

Data Breaches and Identity Theft*

William Roberds[†]
Federal Reserve Bank of Atlanta
1000 Peachtree Street, N.E.
Atlanta, GA 30309-4470, USA

Stacey L. Schreft[‡]
The Mutual Fund Research Center, LLC
7301 College Blvd., Suite 220
Overland Park, KS 66210, USA

This revision: August 12, 2009

Abstract

An environment is analyzed in which agents join clubs (payment networks) in order to facilitate trade. The networks compile personal identifying data (PID) so as to match transactors to transactions histories. Technological limitations cause the networks' data management practices to impact each other's incidence and costs of identity theft. Too much data collection and too little security arise in equilibrium with noncooperative networks compared to the efficient allocation. A number of potential remedies are analyzed: (1) reallocations of data-breach costs, (2) mandated security levels, and (3) mandated limits on the amount of data collected.

Keywords: Identity theft, identity fraud, data breach, fraud, money, search

JEL Codes: D83, E42, G28

* Helpful comments were provided by an anonymous referee and by participants in presentations at the Federal Reserve Banks of Chicago and Kansas City, the 2008 Payments Workshop at the Bank of Canada, the 2008 LAEF Conference on Payments and Networks at UC Santa Barbara, and the 2009 Workshop on the Economics of Information Security at the University of London. The views expressed in this paper are not necessarily those of the Federal Reserve Bank of Atlanta, the Federal Reserve System, or The Mutual Fund Research Center, LLC.

[†] Corresponding author: Tel. +1 404 498 8970

[‡] E-mail addresses: william.roberds@atl.frb.org (W. Roberds); sschreft@mutualfundstore.com (S. Schreft).

1. Introduction

Modern information technology enables the collection and storage of large amounts of personal data. While these activities undoubtedly provide economic benefits, it has proved impossible to keep data completely secure against criminal misuse. Survey data suggest that in 2006 identity thieves obtained about \$49.3 billion from U.S. consumer victims. Add in the time and out-of-pocket costs incurred to resolve the crime, and identity theft may have cost the U.S. economy as much as \$61 billion in 2006 (Schreft 2007).

This looks like a large cost—equivalent to two Bear Stearns rescues in a year—but the central policy question is whether the size of these losses indicates a market failure (Anderson et al. 2008). In the mind of the general public, the answer seems to be a resounding “yes.” Press accounts routinely suggest that too much personal identifying data (PID)¹ is being collected and that this data is being stolen too often, leading to excessive identity theft.² This view is echoed in the legal literature on identity theft and data confidentiality,³ where a recurring message is that the credit industry has failed to deliver “efficient confidentiality” of personal data (Swire 2003). Negative popular sentiment has also contributed to the passage of legislation designed to improve data security practices.⁴

Government reports⁵ and industry sources⁶ have argued against the market failure hypothesis. These arguments often emphasize two stylized facts. First, losses from identity theft are small relative to overall usage of payments and credit in today’s economy (e.g., over \$3 trillion in card transactions in the U.S. each year). Second, much identity theft does not result from any

¹ A.k.a. “personally identifiable information” (PII).

² See e.g., Swartz and Acohidio (2007), Caruso (2007), and Dow Jones and Company, Inc. (2008a, b).

³ See e.g., LoPucki (2001, 2003), Solove (2003, 2004), Swire (2003), and Chandler (2008).

⁴ Including the 2003 U.S. Fair and Accurate Credit Transactions Act (with 30+ pages of implementing regulations) and laws in at least 36 U.S. states.

⁵ See e.g., Synovate (2007) and United States Government Accountability Office (2007).

⁶ See e.g., Cheney (2004), Experian (2006), Kirshbaum (2006), McGrath and Kjos (2006), and Javelin Research (2008).

compromise of data stored by businesses, but from opportunistic, low-tech criminal activity (e.g., stolen wallets). Because this type of fraud can be effectively deterred through intensive data analysis (Greene 2009), the implication is that any problem with identity theft could be best addressed by compiling more (e.g., biometric) data on individuals, not less.

Economists (economic theorists in particular) have remained relatively quiet on issues regarding identity theft and data breaches.⁷ This paper offers an initial exploration, using a model derived from contemporary monetary theory. Monetary theory is informative for this analysis, as it focuses on two key market frictions that may be counteracted through the use of PID: (1) displacement of agents' consumption demands over time, and (2) a limited ability to force agents to repay debts. The benefit of a multilateral recordkeeping arrangement—a credit-based payment system—derives from its ability to overcome these frictions, and knowledge of agents' identities helps provide this benefit. Credit is impossible without knowing who the debtor is.

The environment studied below extends the model of identity theft developed in Kahn and Roberds (2008) to allow for identity theft through data breaches. The paper begins by presenting a game-theoretic model of multiple payment card networks. Payment networks are modeled as club arrangements for the sharing of information for intertemporal trade. Each club must decide how much data on its members to assemble into a database, and each also must choose how thoroughly to secure its database. Collecting more PID imposes costs on card-network participants, but as industry sources assert, yields a benefit in terms of deterring attacks on the network. On the other hand, collecting such data can have negative spillover effects, because one network's data can be stolen and used to open an account with another network. A network can reduce data theft (and therefore suppress identity fraud) by better securing its database, but it might be cheaper to suppress fraud by increasing the amount of PID compiled.

⁷ Some relevant literature is discussed in Section 6 below.

Using the model environment, we then compare networks' noncooperative data and security decisions to the decisions that a planner would implement. This comparison supports some facets of the "popular wisdom": divergences in social and private incentives cause data to inefficiently overcollected and undersecured. However, the net effect of these practices is shown to be an inefficiently *low* rate of identity theft at the expense of privacy, irrespective of the division of identity theft between its low-tech and high-tech forms. In other words, the model shows how inefficiency of equilibrium can be consistent with the facts emphasized in the "industry view." A final section of the paper considers some policy remedies for this inefficiency.

In summary, the model developed here allows for calculation of the efficient levels of data accumulation and data security, and for evaluation of policies meant to attain efficiency. More generally, it illustrates how any such calculation should balance the costs of data misuse against the substantial gains afforded by the relaxation of anonymity.

2. Institutional Background

This section provides a brief overview of the phenomenon of identity theft and its relationship to data security. Recent surveys are given in Schreft (2007) and Anderson et al. (2008).

We begin by defining terms. Identity theft can take many forms in practice. The Federal Trade Commission (Synovate 2007) divides identity theft into two broad categories: *existing-account fraud* and *new-account fraud*. Existing-account fraud occurs when a thief steals an existing payment card or similar account information (e.g., a checking account number) and uses these to purchase goods and services. Traditionally, new-account fraud occurs when a thief uses someone else's PID to open a new account. As will be clear below, new-account fraud is the type

of identity fraud that occurs in the model.⁸

There are no comprehensive statistics on the prevalence of identity theft, or definitive estimates of its cost. In a widely cited survey, the Federal Trade Commission (FTC) estimated that in 2006, 3.7 percent of the U.S. adult population fell victim to some form of identity theft, at a cost of roughly \$16 billion (Anderson et al. 2008). These figures are likely underestimates, however, because they omit certain forms of identity theft as well as many of its indirect costs. Adjusting for some of these effects easily quadruples the cost estimate (Schreft 2007).

A data breach occurs when an unauthorized party is able to access personal data that has been collected by an organization (e.g., business or payment service provider). Data breaches can facilitate either existing-account fraud (as when credit-card information is stolen) or new-account fraud (as when PID is stolen).⁹ There is no definitive estimate of how many cases of identity theft have resulted from data breaches. Certainly, data breaches are numerous and increasing: for example, the information-security website *Attrition.org* lists 553 reported data breach “incidents” for 2008, leading to the compromise of 83 million records of personal data, as compared to 11 reported incidents and 6 million compromised records in 2003. A data breach does not necessarily result in identity theft, as data may be stolen without being used for fraudulent purposes. Nevertheless, there seems to be widespread agreement that data breaches promote identity theft. The United States Government Accountability Office (2007), for example, examined 24 data breaches between 2000 and 2005 in which large amounts of data were compromised, and was able to conclusively link four of these to subsequent outbreaks of fraud.

⁸ The term “new-account fraud” includes an increasingly prevalent type of fraud, which is *fictitious* or *synthetic identity fraud*. In this type of fraud, a thief combines information taken from a variety of sources with invented information to create a new, fictitious identity. By one recent estimate, more than 80 percent of all new-account identity theft has occurred using synthetic identities (Coggeshall 2007).

⁹ Actually, because many credit-card issuers will open accounts for people who present an existing credit card, a data breach involving the theft of credit-card information also contributes to new-account fraud.

Also, identity theft can occur without data breaches. In consumer surveys, victims of identity theft who know how their information was stolen commonly attribute their loss to low-tech channels such as: lost/stolen wallets (e.g., 33% of cases reported in Javelin 2008), fraud by acquaintances (17%), or stolen mail (6%). Since, however, many consumers do not know how their data was stolen, these surveys probably underestimate the impact of data breaches. Gordon et al. (2007) examine 274 cases of identity theft prosecuted by the Secret Service over 2000-2006, and find that half of these resulted from the compromise of data at a business.

The costs of identity theft must be weighed against the benefits provided by the availability of PID, which lies at the heart of credit-based systems of exchange. There are no precise estimates of these benefits, but the sheer volume and increasing popularity of services such as card-based payments indicates that these are substantial. These benefits extend beyond the traditional “credit” sector to industries such as utilities, wireless communications, and medical care, all of which depend on accurate identification of individuals.

3. Environment

The environment is adapted from the well-known Kiyotaki and Wright (1989) model of decentralized exchange.

3.1 Basic features

The economy exists in continuous time and consists of a continuum of risk-neutral agents. Associated with each agent is a unique fixed vector of personal data known as the agent’s *identity*. This vector has effectively infinite dimension.

Agents are divided into groups G_A and G_B of unit measure, where $G_A \cap G_B = \emptyset$. All trade occurs among agents in the same group. Within each group, agents are congenitally subdivided into *legitimate agents* and *frauds* (i.e., identity thieves). F denotes the fraction of frauds in

the population. Legitimate agents and frauds have the same consumption preferences, but differ in two respects. First, legitimate agents are able to produce tradable goods, while frauds lack this ability. Second, frauds are sometimes able to impersonate other agents, while legitimate agents cannot.¹⁰ An agent's identity, group, and legitimacy are all private information until revealed through costly verification and/or observation of the agent's behavior.

Goods are traded within groups through random matches of buyers and sellers.¹¹ There are no double coincidences and no repeated matches, and money is not available, so trade can only occur using some form of multilateral credit.¹² Any agent with access to credit derives a flow utility $u > 0$ from acquisition of other agents' goods. At some point during each unit time interval, agents may be called upon to supply up to a unit measure of their endowment good to other agents. Legitimate agents can perform this action at a disutility of c per unit, where $u > c > 0$. Whether or not an agent has supplied goods is not observable until the next discrete date $n = 0, 1, 2, \dots$, at which point it becomes public information. Information on agents' consumption behavior is not available without the application of a specific technology, which is described below.

Credit-based exchange in the model requires arrangements for sharing two kinds of information: (1) sufficient knowledge of agents' transaction histories (as in, e.g., Kocherlakota 1998) and (2) sufficient knowledge of agents' identities, in order to associate would-be consumers with histories (as in Kahn and Roberds 2008). These arrangements are modeled as clubs for

¹⁰ The environment studied can be generalized to allow for the endogenous choice of agents to specialize in fraudulent activity; see Kahn and Roberds (2008).

¹¹ Additional details of the model are given in Appendix A.

¹² The model could be modified to allow agents to transact with cash as well as with credit. This generalization is explored in Camera and Li (2008), Martin, Orlando, and Skeie (2008), and Monnet and Roberds (2008).

sharing this information, which we visualize as credit card networks.¹³ The analysis will consider the case where one club exists for each group of agents.¹⁴

To encompass the possibility of identity theft via data breaches, the environment allows for turnover in club membership. Turnover in membership gives each club an incentive to retain data on its members' identities, so as to distinguish existing club members from new applicants. However, the presence of such stored data creates opportunities for data thieves.

To incorporate turnover, agents in the model consist of stochastically lived overlapping generations. At each discrete date $n = 0, 1, 2, \dots$, a randomly selected subset of agents dies and is replaced by newborn agents. Newborn agents have unique identities but are otherwise indistinguishable from the agents they replace. The measure of deaths and births is given by $1 - \beta$, where $0 < \beta < 1$, i.e., β serves as a discount factor in agent's decisions. The deaths of agents and the identities of the dead immediately become public information, so only the living are potential victims of identity theft.

3.2 Benchmark: exchange with costless identification

Consider the case where information on agents' identities can be costlessly assembled and stored, so agents can be perfectly identified. Agents from each group form two clubs at time $t = 0$. An agent joining club i , $i = G_A, G_B$, reveals his identity to the club, and the agent receives an uncounterfeitable credit card that signals his membership in the club. The card can be authenticated by all club members at no additional cost, and allows its holder to enjoy the consumption flow u . At each discrete date n , the club learns whether its members have produced goods during

¹³ Clubs are a natural arrangement given the nonrival nature of the good (information) that is to be allocated (Varian 1998). In practice, such information is managed through the interaction of many parties, including card issuers, credit bureaus, and transaction processors. Data sharing among these parties is subject to public-goods problems (Varian 2004). The analysis below abstracts from such problems in order to focus on spillover effects.

¹⁴ Appendix D presents an extension of the model which allows for endogeneity in club size.

the preceding unit interval. Producers remain in the club; nonproducers are subject to expulsion and to a nonpecuniary penalty (e.g., stigma or criminal sanctions) equal to $X > u$ utils. Subsequently, $1 - \beta$ club members die and membership is opened to newborn agents.¹⁵

For this case, it is straightforward to show¹⁶ that exchange through clubs is self-sustaining: all legitimate agents in the model have an incentive to join the appropriate club and remain in it over their lifetimes, while all frauds are excluded. Legitimate agents' expected value of continued club membership at each discrete date n is

$$V \equiv [(u - c)(1 - F)]/r, \quad (1)$$

where the clubs' "discount rate" $r = \beta^{-1} - 1$. Identification of consumers is key to the viability of the clubs: absent identification, the payoff to someone who consumes and never produces is

$$[u(1 - F)]/r > V, \quad (2)$$

i.e., shirking always beats working, so that the clubs collapse. In general, exclusion of frauds is also necessary for the clubs to exist. If all frauds are admitted to the clubs and provided with consumption goods, then the value of legitimate agents' membership falls to

$$[u(1 - F) - c]/r, \quad (3)$$

which is negative for F sufficiently close to one.

3.3 Exchange with costly identification

More generally, reliable identification of agents requires the use of a costly technology. Clubs accomplish identification by collecting a subset of each agent's identity. For this model, the value of such information, and the costs of managing it, is represented by the amount of identifying information disclosed, not by the type of information. The amount of information

¹⁵ See Boyd and Prescott (1987) for an analysis of clubs with a similar structure.

¹⁶ See Appendix A.

disclosed is given by $d_{i,n}$, referring to the number of elements an agent must disclose from his identity vector to be identified by club i at discrete dates n . For analytical convenience $d_{i,n}$ is taken to be a continuous variable, i.e., $d_{i,n} \in \mathbb{R}_+$. Each club compiles and maintains a database containing the identifying information disclosed by its members. The cost to the two clubs of merging their databases is assumed to be prohibitive.

Identity verification has two costs: (1) a fixed one-time cost of K utils, which is incurred when an agent initially joins club i and is borne pro-rata by all legitimate club members, and (2) a per-discrete-period, per-member cost of processing and maintaining the data record $d_{i,n}$ for each club member. This second cost is given by $kd_{i,n}$, where $k > 0$ and is also borne by all legitimate club i members. Note that the parameters K and k reflect physical costs but perhaps also intangible costs associated with the loss of privacy stemming from identity verification. Also note that $d_{i,n}$ can vary across discrete periods, i.e., a club can vary the amount of identifying data it requires from its members from one period to another. Once a club has collected data at discrete dates $n = 0, 1, \dots$, the data must be maintained until date $n + 1$ if the club is to avoid paying the initial identity verification cost K on all members at time $n + 1$.¹⁷

Following the initial verification of an agent's identity, the agent receives an uncounterfeitable credit card. Credit cards are issued at zero additional cost. Because credit cards are uncounterfeitable, identity theft in the model does not involve the cloning of existing cards or use of existing card numbers: there is no existing-account fraud. Rather, all identity theft involves the opening of a new credit-card account in the name of an apparently legitimate agent.¹⁸

¹⁷ Data that has been retained for one period and is not useful for identification can be costlessly destroyed by the clubs. A data destruction technology could be introduced at the cost of some additional complexity.

¹⁸ The model could be extended to allow for existing-account fraud. Formally, existing-account fraud is quite similar to counterfeiting, which has been analyzed in the money literature (see section 6 below).

Credit cards issued at discrete dates n have a virtual expiration date of $n + 1$. That is, at discrete date $n > 0$, each club compiles a list of agents who have supplied goods during the preceding interval $[n - 1, n)$. Members who have not supplied goods are revealed as impersonators (frauds) and removed from the club, while those who have supplied goods continue their membership. Apart from exclusion from the club, no penalties can be applied to impersonators because their real identities are unknown.¹⁹

3.4 The prevalence of identity theft

Identity theft occurs when a fraud gains access to a club by convincingly impersonating a legitimate agent. Reflecting the distinctions in the policy literature, identity theft in the model can occur through “high-tech” methods (i.e., involving data breaches) or “low-tech” methods (without data breaches). High-tech methods require less effort but more skill, i.e., a successful breach lowers the effort cost of fraud. Because the submission of duplicate PID of an existing club member would be automatically revealed as fraudulent, data observed in a breach of club j 's database is always used to gain access to club i .

The probability of a successful data breach depends on how well the target club secures its database. Suppose that club i maintains member data $d_{i,n-1}$ over the interval $t \in [n - 1, n)$. The club then chooses a security variable $s_{i,n-1} \geq 0$ that determines, for the next discrete date n , the likelihood of a data breach, given the technical skills of would-be data thieves.

More specifically, the variable $s_{i,n-1}$ is the *skill threshold* required to access club i 's database at discrete date $t = n$. The distribution of technical skills s within the population of frauds is

¹⁹ One can conceive of other arrangements for trade within the club. For example, each producer could verify each buyer's identity independently, but this would require that each buyer's verification cost be repeatedly incurred (infinitely often). Or, the club could verify members' identities at the beginning of each discrete period, issue “no-name” credit cards valid for only one period, and dispose of all identifying information on its members. In what follows it is assumed that the value of the initial verification cost K is sufficiently high relative to other costs in the model that the use of anonymous credit cards is not an attractive option.

time invariant, and is given by the probability distribution function $\Phi(s)$, where $\Phi(s) < 1$ for $s < \infty$. Intuitively, by setting a higher skill threshold, the club can lower the proportion of the population of frauds that can potentially steal the club's data. Increasing the skill required for data breaches brings with it increased costs, however. In particular, adopting skill threshold $s_{i,n-1}$ results in a disutility to all legitimate members of club i of $\ell s_{i,n-1}$ incurred at discrete date $n-1$, where $\ell > 0$. Thus, the possibility of a breach is never completely eliminated.

Frauds lacking the technical skills for data theft can attempt to obtain the necessary data for impersonation through other, "low-tech" means. Compiling the data $d_{i,n}$ necessary for entry into club i at discrete date n involves a utility cost $\varepsilon d_{i,n}$, where the "effort cost" $\varepsilon > 0$. ε is assumed to have a time-invariant distribution $\Gamma(\varepsilon)$ over the population of frauds, and Γ does not depend on the security variables s .

Frauds with sufficient skills may reduce their effort costs by stealing data. If a fraud of group i breaches club j 's date $n-1$ database, and obtains data $d_{j,n-1}$, then a fraction $\eta \in (0,1)$ of this data can be applied to gain membership to club i . In this case, the net amount of data the fraud must synthesize to gain access to club i is $\max\{d_{i,n} - \eta d_{j,n-1}, 0\}$, and his net effort cost is reduced to $\varepsilon \max\{d_{i,n} - \eta d_{j,n-1}, 0\}$. Under this specification, spillover effects arise due to the overlap η between the kinds of information in various databases of personal identifying data. The analysis below will concentrate on cases where $\eta \rightarrow 1$, i.e., where this overlap is substantial, though still imperfect.²⁰

To summarize, the prevalence and type of identity fraud committed in club i during

²⁰ This expresses the idea that databases of PID tend to contain many common elements such as name, address, birth date, social security number, etc. Requiring η to be strictly less than unity ensures that positive effort is required for impersonation in equilibrium.

$[n, n+1)$ depends on three factors: (1) the amount of data $d_{i,n}$ needed to gain access to club i at discrete date n , (2) the skill threshold $s_{j,n-1}$ specified by club j at discrete date $n-1$, and (3) the amount of club j 's data obtainable through a breach at date n , $\eta d_{j,n-1}$. More specifically, club i 's equilibrium rate of identity theft from unskilled frauds over $t \in [n, n+1)$ is given by

$$\rho_{i,n}^U \equiv F\Phi(s_{j,n-1})\Gamma\left(\frac{u(1-F)}{d_{i,n}}\right). \quad (4)$$

Given symmetry between the clubs, club i 's equilibrium rate of identity theft from skilled frauds over $t \in [n, n+1)$ is

$$\rho_{i,n}^S \equiv F(1-\Phi(s_{j,n-1}))\Gamma\left(\frac{u(1-F)}{d_{i,n}-\eta d_{j,n-1}}\right). \quad (5)$$

Club i 's total rate of identity theft over $[n, n+1)$ is given by $\rho_{i,n} \equiv \rho_{i,n}^U + \rho_{i,n}^S$.

3.5 The costs of identity theft

In addition to identity verification costs, impersonation of legitimate agents by frauds imposes three other types of costs. All legitimate agents are risk neutral and share the same preferences, so there is no loss of generality in assuming that these costs are equally distributed across legitimate club members.

The first cost is simply the cost of providing goods to frauds, which is given by c utils per period, per identity theft. In principle, this cost derives from a transfer from legitimate agents to frauds, but is nonetheless economically meaningful because widespread fraud can undermine the viability of the card networks (see expression (3)). In practice this cost is considerable. For example, the FTC survey (Synovate 2007) estimates the median value of goods obtained through new-account fraud to be \$1,350 for each stolen identity.

The second type of cost is the cost of resolving an identity theft. That is, discovery of an impersonator in club i imposes a resolution cost of L on the club, which represents both a social and private cost. L may include physical costs, loss of leisure time, and inconvenience. This cost is more difficult to measure but nonetheless significant. In the FTC survey, the median amount of time spent by a consumer to resolve a case of new account fraud was 10 hours, equivalent to hundreds of dollars in monetary value. Another example is given by Douglas (2008), who reports that it costs a card issuer about \$25 to reactivate any compromised card account. Other, less readily quantifiable costs of resolving identity theft are catalogued by Anderson et al. (2008), and can include harassment of victims by debt collectors, denial of utility service, and the costs of deflecting civil lawsuits and criminal investigations.

The third type of cost results only when identity theft results from a data breach, i.e., from skilled identity theft. When a club's data is stolen and used to gain fraudulent access to the other club, the members of the first club are subject to an additional resolution cost, or "breach cost" $B > 0$. Empirically, B may be smaller than c or L , but is still nonnegligible. A report by Ponemon Institute (2006) offers examples of such costs. These include the costs of notifying people whose data has been compromised (\$13 per data record breached), labor costs ("lost productivity", \$30 per record), and the costs of managing potential legal liabilities (\$11).²¹

3.6 Clubs' objectives and steady-state equilibrium

The analysis below will focus on steady states. A steady-state allocation in this economy consists of two ordered pairs $\{(d_i, s_i)\}_{i=1,2}$, where d_i gives the data length and s_i gives the skill

²¹ Typically these costs represent the costs of legal safeguards against potential civil and criminal actions stemming from a breach, rather than reallocations of fraud losses incurred by other parties. Reallocations of fraud losses through the legal system are studied in Section 5 below.

threshold chosen by club i . Taking into account all costs, the steady-state, per-period cost of identity theft to legitimate members of club i is

$$C_i = [(1 - \beta)K + kd_i + \ell s_i] + [\rho_i^U (c + L)] + [\rho_i^S (c + L) + \beta \rho_j^S B] \quad (6)$$

i.e., the sum of data costs (data collection and data security), plus the costs of unskilled ID theft (goods provided to frauds plus resolution costs), plus the costs of skilled ID theft (goods provided and resolution costs, plus the costs of resolving breaches).

The general thrust of the analysis below will be to characterize inefficiencies as discrepancies between Nash outcomes and efficient allocations. In Nash equilibrium, each club i unilaterally chooses (d_i, s_i) to maximize the steady-state continuation value of legitimate club membership, which is

$$V_i^f \equiv V - (C_i / r) . \quad (7)$$

Evidently, this is the same as minimizing C_i . A steady-state allocation $\{(d^*, s^*)\}$ is a symmetric Nash equilibrium if $(d_i, s_i) = (d^*, s^*)$ minimizes C_i when club j chooses $(d_j, s_j) = (d^*, s^*)$.²²

The planner operates under the same informational constraints as the clubs in the decentralized arrangements, but is able to coordinate the choice of d and s across clubs. The planner's objective is to minimize the steady-state costs of identity theft to legitimate agents, including all costs resulting from data breaches, i.e., the planner chooses (d, s) to minimize²³

$$C_p = (1 - \beta)K + kd + \ell s + \rho^U (c + L) + \rho^S (c + L + \beta B) . \quad (8)$$

²² In addition, existence of equilibrium requires that certain incentive conditions (given in Appendix A) be satisfied in order to guarantee legitimate agents' participation in the clubs. These can be shown to hold under mild parametric restrictions (given in Proposition 1 below).

²³ The allocation chosen by the planner represents a constrained-efficient allocation, since the planner places no weight on the utility of either frauds or the initial generation of legitimate agents. "Golden-rule" welfare criteria such as (8) are widely employed in overlapping generation settings but of necessity also arbitrary.

4. Analysis of equilibria

This section considers steady-state equilibria for parametric specifications for Φ and Γ . In particular, frauds' skill endowments s are specified to follow an exponential distribution $\Phi(s)$ with hazard rate $\phi \equiv \Phi'/(1-\Phi)$, and the distribution $\Gamma(\varepsilon)$ of frauds' effort costs is specified as a uniform distribution, normalized to $U[0,1]$. These specifications allow for unique equilibria that can be expressed in closed form.

To develop intuition for the model, we consider some particular cases.

4.1 Case 1: All identity theft stems from data breaches

Suppose that neither club secures its data so that, in effect, all frauds are skilled, i.e. $\Phi = 0$. For this case, clubs' rate of identity theft is not determined by the amount of data they collect, but instead by the amount of additional data an ID thief must come up with (beyond that obtainable through a breach) in order to gain access to a club. That is, from (5), club i 's equilibrium rate of ID theft is determined by $e_i = d_i - \eta d_j$. Changing variables and differentiating C_i , club i 's first-order condition is

$$\frac{uF(1-F)(c+L)}{e_i^2} = k + \frac{uF(1-F)\eta\beta B}{e_j^2}. \quad (9)$$

Each club sets the marginal benefit of fraud deterrence through PID collection [LHS(9)] equal to its marginal cost [RHS (9)], which is the sum of the physical/intangible data cost k and the cost of increased vulnerability to data breaches. Best responses are given by

$$e_i = \sqrt{\frac{uF(1-F)(c+L)e_j^2}{ke_j^2 + uF(1-F)\eta\beta B}}. \quad (10)$$

Since RHS (10) is strictly increasing, the unique solution for d^* is

$$d^* = \frac{1}{1-\eta} \sqrt{\frac{uF(1-F)(c+L-\eta\beta B)}{k}} \quad (11)$$

under the empirically plausible restriction (see section 3.5 above) that $c+L > \eta\beta B$.²⁴ Quantity d^* may be contrasted with the unique solution to the planner's problem, which from (8) is

$$d_p = \frac{1}{\sqrt{1-\eta}} \sqrt{\frac{uF(1-F)(c+L+\beta B)}{k}}. \quad (12)$$

Comparing (11) and (12), it follows that $d^* > d_p$ when

$$(1+r)\eta(c+L) > B, \quad (13)$$

i.e., when breach costs B are less than the costs to the other club of increased ID theft stemming from the breach $c+L$, adjusted for “data overlap” η and present value.

This case of the model offers a classic example of a negative production externality: when (13) holds, the discrepancy between private and social costs results in overcollection of PID in equilibrium. Each club internalizes the deterrence benefit of requiring personal data from its members, but does not internalize the cost to the other club of facilitating future identity theft.

Not surprisingly, overcollection of data increases data breaches relative to the planner's allocation, consistent with the “popular wisdom” discussed in the Introduction. However, from (5), this discrepancy also lowers identity theft rates: inefficiency of the noncooperative equilibrium does not stem from too much identity theft, but instead from too much data being collected. Each club would like to compile less data on its members (i.e., reduce d_i) but, given the actions of the other club, cannot do so without encouraging high rates of fraud.

The spillover parameter η influences the extent of this inefficiency. Under condition (13), both d^* and d_p are increasing in η , but as stolen data becomes increasingly useful for ID theft,

²⁴ In addition, existence of equilibrium requires sufficiently low data costs and a discount factor sufficiently close to unity (i.e., conditions (c) and (d) of Proposition 1 below).

i.e., as $\eta \rightarrow 1$, the clubs acting independently require ever larger multiples of the amount of data that a planner would collect, i.e., $(d_p / d^*) \rightarrow 0$.

4.2 Case 2: Fixed proportions of skilled and unskilled identity thieves

Next consider a slightly more general case with fixed proportions of skilled and unskilled identity thieves within the population of frauds. Differentiating C_i with respect to d_i yields the first-order condition

$$\begin{aligned} uF(1-F) \left[\left(\frac{\Phi(s_j)}{d_i^2} \right) (c+L) + \left(\frac{(1-\Phi(s_j))}{(d_i - \eta d_j)^2} \right) (c+L) \right] \\ = k + \beta \eta u F(1-F) \left(\frac{1-\Phi(s_i)}{(d_j - \eta d_i)^2} \right) B, \end{aligned} \quad (14)$$

where again LHS (14) represents club i 's marginal benefit of increased data collection (reduction in unskilled and skilled ID theft) and RHS (14) represents its cost (physical/intangible cost plus data breach vulnerability). Fixing $\Phi(s_i) = \Phi(s_j) = \underline{\Phi}$, and solving as above for equilibrium data length d^* yields a unique solution

$$d^* = \frac{1}{1-\eta} \sqrt{\frac{uF(1-F) \left(\underline{\Phi}(c+L)(1-\eta)^2 + (1-\underline{\Phi})(c+L-\eta\beta B) \right)}{k}}, \quad (15)$$

when $c+L > \eta\beta B$. For the planner's problem, differentiating C_p with respect to data length d yields

$$uF(1-F) \left[\frac{(c+L)\Phi(s)}{d^2} + \frac{(c+L+\beta B)(1-\Phi(s))}{d^2(1-\eta)} \right] = k, \quad (16)$$

where again marginal benefits (all internalized by the planner) are displayed on the left and marginal costs on the right. Solving (16) for d when $\Phi(s) = \underline{\Phi}$ yields

$$d_p = \frac{1}{\sqrt{1-\eta}} \sqrt{\frac{uF(1-F)(\Phi(c+L)(1-\eta) + (1-\Phi)(c+L+\beta B))}{k}}. \quad (17)$$

Using (15) and (17), it can again be shown that $d^* > d_p$ under condition (13). As in case 1, inefficiency of the equilibrium allocation stems from overcompilation of PID. A key difference between case 1 and case 2, however, is in the quantitative manifestation of this inefficiency: equilibrium rates of unskilled and skilled ID theft (from (4) and (5)) are both below those in the planner's allocation. Which means, depending on the value of Φ , that the principal effect of the overcollection of PID may not be a reduction in skilled identity theft—the underlying source of the inefficiency—but instead a reduction in identity theft by the unskilled. In other words, inefficiency of the symmetric equilibrium persists, even when much observed identity theft does not involve data breaches.

4.3 Case 3: Endogenous skill thresholds

In the most general case, the clubs can limit skilled identity theft by increasing data security, i.e., each club i minimizes its costs C_i by setting both d_i and the security level of its data, given by the skill threshold s_i . Club i 's first-order condition in d_i is given by (14); its first-order condition in s_i is

$$uF(1-F) \left(\frac{\beta B}{d_i(1-\eta)} \right) \Phi'(s_i) \leq \ell, \quad (18)$$

with equality for $s_i > 0$, i.e., the club increases security as long as its marginal benefit in terms of reduced breach costs [LHS (18)] exceeds its marginal cost ℓ . Likewise, from (8), the planner's first-order conditions in d and s are given by (16) and

$$uF(1-F) \left(\frac{\eta(c+L) + \beta B}{d(1-\eta)} \right) \Phi'(s) \leq \ell, \quad (19)$$

with equality for $s > 0$. Comparing (18) and (19), note that for a given data length d , the planner internalizes the benefit $\eta(c + L)$ of each club's data security for the other club (deterrence of skilled identity theft), while in equilibrium the individual clubs do not.

One possibility is that neither club opts to secure its data in equilibrium. Data length d^* is then set as in case 1 above. Substituting (11) into (18), such an equilibrium exists if

$$\beta B \phi \sqrt{\frac{uF(1-F)k}{c+L-\eta\beta B}} < \ell, \quad (20)$$

i.e., if the marginal payoff to security (proportional to the hazard rate ϕ of the skill distribution) is always below its marginal cost. Clearly (20) is satisfied for $\phi > 0$ sufficiently small.

The discussion in the rest of this section focuses on the case where the clubs set a positive security level in equilibrium. Sufficient conditions for existence and uniqueness are given in the following proposition (proofs are in Appendix B):

Proposition 1. *A unique symmetric steady-state equilibrium (d^*, s^*) with positive security effort ($s^* > 0$) exists when*

- a) *the hazard rate ϕ of the skill distribution is sufficiently large;*
- b) *the breach cost B is less than the other costs of identity theft, adjusting for present value, i.e., $B < (1+r)(c+L)$;*
- c) *information and data security are sufficiently cheap (costs $K, k, \ell > 0$ are small);*
- d) *the clubs' discount rate ($r > 0$) is sufficiently small.*

Corollary to Proposition 1. *Under the same conditions, there exists a unique solution to the planner's problem (d_p, s_p) with $s_p > 0$.*

Condition (a) in Proposition 1 guarantees that security technology is good enough for the clubs to find it worthwhile to secure their databases. Condition (b) is a convexity condition.

Conditions (c) and (d) guarantee that the benefits of club membership are not overwhelmed by data costs.

Solutions for equilibrium and optimal allocations in this case are more complicated than in the previous cases (see Appendix B), but can be shown to obey the following properties:

Proposition 2. *(Comparative statics). Under the conditions of Proposition 1,*

- a) s^* increases with data collection costs k , the hazard rate ϕ , and with breach costs B as $\eta \rightarrow 1$, but decreases with ID theft costs c and L , security costs ℓ , and the discount rate r as $\eta \rightarrow 1$;
- b) s_p increases with k , ϕ , c , L , and B , and decreases with ℓ and r ;
- c) d^* increases with c , L , ℓ , and r , and decreases with B , k , and ϕ ;
- d) d_p increases with c , L , and ℓ , decreases with k and ϕ ; and does vary with r or B .

A noteworthy aspect of the comparative statics results is that the planner's choice of data length d_p does not depend on the breach cost B . Instead, as B increases (decreases), the planner responds by increasing (decreasing) data security s_p : along this dimension, in the efficient allocation there is a separation between "data privacy" (the amount of data collected by the planner) and "data security" (the skill threshold applied).²⁵ Due to spillover effects, this separation does not hold for the noncooperative equilibrium. In particular, when breach costs fall, there is less internalization of the costs of clubs' data policies. Hence, the clubs respond to lower breach costs by lowering security levels *and* increasing the amount of data that they compile, in essence substituting data collection for data security.

²⁵ As is shown in Appendix B, a complete separation of privacy and security relies on the specification of a constant hazard rate of hacking skills. Under a more general setup we would still expect the planner's data length decision to be less sensitive to breach costs than the Nash clubs'.

The next results show that, as data overlap $\eta \rightarrow 1$, this same tendency leads to overcollected and undersecured PID in the noncooperative equilibrium.

Proposition 3. *Under the conditions of Proposition 1,*

- a) s^* and s_p are increasing in η ;
- b) As $\eta \rightarrow 1$, $s^* \rightarrow \bar{s} < \infty$ and $s_p \rightarrow \infty$, so that $s^* < s_p$.

Proposition 4. *Under the conditions of Proposition 1,*

- a) d_p does not vary with η , and d^* is increasing in η as $\eta \rightarrow 1$;
- b) As $\eta \rightarrow 1$, $d^* \rightarrow \infty$, so that $d^* > d_p$.

Increasing data overlap leads to greater spillovers between the clubs' decisions, and makes skilled identity theft more likely. The planner again treats this purely as a security issue: the skill threshold s_p increases while PID collected d_p remains the same. The Nash clubs also respond by increasing security, but only up to a point, since they do not receive the full benefits of this action. Instead, they find it cheaper to deter unskilled fraud by compiling additional PID, since they do not internalize all of the resultant costs. In other words, club members' privacy is sacrificed as a defense against insufficient security efforts by the other club.

The clubs' inclination to substitute data collection for data security also shows up in comparisons of identity theft rates. These are given as

Proposition 5. *Under the conditions of Proposition 1,*

- a) The rate of skilled identity theft ρ^S is greater in the symmetric equilibrium than for the planner's allocation;
- b) As $\eta \rightarrow 1$, the rate of unskilled identity theft ρ^U is less in the symmetric equilibrium than for the planner's allocation;

c) For ℓ/k bounded, as $\eta \rightarrow 1$ the total rate of identity theft ρ is less in the symmetric equilibrium than for the planner's allocation.

As in the previous cases, the net result of the Nash clubs' data overcollection is a lower overall rate of identity theft (Proposition 5(c)). However, Propositions 5(a) and 5(b) show that, in contrast to the previous cases, suppression of fraud is not uniform but is concentrated in the unskilled forms of identity theft. Thus, with endogenous data security, apparent success in combating unskilled identity theft can be a symptom of failure to deter its skilled counterpart.

5. Attaining efficiency

This section considers three types of policies that have been proposed as remedies for inefficiencies stemming from data breaches: (1) reallocations of the costs of data breaches through the legal system; (2) mandating improved data security; and (3) regulatory limits on the amount of PID collected.

The first approach would increase each network's civil liability for a data breach, i.e., increase each network's breach costs to $B' = B + \pi$, where $\pi > 0$ represents the network's liability.²⁶ In the simplified cases 4.1 and 4.2 above, efficiency can be restored by choosing a level of liability that causes each club to internalize the full costs of its data collection, i.e., by setting

$$\pi = \pi^* \equiv (1+r)(c+L) - (B/\eta), \quad (21)$$

Note that if B in (15) is replaced with $B' = B + \pi^*$, then it follows from (17) that $d^* = d_p$. Also note that club i 's liability for a data breach π^* is bounded by the "actual loss" or "economic loss" suffered by club j , i.e., $(1+r)(c+L)$, which represents the practical limit of liability under the U.S. and Canadian legal systems (Chandler 2008).

²⁶ In practice it can be difficult to enforce liability due to contracting limitations and uncertainty concerning the source of the stolen data (Schreft 2007 and Chandler 2008). The analysis here abstracts from such constraints. Laws requiring consumer notification when a breach occurs also raise the costs of a data breach and would have comparable incentive effects.

A policy of increasing liability for a data breach does not fare as well in case 4.3 with endogenous security. When security effort is positive, it can be shown that imposing any liability up to π^* improves welfare (see Appendix C), but this type of policy cannot simultaneously correct incentives in d and s , and so does not restore efficiency. Intuitively, such a policy undercorrects security incentives, causing the networks to continue to overcollect personal data.

The second regulatory approach, which has been emphasized in the U.S., is to mandate minimum standards for data security, while allowing for private determination of how much PID should be collected.²⁷ In addition to the obvious benefit of reducing the prevalence of data breaches, the model predicts that improved security also lessens incentives to collect PID. The optimal regulatory choice of security level cannot be expressed in closed form, but in numerical examples (see Table 1 below) it closely approximates the planner's level of security s_p .

The third approach seeks to improve incentives by limiting the amount of data collected, while allowing the networks to choose their levels of security.²⁸ It can be shown (see Appendix C) that the optimal limit on data collected corresponds to the level of data d_p that the planner would collect. Through the substitution effect outlined above, this also increases the clubs' incentives to keep their data secure.

To better gauge the efficacy of the various regulatory approaches, allocations were computed numerically. Table 1 below displays some typical results. Parameter values for the example are $c + L = 10$; $B = 1$; $\beta = 0.9$; $\phi = 2$; $\eta = 0.9$; $k = 0.5$; $\ell = 0.1$. These parameter values satisfy the conditions for Propositions 1-5. They allow for considerable data spillover ($\eta = 0.9$) and place a high value on privacy of PID ($(k/\ell) = 5$). Breach costs B are small relative to the other

²⁷ See Keitel (2008) for a discussion of applicable U.S. laws and regulations.

²⁸ This regulatory approach has not been emphasized in the U.S. However, an example of this type of policy can be found in the European Union Privacy Directive, which restricts the collection of some types of personal data.

direct costs of identity theft $c + L$, reflecting the cost figures cited in Section 3.5. To facilitate comparisons, the normalizations $K = 0$ and $uF(1 - F) = 1$ are adopted.

<Insert Table 1 here>

The Table illustrates the comparisons stated in Propositions 3, 4, and 5. In symmetric equilibrium (row 2), the networks collect more than twice as much data as in the efficient allocation (row 1), but skilled identity theft rises because security effort is also reduced. Overall identity theft is suppressed in the symmetric equilibrium, but the welfare cost of this suppression is high since so much data is collected. Unskilled identity theft predominates in equilibrium.

Imposing liability π^* for data breaches (row 3) increases security effort and reduces skilled identity theft, but does not fully correct incentives. Better results are obtained by constraining security to the efficient level (row 4), which approximates the planner's allocation. Note that the success of this policy requires virtual eradication of data breaches. An apparently less stringent (though, in the U.S. case, least applied) policy of constraining PID collection (row 5) does as well as imposing civil liability, but also has the highest identity theft rates of any of the allocations studied.

6. Relationship to the Literature

The above analysis builds on models of exchange in search-theoretic environments. Many papers in this literature examine fraudulent transactions, including counterfeiting (Green and Weber 1996; Kultti 1996; Monnet 2005; Williamson 2002; Nosal and Wallace 2006; Cavalcanti and Nosal 2007) and various other types of fraud (Kahn et al. 2005, Camera and Li 2008, Kahn and Roberds 2008). What is new here is the consideration of an empirically significant type of transactions fraud stemming from the theft of identifying data.

The framework presented also draws on the literature on the economics of information

security (Anderson and Moore 2006). Varian (2004) presents a game-theoretic model in which “system reliability” (e.g., deterrence of identity theft) is modeled as a public good within a network of agents. Varian’s model is extended by Grossklags et al. (2008) to allow for individual insurance (e.g., security effort) against system failures.

The environment above is similar to these models in the sense that knowledge of PID functions as a club good within each transactions network, supplying a network-wide level of security against fraud. However, the focus here is on potential negative spillovers across networks: provision of the same good (data) that suppresses identity theft for one club increases the likelihood of identity theft for the other. Efficient management of personal data strikes a balance between within-club benefits and cross-club costs.

7. Conclusion

This paper has presented a model in which identity theft arises endogenously and the concept of efficient confidentiality for personal identifying information (PID) has meaning. An allocation provides efficient confidentiality if the amount of PID shared for identity verification and the security of that data allow groups of agents to engage in beneficial transactions at minimal cost. Consistent with the “popular wisdom,” inefficiencies can arise due to spillovers from one group of agents’ decisions along these dimensions to another’s. Inefficient outcomes are compatible with empirical patterns of identity theft that are emphasized in industry discussions. Interventions such as regulation of security practices can improve welfare, but the multidimensional nature of the security problem means that attaining efficiency may be problematic.

These results have been developed in the context of a particular methodology, one that abstracts from many of the complexities of modern institutions. However, the basic idea behind this approach—that the compilation, exchange, and storage of PID, despite its risks and costs,

can enable otherwise infeasible intertemporal exchanges of goods—can be generalized and should provide impetus for further research.

Table 1: Numerical comparison of allocations						
	1. PID collected d	2. Security level s	3. Unskilled ID theft $100 * \rho^U$	4. Skilled ID theft $100 * \rho^S$	5. Total ID theft $100 * \rho$	6. Steady-state costs C_p
1. Planner's allocation	4.52	3.04	22.1	0.5	22.6	4.83
2. Symmetric equilibrium	11.8	1.36	7.9	5.5	13.5	7.43
3. Liability π^* for data breaches	6.75	2.89	14.8	0.5	15.2	5.19
4. Regulated security	4.91	3.04	20.3	0.5	20.8	4.84
5. Regulated data collection	4.52	1.84	21.6	5.5	27.1	5.21

Explanatory notes

Columns 1 and 2 of the Table give the numerical values of the allocation (d, s) in each case.

Columns 3, 4, and 5 display the rates of unskilled, skilled, and total identity theft.

Since $uF(1-F)$ is normalized to one in the examples, identity theft rates in columns 3-5 do not represent gross identity theft rates, but instead represent the percentage of frauds who are successful at impersonation.

Allocations are welfare ranked according to the planner's costs C_p (see equation (8)), displayed in column 6.

References

- Anderson, K.B., Durbin, E. Salinger, M.A., 2008. Identity theft. *Journal of Economic Perspectives* 22, 171-192.
- Anderson, R. Moore, T., 2006. The economics of information security. *Science* 314, 610-613.
- Boyd, J.H., Prescott, E.C., 1987. Dynamic coalitions: engines of growth. *American Economic Association Papers and Proceedings*, 63-67.
- Camera, G., Li, Y., 2008. Another example of a credit system that co-exists with money. *Journal of Money, Credit, and Banking* 40, 1295-1308.
- Cavalcanti, R. and E. Nosal, 2007. Counterfeiting as private money in mechanism design. Working paper, Federal Reserve Bank of Cleveland.
- Caruso, D., 2007. Securing very important data: your own. *New York Times*, October 7.
- Chandler, J.A., 2008. Negligence liability for breaches of data security. *Banking and Finance Law Review* 23: 223-247.
- Cheney, J., 2004. "Identity Theft: Where Do We Go From Here?" Payment Cards Center Conference Center, Federal Reserve Bank of Philadelphia.
- Coggeshall, Stephen, 2007. ID Theft knows no boundaries. *eCommerce Times*, April 13.
- Dow Jones and Company Inc., 2008a. New payment card data mantra is "Don't need it, don't store it." *Wall Street Journal*, September 16.
- Dow Jones and Company Inc., 2008b. Data breaches surpass 2007 level, but businesses rarely are penalized. *Wall Street Journal*, September 9.
- Douglas, D.D., 2008. Merchant liability for payment card security breaches. *Electronic Banking Law & Commerce Report* 13, 1-7.
- Experian, 2006. *PreciseID: An integrated approach to the world of identity risk management*. Available online at www.ftc.gov/bcp/workshops/techade/pdfs/Kirshbaum1.pdf.
- Gordon, G.R., Rebovich, D.J., Choo, K.-S., Gordon, J.B., 2007. Identity fraud trends and patterns: Building a data-based foundation for proactive enforcement. Working Paper, Center for Identity Management and Information Protection, Utica College.
- Green, E. J., Weber, W., 1996. Will the new \$100 bill decrease counterfeiting? *Federal Reserve Bank of Minneapolis Quarterly Review* 20(3), 3-10.
- Greene, M.N., 2009. Divided we fall: Fighting payments fraud together. *Federal Reserve Bank of Chicago Economic Perspectives (First Quarter)*, 37-42.

- Grossklags, J., Christin, N., Chuang, J., 2008. Security investment (failures) in five economic environments: A comparison of homogeneous and heterogeneous user agents. Accessed online at weis2008.econinfosec.org/papers/Grossklags.pdf.
- Javelin Research, 2008. 2008 Identity Fraud Survey Report. Available online at www.javelinstrategy.com.
- Kahn, C. M., McAndrews, J., Roberds, W., 2005. Money is privacy. *International Economic Review* 46, 377-400.
- Kahn, C. M. Roberds, W., 2008. Credit and identity theft. *Journal of Monetary Economics* 55, 251-264.
- Keitel, P., 2008. Legislative responses to data breaches and information security failures. Payment Card Center Discussion Paper,, Federal Reserve Bank of Philadelphia.
- Kirshbaum, M.D., 2006. Protecting Against Fraud in the Next TechAde. Available online at www.ftc.gov/bcp/workshops/techade/pdfs/presentations/kirshbaum.pdf.
- Kiyotaki, N., Wright, R., 1989. On money as a medium of exchange. *Journal of Political Economy* 97, 927-954.
- Kocherlakota, N. R., 1998. Money is memory. *Journal of Economic Theory* 81, 232-251.
- Kultti, K., 1996. A monetary economy with counterfeiting. *Journal of Economics* 63, 175-186.
- LoPucki, L., 2001. Human identification theory and the identity theft problem. *Texas Law Review* 80, 89-136.
- LoPucki, L., 2003. Did privacy cause identity theft? *Hastings Law Journal* 54, 1277-1298.
- Martin, A., Orlando, M., Skeie, D., 2008. Payment networks in a search model of money. *Review of Economic Dynamics* 11, 104-132.
- McGrath, J.C. Kjos, A., 2006. Information security, data breaches, and protecting cardholder information: facing up to the challenges. Payment Cards Center Conference Summary, Federal Reserve Bank of Philadelphia.
- Monnet, C., 2005. Counterfeiting and inflation. Working paper, European Central Bank.
- Monnet, C. Roberds, W., 2008. Optimal pricing of payments services. *Journal of Monetary Economics* 55, 1428-1440.
- Nosal, E. Wallace, N., 2006. A model of the (threat of) counterfeiting. *Journal of Monetary Economics* 54, 994-1001.
- Ponemon Institute, LLC, 2006. 2006 Annual study: Cost of a data breach. Available online at www.computerworld.com/pdfs/PGP_Annual_Study_PDF.pdf.

- Schreft, S. L., 2007. Risks of identity theft: can the market protect the payment system? Federal Reserve Bank of Kansas City Economic Review (Fourth Quarter), 5-40.
- Solove, D., 2003. Identity theft, privacy, and the architecture of vulnerability. *Hastings Law Journal* 54, 1227-1253.
- Solove, D., 2004. The new vulnerability: data security and personal information. Working paper, George Washington University Law School.
- Swartz, J., Acohidio, B., 2007. Who's guarding your data in the cybervault? ChoicePoint re-deemed itself but not all brokers as careful. *USA Today*, April 2.
- Swire, P. P., 2003. Efficient confidentiality for privacy, security and confidential business information. *Brookings-Wharton Papers on Financial Services*, 273-310.
- Synovate, 2007. Federal Trade Commission—2006 Identity Theft Report. Available online at www.ftc.gov.
- Varian, H., 1998. Markets for information goods." Available online at people.ischool.berkeley.edu/~hal/Papers/japan/.
- Varian, H., 2004. System reliability and free riding. Available online at people.ischool.berkeley.edu/~hal/Papers/2004/reliability.
- United States Government Accountability Office, 2007. Personal information: Data breaches are frequent, but evidence of resulting identity theft is Limited; However, the full extent is unknown. Report GAO-07-737.
- Williamson, S.D., 2002. Private money and counterfeiting. *Federal Reserve Bank of Richmond Economic Quarterly* 88(3), 37-57.

Appendices (supplementary material/ not for publication)

Appendix A. Transactions in the model

This appendix provides more detail on how transactions occur in the model.

Background

Buyers and sellers are matched according to a simple search process. The search process is similar to that employed in standard first-generation search models (Kiyotaki and Wright 1989), but differs slightly in that it forces every possible type of match to occur within a finite time interval. This feature is convenient for the analysis above because it separates *fraud risk* (the risk that an agent engages in a transaction with no intent to repay, which is the focus of the paper) as opposed to *credit risk* (the risk that a known agent cannot repay). Under the matching specification described below, an agent's fraudulent intent is always revealed, once a certain amount of time has passed. Agents' decision problems can then be reduced to a sequence of static decision problems, which reduces model complexity.

It is clear from credit industry discussions (e.g., Experian 2006, Greene 2009) that the separation of fraud and credit risk represents an abstraction. In practice, there is always some overlap of these types of risk. Consider the case where a person applies for a credit card, receives the card, uses the card to make purchases, and then never makes a payment on the bill. Because the cardholder's income, identity, and inclination towards fraudulent activity are not perfectly known to the card issuer, it is not always clear whether such a loss should be classified as a fraud loss or a credit loss. A cardholder may fraudulently claim to have been defrauded as a way of evading credit limits, further confounding credit and fraud risk. Nonetheless, it is customary

within the credit industry to conceptually (and statistically) separate these two types of risk. A consumer who applies for a credit card, for example, may be assigned an “identity risk score” as well as the more familiar credit score.

Matching specification

Agents in the model are matched according to their *types*. It is convenient to think of an agent’s type as his “location,” although the model does not rely on geography. Within each group G_i , types are distributed uniformly over the unit interval. There is a unit measure of agents of each type. Legitimate agents reside only on a measurable subset of locations Ω , and frauds reside at locations Ω^c , where $\mu(\Omega^c) = F$. At the end of each discrete period, a randomly selected subset of types vanish and are replaced with agents of the same type. The probability of replacement is β for both frauds and legitimate agents.

Agents within each group wish to consume the goods produced by all other types of agents of the same group. Time begins at date $t = 0$. During the initial interval $t \in [0, 1)$, nondurable goods of type y , $y \in [0, 1)$, are available for purchase and consumption at time y , when each type- y agent can supply a unit measure of good y . Intuitively, potential consumers of type $y' \neq y$ “journey” to location y to purchase and consume good y . This process is repeated during subsequent unit intervals; i.e., at any time $t \geq 0$, goods of type $y(t) \equiv t - \llbracket t \rrbracket$ are available for purchase and consumption, where $\llbracket t \rrbracket$ denotes the integer part of t .

Over all times $t \geq 0$, production within group i imposes an instantaneous disutility of $mc\delta(y(t) - y)dt$ on type- y agents, where $c > 0$, δ is Dirac’s delta function, and m is the measure of goods each agent supplies. For type- y' agents, where $y' \neq y$, time t consumption of one unit

of a type- y good yields instantaneous utility $u dt$, where $u > c > 0$. At each time t , potential consumers of type $y' \neq y(t)$ are randomly matched with one (and only one) producer within the same group of type $y(t)$, with i.i.d. matching over time, so that all transactions are between agents without any previous contact.

Trade among agents within a group is facilitated by a central authority (or “court”) with three limited and specific powers. First, the central authority can observe an agent’s actions *as a producer* (i.e., whether an agent has supplied goods during a time interval $[0,1)$, $[1,2)$, ...). Second, at discrete dates $n = 0, 1, 2, \dots$, the court can publicly announce the observed action. Third, the court can, when making this announcement, impose a nonpecuniary penalty of $X > 0$ utils on an agent who has refused to supply a good, *provided that the agent can be identified*.

Sustaining exchange with costless identification

As described above, agents in each group form a transactions club. Club membership entitles the agent to a (flow) unit of a consumption good from any other club member in return for agreeing to provide his own type of good to other club members, at some point during each unit interval of time. At subsequent discrete dates $n = 1, 2, \dots$, the center publicly announces the default of any club members who have not supplied goods and imposes penalty on nonproducers (a penalty of X utils) who are then excluded from the club. Membership in each club subsequently is opened to newborn agents.

Suppose that all legitimate agents of group i and type $y' \neq y$ decide to join club i , and that frauds do not. For a legitimate agent of type $y \in [0,1)$ in group i , the value of club membership during the interval $t \in [n, n+1)$ is given by

$$\int_{y' \in \Omega} (u - mc\delta(y - y')) dy' = u(1 - F) - mc \quad (22)$$

for $n = 0, 1, 2, \dots$. Market clearing requires $m = 1 - F$. Hence, from (22), if all legitimate agents of group i join club i , then the steady-state value of club membership is given as V in equation (1).

Ongoing membership in the club requires that a type- y agent be willing to supply a unit measure of goods at time $n + y$. This requires that the disutility of producing goods, combined with the disutility of the penalty X , be less than the value of continued club membership, i.e.,

$$c - X \leq V, \quad (23)$$

which is the same as

$$(1 - F + r)c \leq (1 - F)u + rX. \quad (24)$$

Under condition (24), no legitimate agent who has joined a club ever has an incentive to defect.

If, in addition,

$$u < X, \quad (25)$$

no fraud ever has an incentive to join the club. Since $c < u$, (24) is implied by (25). It follows that under (25), an equilibrium exists in which all legitimate agents join the transactions club formed by agents in their group, and all frauds remain outside the club.

Sustaining exchange with costly identification

When identification of agents is costly the steady-state value of a legitimate agent's membership in club i is given by V_i^f in (7). In this case, V_i^f must satisfy three conditions for exchange to occur through the clubs.

1. *Individual rationality*: a legitimate agent prefers joining a club to autarky. This requires

$$0 \leq V_i^f; \quad (26)$$

2. *No defection*: legitimate agents in each club have an incentive to produce goods for other club members. This requires

$$c - X \leq V_i^f; \quad (27)$$

Note that under (25), (27) is redundant given (26).

3. *No exclusion*: a club has an incentive to admit new generations of members. This requires

$$\underline{V} \leq V_i^f, \quad (28)$$

where \underline{V} is the value of maintaining the club without admitting new members, i.e.,

$$(1+r)\underline{V} = (u-c)(1-F) \sum_{n=0}^{\infty} \beta^{2n} = \frac{(u-c)(1-F)}{1-\beta^2}. \quad (29)$$

A steady-state allocation is *incentive compatible* if (26) and (28) are satisfied for both clubs.

Appendix B: Proofs of Propositions 1-5

Proof of Proposition 1.

The proof proceeds in four steps. First, we show that any solution (d, s) to first-order conditions (14) and (18) at equality represents a locally optimal and unique response by each club when the other club plays (d, s) . Second, we first show that under the hypotheses of the Proposition, there is only one such solution (d^*, s^*) . Third, we verify that there is no equilibrium with $s = 0$. Fourth, we show that (d^*, s^*) is incentive compatible.

Step 1. First-order conditions for club i 's problem are given in (14) and (18). Second-order conditions are given by

$$\frac{2\Phi(s_j)(c+L)}{d_i^3} + \frac{2(1-\Phi(s_j))(c+L)}{(d_i - \eta d_j)^3} > 0, \quad (30)$$

$$\frac{\beta B \Phi''(s_i)}{d_j - \eta d_i} < 0, \quad (31)$$

$$\left[\frac{2\Phi(s_j)(c+L)}{d_i^3} + \frac{2(1-\Phi(s_j))(c+L)}{(d_i - \eta d_j)^3} \right] \left[\frac{\beta B \Phi''(s_i)}{(d_j - \eta d_i)} \right] + \left[\frac{\beta B \Phi'(s_i)}{(d_j - \eta d_i)^2} \right]^2 < 0. \quad (32)$$

Conditions (30) and (31) are readily seen to hold when $(d_i, s_i) = (d_j, s_j)$. Sufficient conditions for (32) to hold are symmetry and $\beta B < 2(c+L)$, which is implied by condition (b) of the Proposition.

Step 2. Rewrite first-order condition (18) at equality as

$$d = D(s) \equiv \frac{uF(1-F)\beta B\phi(1-\Phi(s))}{\ell(1-\eta)}. \quad (33)$$

Substituting (33) into (14), imposing symmetry, and rearranging gives the following quadratic equation

$$Q(z) \equiv A_0(1-z) + A_1z + A_2z^2 = 0, \quad (34)$$

where $z = 1 - \Phi(s)$ and

$$A_0 = c + L, \quad (35)$$

$$A_1 = \frac{c + L - \beta B \eta}{(1 - \eta)^2}, \quad (36)$$

$$A_2 = -kuF(1-F) \left(\frac{\beta B \phi}{\ell(1-\eta)} \right)^2. \quad (37)$$

From the above, $Q(0) = A_0 > 0$ and $Q(1) = A_1 + A_2 < 0$ under conditions (a) and (b) of the Proposition. $Q(z)$ therefore has a unique root $z^* \in (0, 1)$; in particular, $z^* =$

$$\frac{c + L - \beta B \eta - (1 - \eta)^2(c + L) + \sqrt{\left(c + L - \beta B \eta - (1 - \eta)^2(c + L) \right)^2 + 4(c + L)kuF(1-F) \left(\frac{\beta B \phi}{\ell} \right)^2}}{2kuF(1-F) \left(\frac{\beta B \phi}{\ell} \right)^2}. \quad (38)$$

Now define

$$(d^*, s^*) = \left(D\left(\Phi^{-1}(1 - z^*)\right), \Phi^{-1}(1 - z^*) \right). \quad (39)$$

By construction, (d^*, s^*) satisfies (14) and (18) under symmetry, and $s^* > 0$.

Step 3. From the discussion in the text, there can be no equilibrium with $s = 0$ if (20) is violated, which occurs for $\phi > 0$ sufficiently large (condition (a) of the Proposition).

Step 4. To show incentive compatibility, suppose initially that $F = 0$, so that $V_i^f = V$

Then the individual-rationality and no-exclusion conditions are clearly satisfied with strict inequality for β sufficiently close to unity (condition (d) of the Proposition). Now, for $F > 0$, let

K, k , and ℓ approach zero (Proposition condition (c)); more specifically let $\|(K, k, \ell)\| < \theta$

where $\theta > 0$ and $\|\cdot\|$ is the sup norm. Then it can be shown that as $\theta \rightarrow 0$, d^* and s^* as de-

fined in (39) are bounded by $\theta^{-1/2}$ and $-\ln \theta$, respectively. This, in turn, implies that $V_i^f \rightarrow V$ as $\theta \rightarrow 0$, as fraud rates and all costs of fraud deterrence are driven to zero. Hence, by continuity, incentive compatibility must hold for K, k , and ℓ all positive and sufficiently small.

Proof of the Corollary to Proposition 1.

Begin by solving for (d_p, s_p) . Rewrite first-order condition (19) as

$$d = \underline{D}(s) \equiv \frac{uF(1-F)(\eta(c+L) + \beta B)\phi(1-\Phi(s))}{\ell(1-\eta)}. \quad (40)$$

Substituting (40) into condition (16) and rearranging gives the following quadratic equation

$$\underline{Q}(z) \equiv \underline{A}_0(1-z) + \underline{A}_1z + \underline{A}_2z^2 = 0, \quad (41)$$

where $z = 1 - \Phi(s)$ and

$$\underline{A}_0 = c + L, \quad (42)$$

$$\underline{A}_1 = \frac{c + L + \beta B}{1 - \eta}, \quad (43)$$

$$\underline{A}_2 = -kuF(1-F) \left(\frac{\phi(\eta(c+L) + \beta B)}{\ell(1-\eta)} \right)^2. \quad (44)$$

Proceeding as in the proof of the Proposition, $\underline{Q}(z)$ has a unique root z_p in $(0,1)$ for ϕ sufficiently large (Proposition condition (a)). In particular, $z_p =$

$$z_p = \frac{(1-\eta) \left(1 + \sqrt{1 + 4(c+L)kuF(1-F) \left(\frac{\phi}{\ell} \right)^2} \right)}{2kuF(1-F) \left(\frac{\phi}{\ell} \right)^2 (\eta(c+L) + \beta B)}. \quad (45)$$

The planner's allocation is then given as $(d_p, s_p) = \left(\underline{D}(\Phi^{-1}(1-z_p)), \Phi^{-1}(1-z_p) \right)$.

Second-order conditions for the planner's problem are given by

$$\frac{2(c+L)\Phi(s)}{d^3} + \frac{2(c+L+\beta B)(1-\Phi(s))}{d^3(1-\eta)} > 0, \quad (46)$$

$$\left[-\frac{c+L}{d} + \frac{c+L+B}{d(1-\eta)} \right] \Phi''(s) > 0, \quad (47)$$

$$2 \left[\frac{(1-\eta)(c+L)\Phi(s) + (c+L+B)(1-\Phi(s))}{d^4(1-\eta)^2} \right] (\eta(c+L)+B)\Phi''(s) + \frac{(\eta(c+L)+B)^2 (\Phi'(s))^2}{d^4(1-\eta)^2} < 0, \quad (48)$$

which can be shown to hold for all positive d and s and hence for (d_p, s_p) .

Proof of Proposition 2.

(Sketch). Solutions for d^* and d_p are given in (56) and (57), respectively; solutions for s^* and s_p are given as $-\phi^{-1} \ln[\text{RHS}(38)]$ and $-\phi^{-1} \ln[\text{RHS}(45)]$. The Proposition follows from straightforward differentiation of these expressions.

Proof of Proposition 3.

Part (a). From (38) and (45), both z^* and z_p are clearly decreasing in η , so skill thresholds s^* and s_p must be increasing in η .

Part (b). From (38) and (45), as $\eta \rightarrow 1$, $z_p \rightarrow 0$ while z^* converges to

$$z \equiv \frac{c+L-\beta B + \sqrt{(c+L-\beta B)^2 + 4(c+L)kuF(1-F)\left(\frac{\beta B\phi}{\ell}\right)^2}}{kuF(1-F)\left(\frac{\beta B\phi}{\ell}\right)^2} > 0. \quad (49)$$

Hence, as $\eta \rightarrow 1$, $s^* \rightarrow \bar{s} = \Phi^{-1}(1 - \underline{z})$ while s_p diverges.

Proof of Proposition 4.

To analyze d_p and d^* , we first derive closed-form expressions for these quantities. To solve for d^* , invert $D(s)$ in (33) and substitute into first-order condition (14) to obtain the following condition in d :

$$R(d) = R_0 + R_1 d + R_2 d^2 = 0, \quad (50)$$

where

$$R_0 = uF(1 - F)(c + L), \quad (51)$$

$$R_1 = \ell \left[\frac{(c + L - \beta\eta B) - (1 - \eta)^2 (c + L)}{\beta B \phi (1 - \eta)} \right], \quad (52)$$

$$R_2 = -k. \quad (53)$$

Similarly, to solve for d_p , invert $\underline{D}(s)$ in (40) and substitute into the planner's first-order condition (16) to obtain the condition

$$\underline{R}(d) = \underline{R}_0 + \underline{R}_1 d + \underline{R}_2 d^2 = 0, \quad (54)$$

where $\underline{R}_0 = R_0$, $\underline{R}_2 = R_2$, and

$$\underline{R}_1 = \frac{\ell}{\phi}. \quad (55)$$

Evidently, d^* and d_p may be expressed as (positive) roots of $R(d)$ and $\underline{R}(d)$, respectively. In particular, d^* is given by

$$(2k(1-\eta))^{-1} \times \left[\left(\frac{\ell}{\phi} \right) \left(\frac{(c+L-\beta B\eta) - (1-\eta)^2(c+L)}{\beta B} \right) + \sqrt{\left(\frac{\ell}{\phi} \right)^2 \left(\frac{(c+L-\beta B\eta) - (1-\eta)^2(c+L)}{\beta B} \right)^2 + 4kuF(1-F)(c+L)(1-\eta)^2} \right], \quad (56)$$

and

$$d_p = (2k)^{-1} \left[\left(\frac{\ell}{\phi} \right) + \sqrt{\left(\frac{\ell}{\phi} \right)^2 + 4kuF(1-F)(c+L)} \right]. \quad (57)$$

Part (a). From (57), d_p does not depend on η . From (56), d^* grows as

$$\tilde{d} = (k(1-\eta))^{-1} \left[\left(\frac{\ell}{\phi} \right) \left(\frac{(c+L-\beta B\eta)}{\beta B} \right) \right], \quad (58)$$

as $\eta \rightarrow 1$, which is increasing in η for $c+L > \beta B$.

Part (b). From (58), $\tilde{d} \rightarrow \infty$ as $\eta \rightarrow 1$, whence d^* also diverges.

Additional discussion of Proposition 4 and Proposition 2(c).

The planner's first-order condition in data length d (equation (16)) equates the marginal cost of data collection k to the sum of its marginal benefits, i.e., a reduction in the costs of unskilled fraud

$$uF(1-F) \frac{(c+L)\Phi(s)}{d^2}, \quad (59)$$

plus a reduction in the cost of skilled fraud

$$uF(1-F) \frac{(c+L+\beta B)}{d(1-\eta)}. \quad (60)$$

Another way of expressing (16) is as follows

$$uF(1-F) \left[\frac{(c+L)}{d^2} + \frac{(\eta(c+L) + \beta B)(1-\Phi(s))}{d^2(1-\eta)} \right] = k. \quad (61)$$

Condition (61) decomposes the marginal benefits of data collection into a benefit that would accrue if all identity theft were unskilled,

$$uF(1-F) \frac{c+L}{d^2}, \quad (62)$$

plus an additional benefit that occurs when some identity theft is skilled,

$$uF(1-F) \frac{(\eta(c+L) + \beta B)(1-\Phi(s))}{d^2(1-\eta)}. \quad (63)$$

Under a constant hazard rate of hacking skills specification, expression (63) is proportional to the inverse of data length, multiplied by the marginal benefit to increased data security

$$\left(\frac{1}{\phi d} \right) \left(uF(1-F) \frac{(\eta(c+L) + \beta B)\Phi'(s)}{d(1-\eta)} \right) = \frac{\ell}{\phi d}, \quad (64)$$

where the last equality follows from the planner's first-order condition in data security (19).

Using (64), the planner's first-order condition in d may be rewritten as

$$uF(1-F) \left[\frac{(c+L)}{d^2} \right] + \frac{\ell}{\phi d} = k, \quad (65)$$

which does not involve B or η , i.e., the planner adjusts data security so that the marginal impact of data length on skilled identity theft (beyond that which would occur without data theft skills) does not depend on these parameters. The same result does hold for the Nash equilibrium, basically because the Nash clubs perceive both a higher marginal benefit to data collection (compare LHS (14) to LHS(16)) and a lower payoff to data collection (LHS (18) to LHS (19)).

Proof of Proposition 5.

The calculations in this section simplify notation by setting $uF(1-F) = 1$.

Part (a). From first-order condition (18), the rate of skilled identity theft in the symmetric equilibrium is

$$\frac{1 - \Phi(s^*)}{d^*(1-\eta)} = \frac{\ell(1 - \Phi(s^*))}{\beta B \Phi'(s^*)} = \frac{\ell}{\beta B \phi}. \quad (66)$$

Similarly, the rate of skilled identity theft in the planner's allocation can be calculated using (19):

$$\frac{1 - \Phi(s_p)}{d_p(1-\eta)} = \frac{\ell}{\phi[\eta(c+L) + \beta B]}. \quad (67)$$

Comparing (66) and (67), skilled identity theft must be lower under the planner's allocation.

Part (b). The rate of unskilled identity theft in the symmetric equilibrium is given by $\Phi(s^*)/d^*$. From the Propositions 3 and 4, $\Phi(s^*) \rightarrow \Phi(\bar{s}) > 0$ and $d^* \rightarrow \infty$ as $\eta \rightarrow 1$, implying that unskilled identity theft is driven to zero as $\eta \rightarrow 1$.

The rate of unskilled identity theft under the golden-rule allocation is given by $\Phi(s_p)/d_p$. From the proof of the Corollary to Proposition 1, $\Phi(s_p) \rightarrow 1$ as $\eta \rightarrow 1$ but d_p is positive and does not depend on η . Hence the rate of unskilled identity theft converges to $1/d_p > 0$ as $\eta \rightarrow 1$.

Part (c). The calculations in parts (a) and (b) show that, as $\eta \rightarrow 1$, $\rho(d^*, s^*) < \rho(d_p, s_p)$ iff

$$\frac{\ell}{\phi\beta B} < \frac{1}{d_p} + \frac{\ell}{\phi(c+L+\beta B)}. \quad (68)$$

Substituting for d_p from the proof of the Corollary to Proposition 1, inequality (68) reduces to

$$\frac{2\phi^2}{\ell + \sqrt{\ell^2 + 4(c+L)k\phi^2}} > \left(\frac{\ell}{k}\right) \frac{c+L}{\beta B(c+L+B)}, \quad (69)$$

which must hold for ℓ/k bounded and $k, \ell > 0$ sufficiently small.

Appendix C. Policy analysis

The calculations in this section simplify notation by setting $uF(1-F) = 1$.

1. Imposing liability for a breach

Suppose that penalty $\pi \leq \pi^*$ is in effect. Then in symmetric equilibrium, the clubs choose data length d_π , given by (56) where $B' = B + \pi$ replaces B , and security level s_π , given by $-\phi^{-1} \ln [\text{RHS}(38)]$, where again B' replaces B . From Proposition 2, $d_\pi' < 0$ and $s_\pi' > 0$ for η sufficiently close to one; hence $d_\pi < d^*$ and $s_\pi > s^*$.

Using the Chain Rule,

$$\frac{dC}{d\pi} = \frac{\partial C_p}{\partial d} d_\pi' + \frac{\partial C_p}{\partial s} s_\pi', \quad (70)$$

where

$$\frac{\partial C_p}{\partial d} = k - \left[\frac{(c+L)\Phi(s)}{d^2} + \frac{(c+L+\beta B)(1-\Phi(s))}{d^2(1-\eta)} \right], \quad (71)$$

$$\frac{\partial C_p}{\partial s} = \ell - \left(\frac{\eta(c+L) + \beta B}{d(1-\eta)} \right) \Phi'(s). \quad (72)$$

(cf. the planner's first-order conditions (16) and (19)). But in equilibrium, d_π and s_π must satisfy first-order conditions (14) and (18) where B is replaced with B' , from which it can be shown that

$$\frac{\partial C_p}{\partial d} \geq 0 \quad \text{and} \quad \frac{\partial C_p}{\partial s} < 0 \quad (73)$$

for $(d, s) = (d_\pi, s_\pi)$. Since $d_\pi' < 0$ and $s_\pi' > 0$, it follows from (70) that $dC/d\pi < 0$.

2. Analysis of the regulator's problem when the regulator only sets skill thresholds s

As in the proof of Proposition 1, let $z = 1 - \Phi(s)$. The problem of a regulator who only chooses s is equivalent to the following: minimize steady-state fraud costs C_p over $z \in (0,1)$, i.e., minimize

$$\frac{(c+L)(1-z)}{d} + \frac{(c+L+\beta B)z}{d(1-\eta)} + kd - \frac{\ell}{\phi} \ln z, \quad (74)$$

subject to the clubs' first-order condition (15), which we write as $d = G(z)$ where

$$G(z) = \frac{1}{(1-\eta)\sqrt{k}} \left[(c+L)(1-\eta)^2(1-z) + (c+L-\beta\eta B)z \right]^{\frac{1}{2}}. \quad (75)$$

This regulator's problem may be compared to the planner's problem, which is equivalent to minimizing (74) over $z \in (0,1)$ subject to (17), which we write as $d = P(z)$ where

$$P(z) = \frac{1}{\sqrt{(1-\eta)k}} \left[(c+L)(1-\eta)(1-z) + (c+L+\beta B)z \right]^{\frac{1}{2}}. \quad (76)$$

Substituting (75) into (74) and simplifying, the regulator's problem is to minimize

$$k \frac{(P(z))^2}{G(z)} + kG(z) - \frac{\ell}{\phi} \ln z. \quad (77)$$

This contrasts with the planner's problem, which, substituting (76) into (74), simplifies to the following: minimize

$$k \frac{(P(z))^2}{P(z)} + kP(z) - \frac{\ell}{\phi} \ln z = 2kP(z) - \frac{\ell}{\phi} \ln z. \quad (78)$$

The first-order condition for the regulator's problem is

$$k \left[2P(z)P'(z) + G'(z) \left(1 - \left(\frac{P(z)}{G(z)} \right)^2 \right) \right] - \frac{\ell}{\phi z} = 0, \quad (79)$$

which after some manipulation can be written as

$$\left[\left(\frac{\eta(c+L) + \beta B}{1-\eta} \right) z - \frac{\ell}{\phi} \right] (G(z))^3 = \frac{((c+L)\eta(2-\eta) - \beta\eta B)}{2(1-\eta)^2} z \left[(G(z))^2 - (P(z))^2 \right]. \quad (80)$$

Squaring both sides of (80) to eliminate radicals, a solution to the regulator's problem requires finding the roots of a fifth degree polynomial, a problem for which there is no general analytical solution. Hence this problem is analyzed numerically.

3. Analysis of the regulator's problem when the regulator only sets data length d

A regulator who can only determine data length sets d to minimize C_p subject to the clubs' first-order condition in s , which in symmetric equilibrium is given by (18). Using (18) and imposing symmetry, we can eliminate s and simplify the regulator's problem to the following: choose d to minimize

$$C_p = kd - \frac{\ell}{\phi} \ln d + \frac{c+L}{d} + \langle \text{constant terms} \rangle, \quad (81)$$

which has solution $d_c = d_p$.

Evaluating (18) at $d = d_p$ and comparing to (19), it follows that $s < s_p$ must hold in the regulated equilibrium. From (18) and the fact that $d_p < d^*$ (Proposition 4) as $\eta \rightarrow 1$, it follows that $s > s^*$. Hence, in the regulated equilibrium, security effort s is intermediate between the security effort of the (unregulated) symmetric equilibrium s^* and the security effort chosen by the planner s_p .

Appendix D: Extension with endogenous network size

An alternative method for controlling data breaches is to allow for the sharing of data residing in the databases of the two separate clubs (networks). In the model, sharing data across clubs eliminates the incentive for data breaches because any stolen identifying information duplicates existing information and is automatically revealed as fraudulent. Exchanging data across clubs can thus be beneficial even though agents in each club never interact in commerce with agents of the other group.

In principle, data sharing could be implemented in a number of ways. LoPucki (2001) proposes the creation of a government agency that would manage a consolidated database of PID. Inclusion in the database would be optional. This appendix considers an alternative channel for data sharing, which is the voluntary preference of agents in the two groups to share data across groups. This is done by a slight generalization of the environment studied above.

In this generalized environment, agents have the option of transacting through a single club or dual clubs (one for each group of agents). The two groups of agents may be of different size, i.e., let $\mu(G_A) = \mu_A$ and $\mu(G_B) = \mu_B$. If all legitimate agents decide to form a single club, no data breaches occur in equilibrium, so the club simply compiles data of length d on all its members to maximize the average per-capita net benefit of legitimate club membership. That is, the single club chooses d to maximize (cf. expression (7))

$$V_s = \left(\frac{1}{r}\right) \times \left[(\underline{u} - \underline{c})(1-F) - (1-\beta)\underline{K} - \underline{k}d - \mu_A \frac{u_A F(1-F)}{d} (c_A + L) - \mu_B \frac{u_B F(1-F)}{d} (c_B + L) \right], \quad (82)$$

where the underlines indicate average values, i.e., $\underline{u} = \mu_A u_A + \mu_B u_B$ etc. Let d_s denote the choice of data length that maximizes (82), and let $V_{A,s}$ ($V_{B,s}$) denote the steady-state value of legitimate

club membership for agents of group G_A (G_B) when PID of length d_s is collected. A *steady-state equilibrium with a single club* exists when the following incentive constraints (analogous to (26), (27), and (28)) are satisfied

1. *Individual rationality*, $0 \leq V_{i,s}$ for $i = G_A, G_B$;
2. *No defection*, $c_i - X \leq V_{i,s}$ for $i = G_A, G_B$;
3. *No exclusion*, $\underline{V}_i \leq V_{i,s}$ for $i = G_A, G_B$, where \underline{V}_i is the value of maintaining the club without admitting new members.

If, as in Section 3 above, agents' preferences are symmetric across groups, it is immediate that an equilibrium with a single club exists whenever a symmetric steady-state equilibrium exists. Moreover, the single-club equilibrium dominates the dual-club equilibrium. For any value of d chosen by the dual clubs, the single club can do better with this same data because the single club's database provides a greater benefit in terms of fraud reduction (all frauds must now attempt the more costly unskilled identity theft) at a lower cost (since the single club incurs no costs of securing data against breaches and no breach costs).

In the absence of unanimity, however, conflicts of interest can arise as to the amount of data the single club should compile and retain. Sufficient heterogeneity in preferences can limit potential efficiency gains achievable through voluntary consolidation of data. To demonstrate this point, consider the following parameterization of the model. Suppose that the per-unit physical cost of compiling and storing data is negligible, so that the cost parameter k essentially reflects intangible costs associated with the loss of privacy. Agents in the two groups G_A and G_B have identical preferences, except that agents in group G_A are indifferent to the privacy of their stored personal data ($k_A = \varepsilon$, where $\varepsilon > 0$ is arbitrarily small), while agents in group G_B place a

higher value on confidentiality ($k_B > k_A$). The two groups are of unequal size: group G_A has unit measure as before, while group G_B has measure $\mu_B = \mu > 0$.

Suppose that agents in the two groups decide to form a single club. The optimal data length for the single club is given by (cf. equation (12))

$$d_s = \sqrt{\frac{uF(1-F)(c+L)}{\underline{k}}}, \quad (83)$$

and from (82), the equilibrium per-capita net benefit of club membership for an agent of group i is

$$V_{i,s} = \left(\frac{1}{r}\right) \left[(u-c)(1-F) - (1-\beta)K - \left(k_i + \sqrt{\frac{\underline{k}}{1+\mu}} \right) \sqrt{\frac{uF(1-F)(c+L)}{\underline{k}}} \right], \quad (84)$$

for $i = G_A, G_B$.

Now suppose each group decides to form its own club. In this case, agents in group A are willing to surrender virtually limitless amounts of personal information to club G_A , which effectively precludes the possibility of fraudulent entry into their club. Once assembled, however, club G_A 's database is subject to data breaches committed by skilled frauds seeking access to club G_B . Thus, with dual clubs, club G_A chooses d_A arbitrarily large as $k_A \rightarrow 0$ and chooses s_A to maximize

$$V_{A,d} = \left(\frac{1}{r}\right) \left[(u-c)(1-F) - (1-\beta)K - \ell s_A - \mu F (1 - \Phi(s_A)) \beta B \right]. \quad (85)$$

Differentiating (85), for sufficiently large ϕ , the optimal skill threshold for club G_A is given by

$$s_A = \phi^{-1} \ln((\mu F \beta B \phi) / \ell), \quad (86)$$

which implies that, with dual clubs, the equilibrium net benefit of membership in club G_A is

$$V_{A,d}^* = \left(\frac{1}{r}\right) \left[(u-c)(1-F) - (1-\beta)K - \left(\frac{\ell}{\phi}\right) \left(\ln\left(\frac{\ell}{\mu FB\phi}\right) + 1 \right) \right]. \quad (87)$$

Because the PID stored in club A 's database is so extensive, club G_B cannot control its rate of skilled identity theft: any amount of data d_B that club G_B might require for entry can be stolen from club G_A with sufficient skill. Knowing this, club G_B chooses a data length d_B that balances the benefits of reduced unskilled identity fraud against the costs associated with the loss of privacy. This data does not need to be well secured because data stolen from club G_B 's database is insufficient to gain access to club G_A ; that is, in the limit there are no breach costs for club G_B . Hence, with dual clubs, club G_B 's problem reduces to choosing d_B to maximize

$$V_{B,d} = \left(\frac{1}{r}\right) \left[\begin{aligned} &(u-c)(1-F) - \\ &(1-\beta)K - k_B d_B - \frac{uF(1-F)}{d_B} (c+L) - F(1-\Phi(s_A))(c+L) \end{aligned} \right]. \quad (88)$$

Differentiating (88) and solving yields

$$d_B = \sqrt{\frac{uF(1-F)(c+L)}{k_B}}. \quad (89)$$

Using (86) and (89), the equilibrium per-capita net benefit of membership in club G_B in the case of dual clubs can be expressed as

$$V_{B,d}^* = \left(\frac{1}{r}\right) \times \left[(u-c)(1-F) - (1-\beta)K - \sqrt{k_B uF(1-F)(c+L)} - \left(\frac{\ell}{\phi}\right) \left(\frac{c+L}{FB} \right) \right]. \quad (90)$$

For this parameterization, the comparison between the single club and dual clubs is stated in the following:

Proposition 6. *Suppose that groups G_A and G_B have heterogeneous preferences over the privacy of stored data ($k_b > k_a$ arbitrarily small) and that the measure of each group is*

$\mu_A = 1 > \mu_B > 0$. Then for ϕ sufficiently large and $K, k_B, \ell, \mu_B > 0$ sufficiently small,

- a) A steady-state equilibrium exists for both the single club and dual clubs;*
- b) Legitimate agents in both groups are better off under dual clubs than under the single club.*

Proof. The proof of Part (a) follows that of Proposition 1. To show Part (b), let $\ell / \phi \rightarrow 0$. Then, comparing (84) and (87), $V_{A,d}^* > V_{A,s}$ for $\mu_B > 0$ sufficiently small. Comparing (84) and (90),

$V_{B,d}^* > V_{B,s}$ under the same conditions.

Intuitively, Proposition 6 says that, given sufficient heterogeneity, agents may prefer to tolerate a certain amount of data theft, as occurs under dual clubs, rather than attempt to eliminate the problem by forming a single club. Agents with a low value on privacy allow their club to compile large amounts of personal data because this deters fraud, even though this data is subject to breach and misuse. By contrast, agents who place a high value on privacy will tolerate a higher rate of identity theft, as the cost of keeping more of their PID private. Merging the two clubs can result in a level of personal data collection that seems excessive to the high-privacy group but insufficient to the low-privacy group.

More generally, Proposition 6 illustrates how heterogeneity can limit the efficiency gains from consolidation of PID. So long as this information is shared through voluntary associations (rather than mandatory participation in a single arrangement), disparate groups of agents in an

economy may prefer to sort into separate alliances with differing levels of personal privacy and data security. Clearly, heterogeneity can also limit efficiency gains attainable through other means as well. Regulatory limits on data collected, for example, might constrain groups who place low value on their privacy.