

Kahn, Charles M.; Rivadeneyra, Francisco; Wong, Tsz-Nga

Working Paper

Eggs in one basket: Security and convenience of digital currencies

Bank of Canada Staff Working Paper, No. 2021-6

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Kahn, Charles M.; Rivadeneyra, Francisco; Wong, Tsz-Nga (2021) : Eggs in one basket: Security and convenience of digital currencies, Bank of Canada Staff Working Paper, No. 2021-6, Bank of Canada, Ottawa, <https://doi.org/10.34989/swp-2021-6>

This Version is available at:

<https://hdl.handle.net/10419/241229>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Eggs in One Basket: Security and Convenience of Digital Currencies

by Charles M. Kahn,¹ Francisco Rivadeneira² and Tsz-Nga Wong³

¹ University of Illinois at Urbana-Champaign

Federal Reserve Bank of St. Louis

² Funds Management and Banking Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9

³ Federal Reserve Bank of Richmond

c-kahn@illinois.edu, riva@bankofcanada.ca, russell.wong@rich.frb.org

Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.



Acknowledgements

We thank participants in the 2019 Canadian Economics Association Meeting and the 2020 ETH Zürich conference "Future Money," as well as in seminars at the Bank of Canada, FDIC, University of Illinois Urbana-Champaign and University of Illinois Chicago. We especially thank Rainer Böhme for valuable comments. The opinions here are of the authors and do not necessarily reflect those of the Bank of Canada or the Federal Reserve System. All errors remain our own.

Abstract

Digital currencies store balances in anonymous electronic addresses. We analyze the trade-offs between the safety and convenience of aggregating balances in addresses, electronic wallets and banks. In our model, agents balance the risk of theft of a large account with the cost to safeguarding a large number of passwords for many small accounts. Account custodians (banks, wallets and other payment service providers) have different objectives and trade-offs along these dimensions; we analyze the welfare effects of differing industry structures and interdependencies. In particular, we examine, the consequences of "password aggregation" programs, which, in effect, consolidate risks across accounts.

Bank topics: Digital currencies and fintech; Financial services; Payment clearing and settlement systems; Central bank research

JEL codes: E, E4, E42, E5, E51, E58

1 Introduction

Custodians of financial assets have always needed to weigh customers’ conflicting goals of security and ease of access. However, the rapid development of new electronic arrangements for storing customer balances has made these trade-offs more salient. Digital currencies, in contrast to traditional bank accounts, store balances in *addresses* that do not need to be associated with the identity of the owner,¹ making them, at least potentially, more vulnerable to loss. Without an ability to identify individuals, the users, payments service providers and regulators face different trade-offs of safety and convenience. Analyzing these trade-offs is necessary to guide the design of publicly issued digital currencies and the regulation of private ones.²

In this paper, we consider the security risks of transaction balances and how these risks are shared between users and providers. We ask whether there are externalities from storing large balances and, if so, what the role of regulators should be. For example, should regulators establish standards for passwords and other safeguards, or should we expect competition to achieve an efficient outcome? As new techniques appear for improving safety or convenience, will they be adopted or blocked?

Balances of digital currencies are stored and transacted using addresses and wallets. An address is a location to store valuables. In the case of digital currencies, the valuables are uniquely identified digital objects called coins. Addresses can contain multiple coins or none at all, but each address has one associated private key: a password. Coins inside an address are kept separate and individually identified. In other words, the balance of an address is composed of individual coins that may have different denominations, similar to depositing several physical coins or bank notes in a safe deposit box. In this sense an address is an aggregator of coins.

Another way to maintain balances of digital currencies is to use a wallet, which manages the private keys of the addresses.³ In Bitcoin, the most prominent digital currency, the system itself establishes the safety protocols of addresses, which are extremely good.⁴

¹In this sense, digital currencies are a type of token-based payment system like cash. For more details on the distinction between token-based and account-based systems, see [Green \(2008b\)](#), [Kahn and Roberds \(2009\)](#), and [Kahn \(2016\)](#). For a contrasting view, see [Milne \(2018\)](#).

²A central bank could issue public digital currencies in a scheme where it tracks user identity. However, [Kahn et al. \(2018\)](#) argue that central banks are unlikely to choose this form for retail transactions. [Barontini and Holden \(2018\)](#) survey central banks’ interest in issuing digital currency; [Mancini-Griffoli et al. \(2018\)](#) discuss the trade-offs surrounding the issuance of central bank digital currencies.

³Different wallets provide various services, but at a minimum they manage the private keys and basic transaction functions—for example, choosing which coins to use in a transaction.

⁴To be able to crack a private key of a Bitcoin address by brute force, a hacker would need to try

In contrast, wallets can provide their own level of security and potentially other identification requirements, alternative means of recovering the private keys in the wallet, management of transactions and other convenient features. In most digital currencies, wallets are provided privately and competitively, offering convenience to the users.

Access to the balances in the addresses and wallets is protected by passwords, but with different levels of strength. Although the level of security of Bitcoin private keys is extremely high, directly managing them can be quite cumbersome. Since a wallet is a collection of addresses and its respective private keys, a compromised password for a wallet endangers the entire balance in the wallet.⁵ We call this *contagion*.

We propose a simple framework to analyze several aspects of digital currencies. In our environment, agents try to economize on the cost of remembering and managing passwords while maintaining the safety of their balances. One of the problems agents face when carrying out this task is the question of optimal aggregation: among how many accounts should agents divide their balances? Is it a case of “don’t put all your eggs in one basket,” or, as Mark Twain said, “put all your eggs in the one basket and—*watch that basket*.”⁶

To analyze the trade-off of security and convenience we focus on the transactions motive for holding balances. A customer conducts purchases by withdrawing from the accounts provided by competing “banks” (think of wallet providers or crypto exchanges). There are two types of malefactors: “Hackers,” who try to steal balances directly by brute force, and “thieves,” who attempt to obtain customer passwords. Theft in our framework is conducted using “man in the middle” attacks, social engineering or exploitation of software vulnerabilities. Every time an account is accessed to check the balance or make a withdrawal, there is a risk of a thief intercepting the communication between the investor and the bank, thereby obtaining the password. Contagion puts the rest of the balance of the account at risk. Therefore, in equilibrium, customers choose the best trade-off between convenience and risk, which depends on the probability of contagion, the frequency of withdrawals and the costs of keeping track of passwords.

Theft also occurs because of the consumer’s inability to properly follow the protocols to ensure safety. One example is weak private keys that are generated by mistake or

⁵ $2^{256} - 1$ different keys. There are just under 2^{160} addresses. Some addresses therefore have more than one corresponding public key and thus more than one corresponding private key.

⁵In their shortest human-readable form, private keys can be represented by a string of 30 case-sensitive characters. Typical wallets in turn can be secured with human-readable 12-word pass phrases.

⁶Clemens (1894) *Pudd’nhead Wilson*, Chapter 15, emphasis Twain’s. He takes the maxim from a speech by Andrew Carnegie in 1885.

because the protocol is too complex to use.⁷ This highlights the interaction between the design of the protocols and the effort that customers put into securing their private keys.

Theft of private digital currencies is quite common. Although their cryptography makes them virtually impossible to hack by brute force, the wallets and exchanges that hold them have been subject to frequent hacking and malfeasance. One prominent case is the failure in 2014 of the Bitcoin exchange Mt. Gox, which lost an equivalent of \$473 million at the time of bankruptcy. Since then, cases of hacking of private keys and exchanges have become more common. A report suggests that in 2018, an equivalent of \$950 million was stolen from exchanges ([CIPHERTRACE 2019](#)).

Although these questions are similar to the questions around security and convenience of cash or traditional bank accounts, new issues emerge from the electronic nature of digital currencies. First, contagion is not a serious concern for cash: a stash under a mattress is not immediately put at risk of theft when spending a portion of it at some other location. In the case of traditional bank accounts, identification requirements permit, to a certain extent, the reversal of fraudulent transaction. Moreover, issues of scale pose a greater challenge for new technologies. Cash is bulky and costly to manage in large quantities, but for digital currencies the same private key can manage \$1 or \$1 billion.⁸ In addition, managing the private key of an address requires more attention and knowledge than protecting a bank note. Finally the digital environment enables the arms race between payments system providers and customers: providers include cumbersome security features, customers find software apps that allow evasions, and providers in turn develop deterrents for these evasions.⁹

Literature

Public key cryptography ([Diffie and Hellman 1976](#) and [Merkle 1978, 1980](#)) is the foundation of digital currencies because it allows a receiver of money to identify the sender as the rightful owner of the money being sent. Computer science evaluates a cryptographic protocol by testing the probabilistic safety of its mathematical primitives assuming the protocol is used as intended. In the case of the public key cryptography implemented in most digital currencies, the private key is assumed to be secret. This is the weakest link we examine in this paper.

⁷[Independent Security Evaluators \(2019\)](#) show that in spite of the statistical impossibility of randomly generating the private key of an existing address, typical mistakes in key generation allowed them to find 732 Ethereum private keys in a short period of time.

⁸Nowadays it is hard to find traditional physical bearer instruments of large denominations. In 1982, the U.S. congress restricted the issuance of bearer bonds ([Briner 1983](#)).

⁹[Herley \(2009\)](#) describes some of the practical consequences of this adversarial relationship.

Of course, computer scientists have recognized that the effectiveness of a security protocol depends as much on economic incentives for users as on its mathematical foundations (Moore 2010, Anderson and Moore 2007, 2009, and Herley 2009). There is also literature on interdependent security games that straddles the boundary between economics and computer science.¹⁰ Papers in this literature typically consider games among firms whose risk of a security breach depends on both their own and other firms' defensive actions. In general, the games examined in our paper can be also be classified as part of this literature. However, we focus on the differing incentives of providers and customers as well as malefactors in a digital currency environment.¹¹

This paper is also related to the broader literature that discusses the role of central banks in providing safe and efficient payment systems. There is consensus on the importance of the role central banks play in high-value payment systems (Green 2008a and Lacker 2008) and in the provision of physical currency, an arrangement that has allowed for many forms of private payments systems to emerge. There is no consensus, however, about whether central banks or private firms should provide digital currencies and, if so, how these digital currencies should be designed and how should they interact.¹²

The fact that digital currencies allow storing balances and transacting anonymously, raises the issue of privacy. Kahn et al. (2005) show the value of anonymity in an environment with moral hazard. When online transactions cannot be conducted with cash, credit-based methods expose the user to identity theft from the seller or an intermediary. This highlights the value of electronic money alternatives. In contrast, our paper shows the emergence of another form of moral hazard when the safety of digital currencies is at risk from a potentially inefficient level security chosen by users or wallet providers.

We proceed as follows. Section 2 discusses the technological details of security and account management of digital currencies. Section 3 describes the framework. In section 4 we model the equilibrium password strength and associated number of accounts under various management strategies. In section 5 we explore password reuse, password aggregation software and the implications of competition between providers. Section 6 discusses policy interventions. We conclude in section 7 and suggest some open research

¹⁰For a survey, see Laszka et al. (2014). Early contributions on the economics side focus on security as a public good. See Hirshleifer (1983) and Varian (2004).

¹¹Most of the theoretical work in economics on digital currencies has focused on their effects for monetary policy and their efficiency as means of payments. See Davoodalhosseini (2018) and Chiu and Koepl (2017) and citations therein. All of these analyzes assume away the risks associated in the management of the private keys.

¹²For a wide variety of opinions, see, He et al. (2016), Bordo and Levin (2017), Mancini-Griffoli et al. (2018), Brunnermeier et al. (2019), and Duffie (2019).

questions. The Appendix presents additional model details.

2 Aggregation and Security Risks

To understand the risks in digital currencies we need to describe how the different levels of aggregation work. We also provide some empirical evidence of the security risks. We use the example of Bitcoin, which is the most popular digital currency and a model for many others. The unit of account in Bitcoin is the satoshi, which is equal to 1×10^{-8} bitcoin. However, the Bitcoin system does not keep track of every satoshi location but of *bundles* of them. These distinctly identified bundles are called unspent transaction outputs (UTXOs).


Figure 1 depicts the different levels of aggregation. Addresses form the first level of aggregation. Bitcoin addresses have unique private-public key pairs, randomly generated at no cost. More precisely, an address in Bitcoin is an alphanumeric string that represents a *possible* destination for a Bitcoin payment, i.e., a location to store UTXOs. The public address is broadcast to the network, while the private key should be kept secret. Addresses are like safe deposit boxes that can hold an amount of Bitcoin. In fact, a few addresses hold billions of dollars worth of bitcoins. By combining the private and public keys, an individual can take the coins out of that box to send to a new box. The security of the *entire* balance of an address depends on the secrecy of the private key. In other words, the private key is the password to whatever is in the box.

The next level of aggregation of balances is provided by wallets. A wallet is software that manages public and private keys and is in turn protected by its own master key. In addition, wallets may allow additional means for users to protect and recover the access to them. Wallets perform the functions of selecting the UTXOs when spending and manage the new addresses for change transactions.¹³ The user is responsible for keeping the master key safe. If the wallet's key is lost or stolen, the wallet developer bears no responsibility for any resulting loss of funds.

The third level of aggregation is provided by firms holding balances on behalf of cus-

¹³“Hierarchical Deterministic” (HD) wallets is a standardized protocol used by many private digital currencies. Such wallets have the ability to hold many types of coins. HD wallets handle change transactions by creating a tree of addresses. For example, suppose a person owns 2 UTXOs: one with 2 bitcoins and another with 1 bitcoin. If this person sends 1 bitcoin to a new address, the wallet can choose to use the first UTXO and create “change” of 1 bitcoin, which will be sent back to the sender to a new address as a new UTXO. The other alternative is to send the second UTXO for which there would be no “change.” HD wallets are secured by a master password that can be recovered with a 12-word phrase.

Figure 1: Levels of aggregation of digital currencies and their associated security setup. Tokens are contained in addresses that are secured by private keys. Wallets can manage the private keys of multiple addresses and are secured themselves by a master key. Providers, like exchanges, aggregate funds of customers into accounts and secure them by passwords. Typically, as the stored value increases, customers use higher levels of aggregation.



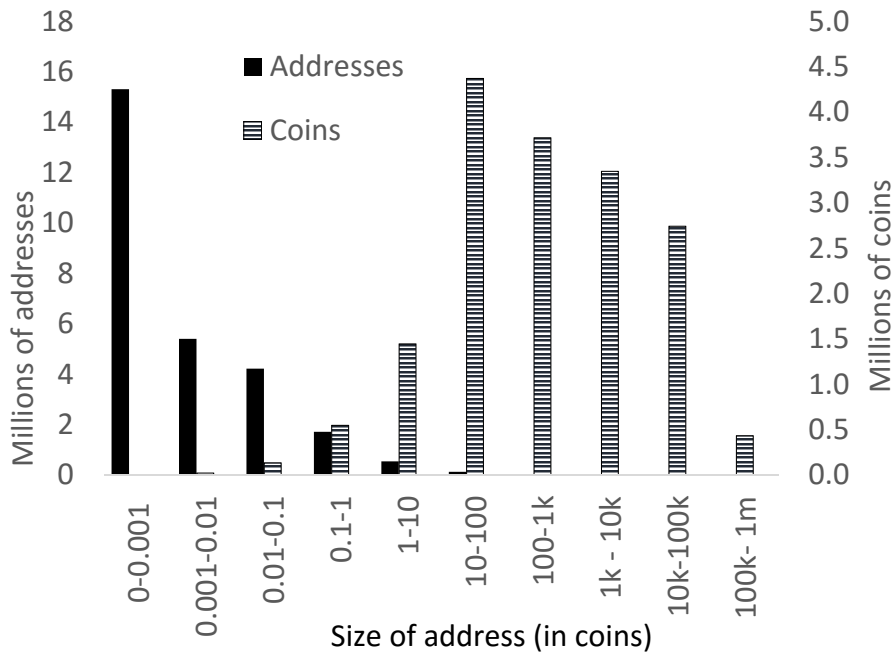
tomers. These firms, known as exchanges, can provide either wallet services or accounts that in fact commingle the funds of customers. In the latter case, the customer surrenders the private keys to the firm and receives a deposit.¹⁴ The exchanges can provide on-us payments (reducing the transaction cost and delay of digital currencies) or liquidity between various digital currencies earning spreads on these transactions. Some of the exchanges further aggregate the balances of their clients into new addresses holding large balances managed in “cold” wallets that are not connected to the internet. In summary, as balances are aggregated, security risk become concentrated at the highest level.

We can gauge the magnitude of the security risks we are discussing by examining the Bitcoin distribution of addresses and coins by amount of coins in each address. Figure 2 shows that the large majority of addresses have very small amounts of coins. These might be “change” transactions or lost addresses. Most of the value of the Bitcoin system is held in the addresses containing between 10 and 100 coins. At the other extreme, a handful of addresses each hold more than \$500 million dollars and one holds \$1.2 billion dollars (at \$10,500 U.S. dollars per bitcoin). Frequently, the largest addresses are owned by the exchanges that manage funds on behalf of clients. As of this writing, six of the ten largest addresses are cold wallets owned by these firms.

All of the theft in Bitcoin and other major digital currencies has occurred because the security of the wallets or exchanges was compromised, and not by the brute force hacking of the private keys associated with addresses. For example, private keys can be stolen from individuals’ computers by exploiting simple software vulnerabilities. Also, some exchanges that provide wallet services have had their security compromised or had

¹⁴Note that this potentially alters the legal liability in case of theft.

Figure 2: Distribution of the number of addresses and of bitcoins by size of address (in number of coins) as of 2018. The vast majority of addresses contain small number of coins, either from change transactions or lost coins. The bulk of coins are contained in addresses that have between 10 and 100 coins, between \$100,000 and \$1,000,000 dollars at today’s prices. A handful of addresses hold more than \$1.6 billion dollars. The security of the coins in a given address relies entirely on the secrecy of that address’s private key. Source Bitinfocharts.com.

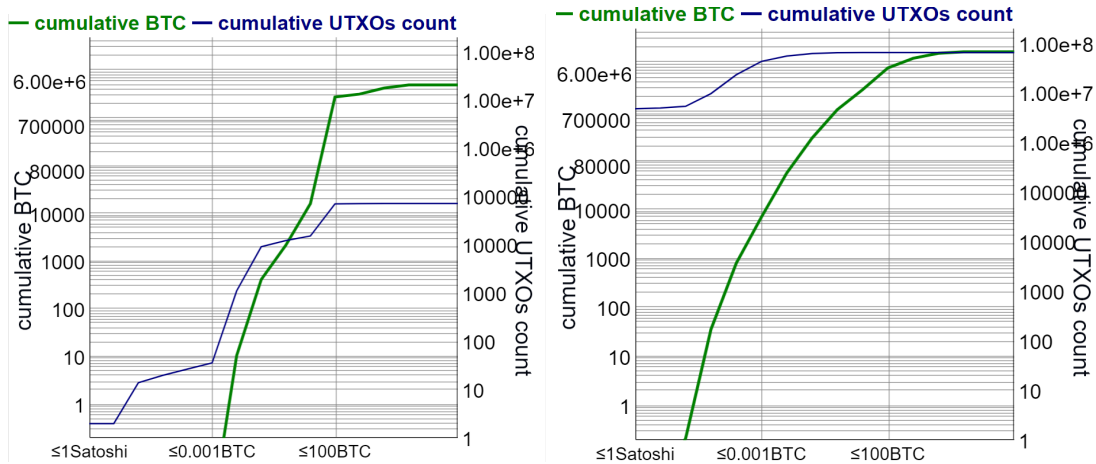


rogue employees (for example, Mt. Gox and Coincheck).

In 2018, \$950 million dollars worth of digital currencies was stolen from established exchanges (CIPHERTRACE 2019). In addition, it is quite likely that a large portion of the theft of digital currencies (in wallets or exchanges) via the capture of their private keys is not reported at all because the exchanges are not able to distinguish a fake from a real claim of theft, leaving customers to bear the entire loss. In general, it is hard to know what share of the theft is due to vulnerabilities in security, scams, or carelessness on the part of the exchange, but analysis of the self-reported breaches by exchanges suggests that the majority of them are due to a hack to exchanges’ hot wallets (those connected to the internet) or to an application vulnerability (a coding error).¹⁵

¹⁵See <https://magoo.github.io/Blockchain-Graveyard/> for a breakdown of causes of the hacks.

Figure 3: Cumulative distribution of unspent transaction outputs (UTXO) by number of bitcoins in each address as of 2011 (left) and 2018 (right). An UTXO is a collection of coins in a given Bitcoin address. Source Bitinfocharts.com.



3 Framework

This section briefly outlines our framework; in subsequent sections we examine a variety of simplified versions of this general structure to consider specific issues. Section 4 examines password length and number of accounts. Section 5 examines how security protocols interact with the effort exerted by customers.

The world consists of *customers*, *banks* and *malefactors*. Customers divide their wealth among one or more accounts, which they access with some frequency. We will focus on the transactions motive for holding balances. Customers need to make a fixed number of payments per period. To make a payment, the individual makes a withdrawal from an account.¹⁶

Banks hold customers' accounts and establish protocols to maintain a level of safety in accessing accounts. We can interpret banks as wallet providers and an account as a particular wallet.¹⁷ In this world there is no way to identify individuals when they make withdrawals. Instead, we assume that for each account the bank provides a q -bit password to be used as identification. If the proper password is given, the individual may make a withdrawal of funds. If the individual has more than one account, he will have a different password for each account; this assumption is relaxed in the later sections.

¹⁶We assume holding funds at home is prohibitively inconvenient or dangerous, thus withdrawals only occur to make mandatory payments (or to transfer to another account).

¹⁷There are five types of wallets for keys: online, mobile, desktop, hardware and paper. Our definition of bank is broad enough so that all these types of wallets fit in our model. See <https://medium.com/@fastinvest/the-most-comprehensive-cryptocurrency-wallet-guide-5e820a26ed44> for an explanation of their differences.

In addition to fees charged by the banks, customers suffer disutility from the effort they take to protect their accounts and the need to keep track of passwords, as well as from any other difficulties in accessing the account, which depend on the level of security imposed by the bank.

We will denote the bank’s cost function by $K(\cdot)$ and the customer’s cost function by $c(\cdot)$. Different sections of the paper will focus on different aspects of these costs. We assume that both banks and customers are risk neutral in monetary losses. This allows us to present results as starkly as possible; extensions when customers are risk averse will be clear.

We distinguish two types of malefactors. *Hackers* focus on banks and make brute force attacks on customer accounts in banks. For simplicity, we will assume that it is length of password that serves as the deterrent to hackers. *Thieves* focus on customers, gaining access by exploiting customers’ transactions activity—in effect using “man-in-the-middle” exploits. Thieves are deterred by the care with which customers act and by the complexity of the protocols maintained by the banks. The crucial distinction is that thieves are aided by the frequency with which a customer uses a password, while hackers are aided by the simplicity of the password. In welfare calculations, we will consider the costs to banks and customers, ignoring any cost imposed on hackers and thieves.

4 Password Strength and Account Management

A customer receives a fixed amount of income I each period, with which the customer makes a fixed number T of equal-sized payments. Meanwhile the customer holds funds in accounts at one or more banks. Let n be the number of accounts the customer holds and let Q be the total number of bits in all of the customer’s passwords. Then, we represent the customer’s direct disutility from dealing with the accounts by $c(n, Q) = n\alpha + C(Q)$, the first term being the fixed cost per account and the second the (increasing) cost of keeping track of all the passwords. Later, we will extend the analysis to consider costs associated with customer effort at account protection and variation in inconvenience associated with different accounts.

For this section, a bank’s costs $K(y, \mathcal{N}, \mathcal{B}, \mathcal{Q})$ depend on the number of customers, y , the total number of accounts, \mathcal{N} , aggregate balances in all the accounts, \mathcal{B} , and the aggregate level of account security, summarized by \mathcal{Q} , the total number of bits in all passwords of all accounts at the bank. We assume costs are increasing in y , \mathcal{N} , and \mathcal{Q} .

Banks pass their costs through to the account holders via fees.¹⁸

4.1 Password Hacking

First, we consider the hacker, who attacks a bank by generating passwords at random in an attempt to obtain funds. Let h be her cost for each attempt to gain access to an account. If successful, the hacker will obtain the funds in the account. If unsuccessful, she can always try again.

Suppose the average balance in an account is b . Given that the bank services \mathcal{N} accounts, each protected by a password of length of q bits, the expected payoff to a hacker from an attempt to access funds at the bank is $2^{-q}\mathcal{N}b - h$. If this is positive, then hackers will enter.¹⁹ Thus, for banks to be viable, password length has to be $q \geq q^*$, where

$$q^*(\mathcal{B}) = \log_2(\mathcal{B}/h). \quad (1)$$

Note that the temptation to hack, and therefore the required size of the password, is associated with the total value of funds in the bank, not with the amount available in any particular account.²⁰ Henceforth, we will assume that the bank sets q to the deterrent

¹⁸This is a reasonable assumption in the environment of private digital currencies, which are typically open-source platforms, so that there is free entry to develop wallet technology. For example, as of 2019 there were at least two dozen approved mobile wallets in Apple’s app store for Bitcoin. Other digital currencies, such as Ethereum, also have many wallets available.

¹⁹This claim is making several large simplifying assumptions. First, we are assuming that all passwords are equally likely. This is a reasonable assumption when the bank provides passwords (provided the bank’s password randomizing procedure is not deficient) but not when customers choose their own passwords. See for example [Bonneau et al. \(2012\)](#).

Second, we are assuming that repeated failures do not alert the bank to shut down access to the accounts. This would certainly be an unrealistic assumption in the case of multiple attempts on a single account. It might be somewhat less distasteful if we are considering attempts on random accounts at the institution, although even then alert IT personnel would be expected to see the pattern of attempted intrusions. One alternative is to assume that the repeated attempts are not instantaneous but carried out with sufficient lag as to reduce the possibility of detection (although such delay would significantly increase the cost of each attempt). On the other hand, repeated attacks without delay are feasible in the bitcoin environment; for example [Vasek et al. \(2017\)](#), who study the hacking of bitcoin “brain wallets” (a method of storing private keys in password-protected publicly accessible locations), conclude that passwords for brain wallets are regularly successfully hacked.

The third simplification is to ignore the fact that as a hacker try more passwords, failures increase her subsequent chances of success, since she amasses a list of passwords that *don’t* work. Thus, the chance of success of each attempt is not constant, even if the distribution of passwords is uniform. However, with large values of q this extra complication is of minor importance. For calculation of bounds on the expected number of attempts needed to hack a single account, see [Massey \(1994\)](#)

²⁰As predicted, attacks on exchanges increased dramatically with the rise in the value of cryptocurrencies in 2017 ([EY 2017](#) and [Stecklow et al. 2017](#)).

level q^* . Thus the bank's costs are

$$K(y, \mathcal{N}, \mathcal{B}, \mathcal{N}_{q^*}(\mathcal{B})).$$

Suppose we vary the number of accounts per customer while holding the number of customers and the total balances unchanged—that is, we let \mathcal{N} vary while holding y and \mathcal{B} constant. Bank costs are minimized by setting \mathcal{N} as small as possible—one account per customer. Meanwhile, if all the customer's accounts are at the same bank, then the customer's costs are

$$n\alpha + C(nq^*)$$

which again is minimized by reducing n .

In other words, customers and the bank agree that when hacking is the sole threat, there is no advantage to holding multiple accounts at a single bank. For example, consider a customer who has three accounts at a single bank, each protected by a different four-letter (essentially 20-bit) password. Suppose the customer is asked to consolidate those accounts into a single account with a single 12-letter password. By our assumptions, the disutility to the customer would decrease, and so if the safety of the arrangement were the same, then the customer would willingly agree to the change. Even if α were zero (that is even if there were no fixed disutility attached to additional accounts), the customer would still be indifferent regarding the change.

Nonetheless, even if the customer were indifferent, the bank would not be indifferent. One large account under a 12-letter password is a greater deterrent to hackers than three accounts under different four-letter passwords. Hackers are tempted by the amounts held as a whole in the bank. If those amounts are split into smaller accounts, then the hacker gets less money on any successful attempt, but this is exactly made up for by the greater chance of hitting on an account. In other words, if the bank is big enough that a 12-letter password is necessary to deter the hacker from attacking and if the accounts are split into smaller accounts, then each of those smaller accounts will still have to have its own unique 12-letter password. That is to say, consolidating a customer's accounts within the bank is cheaper for everybody (except the hacker).

If costs are passed through to the customer, then the customer will prefer a single large account at one bank to multiple small accounts across banks. In a world with identical customers, all banks will pick the same password length, and each individual will concentrate his holdings in a single account at a single bank. With heterogeneous wealth, different banks will specialize in different sized accounts, but again customers

will deal with just one bank. In other words, if password hacking is the only threat, bank accounts should have large balances and passwords should be correspondingly long.

Now to compare the optimal with the actual password length, take the case of Bitcoin. The value of all bitcoins is US\$100 billion. We can approximate the hacking attempt cost with the cost of a computing operation. Currently, the best specialized processors can perform 16 Terahashes per second or 16×10^{12} operations/sec. Assuming the only cost is electricity, with a consumption of 120W/h and an electricity cost of US\$0.15kW/h, each operation would cost US\$ 1.86×10^{-17} . In reality, hacking attempt costs would be higher because of the fixed costs of hardware and because banks will surely respond to the hacking attempts. With these assumptions, the minimum password length required is $q^* = \log_2(10^{11}/1.8^{-17}) = 92$, compared with the current length of 256. This shows that current password length is many orders of magnitude larger than the minimum necessary.²¹

4.2 Password Theft and Contagion

The second type of danger is password theft. The distinguishing feature of theft is that it becomes more likely the more frequently the password is used.

Recall that money is withdrawn T times, and suppose each withdrawal leads to a fixed probability π of a disclosure of the password to a thief.²² To minimize the risk of theft, it is optimal to withdraw from the same account until that account is exhausted. This is because it is the withdrawal that exposes the account to risk of theft. If the agent withdrew from various accounts concurrently, then this would expose more balances for longer periods of time until the balances of all accounts are exhausted from transactions.

Thus, each account would have T/n withdrawals in succession. Each account contains at the start I/n in income. After one withdrawal has been made, then there is $(I/n) - I/T$ left in the account. After the second withdrawal, there is $(I/n) - 2(I/T)$ and so forth until after T/n withdrawals nothing is left. After each withdrawal, whatever remains in the account is the amount that can be stolen, and theft occurs with probability π on

²¹The calculation of the market value that would make a hacker break even is $\$1.015 \times 10^{77}$, an absurd value, or the cost of hacking attempts $\$1.8 \times 10^{-247}$, effectively zero.

²²Probability of theft in our model comes from the ability of thieves to intercept communications between the bank and the customer. In reality this power can come through a variety of attacks, such as keystroke logging or actual interception of communications. If we assume the draws are independent, then the probability of the first disclosure on the first withdrawal from an account is π , on the second withdrawal is $(1 - \pi)\pi = \pi - \pi^2$ etc. For simplicity we approximate the situation by assuming that the probability remains π each time, and higher order terms are ignored.

every occasion. Since there are n accounts, the expected amount lost to theft is:

$$n\pi \sum_{k=1}^{(T/n)} \left(\frac{I}{n} - k \frac{I}{T} \right) = \frac{\pi I}{2} \left(\frac{T}{n} - 1 \right). \quad (2)$$

When banks are competitive, the costs of banks are passed on to the customer. Thus the competitive equilibrium can be derived by solving the cost minimization problem of the customer:

$$\min_n \underbrace{\alpha n + C(nq^*)}_{\text{private costs}} + \underbrace{K(y, \mathcal{N}, \mathcal{B}, \mathcal{Q})}_{\text{bank cost}} + \underbrace{\frac{\pi I}{2} \left(\frac{T}{n} - 1 \right)}_{\text{expected theft}} \quad (3)$$

where increasing n by 1 increases \mathcal{N} by 1 and \mathcal{Q} by q^* . Treating n as a continuous variable, the solution is:

$$n = \sqrt{\frac{(\pi IT)/2}{\alpha + K_{\mathcal{N}} + q^*(C' + K_{\mathcal{Q}})}}, \quad (4)$$

where $K_{\mathcal{N}} + q^*K_{\mathcal{Q}}$ is the marginal cost to the bank of opening an additional account for an existing customer without receiving any additional funds; this consists of a fixed component and a component that depends on the complexity of the account security, here proxied by password length.

Consider the case where the various marginal cost components are constant. As is intuitive, the optimal number of accounts for the customer increases with T , the number of withdrawals to be made; with π , the probability of success of theft; and with I , the total amount initially deposited by the customer. Thus, holding the number of withdrawals fixed, the number of accounts increases with the size of the individual withdrawals I/T . On the other hand, the number of accounts decreases with the costs associated with an account. So, for example, as h decreases (hacking gets easier) the necessary password length q^* increases, and customers respond by reducing the number of separate accounts they hold.²³

²³As in the previous section, when customers are heterogeneous then, under competition, banks will specialize in accounts of different sizes: larger banks will hold larger accounts. Results are similar if the banks are able to extract all consumer surplus through price discrimination (on, for example, size of balance). The general case of heterogeneous customers under oligopoly merits further examination. It is interesting to note the similarity of equation (4) with the classic cash demand result of Baumol (1952) and Tobin (1956).

4.3 Account Management

So far, customers deplete their accounts one by one. Another way to manage the risk is to maintain a hierarchy of accounts such that a large sum of money is kept in a high-level account, occasionally tapped for transfers to a lower-level account, from which the customer makes more frequent withdrawals. This management strategy combines two benefits: maintaining large accounts with high security to control the risk of theft and keeping only small amounts in the frequently accessed accounts to limit the loss to theft. In what follows, we consider the case of a depositor who maintains only two accounts—an “investment” account and a “transaction” account; the case of a more general hierarchy is considered in the appendix.

Label the individual’s two accounts with subscripts i and t respectively. In any period, I is deposited into the investment account. As before, the individual will make T payments during the course of the month. In this setup, the customer’s problem is to determine the optimal number of transfers between accounts. In what follows, n will be the number of transfers from the investment account to the transaction account in a period.

As before, withdrawals incur a probability of theft: π_i and π_t for the investment and transaction accounts respectively. This probability is the risk of theft of the password of the account from which the amount is being withdrawn and not the account to which the funds are being transferred. In other words, we assume that there is no risk of theft associated with depositing funds.²⁴

For concreteness, we assume that when I/n is withdrawn from the investment account, the portion I/T is used directly for payment and the remnant is deposited into the transaction account. Then the transaction account is used to make payments until it is exhausted, at which point the next withdrawal is made from the investment account. Thus, in this environment, the total number of withdrawals is fixed. The more frequently money is moved from the investing to the transactions account, the higher the average balance in the transactions account relative to the average balance in the investment account. As a function of the number of transfers n , the expected loss from theft is:

$$\frac{I}{2} \left((n-1) \pi_i + \left(\frac{1}{n} - \frac{2}{T} \right) (T-n) \pi_t \right).$$

Intuitively, the average holdings in the investment account are $\frac{I}{2}(1 - n^{-1})$, and the

²⁴This would be the case, for example, if the account numbers for depositing in an account were kept distinct from the account numbers for withdrawals, as is the case in some countries.

probability of a breach of the account is $\pi_i n$. Similar calculations apply to the transactions account. Therefore, the choice of the frequency of withdrawals from the investment account is:

$$n = \sqrt{\frac{T}{\pi_i/\pi_t - 2T^{-1}}} \simeq \sqrt{\frac{T\pi_t}{\pi_i}}. \quad (5)$$

As expected, the frequency at which the funds should be moved between the two accounts goes up as the square root of the number of transactions to be made and with relative likelihood of theft in the two accounts.

The sequential account management strategy described by (4) and the tiered strategy described by (5) have close counterparts in the common practices in cryptocurrencies. The sequential account management strategy suggests that although the re-use of addresses for transactions is strongly discouraged by user guidelines of many cryptocurrencies due to security (and privacy) risks, users find it optimal to use a single address for a few transactions.²⁵ Of course, one way to avoid re-using addresses and managing their keys is to use key management software, such as a wallet. We analyze this case in the next section.

If we add in terms d_i and d_t , for the customer's cost of inconvenience per withdrawal from the two accounts respectively, then the objective becomes the minimization of:

$$\min_n \frac{I}{2} \left(\left(\left(1 - \frac{1}{n} \right) \pi_i + d_i \right) n + \left(\left(\frac{1}{n} - \frac{1}{T} \right) \pi_t + d_t \right) (T - n) \right)$$

and the solution in this case is:

$$n = \sqrt{\frac{T}{(\pi_i + d_i - d_t)/\pi_t - T^{-1}}} \simeq \sqrt{\frac{T\pi_t}{\pi_i + d_i - d_t}}. \quad (6)$$

The intuition is similar to that for equation (5). In Appendix A.3 we derive exact versions of these equations with an alternative assumption of the timing of the first withdrawal.

The two-level hierarchy strategy is a common practice among online exchanges of cryptocurrencies. These firms use two types of private key storage, colloquially called hot and cold storage. Hot storage refers to key management databases that are connected to the internet and therefore at a higher risk of theft from hackers. Cold storage, in contrast, refers to key management databases that are connected to the internet only on the occasions when transfers are required and therefore at lower risk. This practice is intended to limit the potential theft in case exchanges are hacked. Jain et al. (2018)

²⁵See the guidelines of Bitcoin about address reuse: https://en.bitcoin.it/wiki/Address_reuse.

analyze a stochastic version of a model for optimizing an exchange's transfers between hot and cold wallets, arriving at an equation analogous to (5).

5 Security Protocols and Customer Response

In reality, the above strategies are stark simplifications because customers interact with banks on several dimensions beyond the choice of the number of accounts or frequency of transfers among them. For example, banks can choose from different security protocols, and customers undertake different levels of care to keep accounts secure. In this section, we consider the interaction between the security protocols established by banks, customer effort, and password aggregation programs. First, we focus on a single bank/customer pair and endogenize the probability of theft. In the following subsections we consider the case where two banks interact by choosing security protocols knowing that customers may reduce their costs of compliance with bank requirements by, for example, reusing passwords or using password aggregation programs. To concentrate on these other dimensions, we will assume throughout this section that a customer has at most one account with each bank and that each account is accessed once a period.

5.1 Customer Response

To begin with, assume that the customer has a single bank account. Theft results in a loss of L . The probability $\pi(s, e)$ of a loss is now a convex, decreasing function of the level of security s chosen by the bank and the effort e (level of care) chosen by the customer. In general, the cost depends on both parties' actions; we will assume that the cost to the bank $K(\cdot) = s$ and the cost to the customer is a convex function $c(s, e)$. We will focus on situations where increasing the level of security increases cost of effort to the customer.

The socially optimal arrangement would minimize

$$L\pi(s, e) + s + c(s, e)$$

so that first order conditions are:

$$\begin{aligned} L|\pi_s| &= 1 + c_s \\ L|\pi_e| &= c_e. \end{aligned}$$

Intuitively, the customer should increase the level of care until on the margin the cost is

equal to marginal benefit in terms of reduced probability of loss. The bank should do the same, taking into account that increased security could increase costs to the customer as well.

In order to choose correctly, the customer must take into account the joint benefits from deterring malefactors. In practice, however, the customer is unlikely to bear the entirety of these costs (both because of legal protections and because of the fact that some of the costs to the bank of a theft will be hard to verify).²⁶ Let us suppose that the losses are divided between the bank and the customer: $L^b + L^c = L$. If the customer chooses e on his own, then the customer sets

$$L^c |\pi_e| = c_e,$$

choosing too little care.

If it were behaving non cooperatively, the bank would also be expected to underprovide security, since it too does not face the full cost. But it is more likely that, with the bank moving first, it is able to establish terms of agreement and fees to help internalize the full costs. If in doing so it could impose the full costs of a loss on the customer, it would be possible for the bank to induce the customer to take the proper care, and in anticipation of that, for the bank to take the efficient level of care as well. Since we assume instead that the customer's losses are limited to L^c , the bank faces a simple moral hazard problem with respect to the customer: it cannot control the customer's actions or condition payments on them, and it knows that the customer will underprovide care. Its only option is to adjust its own level of security accordingly:

$$\min_s L\pi(s, e) + s + c(s, e)$$

s.t.

$$e \in \arg \min L^c \pi(s, e) + c(s, e).$$

The envelope principle implies that for small reductions in L^c away from L , the choice of s remains unchanged to first order and e decreases. The response becomes

²⁶For example, in the U.S., Regulation E places stiff limitations on the monetary losses to a bank account holder if unauthorized withdrawals occur (see the webpage of the Consumer Financial Protection Bureau, <https://www.consumerfinance.gov/policy-compliance/rulemaking/regulations/1005/6/>). It is possible, indeed likely, that the bank will impose a limited amount of non-pecuniary costs on the consumer in the event of a loss, through, for example, bureaucratic hurdles (filling forms, waiting periods for refunds, and the like). But to the extent that these are dissipative costs, they will not be used to reach the full first best outcome, and the results will be similar to those described in the text.

more complicated for larger moves of L^c . The customer's level of care is always smaller, but depending on degrees of substitutability the bank's optimum s can either increase, in order to make up for the customer's reduction, or decrease, in order to entice the customer to increase the level of care. However the following example will be relevant for the subsequent analysis.

5.1.1 Example

Let π be additively separable in s and e

$$\pi(s, e) = \psi(s) + \phi(e)$$

and let the convex function $c(\cdot, \cdot)$ take the form

$$\begin{aligned} c(s, e) &= C(e) \text{ if } e \geq s \\ &= \infty \text{ otherwise.} \end{aligned}$$

In other words, by choosing s , the bank forces the customer to take effort e at least equal to s .

Define s^* and e^* by

$$\begin{aligned} L\phi'(e^*) + C'(e^*) &= 0 \\ L\psi'(s^*) + 1 &= 0. \end{aligned}$$

If $e^* > s^*$ then the first-best outcome is (s^*, e^*) . Otherwise it is (z^*, z^*) , where

$$L\psi'(z^*) + L\phi'(z^*) + 1 + C'(z^*) = 0.$$

In the second best, the customer chooses $e = \min\{s, \hat{e}\}$, where \hat{e} is defined by

$$L^c\phi'(\hat{e}) + C'(\hat{e}) = 0.$$

Thus, the second-best (s, e) is found by comparing the social welfare of the two candidates: (s^*, \hat{e}) , or (z^*, z^*) . Three outcomes are possible:

1. The first best equals the second best (in which case both are (z^*, z^*)).
2. The first best has (s^*, e^*) and the second best has (s^*, \hat{e}) where $\hat{e} < e^*$.

3. The first best has (s^*, e^*) and the second best has (z^*, z^*) , where $z^* < e^*$ and $z^* > s^*$.

In short, in this example, moving from first to second best, customer effort e either remains constant or decreases and bank security choice s either remains constant or increases.

5.2 Competing Banks

We now consider the situation where a customer can have one bank account at each of two banks $k = 1, 2$. When there is no interdependence between the customer's accounts, the analysis from the previous section applies to each account separately. In reality, however, there are significant sources of interaction. When banks offer the customer accounts in a non-cooperative game, then the externalities can lead to inefficient outcomes. In our examples below, the banks choose protocols that are too lenient.²⁷

The game we will consider takes the following form:²⁸ Each bank offers one account, choosing a level of security s_k and a price f_k . The customer chooses whether to accept each offer and then chooses an effort level e . We consider subgame perfect Nash equilibria of this game.

If bank k 's offer is accepted, its expected profit is

$$f_k - L^b \pi^k(\cdot) - s_k,$$

otherwise it is zero. The probability of loss $\pi^k(\cdot)$ can in general depend on levels of security and effort; we will consider specific cases. Let the gross value to the customer of a single account be V , and the gross value of adding a second account be $\Delta V \leq V$. The total cost of care to the customer $c(\cdot)$ will depend in general on the number of accounts accepted, the security measures provided for those accounts and the customer's chosen effort. The customer's utility is

$$V + \Delta V - f_1 - f_2 - L^c \pi^1(\cdot) - L^c \pi^2(\cdot) - c(\cdot)$$

²⁷Using our earlier terminology, the first of the two examples will be relevant when considering defense against hackers, the second when considering defense against thieves.

²⁸The setting is a simple example of a game of non-exclusive contracting under moral hazard. For further analysis of such environments see, for example, [Kahn and Mookherjee \(1995\)](#) or [Attar and Chassagnon \(2009\)](#). Note that in the absence of the moral hazard element, our game would be a standard Bertrand duopoly game.

if the customer accepts both accounts, and

$$V - f_k - L^c \pi^k(\cdot) - c(\cdot)$$

if the customer only accepts bank k 's account, and 0 otherwise.

Again, specific cases will be considered below.

5.2.1 Password Reuse

An important form of interdependence arises through the potential for password reuse. Customers find it easier to remember one password rather than separate passwords for each account. In many respects, this is the analogue of a decision to consolidate two accounts into a single, larger account. Doing so affects both the costs to the participants and the potential loss.

Suppose that customer effort is identified with password length (or more generally, password quality). Then a customer left on his own will equate the marginal cost of the password length with the benefit gained from protecting *both* accounts:

$$L^c |\pi_e^1 + \pi_e^2| = c_e,$$

As before, if $L^c = L$ the customer chooses e efficiently, and if $L^c < L$ the customer underinvests in effort.

When the bank dealt with the customer in isolation, as noted in the previous section, the bank might respond to the underinvestment by toughening its own protocols. However, when passwords are reused, the bank's tendency to do this is dampened by the fact that part of the benefit now accrues to the other bank.

Suppose then that each bank sets a minimum length required for passwords, s_k , but then allows customers to choose their own password meeting that requirement. When customer effort is identified with password length, this is equivalent to the customer bearing a total cost of $C(e)$, subject to

$$e \geq \max\{s_1, s_2\}. \tag{7}$$

Effectively, then, this is an extension of the example from the previous section; with restriction (7), the total expected social cost if both accounts are opened is $W(s_1, s_2, e)$, where

$$W(s_1, s_2, e) = L\pi^1(s_1, e) + s_1 + L\pi^2(s_2, e) + s_2 + C(e).$$

We will assume that ΔV is large enough that it is always efficient to open two accounts.

The second-best allocation (s_1, s_2, e) minimizes this social cost subject to the incentive condition

$$e \in \arg \min_{e \geq s_1, e \geq s_2} L^c \pi^1(s_1, e) + L^c \pi^2(s_2, e) + C(e).$$

Assuming again that each π^k is additively separable in its arguments,

$$\pi^k(s_k, e) = \psi_k(s) + \phi_k(e),$$

the consumer's incentive problem is solved by $e = \max\{s_1, s_2, \hat{e}\}$, where

$$L^c \phi_1'(\hat{e}) + L^c \phi_2'(\hat{e}) + C'(\hat{e}) = 0.$$

Because of the non-convexity of the problem, calculating the second-best allocation is somewhat complex. Description of the solution and conditions for the various cases are presented in Appendix B.1. Depending on parameter values, the second-best allocation may have both banks set minimum password requirements below \hat{e} , the level the customer would choose on his own. In this case, the second best is equal to $(s_1^*, s_2^*; \hat{e})$, where s_k^* satisfies the first-best condition

$$L\psi_k'(s_k^*) + 1 = 0.$$

As a consequence in this case, the banks' strategic actions are uncoupled, and the second best can be achieved by non-cooperative play by the banks.

Another possibility is that the standards of the stricter bank (say, bank 1) are higher than what the customer would independently choose. In this case, the second best takes the form $(\tilde{s}_1, s_2^*; \tilde{s}_1)$, where \tilde{s}_1 satisfies the following condition:

$$L\psi_k'(\tilde{s}_1) + L\phi_1'(\tilde{s}_1) + L\phi_2'(\tilde{s}_1) + C'(\tilde{s}_1) + 1 = 0,$$

and as a result, the second-best outcome cannot be achieved as an equilibrium outcome. To see this, suppose on the contrary that the equilibrium strategy choices for the banks were $(\tilde{s}_1, f_1), (s_2^*, f_2)$ and that the customer responded by accepting both offers and (perforce) choosing $e = \tilde{s}_1$. Since the initial position forces the customer to choose a level of care exceeding what he would choose on his own, slightly relaxing in the bank's requirement to $\tilde{s}_1 - \epsilon$ allows the customer to reduce effort to $\tilde{s}_1 - \epsilon$. From the definition of \tilde{s}_1 ,

we know that

$$L\psi'_1(\tilde{s}_1) + L\phi'_1(\tilde{s}_1) + C'(\tilde{s}_1) + 1 > 0.$$

Thus, the relaxed requirement, increases the joint benefit of bank 1 and the customer and forms the basis of a deviation. If the customer is currently receiving some surplus from the arrangement with bank 2, then the customer continues to accept the second account in the deviation. If the customer is receiving no surplus from bank 2, then the customer simply drops this account as part of the deviation. Either way, there exists a small adjustment in fees such that the deviation results in a gain for both bank 1 and the customer at the expense of bank 2. The equilibrium generally has bank 2 choosing a lower level of care than second best. Details for a particular class of cases are provided in Appendix B.2.

Note that this form of the costs encourages convergence between standards for the two banks: the bank with lower standards is not deterred from raising them by fear of imposing costs on the customer, while the bank with higher standards faces pressure to lower them.²⁹

In this example, if each bank could impose full costs of a loss on the customer (that is, if L equaled L^c), there would be no need for banks to discourage any cost-saving activities by customers. Customer choices would be efficient, and given this, banks would fully internalize costs to choose efficient standards.

5.2.2 Password Aggregation Programs

If a bank wants to forestall password reuse (or use of too-simple passwords), it might issue randomly generated passwords to its customers. However, such a strategy will be ineffective in the presence of *password aggregation programs*, which reduce the cost to customers by storing passwords for them and allowing them to use a master password to access all their accounts. Like the case of password reuse, this is the analogue of a decision to consolidate two accounts into a single, larger account. However, the arms race doesn't stop there: banks can and do choose a variety of other security protocols, some of which are designed to thwart customers' attempts to bypass security requirements. So, for example, the question of password aggregation becomes more complex when competing banks consider choosing different security protocols (for example, two-factor

²⁹The resultant externalities would be a reason for banks to internalize these effects by having customers have both accounts with the same bank. In a more general framework, the feasibility of this will depend on banks' abilities to offer a variety of accounts with specialized purposes and whether there will be a temptation for an account holder to open additional accounts with another bank.

authentication), which in turn affect the usefulness of aggregation programs.

To examine these possibilities, we further simplify the model by limiting the bank to a choice between two levels of service: $s_k \in \{u, p\}$ (for “unprotected” and “protected”), associated with different probabilities of loss to the bank. For protected service, the probability of the loss is π_p . For unprotected service, the probability depends on whether the customer is also accepting unprotected service from the other bank. If the customer uses only one bank account, then the probability is π_u . If the customer is using two unprotected accounts, then each has the probability of loss π_{uu} . We assume

$$\pi_{uu} > \pi_u > \pi_p,$$

which has the interpretation that use of two unprotected accounts entails the possibility of a thief discovering the password.³⁰

The cost to the customer also depends on the security chosen by the banks. The protected service imposes a cost c_p on the customer, while the unprotected service imposes costs c_u or c_{uu} depending on whether the customer is also obtaining unprotected service from the other bank, where

$$c_{uu} < c_u < c_p.$$

This setup reflects an environment where a more secure protocol is burdensome to the bank and customer. However, the benefit of the more secure protocol is that it limits the risk of loss from the choice of the customer with respect to the customer’s second account or the choice of the other bank in terms of which level of service to offer. An interpretation of this environment is where the protected level of service precludes password aggregation programs through protocols like two-factor authentication. The protocol of the unprotected level of service, on the contrary, allows password aggregation programs, thus increasing the risk of loss if both accounts are unprotected.³¹

³⁰For simplicity, we make these probabilities symmetric; effectively, we are assuming that the order of access of the accounts is random. Considering the accounts in a hierarchy, so that likelihood of contagion is no longer symmetric, will be a natural direction for further work.

³¹However, we shouldn’t complacently assume that techniques like two-factor authentication are the finish line in the arms race. Technologies will continue to develop to enable customers to evade the new restrictions. We might postulate the invention of hardware solutions aggregating the various dongles and physical keys required by different banks for two-factor authentication. For a prediction in science fiction more than a quarter century ago about the inevitability of such developments, we turn to Douglas Adams:

...It was an Ident-I-Eeze, and was a very naughty and silly thing for Harl to have lying around in his wallet, though it was perfectly understandable. There were so many different ways in which you were required to provide absolute proof of your identity these days that

Without loss of generality we set $L = 1$ and $L^c = 0$. (Any direct costs to the customer from theft could be reinterpreted as part of c). We place the following restrictions on parameters:

$$V = \Delta V > c_p + \pi_p > c_u + \pi_u. \quad (8)$$

That is, the second account provides the same benefit to the customer as the first account, and if only one service is adopted by the customer, then the unprotected level of service is more efficient than the protected level. Further:

$$c_{uu} + \pi_{uu} > c_p + \pi_p. \quad (9)$$

That is, if the customer adopts the same type of service from each bank, then the protected service is better than unprotected service. An implication of the two sets of restrictions is that the most efficient arrangement is to obtain protected service from one bank and unprotected service from the other bank. (In fact, this weaker requirement is really all we need for the discussion.)

But is such an arrangement a Nash equilibrium?

Theorem. *If*

$$c_u + c_p + \pi_p > 2c_{uu} + \pi_{uu}$$

then there is no fully efficient Nash equilibrium of the game. In any equilibrium, contagion occurs with positive probability.

Proof. If bank 1 offers the protected service with probability 1, then the strictly dominant strategy for bank 2 is to offer the unprotected service at the price $f_2 = V - c_u$, giving it a profit of $V - c_u - \pi_u$.

If bank 2 offers the unprotected service at this price, then the candidates for best response by 1 are to offer the protected service at a price $V - c_p$, or to offer the unprotected

life could easily become extremely tiresome . . . Just look at cash machines, for instance. Queues of people standing around waiting to have their fingerprints read, their retinas scanned, bits of skin scraped from the nape of the neck and undergoing instant (or nearly instant — a good six or seven seconds in tedious reality) genetic analysis, then having to answer trick questions about members of their family they didn't even remember they had and about their recorded preferences for tablecloth colors. And that was just to get a bit of spare cash for the weekend. If you were trying to raise a loan for a jetcar, sign a missile treaty or pay an entire restaurant bill, things could get really trying.

Hence the Ident-I-Eeze. This encoded every single piece of information about you, your body and your life into one all-purpose machine-readable card that you could then carry around in your wallet, and it therefore represented technology's greatest triumph to date over both itself and plain common sense. (Adams (1992) *Mostly Harmless*, p. 72.)

service at the maximum price possible:

$$2V - 2c_{uu} - f_2.$$

Thus, if

$$V + c_u - 2c_{uu} - \pi_{uu} > V - c_p - \pi_p,$$

there is no fully efficient Nash equilibrium for the duopoly game. □

The condition of the theorem is consistent with conditions (8 - 9) above. For example, the following table of parameters satisfies all inequalities (for any V greater than 9):

	π_{uu}	π_u	π_p
π	9	4	3
c	1	4	6
$\pi + c$	10	8	9

The crucial feature of the condition is that contagion be important, that is, that π_{uu} is significantly larger than π_u . If the inequality of the theorem is reversed, there is a fully efficient Nash equilibrium.

6 Policy Discussion

What is the role of a welfare maximizing regulator with respect to the security protocols that banks establish? This depends on which protocols each bank is offering to their respective customers and whether password aggregation programs or other workarounds are available. The most common protocols are: i) requiring customers to have strong passwords (for example, by establishing a particular format, length, or use of special characters); ii) requiring customer to use random passwords generated by the bank; iii) requiring customers to change their passwords at certain intervals; iv) two-factor authentication (requiring a special dongle or app to generate a single-use code, or a phone or verified email address to receive those codes); and v) protocols (invisible to the customer) that interfere with password aggregation programs.³²

We analyze the effects of the frequency of change of passwords and two-factor au-

³²An example of interference is a bank website design that sabotages the ability of browsers to automatically locate the input box for passwords. Another less common protocol is the lookup in public websites for “pwned” passwords in order to alert customers or to reject such passwords outright.

thentication when there are two banks and password aggregation might be available to customers from a third party. The effects of the policies can be examined through their changes in the probability of theft and cost to customers.

6.1 Frequency of Change of Passwords

The first intervention is the requirement of the regulator that banks demand customers to change their passwords with certain frequency. The probability of theft falls proportionally with the frequency of password change mandated by the regulator because the likelihood of the thief getting hold of a useful password falls by half if the password is changed twice as often. Let e_{b_1} here be the number times in a given period of time that a password has to be changed, as mandated by the regulator. Then we can write the probability of theft as:

$$\pi(e_{b_1}, e_c) = e_{b_1}^{-1} \tilde{\pi}(e_c),$$

and $\partial \tilde{\pi} / \partial e_c < 0$. Further assume that the cost to customers is proportional to the times they are required to change the password:

$$c(e_{b_1}, e_c) = e_{b_1}^{\beta} \tilde{c}(e_c),$$

with $\beta \geq 1$ and $\partial \tilde{c} / \partial e_c > 0$. Here, the interpretation of the effort exerted by customers is, for example, the need to write down and protect the list of passwords.

The simplest case is when both banks require random passwords and no password aggregation is available. In this case, the requirement of the regulator will not change the probability of theft of the other bank because the thief does not know the timing of the passwords changes. If the cost to the customer increases with the frequency of change linearly, the level of care that the customer exerts will not change, since every password is equally difficult as it is randomly generated by the bank.

Now consider the opposite extreme case, in which customers, to reduce their burden, reuse passwords for both banks. In this case, the policy of demanding password changes on one bank will have the positive effect of reducing the probability of theft on the second one, as the customer will change both at the same time.

Consider further the situation when there are password aggregation programs. The probability of theft would not be affected because, in spite of the customers changing the passwords used with the banks, the risk of theft of password aggregator would remain the same. This is because the user or program cannot be compelled to comply with

the policy. Likewise, the policy would not have a significant effect on the cost to the customer, as the program would handle most of the additional costs associated with the changes.

6.2 Two-Factor Authentication

Today, many traditional banks, wallets, and exchanges offer two-factor authentication as an option to increase security. The intervention in this case would be to require users to use it. Two-factor authentication methods require a second, additional method of identification independent of the first. The bank only releases the funds if both factors are correct at the time of the withdrawal request. Then the probability of theft is:

$$\pi(e_b^{1F}, e_b^{2F}, e_c^{1F}, e_c^{2F}) = \pi_1(e_b^{1F}, e_c^{1F})\pi_2(e_b^{2F}, e_c^{2F}),$$

where, as before, π is decreasing in all of its arguments. The efforts exerted for each factor are not independent, as the customer will equate both at the margin. But in general, for the same level of total effort ($e_c = e_c^{1F} + e_c^{2F}$), this protocol would result unambiguously in a lower or equal probability of theft.

The cost to customers is in general additive in each factor of authentication and convex in the effort to handle each:

$$c(e_b^{1F}, e_b^{2F}, e_c^{1F}, e_c^{2F}) = c_1(e_b^{1F}, e_c^{1F}) + c_2(e_b^{2F}, e_c^{2F}).$$

Because the probability of theft decreases more than the increase in the cost to the customer, two-factor authentication, provided that it is feasible, will make customers better off with a combination of reduced expected theft and a lower optimal level of effort. The reason why banks might not require customers to always use it might be because for customers with small balances, the reduction in expected theft will not compensate for the additional effort. Banks and exchanges might have difficulty ex-ante determining which customers should be required to use two-factor authentication. A regulator would face the same issues, as balances could vary greatly over time.

7 Concluding Remarks

We have studied the trade-off between the safety and convenience of storing balances in anonymous addresses. This type of aggregation is the foundation of all private digital

currencies, such as Bitcoin. Security risks arise from hackers, focusing on banks and exchanges, and from thieves, attempting to steal private keys and account passwords from customers. The extent of loss depends the technological choices of banks and the effort of customers, giving rise to a moral hazard problem. With shared liability we find that in general customers will take too little care. Even when managing their balances individually and facing the entire risk of loss, customers will find some level of aggregation desirable and so will prefer to use wallets, reuse addresses, and rely on password aggregation program.

Our findings have implications for the design of central bank digital currencies (CBDC) and their ecosystem. If the central bank can establish liability rules for loss of digital currency similar to those for bank notes, customers will have the incentives to exert the appropriate level of care. Enforcing these rules, however, might not be as straightforward. Moreover, determining the liability in case of loss would be even more complicated if the design of the CBDC allows any individual or firm to hold the digital tokens. If this is possible, we would expect customers to aggregate balances in accounts held at intermediaries, such as exchanges. This would give rise to deposits in unregulated entities, which might be out of reach of domestic authorities. Designing a CBDC that is universally accessible but cannot be held by certain firms is a technological challenge.

Our framework applies to any digital asset that functions as a bearer instrument. As institutional investors consider holding digital currencies, it would be relevant to analyze other protocols like multisignature and key sharding. Another avenue is to analyze the incentives of large technology companies considering issuing digital currencies, which could provide custody services and earn revenue from the online activity of their customers. We leave this for future papers.

Bibliography

- Adams, D. (1992). *Mostly Harmless (Hitchhiker's Guide to the Galaxy, Book 5)*. Harmony Books.
- Anderson, R. and T. Moore (2007). Information security economics – and beyond. In A. Menezes (Ed.), *Advances in Cryptology - CRYPTO 2007*, Berlin, Heidelberg, pp. 68–91. Springer Berlin Heidelberg.
- Anderson, R. and T. Moore (2009). Information security: where computer science, economics and psychology meet. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 367(1898), 2717–2727.
- Attar, A. and A. Chassagnon (2009). On moral hazard and nonexclusive contracts. *Journal of Mathematical Economics* 45(9–10), 511–525.
- Barontini, C. and H. Holden (2018). Proceeding with caution - a survey on central bank digital currency. *BIS Papers* (101).
- Baumol, W. J. (1952). The transactions demand for cash: An inventory theoretic approach. *The Quarterly Journal of Economics*, 545–556.
- Bonneau, J., S. Preibusch, and R. Anderson (2012). A birthday present every eleven wallets? the security of customer-chosen banking pins. In A. Keromytis (Ed.), *Financial Cryptography and Data Security, FC 2012, Lecture Notes in Computer Science, vol 7397*. Springer, Berlin, Heidelberg.
- Bordo, M. D. and A. T. Levin (2017). Central bank digital currency and the future of monetary policy. NBER Working Paper Number 23711, <https://www.nber.org/papers/w23711>.
- Briner, M. G. (1983). Tax Equity and Fiscal Responsibility Act of 1982. *Akron Tax J.* 1, 29.
- Brunnermeier, M. K., H. James, and J.-P. Landau (2019). The digitalization of money. https://scholar.princeton.edu/sites/default/files/markus/files/02c_digitalmoney.pdf.
- Chiu, J. and T. Koepl (2017). The economics of cryptocurrencies - bitcoin and beyond. Technical report.

- Ciphertrace (2019). Q4 2018 cryptocurrency anti-money laundering. <https://ciphertrace.com/crypto-aml-report-2018q4/>.
- Clemens, S. (1894). *The Tragedy of Pudd'nhead Wilson*. Charles Webster and Co.
- Davoodalhosseini, M. (2018). Central bank digital currency and monetary policy. *Bank of Canada Working Papers* (2018-36).
- Diffie, W. and M. Hellman (1976, November). New directions in cryptography. *IEEE Transactions on Information Theory* 22(6), 644–654.
- Duffie, D. (2019). Digital currencies and fast payment systems: Disruption is coming. presentation to the Asian Monetary Policy Forum, May, 2019.
- EY (2017). EY research: initial coin offerings (ICOs). [https://www.ey.com/Publication/vwLUAssets/ey-research-initial-coin-offerings-icos/\\$File/ey-research-initial-coin-offerings-icos.pdf](https://www.ey.com/Publication/vwLUAssets/ey-research-initial-coin-offerings-icos/$File/ey-research-initial-coin-offerings-icos.pdf).
- Green, E. J. (2008a). The role of the central bank in payment systems. In S. Millard, A. Haldane, and V. Saporta (Eds.), *The Future of Payment Systems*, pp. 45–56. Routledge.
- Green, E. J. (2008b). Some challenges for research in payments. In S. Millard, A. Haldane, and V. Saporta (Eds.), *The Future of Payment Systems*, pp. 57–67. Routledge.
- He, D., K. Habermeier, R. Leckow, V. Haksar, Y. Almeida, M. Kashima, N. Kyriakos-Saad, H. Oura, T. S. Sedik, N. Stetsenko, and C. Verdugo-Yepes (2016). Virtual currencies and beyond: Initial considerations. *IMF Staff Discussion Notes* (SDN/16/03).
- Herley, C. (2009). So long, and no thanks for the externalities: The rational rejection of security advice by users. New Security Paradigms Workshop.
- Hirshleifer, J. (1983). From weakest-link to best-shot: the voluntary provision of public goods. *Public Choice* 41, 371–386.
- Independent Security Evaluators (2019, April). Ethercombing: Finding secrets in popular places. <https://www.securityevaluators.com/casestudies/ethercombing/>.
- Jain, S., E. Felten, and S. Goldfeder (2018, 08). Determining an optimal threshold on the online reserves of a bitcoin exchange. *Journal of Cybersecurity* 4(1), 1–12. ty003.

- Kahn, C. M. (2016). How are payment accounts special? *Payments Innovation Symposium Federal Reserve Bank of Chicago*.
- Kahn, C. M., J. McAndrews, and W. Roberds (2005). Money is privacy. *International Economic Review* 46(2), 377–399.
- Kahn, C. M. and D. Mookherjee (1995). Market failure with moral hazard and side trading. *Journal of Public Economics* 58(2), 159–184.
- Kahn, C. M., F. Rivadeneyra, and T.-N. Wong (2018). Should the central bank issue e-money? *Bank of Canada Staff Working Paper* (2018-58).
- Kahn, C. M. and W. Roberds (2009). Why pay? an introduction to payments economics. *Journal of Financial Intermediation* 18(1), 1–23.
- Lacker, J. (2008). Payment economics and the role of central banks. In S. Millard, A. Haldane, and V. Saporta (Eds.), *The Future of Payment Systems*, pp. 68–72. Routledge.
- Laszka, A., M. Felegyhazi, and L. Buttyan (2014, August). A survey of interdependent information security games. *ACM Comput. Surv.* 47(2), Article 23.
- Mancini-Griffoli, T., M. S. Martinez Peria, I. Agur, A. Ari, J. Kiff, A. Popescu, and C. Rochon (2018). Casting light on central bank digital currencies. *Staff Discussion Notes* (18/08).
- Massey, J. L. (1994). Guessing and entropy. In *Proceedings of 1994 IEEE International Symposium on Information Theory*, pp. 204.
- Merkle, R. C. (1978, 04). Secure communication over insecure channels. *Commun. ACM* 21, 294–299.
- Merkle, R. C. (1980, April). Protocols for public key cryptosystems. In *1980 IEEE Symposium on Security and Privacy*, pp. 122–122.
- Milne, A. K. L. (2018). Argument by false analogy: the mistaken classification of bitcoin as token money. <https://ssrn.com/abstract=3290325>.
- Moore, T. (2010). The economics of cybersecurity: Principles and policy options. *International Journal of Critical Infrastructure Protection* 3(3-4), 103–117.

- Stecklow, S., A. Harney, A. Irrera, and J. Kelly (2017). Virtual mayhem: chaos and hackers stalk investors on cryptocurrency exchanges, Reuters, Sept. 29, 2017. <https://www.reuters.com/investigates/special-report/bitcoin-exchanges-risks/>.
- Tobin, J. (1956). The interest-elasticity of transactions demand for cash. *The Review of Economics and Statistics* 38(3), 241–247.
- Varian, H. (2004). System reliability and free riding. In L. Camp and S. Lewis (Eds.), *Economics of information security*, Volume 12 of *Advances in Information Security*, pp. 1–15. Springer.
- Vasek, M., J. Bonneau, R. Castellucci, C. Keith, and T. Moore (2017). The bitcoin brain drain: examining the use and abuse of bitcoin brain wallets. In J. Grossklags and B. Preneel (Eds.), *Financial Cryptography and Data Security. FC 2016. Lecture Notes in Computer Science, vol 9603*, pp. 609–618. Springer, Heidelberg.

A Appendix

A.1 Expected Loss from Theft

Here we derive the expected value of theft in the two-account version of the model when the number of transfers in a period is n . We have assumed implicitly that $n \leq T$ and that at the moment of transfer, the amount I/T is used for a purchase. After one transfer of I/n out of the investment account, the remaining balance is $I - I/n$. Every time this transfer occurs, the likelihood of theft is π_i . Therefore, the probability of theft after n transfers is:

$$\begin{aligned} & \pi_i \sum_{k=1}^n \left(I - k \frac{I}{n} \right) \\ &= \frac{I\pi_i n}{2} \left(1 - \frac{1}{n} \right). \end{aligned} \tag{10}$$

Similar calculations apply for the transaction account. After the transfer into this account and the first payment, the balance of the account is $I/n - I/T$. After the next payment, $I/n - 2I/T$ remains, and so forth. Every time this purchase occurs, the likelihood of theft is π_t . Therefore, the expected value of theft after $T/n - 1$ payments from every occasion a transfer is made:

$$\begin{aligned} & \pi_t \sum_{k=1}^{T/n-1} \left(\left(\frac{I}{n} - \frac{kI}{T} \right) - k \frac{I}{T} \right) \\ &= \frac{I\pi_t}{n} (T - n) \left(\frac{1}{2n} - \frac{1}{T} \right), \end{aligned}$$

which occurs n times when the transfer from the investment account is made; therefore the probability of theft from using the transaction account is:

$$= \frac{I\pi_t}{2} (T - n) \left(\frac{1}{n} - \frac{2}{T} \right). \tag{11}$$

Adding (10) and (11), the total expected lost from theft adding from both accounts is

$$\frac{I}{2} \left(\left(1 - \frac{1}{n} \right) \pi_i n + \left(\frac{1}{n} - \frac{2}{T} \right) \pi_t (T - n) \right), \tag{12}$$

as in the text.

A.2 General Hierarchy of Accounts

Here we describe a version of the model with a hierarchy of accounts. For this section, we set consumer costs to $c(n, Q) = \alpha n + CQ$ and bank costs $K(\cdot)$ to zero without significant effect on the results.

Suppose that the depositor maintains a hierarchy of n accounts. The transactions amounts are withdrawn in N chunks, each of amount I/N , and each withdrawal leads to a probability π of a disclosure of the password to a thief. Let m_1 denote the number of withdrawals from the smallest account/ first level of the hierarchy. The expected amount lost to theft on the first level is

$$L_1 = \pi \sum_{i=1}^{m_1} \left(m_1 \frac{I}{N} - i \frac{I}{N} \right) = \pi \frac{I}{N} \frac{m_1(m_1 - 1)}{2}.$$

That is to say, initially the depositor transfers to the first account $m_1 \frac{I}{n}$ units of money from the second, bigger, account. After m_1 withdrawals, nothing is left, and the depositor again transfers $m_1 \frac{I}{n}$ units of money from the second account. Let m_2 denote the number of transfer from the second account to the first account. The expected amount lost to theft on the second hierarchy is

$$L_2 = m_2 L_1 + \pi \sum_{i=1}^{m_2} \left(m_2 \frac{m_1 I}{N} - i \frac{m_1 I}{N} \right) = \pi \frac{m_1 I}{N} \frac{m_2(m_1 + m_2 - 2)}{2}.$$

The expected loss on the second hierarchy consists of the total loss from using the first account, and the loss to theft from making the m_2 transfers to the first account. In general the expected amount lost to theft from all n accounts is

$$L_n = (\prod_{i=1}^n m_i) \frac{\pi}{2} \frac{I}{N} \left(\sum_{i=1}^n m_i - n \right).$$

Thus, the cost minimization problem is

$$\min_{n, m_i} (\prod_{i=1}^n m_i) \frac{\pi}{2} \frac{I}{N} \left(\sum_{i=1}^n m_i - n \right) \text{ s.t. } I = \frac{I}{N} \prod_{i=1}^n m_i,$$

where the budget constraint requires that the total amount deposited is equal to the total

amount withdrawn. The first order condition with respect to m_i is

$$\pi = \frac{\lambda}{m_i},$$

where λ is the Langrange multiplier. Hence we have

$$m_i = N^{1/n}.$$

Thus, the optimal number of withdrawals from each account is the same and is increasing in the number of accounts if and only if $I \geq N$. By substituting the above FOCs, the agent's problem becomes

$$\min_n \alpha n + Cqn + \frac{\pi}{2} I (nN^{1/n} - n).$$

The first order condition with respect to n is

$$1 - \frac{2(\alpha + Cq)}{\pi I} = N^{1/n} (1 - \log N^{1/n}),$$

$$\alpha + Cq = \frac{I\pi}{2} \frac{N}{n^2}$$

where the right side is strictly decreasing in $N^{1/n}$ for any $N > 1$. As is intuitive, the number of accounts increases with the number of withdrawals to be made and with the probability of success of theft. It decreases with the costs associated with an account, including the necessary length of a password for each account (again, large banks will have larger accounts). It increases with customer's total income.

Which type of account management is better at minimizing the loss to theft? Define $H(n) \equiv \alpha n + kqn + \frac{I\pi}{2} n (N^{1/n} - 1)$ and $S(n) \equiv \alpha n + kqn + \frac{I\pi}{2} (\frac{N}{n} - 1)$. Hierarchy management of accounts has lower expected loss to theft than serial management of accounts if and only if

$$\min_n H(n) \leq \min_n S(n).$$

Define $\Delta(n) \equiv H(n) - S(n)$. Notice that $\Delta(n) < 0$ for all $n \geq 1$ when N is sufficiently large. Thus, the hierarchy management has lower loss than the serial management if the withdrawal is sufficiently frequent.

Notice that $\Delta(1) = 0$, $\Delta(\infty) > 0$, $\Delta'(\infty) = 0$ and

$$\begin{aligned}\Delta'(n) &= \frac{I\pi}{2} \left[N^{1/n} (1 - \log N^{1/n}) - 1 + \frac{N}{n^2} \right], \\ \Delta''(n) &= \frac{I\pi}{2n^3} [N^{1/n} (\log N)^2 - 2N].\end{aligned}$$

Since $(\log N)^2 < 2N$ for all $N \geq 1$, we have $\Delta''(\infty) < 0$. If $\log N \leq \sqrt{2}$, i.e., $N \leq 3.69$, then $\Delta''(n) < 0$ for all n and hence $\Delta(n) > 0$ for all $n \geq 1$. If $\log N > \sqrt{2}$, then $\Delta''(n)$ single-cross zero from above and $\Delta'(n)$ is inverted-U shape in n . Furthermore, if $\Delta'(1) = N(2 - \log N) - 1 \geq 0$ i.e., $N \leq 6.31$, then $\Delta(n) > 0$ for all $n \geq 1$. Thus, the serial management has lower loss than the hierarchy management if the withdrawal is sufficiently infrequent that $N \leq 6.31$.

A.3 Expected Loss from Theft with Alternative Timing of Withdrawal

If we remove the assumption that payment I/T is made immediately after the transfer is made to the transaction account, then the expected loss of theft calculation is identical to that in the n sequential accounts:

$$n\pi_t \sum_{k=1}^{(T/n)} \left(\frac{I}{n} - k \frac{I}{T} \right) = \frac{I\pi_t}{2} \left(\frac{T}{n} - 1 \right). \quad (13)$$

Total expected theft is simply:

$$\frac{I}{2} \left((n-1)\pi_i + \left(\frac{T}{n} - 1 \right) \pi_t \right), \quad (14)$$

and the solution to the problem is exact in this case:

$$n = \sqrt{\frac{T\pi_t}{\pi_i}}. \quad (15)$$

If costs d_i and d_t are proportional to I , total cost of account management including expected theft is:

$$\frac{I}{2} \left(\left(\left(1 - \frac{1}{n} \right) \pi_i + d_i \right) n + \left(\frac{T}{n} - 1 \right) \pi_t + d_t T \right), \quad (16)$$

and the solution is:

$$n = \sqrt{\frac{T\pi_t}{\pi_i + d_i}}. \quad (17)$$

In this version, the inconvenience cost of the transaction account does not enter the solution to the number of transfers because the number of payments T is fixed.

B Calculations for Password Reuse (Section 5.2.1)

B.1 Second Best

Recall

$$W(s_1, s_2, e) \equiv L\pi^1(s_1, e) + L\pi^2(s_2, e) + s_1 + s_2 + C(e),$$

where

$$\pi^k(s, e) = \max\{\psi_k(s) + \phi_k(e), 0\}$$

for $k = 1, 2$, and $C(e) = \infty$ if $e < \max\{s_1, s_2\}$.

The first-best problem is convex, and its first-order conditions are:

$$\begin{aligned} L\psi'(s_1) + 1 &\leq 0 \\ L\psi'(s_2) + 1 &\leq 0 \\ L\pi'(e) + L\pi'(e) + C'(e) &\geq 0 \\ L\pi'(e) + L\pi'(e) + C'(e) + L\psi'(s_1) + 1 + L\psi'(s_2) + 1 &= 0. \end{aligned}$$

Defining s_k^*, e^* by

$$\begin{aligned} L\psi'(s_k^*) + 1 &= 0 \\ L\pi'(e^*) + L\pi'(e^*) + C'(e^*) &= 0 \end{aligned}$$

we have that if $e^* > \max\{s_1^*, s_2^*\}$ then (s_1^*, s_2^*, e^*) is the first-best allocation.

The second-best problem is

$$\min_{(s_1, s_2, e) \in \mathbb{R}_+^3} W(s_1, s_2, e)$$

subject to

$$e \in \arg \min_{e \geq s_1, e \geq s_2} L^c\pi^1(s_1, e) + L^c\pi^2(s_2, e) + C(e).$$

Define \hat{e} by

$$L^c \phi'_1(\hat{e}) + L^c \phi'_2(\hat{e}) + C'(\hat{e}) = 0$$

(Note that $\hat{e} \leq e^*$.) The second-best restriction on e is equivalent to

$$e = \max\{s_1, s_2, \hat{e}\}$$

The set of second-best feasible triples (s_1, s_2, e) is *not* convex. However, it is made up of convex sub-problems and for very mild conditions solutions exist. If in the first-best allocation $e = \max\{s_1, s_2\}$, then the allocation is also second best. If $\hat{e} < e^*$, and $e^* > \max\{s_1^*, s_2^*\}$, then the first-best allocation is not second best.

For convenience define the functions

$$\begin{aligned} A^k(x) &= L\psi_k(x) + x \\ G(x) &= L\phi_1(x) + L\phi_2(x) + C(x). \end{aligned}$$

Then

$$W(s_1, s_2, e) = A^1(s_1) + A^2(s_2) + G(e)$$

and the second-best problem is to minimize

$$W(s_1, s_2, \max\{s_1, s_2, \hat{e}\}).$$

Assume the functions A^1, A^2, G are each strictly convex and have minima s_1^*, s_2^* , and e^* respectively. Let

$$\begin{aligned} \hat{s}_k &= \arg \min_{s \in \mathbb{R}_+} A^k(s) + G(s) \\ \hat{z} &= \arg \min_{s \in \mathbb{R}_+} A^1(s) + A^2(s) + G(s). \end{aligned}$$

Note that \hat{s}_k lies between s_k^* and e^* . Consequently, if $s_k^* > \hat{e}$ then $\hat{s}_k^* > \hat{e}$. Also note that \hat{z} lies between \hat{s}_k and s_{-k}^* , $k = 1, 2$, so that of the four values $(s_1^*, \hat{s}_1, s_2^*, \hat{s}_2)$, two are greater than \hat{z} and two less.

Define

$$\begin{aligned} R_1 &= \{(s_1, s_2) \in \mathbb{R}_+^2 \mid s_1 > s_2, s_1 > \hat{e}\} \\ R_2 &= \{(s_1, s_2) \in \mathbb{R}_+^2 \mid s_2 > s_1, s_2 > \hat{e}\} \\ R_3 &= \{(s_1, s_2) \in \mathbb{R}_+^2 \mid \hat{e} > s_1, \hat{e} > s_2\} \end{aligned}$$

so that

$$\bar{R}_1 \cup \bar{R}_2 \cup \bar{R}_3 = \mathbb{R}_+^2$$

and

$$\begin{aligned} & \min_{(s_1, s_2) \in \mathbb{R}_+^2} W(s_1, s_2, \max\{s_1, s_2, \hat{e}\}) \\ &= \min\{\min_{(s_1, s_2) \in \bar{R}_1} W(s_1, s_2, s_1), \min_{(s_1, s_2) \in \bar{R}_2} W(s_1, s_2, s_2), \min_{(s_1, s_2) \in \bar{R}_3} W(s_1, s_2, \hat{e})\} \end{aligned}$$

Thus, we can solve the sub-problem on each region and compare results to find the global minimum. The search is simplified by strict concavity in each region. For example, the first order conditions for region 1 are satisfied by $(s_1, s_2) = (\hat{s}_1, s_2^*)$. Thus, there is an interior minimum in region 1 (namely, (\hat{s}_1, s_2^*)) if and only if

$$(\hat{s}_1, s_2^*) \in R_1.$$

Similarly, the following conditions are necessary and sufficient for interior minima in regions R_2 and R_3 , respectively:

$$\begin{aligned} (s_1^*, \hat{s}_2) &\in R_2 \\ (s_1^*, s_2^*) &\in R_3. \end{aligned}$$

The search is further simplified by the following

Observation. *If there is an interior minimum in any of the regions, then the global minimum is an interior minimum.*

Proof. If each of the three regions has an interior minimum, then the conclusion follows immediately by continuity. If two of the three regions have interior minima, then two of the three sets of conditions for interior minima must hold, and therefore the minimum of the third region is on a boundary with one of the other two, and therefore dominated by that region's interior minimum. If exactly one of the regions has an interior minimum, then the only way that the global minimum can be on the boundary is for it to be on the boundary of the other two regions (and it cannot be (\hat{e}, \hat{e})). There are two cases to consider. First, suppose that region 3 has an interior minimum and there is a common boundary minimum for regions 1 and 2 of (\hat{z}, \hat{z}) . (See Figure 4). Since the minimum for region 3 is interior, s_1^* and s_2^* are both less than \hat{e} . But the first-order condition for s_2 is the same in region 1 and in region 3. That means that on any vertical line in region 1 the objective attains the minimum at $s_2 = s_2^*$, which in turn means that the objective




Figure 4: A single interior minimum, case 1

decreases as we move vertically down from the boundary in region 1 (and correspondingly as we move left from the boundary in region 2) contradicting the claim that there is a minimum on the boundary.

In the second case, suppose that the interior minimum is not in region 3 but in, for example, region 1, and there is a common boundary minimum for regions 2 and 3 (Figure 5). Such a minimum must be (s_1^*, \hat{e}) with $s_1^* < \hat{e}$. This means that $s_2^* \geq \hat{e}$; otherwise there would be an interior minimum in region 3. Since there is an interior minimum in region 1, $\hat{s}_1 > s_1^*$, and so it must be the case that $e^* > \hat{s}_1 > s_2^*$. But \hat{s}_2 lies between e^* and s_2^* ; thus $\hat{s}_1 > s_2^*$ and (s_1^*, \hat{s}_2) is interior to region 2, contradicting the assumption, and completing the proof. \square

As a result, we only have to search for a global minimum on the boundary of a region when none of the regions have interior minima. We have the following result:

Observation. *The global minimum is on the boundary of regions 1 and 2 if and only if*




Figure 5: A single interior minimum, case 2

the following four conditions hold

$$\begin{aligned}
 s_1^* &\geq \hat{e} \\
 s_2^* &\geq \hat{e} \\
 s_1^* &\geq \hat{s}_2 \\
 s_2^* &\geq \hat{s}_1,
 \end{aligned}$$

in which case the global minimum is $(\max\{\hat{e}, \hat{z}\}, \max\{\hat{e}, \hat{z}\})$. Otherwise, the global minimum is interior to one of the regions.

Proof. If any of the conditions is violated, there is an interior minimum. (If the first is violated, either there is a minimum in R_3 or $s_2^* > \hat{e}$, in which case $\hat{s}_2 > \hat{e}$, and there is a minimum in R_2 . If the second is violated, the argument is symmetric. If neither the first nor the second is violated, there is a minimum in R_2 if the third is violated, and a minimum in R_1 if the fourth is violated.)

Conversely, if the minimum for every region is on the boundary, there cannot be three of them, as it would violate transitivity. Therefore, the global minimum must be shared by at least two regions. First, consider the possibility that it lies on the border of R_3 but

not at the corner. Suppose, for example, it is (\hat{e}, s_1^*) , then since it is interior to neither of the two regions, $s_2^* \geq \hat{e} \geq \hat{s}_2$. But this means $e^* \leq \hat{s}_2 \leq \hat{e}$, which is a contradiction.

The only other possibility is that it lies on the border of R_1 and R_2 (Figure 6). In




Figure 6: A boundary minimum

the vicinity of such a point (z, z) , on the R_1 side of the boundary, increases in s_2 and decreases in s_1 must result in decreases in the costs. In other words, $(\hat{s}_1 < z < s_2^*)$ and correspondingly on the R_2 side it must be that $(\hat{s}_2 < z < s_1^*)$. Such a point would also be a local minimum on the ray $\{(z, z) | z \geq \hat{e}\}$, which establishes the final condition. Moving into the region R_3 cannot reduce costs, so it must be that s_1^* and s_2^* are greater than \hat{e} . \square

B.2 Equilibrium

Note: For expositional purposes, in this portion of the appendix we will assume that $L^c = 0$ and $\psi_k(\cdot) = 0$. The results continue to hold for positive L^c , but the notation becomes more complex. Additional considerations that arise in the general case will be pointed out in the footnotes. Reintroducing $\psi_k(\cdot)$ makes no substantive difference; the conditions described in the main text are adjusted accordingly.

Suppose bank k offers (s_k, f_k) . Let ι_k be the indicator variable denoting the customer's acceptance of an account with bank k . The customer chooses (ι_1, ι_2, e) to maximize

$$U(\iota_1, \iota_2, e) \equiv \iota_1 \iota_2 (\Delta V - V) - C(e) + \sum_{k=1,2} \iota_k (V - f_k)$$

subject to

$$e \geq \max\{\iota_1 s_1, \iota_2 s_2\}.$$

Since C is increasing, the constraint binds.³³ We denote the maximand to the consumer's problem by $\hat{U}((s_1, f_1), (s_2, f_2))$, a function of the banks' offers. The following table provides necessary conditions for (ι_1, ι_2) to be a maximizing choice:

If $(\iota_1, \iota_2) =$	then $\hat{U} =$
(0, 0)	0
(1, 0)	$V - C(s_1) - f_1$
(0, 1)	$V - C(s_2) - f_2$
(1, 1)	$V + \Delta V - \max\{C(s_1), C(s_2)\} - f_1 - f_2$

We denote the set of maximizers (ι_1, ι_2) by $I((s_1, f_1), (s_2, f_2))$.

Next consider bank k 's best response given the rival bank's strategy. For instance, bank 2's profits are

$$P_2(s_2, f_2, \iota_1, \iota_2) \equiv \iota_2 (f_2 - L\phi(\max\{s_2, \iota_1 s_1\}) - s_2).$$

Banks are interdependent because when the customer accepts bank 1's offer it can reduce the probability of a loss to bank 2.³⁴ Given bank 1's offer, bank 2 faces the following *second-best maximization problem*:

$$\max P_2(s_2, f_2, \iota_1, \iota_2)$$

subject to

$$(\iota_1, \iota_2) \in I((s_1, f_1), (s_2, f_2)).$$

The strategies $(s_1, f_1), (s_2, f_2)$ belong to a subgame perfect Nash equilibrium in pure

³³If customers bore part of the cost of account loss, then they would sometimes choose e greater than the minimum requirement from the bank.

³⁴In the case at hand, bank 2's probability of loss is reduced whenever bank 1's standards are higher. In the general case, interdependence can arise through a second channel: Even if bank 1's standards are not higher, the customer may increase effort in order to protect both accounts.

strategies if there exists (ι_1, ι_2) such that given bank 1's strategy, the pair $((s_2, f_2), (\iota_1, \iota_2))$ solves bank 2's second-best maximization problem and vice versa. Solving the two second-best problems simultaneously is not the same as finding the overall second-best outcome, since each bank in its problem ignores the portion of the costs of losses faced by the other bank.³⁵

Define \hat{s}_k by

$$L\phi'(\hat{s}_k) + C'(\hat{s}_k) + 1 = 0.$$

This is the first-order condition for the bank's second-best problem when the bank can offer an exclusive contract.

Theorem. *In a pure strategy equilibrium, either one bank chooses \hat{s}_k satisfying the condition and the other chooses security level 0, or else both choose 0.*

Proof. Without loss of generality suppose that in a pure strategy equilibrium $s_2 \leq s_1$. Then, if $s_1 > 0$, it must be the case that the customer is accepting bank 1's contract. (Otherwise, reducing s_1 would reduce the bank's losses.) If the customer is accepting the contract, then unless the condition is satisfied, there is an adjustment of the level of security and accompanying adjustment of the fee such that the customer continues to accept the same menu of contracts and bank 1's profits increase. Meanwhile, if bank 1's contract is accepted in equilibrium, bank 2 always improves its profits by reducing s_2 . \square

Depending on the difference between the value the customer attaches to the first and second account, there may be a unique set of prices offered in a pure strategy equilibrium, or there may be a continuum of possible prices for a given security level. If

$$V - \Delta V > C(0) \tag{18}$$

then in equilibrium:

$$f_1 = \Delta V + C(0) - C(s_1); \quad f_2 = \Delta V.$$

These prices make the customer indifferent about acquiring the second account; nonetheless the customer retains some surplus from the first account acquired. If condition (18)

³⁵In the general case, when the customer accepts both accounts, the solution to one bank's second-best problem is not the same as the solution would have been were the rival's offer not available, because the bank will take into account the effect on the customer of the portion of losses the customer faces from *both* bank accounts.

is reversed, then

$$f_1 \leq \Delta V + C(0) - C(s_1); \quad f_2 \leq \Delta V$$

and the sum of the two prices appropriates the entirety of the customer's surplus from the two accounts:

$$f_1 + f_2 = V + \Delta V - C(s - 1).$$

The following theorem provides sufficient conditions for a non-trivial pure strategy equilibrium to exist. Slight changes are needed if condition (18) is reversed.

Theorem. *Assume (18) holds. Then the following conditions are sufficient for the existence of a pure strategy equilibrium in which bank 1 offers \hat{s}_1 , bank 2 offers the an account with the minimum security 0, and both offers are accepted.*

$$\Delta V > L\phi_2(\hat{s}_1) \tag{19}$$

$$\Delta V > L\phi_1(\hat{s}_1) + \hat{s}_1 \tag{20}$$

$$C(\hat{s}_2) + L\phi_2(\hat{s}_2) + \hat{s}_2 \geq C(0) + L\phi_2(\hat{s}_1) + C(0) \tag{21}$$

Proof. Conditions (19)-(20) guarantee that both banks are making positive profits, so neither will choose to offer a contract that will be rejected by the customer. Holding fixed the offer by bank 2, the offer from bank 1 maximizes the joint payoff to bank 1 and the customer; and holding fixed the offer by bank 1, the offer from 2 maximizes the joint payoff to bank 2 and the customer. Thus, there is no feasible deviation in which the customer continues to take both contracts. Suppose a deviation by bank 1 causes the customer to drop bank 2. Then the best possibility will have the same level of security as before, because bank 1's costs are unaffected by the lower security level of bank 2. As a result, the joint profits of bank 1 and the customer can only decrease if the contract with bank 2 is abandoned. Now consider a deviation by bank 2. The best possible offer would have a security level \hat{s}_2 , but by condition (21) the surplus under that contract is lower than the total of profit and utility to the customer and bank 2 under the existing contract. \square

The conditions in this theorem are easily satisfied by making V and ΔV large and making bank 2 inefficient at account protection. The conditions are compatible with conditions that make the second-best outcome one in which both banks provide accounts. Finally, note that the general result in the text applies to this specific case: the equilibrium outcome is not second best.