Is Everything We Know About Password-Stealing Wrong?

Dinei Florêncio and Cormac Herley
Microsoft Research
One Microsoft Way
Redmond, WA, USA
{dinei,cormac}@microsoft.com

ABSTRACT

Federal Reserve Regulation E guarantees that US consumers are made whole when their bank passwords are stolen. The implications lead us to several interesting conclusions. First, emptying accounts is extremely hard: transferring money in a way that is irreversible can generally only be done in a way that cannot later be repudiated. Since password-enabled transfers can always be repudiated this explains the importance of mules, who accept bad transfers and initiate good ones. This suggests that it is the mule accounts rather than those of victims that are pillaged. We argue that passwords are not the bottle-neck, and are but one, and by no means the most important, ingredient in the cyber-crime value chain. We show that, in spite of appearances, password-stealing is a bad business proposition.

Consumers are not liable for emptied accounts

> *"It's not what you don't know that kills you, it's what you know for sure that ain't true."* - Mark Twain

It is worth, at the outset, dispelling a widely-held misapprehension about password-stealing. Thieves certainly steal passwords, and money is certainly a large part of their motivation, but when they successfully extract money from financial accounts individual consumers do not pay. In the US, Regulation E of the Federal Reserve [1] limits consumer liability, in the event of fraud, to $50 (this is separate from the $50 limit for credit-card fraud, Regulation CC) and covers "any electronic transfer that is initiated through an electronic terminal, telephone, computer or magnetic tape." In the US banks, brokerages, and credit unions are governed by this regulation and most go beyond it and offer a zero liability policy to consumers. Bank of America, for example, "guarantees zero liability for any unauthorized activity originating from Online Banking or Bill Pay." Wells Fargo says "We guarantee that you will be covered for 100 percent of funds removed from your Wells Fargo accounts in the unlikely event that someone you haven't authorized removes those funds through our Online Services." Fidelity "will reimburse your Fidelity account for any losses due to unauthorized activity" and "under HSBC's $0 Liability, Online Guarantee, you're covered 100% and liable for $0." Even non-traditional financial institutions offer this guarantee. For example in its Dec. 2009 10-K filing eBay states: "PayPal currently voluntarily reimburses consumers for all financial losses from transactions not authorized by the consumer, not just losses above $50."

Thus, in the US, *individual consumers* are largely insulated from the *direct financial consequences* of credential theft (losses of small businesses and indirect losses are briefly mentioned below).[1] Consumers who have their accounts emptied through stolen credentials are made whole. Of course, the cost of the fraud doesn't just go away: covering fraud is a cost which gets passed back to consumers in the form of increased fees. However, the idea that consumers are "just a few clicks away" from having their accounts irretrievably emptied is simply incorrect. There is a world of difference between being personally liable for losses, and sharing losses that are diluted across the whole population. While "we all pay for cyber-crime" is true in a general sense, it is not the case that individual users face grave financial risk.

We begin with this misconception because it is widely-held and generates enormous confusion. It also has far-reaching consequences for who loses money, how much is lost and where the bottle-necks lie in the password-stealing pipeline. This is the subject of the remainder of this paper. While Regulation E is not secret and occasional references to it appear (*e.g.*, "Zero Liability Policy Protects Bank Customers," NY Times, Nov. 28, 2009 and Krebs' blog [5, (Oct. 7, 2010)]) the implications are seldom pursued in the academic security literature.

For fear of misunderstanding, it is worth explicitly stating that we limit the scope of our remarks to financially-

[1] While consumer protections in the US are good, they are by no means unique. EU Directive 2007/64/EC of the European Parliament limits consumer liability to 150€ and many banks go beyond this. Mannan and van Oorschot [13] find that most major Canadian banks offer a "100% reimbursement guarantee for online banking fraud losses" (but also suggest that most consumers are unlikely to meet the standard of care required to be eligible).

motivated password-stealing attacks against the bank accounts of US consumers. We do not examine password-theft from email or social-networking sites. We don't explore other aspects of cyber-crime. We touch only briefly on non-consumer losses below. We make no mention of corporate accounts or other vital assets protected by passwords. While we make no claim that our conclusions generalize beyond banking losses to US consumers this case is large enough to be interesting and instructive. Online banking is done almost exclusively with passwords in the US. Thus, using what might be considered the lower bound in terms of security, US banks offer the upper bound in protection: zero liability. Anderson [3] observes that while dumping liability on consumers is far more common in the UK this has not resulted in less fraud or savings in the amount spent on security.

Emptying accounts is hard

That US banks offer zero liability allows us to infer much about losses and the fraud protection mechanisms already in place. We now argue that Regulation E implies that emptying accounts is far from simple.

Once he has stolen passwords a thief clearly needs a way to transfer money from victim accounts. This needs to be irreversible and ideally should also be untraceable. There's little point in doing the transfer if it can be rolled back, and there's high risk if it leaves a trail that leads to the thief's door. We now argue that this is hard. Suppose not, *i.e.*, suppose that doing irreversible untraceable money transfers from a bank account (armed only with a password) is easy. If so this opens an enormous vulnerability to self-theft. Any Internet banking user can transfer her money to another account (that she controls) and then claim fraud and demand reimbursement. Repudiation is easy if the transfer is untraceable. The money can't be recovered if the transfer is irreversible. Regulation E compels the bank to make the (dishonest) customer whole.

The answer, of course, is that getting away with this scam isn't as easy as it looks. Banks can't allow easy repudiation of irreversible transfers for which they have offered zero liability guarantees. They must be able to distinguish fraudulent transfers initiated by a thief and repudiated transfers initiated by the account holder. Otherwise every single one of their customers has a wide open opportunity for fraud without even needing to steal a password. It's sufficient that this determination can be made after the fact if the account holder tries to repudiate the transfer. For example, repudiation of an ATM cash withdrawal requires that one is not captured by the ATM camera and has a plausible argument as to how the thief acquired both the card and the PIN. Repudiation of an online transfer requires that the receiving account cannot be linked in any way to the customer. Ideally, this is done with a "John Doe" stepping-stone bank account that can be used to relay money to cash. This is hard since anti-money laundering provisions of the Bank Secrecy Act (1970) and the Title III of the USA Patriot Act (2001) make it difficult to set up a bank account under an assumed name. To comply with Customer Identification Provisions of these laws US banks require a government-issued ID, a Social Security Number and an in-person appearance at a (generally camera-monitored) bank branch to open an account. The ChexSystems database, a catalog of the disputed transactions, and checks that consumers have bounced is consulted by most banks before opening a new account. Documents can, of course, be forged but this raises the effort, expense and risk. It also limits the throughput: accounts that receive fraudulent transfers are quickly frozen and can't be used further. Thus, at scale, this approach requires not one or two "John Doe" accounts but many. If he can clear five transfers through an account before it is frozen then the thief requires one fifth as many accounts as he has stolen credentials.

Krebs has documented a number of thefts from small businesses, and these accounts make clear that banks have considerable success in reversing transfers. For example, of $2 million stolen from Global Title Services $1.8 million was reversed [5, (Nov. 11, 2011)]; of $217,000 stolen from MECA $147,000 was reversed [5, (Aug. 11, 2011)]; of $1.9 million stolen from Experi-Metal Inc. $1.34 million was reversed [5, (June 17, 2011)]; of $110,000 stolen from United Way of Massachusetts Bay all was blocked or reversed [5, (Feb. 3, 2010, 2010)] and of $801,495 stolen from Hillary Machinery $600,000 was reversed [5, (Jan. 26, 2010)]. Since small businesses typically have a greater number and diversity of transactions than consumers do, fraud is almost certainly far harder to detect there. That emptying accounts is hard is further corroborated by the observation that stolen credentials are offered for sale on underground markets at fractions of a penny on the dollar [14, 12].

This analysis assumes that fraudulent transactions are noticed and reported. For large transfers this seems safe: we assume that few consumers would fail to notice a $10,000 transaction. But what of smaller amounts? Might it not be easier to transfer small amounts regularly in the hope that they are never detected? We argue that this is unlikely. Suppose that a thief does many $10 transfers. If p is the probability of any individual transfer being detected, then to have a 50% chance of extracting $10,000 the thief needs $(1-p)^{1000} > 0.5$ which gives $p < 0.0007$, meaning that the chance of each transfer being detected should be lower than 0.07%. Attempting large transfers seems the better approach.

The different procedures and legal status involved also suggests that banks understand non-repudiation and are keenly aware of the different risks posed by reversible and irreversible transactions. If banks are to protect themselves transactions that are irreversible must be made hard to repudiate. Western Union transfers and cashier checks are inherently irreversible but require in-person appearance, presentation of ID and a signature (and are not covered by Regulation E). ATM withdrawals are covered, but since they require a two-factor authentication (*i.e.*, possession of card and knowledge of PIN) at a camera-protected machine they are hard to repudiate. Getting cash in exchange for a check is extremely hard unless the recipient has an account at the bank (*i.e.*, the bank knows who to debit if things go wrong). Any transfer that requires only a password to initiate is easy to repudiate. If it is covered by Regulation E it must also be reversible. We assume that banks don't simply transfer funds covered by Regulation E and hope for the best. US banks handled about 100 million checks and 75 million automated transfers per day in 2009 [6]. They are very familiar with fraud, money-laundering, check-washing, counterfeiting, the possibility of insufficient funds, and counterparty risk in all its forms. They know that, once they hand over cash, any subsequent problem becomes their problem. They limit their risk by offering zero-liability only where an easily repudiated transaction can be reversed.

Mules, not victims lose money

It is very difficult to get a bank to transfer money irreversibly in a way that can later be repudiated. And password-enabled transfers can always be repudiated. Thus password-thieves have a problem: as things stand the fruits of their labor are worthless. They can see the money, they just can't get it out in a way that ensures that it stays out. Thieves respond by taking to heart a wise saying attributed to Butler Lampson: "every problem in computer science can be solved by adding another layer of indirection." Rather than incur the effort, expense and risk of opening "John Doe" stepping-stone accounts, the thief simply enlists others who already have accounts. The solution is to use a human proxy, *i.e.*, convince someone (with rather less experience of fraud than a bank) to act as a relay. It is largely for this reason that draining accounts is usually done through money mules. A money mule is sent stolen money from compromised accounts and forwards, minus "commission", to the thief. The money mule's role is to turn a traceable reversible transaction into an untraceable irreversible one [7]. Using a stolen password, the thief transfers money (traceably and reversibly) to the mule's account, using, *e.g.*, online billpay. Upon receipt, the mule sends this money (untraceably and irreversibly), minus "commission", to the thief. By using, *e.g.*, Western Union for this transfer the mule has made it irreversible and untraceable. By authorizing the withdrawal with a signature, he gives up any ability to repudiate. He has thus given up any consumer legal protections that he might have enjoyed. He accepts a bad transfer and initiates a good one.

Consider a fraudulent transfer of $9000 from a compromised account. The thief sends $9000 using online billpay to the mule. The mule sends $8100 to the thief and keeps $900 commission for himself. Once fraud is discovered the victim is reimbursed, and reversal is attempted from the mule account. Thus before discovery the victim, mule and thief have -$9000, $900 and $8100 respectively. After discovery and reimbursement they have $0, -$8100 and $8100 respectively.

Notice that the thief is up precisely the amount that the mule is down (or in debt). Thus, the thief is really stealing from the mule, not the compromised account, though that fact does not become clear until the dust settles. Thus money mules are not merely unwitting accomplices, they are the true victims in credential theft fraud. Their accounts are not simply vital stepping stones in the evacuation of funds, their accounts (not the victims) are the ones to be pillaged. If the transaction cannot be reversed (*e.g.*, the mule has insufficient funds) then the bank (either the victim's or the mule's) is left with uncollectible debt.

If the thief really steals from the mule, what need has he of the original victim account? Recall that victim-to-mule transfer was necessary to create the illusion of a legitimate task for the mule, and the temporary availability of funds for the critical mule-to-thief transfer. Mules are recruited with semi-plausible stories of work-at-home schemes. Often the mule is led to believe that this is a real job acting as "clearing agent," or "account manager" for a foreign firm or a "secret shopper." Transfers just below $10k (above which a Currency Transaction Report (CTR) must be filed under the Banking Secrecy Act) are the most popular amount, and Western Union and Moneygram are popular channels of payment [5, (May 11, 2010)]. That all transactions must be handled with urgency is, unsurprisingly, a common theme.

Passwords are not the bottleneck

If emptying accounts armed only with a password is hard then are they truly the keys to the kingdom? There's ample evidence that banking passwords are being stolen at a considerable rate. For example Holz *et al.* [16] discover 10,770 banking passwords in a seven month examination of keylogger dropzones (*i.e.*, locations where keyloggers send their findings for later collection by the thief). Stone-Gross *et al.* [4] find 8,310 in ten days examination of the Torpig botnet (an annualized rate of 303,000). RSA reports finding 300,000 banking credentials in an examination of the Sinowal trojan. The Zeus botnet, which some accounts credit with infecting over

3 million machines, has the theft of financial data as one of its primary goals. Theft of non-financial credentials, occur at even greater rates. Within the last two years Rockyou leaked 32 million, Gawker 1.3 million, and the Waledac botnet was found to be in possession of 489,000 email passwords. It is often claimed that these non-financial credentials can be leveraged into access to more valuable accounts.

Thus, banking passwords are being stolen in considerable numbers. We have seen that emptying accounts is hard, and that mules, not victims, lose money. The password merely provides a way of offering something of apparent value (the victim-to-mule transfer) that will persuade the mule to part with something of real value (the mule-to-thief transfer). The victim's password is only one small part of that elaborate process of socially engineering the mule into parting with money.

If passwords are not the bottleneck what is? Backend fraud detection by banks is a good candidate. This reduces the number of compromised accounts that can be emptied. Mule recruitment is another good candidate for bottleneck. Studies of underground economy markets indicate great demand for mules [14, 12]. The Cisco 2010 Annual Security report claims "the ratio of stolen credentials to available mule capacity could be as high as 10,000 to 1." The RSA blog puts it succinctly [15, (Oct. 6, 2010)] "no mules = no cash." Krebs, who claims to have interviewed over 150 money mules [5, (May 11, 2010)] says "most money mules get a single transfer" and "Each mule is worth slightly less than $10,000 to the cyber gangs." It's difficult to imagine mule recruitment keeping pace with the level of credential theft mentioned above. Annualized, the Torpig [4] data alone (*i.e.*, one single botnet) would imply the need for a third of a million mules (at Krebs estimate of one transaction per mule).

Underground markets are not thriving

The cyber-crime underground economy is often portrayed as a criminal Utopia, rivalling above ground markets in activity and sophistication: it is claimed that illicit goods trade freely, and there is great specialization [14, 12]. The accounts suggest that some offer credentials for sale, some offer kits, and newcomers can buy what they need, and sell what they produce.

The parallels between over- and underground economies go only so far, however. One major point of difference is price: credentials are apparently offered for sale at pennies on the dollar on underground markets [12, 14]. Thomas and Martin report credentials with face value of $10 million being offered for $500. Symantec [2] reports accounts which it estimates as being worth $5.3 billion offered for $163 million. This is enormously puzzling if cashing out is easy and simple. Why would anyone sell the credentials that unlock an account with a $5000 balance for $5? It makes a lot more sense if emptying accounts is hard and stealing passwords is merely the first step in a difficult and error prone process which only occasionally succeeds [11]. If credentials are offered for sale at, *e.g.*, 5% of face value (this is a loose upper bound on the asking prices [14, 12, 2]), then 5% of the value goes to the person who steals the password, and 95% to the person who empties the account. This makes clear that emptying the account is by far the more valuable task. It defies common sense that those who steal passwords would give up 95% of the value of the finished product if they had any means of finishing the raw materials themselves.

In the chain of events that begins with stealing a password, and ends with the untraceable irreversible receipt of cash, passwords are merely one raw material that goes into the creation of the finished product. Every transaction requires a mule who is recruited and socially engineered into laundering the transaction. While passwords can be stolen on an industrial scale the same does not appear to be true of mule recruitment. The premium that emptying accounts enjoys shows that passwords are largely a commodity. This suggests that only a small fraction of the banking passwords that are stolen actually result in the successful extraction of money. If mules are scarce and stolen passwords are plentiful, then only the best prospects among the compromised accounts will be selected for evacuation.

What of the reports of easy money being made on underground markets? As far as we are aware no published account has claimed to have observed a single transaction closing, or a single dollar changing hands on underground markets [11]. The observations we have are of offers to buy and offers to sell [14, 12]. Reports that, *e.g.*, banking credentials are selling for $10 [2] does not mean that any transaction at that price has actually been observed. It merely means that at least one person who claimed to have those credentials offered to sell at that price at least once. Participants are anonymous, posting is free and can be automated, cheating is easy (and common [14]) and there is no contract enforcement. There is little to prevent, and every incentive to deal dishonestly. It is fair to say that many people in Internet chat rooms, bulletin boards and dating sites do not represent themselves truthfully. There is no reason to believe why this should be better on underground markets, and many why it might be worse. We have simply no idea what fraction of advertised transactions close. We have no idea what fraction represent real credentials as opposed to boastful claims, repeat sales, or attempts to cheat. Thus, estimating cyber-crime by taking activity on these channels at face value is, to put it no stronger, unsound. The view that underground markets are an easy-money Utopia is based on a rather credulous interpretation of the observations. Rather,

	Before Discovery	After Discovery
Victim	-$9000	$0
Bank	$0	$0
Mule	+$900	-$8100
Attacker	+$8100	+$8100

Table 1: Gains and losses of the various parties for a $9000 fraudulent transfer via a mule. Before discovery the victim is down the full amount and the mule receives 10%. After discovery the bank makes the victim whole (as required by Regulation E), and reverses the payment to the mule. The attacker is in effect stealing from the mule and not from the account he has compromised. If the mule has insufficient funds to cover the reversal, the bank is left with a (perhaps uncollectible) debt.

we suggest, it is the dumping ground for unused (and in many cases unusable) credentials that have little value.

Credential-Stealing is a terrible business

Suppose we ignore the illegal and unethical nature of credential stealing and evaluate it strictly as a business prospect. Is this "a business with some intrinsic durable competitive advantage"as Charlie Munger says that Berkshire-Hathaway demands of an investment? The advantages have been often discussed: stealing can be done remotely, it can be automated, there is little training or capital outlay required, prosecution is extremely rare. Almost anyone can do it. The popular and trade presses frequently run stories telling of easy cyber-crime riches.

Yet, there are also disadvantages. First, there is no barrier to entry and there is open access to the opportunity. New entrants keep arriving so long as the opportunity is profitable, which leads to the tragedy of the commons [10]. If a fixed pool of money is shared among many thieves the average return drops as more and more thieves arrive. This continues until the opportunity is no better than those elsewhere. However, the pool doesn't remain fixed. It shrinks *as a consequence of the efforts of the thieves* [10]. When stealing becomes common countermeasures increase: browsers deploy phishing warnings and blacklists, service providers do more rapid takedown of malicious sites, and banks place increased effort on back-end fraud detection. A steady stream of phishing emails may alert even unsophisticated users to the phenomenon. The average return thus has increasing denominator (thieves continue to arrive) with decreasing numerator (the pool shrinks).

There is no protection for intellectual or other property. Successful innovations are quickly copied by others, limiting their value to the originator. There is no lock-in, brand loyalty or other factor that helps maximize the revenue from a customer. Thus, stealing credentials meets none of Munger's criteria. Competitive advantages when they arise are neither intrinsic nor durable and the pace of change is relentless. Finally, there is no contract enforcement. Credential-stealing businesses cannot rely on even the most basic tool of commerce. This means that dishonesty is a way of life and dealing with anyone you don't know personally is fraught with risk. The lack of such a mechanism poses a profound difficulty in the development of a mature economy [9]. None of the ingredients that we typically associate with good businesses are present.

What of the size of the market? How big a pot of dollars is shared among password-stealing thieves? As we've shown, in the consumer space Regulation E implies that it is not victims who lose money but mules. Rather than targeting the account balances of all Internet users, credential-stealers are taking from those that can be persuaded to act as human relays. This is a small fraction of the population, and almost certainly concentrated among the poorest. This considerably limits the opportunity. We show elsewhere [8] that widely-circulated estimates that place cyber-crime losses in the billions are based on absurdly bad statistics and are entirely unreliable.

What of small businesses? Their losses are not covered by Regulation E and they are frequently targeted. Krebs, for example, has covered numerous cases of small businesses being successfully attacked [5, (June 28, 2010)]. The amounts are larger than in the consumer space, and banks are reluctant to shoulder the losses. Having more money than consumers, but lacking the security and audit controls that large organizations might have, small businesses might represent the ideal targets for credential-stealing criminals. While small businesses are better targets there are far fewer of them. The US Census Bureau finds that there were 1.25 million businesses with between 10 and 1000 employees in 2008. This is almost a factor of 200 lower than the number of consumers. If 1% of small businesses were successfully targeted annually, and the average haul was $10,000 (*i.e.*, the maximum amount to avoid a transaction report and the amount that Krebs claims a mule is worth) then this opportunity would be $125 million. While by no means small, this is a long way from the billions, or even trillion dollar losses that password-stealing is often claimed to generate.

Finally, in avoiding mule recruitment and other scams users are often cautioned that "if it sounds too good to be true then it is." This is sensible advice. However, it is no less sensible in examining the plausibility of stories of cyber-crime riches. It is naive, indeed, for a user to believe that a stay-at-home job requiring no training, skill, or experience will pay handsomely. However, it is not less naive to believe that a cyber-crime job requiring no training, skill, or experience will do the same. It makes no sense that a script downloaded from the Internet can generate a steady stream of income. It defies common sense that [14] "those without great skills can barter their way into large quantities of money they would never earn in the physical world." Legal or not, above ground or below, the demand for easy money seems likely to always exceed the supply. It is ironic that the magical thinking that we caution users against in their online affairs is baked into the consensus view of password-stealing.

Concluding remarks

In July 2009 a teller at a Key Bank branch in Seattle pursued a would-be robber after a botched hold-up attempt (Seattle Times, Aug. 1, 2009). He leapt over the counter, chased the man for several blocks, knocked him down, and held him until the police arrived. Two days later Key Bank fired the teller. He had violated long-standing bank policy to cooperate in every way and never resist a robbery. The reason for this policy, we suggest, is that banks understand a very simple principle: fear is bad for business. It is far better to comply with the demand than to risk a brawl, or a gunfight in the bank lobby. No bank wants the perception that they valued money more than customer and employee safety. The $40 million that traditional bank robbers in the US steal per year (FBI Bank Crime Statistics) is entirely manageable. Similarly, Regulation E and zero liability guarantees are not the result of altruism, they're just good for business. Limiting consumer liability lessens the anxiety about banking online. It costs very little if covered repudiable transactions are reversible.

The idea that consumers are "just a few clicks away" from grave financial harm makes a compelling narrative, but it is simply incorrect. This does not mean, however, that password-stealing is a minor problem. The indirect costs of cyber-crime almost certainly dwarf the direct losses by orders of magnitude. While password-stealing victims are spared direct losses, they may spend considerable time and energy resolving the mess. Mules bear the full brunt of successful cash-out operations, and are probably those least-able to handle the losses. The entire Internet-using public pays an enormous indirect cost in being compelled to adopt security measures that would not otherwise be necessary. Those who have had an email password stolen to send spam know what a miserable experience that is, and it is little consolation to hear that the hacker probably earned very little.

We acknowledge that the picture we paint in this article is not a proof; we present it as a plausible, rather than definitive, explanation of observations. However, the conventional view appears to require that banks do not understand the importance of non-repudiation. We suggest that our view is far more consistent with what we know.

Our challenge to the conventional view also raises interesting questions. First, if passwords are not the bottleneck, would replacing them, or making them harder to steal, have any influence on the total harm done by credential thieves? If a large lake of credentials is drained by a narrow pipe of mules then reducing the inflow to the lake might have no effect on the net harm done. Enormous energy has been devoted to the task of replacing passwords with something more secure. Yet, there is no clear picture of how much harm this would eliminate.

Second, it is sometimes assumed that banks wish to pass the liability for fraud to consumers. Is this really so? If emptying accounts is hard then credential theft losses may be far smaller than imagined, and borne by mules rather than banks. When perceived risk is greater than actual risk it can be profitable to absorb the risk and charge for it. Rental car companies are not merely willing, but anxious to accept liability for any damage to the car for $35 a day; various companies aggressively market identity theft protection for $12 a month. Banks enjoy a huge information advantage over consumers: they know how much fraud costs them, while consumers merely hear horror stories of cyber-crime losses [8]. Passing liability to consumers (as Anderson argues UK banks do [3]) would seem to be wasting a profitable opportunity.

Finally, many suggest that the switch in recent years from hacking-for-sport to hacking for financial gain represents an extremely serious escalation. This is sometimes offered as evidence that users must finally get serious about security, passwords must be done away with, *etc.* We offer the somewhat provocative thought that this switch is good news, not bad. The banking system has been hardened by centuries of exposure to fraud and money laundering. In spite of the enormous effort devoted to password-stealing, banks offer zero liability guarantees to customers and keep losses manageable. A fixed population of hackers will almost certainly do less harm by attacking hardened targets like banks than if they applied the same energy elsewhere. Getting in and getting out with money is a far harder problem than simply causing destruction. If the goal were mayhem and destruction rather than money-making we might be a great deal worse off.

Acknowledgements: the authors would like to thank the anonymous reviewers, and Joseph Bonneau, Paul C. van Oorschot and Frank Stajano for comments that

helped improve the article.

1. REFERENCES

[1] Regulation E of the Federal Reserve Board. http://ecfr.gpoaccess.gov/cgi/t/text/text-idx?c=ecfr&sid=0283a311c8b13f29f284816d4dc5aeb7&rgn=div9&view=text&node=12:2.0.1.1.6.0.3.19.14&idno=12.

[2] Symantec Report on the Underground Economy XII. http://eval.symantec.com/mktginfo/enterprise/white_papers/ent-whitepaper_internet_security_threat_report_xii_09_2007.en-us.pdf.

[3] R. Anderson. Closing the phishing hole–fraud, risk and nonbanks. In *Federal Reserve Bank of Kansas City, Conference on Nonbanks in the Payments System*, 2007.

[4] Brett Stone-Gross, Marco Cova, Lorenzo Cavallaro, Bob Gilbert, Martin Szydlowski, Richard Kemmerer, Christopher Kruegel, and Giovanni Vigna. Your Botnet is My Botnet: Analysis of a Botnet Takeover. *CCS*, 2009.

[5] Brian Krebs. KrebsonSecurity Blog. http://krebsonsecurity.com.

[6] Federal Reserve. Federal Reserve Payments Study. 2010. http://www.frbservices.org/files/communications/pdf/press/2010_payments_study.pdf.

[7] D. Florêncio and C. Herley. Phishing and Money Mules. *Proc. WIFS, 2010.*

[8] D. Florêncio and C. Herley. Sex, Lies and Cyber-crime Surveys. *WEIS, 2011, Fairfax.*

[9] A. Greif. Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition. *American Economic Review*, 1993.

[10] C. Herley and D. Florêncio. A Profitless Endeavor: Phishing as Tragedy of the Commons. *NSPW 2008, Lake Tahoe, CA.*

[11] C. Herley and D. Florêncio. Nobody Sells Gold for the Price of Silver: Dishonesty, Uncertainty and the Underground Economy. *WEIS 2009, London.*

[12] J. Franklin and V. Paxson and A. Perrig and S. Savage. An Inquiry into the Nature and Causes of the Wealth of Internet Miscreants. *Proc. CCS*, 2007.

[13] M. Mannan and P.C. van Oorschot. Security and Usability: The Gap in Real-World Online Banking. *NSPW*, 2007.

[14] R. Thomas and J. Martin. The Underground Economy: Priceless. *Usenix ;login:*, 2006.

[15] RSA Fraud Action Research Labs Blog. Follow the money, and go for the mules! http://blogs.rsa.com/rsafarl/follow-the-money-and-go-for-the-mules/.

[16] Thorsten Holz, Markus Engelberth and Felix Freiling. Learning More About the Underground Economy: A Case-Study of Keyloggers and Dropzones. *Reihe Informatik. TR-2008-006*, 2008. http://honeyblog.org/junkyard/reports/impersonation-attacks-TR.pdf.