Data Breaches and Identity Theft

William Roberds
Federal Reserve Bank of Atlanta

Stacey L. Schreft*
The Mutual Fund Research Center, LLC

August 15, 2008

Abstract

This paper presents a monetary-theoretic model to study the implications of networks' collection of personal identifying data and data security on each other's incidence and costs of identity theft. To facilitate trade, agents join clubs (networks) that compile and secure data. 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) mandated limits on the amount of data collected, (2) mandated security levels, (3) reallocations of data-breach costs, and (4) data sharing through a merger of the networks.

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

JEL Codes: D83, E42, G28

*Roberds, Research Department, Federal Reserve Bank of Atlanta, 1000 Peachtree Street, Atlanta, GA 30309-4470, 404-498-8970, william.roberds@atl.frb.org; Schreft, The Mutual Fund Research Center, 7301 College Blvd., Ste. 220, Overland Park, KS 66210, 913-319-8167, sschreft@mutualfundstore.com. 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.

1. Introduction

More and more personal data is being collected as the cost of information technology falls. While collecting such data undoubtedly provides 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 cost the U.S. economy $61 billion in 2006. Even this estimate, however, omits many contributors to the true economic cost.[1]

Dollar estimates of the cost of identity theft do not by themselves indicate that too much identity theft is occurring. However, press accounts of data breaches suggest that personal identifying data (PID) is being stolen too frequently, and that the data thefts are unduly facilitating various kinds of identity theft.[2] This view is echoed in the legal literature on identity theft and data confidentiality.[3] There is also a general sense that "too much" PID is being collected, though some suggested policy fixes imply that more, not less, PID should be collected as a deterrent against its potential misuse.

Economists (economic theorists in particular) have remained relatively quiet on issues

[1] These estimates are derived in Schreft (2007). It is difficult to gauge the extent and direction of identity theft from available data. The Federal Trade Commission (FTC) has conducted surveys of consumers to determine the incidence of identity theft. A superficial reading of the FTC's 2006 survey, released November 2007, suggests that rates of identity theft might have stabilized in the last few years, but the FTC acknowledges that methodological changes to the 2006 survey make the survey's results noncomparable to those from earlier surveys, thus preventing the survey from being used to identify trends in the incidence of identity theft (Synovate, 2007). Javelin Strategy and Research conducted the survey in years when the FTC did not and in 2006, and made the same methodological changes in its 2006 survey as did the FTC. Hence Javelin's 2006 results are also noncomparable to its earlier survey results.

Regardless of methodology, the FTC and Javelin surveys, as well as any other surveys of consumers, at best yield estimates of acts of identity theft known to consumers. Acts of identity theft not known to consumers obviously cannot be captured by surveys of consumers. (Schreft 2007, footnotes 13-15)

[2] Recent examples include "To Fight Identity Theft, A Call for Banks to Disclose All Incidents," *New York Times*, March 21, 2007; "Who's guarding your data in the Cybervault? ChoicePoint redeemed itself but not all brokers as careful," *USA Today*, April 2, 2007; "Securing Very Important Data: Your Own," *New York Times,* October 7, 2007; and many others.

[3] See e.g., LoPucki (2001, 2003), Solove (2003, 2004), Swire (2003), and Chandler (2007).

regarding identity theft and data breaches.[4] Swire (2003) attributes this lack of interest to the commonly held belief among economists that information revelation generally promotes efficiency, leading economists to systematically overemphasize the costs and underestimate the benefits of data security. Reliance on economic theory can therefore lead to a serious underestimation of the efficient degree of data confidentiality, according to Swire.

Swire's argument is a challenge to economists to develop more precise notions of what constitutes an efficient level of PID accumulation and security. This paper is one response to this challenge. The formal model presented below uses contemporary monetary theory to evaluate the costs and benefits of amassing and securing PID as key elements of a credit-based transactions arrangement. This framework allows exploration of what is gained and lost through the accumulation, sharing, and theft of PID.

The application of monetary theory is fundamental to this task, as it explicitly delineates two key market frictions that might 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 economic benefit of a credit-based payment system derives from its ability to counteract these frictions, and sufficient knowledge of agents' identities is indispensable to the provision of this benefit—credit is impossible without knowing who the debtor is.

The environment in this paper extends the model of identity theft developed in Kahn and Roberds (2008) to incorporate the possibility of identity theft through data breaches. The paper begins by presenting a game-theoretic model of multiple credit-card networks. Card networks are modeled as club arrangements for the sharing of essential information for intertemporal trade: sufficient knowledge of agents' identities and credit histories. Each club must decide how much data on its members to assemble into a database, and each also must choose how

[4] Some relevant literature is discussed in Section 5 below.

thoroughly to secure its database. Collecting more PID imposes costs on card-network participants, but yields a benefit in terms of identifying the most casual, opportunistic, and simplistic attempts at fraudulent access to 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 deter 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.

The paper proceeds to compare the networks' data and security decisions to the decisions that a planner would implement. Under a mild set of technical conditions, this analysis confirms the popular wisdom concerning data breaches: in equilibrium, too much PID is collected, and the data is insufficiently secured. The paper then considers a number of regulatory remedies for this inefficiency.

The model environment is initially developed for networks of fixed size. An extension allows for networks of variable size. Merging networks internalizes the benefits of fraud deterrence and can reduce the scope for identity theft. For sufficiently heterogeneous preferences, however, it is shown that agents may prefer to separate into multiple networks, even when this facilitates identity theft through data breaches. This analysis, while exploratory, illustrates bounds on efficiency gains achievable from consolidation of PID.

In summary, the approach here allows for explicit calculation of the efficient levels of data accumulation and data security, and for straightforward evaluation of policies meant to attain efficiency. More generally, it offers an illustration of how any such calculation should balance the costs associated with 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. More extensive surveys are given in Schreft (2007) and Anderson et al. (2008).

It is constructive to begin by defining terms. Identity theft can take many forms in practice and need not involve data breaches. The Federal Trade Commission (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. An increasingly prevalent type of identity fraud is *fictitious* or *synthetic identity fraud*, in which a thief combines information taken from a variety of sources with invented information to create a new, fictitious identity (Schreft 2007). Synthetic identity theft is actually a type of new account fraud, with the new account being in the name of a real or fictitious person.[5] By one recent estimate, more than 80 percent of all new-account identity theft has occurred using synthetic identities (Coggeshall 2007). As will be clear below, new-account fraud is the type of identity fraud that occurs in the model.

Data breaches can facilitate either existing-account fraud (as when credit-card

[5] It has been noted that the payment-card industry uses some additional terminology in discussing identity theft. Cheney (2005) distinguishes *payment-card fraud*, which refers to the theft of information about an existing payment card and use of the information to make fraudulent card purchases, and *account-takeover fraud*, where the identity thief changes the address on an existing financial account, which allows the thief to more fully control the account and to deter capture longer. Both *payment-card fraud* and *account-takeover fraud* are cases of *existing account fraud* under the FTC definition.

information is stolen) or new-account fraud (as when PID is stolen).[6] There is no definitive estimate of how many cases of identity theft have resulted from data breaches. Certainly, data breaches are numerous and increasing: although no comprehensive surveys are available, the information-security website *Attrition.org* lists 326 reported data breach "incidents" for 2007, leading to the compromise of 162 million records of personal data, as compared to 11 reported incidents and 6 million compromised records in 2003. These figures are likely underestimates as many breaches are not reported.

Of course, a data breach is neither a necessary nor a sufficient condition for identity theft. Data can be stolen without being used for fraudulent purposes.[7] Nevertheless, there seems to be widespread recognition that data breaches can promote identity theft, particularly in its more costly and pernicious forms.[8]

The costs of identity theft must be weighed against the benefits provided by the availability of PID, which lie at the heart of modern credit-based systems of exchange. There are no direct estimates of these benefits, but the sheer volume and increasing popularity of services such as card-based payments indicates that these are substantial. In 2005 in the U.S. alone there were 43 billion card transactions worth $2.6 trillion (Bank for International Settlements (2007)).

3. The Model

3.1 Modeling choices

As discussed in the Introduction, the central policy issue concerning identity theft is

[6] 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.

[7] For example, the Javelin 2008 *Identity Fraud Survey Report* finds that 7 percent of consumers surveyed who knew how their identifying information was stolen reported a data breach as the culprit (Javelin 2008). However, year after year, a large majority of consumers surveyed by Javelin do not know how their identifying information was stolen.

[8] Such data breaches include the 2005 breach at TJX Companies, with an estimated total cost in the hundreds of millions of dollars (Schreft 2007).

whether, under current arrangements, PID is being efficiently collected and secured. Progress on this basic issue requires answering several difficult questions: Who should bear the costs of identity theft? Should there be limits on how much PID can be compiled and shared? Is there a "market failure"? Should criminal penalties for identity theft be increased? These questions, however, cannot be meaningfully addressed without reference to efficiency.

There are two obstacles to analyzing this issue. The first is that with modern information technology, knowledge of PID and control of access to it has been effectively transformed into a type of nonrival good, whose efficient allocation is bound to be less straightforward than that of standard, rival goods (Varian 1998, 2004). The second is that in the marketplace, these nonrival goods are provided through the interaction of many disparate parties (e.g., consumers, merchants, credit bureaus and other information brokers, credit-card issuers, financial intermediaries, and firms that provide transaction processing and information-security services) whose actions are subject to complex laws, regulations, and contractual obligations.

To shed some light on the relevant policy questions, the analysis below abstracts from the second difficulty to concentrate on the first. That is, in the model environment, PID is accumulated and shared through simple club arrangements. By forming and dividing across multiple clubs, agents can facilitate exchange in the presence of uncertainty about agents' identities. For concreteness, a club is referred to as a "credit-card network," and there is sufficient homogeneity so that each club *qua* network can be sustained through straightforward agreements among club members. Collecting and maintaining a database of personal information provides benefits to club members by assuring that debts will be repaid and deterring frauds. However, if a club's database is not adequately secured, it also can facilitate identity theft.

The model environment does not incorporate existing account fraud. This is done to maintain tractability and concentrate on the more costly varieties of identity theft, i.e., new account and synthetic identity theft. Existing account fraud is actually quite similar to counterfeiting, which already has been formally modeled and analyzed in the money literature (discussed in Section 5 below).

3.2 Basic environmental 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 known as the agent's *identity*. The dimension of this vector is sufficiently high as to be effectively infinite. An agent's identity is private information and never automatically revealed. Each agent is congenitally either a *legitimate agent* or a *fraud* (i.e., an identity thief).[9] F denotes the fraction of frauds in the population. The next subsection describes frauds in more detail, while this subsection further describes legitimate agents.

A measure one of legitimate agents is of type α, where α denotes the agents' production *types*, meaning the consumption good the agents can produce. It is convenient to think of an agent's type as his "location," although the model does not rely on geography. Also, the production types fall into two distinct groups, G_A and G_B, where $G_A \cap G_B = \varnothing$. The measure of group G_A (G_B) is given by μ_A (μ_B). In this section, $\mu_A = \mu_B = 1$.

Within each group, production types are distributed uniformly over the unit interval. Agents within each group wish to consume the goods of all other agents of the same group. Time begins at date $t = 0$. During the initial interval $t \in [0,1)$, nondurable goods of type *y,*

[9] As modeled here, only certain agents have the option to engage in identity theft. The environment studied can be generalized to allow for endogeneity along this margin (see Kahn and Roberds (2008)).

$y \in [0,1)$, 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 $y(t) \equiv t - [\![t]\!]$ are available for purchase and consumption, where $[\![t]\!]$ denotes the largest whole number less than or equal to *t*.

Over all times $t \geq 0$, production within group *i* imposes an instantaneous disutility of $c_i \delta\left(y(t) - y\right) dt$ on type-*y* agents, where $c_i > 0$ and *δ* is Dirac's delta function. For type-y' agents of group *i*, where $y' \in [0,1)$ and $y' \neq y$, time *t* consumption of one unit of a type-*y* good yields instantaneous utility $u_i \, dt$, where $u_i > c_i > 0$. 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.[10]

Both groups of agents consist of stochastically lived overlapping generations. At each discrete date $n = 0, 1, 2, \ldots$, a randomly selected subset of types die and are replaced by newborn agents of the same type. The measure of deaths and births is given by $1 - \beta$, where $0 < \beta < 1$. The deaths of agents and the identities of the dead immediately become public information, so only the living are potential victims of identity theft.

By construction, barter cannot occur in this economy, and money does not exist.

[10] It might be helpful to think of this environment as a limiting version of the following finite economy. Suppose that the economy consists of two groups of *NM* agents, where agents in each group are distributed over *N* locations, and *M* agents live at each location. At each discrete time $t + (y / N)$, where $t = 0, 1, \ldots,$ and $y = 0, \ldots, N - 1$, each type *i* agent can provide $N - 1$ goods to agents of type $y' \neq y$ who visit location *y*. In the limit, $N, M \to \infty$ as $M / N \to 0$. Credit in the limit economy more closely corresponds to Kocherlakota's (1998) concept of credit as "memory," which essentially requires that all transactions be between agents who have not encountered one another in any previous transaction.

Exchange thus depends on the existence of some sort of credit arrangement, and therefore on sufficient knowledge of agents' credit histories (Kocherlakota 1998).[11] A difficulty in constructing such histories is private information: in addition to an agent's identity, an agent's group and type are private information ex ante. Without some arrangement to overcome these frictions, no one would have an incentive to supply a good, and trade would not occur.

To enable trade to occur in some circumstances, a central authority (or "court") exists 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 a unit measure of goods during a time interval [0,1), [1,2), …). Second, at discrete dates $n = 0,1,2,\ldots$, the court can publicly announce the observed action.[12] 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*.

3.3 Benchmark: exchange without identity theft

As a benchmark, this subsection considers the case where there are no frauds ($F = 0$) and thus no identity theft. For ease of comparison, it is shown that the conditions under which exchange occurs here are the same as in many well-known models of decentralized exchange (e.g., Kiyotaki and Wright (1989) or Diamond (1990); see Kahn and Roberds (forthcoming) for a survey).

An agent's actions as a consumer (purchases of goods away from the agent's "home

[11] At the cost of considerable added complexity, the model could be modified to allow agents the option of transacting with cash as well as by credit. This generalization is explored in, for example, Martin, Orlando, and Skeie (2008) and Monnet and Roberds (forthcoming). Here it might be useful to think of agents' utility from consumption, u, as their "credit benefit," i.e., the utility from additional consumption (or convenience) beyond that which would be available if cash were the only means of transacting. The analysis below implicitly assumes that this credit benefit outweighs the privacy advantages of using cash (see Kahn et al. (2005)).

[12] In practice, such announcements (or close approximations thereto) are provided by credit bureaus.

location") are not observable by the center. Exchange will require some arrangement for associating the identities of would-be buyers with histories. These arrangements are modeled as clubs for sharing information on buyers' identities.[13] The analysis will initially consider the case where one club exists for each group of agents. Each club *i*, $i = G_A, G_B$, is formed at time $t = 0$. An agent joining the club agrees to reveal a subset of his identity, sufficient to distinguish him from all other agents.[14] Having revealed part of his identity, the agent receives an uncounterfeitable credit card that signals his membership in the club. The card can be costlessly authenticated by all club members.

By joining club *i*, an agent also reveals his group, though not his type. Club membership entitles the agent to a (flow) unit of a consumption good from any other club member in return for agreeing to provide a unit measure of his own type of good to other club members, at some point during each unit interval of time. At subsequent discrete dates $n = 1, 2, \ldots$, the center publicly announces the default of any club members who have not supplied goods and imposes penalty *X* on them, and they are excluded from the club. Membership in each club subsequently is opened to newborn agents.[15]

For an 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

$$\int_{t=n}^{n+1} \left(u_i - c_i \delta(y(t) - y)\right) dt = u_i - c_i \qquad (1)$$

for $n = 0, 1, 2, \ldots$. The date *n* discounted present value of club membership is

[13] As in Boyd and Prescott (1987), membership in the clubs will vary over time even as the clubs persist.

[14] Legitimate agents have no talent for obtaining goods through fraudulent activity. For the purposes of this introductory section only, it is assumed that data on legitimate agents can be assembled at a negligible cost.

[15] In a more general setup, the club would need to keep track of each agent's detailed consumption history as well. In the structure considered here, these histories would be essentially identical (differing only in the instants when goods are supplied), so that an agent's history is automatically revealed through his decision to supply goods.

$$V_{i,n} = u_i - c_i + \beta V_{i,n+1}. \tag{2}$$

In steady state, $V_{i,n} = V_{i,n+1} = V_i$, which implies

$$V_i = \left(\frac{1+r}{r}\right)(u_i - c_i), \tag{3}$$

where $r = \beta^{-1} - 1$. Ongoing membership in the club requires that a type-*y* agent be willing to supply a unit measure of goods at time $n + y$. Absent nonpecuniary penalties, this requires that the disutility of producing goods be less than the value of continued club membership, i.e.,

$$c_i \leq \beta V_i, \tag{4}$$

or equivalently that $\beta u_i \geq c_i$. If a nonpecuniary penalty *X* is available, condition (4) becomes

$$c_i - X \leq \beta V_i, \tag{5}$$

which implies that the club can always sustain the efficient allocation, in which everyone trades, for *X* sufficiently large. The analysis below will assume that condition (5) holds, so that all difficulties in organizing exchange stem from the presence of fraud.

3.4 Fraud and frauds

In general, a subset of the agents within each group at each date are frauds. Frauds resemble legitimate agents, except in the detail that they are unable (or for unknown reasons unwilling) to supply goods to other agents. However, frauds still enjoy consuming the goods produced by others. Thus, they cannot gain legitimate access to their preferred club without incurring penalties and subsequent exclusion. As a result, they have an incentive to obtain consumption goods by posing as legitimate agents. Within each group, frauds are not distributed uniformly over production types but are concentrated over a measurable set of known "locations," where the measure of this set is given by $F > 0$.

For legitimate agents, the presence of frauds reduces the value of club membership. In particular, if all frauds in group *i* are able to pose as legitimate agents, then the value of legitimate membership becomes

$$V_i = \left(\frac{1+r}{r}\right)\left(u_i(1-F)-c_i\right), \tag{6}$$

which is negative for *F* sufficiently close to one. A sufficiently high rate of fraud undermines legitimate agents' incentives to participate in a club, which can preclude trade. Legitimate agents thus will have an incentive to exclude frauds from their clubs.

3.5 Identification of agents

To distinguish legitimate agents from frauds, agents must be reliably identified. For this model, the amount of identifying information disclosed, not the type of information, matters. Hence, the information disclosed is represented by $d_{i,n} \in \mathbb{R}_+$, referring to the number of elements an agent must disclose from his identity vector to be identified by club *i* at discrete dates *n*. 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. (This assumption is relaxed in a subsequent section.)

Identification of agents is costly, and there are two components to the cost. The first component is a fixed one-time cost of K_i utils, which is incurred when an agent initially joins club *i* and is borne pro-rata by all legitimate club members. The second component is a per-discrete-period, per-member cost of processing and maintaining the data record $d_{i,n}$ for each club member. This cost is given by $k_i d_{i,n}$, where $k_i > 0$ and is also borne by all legitimate club *i* members. Note that the parameters K_i and k_i 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. That is, 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 $n = 0,1,\ldots$, 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 $n+1$.[16]

3.6 Identity theft

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.

Credit cards issued at discrete dates *n* have a virtual expiration date of $n+1$. That is, at discrete date $n > 0$, each club receives from the center a list of agents who have supplied goods during the preceding interval $[n-1, n)$. Members who have not supplied goods are revealed as frauds, penalized if their identities are known, and removed from the club, while those who have supplied goods continue their membership.[17] Apart from exclusion from the club, no penalties can be applied to impersonators because their real identities are unknown.

[16] In other words, data compiled at discrete date *n* and not needed at *n*+1 can be costlessly and securely disposed of at *n*+1, but must be held over the interval [*n*, *n*+1] to avoid incurring the fixed cost *K*. A more general setup could incorporate a flexible cost function for secure data disposal.

[17] 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 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.

Discovery of an impersonator in club *i* imposes a fixed resolution cost of *L*, which is borne equally by all legitimate members of club *i*.[18] *L* can include various kinds of costs, including physical costs, loss of leisure time, inconvenience, and simply loss of privacy.[19] Note that this cost is in addition to the fraud loss, *c*, incurred when a fraud illicitly obtains a good.

To gain access to a club, frauds must convincingly impersonate a legitimate agent. A fraud has two means of obtaining the necessary data: he can steal (i.e., observe) at least some of the data needed for the impersonation, or simply manufacture sufficient data to construct a convincing identity. Because the submission of duplicate PID of an existing club member would be automatically revealed as fraudulent (i.e., there is no existing-account fraud in the model), data observed in a breach of club *j*'s database is always used to gain access to club *i.*

The amount of data lost through a data breach depends on how well the target club secures its database. Suppose that club *i* decides to maintain member data $d_{i,n-1}$ over the interval $t \in [n-1, n)$, where $n > 1$. The club then chooses a 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 the would-be data thief.

More specifically, the variable $s_{i,n-1}$ is the technical *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 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 gain access to the club's database.

[18] Because all legitimate club members are risk neutral and have the same preferences, there is no loss of generality in assuming these costs are equally distributed.

[19] For example, according to Douglas (2008), it costs a card issuer about $25 to reactivate a compromised credit card account. Other, less readily quantifiable costs of resolving identity fraud are catalogued by Anderson et al. (2008), and include the time cost of resolution, harassment of victims by debt collectors, denial of utility service, and being subject to misplaced civil lawsuits and criminal investigations.

Increasing the skill required for database breaches brings with it increased costs, however. In particular, adopting skill threshold $s_{i,n-1}$ results in a cost to all legitimate members of club *i* of disutility $\ell s_{i,n-1}$ 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 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 $\varepsilon > 0$. As with the technical skills, the "fraud effort cost" ε is assumed to have a time-invariant distribution $\Gamma(\varepsilon)$ over the population of frauds, where Γ is independent of the skill distribution Φ.[20]

Frauds who possess sufficient skill 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 proportion η 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

$$\max\left\{d_{i,n} - \eta d_{j,n-1}, 0\right\}, \tag{7}$$

and his net effort cost is given by[21]

$$\varepsilon \max\left\{d_{i,n} - \eta d_{j,n-1}, 0\right\}. \tag{8}$$

To summarize, the prevalence and type of identity fraud committed in club *i* during $[n, n+1)$ 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

[20] Unskilled fraud can occur through opportunistic (low-tech) data theft, data synthesis, or a combination of the two.
[21] Note that skilled data theft always involves the theft of some data, which is then (in general) combined with synthesized data to construct a false identity.

amount of club *j*'s data obtainable through a breach at date *n*, $\eta d_{j,n-1}$.

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 a "breach cost" $B > 0$ borne equally by all members. As with the resolution cost *L*, *B* can include physical, time, and intangible costs.

3.7 Symmetric steady-state equilibrium

Suppose that at discrete date $t = n - 1$, club *j* decides to maintain data $d_{j,n-1}$ on its members and specifies a skill threshold $s_{j,n-1}$. For an unskilled fraud (one unable to attempt a data breach) from group *i*, the payoff to committing identity theft at $t = n$ and gaining access to club *i* over $t \in [n, n+1)$ is given by

$$u_i(1-F) - \varepsilon d_{i,n}. \tag{9}$$

From (9), club *i*'s equilibrium rate of identity theft from unskilled frauds over $t \in [n, n+1)$ is given by

$$F\Phi(s_{j,n-1})\Gamma\left(\frac{u_i(1-F)}{d_{i,n}}\right). \tag{10}$$

For a skilled fraud of group *i*, the payoff from fraud over $t \in [n, n+1)$ is given by

$$u_i(1-F) - \varepsilon \max\left\{d_{i,n} - \eta d_{j,n-1}, 0\right\}. \tag{11}$$

Hence the set of successful skilled frauds in the population of group *i* is those for whom

$$\varepsilon \le \bar{\varepsilon} = \begin{cases} \dfrac{u_i(1-F)}{d_{i,n} - \eta d_{j,n-1}}, & \text{when } d_{i,n} - \eta d_{j,n-1} > 0, \\ \infty, & \text{when } d_{i,n} - \eta d_{j,n-1} \le 0. \end{cases} \tag{12}$$

If preferences are symmetric across clubs ($u_A = u_B = u$; $c_A = c_B = c$; $K_A = K_B = K$; $k_A = k_B = k$), then in steady-state equilibrium it must be the case that $d_{i,n} - \eta d_{j,n-1} > 0$. Hence, for the symmetric case, the measure of skilled frauds who gain access to club *i* at discrete date *n* can be stated as

$$F\left(1 - \Phi(s_{j,n-1})\right)\Gamma\left(\frac{u(1-F)}{d_{i,n} - \eta d_{j,n-1}}\right). \tag{13}$$

Each club *i* chooses a data record length $d_{i,n}$ and a skill threshold $s_{i,n}$ for each discrete date *n* so as to maximize the discounted future utility of its club members, taking into account the choices of the other club.[22] Club *i*'s date-*n* objective (the continuation value of club membership) can be represented as

$$V_{i,n}^{f} = \sum_{m=0}^{\infty} \beta^{m} U_{i,n+m} , \tag{14}$$

where $U_{i,n}$ gives each legitimate agent's payoff to membership in club *i* over $[n, n+1)$, i.e.,

$$\begin{aligned} U_{i,n} &= (u-c)(1-F) - (1-\beta)K - kd_{i,n} - \ell s_{i,n} \\ &\quad - F\Phi(s_{j,n-1})\Gamma\left(\frac{u\left(1-F\right)}{d_{i,n}}\right)(c+L) - F\left(1-\Phi(s_{j,n-1})\right)\Gamma\left(\frac{u\left(1-F\right)}{d_{i,n} - \eta d_{j,n-1}}\right)(c+L) \\ &\quad - F\left(1-\Phi(s_{i,n-1})\right)\Gamma\left(\frac{u\left(1-F\right)}{d_{j,n} - \eta d_{i,n-1}}\right)B. \end{aligned} \tag{15}$$

In words, a legitimate agent's per-period payoff is given by the net benefits of trade, minus the costs associated with administering data and keeping it secure, minus the costs associated with identity theft by the unskilled and skilled, minus the costs of resolving data breaches.

A *symmetric steady-state allocation* in this economy is an ordered pair (d, s), where *d*

[22] The equilibrium studied here is that of an "open loop" Nash equilibrium (i.e., an equilibrium in sequences of decisions) in a dynamic game between clubs *A* and *B* (Başar and Olsder 1998). The open-loop equilibrium concept is appealing here because, by restricting the clubs to relatively simple strategies, it allows for straightforward analytical characterization of the strategic interactions between clubs. This equilibrium is time-consistent in the usual sense.

gives the data length and *s* gives the skill threshold for both clubs. In symmetric steady state, the continuation value of membership in each club is given by

$$V^f(d,s) = \left(\frac{1+r}{r}\right)\begin{bmatrix} (u-c)(1-F)-(1-\beta)K-kd-\ell s-F\Phi(s)\Gamma\left(\dfrac{u(1-F)}{d}\right)(c+L) \\ -F\left(1-\Phi(s)\right)\Gamma\left(\dfrac{u(1-F)}{d(1-\eta)}\right)(c+L+\beta B) \end{bmatrix}. \qquad (16)$$

A symmetric steady-state allocation (d,s) is *incentive compatible* if it satisfies the following conditions:

1. *Individual rationality*, i.e., $V^f(d,s)\geq 0$;
2. *No defection* (legitimate agents in each club have an incentive to produce), i.e., $\beta V^f(d,s)\geq c-X$;
3. *No exclusion* (each club has an incentive to admit new members), i.e., $V^f(d,s)\geq \underline{V}$, where $\underline{V}$ is the value of maintaining the club without admitting new members:

$$\underline{V} = (u-c)(1-F)\sum_{n=0}^{\infty}\beta^{2n} = \frac{(u-c)(1-F)}{1-\beta^2} \qquad (17)$$

A symmetric steady-state allocation (d^*,s^*) is an *equilibrium* if

1. It is incentive compatible, and
2. The infinite sequence $\left\{(d_{i,n},s_{i,n})\right\}_{n=0}^{\infty} = \left\{(d^*,s^*),(d^*,s^*),\,\ldots\right\}$ represents a best response for each club in steady state, i.e., $\left\{(d^*,s^*),(d^*,s^*),\,\ldots\right\}$ maximizes $V_{i,0}^f$ for each club *i*, when club *j* also chooses $\left\{(d_{j,n},s_{j,n})\right\}_{n=0}^{\infty} = \left\{(d^*,s^*),(d^*,s^*),\,\ldots\right\}$, and both clubs have "initial conditions" $(d_{i,-1},s_{i,-1})=(d^*,s^*)$.

The analysis below considers illustrative equilibria for two candidate distributions for Φ and Γ. In particular, frauds' skill endowments *s* are specified to follow the exponential distribution $\Phi(s)=1-e^{-\phi s}$, and the distribution $\Gamma(\varepsilon)$ of frauds' effort costs is specified as a uniform distribution, normalized to $U[0,1]$. These choices can be rationalized as follows. In the case of Φ, the set of equilibria considered will be determined by the hazard function $f(s)=\Phi'(s)/\left(1-\Phi(s)\right)$. The analysis below focuses on the case of a constant hazard rate $f(s)=\phi$, which is equivalent to assuming an exponential distribution for *s*. Note that ϕ determines the incremental benefit of a small increase in data security. In the case of Γ, a sufficiently "flat" distribution ($|\Gamma''|$ small) is necessary to ensure the intuitive property that each club's optimal data length *d* decreases with a falling cost of maintaining such data *k*. This requirement is clearly satisfied if Γ is uniform.[23] Together, these specifications for Φ and Γ can be shown to guarantee sufficient concavity of $V^f_{i,n}$ so that each club's objective is well defined.[24]

When Γ is $U[0,1]$, first-order conditions in $d_{i,n}$ and $s_{i,n}$ are given by

$$\begin{aligned}&\frac{uF(1-F)(c+L)\Phi(s_{j,n-1})}{d^2_{i,n}}+\frac{uF(1-F)(c+L)\left(1-\Phi(s_{j,n-1})\right)}{\left(d_{i,n}-\eta d_{j,n-1}\right)^2}\\&=k+\frac{uF(1-F)\beta B\eta\left(1-\Phi(s_{i,n})\right)}{\left(d_{j,n+1}-\eta d_{i,n}\right)^2},\end{aligned}\tag{18}$$

and

$$\frac{uF(1-F)\beta B\Phi'(s_{i,n})}{d_{j,n+1}-\eta d_{i,n}}\le \ell,\tag{19}$$

[23] This specification is implicit in the model of Kahn and Roberts (2008).
[24] See the proof of Proposition 1 below.

respectively, where (19) holds with equality for $s_{i,n} > 0$. Note that the left-hand side of condition (18) [(19)] gives the marginal benefit of an increase in $d_{i,n}$ [$s_{i,n}$] while the right-hand side gives its marginal cost. In symmetric steady state these conditions reduce to

$$uF(1-F)\left[\frac{(c+L)\Phi(s)}{d^2}+\frac{\left(c+L-\beta B\eta\right)\left(1-\Phi(s)\right)}{d^2\left(1-\eta\right)^2}\right]=k\,; \tag{20}$$

$$\frac{uF(1-F)\beta B\Phi'(s)}{d\left(1-\eta\right)}\leq \ell\,. \tag{21}$$

Thus, for this particular specification, a symmetric steady-state allocation (d,s) is an equilibrium if it is incentive compatible, and satisfies (20) and (21). The following proposition may now be stated (proofs in this section are given in the Appendix):

***Proposition* 1**. A unique symmetric steady-state equilibrium (d^*,s^*) exists with $s^*>0$ under the following conditions:

1. the hazard rate ϕ of the skill distribution is sufficiently large;
2. the breach cost B is less than the other costs of identity theft, i.e., $\beta B < c + L$;
3. verification costs $K,k,\ell > 0$ are sufficiently small;
4. β is sufficiently close to unity (agents are sufficiently long-lived).

3.8 Comparison with the efficient allocation

The data record length d^* and the skill threshold s^* in the symmetric equilibrium allocation can be usefully compared to the values of d and s that would be chosen by a planner. The planner operates under the same informational constraints as the decentralized arrangements. PID must be freely surrendered and cannot be shared across groups. A separate club is formed for each group, and agents have the option of joining the appropriate club. Also, allocations chosen by the planner are subject to the same incentive-compatibility constraints as in the noncooperative allocation.

The planner's objective is taken as the steady-state value of legitimate agents' club membership, $V^f(d,s)$. A *golden-rule allocation* is a steady-state allocation (d_p, s_p) that maximizes the planner's objective. Note that a golden-rule allocation represents a constrained-efficient allocation because the planner places no weight on either the utility of the initial generation of legitimate agents or the utility of frauds.

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

$$uF(1-F)\left[\frac{(c+L)\Phi(s)}{d^2}+\frac{\left(c+L+\beta B\right)\left(1-\Phi(s)\right)}{d^2\left(1-\eta\right)}\right]=k, \tag{22}$$

$$uF(1-F)\left(-\frac{c+L}{d}+\frac{c+L+\beta B}{d\left(1-\eta\right)}\right)\Phi'(s)\le \ell, \tag{23}$$

where (23) holds with equality for $s>0$. These conditions differ from equilibrium conditions (20) and (21) because the planner fully internalizes the fraud-suppression benefits of setting both the required data length *d* and the skill threshold *s*. The following result is shown in the Appendix.

Proposition 2. Under the conditions of Proposition 1, there is a unique golden-rule allocation (d_p, s_p) where $s_p > 0$.

The next two results compare the solution to the planner's problem to the symmetric steady-state equilibrium:

Proposition 3. Under the conditions of Proposition 1,

1. s^* and s_p are increasing in η (skill thresholds increase as stolen data becomes more useful for identity theft);
2. As $\eta \to 1$ (stolen data becomes more useful), $s^* \to \overline{s} < \infty$ while $s_p \to \infty$, whence $s^* < s_p$ (the skill threshold in the symmetric equilibrium is lower than that chosen by the planner).

Proposition 4. Under the conditions of Proposition 1,

1. The amount of data collected by the planner, d_p, does not vary with η, while the amount of data collected in the symmetric equilibrium, d^*, is increasing in η as $\eta \to 1$;
2. As $\eta \to 1$ (stolen data becomes more useful), $d^* \to \infty$, whence $d^* > d_p$ (the planner collects less data than is collected in the symmetric equilibrium).

Not surprisingly, rates of identity theft differ across the two allocations. For a steady-state allocation (d, s), the rate of identity theft (measure of successful frauds) $\rho(d, s)$ is given by the sum of the rate of identity theft by unskilled and skilled frauds, and can be computed as

$$\frac{\rho(d,s)}{uF(1-F)} = \frac{\Phi(s)}{d} + \frac{\left(1-\Phi(s)\right)}{d(1-\eta)}. \tag{24}$$

Proposition 5. Under the conditions of Proposition 1,

1. The rate of skilled identity theft is greater in the symmetric equilibrium than under the golden-rule allocation;

2. As $\eta \to 1$, the rate of unskilled identity theft is greater under the golden-rule allocation than in the symmetric equilibrium;
3. For ℓ / k bounded, as $\eta \to 1$ the total rate of identity theft is greater under the golden-rule allocation than in the symmetric allocation.

3.9 Discussion

Proposition 3 establishes, under mild conditions, that when each card network independently determines the amount and security of data compiled on its members, networks insufficiently secure their data relative to the golden-rule allocation. The clubs attempt to compensate for insufficient security by overaccumulating identifying data on their members (Proposition 4).

Insufficient security is applied because each network's cost of a data breach B is less than its social cost $c + L + \beta B$. Lax security leads, in turn, to a suboptimally high rate of identity theft by skilled frauds (Proposition 5). Because each network cannot control the rate of data theft from the other network's database, its best response is to accumulate more PID, thereby suppressing the rate of unskilled identity theft, and driving the overall rate of identity theft below that of the efficient allocation (Proposition 5). Despite its lower rate of identity theft, the equilibrium allocation is inefficient due to its higher "privacy" costs, i.e., the costs of assembling and maintaining personal data necessary to keep fraud in check.

4. Attaining efficiency

This section discusses three means for improving on the inefficient steady-state equilibrium allocation: (1) mandated limits on the amount of data collected and security levels,

(2) reallocations of data-breach costs, and (3) data sharing through a merger of the networks.

4.1 Direct regulation

One possibility would be direct regulation of entities engaged in the collection of personal data, such as the clubs in the model. The strategic interplay between the data compiled and its security imposes a high informational burden on this type of regulation. In practice it may be difficult for policymakers to enforce standards along both of these dimensions. Consequently, this section analyzes the effects of policies that regulate data collection or data security, but not both.

Suppose, for example, that a regulator observes that excessive PID is being collected, and decides to constrain the amount of data that each network collects, i.e., the regulator requires $d = d_c < d^*$. Security levels would still be set noncooperatively: let s_c be the equilibrium skill threshold chosen by the clubs under this constraint.

From the equilibrium condition (21), s_c can be expressed as

$$s_c = \frac{1}{\phi} \ln\left(\frac{\phi \beta B u F(1-F)}{\ell d_c (1-\eta)} \right), \tag{25}$$

which can be compared to condition (23) evaluated at the solution to planner's problem

$$s_p = \frac{1}{\phi} \ln\left(\frac{\phi\big(\eta(c+L)+\beta B\big) u F(1-F)}{\ell d_p (1-\eta)} \right). \tag{26}$$

A benevolent regulator who only regulates data length would choose d_c to maximize $V^f(d_c, s_c)$ subject to (25). The solution to the regulator's problem is given as (calculations are in the Appendix)

Proposition 6. A regulator who can only regulate data length chooses the same data length as its golden-rule value, i.e., the regulator sets $d_c = d_p$. Under this policy, as $\eta \to 1$, clubs choose a skill threshold s_c greater than its value in the symmetric equilibrium, but less than its golden-rule value, i.e., $s^* < s_c < s_p$.

Thus, relative to an unregulated outcome, a policy of constraining data collection improves welfare by (1) reducing the costs of data collection (including intangible costs) and (2) encouraging networks to increase security and therefore make skilled identity theft more difficult. The potential benefit of increased security can be partly undone by two effects, however. First, there is substitution into unskilled identity theft, since unskilled identity theft becomes both easier (less PID is required for an impersonation) and more popular, as some skilled frauds shift into low-tech forms of identity theft. Second, for frauds with sufficient technical abilities, skilled identity theft becomes easier as it requires less PID. Relative to the efficient allocation, a policy of constraining data collection does not completely correct the inefficient pattern observed in the unregulated equilibrium, of over-suppression of unskilled and under-suppression of skilled identity theft (cf. Proposition 5).

Likewise, a regulator might require that networks increase security levels. Consider the case where a regulator requires each network to implement $s = s_C$, but allows networks to privately determine the amount of data that they collect. Let d_C be the equilibrium amount of data chosen by the clubs under this constraint. From equilibrium condition (20), d_C can be expressed as

$$d_C = \left[\frac{uF(1-F)}{k(1-\eta)^2}\right]^{\frac{1}{2}} \left[(c+L)(1-\eta)^2\Phi(s_C)+(c+L-\beta\eta B)\left(1-\Phi(s_C)\right)\right]^{\frac{1}{2}}, \tag{27}$$

which can be compared to condition (22) evaluated at the solution to the planner's problem

$$d_p = \left[\frac{uF(1-F)}{k(1-\eta)}\right]^{\frac{1}{2}} \left[(c+L)(1-\eta)\Phi(s_p)+(c+L+\beta B)\left(1-\Phi(s_p)\right)\right]^{\frac{1}{2}}. \tag{28}$$

A benevolent regulator who regulates only security would choose s_C to maximize $V^f(d_C, s_C)$ subject to (27). In this case, the solution to the regulator's problem is given as

Proposition 7. A regulator who can only regulate skill thresholds chooses a skill threshold higher than its golden-rule value as $\eta \to 1$, i.e., the regulator sets $s_C > s_p$. By Proposition 3, $s_C > s^*$. Under this policy, clubs choose a data length less than its symmetric equilibrium value, i.e., $d_C < d^*$, but greater than its golden-rule value as $\eta \to 1$, i.e., $d_C > d_p$.

Proposition 7 says that a policy of regulating data security only tends towards overcorrection in data security, imposing even higher security standards than a planner would choose. This policy also reduces the amount of data collected by the clubs. Relative to the symmetric equilibrium, a higher rate of unskilled identity theft is tolerated in return for a reduction in the costs of collecting and maintaining data.

4.2 Increasing liability for a data breach

An alternative regulatory approach would be to increase each network's costs for a data breach so as to better align the private and social costs of a breach, i.e., raising each network's breach costs to $B' = B + \pi$ where $\pi > 0$. This might occur in a number of ways. One possibility would be to increase each network's civil legal liability for the costs resulting from theft of its data. Another possibility would be for regulators to levy penalties in the case of a data breach. Such penalties have been de facto imposed, for example, by at least 35 state legislatures through the passage of laws that require consumers be notified (at some cost to the data collector) when their data is subject to unauthorized access (Anderson et al. 2008).

There are some significant practical restrictions on this type of policy. For example, under U.S. law it is difficult to establish liability for identity theft because many entities have access to payment data, which tends to constrain the use of contractual agreements to allocate the risk of harm from identity theft (Schreft 2007). Awards for damages, when they do occur, are limited to the economic loss resulting from a breach, rather than the larger amounts that might result from application of a negligence standard (Chandler (forthcoming)).

Translating these constraints in the context of the model, an "economic loss standard" would require each club to pay a pro-rata share of the losses of the other when it experiences identity theft stemming from a data breach. That is, an economic loss standard would impose a penalty

$$\pi = \pi_{EL} = (1+r)\eta(c+L) . \tag{29}$$

An economic loss standard achieves efficiency for the special case where clubs are constrained to collect the efficient amount of data ($d = d_p$). To see this, note that if we replace *B* in the clubs' first-order condition (25) with $B' = B + \pi_{EL}$, this is the same as the planner's first-order condition (26) so long as *d* is identical in both conditions. Where data length is endogenous, however, an economic loss standard does not correct clubs' incentives in data collection (cf. conditions (27) and (28)); hence efficiency does not obtain for the general case.

4.3 Numerical Example

To better gauge the relative efficacy of the various regulatory approaches, allocations were computed numerically. Table 1 below displays some typical results. Parameter values for the example are:

$$c+L=25;\ \ B=1;\ \beta=.9;$$
$$\phi=2;\ \eta=.5;\ k=1;\ \ell=.1\,.$$

Note that these parameter values allow for a moderate spillover ($\eta=.5$) from one club's data practices to the other's. The ratio $(k/\ell)=10$ places a relatively high value on the privacy of personal information. To facilitate comparisons, the normalizations $K=0$ and $uF(1-F)=1$ are adopted. Columns 1 and 2 of the Table give the numerical values of the allocation (d,s) in each case. Column 3 gives the percentage of skilled identity thieves, i.e., the proportion of frauds who are able to attempt data breaches. Column 4 gives the identity theft rate $\rho(d,s)$ of each allocation. Since *uF*(1-*F*) is normalized to one in the examples, the identity theft rates in Table 1 do not represent gross identity theft rates, but instead represent the proportion of frauds who are successful in their attempts at impersonation. Column 5 gives the steady-state variable cost of identity theft for each allocation, including the cost of data collection and security, i.e.

$$\begin{aligned} C(d,s) &= -\left(\frac{r}{1+r}\right)V^{f}-(u-c)(1-F)-(1-\beta)K \\ &= \frac{c+L}{d}\Phi(s)+\frac{(c+L+\beta B)}{d(1-\eta)}\big(1-\Phi(s)\big)+kd+\ell s\,. \end{aligned} \tag{30}$$

Table 1: Comparison of Allocations

	Personal data collected d	Security level (skill threshold) s	Percentage of skilled frauds $100*(1-\Phi(s))$	Identity theft rate $100*\rho(d,s)$	Steady-state costs of ID theft $C(d,s)$
1. Golden-rule (efficient) allocation: (d_p, s_p)	5.03	2.53	0.6	20.0	10.3
2. Symmetric equilibrium: (d^*, s^*)	33.2	0.04	92.3	5.78	34.7
3. Regulated data collection: (d_c, s_c)	5.03	0.984	14.0	22.7	10.8
4. Regulated security level: (d_C, s_C) (approximate)	5.04	2.5	0.7	20.0	10.3
5. Economic loss standard: (d^*, s^*) when $B' = B + \pi_{EL}$	17.3	1.72	3.2	6.00	19.0

Allocations 1 and 2 illustrate the comparisons derived in Propositions 3, 4, and 5. In symmetric equilibrium, the networks collect over six times as much data as in the efficient allocation, and the equilibrium security effort (skill threshold) is very low. Identity theft is effectively suppressed in the symmetric equilibrium, but the welfare cost of this suppression is high since so much data is collected.

Of the three regulatory policies, regulation in security levels (i.e., skill levels, allocation 4) is the most effective, virtually replicating the efficient allocation for this example. This policy is successful because requiring clubs to increase their security efforts simultaneously diminishes their incentives to overcollect personal data. A policy of limiting data collection (allocation 3) does nearly as well, since placing limits on PID also improves clubs' security incentives. The least effective policy is the implementation of an economic loss standard (allocation 5). While this policy increases data security and improves welfare, it does not fully eliminate clubs' incentives to inefficiently substitute data collection for data security.

4.4 Variable 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 governmental agency that would manage a consolidated database of PID. Inclusion in the database would be optional. This section 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 in the environment studied in Section 3.

In this generalized environment, agents have the option of transacting through a single club or dual clubs (one for each group of agents). If agents decide to form a single club, no data breaches can occur in equilibrium, so the club simply compiles data of length *d* on all its members[25] to maximize the average per-capita net benefit of legitimate club membership. That is, the single club chooses *d* to maximize (cf. expression (15))

$$V_s = \left(\frac{1+r}{r}\right)\times \left[(\underline{u}-\underline{c})(1-F)-(1-\beta)\underline{K}-\underline{k}d-\mu_A\frac{u_A F(1-F)}{d}(c_A+L)-\mu_B\frac{u_B F(1-F)}{d}(c_B+L)\right], \tag{31}$$

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 (31), 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 are satisfied

1. *Individual rationality*, $V_{i,s}\geq 0$ for $i=G_A,G_B$;
2. *No defection*, $\beta V_{i,s}\geq c_i - X$ for $i=G_A,G_B$;
3. *No exclusion*, $V_{i,s}\geq \underline{V}_i$ *for* $i=G_A,G_B$, where $\underline{V}_i$ is the value of maintaining the club without admitting new members, analogous to (17).

If, as in Section 3, 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 equilibrium with the single-club equilibrium dominates the equilibrium with dual clubs. 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

[25] Recall that an agent's group is private information, so the club cannot require different amounts of data from agents of different groups. Note also that the information provided by the court does not allow for separation of agents by groups ex post.

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* reflects only 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 essentially 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 $\mu_B = \mu \in (0,1)$.

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

$$d_s = \sqrt{\frac{\underline{u}F(1-F)(c+L)}{\underline{k}}} \quad , \tag{32}$$

and the equilibrium per-capita net benefit of club membership for an agent of group *i* is

$$V_{i,s} = \left(\frac{1+r}{r}\right)\left[(u-c)(1-F)-(1-\beta)K-\left(k_i+\sqrt{\frac{\underline{k}}{1+\mu}}\right)\sqrt{\frac{\underline{u}F(1-F)(c+L)}{\underline{k}}}\right], \tag{33}$$

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 $d_A \to \infty$ as $k_A \to 0$ and chooses s_A to maximize

$$V_{A,d} = \left(\frac{1+r}{r}\right)\left[(u-c)(1-F) - (1-\beta)K - \ell s_A - \mu F\left(1-\Phi(s_A)\right)\beta B\right]. \tag{34}$$

For sufficiently large ϕ, the optimal skill threshold for club G_A will be given by

$$s_A = \phi^{-1}\ln\left((\mu F\beta B\phi)/\ell\right), \tag{35}$$

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

$$V_{A,d}^{*} = \left(\frac{1+r}{r}\right)\left[(u-c)(1-F) - (1-\beta)K - \left(\frac{\ell}{\phi}\right)\left(\ln\left(\frac{\ell}{\mu FB\phi}\right)+1\right)\right]. \tag{36}$$

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, 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

$$\begin{gathered} V_{B,d} = \left(\frac{1+r}{r}\right)\times \\ \left[\begin{array}{l} (u-c)(1-F) - \\ (1-\beta)K - k_B d_B - \dfrac{uF(1-F)}{d_B}(c+L) - F\left(1-\Phi(s_A)\right)(c+L) \end{array}\right], \end{gathered} \tag{37}$$

which yields

$$d_B = \sqrt{\frac{uF(1-F)(c+L)}{k_B}}\,. \tag{38}$$

From (35) and (38), the equilibrium per-capita net benefit of membership in club G_B in the case of dual clubs is given by

$$\begin{aligned} V^*_{B,d} &= \left(\frac{1+r}{r}\right)\times \\ &\left[(u-c)(1-F)-(1-\beta)K-\sqrt{k_B uF(1-F)(c+L)}-\left(\frac{\ell}{\phi}\right)\left(\frac{c+L}{FB}\right)\right]. \end{aligned} \tag{39}$$

For this parameterization, the comparison between the single club and dual clubs is stated as

Proposition 8. 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,

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

Proof. The proof of Part 1 follows that of Propositions 1 and 2. To show Part 2, let $\ell / \phi \to 0$. Then, comparing (33) and (36), $V^*_{A,d} > V_{A,s}$ for $\mu_B > 0$ sufficiently small. Comparing (33) and (39), $V^*_{B,d} > V_{B,s}$ under the same conditions.

Intuitively, Proposition 3 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 entirely by forming a single club. Agents who place little value on privacy allow their club to compile large amounts of personal data, since this deters fraud, even though this data is subject to occasional breaches and misuse. By contrast, agents who place a higher value on privacy will tolerate a higher rate of identity theft, including theft that arises through data breaches, 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 8 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.

5. Relationship to the Literature

The above analysis builds on well-known models of exchange in search-theoretic environments. Numerous papers in this literature allow for the possibility of fraudulent transactions, both through counterfeit currency (Green and Weber (1996), Kultti (1996), Monnet (2005), Williamson (2002), Nosal and Wallace (2007), Cavalcanti and Nosal (2007)) and through various types of fraud in credit-based payments (Camera and Li (2003), Kahn et al. (2005), Kahn and Roberds (2008)). What is new here is the consideration of a new, empirically significant, type of transactions fraud stemming from the theft of identifying data.

The framework presented also shares features with some papers in the literature on the economics of information security (surveyed in Anderson and Moore (2006)). Varian (2004) presents a game-theoretic model in which “system reliability” (here, corresponding roughly to deterrence of identity theft) is modeled as a public good produced by the interaction between individual efforts at reliability provision (corresponding to PID collection and storage). The Varian model is extended by Grossklags et al. (2008) to allow for individual efforts at insurance (corresponding to setting skill thresholds) against system failures.

The environment above is similar to these models, in the sense that knowledge of PID functions as a nonrival good within each group of agents, supplying a club-wide level of deterrence against identity theft. Here, unanimity of preferences ensures that public-goods problems do not arise within groups. Instead the focus is on potential negative spillovers across groups: 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 therefore involves a balance between its positive (within-club) and negative (cross-club) effects.

6. Conclusion

This paper has presented a formal model in which identity theft arises endogenously and the concept of an efficient degree of 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. In noncooperative settings, inefficiencies can arise due to spillovers from one group of agents’ decisions along these dimensions to another’s. Interventions such as direct regulation of security practices can increase efficiency, but the multidimensional nature of the security problem means that attaining full efficiency may be problematic. Sharing data across groups also can improve efficiency, but heterogeneity in preferences may limit welfare gains attainable through this channel.

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 exchange 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.

Appendix: Proofs.

Proof of Proposition 1.

The proof proceeds in three steps. First, we show that any solution (d,s) to conditions (20) and (21) 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^*), so that this solution satisfies the second requirement for an equilibrium. Third, we verify that (d^*,s^*) is incentive compatible.

Step 1. In an open-loop Nash equilibrium, at each discrete date *n*, each club *i* maximizes its objective $V_{i,n}^{f}\ \{(d_{i,n},s_{i,n})\}$ by choosing a strategy $\{(d_{i,n+m},s_{i,n+m})\}_{m=0}^{\infty}\in C$, taking the strategy of the other club $\{(d_{j,n+m},s_{j,n+m})\}_{m=0}^{\infty}$ as given. Each club's strategy space *C* is the product space $\{(D\times S)\times(D\times S)\times\ldots\}$, where *D* and *S* are the set of feasible choices for *d* and *s* at each discrete time period. To guarantee that this problem is well-defined take $D=(\underline{d},\overline{d})$ and $S=(\underline{s},\overline{s})$ for "small" $\underline{d},\underline{s}>0$ and "large" $\overline{d},\overline{s}<\infty$.

In general, necessary conditions for an interior optimum for club *i*'s problem are given by functional (i.e., difference) equations (18) and (19) at equality (e.g., Luenberger (1969), chapter 7). But in symmetric steady state, each club *i* knows that the other club will play a time-invariant strategy, which implies, from (18) and (19), that club *i*'s best response will also be time invariant. Club *i*'s optimization problem can therefore be reduced to the following ordinary calculus problem: choose (d_i,s_i) to minimize club *i's* steady-state cost of identity theft, given (d_j,s_j), i.e., choose (d_i,s_i) to minimize

$$\begin{aligned} &kd_i + \ell s_i + F\Phi(s_j)\left(\frac{u(1-F)}{d_i}\right)(c+L) + F\big(1-\Phi(s_j)\big)\left(\frac{u(1-F)}{d_i-\eta d_j}\right)(c+L) \\ &\quad + \beta F\big(1-\Phi(s_i)\big)\left(\frac{u(1-F)}{d_j-\eta d_i}\right)B. \end{aligned} \tag{40}$$

For ϕ sufficiently large, first-order conditions for the simplified problem are given by

$$\begin{aligned} &k - F\Phi(s_j)\left(\frac{u(1-F)}{d_i^2}\right)(c+L) - F\big(1-\Phi(s_j)\big)\left(\frac{u(1-F)}{(d_i-\eta d_j)^2}\right)(c+L) \\ &\quad + \beta\eta F\big(1-\Phi(s_i)\big)\left(\frac{u(1-F)}{(d_j-\eta d_i)^2}\right)B = 0, \end{aligned} \tag{41}$$

$$\ell - \frac{\beta BuF(1-F)\Phi'(s_i)}{d_j-\eta d_i} = 0\,, \tag{42}$$

which correspond to (20) and (21) when $d_i = d_j$. Second-order conditions are given by

$$\frac{2\Phi(s_j)(c+L)}{d_i^3} + \frac{2\big(1-\Phi(s_j)\big)(c+L)}{(d_i-\eta d_j)^3} > 0\,, \tag{43}$$

$$\frac{\beta B\Phi''(s_i)}{d_j-\eta d_i} < 0\,, \tag{44}$$

$$\left[\frac{2\Phi(s_j)(c+L)}{d_i^3} + \frac{2\big(1-\Phi(s_j\big)(c+L)}{(d_i-\eta d_j)^3}\right]\left[\frac{\beta B\Phi''(s_i)}{(d_j-\eta d_i)}\right] + \left[\frac{\beta B\Phi'(s_i)}{(d_j-\eta d_i)^2}\right]^2 < 0. \tag{45}$$

Conditions (43) and (44) are readily seen to hold when $(d_i, s_i) = (d_j, s_j)$. Sufficient conditions for (45) to hold are symmetry and $\beta B < 2(c+L)$, which is implied by $\beta B < c+L$.

Step 2. Under the assumption of a constant hazard rate for Φ, rewrite (21) at equality as

$$d = D(s) \equiv \frac{uF(1-F)\beta B\phi\big(1-\Phi(s)\big)}{\ell\big(1-\eta\big)}. \tag{46}$$

Substituting (46) into (20) and rearranging gives the following quadratic equation

$$Q(z) \equiv A_0(1-z) + A_1 z + A_2 z^2 = 0\,, \tag{47}$$

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

$$A_0 = c + L\,, \tag{48}$$

$$A_1 = \frac{c + L - \beta B\eta}{(1-\eta)^2}\,, \tag{49}$$

$$A_2 = -kuF(1-F)\left(\frac{\beta B\phi}{\ell(1-\eta)}\right)^2. \tag{50}$$

From the above, $Q(0) = A_0 > 0$ and $Q(1) = A_1 + A_2 < 0$ for ϕ sufficiently large. $Q(z)$ therefore has a unique root $z^* \in (0,1)$; in particular, $z^* =$

$$\frac{c + L - \beta B\eta - (1-\eta)^2(c+L) + \sqrt{\left(c + L - \beta B\eta - (1-\eta)^2(c+L)\right)^2 + 4(c+L)kuF(1-F)\left(\frac{\beta B\phi}{\ell}\right)^2}}{2kuF(1-F)\left(\frac{\beta B\phi}{\ell}\right)^2}. \tag{51}$$

Now define

$$(d^*, s^*) = \left(D\left(\Phi^{-1}(1-z^*)\right), \Phi^{-1}(1-z^*)\right). \tag{52}$$

By construction, (d^*, s^*) satisfies (20) and (21), and $s^* > 0$.

Step 3. To show incentive compatibility, suppose initially that $F = 0$, so that $V^f = V$ where V is given in (3). Then the individual-rationality, no-defection, and no-exclusion conditions are clearly satisfied with strict inequality for β sufficiently close to unity. Now, for $F > 0$, let $K, k,$ and ℓ approach zero; more specifically let $\|(K, k, \ell)\| < \theta$ where $\theta > 0$ and $\|\bullet\|$ is the sup norm. Then it can be shown that as $\theta \to 0$, d^* and s^* as defined in (52) are bounded by $\theta^{-1/2}$ and $-\ln\theta$, respectively. This, in turn, implies that $V^f(d^*, s^*) \to V$ as $\theta \to 0$, 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 Proposition 2.

Begin by solving for (d_p, s_p). Rewrite (23) at equality as

$$d = \underline{D}(s) \equiv \frac{uF(1-F)\big((c+L+\beta B)-(c+L)(1-\eta)\big)\phi\big(1-\Phi(s)\big)}{\ell\big(1-\eta\big)}. \tag{53}$$

Substituting (53) into (22) and rearranging gives the following quadratic equation

$$\underline{Q}(z) \equiv \underline{A}_0(1-z)+\underline{A}_1 z+\underline{A}_2 z^2 = 0\,, \tag{54}$$

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

$$\underline{A}_0 = c+L\,, \tag{55}$$

$$\underline{A}_1 = \frac{c+L+\beta B}{1-\eta}\,, \tag{56}$$

$$\underline{A}_2 = -kuF(1-F)\left(\frac{\phi\big((c+L+\beta B)-(c+L)(1-\eta)\big)}{\ell(1-\eta)}\right)^2. \tag{57}$$

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

$$z_p = \frac{(1-\eta)\left(1+\sqrt{1+4(c+L)kuF(1-F)\left(\dfrac{\phi}{\ell}\right)^2}\right)}{2kuF(1-F)\left(\dfrac{\phi}{\ell}\right)^2} \tag{58}$$

The golden-rule allocation is then given as $(d_p, s_p) = \Big(\underline{D}\big(\Phi^{-1}(1-z_p)\big), \Phi^{-1}(1-z_p)\Big)$.

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

$$\frac{2(c+L)\Phi(s)}{d^3}+\frac{2(c+L+\beta B)\left(1-\Phi(s)\right)}{d^3(1-\eta)}>0\,, \tag{59}$$

$$\left[-\frac{c+L}{d}+\frac{c+L+B}{d(1-\eta)}\right]\Phi''(s)>0\,, \tag{60}$$

$$\begin{gathered}2\left[\frac{(1-\eta)(c+L)\Phi(s)+(c+L+B)\left(1-\Phi(s)\right)}{d^4(1-\eta)^2}\right]\left(\eta(c+L)+B\right)\Phi''(s)+\\ \frac{\left(\eta(c+L)+B\right)^2\left(\Phi'(s)\right)^2}{d^4(1-\eta)^2}<0\,,\end{gathered} \tag{61}$$

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

Incentive compatibility of (d_p, s_p) follows from the same arguments as given in the proof of Proposition 1.

Proof of Proposition 3.

Part 1. From (51) and (58), both z^* and z_p are clearly decreasing in η, so skill thresholds s^* and s_p must be increasing in η.

Part 2. From (51) and (58), as $\eta \to 1$, $z_p \to 0$ while z^* converges to

$$\underline{z} \equiv \frac{c+L-\beta B+\sqrt{\left(c+L-\beta B\right)^2+4(c+L)kuF(1-F)\left(\dfrac{\beta B\phi}{\ell}\right)^2}}{kuF(1-F)\left(\dfrac{\beta B\phi}{\ell}\right)^2}>0\,. \tag{62}$$

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

Proof of Proposition 4.

To compare d_p and $d*$, first invert $D(s)$ in (46) and substitute into first-order condition (20) to obtain the following condition in *d*:

$$R(d) = R_0 + R_1 d + R_2 d^2 = 0\,, \tag{63}$$

where

$$R_0 = uF(1-F)(c+L)\,, \tag{64}$$

$$R_1 = \ell\left[\frac{(c+L-\beta\eta B)-(1-\eta)^2(c+L)}{\beta B\phi(1-\eta)}\right], \tag{65}$$

$$R_2 = -k. \tag{66}$$

Similarly, invert $\underline{D}(s)$ in (53) and substitute into the planner's first-order condition, (22), to obtain the condition

$$\underline{R}(d) = \underline{R}_0 + \underline{R}_1 d + \underline{R}_2 d^2 = 0\,, \tag{67}$$

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

$$\underline{R}_1 = \frac{\ell}{\phi}\,. \tag{68}$$

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

$$\begin{gathered}\left(2k(1-\eta)\right)^{-1}\times \\ \left[\left(\frac{\ell}{\phi}\right)\left(\frac{(c+L-\beta B\eta)-(1-\eta)^2(c+L)}{\beta B}\right)+ \sqrt{\left(\frac{\ell}{\phi}\right)^2\left(\frac{(c+L-\beta B\eta)-(1-\eta)^2(c+L)}{\beta B}\right)^2+4kuF(1-F)(c+L)(1-\eta)^2}\right],\end{gathered} \tag{69}$$

and

$$d_p = \left(2k\right)^{-1}\left[\left(\frac{\ell}{\phi}\right) + \sqrt{\left(\frac{\ell}{\phi}\right)^2 + 4kuF(1-F)(c+L)}\right]. \tag{70}$$

Part 1. From (70), d_p does not depend on η. From (69), $d*$ grows as

$$\tilde{d} = \left(k(1-\eta)\right)^{-1}\left[\left(\frac{\ell}{\phi}\right)\left(\frac{(c+L-\beta B\eta)}{\beta B}\right)\right], \tag{71}$$

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

Part 2. From (71), $\tilde{d} \to \infty$ as $\eta \to 1$, whence $d*$ also diverges.

Proof of Proposition 5.

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

Part 1. From first-order condition (21), the rate of skilled identity theft in the symmetric equilibrium is

$$\frac{\ell\left(1-\Phi(s*)\right)}{\beta B\Phi'(s*)} = \frac{\ell}{\beta B\phi}. \tag{72}$$

Similarly, the rate of skilled identity theft in the golden-rule allocation can be calculated using (23):

$$\frac{1-\Phi(s_p)}{d_p(1-\eta)} = \frac{\ell}{\phi\left[\eta(c+L)+\beta B\right]}. \tag{73}$$

Comparing (72) and (73), skilled identity theft must be lower under the golden-rule allocation.

Part 2. The rate of unskilled identity theft in the symmetric equilibrium is given by $\Phi(s*)/d*$. From the Propositions 3 and 4, $\Phi(s*) \to \Phi(\overline{s}) > 0$ and $d* \to \infty$ as $\eta \to 1$, implying that unskilled identity theft is driven to zero as $\eta \to 1$.

The rate of unskilled identity theft under the golden-rule allocation is given by $\Phi(s_p)/d_p$. From the proof of Proposition 2, $\Phi(s_p) \to 1$ as $\eta \to 1$ 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 $\eta \to 1$.

Part 3. The calculations in parts 1 and 2 show that, as $\eta \to 1$, $\rho(d^*, s^*) < \rho(d_p, s_p)$ iff

$$\frac{\ell}{\phi\beta B} < \frac{1}{d_p} + \frac{\ell}{\phi(c+L+\beta B)}. \tag{74}$$

Substituting for d_p from the proof of Proposition 2, inequality (74) reduces to

$$\frac{2\phi^2}{\ell + \sqrt{\ell^2 + 4(c+L)k\phi^2}} > \left(\frac{\ell}{k}\right)\frac{c+L}{\beta B(c+L+B)}, \tag{75}$$

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

Calculations for Section 4.1:

(Again we simplify notation by setting $uF(1-F) = 1$.)

Proof of Proposition 6.

A regulator who can only determine data length sets *d* to maximize $V^f(d,s)$ subject to the clubs' first-order condition in *s*, which in symmetric equilibrium is given by (21) or equivalently (25). Using (21) we can eliminate *s* and simplify the regulator's problem to the following: choose *d* to minimize the steady-state fraud costs, i.e., choose d_c to maximize

$$V^f = -kd + \frac{\ell}{\phi}\ln d - \frac{c+L}{d} + <\text{constant terms}> \tag{76}$$

which has solution $d_c = d_p$.

Evaluating (25) at $d_c = d_p$ and comparing to (26), it follows that $s_c < s_p$. From (25) and the fact that $d_p < d^*$ (Proposition 4), it follows that $s_c > s^*$.

Proof of Proposition 7.

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

$$\frac{(c+L)(1-z)}{d} + \frac{(c+L+\beta B)z}{d(1-\eta)} + kd - \frac{\ell}{\phi}\ln z\,, \tag{77}$$

subject to (27), which we write as $d = G(z)$ where

$$G(z) = \frac{1}{(1-\eta)\sqrt{k}}\left[(c+L)(1-\eta)^2(1-z) + (c+L-\beta\eta B)z\right]^{\frac{1}{2}}. \tag{78}$$

This regulator's problem may be compared to the planner's problem, which is equivalent to minimizing (77) over $z \in (0,1)$ subject to (28), which we write as $d = P(z)$ where

$$P(z) = \frac{1}{\sqrt{(1-\eta)k}}\left[(c+L)(1-\eta)(1-z) + (c+L+\beta B)z\right]^{\frac{1}{2}}. \tag{79}$$

Substituting (78) into (77) and simplifying, the regulator's problem is to minimize

$$k\frac{\left(P(z)\right)^2}{G(z)} + kG(z) - \frac{\ell}{\phi}\ln z\,, \tag{80}$$

and substituting (79) into (77), the planner minimizes

$$k\frac{\left(P(z)\right)^2}{P(z)} + kP(z) - \frac{\ell}{\phi}\ln z = 2kP(z) - \frac{\ell}{\phi}\ln z\,. \tag{81}$$

From Proposition 2, (81) is minimized at $z = z_p$. Also, as $\eta \to 1$, $P(z)/G(z) \to e$ where

$$e = \sqrt{\frac{(1-\eta)(c+L+\beta B)}{c+L-\beta\eta B}}\,, \tag{82}$$

and the regulator's minimand (80) is approximately

$$k(e+e^{-1})P(z)-\frac{\ell}{\phi}\ln z\,. \tag{83}$$

Since , as $\eta \to 1$, $(e+e^{-1})>2$, it follows that $z_C < z_p$ and $s_C > s_p > s^*$. From (27) it then follows that $d_C < d^*$.

To compare d^* and d_p, note that from (83), we can approximate z_C by replacing ϕ in equation (58) with

$$\phi' = \frac{2\phi}{e+e^{-1}}\,, \tag{84}$$

from which we deduce that z_C goes to zero at a rate $(1-\eta)^{\frac{3}{2}}$ as $\eta \to 1$. From (27), d_C^2 is given by

$$k^{-1}\left[(c+L)(1-z_C)+\frac{(c+L-\beta\eta B)}{(1-\eta)^2}z_C\right] \tag{85}$$

Now, the first term of (85) converges to $c+L$ as $\eta \to 1$ while the last term eventually grows as $(1-\eta)^{-\frac{1}{2}}$, whence $d_C > d_p$.

References

Anderson, K.B., E. Durbin, and M.A. Salinger, 2008. "Identity Theft," *Journal of Economic Perspectives* 22, 171-192.

Anderson, R. and T. Moore, 2006. "The Economics of Information Security," *Science* 314, 610-613.

Bank for International Settlements, Committee on Payment and Settlement Systems, 2007. *Statistics on Payment and Settlement Systems in Selected Countries.* Bank for International Settlements, Basel.

Başar, T. and G.J. Olsder, 1998. *Dynamic Noncooperative Game Theory.*

Boyd, J.H. and E.C. Prescott, 1987. "Dynamic Coalitions: Engines of Growth," *American Economic Association Papers and Proceedings*, 63-67.

Cavalcanti, R. and E. Nosal, 2007. "Counterfeiting as Private Money in Mechanism Design." Working paper, Federal Reserve Bank of Cleveland.

Chandler, J.A., forthcoming. "Negligence Liability for Breaches of Data Security," *Banking and Finance Law Review.*

Cheney, J., 2005. "Identity Theft: Do Definitions Still Matter?" Discussion Paper, Federal Reserve Bank of Philadelphia Payment Cards Center.

Coggeshall, Stephen, 2007. "ID Theft Knows No Boundaries," *eCommerce Times*, April 13, *www.ecommercetimes.com/story/56864.html*.

Diamond, P., 1990. "Pairwise Credit in Search Equilibrium," *Quarterly Journal of Economics* 105, 285-319.

Douglas, D.D., 2008. "Merchant Liability for Payment Card Security Breaches," *Electronic Banking Law & Commerce Report* 13, 1-7.

Green, E. J. and W. Weber, 1996. "Will the New $100 Bill Decrease Counterfeiting?" *Federal Reserve Bank of Minneapolis Quarterly Review* 20(3), 3-10.

Grossklags, J., N. Christin, and J. Chuang, 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., J. McAndrews, and W. Roberds, 2005. "Money is Privacy," *International Economic Review* 46, 377-400.

Kahn, C. M. and W. Roberds, 2008. "Credit and Identity Theft," *Journal of Monetary Economics.*

Kahn, C. M., and W. Roberds, forthcoming. "Why Pay? An Introduction to Payments Economics." *Journal of Financial Intermediation.*

Kiyotaki, N. and 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.

Luenberger, D.G., 1969. *Optimization by Vector Space Methods*. John Wiley & Sons, New York.

Martin, A., M. Orlando, and D. Skeie, 2008. "Payment Networks in a Search Model of Money," *Review of Economic Dynamics* 11, 104-132.

Monnet, C., 2005. "Counterfeiting and inflation," working paper, European Central Bank.

Monnet, C. and W. Roberds, forthcoming. "Pricing Payments as an Alternative to Cash," *Journal of Monetary Economics.*

Nosal, E. and N. Wallace, 2006. "A model of the (threat of) counterfeiting," *Journal of Monetary Economics* 54, 994-1001.

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.

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 http://people.ischool.berkeley.edu/~hal/Papers/japan/.

Varian, H., 2004. "System Reliability and Free Riding." Available online at http://people.ischool.berkeley.edu/~hal/Papers/2004/reliability.

Williamson, S.D., 2002. "Private Money and Counterfeiting," *Federal Reserve Bank of Richmond Economic Quarterly* 88(3), 37-57.