

Cryptobytes

CONTENTS

- I. **Identity Theft:
An Introduction**
J. Black

- II. **Why Anti-Phishing
Toolbars Do Not Work**
Min Wu and
Robert C. Miller

- III. **Messin' With Texas –
Deriving Mother's Maiden
Names Using Public Records**
Virgil Griffith
Markus Jakobsson

Nowadays, users increasingly rely on the Internet and mobile networks for many of their commercial and financial activities as well as for other daily life activities (typically involving private information and users' credentials). The computing infrastructure has become a major marketplace and thus, naturally, attracts criminal activities, fraud, thefts and various other threats of similar nature. Cybercrime is now well recognized by consumers, businesses, service providers and even in the popular culture where notions like "phishing attacks" have become commonplace.

The goal of the current issue of *Cryptobytes* is to deal with certain aspects related to threats and countermeasures in modern cyberspace. The modern computing environment is vulnerable. First and foremost, users access directly servers in a very open environment. This is an opportunity for attackers who may exploit the natural weaknesses of the average user via social engineering attacks. These attacks trick the user into accessing wrong servers and performing incorrect actions. In addition, the various tools and underlying mechanisms of the networking infrastructure possess technical weaknesses of their own which may ease the task of potential attackers. The open nature of computing activities over the Internet and the mobile networks, and the fact that user interfaces should be easy to use in order to attract clients, make traditional security countermeasures (e.g., a firewall) irrelevant to current users of cyberspace.

Given the above situation, understanding the security threats, the usability trade-offs, the techniques and methods available as countermeasures, and how to assess their effectiveness, are of high priority. This volume covers a few areas related to the above general subject. In the first paper John Black presents an introduction to "Identity Theft," one of the major concerns in modern society where information about individuals has been made more accessible. This problem has been intensified in recent years due to the global computing infrastructure where information in electronic form can be gathered from numerous sources and can be easily searchable by anyone. The world wide web is the major tool for users accessing cyberspace. The second paper by Min Wu and Robert Miller explains some of the issues with web technology; in particular they assess the effectiveness of security toolbars from a usability perspective. The final paper by Virgil Griffith and Markus Jakobsson demonstrates how personal information, often used for protection of online user accounts, is actually susceptible to automated discovery nowadays.

While the volume is not an attempt to cover the entire area of threats and countermeasures in modern computing environments (e.g., the Internet), the hope is that it is, nevertheless, instrumental. The goal is that it will serve as a source of information regarding some crucial and quite prototypical problems and some basic issues in the area of protecting end users in cyberspace.



Identity Theft: An Introduction

J. BLACK*

October 3, 2005

Abstract

We introduce the concept of “identity theft,” starting from an attempt to establish a definition, then reviewing common methods, both traditional and modern, used by identity thieves. We review typical scams (including phishing) and other types of attacks, then conclude with some modest recommendations.

1 Introduction

WHAT IS IDENTITY THEFT? One can hardly open a newspaper or magazine today without seeing the term “Identity Theft.” A missing purse, a fake ID, the latest virus, a stolen laptop, etc. are often all assigned the sensational label “Identity Theft.” But what, really, is identity theft?

We, as scientists, would like to have a precise definition, but unfortunately this is nearly impossible due to the complex social aspects of this concept. So instead we will develop an intuitive definition and work from there.

“Theft” is clear enough, although even this concept has a complex legal meaning. “Identity” is much harder to get a handle on. What is an “identity?”

We will think of identity as your physical being; your body. Millenia ago, the only means

of identifying a person was to examine an individual’s external appearance and verifying it matched some mental image you had recorded. If the image was sufficiently close, you would accept this person as authentic. If not, he was an imposter.

Other animals might use other means of identification: smell, for example. As we humans expanded our repertoire as language developed: the sound of someone’s voice became a means of identification (just as it is today when we answer the phone, to some degree of accuracy).

As society grew in complexity, we began to find new ways to identify people. And as the stakes grew higher in *correctly* identifying a person (for security purposes, or for financial settings), methods of masquerading and identifying became correspondingly more sophisticated.

In order to extend our ability to identify someone, we can break those attributes we might measure into three categories:

- Physical Attributes
 - Appearance (face, stature, sex)
 - Behavior (personality traits, posture, gait)
 - Biology (voice, DNA, fingerprints, retinae)
 - Augmentation (dental records, implanted RFIDs)
- Assigned Attributes
 - Name, Address

*Department of Computer Science, 430 UCB, Boulder, Colorado 80309-0430 USA.

E-mail: jrblack@cs.colorado.edu WWW:
www.cs.colorado.edu/~jrblack

- Identifying number (SSN, National ID, Driver’s License Number)
- Ad hoc Account Number (Credit cards, Bank Accounts, etc)
- Computer Related (Passwords, Cryptographic Keys, etc)
- Abstract Attributes
 - What you know (Mother’s Maiden Name, Name of first pet, Who won the 1926 World Series)

Virtually *all* of these items have been used at some point to identify most of us, with the exception perhaps of the attributes that are more difficult to extract (eg, DNA or retinal patterns). Many of the “assigned attributes” require some associated documentation (birth certificate, ID card, passport, credit cards, etc), that many of us carry with us much of the time. And, ironically, most of these documents were not intended to act as identification instruments.

Given that today we often wish to identify someone from a distance, and often someone we have never before met, we must rely on the assigned and abstract attributes given above. Since these attributes are more easily transferred to another individual than are the physical attributes, theft suddenly becomes an easier task for criminals.

But still with all of this in mind, it remains difficult to state precisely what “stealing an identity” really means. If you obtain another person’s Social Security Number, have you stolen his identity? If you steal another person’s credit card and use it to purchase items for yourself, does that qualify as “identity theft.” Or perhaps you have stolen a laptop that, somewhere on it, contains the owner’s name, address, and `gmail` password. Perhaps this is “identity theft?”

Each of the above situations—to name a few—have been termed “identity theft” in the mainstream press. The experts in this domain are not cryptographers, nor even computer security folks: they are the specialists in law enforcement

and in the financial sector who have dealt with these problems for decades. Most experts do not consider any of the above incidents as “identity theft.” For example, using someone else’s credit card is called “credit card fraud.”

Instead, “identity theft” is defined (roughly) as the act of one person assuming the identity of another by means of simulating or acquiring as many of the above-mentioned identifying instruments as possible, without the consent of the legitimate owner. One use of a credit card does not qualify: there must be an act where the attacker attempts to “become” another individual in virtually all aspects.

2 Why Steal an Identity?

Theft is an old crime: stealing valuable items has obvious rewards. But why would a criminal want to assume another person’s identity? There are many common motivations: the most common is in order to gain access to that person’s assets (money, physical possessions, etc). Another common reason is to violate physical security (the Trojan War being an older example, Kevin Mitnick’s antics being a newer one). These examples are temporary assumptions of another’s identity, but there are cases where people wish to *permanently* change their identity: someone attempting to illegally immigrate to a new country, a wanted criminal attempting to avoid apprehension. Some governments will (legally) provide new identities in order to protect individuals: in the US this is the “Federal Witness Protection Program.” And of course there is the voluntary transferral of identity (usually temporary) such as when a friend lends a driver’s license (say) to another person, typically as a false proof-of-age. These latter two examples are, of course, not “theft.”

Most of the above changes in identity entailed channels outside of the modern Internet, and indeed they predate its existence. So “identity theft” has been a problem for a very long time,

and yet is just now becoming a common topic in the mainstream press and in research venues. Why?

Partly because the term is being misused, as discussed above, but also because information is now becoming available freely and rapidly due to the emergence of the Internet. A recent project [3] showed that using a variety of freely-available online databases, mother's maiden names could be accurately obtained for a majority of Texas citizens. Although this would probably have been possible before the existence of the Internet, it clearly has been made vastly easier by the network's presence.

Because an individual's entire personal profile (including enough information to effectively assume that person's identity) is commonly stored in corporate databases, computer security has become more and more important as these corporations seek to avoid the lawsuits and bad publicity resulting from break-ins. The old technique for companies wishing to avoid these embarrassments was to not admit them. However, certain jurisdictions (California, USA, being a well-known one) now require that companies report all such break-ins. This type of law is likely to become more widespread in the near future.

The good news is that many government agencies and credit-tracking companies now have divisions exclusively devoted to addressing issues of fraud and identity theft.

3 Identity Theft in a Computer Context

Once again depending on the definition of "identity theft," we can view computer-based identity theft as an old problem. Simple "shoulder surfing" has been around as long as login ids and passwords. Using a "spoofed login screen" to entice computer lab users to enter their login id and password is an idea at least 30 years old. More elaborate is the example of using a fake Auto-

mated Teller Machine to collect bank cards and PINs from unsuspecting bank clients.

Few of these attacks have been addressed by the research community. Most notably, a move away from passwords has been repeatedly advocated (using hardware tokens is a common way to accomplish this), as has the use of new kinds of passwords cast as a challenge-response protocol.

The modern criminal using the Internet now has new and powerful advantages available to him: anonymity, spamming, and clever impersonation attacks. This has resulted in new scams, or more accurately, old scams done in slightly new ways.

NIGERIAN 419'S. Probably one of the most well-known scams is the "Nigerian 419 Scam." This is named after Section 419 of the Nigerian penal code which addresses fraud. This scam, probably known to every reader of this article, involves the perpetrator's promising a commission if you would only help him extract money (usually millions of US dollars) from some problematic institution in his home country. The scammer eventually needs your bank account and other identifying information in order to complete the transfer. This scam has been overwhelmingly successful in Nigeria and other African countries, with estimated profits in the billions of US dollars. It was recently listed as the 2nd largest industry for Nigeria, after oil.

The 419 scam has been so successful largely because of spam: scammers are able to cheaply reach many thousands of targets, and only a few responses can net a substantial profit. If these con-artists had to use postal mail or the telephone, the costs would go way up.¹

PHISHING. Another well-known scam has been termed "phishing." This is a l33t-spelling² of the word "fishing." It is called "fishing" because the

¹I should note that, despite this, I have personally received both a physical letter from Nigeria and a follow-up phone call in just the past week.

²See <http://en.wikipedia.org/wiki/Leet>

ploy involves “casting a fishing line and seeing who bites.”

Phishing is a practice where the attacker creates a fake web page that looks identical to a legitimate one (eg, eBay’s login screen, PayPal’s login screen, etc) and then entices customers to provide their personal information, typically username and password, to this fake web site. In this sense, it is just another version of the spoofed login screens mentioned above. But once again, the Internet comes into play to vastly increase the efficacy of this attack: because spamming is so cheap, tens of thousands of emails can be sent in an attempt to entice users to visit the illegitimate site. And, once again, only a few successes are needed to produce a windfall profit.

In order to succeed in a phishing attack, the attacker must convince the email recipient to visit the illegitimate website. Usually the email claims that an update is required to your account, or more deviously, that fraud is suspected and you must go update your account information. In order to make the URL look legitimate (since obviously the attacker is hosting his fake web page elsewhere), there are a variety of well-known tricks employed, most of which are fairly effective.

Estimated losses from phishing attacks vary quite a bit, and it’s unclear how effective they are. However, if the number of attacks is any indication of prior successes, antiphishing.org reported a 15% average monthly growth in the number of active phishing sites between July 2004 and April 2005. As a point of reference, it found 2,625 active phishing sites during the month of February 2005.

The most common administrative reaction to phishing is to shut down the site: since the web sites are easily tracked down, this involves contacting the ISP and pointing out the illegal site. Ironically, the most expeditious route to convincing the ISP that some site is illegal is to point out that trademarked logos are illegally being used on the site. The antiphishing.org report states that the average time online for a phishing site

from July 2004 to April 2005 was 5.8 days. If phishers begin distributing their sites [4], shutting the sites down will become for more difficult.

Research attempting to combat phishing has boomed recently. Since phishing has a distinct technical flavor to it (as opposed to traditional identity theft which is more “social engineering”), many approaches have been proposed to detect, block, or otherwise prevent phishing. Perhaps the best-known project is SpoofGuard [2], produced as part of the PORTIA project. SpoofGuard is a complex program that looks at a variety of attributes in a mail message in order to assess the likelihood that the email is fraudulent. It looks for typical tricks used by phishers such as the “borrowing” of logos from the legitimate host site, and matching well-known logos to a database. If SpoofGuard believes there is a high probability that the email is illegitimate, it displays a red light icon on the browser’s toolbar.

As with many security products, SpoofGuard is not perfect (for example, an attacker can cut the logo image into several parts and overlay them on the screen so the logo looks correct, but won’t be found in the database). But it detects the majority of fraudulent emails used by phishers. The biggest problem, of course, is that few people use it and those who do are probably those who are least likely to fall for phishing attacks in the first place.

4 CrimeWare

A disconcerting new trend is the number of recent viruses that install “Spyware” or “CrimeWare” on the user’s computer. This software logs keystrokes (capturing sensitive information such as passwords), looks through various files, and sends information back to some central repository for the criminal to collect later. Some of these pieces of malware will attempt to propagate to other computers on the same network, and most will become “zombies:” computers un-

der central control for use in launching Denial-of-Service attacks or engaging in spam-generation.

Finding and removing this type of software can be very challenging. “Cool Web Search” is a well-known piece of spyware that redirects web-based searches to sellers who have paid a fee to the site’s proprietor. The spyware removal tool is known as “CWShredder.” The original CW-Shredder author, after months of trying to keep up with this intrusive piece of spyware, finally gave up.

5 Security is the Bottom Line

For preventing most computer-based intrusions, security is paramount. However, for the lay user, maintaining good security practices is already a daunting task. At a minimum we ask users to run a virus checker, run a firewall, and maintain their operating systems with the latest patches. Additionally, they should keep in mind several rules to abide by at all times: don’t run executables attached to email, don’t open attachments containing macros, don’t check your bank balances from an Internet Cafe. And always check the name in the X.509 certificate before you trust the SSL connection you have just established.

It’s a daunting task just to follow these few recommendations, and most users are (understandably) mystified by it all. Some have resorted to never sending any of their personal information over the Internet. But no one can escape the fact that their information *is* on the Internet somewhere. We only hope that those companies who have it take adequate steps to protect it, despite copious evidence to the contrary.

So what can we, as security researchers, do to help the situation? The current state of network security is embarrassing: ARP is completely insecure, DNS is still insecure, denial-of-service attacks are commonplace and we have no cure. Man-in-the-middle attacks on SSL are very practical due to the fact that no one checks certifi-

cates. There are endless instances of vulnerable pieces of web scripting software, often because Perl, PHP, and ASP are misused. WEP was trivially insecure from the outset.

We need to fix these things. Although the latest results on Trapdoor Pairing-Based Public-Key Encryption Schemes [1] are very cool, one could convincingly argue that fixing any of the above widespread vulnerabilities would be far more relevant for the real world. We should make it a priority.

6 Conclusion

While much fraud is being perpetrated online, and the numbers are increasing, the good news is that other types of fraud still outstrip online fraud, and these attacks are more easily prevented. Use a shredder (to prevent “dumpster diving”), don’t place your outgoing mail in an insecure mailbox, get regular credit checks, don’t readily give out your personal information, etc. And even if you do end up the victim of this type of attack, there is help. Credit Bureaus are now well-equipped to deal with identity theft and fraud.

Perhaps the best advice is to just think suspiciously: we, as security people, already do this most of the time. So we would never, for example, sell your home computer on eBay even after erasing the hard disk. Right?

References

- [1] BELLARE, M. A trapdoor, pairing-based public-key encryption scheme. See <http://www-cse.ucsd.edu/users/mihir/crypto-topic-gener>
- [2] CHOU, N., LEDESMA, R., TERAGUCHI, Y., BONEH, D., AND MITCHELL, J. C. Client-side defense against web-based identity theft. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)* (2004), The Internet Society.

- [3] GRIFFITH, V., AND JAKOBSSON, M. Messin' with Texas: Deriving mother's maiden names using public records. In *Proceedings of the Third International Conference on Applied Cryptography and Network Security* (June 2005), Ioannidis, Keromytis, and Yung, Eds., vol. 3531 of *Lecture Notes in Computer Science*, pp. 91–103.
- [4] JAKOBSSON, M., AND YOUNG, A. Distributed phishing attacks. Cryptology ePrint archive, reference number 2005/091, submitted Mar 25, 2005. See eprint.iacr.org.

Why Anti-Phishing Security Toolbars Do Not Work

Min Wu and Robert C. Miller

MIT Computer Science and Artificial Intelligence Laboratory

{minwu, rcm}@mit.edu

Abstract

Security toolbars in a web browser show security-related information about a website in order to help users detect phishing websites. We conducted a user study of three security toolbars and found them all ineffective at preventing phishing attacks. Even though subjects were asked to pay attention to the toolbar, many failed to look at it; others disregarded or explained away the toolbars' warnings if the content of web pages looked legitimate. We found that many subjects do not understand phishing attacks or realize how sophisticated such attacks can be.

1 Introduction

As people increasingly rely on the Internet for business, personal finance, and investment, Internet fraud becomes a greater threat. Internet fraud takes many forms, from phony items offered for sale on eBay, to scurrilous rumors that manipulate stock prices, to scams that promise great riches if you will help a foreign financial transaction through your own bank account.

One interesting and fast-growing species of Internet fraud is phishing. Phishing attacks use email messages and web sites designed to look as if they come from a known and legitimate organization, in order to deceive users into disclosing personal, financial, or computer account information. The attacker can then use this information for criminal purposes, such as identity theft, larceny, or fraud. Users are tricked into disclosing their information either by providing it through a web form or by downloading and installing hostile software. According to the Anti-Phishing Working Group (APWG), 2870 active phishing sites appeared in

March 2005, a 28% increase per month since July 2004. [3] A survey sponsored by TRUSTe showed that more than seven out of ten respondents have visited a phishing site, over 15% admit to having provided personal data to a phishing site, and estimated that US consumer have lost about \$500 million as a result of phishing attacks. [16]

Examples of phishing attacks are collected and archived by APWG. A typical example is an attack against eBay customers, first reported on March 9, 2004. [2] The attack starts with an email claiming that the recipient's eBay account information is invalid and needs to be updated by visiting the link embedded in the email. This message looks like a legitimate email from S-Harbor@eBay.com, and the link apparently points to cgi1.ebay.com, but actually leads to 210.93.131.250 - a server in Seoul, South Korea with no relationship to eBay at all. Following the link produces a fake web page that looks legitimate, with an eBay logo and page design, and asks for the victim's credit card, Social Security number, eBay username and password. When the user clicks the submit button, the data goes to the hostile server, where it is collected and used by the attackers.

Since most current attacks use broadcast email (spam) to lure victims to a phishing website, one approach is to stop phishing at the email level. Spam filters (e.g., [17]) try to get rid of spam and possibly phishing emails from user's inbox. Yahoo's DomainKeys [8] verifies the delivery path of the incoming email so that the sender cannot be spoofed. Identity-based encryption [1] and key continuity management [13] have been proposed to help users correctly recognize the true email sender. Eudora's ScamWatch [10] analyzes embedded links in emails and warns users when they are about to click a suspicious link.

Another approach for stopping phishing attacks

relies on a security toolbar that displays warnings or security-related information in the web browser's interface. Figure 1 shows some existing security toolbars:



Figure 1: Existing security toolbars

- SpoofStick [7] displays the website's real domain name, in order to expose phishing sites that obscure their domain name. For example, an attack might use a legitimate-looking domain name as a sub-domain, e.g., `www.paypal.com.wws2.us`, which might fool a user looking at the URL, but SpoofStick would display it as `wws2.us`.
- Netcraft Toolbar [18] displays information about the site, including its registration date, hosting country, and popularity among other toolbar users. This information is expected to be helpful in detecting phishing sites because most phishing sites are short-lived compared to the legitimate sites they imitate, and a large number of phishing sites spoof US-based corporations but are in fact registered in other countries.
- Trustbar [14] makes secure web connections (SSL) more visible by displaying prominently the logos of the website and its certificate authority (CA). This is useful against phishing because most legitimate websites use SSL to encrypt the user's sensitive data transmission, but most phishing sites do not. Phishing attackers avoid SSL because on the one hand obtaining an SSL certificate from a well-known

CA, such as VeriSign, requires site identity information that can be traced, and on the other hand using a CA that is not known to the browser will trigger a warning and thus raise the user's suspicion.

- eBay's Account Guard [9] uses a green icon to indicate that the current site belongs to eBay or PayPal, a red icon to indicate a known phishing site found on a blacklist maintained by eBay, and a gray icon for all other sites.
- SpoofGuard [6] calculates a spoof score for the current web page using a set of heuristic rules derived from common characteristics of known phishing attacks. It then translates this score into a traffic light: red for spoof scores above a threshold, indicating that the page is probably hostile; yellow for scores in the middle; and green for low scores, indicating that the page is probably safe.

In addition to these toolbars, browsers already have other indicators that can help users to detect phishing attacks. For example, the address bar displays the URL of the current web page, and the status bar displays a lock icon to indicate whether the page was downloaded with SSL. To further differentiate SSL-downloaded pages, Mozilla Firefox changes the address bar's background color to yellow and adds a lock icon to the address bar. Users are commonly advised by online security tips to pay attention to these indicators whenever they access a web site. [11]

Are these security toolbars and visual indicators actually effective at preventing phishing attacks? There are several potential drawbacks to the security-toolbar approach:

- A toolbar is a small display in the peripheral area of the browser, compared to the large main window that displays the web content. Users may not pay enough attention to the toolbar at the right times to notice an attack.
- A security toolbar shows security-related information, but security is rarely the user's primary goal in web browsing. Users may not care about the toolbar's display even if they do notice it.

- If a toolbar sometimes makes mistakes and identifies legitimate sites as phishing sites, users may learn to distrust the toolbar. Then, when the toolbar correctly identifies a phishing site, the user may not believe it.

This paper describes a user study we performed to evaluate the security toolbar approach to fighting phishing. More generally, we also sought to find out why users get fooled by phishing attacks. More detailed description and discussion of this user study can be found in [22].

2 Study design

Based on the kind of information displayed, we grouped the features of the five existing toolbars into three abstract security toolbars and simulated their displays as shown in figure 2.



Figure 2: Simulated security toolbars

- The Neutral-Information toolbar combines the SpoofStick and Netcraft toolbars, displaying the website’s domain name, registration date, and hosting country.
- The SSL-Verification toolbar imitates Trustbar and displays confirmation information for secure sites with the site’s logo and its CA; a general warning message is displayed for other sites.
- The System-Decision toolbar simulates eBay Account Guard and SpoofGuard, presenting a

judgment about the site’s trustworthiness with a red, yellow, or green light. When the toolbar displays the red light, it also displays “Potential Fraudulent Site” in red as an additional explanation.

2.1 Study implementation

In order to simulate attacks against users, we needed to completely control the display of the toolbars and other security indicators. Users in the study interacted with a simulated Internet Explorer built inside an HTML Application running in full screen mode (figure 3). Different HTML frames displayed different browser components, including the security toolbars. The main frame always connected to the real website, regardless of whether the site was supposed to be phishing or not. To simulate phishing attacks, we changed the appearance of the HTML frames that displayed the browser’s security indicators, including the security toolbar, the address bar and the status bar, to indicate that the web page was served by an unusual source, e.g., tigermail.co.kr rather than paypal.com.



Figure 3: Browser simulation using HTML frames

Our study simulated ideal phishing attacks whose content is a perfect copy of the actual website. This is realistic, since an attacker might not bother mirroring the entire site, but might simply act as a man-in-the-middle between the user and the real site. The attackers would pass the real web pages to the user and the user’s submitted data to the real site and in the meantime capture the user’s sensitive data during the online transaction. Moreover, by simulating a web browser, we can simulate

different phishing attacks without creating phishing sites and we can integrate all security features into one browser since some features (e.g., FireFox address bar feature and TrustBar) have not been implemented within IE yet.

2.2 Study scenario

Phishing is an attack that directly targets the human being in the security system. Simulating these kinds of attacks for study purposes raises some special problems. Chief among them is the secondary goal property first articulated by Whitten and Tygar [21]: in real life, security is rarely a user’s primary goal. The user is primarily concerned with other tasks, such as checking email, buying a book online, or editing a document. Avoiding disclosure of passwords or personal information may be important, but it isn’t foremost in the user’s mind.

In order to produce generalizable results, then, a lab study must be designed to preserve this behavior as much as possible. If we simply asked subjects to “identify the fake web pages,” security would become their primary goal and hence lead them to pay attention and take precautions that they would be unlikely to take in real life.

We addressed this problem by creating a scenario which gave the user tasks to attend to other than security. We set up dummy accounts in the name of “John Smith” at various websites. The subject played the role of John Smith’s personal assistant, and the task was to handle 20 email messages that John had forwarded to them, mostly about managing John’s wish lists at various e-commerce sites. Each email contained a link that that the user had to click to visit the site. The user also received a printout of John’s profile, including his fictitious personal and financial information and a list of his usernames and passwords.

2.3 Simulating phishing attacks

Five of the 20 forwarded emails were attacks, whose links directed the users to simulated phishing websites. Each of the five phishing attacks in the study represents a real phishing attack technique that has been recorded by APWG:

- Similar-name attack: Since one way that users authenticate web sites is by examining the

URL displayed in the address bar, attackers can use a hostname that bears a superficial similarity to the imitated site’s hostname. For example, in this study the phishing hostname `www.bestbuy.com.ww2.us` was used to spoof `bestbuy.com`.

- IP-address attack: Another way to obscure a server’s identity is to display it as an IP address. For example, `http://212.85.153.6/` was used to spoof `bestbuy.com`.
- Hijacked-server attack: Attackers sometimes hijack a server at a legitimate company and then use the hijacked server to host phishing attacks. For example, a hijacked site `www.btinternet.com` was used to spoof `bestbuy.com`.
- Popup-window attack: A popup-window attack displays the real site in the browser but pops up an undecorated window from the phishing site on top to request the user’s personal information. In this study, the phishing site displayed the true `hollywoodvideo.com` site in the browser and popped up a window on top requesting the username and password. Although this pop-up window lacked an address bar and status bar, it nevertheless included the security toolbar.
- PayPal attack: This attack is based on existing phishing attacks against PayPal. The email message warned that John’s account has been misused and needs to be reactivated, and points to a phishing website with hostname `tigermail.co.kr`. Unlike the other attacks, which simulated man-in-the-middle behavior while displaying the real web site, this attack used static web pages saved from PayPal and modified to request not only a PayPal username and password, but also credit card and bank account information.

We consider the PayPal attack different from the other four attacks, which we call wish-list attacks because they merely asked the user to log in and modify a wish-list. The PayPal attack is a current phishing attack, which users in our study may know about from media reports or personal experience; wish-list attacks have not appeared in the

wild. Most current phishing attacks target online banks and financial services, like the PayPal attack. The wish-list attacks target online retailers, which is not yet common but growing. [12] The PayPal attack is greedy, asking for lots of sensitive information; the wish-list attacks can only steal usernames and passwords. The PayPal attack is far more intimidating, urging users to reactivate their account and threatening to suspend their account if they did not do so immediately. Experienced Internet users may be suspicious about such types of emails. Finally, the wish-list attacks simulate man-in-the-middle attacks, in which the user’s personal information is stolen while interacting normally with a real web site through the attacker, while the PayPal attack did not involve the live PayPal site, but some static pages that merely looked like PayPal.

All three toolbars were configured to differentiate the legitimate sites from the phishing sites. None of the phishing sites used SSL so that the SSL-Verification toolbar always displayed a warning on them. On the System-Decision toolbar, all legitimate sites were displayed as trustworthy (green) but all the phishing sites were displayed as phishing (red) or unsure (yellow). On the Neutral-Information toolbar, all phishing sites but the hijacked servers were displayed as a “New Site” and some of them were displayed as they were hosted in other countries outside the US.

2.4 Security toolbar tutorial

Another question in this study protocol design is when and how to give users a tutorial about the security toolbar. We discovered in a pilot study that the presence or absence of a tutorial has a strong effect on performance. Presenting the subjects with a printed tutorial that explained the security toolbar in detail gave them too strong a clue that security was the primary goal in the study. On the other hand, without a tutorial the subjects had no idea what the security toolbar meant, or that its purpose was to prevent phishing attacks. As a result, we introduced the tutorial into the scenario, as one of the email messages handled by the subject. Sent by a system administrator, the email announced that a security toolbar had been installed on the company’s computers to prevent phishing attacks. The message contained a link to the tutorial. In this way, we (the experimenters) did not introduce

the toolbar, but used a third party embedded in the scenario instead. When John Smith forwarded this email to the subject, he explicitly requested that they be careful with his personal information. The toolbars continued to have a “What’s this?” link, and seven users in the study did in fact click on it before viewing the tutorial email.

3 Results and discussion

A total of 30 subjects with previous experience in online shopping, 14 females and 16 males, were recruited by online and poster advertising at a college campus. Twenty subjects were college students from 10 different majors. Each of the three security toolbars was tested on 10 subjects.

3.1 The wish-list attacks

We define the spoof rate as the fraction of simulated attacks that successfully obtain John’s username and password or other sensitive information without raising the subject’s suspicion. Figure 4 shows the spoof rates of wish-list attacks for each toolbar. These spoof rates, 45% for the Neutral-Information toolbar, 38% for the SSL-Verification toolbar, and 33% for the System-Decision toolbar, are all significantly higher than 0%, the ideal. No significant difference was found between the toolbars. But this hardly matters since all the toolbars have high spoof rates. Among the 30 subjects, 20 were spoofed by at least one wish-list attack (7 used the Neutral-Information toolbar, 6 used the SSL-Verification toolbar, and 7 used the System-Decision toolbar). We interviewed these subjects to find out why they did not recognize the attacks by going over these unrecognized phishing sites again:

- Many subjects mentioned in the interview that the web content looked professional or similar to what they had seen before. They were correct in this case because the content was the real web site, but a high-quality phishing attack or man-in-the-middle can look exactly the targeted website as well. Seven of these subjects were observed to use security-related indicators, for example, clicking the security-related links, on the site itself to decide if a site was legitimate or not. These indicators

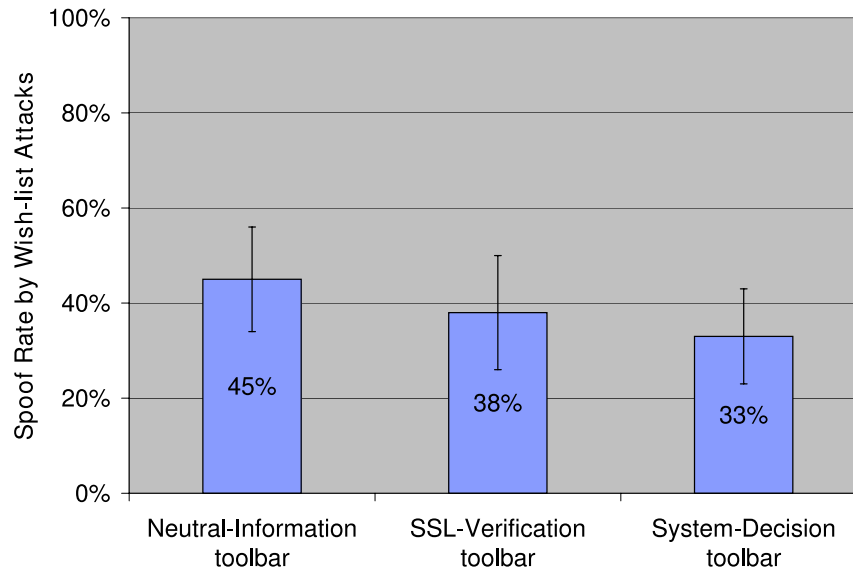


Figure 4: Spoof rates with different toolbars

included a Verisign seal, privacy policy, site contact information and customer service information, a credit card security claim, copyright, and a submit button reading “sign in using our secure server.” Of course, attackers can and do fake these indicators easily.

- Many subjects used rationalizations to justify the attacks. Nine subjects explained away the odd URLs, with comments like:
 - *www.ssl-yahoo.com is a subdirectory of Yahoo!, like mail.yahoo.com.*
 - *sign.travelocity.com.zaga-zaga.us must be an outsourcing site for travelocity.com.*
 - *Sometimes the company [Target] has to register a different name [www.mytargets.com] from its brand.*
 - *What if target.com has already been taken by another company?*
 - *Sometimes I go to a website and the site directs me to another address which is different from the one that I have typed.*

- *I have been to other sites that used IP addresses [instead of domain names].*
- Some subjects explained away the popup window that asked for the username and the password. One subject commented that the popup window was triggered by mistake in a way that she must have clicked “Register for new account” instead of “Sign in for existing account”.
- One subject explained away a toolbar message showing that Yahoo! was a “New Site” and located in Brazil by reasoning that Yahoo must have a branch in Brazil. Another explained away the warning on the System-Decision toolbar by saying that it was triggered because the web content is “informal” just like the spam filter says that “this email is probably a spam.”

- Some subjects said that the reason they were spoofed was that they were focused on finishing the study tasks - i.e., dealing with John Smith’s email requests. Three explicitly mentioned that in order to get the tasks done they

had to take some risks even though they did notice the suspicious signs from the toolbar. Simply warning these subjects that something is wrong is not enough. They need to be provided an alternative safe way to finish the given tasks.

- Some subjects claimed that they did not notice the toolbar display at all for some attacks.
- One subject extensively clicked links on the web pages to test whether the web site worked properly. By relying on the site’s behavior as an indication of its authenticity, this subject was fooled by all of the wish-list attacks.

3.2 The PayPal attack

As discussed above, the PayPal attack is very different from the wish-list attacks. The difference is reflected in the study. The PayPal attack had a significantly lower spoof rate than the wish-list attacks. Several subjects had already seen similar phishing emails in the real world, so they could detect the PayPal attack just by reading the email message, without even clicking through to the phishing site. Other subjects did not feel comfortable providing John Smith’s credit card and bank account information, and eventually noticed the suspicious signs from the toolbar or the suspicious URL from the address bar and thus avoided the attack.

However, there were still some subjects who were tricked by the PayPal attack (at least one using each toolbar). Most of them were PayPal users in real life. They were spoofed because the content of the site looked authentic. One typical comment was “I’ve used PayPal before and this site looks exactly the same. If I trust a site from my experience, I am not suspicious.” They also justified the request as being reasonable. One subject said that “they need this information [the credit card and the bank account information] to charge me.” Thus, familiar phishing attacks can continue to be persuasive and effective, even with security toolbars to warn the user.

3.3 Subjective ratings and comments on toolbars

Subjects were asked at the conclusion of the study to rate the address bar, the status bar, and the security toolbar that they used in terms of their effectiveness at differentiating authentic web sites from phishing web sites, on a scale from -2 (very ineffective) to 2 (very effective). Figure 5 shows the mean ratings.

Among the three toolbars, the SSL-Verification toolbar was rated as less effective, although the difference was not significant. One reason might be because the SSL-Verification toolbar cannot distinguish phishing sites from good sites that do not use SSL. Sadly, many such sites exist in the wild, and many were used in our study. But even when the toolbar functioned properly, it was often ignored. One subject commented that the toolbar looked like an advertisement banner, so it was unclear whether it was put there by the browser or by the site.

The other two toolbars were thought more effective than the browser’s own address bar. A common remark on the security toolbar is that the toolbar combining with the address bar affects the decision about phishing in the way that the toolbar emphasizes and thus makes the subjects to pay more attention to the address bar.

But some subjects did not know how to interpret the information that the toolbars displayed. This was especially true of the Neutral-Information toolbar. One subject said: “How do I have any idea about the [registration] time and location of a site?”

3.4 Why don’t the security toolbars work?

Many users relied on the web content to decide if a site is authentic or phishing. The web content has a large display area and is in the center of the user’s attention. It can make itself very convincing. Most of the time, the web appearance does reflect the site’s identity because of the low phishing rate in the real world. What’s more, in the early days of phishing, phishing attacks frequently had poor grammar and spelling mistakes. In our study, simulated phishing sites had high-fidelity content. As a result, even though the security toolbar and other security indicators in the browser tried to

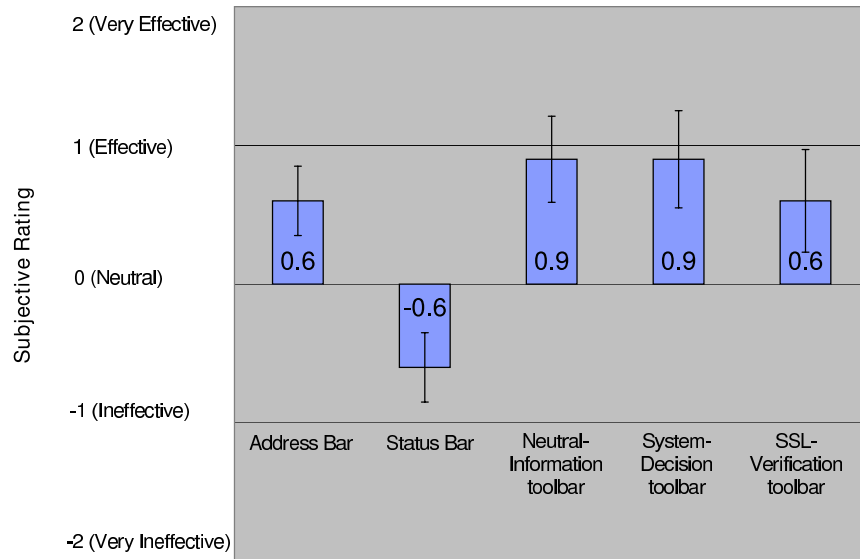


Figure 5: Subjective ratings of the address bar, the status bar and the toolbars

alert the user, many users disregarded the security toolbars because the content looked so good. Poor e-commerce web practices that are common today make phishing attacks even more likely to succeed. For example, many legitimate companies do not use SSL to protect their login page, which was a serious problem for the SSL-Verification toolbar. Many do not have consistent domain names for their web sites. Some use domain names that are vague or unrelated to their brands. Many organizations rely on outsourcing. These practices make it even harder for users to distinguish legitimate websites from malicious attacks.

4 Related work

A growing number of user studies are investigating why phishing attacks are so effective against computer users.

Anti-spam firm MailFrontier Inc did a web survey on how well people can distinguish phishing emails from legitimate ones. [19] Subjects saw screenshots of ten emails but could not interact

with them. About 28% of the time, subjects incorrectly identified the phishing emails as legitimate.

In April 2004, a study in London found that 34% of the respondents would give the researchers their password in exchange for a bar of chocolate. [5] The researchers did not test the passwords to see if they were accurate, however.

In April 2005, a study at Indiana University Bloomington showed that social context makes phishing attacks far more effective. [15] The researchers sent out phishing emails to university students, claiming to be from a friend, having mined friendship relations from a social networking site used on campus. The email led to a phishing site that asked for the subject’s university username and password, and validated them. 72% of the subjects provided valid usernames and passwords to the phishing site.

Whalen and Inkpen used an eye-tracker to study the user’s attention to browser security indicators when doing secure online transactions. [20] Their study found that subjects often looked at the lock icon in the status bar (but rarely clicked on it, and

so didn't learn anything about the site's certificate). By contrast, subjects in our studies rated the status bar least effective at preventing phishing attacks. We think that the difference is due to the fact that Whalen and Inkpen's subjects were explicitly told to pay attention to the security indicators in the browser, while our subjects were asked to detect fake websites, so the address bar and the URL were a more useful indicator.

At least two organizations have initiated phishing attacks against their own members, with the goal of teaching them to protect themselves. [4] West Point found that more than 80% of its cadets succumbed to a phishing attack by a fictional colonel. The State of New York mounted two attacks on its 10,000 employees; 15% were spoofed by the first attack, but only 8% by the second, which came three months later.

5 Conclusion

We have tested three types of security toolbars, as well as the browser's address bar and the status bar, to evaluate their effectiveness at preventing phishing attacks. All the security indicators failed to prevent the users from being spoofed by high-quality phishing attacks.

Users fail to continuously check the browser's security indicators, since maintaining security is not the user's primary goal. Although users sometimes notice suspicious signs coming from the indicators, they either do not know how to interpret the signs or they explain them away. Many users have no idea how sophisticated an attack could be, and do not know good practices for staying safe online.

References

- [1] B. Adida, S. Hohenberger, and R. Rivest. Fighting Phishing Attacks: A Lightweight Trust Architecture for Detecting Spoofed Emails. 2005.
- [2] Anti-Phishing Working Group. *eBay - NOTICE eBay Obligatory Verifying - Invalid User Information*, 2004.
- [3] Anti-Phishing Working Group. *Phishing Activity Trends Report, March 2005*, 2005.
- [4] D. Bank. 'Spear Phishing' Tests Educate People About Online Scams. *The Wall Street Journal*, 2005.
- [5] BBC. *Passwords revealed by sweet deal*, 2004.
- [6] N. Chou, R. Ledesma, Y. Teraguchi, and J. Mitchell. Client-Side Defense Against Web-Based Identity Theft. In *11th Annual Network and Distributed System Security Symposium*, 2004.
- [7] CoreStreet. *SpoofStick*, 2004.
- [8] M. Delany. Domain-based Email Authentication Using Public-Keys Advertised in the DNS (DomainKeys). 2004.
- [9] eBay. *eBay Toolbar and Account Guard*.
- [10] Eudora. *ScamWatch*.
- [11] Federal Bureau of Investigation, Department of Justice. *FBI Says Web 'Spoofing' Scams are a Growing Problem*, 2003.
- [12] S. Fluendy. Phishing targeting online outlets. *Computer Crime Research Center*, 2005.
- [13] S. Garfinkel and R. Miller. Johnny 2: A User Test of Key Continuity Management with S/MIME and Outlook Express. In *Symposium On Usable Privacy and Security*, 2005.
- [14] A. Herzberg and A. Gbara. TrustBar: Protecting (even Naive) Web Users from Spoofing and Phishing Attacks. 2004.
- [15] T. Jagatic, N. Johnson, M. Jakobsson, and F. Menczer. Social Phishing. 2005.
- [16] J. Leyden. US phishing losses hit \$500M. *The Register*, 2004.
- [17] J. Mason. Filtering Spam with SpamAssassin. In *Proceedings of HEANet Annual Conference*, 2002.
- [18] Netcraft. *Netcraft Toolbar*, 2004.
- [19] B. Sullivan. Consumers still falling for phish. *MSNBC*, 2004.
- [20] T. Whalen and K. Inkpen. Gathering Evidence: Use of Visual Security Cues in Web Browsing. In *Graphics Interface*, 2005.

- [21] A. Whitten and J. Tygar. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. In *8th Usenix Security Symposium*, 1999.
- [22] M. Wu, R. Miller, and S. Garfinkel. Do Security Toolbars Actually Prevent Phishing Attacks? In *Submission to Conference on Human Factors in Computing Systems*, 2005.

Messin' with Texas

Deriving Mother's Maiden Names Using Public Records

Virgil Griffith Markus Jakobsson
School of Informatics - Indiana University at Bloomington

Abstract

We have developed techniques to automatically infer mother's maiden names from public records. We demonstrate our techniques using publicly available records from the state of Texas, and reduce the entropy of a mother's maiden name from an average of close to thirteen bits down to below seven bits for more than a quarter of the people targeted, and down to a zero entropy (i.e., certainty of their mother's maiden name) for approximately 3,142,771 Texans. To illustrate the ease of derivation, we are making a demo available for online access. We also describe enhancements to the technique that can be used to *further* reduce the entropy – i.e., increase the accuracy of the derivation – using heuristic rules and public records not used in our experiment. The existence of techniques like the ones we describe pose a significant risk to the individuals whose mother's maiden name can easily be inferred, or guessed from a small number of choices. However, and more importantly, our findings highlight the vulnerability of the system as such, given the traditional reliance of authentication by mother maiden names for financial services. Namely, while our techniques and approach are novel, it is important to note that these techniques – once understood – do not require any insider information or particular skills to replicate and use. Still, we defend the importance of publishing our findings, given the risk for independent discovery of similar techniques by individuals who rather would turn their findings into profit than a publication.

1 Introduction

Within the security community, the secrecy of your mother's maiden name (MMN) is known to not to be the strongest form of authentication. However, the MMN is frequently used as an authentication credential by the commercial sector, including banks, credit cards agencies, internet service providers, and many web services. This may be largely for convenience, but by and large the MMN is considered to be suitably secure against all but the most targeted attacks or those performed by close family friends. However, our study shows that by mining and cross-correlating public records information (which is required by US law to be public), an attacker can determine or “compute” MMNs with startling accuracy.

The ubiquity of birth and marriage information constitutes the most direct threat of MNN compromise through public records. Marriage records are a reliable way of obtaining large numbers of maiden names, while birth records provide the identities of offspring. By using them in conjunction, all that remains for a successful compromise is linking a child to the appropriate parents, and then outputting the bride's maiden name as listed within the marriage record.

The cross-correlation of birth and marriage data is not only effective as a general approach to MMN compromise, but also has numerous non-obvious special cases that make MMN derivation alarmingly easy. For example, if a groom has a very uncommon last name,

then it becomes very easy to match him with any of his children simply by their uncommon last name. Secondly, if the birth record denotes that the child is suffixed “Jr.”, “III”, etc., an attacker can drastically narrow down the number of candidate parents. Third, if the child’s last name is hyphenated, an attacker will seldom have any trouble matching the child with the appropriate marriage. While each of these special cases make up only a relatively small portion of the population, as we increase in scale, even the smallest tricks will result in thousands of compromises.

Our results are important not only to systems that explicitly rely on mother’s maiden names, but also to systems where the MMN is used as back-up authentication. Moreover, the vulnerability we expose poses a threat even to services that do not normally rely on mother’s maiden names. For example, while a given eBay user may not have divulged his MMN to eBay, a sophisticated phisher posing as eBay could include the derived MMN in a spoofed email sent to the user, as a way to gain the trust of the victim. Related context aware attacks are described more in detail in [7].

It is well understood that security is only as strong as the weakest link. The ability to deduce secret information from supposedly innocuous information has been discussed previously, see e.g., [9]. However, we are not aware of any previous instances of deduction of authenticating information on this scale. Although no extensive survey has been made, the use of mother’s maiden name as a security authenticator seems to be a practice unique to Canada and the United States.

As for all studies involving vulnerabilities, the authors were posed with the dilemma of whom to tell, and when. Given the risk that these techniques would be understood and silently used, we believe it is preferable that the vulnerabilities are made public, allowing corporations, government entities and individuals to consider possible countermeasures. We believe that the restriction of publicly available data is not a

meaningful countermeasure; most notably because of the inherent impossibility of undoing the damage associated with already having made the data public. We also recognize that it is important to see the problem in a bigger context; given recently developed methods for deriving banking relationships [8] of victims simply by luring the victim to a rogue site (which can be done as described in [6]), attackers could easily determine where to use social security numbers that he has derived.

The availability and exact information contained within birth and marriage records varies slightly from state to state. So, for purposes of illustration, we decided to focus on only one. Naturally, we wanted as large a sample size as possible to ensure that our methods scaled well to very large datasets, but also to assure that any conclusions pertaining to the sample would be worthy of attention in their own right. This left us with two prominent choices for in-depth analysis: California and Texas. The most recent US Census [4] indicates that Texas is substantially more representative of the entire country than California. Particularly, the ethnic composition of Texas is closer to that of the nation than California. This is of special relevance considering that marriage patterns as well as last names (and therefore maiden names) are strongly influenced by ethnicity. Texas is also more representative in the percentage of foreign-born residents, and the frequency of households moving to other states. Overall, this made Texas a natural choice for our studies. It should be clear that although we chose Texas because of its statistical proximity to the national averages, these same techniques can be used to derive MMNs in other states (especially large states with digitized records) with success rates likely to be on the same order as our findings.

Our success rates are bound to substantially increase if we perform our analysis on a national level, since we only include people *born and married* in the state of Texas, and so, exclude anybody married in a state other than where they were born.

2 Availability of Vital Information

In smaller states, vital information is usually held by the individual counties in which the recorded event took place, and in larger states there is an additional copy provided to a central state office. Texas is no exception to this pattern. Yet, regardless of where the physical records happen to be stored, all such records remain public property and are with few exceptions fully accessible to the public. The only relevance of where the records are stored is that of ease of access. State-wide agencies are more likely have the resources to put the information into searchable digital formats, whereas records from smaller local counties may only be available on microfilm (which they will gladly ship to you for a modest fee). However, as time progresses, public information stored at even the smallest county offices will invariably become digitally available.

The Texas Bureau of Vital Statistics website [17] lists all marriages state-wide from 1966 to 2002; records from before 1966 are available from the individual counties. Texas birth records from 1926 to 1995 are also available online but the fields containing the names of the mother and father (including the MMN) are “aged” for 50 years (meaning they are withheld from the public until 50 years have passed). This means that for anyone born in Texas who is over 50, a parent-child linking has conveniently already been made. It may seem obvious if we think about it, but it’s worth mentioning that the average American lives well beyond the age of 50, making this security measure insufficient. In our analysis we were able to find the unredacted or “aged” birth records for 1923 to 1949.

From the aged birth records alone we are able to fully compromise 1,114,680 males. Married females¹ are more difficult to attack, because they will currently

¹We make the assumption that traditional naming conventions are used, i.e., that all women changed their last name to that of their husband.

have a different last name than the one they were born with. However, the connection can still be made: We matched up females born from 1923 to 1949 with brides married from 1966 to 2002, using first and middle names together with the age of the person. Using this method we were able to compromise 288,751 women (27% of those born in Texas between 1923 and 1949). It is worth noting that MMN compromise from aged records is not only easier, but that older people are also more lucrative targets for fraud as they are likely to have more savings than younger adults.

Here it is worth mentioning that in October 2000, Texas officially removed *online access* to their birth indexes due to concerns of adopted children discovering the identities of their biological parents [2]. Death indexes were similarly taken down as of June 2002 [3]. Furthermore, they increased the aging requirement for both the partially redacted and full birth records to 75 years. However, before the online access was taken down, partial copies of the state and county indexes had already been mirrored elsewhere, where we were able to find and make use of them. We found two sizeable mirrors of the birth and death information. One was from Brewster Kahle’s famous *Wayback Machine* [1], and the other from the user-contributed grassroots genealogy site Rootsweb.com [12], which had an even larger compilation of user-submitted birth indexes from the state and county level. Oddly, despite these new state-level restrictions, county records apparently do not require aging and many county level birth and death records all the way up to the present remain freely available on microfilm or through their websites [14]. Of particular amusement, over three years after being “taken down”, the full death indexes are *still* available (although not directly linked) from the Texas Department of Vital Statistic’s own servers, and at *exactly the same URL they were at before* [20]! All of this is particularly relevant because even though Texas is now doing a better job protecting their public records (although largely for unrelated reasons), the public is just as vulnerable as they were before.

3 Heuristics for MMN Discovery

We have described how a cursory glance at birth and marriage records reveals an ample supply of low-hanging fruit. However, the correlation of marriage data (perhaps the best source of MMNs) with other types of public information comprises an effective and more general approach to linking someone to his or her mother's maiden name. When given a list of random people - whether it be produced by partially redacted birth records, phonebooks, or your favorite social networking service - there are at least seven general observations that an attacker could use to derive someone's MMN with high probability. Naturally, as each heuristic is applied, the chance of MMN compromise will be increased.

1. Children will generally have the same last name as their parents.
2. We do not have to link a child to a particular marriage record, only to a particular maiden name. There will often be cases in which there are repetitions in the list of possible maiden names. This holds particularly true for ethnic groups with characteristic last names. An attacker does not have to pick the correct parents, just the correct MMN. This technique is described in more detail in Section 4.
3. Couples will typically have a child within the first five years of being married. This is useful because given a child's age, it allows a narrowing of the range of years in which to search for the parents' marriage.
4. Children are often born geographically close to where their parents were recently married, i.e., the same or a neighboring county. Again, this is useful because if an attacker knows in which county a child was born (something readily available from birth records), she can restrict the search for marriage records to neighboring counties.
5. Parts of the parents' names are often repeated within a child's first or middle name. Conveniently, this is especially true for the mother's maiden name and the child's middle name.
6. Children are rarely born after their parents have been divorced. In addition to this rule, all Texas divorce records [19] list the number of children under 18 bequeathed within the now dissolved marriage. So, divorce records are helpful not only by eliminating the likelihood of children being born to a couple beyond a divorce date, but they also tell us how many children (if any) we should expect to find, as well as the general birth range to expect for them. In Texas, every divorce affects on average 0.79 children [18]. As nationwide divorce rates average about half that of marriage rates, divorce data can significantly complement any analysis of marriage or birth records.
7. Children cannot be born after the mother's death nor more than a year after the father's death. Texas death indexes are aged 25 years before release (full state-wide indexes for 1964–1975 are available online [20]). Death records are useful in that they not only contain the full name (First/Last/Middle/Suffix) of the deceased, but also the full name of any spouse. This seemingly innocuous piece of information is useful for easily matching up deaths of husbands and wives to their marriages, thus narrowing the list of possible marriages that can still produce offspring by the time of a victim's birth.

For our preliminary statistics, we have taken into account observations 1, 2, 3, and 4. The heuristics listed above certainly are not the only viable attacks an attacker could use, but they serve as a good starting point for the automated derivation of MMNs.

4 Experimental Design

With easy access to public records and no easy way to put the cat back in the bag, we should now be asking ourselves, “How effective are the above described attacks/heuristics in leading to further MMN compromise?”, and “What percent of the population is at risk?” To answer these questions, we will use data entropy to measure the risk of MMN discovery from our attacks. Comparing the entropy of different sets of potential MMNs is a suitable and illustrative measurement for accessing the vulnerability to these attacks. Data entropy measures the amount of unpredictability within a distribution. Its primary benefit over simply listing the number of possible marriage records after filtering is that entropy takes into account repetitions within the set of possible MMNs. For example, after applying all of our derivation rules there could be a set of 40 possible marriages from which the child could have come from. However, 30 of these marriages may have the maiden name “Martinez”, and 5 of the remaining 10 marriages have the maiden name “Lopez.” Clearly, in this case there is a far greater than a 2.5% chance (1/40) of correctly guessing the MMN.

More precisely, after we have applied each of our heuristics, we can measure the chance of MMN compromise as follows. Let x be defined as the list of remaining marriage records, taking only the bride’s last name. Define $|x|$ as the number of items in list x and R_i as the number of occurrences of a given x_i in x . Then, we can define the chance of guessing the correct MMN as:

$$\text{Chance of guessing MMN} = \frac{1}{2^{\text{Entropy}(x)}}$$

$$\text{Entropy}(x) = \sum_i \frac{R_i}{|x|} \log_2 \frac{R_i}{|x|}$$

To provide a baseline comparison for assessing the increased vulnerability due to attacks using public records, we calculated the data entropy across all maiden names in our database (1966–2002). By simply calculating the entropy across all maiden names in our marriage records, we assess that the entropy for randomly guessing the MMN is 12.91 bits.

5 Assessing the Damage

By our methods, we get the following graph (Fig. 1) gauging the risk of MMN compromise from an attacker who makes use of marriage data and makes the assumption that the parents’ marriage took place anytime from 1966 to 2002, but who knows nothing more than the victim’s last name (i.e., has no knowledge of the victim’s age, first or middle name, place of birth, etc.).

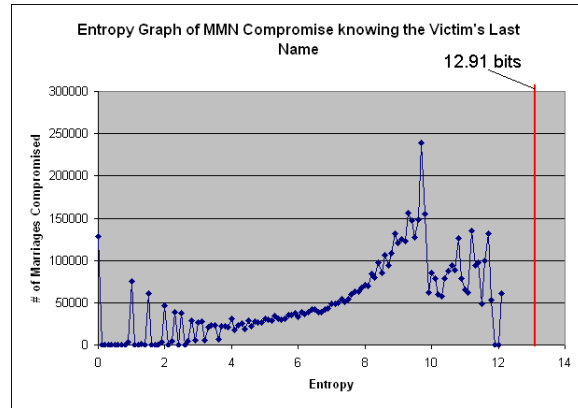


Figure 1: Marriages compromised using marriage records and knowing the victim’s last name compared to not using any records (12.91 bits).

Public records allow the attacker to take advantage of the fact that they know the victim’s last name. Therefore, we will have different entropies, one for each last name. Naturally, deriving someone’s MMNs based solely on their last name will be more difficult for common last names than for uncommon last names given the larger pool of possible parents.

For example, if the attacker *only* knows that the intended victim’s last name is “Smith” (resulting entropy = 12.18 bits), this reduces the entropy only by 0.74 bits from the original 12.91 bits. However, if it is a less common last name like “Evangelista” (resulting entropy = 5.08 bits), or “Aadnesen” (resulting entropy = 0 bits), the attacker is immensely increasing the chances of correctly guessing the MMN. Note that for the absolute worst cases like “Smith” or “Garcia” (9.811 bits), these entropies will still be too high to compromise their bank accounts over the phone.

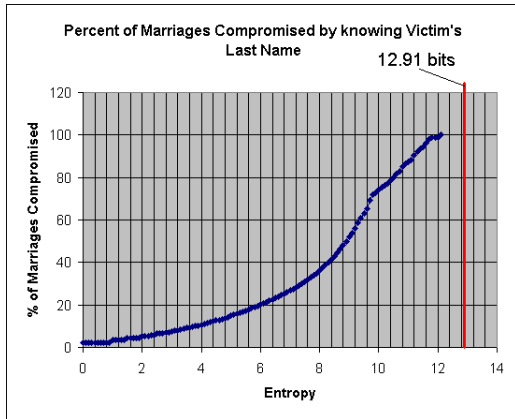


Figure 2: Cumulative percent of marriage records compromised from knowing the victim’s last name.

However, if an attacker has knowledge of the victim beyond his or her last name (such as age, place of birth, etc.), the attacker can eliminate large pools of candidate parents, and therefore improve the chances of determining the MMN. To allow effective comparison of different attacks, we will redraw Fig. 1 as a cumulative percentage of marriage records compromised. We will then take the lowest entropies from the marriage analysis and look for children to compromise within birth records.

A full zero-entropy compromise of approximately 1% of marriages may not initially seem so terrible, but the table above shows that even the smallest percentages will lead to massive compromise.

Entropy	# Children Compromised	% Birth Records Compromised
= 0 bits	82,272	1.04
≤ 1 bit	148,367	1.88
≤ 2 bits	251,568	3.19
≤ 3 bits	397,457	5.04

Table 1: Using birth records from 1966–1995 to search for children with highly unusual last names. For each entropy the percentage of marriage records compromised (the graphs) does not necessarily reflect the percent of birth records compromised (the tables).

5.1 Time and Space Heuristics

Although the first attack is the easiest and most assured route to MMN compromise, in efforts to gain a greater yield there are times when an attacker would be willing to apply further heuristics, such as creating a “window” of time in which it is reasonable to assume the victim’s parents were married. This window of time could be as long or as short as the attacker desires. Naturally, longer windows increase the chances of including the parents’ marriage record, while shorter windows yield higher percentages of compromised MMNs. In this example we assume the attacker knows not only the victim’s last name, but also his or her age (this information can be obtained from birth records or online social networks), and the county in which the victim was born (which can be obtained from birth records). This attack uses a five year window up to and including the year the victim was born when selecting marriage records. Thus, it deduces MMNs in accordance with the observation that couples frequently have children within the first five years of being married. Naming statistics do vary from year to year, but for the reader’s convenience we have averaged all years.

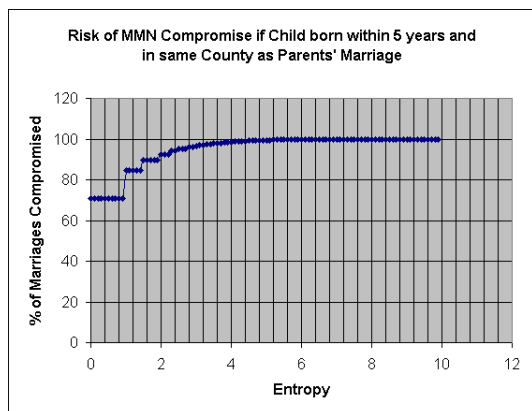


Figure 3: 1970-2002: Marriage records compromised when children have the same last name as parents, the parents' marriage county is known, and the marriage year known within five years.

Entropy	# Children Compromised	% Birth Records Compromised
= 0 bits	809,350	11.6
≤ 1 bit	1,278,059	18.3
≤ 2 bits	1,844,000	26.5
≤ 3 bits	2,459,425	35.3

Table 2: Using birth records from 1966–1995 to look for children – applying heuristics 1, 2, 3, and 4. For each entropy the percentage of marriage records compromised does not necessarily reflect the percent of birth records compromised.

By narrowing our window in which to look for candidate marriages, the resulting entropies drop substantially. An attacker can increase or decrease the window size based upon the uncertainty that the marriage year. As the window increases, there are fewer zero-entropy compromises, but any compromises are more reliable as there is a better chance of the correct marriage record being included within the window.

5.2 MMN Compromise in Suffixed Children

Our final quantitative analysis is for an attack using public records in which the attacker has no knowledge of the victim's age but instead knows the victim's first name, last name, and suffix. Knowing that the victim has a suffix is immensely valuable as it tells the first name to look for within the parents' marriage record.

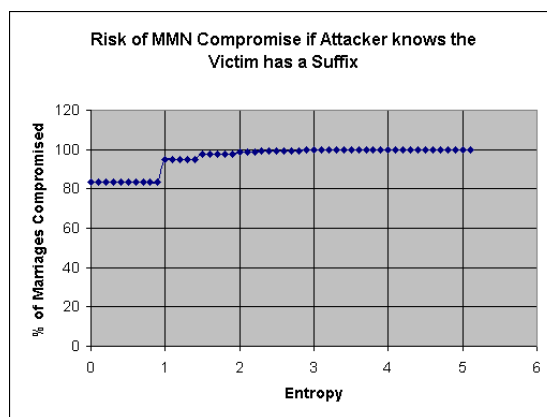


Figure 4: Marriages compromised by knowing the first and last name of the groom.

Entropy	# Children Compromised	% Suffixed Children Compromised
= 0 bits	78,197	13.7
≤ 1 bit	126,153	22.1
≤ 2 bits	178,234	31.3
≤ 3 bits	231,678	40.7

Table 3: Using birth records from 1966–1995 to look for suffixed children. For each entropy the percentage of marriage records compromised does not necessarily reflect the percent of birth records compromised.

5.3 Other Ways to Derive Mother's Maiden Names

Hereto we have focused on the use of birth and marriage records in compromising MMNs. Although birth and marriage information constitute the greatest threat to large-scale MMN discovery, it is by no means the only viable route. The following is a list of more creative public records attacks that we have confirmed to work in sample cases, but which so far remain largely unexplored.

Social Security Death Index. The Social Security Death Index (SSDI) [15] provides up-to-date information of people who have passed away. The SSDI was created as a security measure to prevent the mafia from selling the identities of deceased infants to illegal immigrants. As such, it is comprehensive, digitally available, and fully searchable. In the case of Texas, the SSDI can be used to verify the connection between a groom's death and marriage record. The state death record provides the full name of the deceased person and his or her spouse. (However, there is still always the possibility for name overlap, particularly as you increase in scale.) By taking information from the Texas state death index information and plugging the it into the SSDI, we are able to learn the groom's date of birth, a fact that was unknowable from the state records alone. By knowing the groom's date of birth, an attacker is able to strongly verify the connection to a particular marriage as the marriage record contains the bride and groom's age. This is a reminder of the ability of different records to "interlock" (also called database aggregation) which allows for much stronger conclusions.

Voter Registration Records. In efforts to prevent voter fraud (a concern especially of late) voter registration records are by U.S. law [10] required to be public. But despite the good intentions, next to marriage

and birth information, voter information constitutes the greatest threat to automated MMN discovery and can perhaps fill in the place of either a birth or marriage record. They contain the full name, "previous name" (the maiden name), date of birth, and county of residence [21]. Texas voting records for individual counties are sometimes available from the county websites, but for any significant coverage an attacker would have to purchase them from the state bureau. The database for voter registration records across the entire state costs approximately \$1,100. As of 2000, 69% of voting-age Texans were registered to vote; this percentage has almost certainly increased since then due to efforts to "get-out-the-vote" during the 2004 elections.

Genealogy Websites. Not only a source for mirrored public records data, Rootsweb [13] is an all-purpose user-contributed genealogy website. Amazingly, more often than not, MMNs of currently living people can be read directly from the submitted family trees with no further analysis required for successful MMN compromise. In the off-chance that a security conscious genealogy researcher lists a mother under her husband's last name (as opposed to her maiden name), an attacker can simply look at the last name of the bride's father or one of her brothers. If for some reason this information is not listed, the bride's first name, middle name, marriage date, date and place of birth are always given. With this much information already in hand, a marriage or birth record will allow for certain recovery of the maiden name. Online user-contributed family trees currently do not cover a large fraction of the population, but the submitted trees are still a complete map for MMN compromise and are available to anyone with Internet access. In our analysis we found Rootsweb.com to contain full family trees for 4,499 living Texans. Some genealogy resources such as the Church of Latter-day Saints' FamilySearch.org avoids listing information about living people.

Newspaper Obituaries. Local newspapers frequently publish, both in print and online, obituaries of those who have recently died. Regardless of whether these obituaries happen to be analyzed by hand or via some clever natural language analysis, an obituary entry will generally give an attacker the deceased's name, date of birth, name of spouse, as well as the names of any children. The recently deceased is of no interest to an attacker, but the recent departure of a parent is a convenient opportunity for attacking any children. With the information contained in an obituary, the maiden name can be gotten easily from either the marriage or voting record. However, the children may have moved to other parts of the country, so simply looking them up in the local phonebook may not work. However, an attacker can look up the deceased's SSDI entry which lists a "zipcode of primary benefactor," which will almost invariably be the zipcode of one of the children. The combination of a name and zipcode is a surprisingly unique identifier and the location of the child can be easily queried using Google Phonebook.

Property Records. At our current scale, property records are of relatively little value. However, if we wanted to expand these techniques to a national scale, property records are a good option for tracking people who have moved to another state. Property records are required by law to be public and are usually freely available online [16]. For the purpose of deriving maiden names, property records can be thought of as phonebooks that owners are legally required to be in.

6 Conclusion

Our analysis shows that the MMN is vulnerable to the automated data-mining of public records. New data-mining attacks show that it is increasingly unacceptable to use a documented fact as a security authenticator. Facts about the world are not true secrets. As

a society, there are many ways to respond to this new threat. Texas' response to this threat was by legislating away easy and timely access to its public information. This approach has been largely ineffective, and has accomplished exceedingly little in diminishing the threat of MMN compromise. If these actions have accomplished anything of significance, it is only the creation of a false sense of security. Access to public records of all types was created to strengthen government accountability and reduce the risk of government misconduct by allowing the public to watch over the government that it supports with its tax money. We can only speculate as to the long term effects of policies which would routinely restrict access to valuable public information simply because it might also be valuable to those with less-than-noble intentions.

In today's society, the existence of a separate mother's maiden name, much less a secret one, is becoming obsolete. At one time, the mother's maiden name served as a convenient and reasonably secure piece of information. However, as sociological changes have made it socially permissible for a woman to keep her original name, new technologies have made for comprehensive and accurate record keeping and efficient searching for such information.

Using one of our methods (but expanding our search beyond the state of Texas), we established that the mother's maiden name of the current president of the United States is "Pierce," and the mother's maiden name of his two children is "Welch."

Acknowledgements

The first author wishes to thank Henry Strickland for his suggestions on the entropy graph presentations.

References

- [1] Archive.org 21-Jun-2001: Bureau of Vital Statistics General and Summary Birth Indexes
<http://web.archive.org/web/20000621143352/http://www.tdh.state.tx.us/bvs/registra/birthidx/birthidx.htm>
- [2] Archive.org 20-Nov-2001: Bureau of Vital Statistics, General and Summary Birth Indexes
<http://web.archive.org/web/20001120125700/http://www.tdh.state.tx.us/bvs/registra/birthidx/birthidx.htm>
- [3] Archive.org Birth/Death Index mainpages for 19-Nov-2001 and 05-Jun-2002
Comparing <http://web.archive.org/web/20011119121739/http://www.tdh.state.tx.us/bvs/registra/bdindx.htm> to <http://web.archive.org/web/20020605235939/http://www.tdh.state.tx.us/bvs/registra/bdindx.htm>
- [4] Census 2000 Briefs
www.census.gov/population/www/cen2000/briefs.html
- [5] Florida State Constitution, Section 24.
<http://www.flsenate.gov/Statutes/index.cfm?Mode=Constitution&Submenu=3&Tab=statutes#A01S24>
- [6] T. Jagatic, N. Johnson, M. Jakobsson, F. Menczer, “Phishing Attacks Using Social Networks,” <http://www.indiana.edu/phishing/social-network-experiment/>
- [7] M. Jakobsson, “Modeling and Preventing Phishing Attacks,” Phishing Panel at Financial Cryptography '05. 2005. <http://www.markus-jakobsson.com>
- [8] M. Jakobsson, T. Jagatic, and S. Stamm, “Phishing for Clues: Inferring Context Using Cascading Style Sheets and Browser History,” <http://www-browser-recon.info>
- [9] L. Sweeney and B. Malin, “How (not) to protect genomic data privacy in a distributed network: using trail re-identification to evaluate and design anonymity protection systems,” *Journal of Biomedical Informatics*, pp. 179-192
- [10] National Voter Act of 1993
<http://www.fvap.gov/laws/nvralaw.html>
- [11] Texas State Property Records
<http://www.txcountydata.com>
- [12] Rootsweb.com FTP server with complete copies of both the marriage and death indexes
<ftp://rootsweb.com/pub/usgenweb/tx/>
- [13] RootsWeb.com Home Page
<http://www.rootsweb.com>
- [14] SearchSystems.net listing of Texas Counties' online public record offerings
<http://searchsystems.net/list.php?nid=197>
<http://searchsystems.net/list.php?nid=344>
- [15] Social Security Death Index
<http://ssdi.genealogy.rootsweb.com/>
- [16] Texas State Property Records
<http://www.txcountydata.com>
- [17] Texas Department of Health, Bureau of Vital Statistics, Marriage Indexes
<http://www.tdh.state.tx.us/bvs/registra/marridx/marridx.htm>
- [18] Texas Department of Health, Divorce Trends in Texas, 1970 to 1999
www.tdh.state.tx.us/bvs/reports/divorce/divorce.htm
- [19] Texas Department of Health, Bureau of Vital Statistics, Divorce Indexes
<http://www.tdh.state.tx.us/bvs/registra/dividx/dividx.htm>
- [20] Texas Department of Health, Bureau of Vital Statistics, General and Summary Death Indexes
<http://www.tdh.state.tx.us/bvs/registra/deathidx/deathidx.htm>
- [21] TX Secretary of State Voter Information
<http://www.sos.state.tx.us/elections/voter/index.shtml>