ABSTRACT

Title of Thesis: WHO ARE YOU? EXAMINING THE RESTRICTIVE DETERRENT IMPACTS OF CONVERSATIONAL MESSAGES ON SYSTEM TRESPASSERS

Megan Almeda Smith, Master of Arts, 2017

Thesis directed by: Associate Professor David Maimon
Department of Criminology and Criminal Justice

With cyber attacks and computer hacking becoming a growing concern in the United States and abroad, measures that could reduce the frequency of these crimes and mitigate their consequences are of a growing interest. This thesis adds to the growing body of criminological literature on cybercrime by examining how computer system trespassers may alter their online behaviors when their presence has been detected on the system. Using an experimental design to examine how detection cues presented as a series of messages in an attacked computer system influence the trespassers decisions' to enter specific commands to influence or identify aspects of the system, it was hypothesized that there would be differences in the commands entered to attack and explore the system, and in IP Addresses used between treatment and control groups. The results do not support most of the hypotheses, but several implications are still drawn.

WHO ARE YOU? EXAMINING THE RESTRICTIVE DETERRENT IMPACTS OF CONVERSATIONAL MESSAGES ON SYSTEM TRESPASSERS.

by

Megan Almeda Smith

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
2017

Advisory Committee:
Associate Professor David Maimon, Chair
Professor Sally Simpson
Assistant Professor Sarah Tahamont

Table of Contents

List of Figures

List of Tables

1. INTRODUCTION

Cybercrime is a broad concept encompassing numerous illegal activities occurring online or on computer systems (Yar, 2006). It includes behaviors such as system hacking, copying or removing information from a computer, and adjusting or crippling a system entirely (Wilson, 2001; Yar, 2006). This form of crime has become a growing concern in both the United States and abroad (Goodman, 2010; Holt & Bossler, 2013; Morris & Blackburn, 2009; Young & Zhang, 2007; Young, Zhang, & Prybutok, 2007). One form of cybercrime, system trespassing, is of particular concern as it involves unauthorized hacking or access onto a computer system in order to obtain information or alter aspects of the system (Yar, 2006). This has the potential to be very problematic for both citizens and companies considering the harm that offenders trespassing could cause should they be able to breach a system and access or alter the information it holds. CSID, a security division with Experian, notes that in 2014 approximately 8.5 million people were victims of reported illegal data breaches, indicating numerous likely victims of system trespassing who have had their personal or financial data compromised (CSID, 2017). Financial losses resulting from cyber attacks and breaches can approach the billions of dollars in the United States alone (McAffee, 2013). As such, system trespassing has the potential to cause great harm in both the form of confidential information breaches and financial losses. Measures that could potentially reduce the number of these attacks or, at the very least, that could alter the actions or behaviors of trespassers are of growing consideration as possible ways to help mitigate the damage caused.

This thesis seeks to address this challenge by examining how system trespassers respond to the presentation of a series of messages displayed on the computer indicating that they have been detected on the system. Even with a growing body of literature on cybercrime (D'Arcy,Hovav, & Galleta, 2009; Higgins, Wilson, & Fell, 2005; Maimon, et al., 2014; Maimon, et al., 2015; Wilson et al. 2015) it is difficult to determine how system trespassers will respond to notifications that they are being monitored, but it is theoretically plausible that these messages will result in trespassers changing their behavior given that the messages do indicate that they have been detected. Specifically, they may adjust what commands they enter into the computer system, or they may even change their identifying Internet Protocol (IP) Address in order to return to the system later under a different address that has not been detected. Utilizing a restrictive deterrence perspective, this thesis seeks to explore instances of unauthorized system trespassing and how implying detection towards the perpetrators may influence how they proceed with the trespassing event. It will then contribute further to the growing body of literature on restrictive deterrence and cybercrime (Maimon, et al., 2014; Maimon, et al., 2015; Testa et al., 2017; Wilson et al. 2015), through which the information gained can help expand our theoretical ideas of the behavior of system trespassers, and can influence further discussions on possible preventative measures for system operators to take to reduce instances of trespassing.

2. BACKGROUND

2.1 System Trespassing

System, or computer, trespassing involves entering another's computer system without permission (Maimon et al., 2014; Maimon et al., 2015; Wilson, 2001; Yar, 2006). This includes illegal hacking in broader terms, as well as accessing systems to cause more serious damage through the use of installed malware or by obtaining confidential files and information stored on the system (Jordan & Taylor, 1998; Marcum et al., 2014; Morris & Blackburn, 2009; Taylor, 1999). System trespassing generally involves several stages: reconnaissance, network scanning, collecting information to aid in guessing passwords, exploiting the systems' weak areas, and ending with the completion of the intrusion and beginning of the damage or other illicit activities (Maimon et al., 2015; Wilson, 2001). Figure 1 portrays the progression of these events. The beginning stages of reconnaissance and scanning involve would-be trespassers searching online for any information to better understand the targeted system and its users in order to both increase their chances of a successful attack and to identify any valuable information that they may want. The information gained during these stages also allows for trespassers to move to the next stages of their attack, which is again searching online and around the system to identify weak points and information pertinent to passwords and login information. By doing this, they may find ways for quicker and easier access, which could then reduce their chances of being caught or identified before they can complete their intrusion. The real damage comes from the final stage, which is the execution of the attack that occurs once full access to the system is obtained. More damage can occur here as trespassers will attempt to gather information by entering different command keystrokes, install or create malware, or conduct other illicit and potentially damaging acts on a system (Maimon et al., 2015; Wilson, 2001). Completions of cyber-attacks can be costly with damages from malware or ransomware[1] or other attacks potentially reaching in the billions of dollars annually (Norton, 2016).

[1] A form of information or online item trespassers use to obtain a ransom.

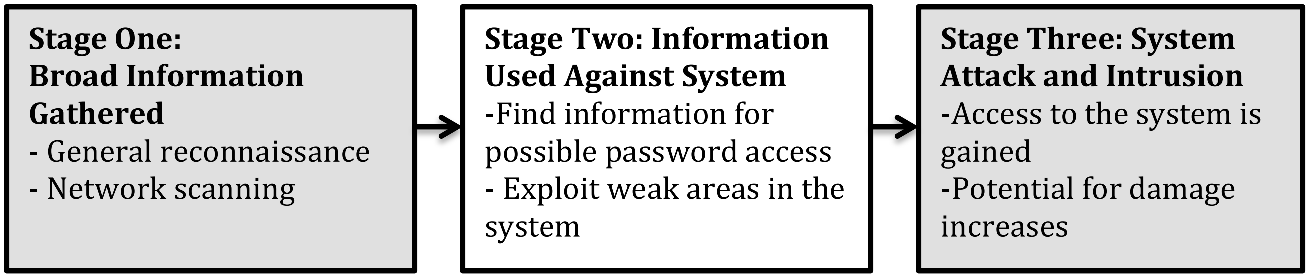
Stage One:
Broad Information Gathered
- General reconnaissance
- Network scanning
Stage Two: Information Used Against System
-Find information for possible password access
- Exploit weak areas in the system
Stage Three: System Attack and Intrusion
-Access to the system is gained
-Potential for damage increases

Figure 1: Stages of a System Trespassing Event

2.2 Theoretical Perspective

While system trespassing can result in serious consequences, preventing attacks is by no means straightforward. However, the deterrence framework in criminology offers some insights on ways to prevent or reduce its occurrence. The concept of deterrence within criminology largely comes from the late 18th century work of Beccaria (1963 [1764]) and Bentham (1970 [1785]), whose writings developed ideas about punishment and how it may deter individuals from engaging in illegal or immoral acts. The deterrence perspective within criminology is often considered broad and sometimes vague (Gibbs, 1975), but it is focused primarily on the perceived indications of the certainty, severity, and celerity of punishment (Gibbs, 1975; Jacobs, 2010; Loughran et al., 2011; Nagin & Paternoster, 1991; Pogarsky, 2002; Wenzel, 2004). Perceptions are of particular importance in many cases given that possible offenders may weigh the costs and benefits of committing a crime. If the perceived costs, whether informal or formal, are deemed too high by the individual offender then they will theoretically avoid committing the offense. Not everyone can be completely deterred however, as some will still engage in offenses but will alter their behavior to an extent in order to reduce their chances of apprehension.

Gibbs (1975) develops this point by identifying two subsets of deterrence:

absolute and restrictive deterrence. Absolute deterrence results in a potential offender being fully deterred from engaging in a crime, resulting in a complete refrain from ever engaging in criminal activities at any point or a complete desistence from any future offending if they have already been caught and punished. Restrictive deterrence however, is of primary interest in this study as we are examining system trespassers who have already illegally entered a system, but who may yet change their offending behaviors. Possible adjustments to behavior highlighted by Gibbs (1975), and expended upon by Jacobs (2010) include reducing the frequency of offending, offending at times and in locations where the likely risks are low, or by reducing the severity of the offenses committed. All of these changes are significant from a restrictive deterrence perspective as they allow for people to still offend, but in such a way that they reduce their chances of incurring punishment, or even reducing how punitive the sanctions are should they get caught.

These behavioral adjustments that are discussed in the restrictive deterrence perspective are often the results of offenders applying what they have observed or learned in order to best avoid detection. Examples of these changes are seen in Jacobs and Cherbonneau's (2014) interviews with auto thieves. The thieves interviewed noted that they take care in selecting their victims, and that they would also make adjustments to their behavior in order to appear as normal as possible to better avoid apprehension. Jacobs (1993) also discusses restrictive deterrence as it applies to drug dealers by noting how they rely on both physical and non- physical cues to help them identify whether or not it is safe to sell to a specific person. These examples show how offenders can still engage in an offense while also paying enough attention to risks to attempt to avoid punishment. This implies deeper analyses and consideration of risks on the offender's part, but it also implies that offenders can be actively learning different behaviors, which then ties in some elements of learning theory in with restrictive deterrence. This is significant for offenses such as cybercrime given that would-be offenders have to first learn a variety of elements of computer systems before they can cause any harm (Skinner & Fream, 1997). What they learn may then translate to whether or not they respond to certain deterrents or not.

Whether or not an offender is deterred completely or partially can depend on different factors. Pogarsky (2002) argues that some people are more influenced by potential punishments than others, implying that there is variation in the "deterrability" of offenders, an idea which Jacobs (2010) also supports and expands upon by suggesting that some individuals will be deterred while others will not. For those who are deterred in part, there is more risk sensitivity and awareness involved, but not enough to fully deter them from committing an offense. A meta- synthesis conducted by Moeller, Copes and Hochstetler (2016) reveals patterns in the literature on deterrability. They note through this examination that some people are indeed less likely to be deterred completely, but they may still be influenced by risks and thus restrictive deterrence may still apply. These risks can influence offenders to find ways to avoid detection or altering behaviors if it seems they may be caught otherwise.

2.3 Deterrence in Cyber Space

An existing argument on what may impact whether or not punishment deters potential offenders, whether completely or not, relates to perceptions. In particular, focus is on the perceived certainty and severity of punishments, with indications that it is the certainty of punishment that is more important (Loughran et al., 2011; Nagin & Paternoster, 1991). Nagin and Pogarsky (2001) in a study focusing on drunk driving, find through surveys administrated among undergraduate students, find that while severity and certainty of punishment had an effect in some models, certainty of punishment was most consistently related to intention inhibition. In relation to cybercrime, Higgins, Wilson and Fell (2005) and Peace, Galletta, and Thong (2003) find the certainty element is most important for digital piracy deterrence. Certainty of punishment then is clearly important in determining whether or not someone engages in crime. For cybercrime however, this is particularly significant given the anonymity offenders enjoy and the relatively low levels of prosecution (Jordan & Taylor, 1998) which can both heavily reduce offender's perceived certainty. Still, even with increases in certainty, there is likely to be variation in how much is required for different people to be deterred completely, but restrictive deterrence could still be applicable in many cases. In particular, and within this study, merely offering an implied risk of discovery even without mention of punishment may be enough to warrant some changes in the behavior of the trespassers.

In the context of cyberspace and system trespassing, elements of restrictive deterrence could be seen in several ways. Trespassers may take the time to analyze potential targets beforehand to assess the risks of attacking that particular system, or they may look into ways to reduce the likelihood of detection. One possible way to avoid detection in cyberspace is through IP spoofing, where hackers will change their identifying Internet Protocol (IP) Addresses in order to fool others of their identity (Jordan & Taylor, 1998; Yar, 2006). This allows for greater anonymity among system users and trespassers, which could then decrease the perceived likelihood of discovery and punishment, even if they are presented with signals indicating that the system is being monitored. More confident offenders may not view system monitoring as a significant risk if they believe themselves to be anonymous or untraceable.

This anonymity allowed by the inherent nature of the Internet and technology makes catching and punishing trespassers challenging (Holt & Kilger, 2012; Jordan & Taylor, 1998). Given the importance of punishment certainty for deterrence, it could be argued that increasing the perceptions of detection and punishment for system trespassers could improve deterrence and reduce the number of trespassing events. Young and Zhang (2007) make this argument for cyber offenders with findings indicating that by increasing the certainty, or even simply the attacker's perception of it, at any stage of the trespassing event may result in the trespasser concluding their attack before causing any damage. However, an opposing argument exists suggesting that this may not apply to all trespassers as some researchers argue that hackers may actually enjoy the challenge of avoiding detection (Goodman, 2010; Young & Zhang, 2007; Young, Zhang, & Prybutok, 2007), making it difficult then to predict or possibly alter their behavior. This creates dueling positions in regards to deterring system trespassers with the introduction of certain security measures, such as a warning banner or direct messages, possibly deterring some offenders while presenting an enticing challenge to others. It is also possible that offenders may simply be confident that they will not be caught or punished regardless of what messages they are presented with.

It is worth noting this opposing position that argues that the deterrence perspective is not always considered the most effective at explaining cybercrime. Rather, other approaches such as social learning are thought to be more effective. The skills required for almost any form of cybercrime need to be taught or conveyed to would-be offenders. These skills may be self-taught in part, but trespassers will likely receive influence from others who have experience hacking in the past, given the complexity of computer systems (Holt, Burgess, & Bossler, 2010). As such, the impacts that peers may have on trespassers are worth discussing as what is taught or conveyed could shape how they respond to certain scenarios. For example, Marcum and colleagues (2014) found when surveying juveniles that having more deviant peers was related to unauthorized access of social media and emails. Morris and Blackburn (2009) also find in a survey of undergraduate students that peer associations were more impactful on the more serious cyber offenders. It is possible that having peers or closer friends that engage in cybercrime can then influence one to also try and engage in offenses of equal or greater severity.

The presence of harmful system trespassers who find that higher risks provide them with a greater test for their abilities also complicates the deterrence argument, particularly considering how the thrill of avoiding detection in high-risk scenarios may actually encourage further attacks. Stiren and Applegate (2012) note in their survey of inmates at a Work Release Center that higher perceived excitement of an offense increases the likelihood that an individual would engage in it. They did not look at cybercrime perpetrators, but the findings are still noteworthy in suggesting that excitement may be a benefit that can possibly overrule risk. Wood and colleagues (1997) do discuss how certain symbolic or physiological responses to committing a crime may serve to reinforce that behavior. If computer trespassers experience a “high” or a sense of accomplishment from conducting their attacks, then that could reinforce that behavior prompting them to do it again, particularly if they feel that they will not be caught or punished. Indeed, several studies do report computer trespassers feeling accomplished or respected when they complete attacks, regardless of what damage is caused (Jordan & Taylor, 1998; Taylor, 1999; Turgeman-Goldschmidt, 2008), indicating less regard for causing harm and more interest in achieving recognition or a "high".

Finally worth noting is that some self-identified hackers do not consider themselves criminals or dangerous people (Turgeman-Goldschmidt, 2008), but rather they just feel that they are good with computers or that there is no harm in exploring or conducting minor offenses such as downloading music or videos without paying. Others however, are not as benign in their activities. Seebruck (2015) highlights the range of trespasser categories in his discussion on various typologies by offering five differing ideologies that fuel why trespassers offend. These five types are social or political ideology, recreation, profit, revenge or prestige. All but perhaps recreation could drive trespassers to cause great harm, and if they are committed enough to their objective, deterring them may prove rather difficult. For example, if the trespasser is motivated by prestige, presenting them with indications that they have been discovered or detected may actually entice them into continuing their attack in order to prove that they are skilled enough to avoid punishment.

Even with these typologies offering strong motivation for some, conveying that the costs outweigh the rewards to trespassers should still be as important for deterring trespassers as for other offenders. If the costs are seen as high enough, then the deterrence argument should still be effective. Detecting system trespassers and actually demonstrating that there is a higher certainty of being caught may also be enough to outweigh the excitement (Goodman, 2010). If however, trespassers can conduct attacks without being identified or punished, then it is likely that they will not be deterred in any way from trespassing on another system. This is more problematic for trespassers with more malicious motivations, such as revenge or profit, as they are likely to cause more harm, but they also more prone to try and avoid detection in any way. Merely threatening sanctions also may not be enough as it may not be taken seriously by hackers if they have no reason to believe any punishment will actually occur, or if they feel that they will earn more recognition by continuing with their attacks (Young & Zhang, 2007; Young, Zhang, & Prybutok, 2007). As such, the chances of all system trespassers being completely deterred from engaging in future attacks are likely low unless the certainty of punishment can be increased, but they may still be influenced to adjust their behavior if the situation requires it. Even for those who have been identified in the past, they may return to trespassing if they believe that they have learned from past encounters and can more successfully avoid detection in future attacks.

In contrast to some of the aforementioned views, other studies do suggest that restrictive deterrent approach is most appropriate when considering system trespassers, and indeed some studies do examine this and possible ways to affect their behavior. Wilson and colleagues (2015) found that the longer one trespassed on computers where a warning banner advising trespassers to disconnect was displayed, the more likely it was to have an impact on the number of commands that they entered, implying that the threat of detection and punishment suggested from the banner had an effect on behavior. However, if trespassers were not deterred by the banner’s presence in their initial attack, they also appeared to be undeterred in any future attacks, indicating again that not everyone responds to deterrents in the same way, but that there could still be some effects. Maimon and colleagues (2014) also found that the presence of a warning banner had some effects by reducing the length of time trespassers spent on the system. Again however, this effect was not seen with everyone, with the incidents that were terminated within the first five seconds showing similarities on systems with and without the banner, indicating that the banner alone may not have had as strong of an effect. However, these studies still suggest that some trespassers will be deterred by security messages, such as warning banners, even if others are not.

A recent study by Testa and colleagues (2017) also looks at how the presence of a warning banner can affect the behavior of trespassers. In particular, they were interested in those who continue with their attacks even after they have been shown a warning message. In examining how the trespassers utilize their access to the system, the authors found no support for their hypothesis that those seeing the warning message while on the system would spend less time exploring the system and instead would end their attack. This is contrary to some theoretical expectations as the warning message should increase the perceived certainty of detection and thus should result in more trespassers leaving the system. However, those who were on the systems with the warning banners did appear to enter different commands, specifically those that focused on changing the file permissions. This suggests that while they may not have been completely deterred, the banner still had some effect on some trespasser's behavior. Also of noteis that the warning banner's presence appeared to reduce the number of list-file commands and navigation-related commands. These examples suggest that restrictive deterrence may be more applicable for system trespassers depending on what the deterrent is. As such, this thesis utilizes this theoretical perspective approach to identify what, if any, effects are seen when system trespassers are presented with a sign of detection in the form a series of messages attempting to initiate a conversation with them. Specifically, this thesis examines how implied discovery alone may impact the behavior of trespassers.

2.4 The Current Thesis

This thesis looks to further explore the behavior of system trespassers and how restrictive deterrence may play a role in their future behavior by potentially guiding them to alter their online behavior and how they approach their attack. In particular, I am studying whether the detection of hacker's presence on the system influences hackers' probability to (1) type the computer commands designed to gather intelligence, retrieve information and end processes on the attacked system, and (2) influences the volume of repeated system trespassing events originating from the same IP address used in their initial attack.

Reconnaissance and obtaining information are noted as being important steps for system trespassing (Maimon et al, 2015; Wilson, 2001), tying into stage one of the trespassing events. It could be theorized through restrictive deterrence that being presented with a series of messages demonstrating to trespassers that they have been discovered on the system would result in a trespasser not entering any commands, including reconnaissance ones, and instead trying to avoid further detection. As such, the first hypothesis (1a) relating to reconnaissance commands is: *System trespassers that are notified that they are detected on the system will utilize fewer reconnaissance-based commands when conducting their attacks compared to trespassers not receiving any* *message.* However, given the propensity some hackers have for risk-taking, the presentation of a detection message may result in the trespasser working to gather more information on the system. By doing so, they may want to better assess whether or not the messages' presence indicates a legitimate risk, or they may be searching for something that they believe the system may be trying to protect. They may then try to enter more reconnaissance- focused commands to see who or what is contacting them and sending the messages. These messages may also be less of a deterrent and more of a motivator for hackers to gather more intelligence on the system, particularly if prestige is their primary drive. Hypothesis 1b then is as follows: *System trespassers that are notified that they are detected on the system will utilize more reconnaissance-based commands when conducting their attacks compared to trespassers not receiving any message.*

Even after entering or evaluating a system, information and files are often sought by attackers, particularly within stages two and three of system trespassing events (Wilson, 2001; Yar, 2006). Again through restrictive deterrence, it is hoped that the detection messages would reduce the number of commands entered by trespassers as they would ideally want to avoid detection. Even if the trespassers want to attempt downloading or retrieving files or information, the detection messages may deter them from doing so and thus reduce the number of commands used to collect files and information (henceforth referred to as fetch commands). The second hypothesis then (2a) is: *System trespassers that are notified that they are detected on the system will utilize fewer fetch-related commands when conducting their attacks compared to trespassers not receiving any message.* As with the previous hypothesis though, the opposite may also be true. Testa and colleagues (2017) did find that the presence of a warning banner did not necessarily deter trespassers, but it did appear to affect the commands entered, specifically those focused on altering file permissions and those related to system navigation. It is possible that hackers may view the risk of detection as a testament to their abilities and thus they may look to leave a mark of some sign or seek to collect a "trophy" from their intrusion (Jordan and Taylor, 1998). As a result, trespassers may be enticed to gather more information on the system that detects their presence. And so an alternate hypothesis (2b) is presented as: *System trespassers that are notified that they are detected on the system will utilize more fetch-related commands when conducting their attacks compared to trespassers not receiving any message.*

While higher risks may tempt some attackers, others may try to eliminate or reduce those perceived risks if it means that they can resume their attack with less pressure. As such I am also interested in exploring the effect of detection on commands seen in stage three of a trespassing attack that are focused on disabling a system, or parts of it, making them especially dangerous should they deactivate something important to the user. Considering the damage these commands may cause, one would hope that the detection messages would serve as an effective deterrent and reduce the number of times these commands are entered. The third hypothesis (3a) is: *System trespassers that are notified that they are detected on the system will utilize fewer removal or kill commands when conducting their attacks compared to trespassers not receiving any message.* An alternate view however is that these disabling commands could also be used by trespassers who have been detected in an attempt to disable or remove any file or application that they think could be responsible for identifying them. Hypothesis 3b then is as follows: *System trespassers that are notified that they are detected on the system* *will utilize more removal or kill commands when conducting their attacks compared to trespassers not receiving any message.*

Another method hackers may use to avoid detection is through spoofing, or changing, IP Addresses in order to hide or alter their identify (Yar, 2006). Those who are returning to the system, but who have already encountered the detection messages, may elect to change the IP address that they use to avoid repeated detection by whomever they believe is initiating the conversation, thereby still offending but demonstrating a possible example of restrictive deterrence. The fourth hypothesis then is as follows: *Computers with the detection messages shown to attackers will have fewer repeat IP Addresses present then those without the messages.* In addition, whether or not offenders return at all may depend on whether or not they witnessed the message, and so a final hypothesis is that *IP Addresses first appearing on computers without the detection messages will be more likely to reappear in subsequent attacks compared to those seen on computers with the messages on their first attack*. Through these proposed hypotheses, further information on restrictive deterrence and cybercrime will be obtained and analyzed.

3. DESIGN AND METHODS

This thesis makes use of research honeypots set up on a university network in Israel. Honeypots are computer systems made to resemble working computers and are defined as a "security resource whose value lies in being probed, attacked, or compromised" (Spitzner, 2003, p. 40). In this instance, the honeypots (which will henceforth be referred to as target computers) used are for research purposes and are set up to resemble legitimate servers on the university network. As part of the experiment, approximately 250 target computers were developed and deployed on the institute's network where they ran consistently from February to September of 2015. The computers utilized IP Addresses matching those given from the University technology team to appear more legitimate, and ran using the Linnux CentOS 6.0 operating system, a free computing platform available to the public. Access to the target computers was available via a susceptible and highly scanned entryway. The computers were designed to deny entry for trespassers until they had either attempted entry several times or they utilized a frequent password seen in computers for entry. Once they had entered the system, the trespassers were randomly assigned and diverted to either a treatment or control computer, with the treatment being the attempted initiation of a conversation through the presentation of a series of messages. Any trespassing instance was then monitored and recorded in order to gain a range of information on the trespasser and their behaviors when approaching the system.

Trespassers sent to both computers were given free reign on the system. In order to examine the impact that the presentation of detection messages would have on system trespassers, treatment computers would present a series of notifications in an attempt to engage in conversation with system trespassers, beginning with “who are you?”. The remaining questions presented to trespassers were as follows: “Why are you here?” “We are currently tracking your original IP address & geographic location. We are fully aware that the IP address that you use to attack our server - [Attacker IP address here] - does not represent your geographic location.” “We are currently recording every communication to and from our server. This is what we recorded from the last couple of minutes....” “Please disconnect immediately”. All of these messages heightened the implication of potential detection, particularly with the display of their IP Address and their recent activity on the server. While leaving out mentions of punishment, the presentation of these continued messages suggested not only that their presence was detected, but also that there was an increase in the certainty of the trespasser being identified. Control computers had no messages presented to the hackers, but everything else remained the same, thus allowing for information on the effect of the messages to be analyzed. Some computers also contained fabricated data resembling University records as well as information on research and false student identification. The presence of this content was intended to make the computers appear more legitimate to hackers once they had begun exploring the system more thoroughly.

System trespassing events were monitored and recorded as well if the hacker responded with any commands or answered any of the questions presented. Information gathered in the logs collected from the computers includes commands typed, any communication to and from the system server, and information on both successful and failed trespassing attempts. Also collected were videos of the target computers' screens when an attack took place to allow for further analysis of hacker behavior. Several changes were made to some of the operating procedures during the initial phases of the experiment in order to increase the number of trespassers accessing the system. No subject was recruited, rather the target computers were left running for the duration allowing for trespassers unaware of the experiment to access the system. Ultimately, over 60,000 system trespassing events were recorded across 249 target computers used in this study. These were divided closely across both the treatment and control computers, with 30,407 events on the treatment computers and 30,042 on the controls.

3.1 Dependent Variables

For this study, the attackers' IP addresses and the commands that they utilized are of primary interest. In particular, usage commands with a reconnaissance, fetch, and disabling focus are examined. The reconnaissance – focused commands of interest here are used to gather the history of commands that have been entered on the system previously (listed as *history_usage*) and to see the details of the computers network configuration (*ifconfg_usage*). Both of these commands are not inherently damaging, but they would commonly be seen in stage one of a trespassing event and they do allow for trespassers to gather more detailed information on the system that they are accessing which could then be used against it**.** Both of these individual commands, *history_usage* and *ifconfg_usage,* can be classified as reconnaissance-focused, and so for this study they were grouped into one count variable reflecting the occurrences of both commands. This was done by generating a new variable that combined the occurrences of both individual commands to reflect the number of times that they appeared in a session. The new variable reflecting these instances was labeled *recon_usage*.

The fetch-related commands include retrieving a listing the files on the system (*ls_usage)*, and downloading files or information (*wget_usage).* While the file listing command could be considered reconnaissance, it could also arguably be directly connected to fetch commands as it provides a list of files or information that could then be taken by the trespassers. As such, it is grouped with the retrieval command in the fetch-related group. This will also be reflected in a count variable of the two commands combined, labeled as *fetch_usage*. Again, to generate this new command, the instances of the variables *ls_usage* and *wget_usage* were combined together into the new variable to reflect the number of times that they appeared.

Finally, commands that could disable or more directly damage system involve removing files or programs (*rm_usage*) or terminating an application or program (*kill_usage*) were of interest and examined. These are the potentially more damaging commands that could be seen in the final stage of a trespassing event. As with the prior two command groupings, a new variable was generated here as well to reflect the number of times these disabling commands appeared across the sessions. The new variable, *disable_usage*, was generated by combining both the *rm_usage* and *kill_usage* variables in order to identify sessions that had these commands appear (see TABLE 1 for a description of all commands).

The total number of honeypots used in this study was 249 with an average of around 224 sessions on each. Across all sessions, the overall number of commands was quite low with no more then 200 of each observed in over 60,000 events, translating to a rate of under 33 per 10,000 sessions. Each command category was tested with the target computer type individually for the first three hypotheses. The rates of commands per 10,000 sessions were also calculated and tested. The dependent variables for the fourth and fifth hypothesis are the number of duplicated IP Addresses appearing on each target computer. Variables counting the number of duplicates per target computer were constructed by tagging and separating any duplicates observed at the session level. The tagged duplicates revealed that the total number of IP Addresses appearing across all sessions was 1,429 with only 398 appearing once during the entire course of the experiment. Of the 398 unique IP Addresses, 200 appeared on control computers while 198 appeared on treatment ones. Using the identified 398 unique IP addresses, I find that the rate of unique IP Addresses across the 60,449 sessions is 65.84 per 10,000 sessions.

TABLE 1: Command Descriptions

Command Name	Description
ifconfg_usage	Interface configuration – a way to observe and configure the network
history_usage	Obtain the history of commands used on the system
ls_usage	List all files in the system
wget_usage	Retrieve or download files
rm_usage	Remove or delete files from the system
kill_usage	End a process – can be done without logging out
fetch_usage	Count variable of the combined instances of *ls_usage* and *wget_usage*
recon_usage	Count variable of the combined instances of *ifconfg_usage* and *history_usage*
disable_usage	Count variable of the combined instances of *rm_usage* and *kill_usage*

3.2 Independent Variables

The experimental conditions for this study serve as the main independent variables. Namely, whether or not the target computer accessed had the detection message or not is the focus as the proposed thesis is meant to identify any differences between the treatment and control groups. The target computers were coded as binary variables representing the treatment and control groups with those having the conversation messages identified as the treatment and all others as the controls. Additional analyses were also conducted separating the target computers into four categories based on whether or not the computer contained fabricated content or not, and whether or not they contained the treatment condition. In total, 30,042 sessions appeared on control computers while 30,407 appeared on treatment computers. These target computers ran consistently for the seven-month duration of the experiment.

3.3 Methods

In order to test the main hypotheses, independent sample t-tests were used to compare the command categories as well as the rates of the commands, the combined command variables, and IP address duplicates across both groups in the experiment. Chi-square testing was also done to compare the individual commands that made up the different commands category variables across honeypot types. While t-tests and chi-square testing and are sufficient for observing the differences between the two groups, poisson and logistic regression models were also included in the analyses to observe any relationships with additional variables that could have an impact on the outcomes.

4. RESULTS

4.1 Commands Entered by Trespassers

An initial t-test to compare session numbers across the two computer types revealed no significant differences, and so this should not affect the consistency of later results. To examine the first three hypotheses regarding the commands entered by trespassers on the control and treatment groups, three separate t-tests were run with the combined command variables. The first, the reconnaissance command variable (*recon_usage*), showed significant results (p<.000) with more commands entered on the treatment computers then on the controls (See TABLE 2 for t-test results). Dividing the *recon_usage* variable into the two separate reconnaissance commands, *history_usage* and *ifconfg_usage*, and running a chi-square analyses revealed that the history command variable was highly significant, while the network configuration variable was not (See TABLE 3 for chi-square results). Testing the rate of the history and configuration commands per 10,000 sessions, reveals significant differences again for the history variable but not the network configuration.

TABLE 2: T-Test Results

Command and Rates	**Means and Standard Errors**
fetch_usage	Control: μ = .010 Std. Dev. = .112 Treatment: μ = .011 Std. Dev. = .110
*recon_usage****	Control: μ = .004 Std. Dev. = .065 Treatment: μ = .084 Std. Dev. = .277
disable_usage	Control: μ = .011 Std. Dev. = .137 Treatment: μ = .013 Std. Dev. = .148

***p<.001

TABLE 3: Chi-square Results [1]

Command	**Command Means for Treatment**	**Command Means for Control**	**Chi-Square Results**
*history_usage****	μ = .083 Std Dev. = .275	μ = .003 Std. Dev. = .059	Pearson Chi^2=2.3e+03 p=.000
ifconfg_usage	μ = .001 Std. Dev.=.029	μ = .001 Std. Dev.=.024	Pearson Chi^2=1.36 p=.243
ls_usage	μ = .008 Std. Dev.=.087	μ = .008 Std. Dev. = .087	Pearson Chi^2=.010 p=.917
wget_usage	μ = .003 Std. Dev.=.055	μ = .003 Std. Dev.=.052	Pearson Chi^2=.566 p=.452
kill_usage	μ = .004 Std. Dev. = .067	μ = .004 Std. Dev.=.061	Pearson Chi^2=2.05 p=.152
rm_usage	μ = .009 Std. Dev.=.093	μ = .008 Std. Dev.=.088	Pearson Chi^2=1.93 p=.164

***p<.001

For the second hypothesis focusing on fetch-related commands, neither the analysis examining the *fetch_usage* variable, or its individual commands (*ls_usage* and *wget_usage*) were significant. This was also the case when examining the rates of the two individual commands. The disable_usage command was also insignificant as were the individual commands to kill programs or remove software or files, although there were slightly more of these commands viewed on the treatment computers. Examining the rates for these commands also did not yield any significance. The initial analyses conducted here show some support for the alternate first hypotheses, but not for either part of the second or third. However, given that additional factors or variables may have an impact on these results, several additional analyses were conducted.

In order to better examine the command variables of interest in the first three hypothesis, and possible relationships between several other variables that could be influencing the results, poisson and logistic regression models were run following the initial analyses. The poisson model was selected for the combined command variables (*fetch_usage, recon_usage*, and *disable_usage*) as they were count variables. The logistic model was used for the binary command indicators that made up the individual command variables. The target computer classification was used as the primary independent variable for all analyses. Two additional variables were also included in the models. The first was the successful attempt number corresponding to how many tries the trespasser took to access the system. This was included as it could theoretically correspond to the skill of the hacker, with fewer attempts possibly indicating more skill, which could then impact what commands the trespasser was aware of and which ones they would then likely use. IP Address duplication was also included in the model given that repeat trespassers may change their tactics, particularly if they are returning to system and have already seen the messages.

For the poisson regression models, only the reconnaissance combined command showed a significant relationship with the target computer type. For the *fetch_usage* and the *disable_usage* categorical variables, only IP duplicates were significant (See TABLE 4 for poisson results). Logistic models were also run subsequently to determine if any of the individual commands were significantly related to any of the other variables included in the model. Unsurprisingly, only the history command showed a significant relationship with the target computer type (p<.000) again with treatment computers being significantly more likely to have the command entered on them (See all logistic regression results in TABLE 5). Models with the outcomes of network configuration, file retrieval, file listing, file or program removal, and system or program termination all show the IP duplication variable is significantly related to the commands (p<.000). Overall however, results of these models though only show some support for the first hypothesis in relation to the history command while the remaining hypotheses are not supported by the results of these regressions. The results here do however suggest that IP addresses are related to command usage.

TABLE 4: Poisson Regression Results

Command Variable and Coefficients	fetch	recon	disable
Honeypot Type	.069	2.89***	.081
Successful Attempt Number	.099	.961***	.136
IP Duplicate	-.000***	.001***	.000***

***p<.001; **p<.05; *p<.01

TABLE 5: Logistic Model Results

Command Variable and Coefficients	history	ifconfg	ls	wget	kill	rm
Honeypot Type	3.25***	.399	.025	.106	.219	.139
Successful Attempt Number	.452***	.442	-.008	.577***	.011	-.371***
IP Duplicate	-.000***	-.005***	.001***	-.003***	.004***	.001***

***p<.001; **p<.05; *p<.01

4.2 Effect of Computer Content on Commands Used

TABLE 6: Chi-square Results – Computers With and Without content

Command	Control w/ content	Control w/o Content	Treatment w/ content	Treatment w/o content	Chi-Square Results
recon_usage***	$\mu = .004$ Std Dev. =.067	$\mu = .004$ Std Dev. = .061	$\mu = .152$ Std Dev. = .359	$\mu = .006$ Std Dev. = .081	Pearson Chi^2=.000 p=.000
fetch_usage*	$\mu = .010$ Std Dev. = .109	$\mu = .010$ Std Dev. = .114	$\mu = .009$ Std Dev. = .102	$\mu = .013$ Std Dev. = .119	Pearson Chi^2=13.5 p=.035
disable_usage**	$\mu = .011$ Std Dev. = .139	$\mu = .012$ Std Dev. = .134	$\mu = .011$ Std Dev. = .135	$\mu = .016$ Std Dev. = .161	Pearson Chi^2=20.3 p=.002
history_usage***	$\mu = .004$ Std Dev. = .061	$\mu = .003$ Std Dev. = .057	$\mu = .151$ Std Dev. = .358	$\mu = .006$ Std Dev. = .074	Pearson Chi^2=.000 p=.000
ifconfg_usage	$\mu = .001$ Std Dev. = .025	$\mu = .001$ Std Dev. = .024	$\mu = .001$ Std Dev. = .029	$\mu = .001$ Std Dev. = .029	Pearson Chi^2=1.43 p=.699
ls_usage	$\mu = .007$ Std Dev. = .085	$\mu = .008$ Std Dev. = .088	$\mu = .007$ Std Dev. = .082	$\mu = .009$ Std Dev. = .092	Pearson Chi^2=3.39 p=.335
wget_usage	$\mu = .003$ Std Dev. = .053	$\mu = .003$ Std Dev. = .051	$\mu = .002$ Std Dev. = .048	$\mu = .004$ Std Dev. = .062	Pearson Chi^2=7.52 p=.057
kill_usage	$\mu = .004$ Std Dev. = .064	$\mu = .004$ Std Dev. = .057	$\mu = .004$ Std Dev. = .062	$\mu = .005$ Std Dev. = .072	Pearson Chi^2=6.63 p=.085
rm_usage**	$\mu = .007$ Std Dev. = .085	$\mu = .008$ Std Dev. = .091	$\mu = .007$ Std Dev. = .084	$\mu = .011$ Std Dev. = .103	Pearson Chi^2=14.3 p=.003

*p<.05, ** p<.01, ***p<.001

As previously mentioned, a number of computers were also set up with fabricated data to mimic legitimate University records while others were left blank. In order to examine if the presence of this fabricated content had any affect on the commands utilized, the previous analyses were re-run to determine if there were any significant differences between control computers with content, control computers without content, treatment computers with content, and treatment computers without content. Running a chi-square analysis revealed significant results for all three categories (See TABLE 6). The reconnaissance commands appeared to heavily favor the treatment computers with content, while the disabling and data-fetching commands appeared most often on the treatment computers without any content. When breaking these down into individual commands, only the *history_usage* and *rm_usage* commands are significant. Interestingly, neither of the fetch-related commands are significant at the .05 level, although the *wget_usage* command is close.

4.3 IP Addresses Used by Trespassers

To examine the fourth hypothesis, that fewer repeat IP Addresses will appear on treatment computers, a t-test was conducted after tagging and identifying the repeat IP addresses. This analysis showed significant results (p<.000) but in the direction opposite of what was expected. More repeat IP addresses seem to have been recorded on the treatment computers then the control computers with means of 978.05 (Std. Dev. = 1167.29) and 855.66 (Std. Dev. = 1123.51) respectively. This indicates a rather substantial difference between the two groups with a good majority of repeat IP Addresses appearing on treatment computers, contrary to the fourth hypothesis. A poisson regression was also run with IP duplicates as the outcome along with the target computer type and successful attempt number included in the model also indicates a significant relationship between the computer type and IP duplication (p<.000). To test the final hypothesis, a poisson regression was again run, but this time including the specific target computer identification as the primary independent variable rather then its classification. All variables were significant in this model (p<.000) although the coefficient of the target computer name (0.000) indicates a low increase in the rate of occurrence. Overall, both of these analyses do not indicate support for the fourth or the fifth hypothesis.

4.4 Effect of IP Duplication on Commands Entered

Given how significant the IP addresses were in the regression models focused on the first three hypotheses, additional analyses were run to examine how the results may differ when duplicate addresses are removed and accounted for. In order to do this, the t-tests and chi-square analyses for the first three hypotheses were re-run again two different times. The first analysis was focused on sessions where IP addresses appeared for the first time only in the experiment. This resulted in the analyses of 1,429 sessions after all repeat occurrences of IP Addresses were dropped. Among these, only 398 IP addresses did not appear again in the experiment and so a second series of analyses focused on these unique addresses.

TABLE 7: Chi-square Results – First IP Address Appearance

Command	Command Means in Treatment	Command Means in Control	Chi-Square Results
history_usage	$\mu = .012$ Std Dev. = .109	$\mu = .002$ Std Dev. = .043	Pearson Chi^2=6.27 p=.120
ifconfg_usage	$\mu = .009$ Std Dev. = .094	$\mu = .004$ Std Dev. = .060	Pearson Chi^2=1.48 p=.224
ls_usage	$\mu = .021$ Std Dev. = .143	$\mu = .023$ Std Dev. = .149	Pearson Chi^2=.045 p=.832
wget_usage	$\mu = .015$ Std Dev. = .121	$\mu = .012$ Std Dev. = .108	Pearson Chi^2=.191 p=.662
kill_usage	$\mu = .003$ Std Dev. = .055	$\mu = .002$ Std Dev. = .043	Pearson Chi^2=.163 p=.686
rm_usage	$\mu = .015$ Std Dev. = .121	$\mu = .009$ Std Dev. = .095	Pearson Chi^2=.826 p=.363

***p<.001

When re-running the analyses for the first three hypotheses with each of these conditions only the reconnaissance variable was significant (p<.01). No significant results were found for any of the individual commands (See TABLE 7 for results). However, it is important to note that far more of the sessions with the IP addresses appearing for the first time were recorded on control computers (1,094) then on treatment ones (335), and so the results of the analyses could certainly be affected by this disparity. These analyses were also re-run for the unique IP Addresses which only appeared once across all recorded sessions. None of these results were significant (See TABLE 8). Given the low number of unique IP Addresses appearing once though across the sessions (398), several commands were not used at all in these sessions making the analyses irrelevant.

TABLE 8: Chi-square Results – Unique IP Addresses Only

Command	Command Means in Treatment	Command Means in Control	Chi-Square Results
history_usage	$\mu = 0$ Std Dev. = 0	$\mu = .005$ Std Dev. = .071	Pearson Chi^2=.992 p=.319
ifconfg_usage	$\mu = .010$ Std Dev. = .100	$\mu = .005$ Std Dev. = .071	Pearson Chi^2=.346 p=.556
ls_usage	$\mu = .021$ Std Dev. = .143	$\mu = .01$ Std Dev. = .099	Pearson Chi^2=.000 p=.992
wget_usage	$\mu = .005$ Std Dev. = .071	$\mu = .02$ Std Dev. = .140	Pearson Chi^2=1.79 p=.181
kill_usage	$\mu = 0$ Std Dev. = 0	$\mu = .005$ Std Dev. = .071	Pearson Chi^2=.992 p=.319
rm_usage	$\mu = 0$ Std Dev. = 0	$\mu = .005$ Std Dev. = .071	Pearson Chi^2=.992 p=.319

***p<.001

5. DISCUSSION

Given the impact of cybercrime on technology and critical systems, further research has indeed been warranted on the subject. While existing research in criminology does look at cybercrime, and in particular at system trespassing, only a handful of studies have examined a restrictive deterrent impact on the behavior of trespassers (Maimon et al., 2014; Testa et al., 2017; Wilson et al. 2015). For these prior studies utilizing a restrictive deterrence approach, the focus was on the impact of a warning banner that implied both surveillance and the possibility of punishment, and how its presence may affect behavior. No study up until this point had examined how attempting to initiate a conversation with a system trespasser, while also implying discovery, can impact their subsequent actions.

Results of this study indicate some support for only the first hypothesis. Namely, hypothesis 1b, which argued for more reconnaissance commands on computers with the detection messages. It appears that trespassers may be more likely to search the history of the computer system when presented with a series of messages. This could be explained theoretically though as the messages shown may have actually driven more motivated trespassers to explore the system's history to identify anything that may indicate where the messages are coming from. It is possible that by doing this, they may not have as much of a need to observe the network configuration if they have all of the information that they want from observing the history. It does however mean that, in this instance, the trespassers were not deterred by the detection messages as was hoped. Rather they appeared to have been more motivated to explore the system, falling in line with the theory that some trespassers are enticed by challenges to avoid detection or to further their prestige (Jordan & Taylor, 1998; Seebruck, 2015; Taylor, 1999; Turgeman-Goldschmidt, 2008).

Unexpectedly, neither part of the second or third hypotheses was supported. Theoretically, trespassers would enter a system to either retrieve information or to cause some damage (Wilson, 2001; Yar, 2006), but the number of commands entered overall relating to these categories was low. Not finding support for the second and third hypothesis was surprising as it was theorized restrictive deterrence would have worked. If it had not worked, then it was expected that conveying to the trespassers that they were discovered would potentially motivate trespassers to not only attempt to retrieve more information and files from the system, but also to potentially attempt to cripple or damage the target computers to ensure that they are not identified further. Neither of these scenarios were supported. One possible explanation is that the trespassers may not have felt that the messages effectively increased the certainty of punishment and that the mere discovery of their presence was not enough to fully deter them. As such they may not have felt the need to damage the system or kill any programs. It is also possible that they had already explored the system's history and layout and did not find anything of interest, and so they subsequently used no additional commands. Another explanation though is that many of the trespassers on the system could have been more recreationally motivated and were merely practicing and not actually interested in obtaining anything other then practice from the system.

Interestingly, when accounting for whether or not the target computers had fabricated content on their systems or not, more of the command variables were found to be significant than when only accounting for whether or not the computer had the treatment or control condition. The combined reconnaissance variable as well as the individual history command both appeared considerably more often on the treatment computers that had the content then on any other computer type. It is possible that by observing the presence of content, the trespassers were more inclined to enter reconnaissance-related commands in order to obtain more information on the content and its history. For the fetch-related and disabling commands however, most appeared on the treatment computers without any content. While this is somewhat surprising, it is possible that the trespassers who entered in these commands did so prior to observing whether or not the computers had any content as they would have had to do some explorations of the system first in order to observe whether or not any content was present. These results here do not change the outcomes for any of the hypotheses though as we still observe the same significant or insignificant differences between the treatment and control computers when not considering the presence of fabricated content. These results do however suggest possible explanations for why so many reconnaissance commands were used on some of the treatment computers, and they also suggest that future studies may want to examine more closely the impact content on computers can have on trespassers' behavior.

Results for the fourth and fifth hypothesis were significant, but in the direction opposite of what was expected. More repeat IP addresses appeared on computers with the detection messages rather then those without them. This was unexpected as those who had been on the system before would have theoretically thought to alter their IP Address, especially if the first computer they encountered had the detection message. Given the very low rate of unique IP Addresses though, it does appear that the vast majority did not change their IP Address with most trespassers returning at least once. It is possible that so many returned because the messages did not serve as an effective deterrent. The messages, while conveying detection and surveillance, may ultimately not have been taken seriously by trespassers and so the certainty of punishment for them was unaffected. It is also quite likely that many of the IP Addresses were fabricated or spoofed allowing for a further increase in the anonymity that trespassers often enjoy (Jordan & Taylor, 1998). This would allow the trespassers to feel that the certainty of punishment was so low it was insignificant, and so they could return to the system without changing their IP Addresses while feeling confident that they would not be caught. A final possibility is that trespassers had set up their own hacking program meant to repeatedly attack a system without regard to any messages or firewalls presented. Should this be the case, the IP Addresses also would not change and would instead continue to reappear as programmed.

5.1 Limitations

The last point made can and should be listed as a possible limitation. It cannot be determined if the trespasser was a person or a bot designed by a person on the other end of the system. This means that we may not see the same rationalizing process in this scenario from the trespassing entity as in those for other crimes with other offenders. As such, restrictive deterrence would not apply as well, if at all. However, we can also argue that even if a bot was being used in the trespassing events it would still have to have been designed and deployed by a human, and so the designer would still likely apply some measures to the programmed bot to avoid detection along with some instructions for what commands should be used.

A limitation relating to the commands entered by trespassers should also be noted. With a rate of under 33 per 10,000 sessions, very few commands were entered relative to the total number of sessions. This could imply several things. First, if the trespassers were mostly bots programmed to access systems, they simply may not have been programmed to utilize these commands. Instead, they could have had other commands and options written in that were used to explore the target computers. Should this have been the case, the sessions in which commands were entered could have then involved an actual person on the other end, rather then a bot, who were actively making decisions on what commands to enter based on what they were seeing in the computers. The results would then hold more significance and implications regarding trespassers decision-making. Unfortunately, it cannot be determined for sure if this is the case.

Another limitation is that trespassers may have deduced that it is a honeypot computer or that the messages were automated and therefore they may have had no fear regarding detection or punishment. However, considering that the target computers were designed to mimic legitimate university computers even when they did not contain any content, the chances of the trespassers recognizing the computers as honeypots are likely low. Finally, the skill of trespassers is also relevant in this discussion. I was only able to observe the commands and addresses of those who accessed the computers in the give time frame. As such, there are a great deal of trespassers who had no interaction with our computers, and so we cannot know or observe how they would have responded. This does create an issue for the generalizability of the results. However, given that it is the first study of its kind to observe these behaviors, the results are still worth considering and are a good starting point for future research.

5.2 Conclusion and Implications

Future research should continue to examine the commands and IP addresses utilized by trespassers when they are presented with messages indicating they have been identified. To allow for more subjects, it may be beneficial to run target computers for longer then seven months as more trespassers may enter the system if it is available for a longer time frame. Spreading out the messages or changing them to resemble a person rather then a bot could also be useful to deceive more trespassers. Finally, the greatest interest for future studies should be finding ways to both imply detection and strongly convey that punishment is likely as it is still very possible to observe restrictive deterrent effects on system trespassers.

Changes to the messages displayed may be particularly important as the messages presented in this study implied only discovery and not what the punishment would be if the trespassers were successfully identified. Prior studies examining the warning banner effects on trespassers included mentions that trespassers' conduct was illegal and thus they implied that the trespasser could then be punished (Maimon et al., 2014; Testa et al., 2017; Wilson et al. 2015). This mention of punishment could convey a bigger threat than what the messages in the current thesis did, which may explain why the results here were mostly insignificant. Further studies may then benefit more from actively conveying possible punishments and the certainty that they will occur in order to potentially deter more system trespassers.

Until we can determine more information about trespassers and the commands that they enter in different situations, we cannot from this study support the utilization of conversational or detection messages. While most hypotheses were unsupported, those that were indicate more commands entered and more return IP Addresses on computers with the detection messages. As a result, we find no support for detection messages serving as effective deterrents even from a restrictive deterrent perspective. While it is good that fetch-related and program killing or removing commands did not appear more often, the goal would have been to have a reduction of all commands in the treatment computers. One recommendation then would be to try to study more effects that messages have on trespassers in different areas other then those examined in this study. Changing what the messages say and how they are displayed may also be helpful, especially if they can serve as better deterrents by mentioning punishments. In this study the messages all appeared at the very beginning of the trespassing event and were finished in less then three minutes. Having messages displayed more often or conveying more severe punishments may yield different results. More research overall is needed on the effects that messages have to better determine if they can be effective deterrents for system trespassers.

Cybercrime and the people who engage in it are surrounded by layers of secrecy (Holt & Kilger, 2012; Taylor, 1999; Turgeman-Goldschmidt, 2008; Yar, 2006), but what has been seen suggests that system trespassers may be less-likely to be deterred by common measures (D'Arcy, Hovav, & Galleta, 2009; Goodman, 2010; Maimon et al., 2014; Wilson et al., 2015). In particular, the idea that many engage in hacking because of the risks or challenges involved suggests that increasing the severity or certainty of punishment may actually generate more of a challenge for hackers to attempt. As such, this study used a restrictive deterrence approach to observe how trespassers may alter their behavior when it is implied that they have been detected and are being monitored. The results do not indicate that the messages had a dramatic effect on trespassers when compared to those on computers with no messages. This is telling in and of itself as it does confirm that trespassers in cyberspace may not behave the same as offenders in person, particularly when responding to restrictive deterrents. Thus, this study can then pave the way for further research and observations of system trespassers to identify and analyze other ways in which they may respond to possible deterrents.

6. REFERENCES

Beccaria, C. (1963 [1764]). *On crimes and punishments*. New York: Macmillan.

Bentham, J. (1970 [1785]). *An introduction to the principles of morals and legislation*. New York: Oxford University Press.

CSID (2017). *Data Breaches are on the rise and businesses should plan accordingly* [Data file]. Retrieved from https://www.csid.com/resources/stats/data-breaches/

D'Arcy, J., Hovav, A., & Galleta, D. (2009). User awareness of security countermeasures and its impact on information systems misuse: A deterrence approach. *Information Systems Research, 20*, 79-98.

Federal Bureau of Investigation Internet Crime Complaint Center. (2015). *2015 Internet Crime Report*. [Data file]. Retrieved from https://www.ic3.gov/media/annualreports.aspx

Gibbs, J. P. (1975). *Crime, punishment, and deterrence.* New York, NY: Elsevier Scientific Publishing Company.

Goodman, W. (2010). Cyber deterrence: Tougher in theory than in practice? *Strategic Studies Quarterly, 4,* 102-135.

Higgins, G. E., Wilson, A. L., & Fell B. D. (2005). An application of deterrence theory to software piracy. *Journal of Criminal Justice and Popular Culture, 12*(3), 166-184.

Holt, T. J., & Bossler, A. M. (2013). Examining the relationship between routine activities and malware infection indicators. *Journal of Contemporary Criminal Justice, 29*(4), 420-436.

Holt, T. J., Burruss, G. W., & Bossler, A. M. (2010). Social learning and cyber deviance: Examining the importance of a full social learning model in the virtual world. *Journal of Crime and Justice, 33*, 15–30.

Holt, T. J., & Kilger, M. (2012). Examining willingness to attack critical infrastructure online and offline. *Crime & Delinquency, 58*(5), 798-822.

Jacobs, B. A. (1993). Undercover deceptions clues: A case of restrictive deterrence. *Criminology, 31*(2), 281-299.

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

Jacobs, B. A., & Cherbonneau, M. (2014). Auto theft and restrictive deterrence. *Justice Quarterly, 31*(2), 344-267.

Jordan, T. and Taylor, P. (1998). A sociology of hackers. *Sociological Review, 46*(4), 757–780.

Loughran, T. A., Paternoster, R., Piquero, A. R., Pogarsky, G. (2011). On ambiguity in perceptions of risk: Implications for criminal decision making and deterrence. *Criminology, 49*(4), 1029-1061.

Maimon, D., Alper, M., Sobesto, B., and Cukier, M. (2014). Restrictive deterrent effects of a warning banner in an attacked computer system. *Criminology, 52*(1), 33-59.

Maimon, D., Wilson, T., Ren, W., and Berenblum, T. (2015). On the relevance of spatial and temporal dimensions in assessing computer susceptibility to system trespassing incidents. *British Journal of Criminology, 55,* 615-634.

Marcum, C.D., Higgins, G. E., Ricketts, M. L., & Wolfe, S. E.(2014). Hacking in high school: Cybercrime perpetration by juveniles. *Deviant Behavior, 35*, 581-591.

McAffee Labs (2013). *The Economic Impact of Cybercrime and Cyber Espionage* [Data file]. Retrieved from https://www.mcafee.com/us/resources/reports/rp-economic-impact-cybercrime.pdf

Moeller, K., Copes, H., & Hochstetler, A. (2016). Advancing restrictive deterrence: A qualitative meta-synthesis. *Journal of Criminal Justice, 46*, 82-93.

Morris, R. G., & Blackburn, A. G. (2009). Cracking the code: An empirical exploration of social learning theory and computer crime. *Journal of Crime and Justice, 32*(1), 1-34.

Nagin, D. S., & Paternsoter, R. (1991). The preventative effects of the perceived risk of arrest: Testing an expanded conception of deterrence. *Criminology,* 29(4), 561-587.

Norton by Symantec. (2016). *2016 Norton Cyber Security Report* [Data file]. Retrieved from https://www.symantec.com/content/dam/symantec/docs/reports/2016-norton-cyber-security-insights-report.pdf

Peace. A.G.. Galletta, D.F. and.Thong, J.L. (2003). Software Piracy in the workplace; A model and empirical test. *Journal of Mmagemeni Information Systems. 20(1),* 153-177.

Pogarsky, G. (2002). Identifying “deterrable” offenders: Implications for research on deterrence. *Justice Quarterly, 19*(3), 431-452.

Pogarsky, G., & Nagin, D. S. (2001). Integrating celerity, impulsivity, and extralegal sanction threats into a model of general deterrence: Theory and evidence. *Criminology, 39*(4), 865-892.

Rogers, M., Smoak, N. D., & Liu, J., (2006). Self-reported deviant computer behavior: A big-5, moral choice, and manipulative exploitive behavior. *Deviant Behavior, 27*, 245- 268.

Seebruck, R. (2015). A typology of hackers: Classifying cyber malfeasance using a weighted arc circumplex model. *Digital Investigations, 14*, 36-45.

Skinner, W.F., & 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.

Sitren, A. H., & Applegate, B. K. (2012). Testing deterrence theory with offenders: The empirical validity of Stafford and Warr's Model. *Deviant Behavior, 33*, 492-506.

Spitzner, L. (2003). *Honeypots: Tracking hackers*. Boston, MA: Pearson Education, Inc.

Taylor, P. A. (1999). *Hackers: Crime in the digital sublime.* New York, NY: Routledge.

Testa, A., Maimon, D., Sobesto, B., Cukier, M. (2017). Illegal Roaming and File Manipulation on target computers: Assessing the effects of sanction threats on system trespassers' online behaviors. *Forthcoming in Criminology and Public Policy.*

Thomas, K. J., Loughran, T. A., & Piquero, A. R. (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*(1), 10-21.

Turgeman-Goldschmidt, O. (2008). Meanings that hackers assign to their being a hacker. *International Journal of Cyber Criminology, 2,* 382–396.

Wenzel, M. (2004) The social side of sanctions: Personal and social norms as moderators of deterrence, *Law and Human Behavior, 28*(5), 547-567.

Whitman, M. E. (2003). Enemy at the gate: Threats to information security. *Communication of the ACM, 46,* 91–5.

Wilson, T., Maimon, D., Sobesto, B., & Cukier, M. (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*(6), 829-855.

Wilson, Z. (2001). *Hacking: The Basics*. SANS Institute.

Yar, M. (2006). *Cybercrime and society*. London: Sage Publication.

Young, R., & Zhang, L. (2007). Illegal computer hacking: An assessment of factors that encourage and deter the behavior. *Journal of Information Privacy and Security, 3*(4), 33-52.

Young, R., Zhang, L., Prybutok, V. R. (2007). Hacking into the minds of hackers. *Information Systems Management, 24*, 281-287.