

Countermeasures: Social Networks Alla Genkina and Jean Camp Indiana University, Bloomington chap:socialnetworksecu  
.tifpng.png'convert 1 'dirname 1'/'basename 1 .tif.png

Phishing is a new category of crime enabled by the lack of verifiable identity information or reliable trust indicators on the web. A phishing attack works when it convinces a person to place trust in a criminally untrustworthy party by masquerading as a trusted party. Better indicators about which parties are trustworthy can enable end users to make informed trust decisions and thus decrease the efficacy of phishing. Physical indicators, as embodied in the marble and brass of a bank, cannot be used on the network. Therefore this chapter describes the theoretical underpinning, design, and human subjects testing of a mechanism to use social networks to create reliable trust indicators.

This chapter will first discuss the role and importance of trust online. Then the chapter will present and evaluate existing technical solutions that attempt to secure trust online in terms of their expectations of user trust behaviors. The chapter concludes with a case study of an application that utilizes social networks as a countermeasure to fraudulent activity online will be introduced.

Overview The core observation in this chapter is that phishing requires an incorrect human trust decision. Understanding phishing requires some understanding of human trust behavior. Therefore this chapter begins with an overview of human trust behavior online. Like trust behaviors offline, people online do not complete a thoroughly rational calculation of risk before taking action. People apply systematic heuristics. These heuristics (sometimes called rules of thumb) lead to systematic outcomes, both good and bad.

The understanding of the role of human trust behaviors online provided by Section section:Role can be used to take a critical look at the currently used mechanism for informing trust behaviors. These technical mechanisms sometimes are at odds with what would be predicted by human trust behaviors. The examination of the dominate methods for trust communication illustrate a need to provide better mechanisms to inform trust behaviors.

One way to provide better trust communication is to leverage social networks. Therefore the next section, Section section:Net Trust introduces a mechanism for embedding social networks into online trust behaviors. The goal of the Net Trust system is to inform trust behaviors and undermine the efficacy of all types of fraud. In particular Net Trust uses a reputation system to rate sites. That reputation system will provide positive information only to those sites that have a history of interacting with the social network. Phishing sites are characterized by a lack of history - they appear, attack, and disappear. Net Trust also effectively integrates and communicates third party ratings of websites.

The Role of Trust Online section:Role Phishing is hard to block because it preys directly on the absence of reliable identifying information or trust signals online. Absent any information other than an email from a bank, the user must to decide whether to trust a website that looks very nearly identical to the site he or she has used previously. In fact, any online transaction requires a leap of faith on the part of the consumer because the individual must provide the website with personal and financial information in order to complete the transaction. Thus the consumer must completely trust an effectively anonymous party to be both honorable in the intention of protecting information and technically capable of fulfilling that intention. The consumers' trust decision must be made entirely on the basis of web interaction.

In the physical realm, individuals can use visual, geographical, and tactile information to evaluate the authenticity and trustworthiness of a service provider niss01. In the virtual realm, transactions are characterized by spatial, temporal, and social separation. grabner02, kala97. This separation simplifies masquerade attacks in part by decreasing the cost of constructing a false business facade. While there exists a range of security protocols that are testament to the brilliance of their creators, Internet-based confidence scams continue to increase in profitability and sophistication.

The Federal Trade Commission has reported that in 2004, 53% of all fraud complaints were Internet-related FTC with identity theft at the top of the complaint list with 246,570 complaints, up 15 percent from the previous year reuters. The Pew Internet & American Life Project of the Pew Charitable Trusts has reported that 68% of Internet users surveyed were concerned about criminals obtaining their credit card information, while 84% were worried that their personal information would be compromised pew05.

Banking institutions, American Federal law, and Basel II BIS:05 capital accords distinguish between these types of fraud risks. There are significant legal distinctions between instantiations of unauthorized use of authenticating information to assert identity in the financial namespace. Yet the risks for the subjects of such information is the same regardless of the mechanisms of disclosure. For example, when Choicepoint

exposed information of 145,000 California residents to an identity thief, it was because that thief had created 43 Choicepoint - authorized accounts and purchased the information. The data sharing of Choicepoint was no legal violation, as data brokers have no privacy constraints. The business model, not the security protocols, of Choicepoint is itself the threat to American consumers. The business model of Choicepoint is to obtain all possible information about individuals from purchase records, browsing records, employment records and public records. Choicepoint correlates the information, maintains files on individuals, and resells the repackaged information. In contrast, when Bank of America lost unencrypted back-up tapes containing 1.2M records of personal account information this was a corporate security failure based on a flawed policy. When personnel information was stolen from a laptop at the University of California at Berkeley, it was theft of property and the existence of personal information on the laptop was a violation of university security policy. None of these data exposures were legal violations.

While policymakers and technologists make distinctions between sources of information exposure (e.g., lack of privacy, process failure and security failure) there is no reason for any victim of theft from these incidents to share this perspective. Each information exposure was from different source, but for the subjects of the data the resulting risks were the same. Similarly, end users have been given a proliferation of tools that identify unacceptable privacy policies (Privacy Bird Cranor:02), identify potentially malicious key hierarchies (PKI Perlman:99), or identify malicious sites of a certain category (e.g., Phishguard Keize:04). But there has not been a unified system to provide consumers integrated information to evaluate risks in all its forms. And, as just described, users face not only discrete operational risks but also cumulative risks.

Risk information is embedded in physical spaces by the very nature of the physical world. In the physical realm many trust cues are integrated into the single experience of visiting a store: including the location of a business, appearance of employees, demeanor of employees, permanence of physical instantiation, and thus level of investment in business creation. In the virtual realm the consumer is expected to parse particular threats (e.g., secondary use of data, validity of identity assertion) and weigh each threat independently. If a consumer were to exclude the possibility of trusting the vendor without weighing all the possible outcomes of conducting the transaction, the calculation would be so characterized by uncertainty that the consumer be unable to act niss01. Therefore, when an individual does make the decision to trust an online entity with their personal information, they inevitably assume an uninformed risk. This assumption of uninformed risk is exacerbated by the fact that there is no mechanism that compiles information across various types of risk. The perceived risk of the transaction can often be mitigated by attacks that leverage out of band sources of trust, such as social networks. (These context aware attacks are presented in another part of the text.) In the physical realm, an individual evaluates the perceived risk in a specific context through familiarity of social and physical clues. People infer knowledge about a vendor's values based on their appearance, behavior, or general demeanor, and will proffer trust to those they find significantly similar to themselves niss01. Familiarity "triggers trusting attitudes" because it is perceived as indicating past actions and thus predicting future behaviors gefen00, which in turn lessens the perceived risk. Brick and mortar businesses must invest in physical infrastructure and trusted physical addresses to obtain high degrees of trust. For example, a business on fiftieth block of Fifth Avenue has invested more in its location than a business in the local mall which has in turn invested more than a roadside stall. The increased investment provides an indicator of past success and potential loss in the case of criminal action. Such investment information is not present on the Internet. For example, "tiffany.us" is currently available as a domain name for tens of US dollars. Creating a physical believable version of Tiffany's requires more investment.

In summary, there are four basic difficulties with trust online: cost, identity, signaling and jurisdiction. All these differences between virtual commerce and brick and mortar commerce make trust more difficult.

Cost is the first and foremost difficulty with trust online. The cost of fraud is lower online. While fraud attempts have existed since transactions have, information systems make it easier. Con artists from Nigeria have tried to lure victims with advance-fee fraud by phone for decades, but mass email allows them to reach each and every wired citizen on the globe (several times a day in some cases). The diminished cost is more than a function of technological increase in efficiency. Not only is computer is cheaper than a storefront, but also electronic transactions do not require a social presence. Committing fraud in person demands, at very least, a presentable demeanor and the absence of overt signals of mischief. These require time and effort not demanded by in an online transaction.

Masquerading is easier online because of the lack of reliable identifiers. One reason trust is hard on

the internet is that the online world does not support robust, widespread identity mechanisms. Identity mechanisms have proven to be a very complex problem, and are tied to questions of authentication and accountability, which are important properties in any transaction Camp:2000. While offline identity systems are far from perfect, they are still embedded in a social and institutional context. A hand-signed document, for example, can be corroborated against anything from a personal familiarity with the signer to evidence from a phone record that the signer was not in that location at that time. Moreover, if a single signature were to fail, every other document physically signed would no longer be invalid, as is the case with some digital signature schemes. The physical signatures are still embedded in their situational context.

Showing or signaling that a party is trustworthy is difficult online. (Illustrations that are hard to falsify in economic terms are called signaling.) With low cost for many behaviors, from posting information to providing negative information about a competitor, and an absence of a strong identity mechanism fraud is easy. When proving you are trustworthy is very hard, signaling your attractive quality is quite difficult. Were effective signals possible, there would still be the question of what to signal. Even when a buyer can trust the vendor not to defraud her, the vendor might have a security flaw that creates fraud opportunities for others. Also, the privacy of the purchase might be independent of the security of the buyer's personal information as the vendor might perform admirably in the transaction and later sell customer information. And while the internet has been a boon to new forms of expression, the number of channels through which one can send information is greatly decreased when compared with physical channels. Online entities are limited to online media. A new retail bank might use marble and gild to signal trustworthiness with a significant sunk cost from which they will not retreat. No existing trusted organization has strongly stepped in to lend their extant reputation to signal trustworthiness in others, and newer firms such as TRUSTe and Verisign have limited impact in actually conveying credibility on behalf of their clients as described below.

Finally, the lack of centralization and diversity in legal jurisdictions mean that fraud in one location often cannot be tied to another party or location, and if connected may not be prosecuted. While key actors can leverage advanced cryptography to become what is known as a trusted third party, no single party has inspired the full measure of trust from the market because no third party has enforcement mechanisms equal to rule of law. When transacting with another party online, sometimes only the shipping cost indicates jurisdiction.

Reliable trust information is difficult to create, and such information is even more difficult in a virtual environment. There are four challenges in securing trust on-line: cost, signaling, identity and jurisdiction. Of course, there is a plethora of mechanisms for securing trust online, despite the difficulties. In the next section the mechanisms for online trust are defined and categorized according to their socio-technical assumptions. Then the efficacy of each type to address these four factors is discussed.

Existing Solutions for Securing Trust Online section:Exist In the previous section the difficulties of trust in the virtual environment were discussed. The need for trust online was documented. Of course multiple solutions to the various dimensions of the problem of trust exist. None of these have been completely successful.

In this section we classify the mechanisms based on their fundamental socio-technical approaches for securing trust online: social and trust networks, third-party certifications, and first-party assertions.

Social and trust networks are system where individuals share information to create a common body of knowledge. Third party certifications include all mechanisms where there is a single party that provides verification to all other individuals. These certification systems may or may not include any mechanism for individual feedback on the quality of the certified party. First party assertions are claims that are made by the party seeking to be trusted. Each of the three are described in more detail below, with examples.

Reputation Systems and Social Networks Social networks are a powerful tool that can be used to enrich the online experience. A social network is a map of relationships between individuals, as shown in Figure figure:network. It is defined by the connections between the individuals, which can range from an informal acquaintance to a close, immediate family member. Individuals may, and often do, have several mutually exclusive social networks; such as a network of co-workers that is not connected to the network of close personal friends or those sharing deep religious convictions. The strength of social networks is a function of the trust between the individuals in the network.

Figure figure:network below shows a social network where each circle or node is a person and each line is a connection. The connection may be strong, such as kinship or friendship, or weak. There are assumptions

about social networks that shape the example below. One commonly made assumption is that if two people both know a third person then they are more likely to know each other than if they share only one friendship. This is reflected in the figure below. Another common assumption is that a few people are highly networked (that is, they know many other people) and two highly networked people are very likely to know each other.

On the internet, social networks are implemented through buddy lists, email lists, or even social networking websites. A social network map can be made by tracing all email sent from one person to others, and then tracing those as they are forwarded.

Referrals made through an existing social network, such as friends or family, “are the primary means of disseminating market information when the services are particularly complex and difficult to evaluate ... this implies that if one gets positive word-of-mouth referrals on e-commerce from a person with strong personal ties, the consumer may establish higher levels of initial trust in e-commerce” (Granovetter as cited in Kim00). In 2001, the Pew Internet & American Life Project (PEW) found that 56% of people surveyed said they email a family member or friend for advice about internet vendors PEW02.

Social networks are used in the market today as powerful sources of information for the evaluation of resources. Several commercial websites, such as Overstock.com and Netflix.com, use social networks by enabling users to share opinions, merchandise lists, and other information in order to increase customer satisfaction. In fact, Overstock.com attracted more than 25,000 listings in six months after the implementation of a friends list. Tedeschi04.

Mechanisms that leverage social networks include reputation, rating and referral networks. A reputation system measures individuals who interact on the basis of that interaction. A reputation system may be automated or it may depend on feedback from the people in the system. A rating system has people who rate artifacts. An important distinction between rating people and artifacts is that an artifact never rates the person who rated it in return. A recommender systems (often called collaborative filtering systems) takes ratings given by different people, and uses them to make recommendations on materials to those users.

Consider Amazon to understand the distinction. Buyers can rate books and offer comments. Buyers can write reviews and rate other peoples’ reviews. Buyers develop reputations based on their reviews, as others rate the reviews offered. If one buyer finds all the reviewers of a second buyer and rates these reviews badly (not helpful in the Amazon rating scheme), then the second buyer might identify the first and similarly downgrade his reviews. The ratings for comments develop reputations for buyers who review books. Amazon also tracks every purchase. If these two enemies in the ratings contest have similar buying patterns then Amazon would nonetheless recommend books for one based on the purchases of another. The idea behind the Amazon recommendation system is that people who have been very similar in the past will be very similar in the future.

Examples of social network systems include the FaceBook, Friendster, Orkut and LinkedIn. These systems all leverage social networks for rating, with a link in a social network implying validation of the user. LinkedIn is designed for business opportunities including consulting opportunities and employment. Orkut allows people to self-organize in discussion groups, while Friendster connects individuals interested in romantic relationships. The Face Book uses social networks to connect college students for whatever purpose.

Furl and Del.icio.us are recommender systems that use social browsing. In Del.icio.us personal browsing information is centralized, provided to the users from whatever browser they might be using, and then integrated into a stream of information that is broadcast across a wide range of users. In Furl, individual mirrors of a persons’ web browsing is generated on a centralized site. Both systems use their own recommender systems based on individual’s browsing to recommend sites. The systems allow users to annotate sites with textual labels. Both systems implement search by using the personal browsing history and textual labels provided by everyone. Sites that have been identified as useful by those with similar browsing patterns are implicitly recommended more strongly by being ranked more highly in a search.

Besides specifically designed reputation and social network systems, there are other mechanisms that use the underlying dynamics of reputation. Public forums (perhaps on a vendor’s site as with Dell) and rating systems provide a natural incentive for vendors to act in an honorable manner or otherwise face economic and social consequences. “Social costs can be imposed on exchange partners that behave in opportunistic ways. For example, a firm that gains a reputation as cheater may bear substantial economic opportunity costs, but it may also lose its social legitimacy,” Barney94. Opportunity costs are the literally the costs

of lost opportunity. Choosing not to buy a book means that the book cannot be read. Choosing to buy a CD instead implies that the book will not be bought, thus the opportunity cost of the CD is not reading the book. In this case the opportunity costs are the lost sales that the vendor would have fulfilled in a trustworthy manner, but never had a chance to complete because the customer had no trust. The cost of misbehavior can be greater if the sources of the reputation information are trusted, because the value of lost trust is presumably greater.

While “the Internet was originally expected to be a great medium for disintermediation or the elimination of people as intermediaries to sources and services,” Olson00 in fact it altered the nature of intermediation. Rather than relying on centralized dispassionate authorities, people are necessary to offer “counsel, guidance to the right sources of information, assessment of quality of the sources, and customized advice” Olson00.

Reputation systems attempt to enforce cooperative behavior through ratings or credits. The opportunity for retaliation (through ratings or direct punishment) is an important predictor of behavior in repeated trust interactions. Axelrod94

A good reputation, i.e. credit in a reputation system, enables particular actions or provides privileges. The correlation between the ratings and the privileges may be direct or indirect. The reputation system may have immediate or delayed impact. An example of a reputation system with an immediate impact is the credit rating system. A low credit rating score will result in lost opportunities for credit, as well as increased cost of borrowed funds. The credit rating system attempts to enforce paying off bills and credit cards where legal enforcement is too costly by assisting lenders in evaluating their risk profiles in a particular transaction. Any lender who is not repaid cannot retaliate directly, but can provide information that will reduce the likelihood of any other lender extending credit.

The value of ratings can be quite hard to determine, but there is evidence that even when the ratings do not correspond to well-defined rights they can still be effective in markets. For example, eBay has a reputation system for vendors and buyers. For vendors a lower reputation corresponds to lower prices and fewer bidders. As eBay accounts are simple to obtain, those who have no history can also receive lower prices and fewer bids for their offerings. Resnick:00 Some vendors exclude low-ranked buyers from auctions, but no vendors exclude new buyers. Thus vendor reputations are more useful to vendors than the buyer reputations are useful to buyers. As a result, vendors invest more in manipulation of the reputation system.

In Mojo Nation, the reputation of the user was specified as a virtual currency, “mojo”. This meant that reputation could be explicitly exchanged for services including the download of material, bandwidth and server time. Similarly Kazaa implemented a system that allowed users who upload and annotate files priority in downloading. Bitorrent has a system whereby to download a file, you must assist others in downloading as well. The result in terms of ratings is that reputation is spent as soon as it is accumulated.

The design of a reputation system is not trivial. Flawed reputation systems can inadvertently promote opportunistic behavior. Consider the case of cumulative ratings, as on eBay. On eBay, a vendor who has many transactions (more than twelve) but cheats 25% of the time will have a higher rating than a vendor who has been 100% honest but has only ten honest transactions Dingleline01. Further on eBay vendors manipulate the reputation system by refusing to rate customers until the vendors themselves are rated, thus having the implicit threat or retaliation. Note that the vendors as soon as the buyer’s payment has clear that the buyer has behaved in a trustworthy fashion. For more than 90% of vendor negative ratings there is a corresponding negative buyer rating, indicating that retaliation is a very real threat. Resnick:00 Of course, this manipulation of the reputation system also serves eBay’s interest by creating an artificially low rate of fraud, thus encouraging buyers to participate. essentially the capacity to manipulate the eBay reputation system is a result of eBay’s desire to appear to provide a more trustworthy marketplace than it does in fact provide. A reputation system that was more effective for the buyer in indicating reliability of vendors would both create a more trustworthy marketplace and create a higher reported fraud rate for eBay. Thus a more trustworthy eBay would have a higher apparent fraud rate than the less trustworthy current instantiation.

Reputation systems may be centralized or decentralized. In a centralized system a single authority ensuring the validity of the reputation by tracking the reputations and the identity of the rating party as with eBay. Reputation systems may be distributed, with multiple parties sharing their own reputation information Dingleline01Feldman04.

While Ebay is centralized, Pretty Good Privacy (PGP) uses a distributed reputation system. The reputation system in PGP is used not to offer claims of honesty in transactions, but rather to confirm the

validity of identity claims. PGP allows people to introduce each other and use public key cryptography to vouch for identity claims. In PGP, a social network instead of centralized authentication is used for verification of claims of identity. Garfinkel:94 PGP has weaknesses just as eBay has flaws. Because there is no gatekeeper there are specious claims of identity; for example, there are multiple claimants to the identity Bill Gates and the corresponding cryptographic keys. There is also an issue of affiliation in PGP because no one can stop another person from confirming his or her claim of identity. Public keys are public, and anyone can add their verification of the identity claim corresponding to a public key. There are no major political figures with signed PGP keys as individuals who were clear political liabilities could sign the PGP key and implicitly assert association.

The Google search algorithm is based on an implicit reputation system. Google was the first search company to evaluate the relevance of a web page by examining how many other pages link to a page, as well as the text in the linking anchor. Linking to a page is not an explicit vote, and linking was not previously conceived of as a reputation mechanism. Google integrates the information across many pages, and leverages the combined opinions embedded in those pages. The exact Google mechanism for search is a protected corporate secret. However, there is no question that the implicit reputation mechanism of searching is a critical element of the efficacy of Google search. Google is a centralized evaluation of a distributed reputation. Few individual links have value; however, Google's combination of that reputation information has placed its value in the billions.

Some reputation systems that work in theory do not work in practice. Reputation systems must be simple enough to understand, but not so simple as to be easily subverted. Complexity of the required user interaction can cause a reputation system to fail. Some complex rating systems require that users quantify inherently qualitative information. Multiple proposals for otherwise useful reputation and ratings systems require individuals to place numerical ratings of the level of trust of each rater participating in a reputation system. For example, BBK is a proposal that allows individuals to select and rate individuals providing everything from emails to opinions. BBK functions by constructing a graph of a social network for each user. Each path or link on the social network has a rating between zero and one that indicates if the person is trusted perfectly (1) or not at all(0). Then recommendations are multiplied by the trust factor before being summed. If a recommendation comes from someone far from the evaluator of trust, all the weighing values in the network are multiplied to obtain the final rating value. The individual extending trust is the one who determines the weights in the graph Beth:1994 and thus must be able to meaningfully assign a decimal to each person. The result is a system that is difficult to use particularly for an innumerate computer user.

Of course, the complexity of the Google algorithm has evolved but the requirements on the user have remained simple. The use of public links to calculate standing make it possible to manipulate Google ratings. In response to the manipulation, the Google algorithm has become more complex. There is a continuous evolution of the Google reputation system as those attempting to manipulate search results create ever more complex attacks. Yet the individual users of Google, both those who contribute links and those who search, interact with an extremely simple interface. It is the complexity of the user interaction not the complexity of the reputation system that drives usability.

As has been demonstrated by game theoretic experiments Axelrod94, data provided from the Federal Trade Commission (FTC) FTC and PEW PEW02 social networks and reputation systems encourage but cannot enforce trustworthy behavior. Reputation systems can effectively condense information, share knowledge, and increase accountability. Reputation systems therefore can support individual decision-making about vendors, users, products and websites PEW04. However, there is no theoretically ideal rating or reputation system. The design of a reputation system depends upon the abilities and roles of the players, the length of interaction, the potential loss if the system is subverted, and the distribution of that loss. Even the best reputation system design cannot work without careful design of the user interaction.

Recall that the four difficulties of online trust are cost, identity, signaling and jurisdiction. How well do reputation systems and social networks solve these fundamental problems?

The cost of designing and implementing a reputation or ratings system can be quite high. A well designed reputation system can provide effective signaling. However, reputation systems do not solve the problem of identity. Reputation systems that are hard to join prevent most from participating. Reputation systems that are easy to join suffer the problem that cheap new identities are not trusted. However, a reputation system can solve the problem of signaling by differentiating good and bad resources.

Reputation mechanisms can provide censure within their domains, but that is not the equivalent of having jurisdiction as a bad party can choose to cease participating in a reputation system or create a new identity to lose bad ratings. Recall that there is a problem of preventing bad parties from escaping punishment by obtaining new identities while allowing good parties to join without being penalized is problematic.

#### Third Party Certifications

Third party certification systems vary widely in technical validity, from those that are cryptographically secure to those based on easily copied graphics. Third party certification is provided by a range of vendors, from traditional brick and mortar organizations (e.g., Better Business Bureau) to Internet-only organizations (e.g., TrustE, Verisign). In any case, the third party organization is paid by the Internet entity that is trying to assure the customer of their inherently trustworthy nature.

The most popular method of securing trust online is through the display of small images called trust seals. Trust seals are meant to facilitate the transfer of trust from an individual to a vendor through a third-party intermediary. Unfortunately, the seals themselves are digital images that can be easily acquired and displayed by malicious websites.

Trust seals also do not inherently expire. Once an online vendor has procured a seal, there is limited auditing to ensure that the vendor continues to comply with all the related policies. Compared with ratings or reputations systems, trust seals are not continually evaluated and updated. Trust seals are limited in their efficacy because any web server can simply include the graphic and thereby claim the seal. A seal provider in one jurisdiction will have a limited ability to act on a vendor in a different jurisdiction. A trust seal by a third party will not make a site “automatically trustworthy in the consumer’s eyes. The third-party itself must be credible”. barney94

The seals are can be inherently meaningless. For example, the TrustE basic seal confirms only that the vendor complies with the vendor-generated privacy policy. Not only is there no confirmation of the quality of security of the vendors site, but it is also the case that the privacy policy may be exploitive or consist of an assertion of complete rights over customer data.

figure[h!] center [TrustE seal] [scale=0.8]SocialNetworkSecurity/TrustEa subfig:TrustEa [Child-friendly TrustE seal] [scale=0.8]SocialNetworkSecurity/TrustEb subfig:TrustEb [EU TrustE seal] [scale=0.8]SocialNetworkSecurity/TrustEc subfig:TrustEc fig:TrustE