained a surveillance banner versus those honeypots that did not retain a surveillance banner. There was a total of 2942 sessions with 1318 of these sessions retaining a keystroke. As depicted in both tables, the average keystroke ratios (as previously defined) are consistently distributed according to whether there is a banner on the system. Both of the conditions with the surveillance banner have average keystroke ratios between 0.449 and 0.468 keystroke sessions per system trespassing incident. Both of the conditions without the surveillance banner retain average keystroke ratios of roughly 0.49 keystroke sessions per system trespassing incidents.

As the data utilized for this analysis come from a randomized controlled trial, it can be assumed that intruder characteristics, both observed and unobserved, will be balanced in expectation across conditions. However, information is still observed regarding the intruder's choice of IP address and username in brute forcing entry to a target computer within this research design and should be assessed for balance across conditions. Significant differences are not found according to the country of origin across conditions (p=0.446) nor for the username employed during the brute force entry (p=0.498). This balance over observables lends support to the assumption that pretreatment observables will be balanced in expectation across the experimental conditions due to randomization.

4.2 Analytic Method

Preliminary investigations suggest that less than 5% of the intruders actually investigated the processes on the system, with this rate dropping to below 2% for those systems actually retaining the process condition. As such, the surveillance process should theoretically have a diminished to negligible effect upon subsequent behavioral outcomes. This led to the decision of this thesis to consolidate the four conditions from the 2x2 factor design to look solely at the effect of the presence of the surveillance banner to take advantage of the added statistical power. Concordantly, the control group and the process only group were consolidated into a single no banner group for subsequent analysis, while the banner only and banner and process groups will be consolidated into a banner group. Sensitivity analyses were conducted to ensure that this methodological decision did not bias the ultimate results.

Each of the hypotheses discussed within this work can be assessed and directly tested with a test of a difference in proportions according to whether the intruder was assigned to a condition retaining a surveillance banner. There are no statistical controls necessary for assessing hypothesis 1, hypothesis 2, and hypothesis 3, as randomization occurred at the point of deployment and any differences observed following the deployment can be discussed as resulting from the imposed treatment. One notable exception to this is with regard to testing hypotheses 3a and 3b wherein it is possible for the intruder to have typed commands into the computer on the session prior to the session presently under study. These previously typed commands may have altered the system environment, and, as such, it is necessary to conduct conditional tests of proportions for the second session, based upon whether or not there had been any commands entered onto the system during the first session. This implies two separate tests of proportions – one for those intruders who had a keystroke on the first session and one for those that did not. Sensitivity analyses were conducted for all aforementioned tests to ensure the decision to consolidate the data according to the presence of a surveillance banner does not impart any undue bias upon subsequent results.

5. Findings

5.1 Absolute Deterrence and Recidivism

The results of a test of the first hypothesis pertaining to absolute deterrence are presented in Table 3. Recidivism is defined as conducting at least one system trespassing incident following the initial display of this surveillance banner. The probability of recidivism amongst those attackers that did not receive a surveillance banner is 0.24 while the probability of recidivism amongst those attackers that did receive a surveillance banner is 0.22. This small difference is not found to be statistically significant (p=0.31) and lends tentative support for the first hypothesis, as I fail to reject the null hypothesis that there is no absolute deterrence effect of this surveillance banner upon the likelihood of recidivism. This first hypothesis of a null effect cannot be proven directly, but the evidence presented in Table 3 does not run counter to expectations of a null effect with regard to absolute deterrence, as the surveillance banner is only visible during the system trespassing incident and no sanctions were imposed within the experiment under study. This finding is not sensitive to the duration of the first system trespassing incident, as identified from several sensitivity analyses not reported in this thesis.

[INSERT TABLE 3 HERE]

5.2 First Session Analysis

Table 4 depicts the findings pertaining to the analysis of a restrictive deterrent effect for the first session on a target computer. Those intruders who received a surveillance banner escalated their offending to incorporate keystroke commands 36.4% of the time while those intruders who did not receive a surveillance banner escalated their offending 38.7% of the time. This reduction of 2.3% due to the presence of a banner is not found to be statistically significant at any level of statistical significance, despite the estimated direction being consistent with a priori expectations. However, sensitivity analyses reveal that this estimate may be biased downward due to the presence of sessions that were too short for an offender to both read the banner and input keystroke commands into the system. Zero second duration and other short sessions could be prompted by connectivity issues, server issues, and/or human error in closing the window prematurely. The distribution of these connectivity issues theoretically should not be related to assignment to receive a surveillance banner, especially with regard to zero second duration sessions as the intruder would not have been able to see nor read the surveillance banner in these instances. These results that remove these potentially downward biasing observations are depicted in columns 3 and 4 of Table 4.

Upon limiting the sample to those attackers that had a first session that was greater than zero seconds in duration, the estimated reduction in the likelihood of a keystroke on the first session due to the presence of a surveillance banner more than doubles to 4.6%. This effect increases further to 5.4% when assessing those attackers that had a first session that was greater than five seconds in duration, which should allot enough time for an attacker to read and process the assigned surveillance banner. This reduction of 5.4% retains a probability of having occurred due to chance of 0.0949 if there was truly no restrictive deterrent process at play. While not statistically significant by conventional standards, this sequence of consistent findings with regard to the restrictive deterrent effect of this surveillance banner in reducing the hacker's likelihood of escalating their offending to incorporate keystroke commands offers some support toward the second hypothesis.

[INSERT TABLE 4 HERE]

5.3 Second Session Analysis

The analysis pertaining to the second session is not as straightforward as the above analyses that simply involved tests of differences in proportions. While the banner in the previous tests could be treated as purely exogenous, this is no longer the case when assessing the second session for an attacker's offending progression. Attackers during the second session can differ from one another in this sample on the basis of whether or not they received a surveillance banner in addition to whether or not they actively engaged with the system by entering keystroke commands during the first session. These previously entered keystrokes can be reasoned to provoke biases in assessing the effect of the surveillance banner on the probability of any keystrokes on the second session. As such, this analysis assesses the effect of the surveillance banner on the probability of a keystroke on the second session for (1) the total sample of attackers having at least two sessions, (2) the subsample of this total that did not enter any keystrokes during the first session, and (3) the subsample of this total that did enter keystrokes during the first session. These results are depicted in Table 5.

The overall, or unconditioned, reduction in the probability of any keystrokes being entered into the system on the second session as a result of this surveillance banner is 4.5%, which is not found to be statistically significant with a test of differences in proportions (p=0.15). However, in directly assessing hypothesis 3a regarding the conditional subsample of those attackers that did not escalate their offending during the first session on the target computer, this restrictive deterrent effect nearly doubles to a roughly 8.7% reduction that is nearly statistically significant by conventional standards (p=0.053). As such, this thesis rejects the null hypothesis that the surveillance banner has no effect upon the likelihood of an offender with no prior escalation during the first session escalating their offending during the second session. Put more simply, the surveillance banner appears to have a restrictive deterrent effect for this subsample of the sample in this analysis, as the probability of an attacker within this group escalating their offending during the second session decreases from 0.468 to 0.38 in the presence of the surveillance banner.

This thesis' attention is then turned to the subsample of attackers that did enter keystrokes into the system during the first session. For this subsample, there is no deterrent effect as the point estimate for the potential reduction due to the surveillance banner is in the incorrect direction with an increase in the probability of an escalation on the second session for this subsample from 63% to roughly 67% in the presence of the surveillance banner. This result lends support to the aforementioned statement that the presence of a keystroke on the first session could be biasing results when evaluating the potential deterrent effect of the surveillance banner on the second session. This apparent disparity in treatment effects for these two subsamples must be treated separately to accurately portray the potential effect of the surveillance banner upon prospective hackers. In sum, some support is found for hypothesis 3b regarding the null effect for the surveillance banner on the probability of a keystroke on the second session for those attackers that previously entered a keystroke during the first session.

[INSERT TABLE 5 HERE]

5.4 Global Keystroke Ratio

The global keystroke ratio, as previously defined, is the ratio of the number of sessions retaining at least one keystroke to the total number of sessions for a given honeypot's 30 day duration within this experimental design. The ratio for those attackers in the control group is 0.495 sessions with keystrokes per session recorded while the ratio for those attackers in the treatment group is 0.459 sessions with keystrokes per session recorded. This is a reduction of 0.035 sessions with keystrokes per session recorded that is not found to be statistically significant by conventional standards with a p-value of 0.1017. A more meaningful interpretation of this effect would be a reduction in the severity of 35 sessions that would otherwise escalate out of a total of 1000 sessions. Although this result does not provide enough statistical evidence to reject the null hypothesis of no effect of the surveillance banner upon the global keystroke ratio, the direction of the estimated effect is consistent with the previous analyses that produced a restrictive deterrent effect for at least part of the sample under study.

[INSERT TABLE 6 HERE]

6. Discussion

This thesis expands upon the prior research into both system trespassing and restrictive deterrence by directly testing for restrictive deterrence with regard to the severity of an attacker’s actions on a target system in the presence or absence of a surveillance banner. Based upon these data and subsequent analyses, it would appear that this surveillance banner had a discernible effect upon reducing the likelihood of an attacker escalating toward actively engaging with the system through the medium of keystrokes.

However, in line with a priori expectations, there was no discernible absolute deterrent effect within these data; the attackers continued to trespass on the system regardless of the display of this surveillance banner. Additionally, this assessment of absolute deterrence is inadequate in accordance with Gibbs’ (1975) conception of absolute deterrence as the complete abstention of a particular type of criminal conduct on the part of an offender. This dataset only retains measures of an intruder’s behavior in relation to a specific target system; there is no record or means of tracking the intruder to assess whether the intruder modified his or her practice on other target computers in response to the display of such a surveillance banner on the system within this experiment. However, given the failure to identify a statistically significant effect of this surveillance banner in terms of absolute deterrence on the system retaining the surveillance banner, it is unlikely for there to be an effect on the intruder’s behavior in relation to other target systems.

The effects observed for the first session are promising from a policy perspective as a delay of keystroke usage to later sessions constitutes added time for an IT practitioner to respond to and ameliorate the compromise on the target system. As the previous discussion and Table 4 relate, the effect for the first session is not statistically significant by conventional standards, but the point estimates are in the hypothesized direction with the presence of very short duration sessions serving as a confound in the analysis. The consistency in the direction of the observed difference in addition to substantive reasons justifying this conditional analysis[9] leads to this thesis' stance of taking a less conservative benchmark with regard to statistical significance and interpreting the observed difference between the point estimates in this context as a restrictive deterrent effect.

The unconditional probability of any keystrokes being entered on the second session presents a likewise confounded analysis of a potential deterrent effect. Unlike the otherwise experimentally sound analysis of the first session, the second session analysis must condition on the presence of any keystrokes entered on the first session due to the potential for both intruders and systems linked to previous keystrokes being different from intruders and systems that are not linked to previous keystrokes. The data bears this out, as those intruders who do not retain any previous keystrokes entered on the first session for a given target system are subsequently deterred from escalating their offending on the second session. However, those intruders who entered any keystrokes on the first session are not deterred from entering subsequent keystrokes into the system during the second session. This lends credence to a priori expectations that those attackers who enter keystrokes during the first session and subsequently receive no sanctions should update their perception of the certainty of detection to a lower value, which is in keeping with the general framework afforded by Loughran and colleagues (2012). This updating procedure is not directly observed by this thesis, however, and is

[9] These substantive issues include connectivity issues, time to read the banner, and time to type into the system. Short sessions of less than 5 seconds do not afford the same level of opportunity for intruders to interact with the system and respond to the surveillance banner.

merely noted as a potential mechanism underlying the observed effect for this subsample of the data. These individuals who entered keystrokes during the first session were also far more likely to engage the system with additional keystrokes than those individuals who did not enter keystrokes during the first session regardless of assigned treatment condition. This lends further support toward some form of splitting occurring within the sample of intruders, which is revealed in relation to temporally earlier behavior on a target system. This heterogeneity within the population of intruders, while not a direct target of this analysis, is worthy of independent study and analysis, and retains substantial potential toward directing future policy and theoretical developments with regard to the efficacy of deterrence in cyberspace.

The final test and result of the above analyses pertained to the global keystroke ratio of sessions wherein at least one keystroke was entered into the system as compared to the total number of sessions for a given deployment. This result was not statistically significant at any conventional level of statistical significance, but the direction of the effect is consistent with previous analyses. This is consistent with a priori expectations as the previously observed effects for the first and second sessions are inherently included within this more global analysis of the effect of a surveillance banner on curtailing an individual's decision to actively engage with the system through entering keystrokes.

These results retain implications for future theoretical development within the deterrence tradition both with regard to the severity component of restrictive deterrence and the continued application of deterrence principles to cyberspace. These results coupled with the findings of Maimon and colleagues (2014) provide the beginning of an empirical foundation upon which to support the continued assessment of and application of deterrence to cyberspace. Absolute deterrence may not be the appropriate means by which to assess the efficacy of the deterrence process in cyberspace, but restrictive deterrence presents itself as a viable alternative. With the relative ease with which individuals can conduct and commit computer-focused crimes, altering the pathways of these criminal behavior patterns may prove to be a fruitful course of action for deterrence theorists and cybercrime researchers. These results contribute to the ongoing discussion in support of recognizing the human element inherent to cybercriminals as an additional modicum with which to prevent and modify cybercriminal activity.

Additionally, policy implementation based upon these results is straightforward, cheap, and low-risk, as SSH banners similar to the one employed in this experiment can be applied with minimal cost and effort. This thesis does not find any potential negative repercussions to altering an entry banner to a computer system in this fashion, as the system conditions are not physically modified to otherwise prevent legitimate user activity. However, replication should still be sought before these results contribute directly to the enactment of policy protocols. Additionally, experimentation with alternative wordings of the presented banner should be conducted to separate the effect of the content of the banner from the presentation of a banner upon entry to a computer system.

7. Conclusion

This work sought to address numerous empirical gaps throughout the literature pertaining to both theoretical and policy development. An absolute deterrent effect for a surveillance banner was not detected, but a restrictive deterrent effect appeared throughout the data in assessing the three hypotheses predicting a deterrent effect. These findings support the restrictive deterrence hypothesis, specifically the severity component, with regard to reducing the likelihood and prevalence of system trespassing incidents containing keystrokes.

7.1 Limitations and Future Directions

This work is not without limitation, though, as the results are not necessarily generalizable beyond the university setting in which the experiment took place. Further research should investigate this deterrence process in alternate venues including government and/or commercial entities toward assessing whether this finding is externally valid. The experiment from which data was obtained was limited to computers running a Linux-based operating system with SSH as the means of entry. Further research is necessary to ascertain whether these results bear fruit toward reducing the severity of attacks targeting systems using alternate operating systems, system configurations, and/or means of access. This experiment made an assumption that intruders would be English-literate in the construction of a banner that was displayed in English. While this limitation would only serve to downward bias any obtained results from such an analysis, additional analyses identifying offenders directly toward ascertaining their knowledge directly would be beneficial to the field. Along this vein, additional research that can more adequately identify intruders toward administering surveys and obtaining demographic and related personal characteristics beyond a proficiency in English is sorely lacking and would prove invaluable to future efforts in targeting innovations and policies derived from social sciences toward cyber security and cybercrime prevention. Additionally, further research should extend this analysis with regard to the potential for differential deterrence and further identifying heterogeneity within the population of system trespassers. Similar analyses should be conducted with regard to phenomena occurring within the physical realm to ascertain whether deterrence cues have an effect in reducing the severity of criminal behavior; the potential policy implications of such analyses in terms of crime reduction and harm reduction are immediate.

Appendix – Tables

Table 1: Descriptive Statistics

Condition	Deployments	Deployments with Sessions	Sessions	Average sessions per deployment	Sessions with Keystrokes	Average Keystroke Ratio
Control	275	172	805	4.68 (5.667)	385	0.492 (0.366)
Banner	263	155	641	4.135 (3.764)	255	0.449 (0.366)
Process	250	164	686	4.183 (3.5)	311	0.497 (0.343)
Both	271	169	810	4.793 (4.505)	367	0.468 (0.358)
Total	1059	660	2942	4.458 (4.465)	1318	0.477 (0.358)

Note: Standard deviations are presented in parentheses where appropriate.

Table 2: Descriptive Statistics by Banner

Type	Deployments	Deployments with Sessions	Sessions	Average sessions per deployment	Sessions with Keystrokes	Average Keystroke Ratio
No Banner	525	336	1491	4.438 (4.734)	696	0.495 (0.354)
Banner	534	324	1451	4.478 (4.174)	622	0.459 (0.362)
Total	1059	660	2942	4.458 (4.465)	1318	0.477 (0.358)

Note: Standard deviations are presented in parentheses where appropriate.

Table 3: Testing Hypothesis 1 – Absolute Deterrence and Recidivism

Treatment	Probability of Recidivism
No Banner	0.241 (0.023)
Banner	0.222 (0.023)
Difference	0.019 (0.033)
z-statistic	0.5737
p-value	0.3091

Note: Standard deviations are presented in parentheses where appropriate.

Table 4: Testing Hypothesis 2 – Probability of any Keystrokes on the First Session with Varying Restrictions on First Session Duration

Treatment	With No Restriction	First Session Longer Than Zero Seconds	First Session Longer Than Five Seconds
No Banner	0.387 (0.027)	0.432 (0.029)	0.443 (0.029)
Banner	0.364 (0.027)	0.386 (0.028)	0.39 (0.029)
Difference	0.023 (0.038)	0.046 (0.04)	0.054 (0.041)
z-statistic	0.60	1.16	1.31
p-value	0.2735	0.1231	0.0949
n	660	607	581

Note: Standard deviations are presented in parentheses where appropriate.

Table 5: Testing Hypotheses 3a and 3b – Probability of any Keystrokes on the Second Session Conditional on Prior Keystrokes on First Session

Treatment	Unconditional Probability	Conditioned on **NO Prior Keystrokes** on First Session	Conditioned on **Prior Keystrokes** on First Session
No Banner	0.522 (0.031)	0.468 (0.038)	0.631 (0.053)
Banner	0.476 (0.031)	0.381 (0.037)	0.667 (0.051)
Difference	0.045 (0.044)	0.087 (0.053)	-0.036 (0.074)
z-statistic	1.02	1.62	-0.48
p-value	0.1534	0.0528	0.6861
n	507	339	168

Note: Standard deviations are presented in parentheses where appropriate.

Table 6: Testing Hypothesis 4 – Global Keystroke Ratio

Treatment	Ratio of Keystroke Sessions to Total Sessions
No Banner	0.495 (0.019)
Banner	0.459 (0.02)
Difference	0.035 (0.028)
t-statistic	1.2733
p-value	0.1017

Note: Standard deviations are presented in parentheses where appropriate.

Bibliography

Anderson, James P. 1980. "Computer Security Threat Monitoring and Surveillance." Technical Report, James P Anderson Co., Fort Washington.

Anwar, Shamena and Thomas A. Loughran. 2011. "Testing a Bayesian Learning Theory of Deterrence among Serious Juvenile Offenders." *Criminology* 49: 667-698.

Beauregard, Eric and Martin Bouchard. 2010. "Cleaning Up Your Act: Forensic Awareness As a Detection Avoidance Strategy." *Journal of Criminal Justice* 38: 1160-1166.

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

Becker, Gary S. 1968. "Crime and Punishment: An Economic Approach." *Journal of Political Economy* 76:169-217.

Bentham, Jeremy. 1970 [1785]. *An Introduction to the Principles of Morals and Legislation*. New York: 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. 2011. "Can information warfare be deterred?" In *Information Age Anthology, Volume III: The Information Age Military*, eds. David S. Alberts and Daniel S. Papp. Washington, DC: Command and Control Research Program.

Furnell, Steven. 2002. *Cybercrime: Vandalizing the Information Society.* Boston, MA: Addison- Wesley.

Gibbs, Jack. 1975. *Crime, Punishment, and Deterrence*. Elsevier Scientific Publishing Company.

Geerken, Michael R. and Walter R. Gove. 1975. "Deterrence: Some Theoretical Considerations." *Law and Society Review* 9: 497-513.

Goodman, Will. 2010. "Cyber deterrence: Tougher in theory than in practice?" *Strageic Studies Quarterly* (Fall): 102-135.

Jacobs, Bruce A. 1993. "Undercover Deception Clues: A Case of Restrictive Deterrence." *Criminology* 31: 281-299.

Jacobs, Bruce A. 1996A. "Crack Dealers Apprehesion Avoidance Techniques: A Case of Restrictive Deterrence." *Justice Quarterly* 13: 359-381.

Jacobs, Bruce A. 1996b. “Crack Dealers and Restrictive Deterrence: Identifying Narcs.” *Criminology* 34: 409- 431.

Jacobs, Bruce A and Jody Miller. 1998. “Crack Dealing, Gender and Arrest Avoidance.” *Social Problems* 45: 550-569.

Jacobs, Bruce A. 2010. “Deterrence and Deterrability.” *Criminology* 48: 417-441.

Jacobs, Bruce A and Michael Cherbonneau. 2012. “Auto Theft and Restrictive Deterrence.” *Justice Quarterly* 1-24

Kahneman, Daniel and Amos Tversky. 1979. “Prospect Theory: An Analysis of Decision under Risk.” *Econometrica* 47: 263-191.

Loughran, Thomas A., Greg Pogarsky, Alex R. Piquero and Raymond Paternoster. 2012. “Re-Examining the Functional Form of the Certainty Effect in Deterrence Theory.” *Justice Quarterly* 29: 712-741.

Maimon, David, Mariel Alper, Bertrand Sobesto and Michel Cukier. 2014. “Restrictive Deterrent Effect of a Warning Banner in an Attacked Computer System.” *Criminology* 52: 33-59.

Mackey, David. 2003. *Web Security for Network and System Administrators*. Cengage Learning.

McQuade, Samuel C. 2006. *Understanding and Managing Cybercrime*. Pearson Education INC.

Nagin, Daniel S., and Greg Pogarsky. 2003. “An experimental investigation of deterrence: Cheating, self-serving bias, and impulsivity.” *Criminology* 41: 167-194.

Pratt, Travis C., Francis T. Cullen, Kristie R. Blevens, Leah E. Daigle, and Tamara D. Madensen. 2006. “The empirical status of deterrence theory: A meta-analysis.” In *Taking stock: The status of criminological theory*, eds. Francis T. Cullen, John Paul Wright, and Kristie R. Blevins, 367–96. New Brunswick: Transaction Publishers.

Paternoster, Raymond. 1987. “The deterrent effect of the perceived certainty and severity of punishment: A review of the evidence and issues.” *Justice Quarterly* 4:173–217.

Paternoster, Raymond. 1989. “Absolute and Restrictive Deterrence in a Panel of Youth: Explaining the Onset, Persistence/Desistence and Frequency of Delinquent Offending.” *Social Problems* 36: 289-309.

Ponemon Institute. 2011. “Second Annual Cost of Cyber Crime Study; Benchmark Study of U.S. Companies.” Research Report available at <http://www.hpenterprisesecurity.com/collateral/report/2011_Cost_of_Cyber_Crime_Study_August.pdf>

Ratcliffe, Jerry, Travis Taniguchi and Ralph Taylor. 2009. “The Crime Reduction Effects of Public CCTV Cameras: A Multi-Method Spatial Approach.” *Justice Quarterly* 26: 746-770.

Salles-Loustau, G., Berthier, R., Collange, E., Sobesto, B., and Cukier, M. 2011. “Characterizing Attackers and Attacks: An Empirical Study.” *Proc. 17th IEEE Pacific Rim International Symposium on Dependable Computing (PRDC 2011)*, Pasadena, California, December 2011.

Stafford, Mark and Mark Warr. 1993. “A reconceptualization of General and Specific Deterrence.” *Journal of Research in Crime and Delinquency* 30: 123-135.

Skinner, W.F., and Fream, A.M. 1997. “A Social Learning Theory Analysis of Computer Crime among College Students.” *Journal of Research in Crime and Delinquency* 34:495-518.

Sobesto, Bertrand, Michel Cukier and David Maimon. 2012. “Are Computer Focused Crimes Impacted by System Configurations? An Empirical Study.” 23rd IEEE International Symposium on Software Reliability Engineering.

Spitzner, L. 2002. *Honeypots: Tracking Hackers.* Addison-Wesley Longman Publishing Co., Inc., Boston, MA.

Stallings, William. 2005. *Wireless Communications and Networks*. Pearson Prentice Hall, NJ.

Welsh, Brandon and David Farrington. 2008 “Effects of Closed Circuit Television Surveillance on Crime.” *Campbell Systematic Reviews* 17.

Whitman, Michael E. 2003. “Enemy at the Gate: Threats to Information Security.” *Communication of the ACM* 46: 91-95.

Zhang, Feng, Shijie Zhou, Zhiguang Qin, and Jinde Liu. 2003. "Honeypot: a supplemented active defense system for network security." In *Parallel and Distributed Computing, Applications and Technologies, 2003. PDCAT'2003. Proceedings of the Fourth International Conference on*, pp. 231-235. IEEE.