"The Uses and Limits of Financial Cryptography: A Law Professor's Perspective"

Peter P. Swire
Ohio State University
College of Law
Columbus, Ohio, USA
swire.1@osu.edu
www.osu.edu/units/law/swire.htm

Abstract:

There is considerable support in the cryptography community for the "Cypherpunk Credo," defined as: "Privacy through technology, not legislation." Much discussion to date has assumed that the U.S. government's opposition to strong cryptography, such as its key escrow proposals, is the primary obstacle to widespread use of anonymous electronic cash. For purposes of this paper, I assume that strong cryptography is legal and readily available. Even in that event, I claim that strong cryptography will be used to preserve anonymity only in a highly restricted subset of financial transactions. Furthermore, because technology often will not assure privacy, legal rules can and should play an important supplementary role in the protection of privacy in financial transactions.

Introduction.

The use of electronic payments will spread widely in coming years. We expect to be able to buy products easily from home over the Internet. Face-to-face transactions will increasingly be done by debit card, credit card, and emerging forms of smart cards. Within the cryptography community, an ongoing research project has been to devise ways of assuring privacy while performing electronic financial transactions. The well-known concern is that many systems of electronic payments would create detailed databases describing each user's transactions. In such systems, users would gain the advantages of electronic payments. But other parties would be able to assemble comprehensive dossiers about their spending patterns.

Faced with this challenge, the cryptographic community has proposed numerous electronic cash (e-cash) systems for enabling users to make anonymous, untraceable payments. [3] The goal has been to create mathematically-rigorous systems that will prevent even the most determined attackers from discovering the user's identity. These systems rely on "strong cryptography" -- a term that for our purposes means that no outside party can crack the code within a useful amount of time. Strong cryptography can allow achievement of the "Cypherpunk Credo": "Privacy through technology, not legislation." [11] As Ian Goldberg states: " The law of the land can be changed by the next administration. The laws of mathematics are more rigid." Id.

From the perspective of the Cypherpunk Credo, there has been one major villain standing in the way of anonymous e-cash systems -- the United States government, and its advocacy of the Clipper Chip and related key escrow schemes. [5] Key escrow means that the government would have access to the secret keys that permit a message to be read. In a world of mandatory key escrow, the government might have the power to read any message, and to trace any encrypted financial transaction to the user. Privacy then would not depend on mathematics -- on the power of encryption. It would depend instead on legislation -- on the legal safeguards that are supposed to prevent the government from abusing its power.

The discussion to date within the cryptography community has thus relied on two premises: (1) strong cryptography and anonymous e-cash are important goals; and (2) the government's opposition to

strong cryptography, such as its key escrow proposals, is the primary obstacle to use of anonymous e-cash. This paper takes issue with the second premise. For purposes of this paper, I *assume* that strong cryptography is legal and readily available. My principal claim is the following: even if we assume that strong cryptography is available, it will be used to preserve anonymity only in a highly restricted subset of financial transactions. My second claim follows from the first: because mathematics often will not assure privacy, the law can and should play an important supplementary role in the protection of privacy.

The paper proceeds as follows. Part I briefly discusses some valuable uses of financial cryptography. Many of the most important uses are to protect the security of the transaction from malicious outsiders. Cryptography is already widely used in banks to assure security, and such uses will likely spread enormously in coming years.

I suggest, however, that we have less reason to think that cryptography will spread nearly as widely in order protect the user's anonymity. Part II provides a detailed examination of three reasons that the use of anonymous financial transactions will be less extensive than the cryptographic community has generally hoped. First, the entire array of lending transactions will be difficult or impossible to conduct with anonymous borrowers. Second, key management poses more fundamental problems for maintaining anonymity than has usually been recognized. Third, the use of cryptography will face a daunting array of market acceptance problems. The most important of these may be the difficulty of explaining to ordinary consumers why they should take on the extensive burdens of protecting their own anonymity.

Part III continues the examination of whether technology, rather than law, will protect privacy. There are additional constraints on the use of cryptography to protect anonymity. Once we relax the assumption that strong cryptography will be legal and readily available, it is possible to imagine significant ways that governments will interfere with the use of anonymous payments. An additional question is whether users of anonymous digital signatures will face disadvantages compared with users who reveal their identity. If so, there will be yet another reason not to expect widespread use of anonymous transactions.

If users do not generally employ cryptography to assure their anonymity, then the question arises whether there are legal or other ways to protect their privacy. I submit that the cryptography community has been scared away from legal solutions because of its experience with the Clipper Chip and related controversies. [5] Where users deploy strong cryptography, legal rules can only have the effect of reducing privacy. The situation is very different, however, where users do not use cryptography to protect their anonymity. In such cases the unregulated market can easily lead to a great deal of disclosure of sensitive information. Legal rules then can play a constructive role in the protection of privacy.

Part IV sketches the main points of how to think about the legal protection of personal financial information. Where technology itself will not protect privacy, I propose a contracts approach, which will be developed at greater length in a forthcoming article.

I. Some Uses and Limits of Cryptography in Protecting Privacy and Security.

For a cryptographic audience, there is little need to explain the wide and growing range of uses for cryptography for protecting financial information. For instance, cryptography is already used in many bank-to-bank transactions. As Internet commerce grows, most people assume that more cryptography will be deployed between users and their banks -- the risks of plaintext transmission for huge numbers of financial transactions can be readily reduced with cryptographic systems. Over time, cryptography will be increasingly used within organizations for security purposes, such as by requiring strong authentication before granting access to sensitive databases. In addition, many people in the cryptography community hope that strong cryptography will become the norm in anonymous e-cash systems.

In order to place the uses of cryptography in broader perspective, I have created a stylized chart of four different types of financial transactions, discussed from the most private (bottom left) to the most public (top right). One simple task here is to remind us of the wide array of transactions in which people do not currently use cryptography. The chart will also help us begin to assess where cryptography is likely to be deployed in the future.

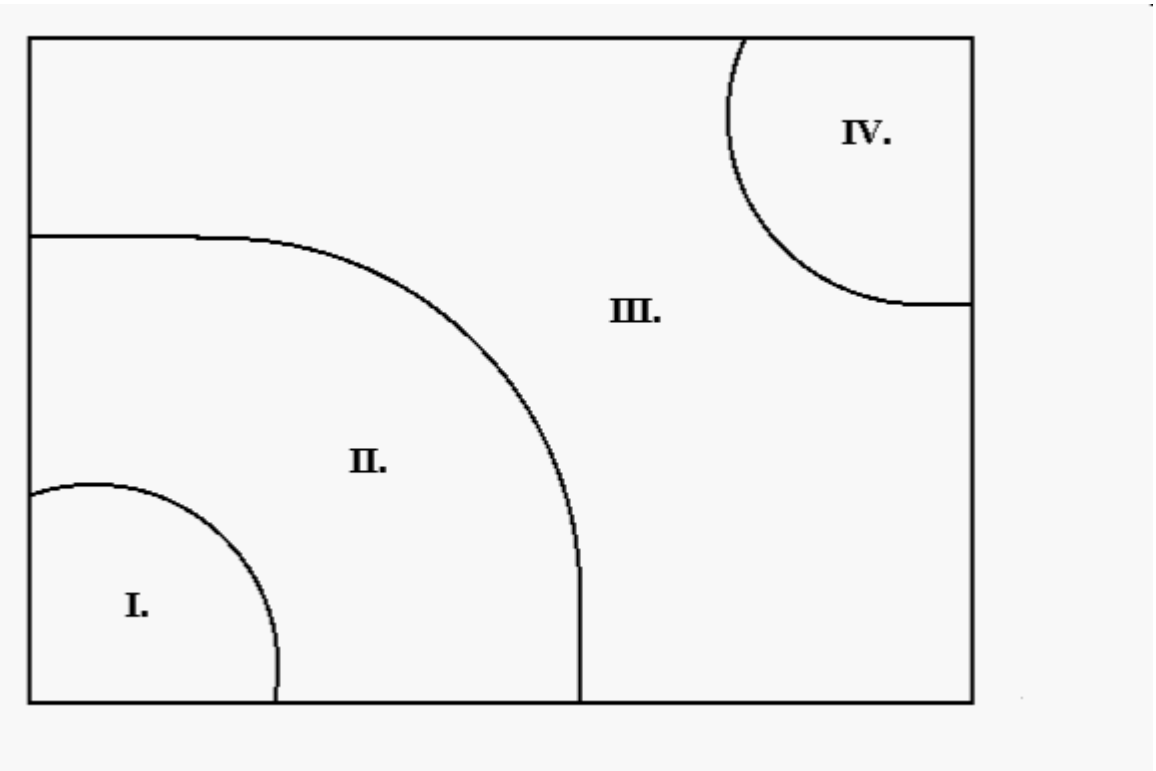

In the chart, Area I represents transactions in which strong cryptography is fully deployed. An example would be a payment made through an e-cash system designed by David Chaum. [3] For these transactions, cryptography is used to create anonymity for the purchaser -- no outside party after the fact can learn the purchaser's identity. Cryptography is also used for security -- if an enemy intercepts the transaction, the enemy will not gain any useful information about the purchaser. Nor will the enemy be able to use the intercepted transaction to get money or any valuable information out of the payment system. I suggest that much of the cryptography debate to date has implicitly assumed that the important issues lie within Area I. The research agenda to date has focused on how to implement payment systems that fully assure both anonymity and security.

As we move to Area II, the purchaser no longer uses cryptography to hide his or her identity from *everyone* else. For instance, a typical consumer "Alice" will make payment through some trusted other party, often a bank with whom she has an account. After a transaction occurs, the bank will indeed have the ability to link Alice with the transaction. In Area II, however, cryptography will be used to assure security -- to prevent the attacker from benefitting from any interception of the transmission. A simple example of an Area II transaction is the use of a credit card over the Internet, where Alice's account number is strongly encrypted. In this example, Alice will have no privacy from her credit card company. That company will know how to link the purchase to Alice's account, so that she can be properly billed for her purchases. Cryptography helps provide security, however. In this instance, cryptography reduces the ability of hackers to profit from the stealing of credit card numbers.

In Area III, cryptography is not deployed for either privacy or security. There is no strong privacy because at least one party, such as the bank, can link the purchaser with the transaction record.

Nor is cryptography used for security. Other parties can see information about the transaction, such as the name of the purchaser, the name of the seller, the purchase amount, the purchaser's account number, and the purchaser's signature. Such information is routinely available in checking and credit card transactions today. The store clerk receiving the check, or the waiter who brings the credit card into the back room, has ample opportunity to copy and misuse the purchaser's signature and account number.

To those schooled in cryptography, the use of Area III transactions may seem hopelessly unprotected. The cryptographer might even question whether any sensible person would participate in such a transaction. In answer, we can observe simply that millions of these transactions occur every day. It is instructive to consider why these transactions in fact take place. Consumers, at least in the United States, can participate in unencrypted transactions in part because of *laws* that limit their losses from theft or fraud. For instance, U.S. consumers face a $50 maximum loss from theft of their credit card, and significant consumer protections also exist for theft of debit cards and checks. For banks and merchants, there are significant losses due to the theft of credit cards, debit cards, and checks. These losses are manageable, but they create an incentive to move toward a more secure system, such as is available in Area II. Banks and merchants also benefit from helpful laws -- bank employees and store clerks are deterred from misusing the information for fear of losing their jobs and facing legal punishment.

Area IV transactions are those in which there is quintessentially the absence of privacy -- public records. The most familiar example is a land sale, for which the purchaser is usually listed at the local courthouse. Area IV is similar to Area III in that cryptography is absent from the transaction. Area IV is different from Area III because of the expectation that public records will indeed be made public. For Area III transactions, even though cryptography is not used to hide information, there is an expectation that the bank employee and store clerk will *not* publicize the transactional information.

The discussion of the chart included here is intended to help us begin to assess the likely future uses of cryptography. We note that consumers today participate in many insecure, un-private transactions. Current laws against theft and fraud provide important protections for consumers in such transactions. Banks and merchants, however, are not as well protected against loss. They consequently have strong incentives to develop more secure forms of electronic transactions, notably including encryption. Our analysis thus gives us good reason to expect widespread adoption of cryptography for Area II transactions. What is far less clear, and what is discussed in the next section, is whether we should expect any similar widespread adoption of cryptography for Area I transactions.

II. Some Limits on Anonymity Even in a World of Strong Cryptography.

This section of the paper will discuss three important constraints on the use of anonymity in financial transactions, even where strong cryptography is legal and available. I am not aware of previous attempts to analyze the conditions under which loans can be made to anonymous borrowers. Key management and market acceptance problems have been more widely discussed in the cryptographic community. The intent here is to use my perspective from law and economics to deepen the analysis of how key management and market acceptance problems are likely to result in limited uses of anonymous financial accounts.

II.A. The Difficulty or Impossibility of Lending to Anonymous Borrowers.

The focus of financial cryptography to date has been on debit systems, in which the customer spends down an existing amount of his or her money. Let us direct our attention, instead, to the wide universe of credit transactions. There is an indefinitely large number of types of lending transactions, from consumer loans through small-business loans to public securities offerings. There are also many different reasons to borrow money. A few reasons include: to relieve temporary shortages of cash; to finance long-term mortgages and other consumer debt; to finance capital investment; and to allow

increased profitability through leverage. In considering this wide range of lending transactions, my claim is the following: most or all lending transactions will be impossible with anonymous borrowers.

The core idea in a credit transaction is that the bank or other lender makes more money available than the borrower has deposited with the bank. Once we understand this simple point, some sorts of "loans" turn out to be debit relationships under another name. A simple example is the secured credit card. Under such an arrangement, the customer might place $1,000 on account at the bank. The bank then allows the customer to have a "credit card," which conveniently happens to have a $1,000 spending limit. Such an arrangement may be quite useful to the customer, who can rent a car and do other transactions that typically require a credit card. But the bank is not lending anything to the customer. As an economic matter, the customer has placed $1,000 on deposit with the bank, and withdrawals are done by credit card rather than by check.

Another strategy for borrowing anonymously could be for the borrower, Alice, to adopt a pseudonym, such as Paula or the XYZ Corporation. My claim is that little or no lending will be done to "Paula." To see why, consider the obvious point that loans are made to borrowers who can assure a bank that it will be repaid. The most common sorts of assurances are to offer security to the bank or to have a good credit rating. We will consider each of these in turn.

(1) The difficulty of offering security anonymously. The usual way that Alice offers security is to allow the bank to repossess her house or other security if she fails to pay the loan. For a home mortgage, the lender requires an appraisal of the worth of the house. The appraisal includes an evaluation of property values in the neighborhood and of the physical properties of the house. Such an appraisal simply cannot be done unless someone besides Alice knows which house is securing the mortgage. There will be no anonymous home loans.[1]

It is not strictly necessary for the bank to know Alice's identity, but I suggest that the bank almost always will. One can imagine that Alice agrees to reveal herself to a third party that she trusts. The third party can perform the appraisal and report the results to the bank. The bank may even trust the same appraiser that Alice has trusted. In this way, one can imagine that Alice can get her home loan without the bank knowing her identity. Notice, however, how little this elaborate effort helps Alice. Most importantly, she has revealed her identity to some other party -- she received the loan only by cooperating with the appraiser and losing her anonymity. Moreover, she accomplishes this loss of anonymity at a significant price. The bank, forced to rely on the appraiser and unable to see the property for itself, will add a risk premium to the loan. That is, the bank will charge a higher interest rate for the anonymous appraisal than it would for a property that it could appraise for itself. Because Alice is going to reveal her identity in any event, she will have a strong incentive to reveal it directly to the lending bank, so that she can get the loan at the lowest price. Taking a pseudonym will make it expensive or impossible to offer security to get a loan.

(2) Credit ratings and the difficulty of pseudonymous loans. Taking a pseudonym will also make it expensive or impossible to create a good credit rating to get a loan. After all, how can a lender evaluate your credit rating when it cannot find out about your previous defaults and bankruptcies? Masking your true identity from a lender raises the strong possibility that you are adopting a pseudonym precisely in order to conceal your actual history as a deadbeat.

Within the cryptographic community, one response would be that a pseudonym can build up a good reputation over time, without anyone necessarily knowing the identity of the person that is using the pseudonym. On this theory, perhaps lenders might be able to lend to those pseudonyms that develop their own good credit history. The possibility of such lending becomes even greater if lenders can develop a statistical profile of the average default rate in the population of persistent pseudonyms. Once the behavior of the population is known, it would seem that lenders could then charge for the level of risk, making widespread pseudonymous loans possible.

Although loans might theoretically be made on this basis, three major reasons suggest that such loans will remain rare. First, the lender cannot control who applies for credit, leading to what economists call an "adverse selection" problem. Second, lenders face a severe "end-game" problem, in which the borrower can profit by absconding with large sums of the lender's money. Third, the practices in today's lending markets suggest that such pseudonymous loans are less likely to succeed than proponents would hope.

(a) Adverse selection. The adverse selection problem is most familiar in the insurance context. Imagine that Equal Price Auto Insurance (EPAI) offered the same rate for all drivers -- the rate pegged to the average losses experienced by all drivers. What will happen? Risky drivers with bad records will rush to buy the company's insurance, attracted by what for them is a low rate. Soon, the company would begin paying out high claims for all of the bad drivers. If the company insists on charging the same rate to all customers, the rates will have to rise for everyone in order to pay off the high claims. Pretty soon the good drivers will become upset. These drivers deserve low rates in the marketplace, and they will shift to other companies that charge low rates to low-risk drivers. Over time, EPAI will be stuck with lots of bad drivers and few good drivers. Customers will self-select, with adverse consequences to the company.

Adverse selection will also happen with anonymous lending. Suppose a credit card company begins to offer cards to anonymous borrowers, and sets the interest rate at the industry average. Who will apply for these credit cards? Some of the borrowers may be creditworthy Internet citizens who are delighted at the chance to have an anonymous credit card. I suggest, however, that people with bad credit histories will have the dominant incentive to get the cards. These people currently lack the ability to get a credit card, or else they pay high rates that reflect their bad credit histories. The chance to pay lower interest rates will itself be a great temptation. So will the realization that there is no effective enforcement against deadbeat borrowers -- the borrowers, after all, are anonymous and untraceable.[2] The lack of effective enforcement might be especially tempting to people with bad credit histories, who already have personal knowledge about the unpleasant effects of bankruptcy or other enforcement actions against borrowers.

(b) End-game problems. The discussion of adverse selection shows the very real possibility that any program of anonymous borrowing will be flooded by the worst credit risks. In response, an advocate of anonymous borrowing might emphasize the need for a pseudonymous borrower to build up a good credit history over time. On this theory, a lender might profitably be able to screen out enough of the bad credit risks by initially offering only secured credit cards or small credit limits. Many of the worst credit risks will soon prove inept at paying their accounts. They will default soon enough so that the lender will experience manageable losses. The lender can then make a good profit on the accounts that gradually build up a good pseudonymous credit rating.

I agree that lending is at least theoretically possible along these lines, but only under very restrictive conditions that are not likely to be important in practice. First, notice that the scenario assumes that the bad credit risks will quickly prove inept at paying their accounts, so that the bank can concentrate its loans on good credit risks. If the bad credit risks only manifest themselves gradually, the lender may still be stuck with unacceptable credit losses. More precisely, the lending will only be profitable if the average profits during the period of creditworthiness outweigh the average losses once the borrower defaults.

A second necessary condition is that the lending must be robust in the face of systematic attack. The lender needs to be highly concerned about the "end-game" problem. The end-game problem arises when a borrower establishes a good pseudonymous credit rating over time, lulling the lender into believing in the borrower's creditworthiness. Over time, the lender will raise the credit limit to its good customer, perhaps to $10,000. When the limit is high enough, the borrower can quickly spend the full $10,000, effectively absconding with that amount. In this scenario, the malicious borrower can

systematically profit at the expense of the lender, if it costs the borrower less than $10,000 to establish a good credit rating.

In contemplating a world of anonymous lending, imagine what will occur once organized crime learns of the opportunity to play this sort of end game. A well-financed organization (perhaps including a government that is hostile toward the lender's country) can set up an indefinitely large number of pseudonymous credit-card accounts. The accounts will not be linkable to each other. The organization then builds up good credit histories on each of the accounts, perhaps varying its behavior a bit so that a pattern does not become too obvious. When credit limits become high enough, the organization can play an end game on each account, giving itself a guaranteed profit on each account. The scheme will fail only under two conditions: (1) where the lender can somehow screen hostile account-holders from getting access to anonymous accounts; or (2) where it costs the borrower more to build up a good credit rating than the borrower can gain from an end-game spending of the credit limit.

(c) Lessons from current practice. Attention to current practices suggests additional reasons to believe that there will be little or no lending to anonymous borrowers. To avoid the end-game problem, we have stated that it must cost more for the borrower to establish a credit rating than the borrower can seize in the last period. From my observation of today's credit-card industry, however, it would take much less than a $10,000 investment to establish a credit limit of $10,000. In today's market, credit limits increase if a borrower simply spends and repays a moderate amount of money for a moderate amount of time. Lenders get a large share of their profits when borrowers have to pay interest on substantial unpaid balances. Lenders thus have a strong incentive to increase credit limits, as shown by the efforts they make to have customers consolidate balances on one card.

Lenders to identifiable borrowers have crucial advantages over anyone that lends to pseudonymous borrowers. For pseudonymous borrowers, lenders can observe only the transactions in that one account. By contrast, lenders to identifiable borrowers can see the borrower's *entire* credit rating. They also have the ability to enforce against the borrower's assets. These lenders can learn a good deal about their customers, and so do not suffer the end-game problem to nearly the same degree as would pseudonymous lenders. Lenders to identifiable borrowers thus face much lower risks. In a marketplace full of thousands of competing credit cards, the advantages to lenders translate to advantages for borrowers. Borrowers who identify themselves can expect to receive much more favorable terms. In short, it seems unlikely that anonymous credit will be attractive enough to tempt creditworthy borrowers and strict enough to exclude the professional thieves.

One other lesson from current practice shows a different reason for suspecting that persistent pseudonyms will be less effective at borrowing than advocates might hope. Consider the lending practices to the most important type of persistent pseudonym, the corporation. Large corporations are able to borrow money from public securities markets, typically after detailed disclosure of their activities and of the names of their principal officers. Smaller corporations often fund themselves with bank loans. A standard part of these loans, however, is that they are done *with recourse*. That is, identifiable individuals often have to make themselves personally responsible for the loans, in the event that the persistent pseudonym (the corporation) does not pay in full. From the lender's perspective, requiring the personal guarantees serves two important purposes: (1) it adds the individual's ability to pay to the corporation's; and (2) it creates a strong incentive for the individual to try to make the corporation successful, so that the individual's personal assets will not be lost. The widespread use of loans with recourse suggests the difficulty that will face any borrowers who wish to operate under a persistent pseudonym.

In conclusion, the discussion here provides a number of compelling reasons why anonymous lending typically will not be profitable, and so will not occur with any frequency. Put somewhat differently, the analysis here states the necessary conditions for anonymous or pseudonymous borrowing. These conditions notably include overcoming the adverse selection and end-game problems.

II.B. The Practical Difficulties of Good Key Management

As we proceed in our exploration of the uses and limits of financial cryptography, the previous section discussed why there is likely to be little or no lending to anonymous borrowers. The next topic concerns key management -- how users will take care of their keys, and what will happen if the keys are lost. In a world of strong cryptography, losing the key is by definition catastrophic -- "strong" cryptography means that no one can break the code. If digital money is encrypted in your computer, and you lose the key, then you lose the money, too.

The importance of key management is well recognized in the cryptographic community. Bruce Schneier, for instance, states: "In the real world, key management is the hardest part of cryptography." [10, p. 169] Schneier perceptively analyzes a range of challenges in key management, and proposes technical solutions to problems that arise in the use, updating, and storage of keys.

While acknowledging the technical difficulties of key management, the cryptographic community has usually given less recognition to certain institutional aspects of the topic. My point here is to emphasize the reasons that individuals and firms are likely to give up anonymity because of the difficulties of key management. Even if the technical problems are solved, there are personal and institutional problems in key management, which mathematics cannot solve.

Before examining these institutional problems, I would like to add a caveat. The topic of key management has become intensely politicized in the wake of the Clipper Chip and related proposals for the government to be entitled to a user's keys. My conclusions here might be seen within the Internet community as giving aid and comfort to government proposals for mandatory key escrow. So let me be explicit. I do not intend to take any position here on the desirability of mandatory key escrow. Instead, the analysis here tries to show the strong incentives that individuals and firms engaging in financial transactions will have to provide their keys to *someone* else, that is, to give up their anonymity.

(1) Key Management for Individuals -- The Myth of Omnicompetence. Before turning to corporate uses, let us first consider how ordinary individuals are likely to manage their keys. The basic answer is: not well. As individuals begin to use strong cryptography for financial transactions and digital signatures, they will begin to learn about the importance of protecting their keys. They will learn that losing the key will mean that the encrypted material, including electronic cash, is unrecoverable. They will learn that revealing the key for their digital signature will allow malicious third parties to forge the individuals' signatures in a potentially vast array of unwanted ways. In short, the costs to individuals of having a private, secure key are potentially enormous.

One response to these observations is simply to encourage people to take responsibility for their own lives. On what we might call the cyber-libertarian view, individuals who desire freedom should learn cryptography. Strong cryptography is liberating precisely because it empowers the individual to protect his or her personal information, even in the face of determined government attack. On this view, the costs of keeping keys private are more than offset by the gains in individual freedom and the limits on the ability of governments and corporations to learn about and manipulate individuals.

Let us grant, for discussion purposes, the attractiveness of this cyber-libertarian view. (Those who do not agree with the viewpoint will be even less likely to bear the costs of keeping a private, secure key.) My point, again, is descriptive. I think that many individuals will be so lousy at protecting their keys that a small percentage, at most, will actually maintain private, secure keys.

To get a sense of the difficulty of key management, let us look at Bruce Schneier's chapter on "Key Management" in his standard reference work on cryptography. [10] By recalling just some of

Schneier's recommendations for effective key management, we can get a feel for the sorts of challenges facing ordinary users who wish to have private, secure keys.

(a) *Key backup.* With strong cryptography, if you lose the key, you lose the data. So keeping a backup of the key is incredibly important. Schneier himself admits to losing a key every five years or so. Most of the rest of us are not likely to outdo Schneier, so we will need to have an excellent system in place for protecting backups.

(b) *Good key selection.* There is one obvious way to reduce the need for key backup. Alice can simply pick a key that is easy to remember, such as her mother's maiden name or the digits of her birthday. There is a slight downside -- the malicious attacker can often guess the key. Schneier reports the efforts of Daniel Klein, who developed a "dictionary attack" for guessing common keys. Klein was able to guess 40 percent of the passwords on the average computer using his approach.

(c) *Secure physical protection of keys.* Some people have a wonderfully obscure key, but keep it written in their wallet or taped under their desk. Not much help when the bad guys come and get you.

(d) *Multiple keys.* Alice takes on great risk if all of her activities are conducted using the same key. Schneier recommends having multiple keys, for different types of uses. If Alice has one key, and it is ever compromised, then all of her encrypted information is simultaneously made open to attack. Just imagine if the key is posted on a Web site or otherwise made public. Alice may have left encrypted files in computers anywhere in the wide world, outside of her current control. Of course, if Alice has multiple keys, she has to have good backup and key selection and physical protection for each of them.

(e) *Updating keys.* Schneier writes: "No encryption key should be used for an indefinite period." [10, p. 183.] Permanent use of a key has a number of disadvantages: greater risk of compromise; greater loss if the key is compromised; greater temptation for someone to spend the effort necessary to break the key; and greater ease of doing cryptanalysis when more text is encrypted with the same key. For Alice, these problems mean that she should have a well-organized routine for updating her keys. Then again, she needs a good routine for protecting her old keys, so that she can use them to open up the old files that were encrypted with them.

We could go on, but the general picture should now be clear enough. It is just plain *hard* for the individual to set up good key management practices. There is an important distinction between ordinary consumers and cryptographic adepts. For hackers and professional cryptographers, key management can be kind of fun. Jude Milhon claims that "the chief characteristic of hackers is wily intelligence." [11] Using wily intelligence on key management can be entertaining, as well as good practice in thinking about the flaws in other people's security systems. For the adept, practicing good key management is one more sign of their outstanding competence at things cryptographic.

The situation is entirely different, however, for most ordinary consumers. For non-cryptographers, key management à la Schneier would be, at best, a colossal annoyance. For many people, keeping a key private would be downright impossible. Imagine a Schneier-type system for an ailing grandparent, a child, a homeless person, or for the many individuals who happen to be chronically disorganized. Imagine how hard good key management will be when we know that huge numbers of people cannot even keep track of their 4-digit PIN!

Here and elsewhere, I suggest that the cryptographic community has fallen prey to what I call the "myth of omnicompetence." Wonderful things can become available to the anonymous Alice -- so long as she is splendidly competent at a wide variety of difficult tasks. Alice can manage her keys without a slip. She can use the anonymous e-cash system described by David Chaum, and can manage elaborate protocols for assuring that her software and hardware agents are operating correctly. [3] Alice can also detect when her physical presence can be identified, and takes elaborate precautions not to allow any of her bank accounts to be linked to her physical self. In short, Alice displays the wily intelligence of the

hacker, and protects her anonymity with the zeal of an avid gamer. In considering omnicompetent Alice, I wonder what will become of bumbling Bob, an ordinary consumer with no interest in or talent for cryptography.

Put another way, the focus of much cryptographic effort to date has been in the form of existence proofs -- there exists some mathematical mechanism that can allow certain transactions to take place without the user revealing his or her identity. A great deal of progress has been made in recent years in showing the existence of financial systems that can, from a technical perspective, protect privacy. But as the use of cryptography expands from the cryptographic adepts out to the general population, existence proofs will no longer be sufficient. If anonymous transactions are to become a large fraction of all consumer transactions, key management and other tasks must come within the competence of ordinary consumers.

Are the difficulties facing consumers overstated? There are two important objections to my claim that key management and other cryptographic tasks will overwhelm many consumers and expose them to costly mistakes. First, consumers will generally not need to be as comprehensive in their protection of privacy as the state-of-the-art approach described by Schneier. Second, key management and other cryptographic tasks will become more user-friendly over time, as experience teaches us which safeguards are most important from an economic and privacy standpoint.

The answer to these two objections turns out to be the same. Let us agree that most individuals would not institute a state-of-the-art key management system for themselves. That is just the point. If Alice cannot manage her keys effectively, even after considerable hassle and risk of mistake, then she might seek a much simpler solution. She might delegate her key management tasks to someone who is expert in the subject. Alice might be very willing to give up her anonymity and reveal her identity to the right sort of third party. To foreshadow some of the discussion below, a crucial factor for Alice will be the quality of her *contract* with the party that manages her keys. She will want assurances that the third party is well-established and financially responsible for any losses that result from cryptographic problems. Alice will also probably want assurances about her privacy -- about how that party will or will not use her private transactional information. For Alice's financial transactions, she might be delighted to turn the problems of key management over to the "bank" or to whatever institution is handling her financial transactions. After all, consumers today do the same thing with their debit cards: If Alice forgets her PIN, the bank just issues her a new one.

As promised, the same analysis helps us understand what a user-friendly system will look like, at least in the eyes of many users. For those who are not adept at cryptography and key management, it will be tremendously tempting to delegate these troublesome issues to some expert institution, such as a bank.[3] If the expert institution misbehaves, the customer can sue under contract law.

(2) Key Management for Corporations: Protecting Against Principal/Agent Problems. Corporations and other institutions face the key-management challenges that individuals do. As we all know from personal experience, people do not become omnicompetent just because they work for a corporation. At least for smaller organizations, there may be no one on the payroll who is good at key management or other cryptographic tasks. Such organizations will have a strong incentive to delegate their key management to a trusted outside party -- to give up their anonymity -- rather than invest in costly in-house expertise.

Corporations also face an important additional category of risk. In the language of lawyers and economists, there are risks to the principal (the owner of the company) of having cryptography in the hands of an agent (such as the employee). Some of these costs of having an agent are obvious. The agent who knows the keys may be malicious. For instance, the malicious agent may enter the files containing electronic cash, and transfer the money to his own account. In addition, a well-meaning agent who knows the keys may be less careful than the principal would wish. In addition, the agent may become

unavailable to the principal, due to a heart attack or car crash. The principal then has no way to make use of encrypted files.

How should the principal guard against the risk of malicious, negligent, or disappearing agents? In answer, notice that the existence of agency costs has very little to do with cryptography itself. Entirely aside from cryptography, principals have forever had to worry that agents would steal the principal's secrets or property. Similarly, principals have had to live with the risk that an important agent would one day be careless, suddenly quit, or otherwise disappear. Indeed, an organizing theme of all of corporate law is how to reduce the costs of using agents: how can shareholders ever trust directors, directors trust top management, and top management trust lower-level employees? [8]

How, then, do corporations generally protect against these sorts of agency costs? One way is to try to have the incentives of the agent match the goals of the principal. A simple example is when an agent gets a commission or bonus when sales go up. For cryptography, the lesson is to look for situations where the agent no longer has the same incentives as the corporation. For instance, protections must be in place to protect against damage by disgruntled or former employees -- precisely those people who no longer care about the success of the corporation as a whole.

Even more importantly, principals try to reduce agency costs by closely monitoring agents. Typically, no one employee can authorize large checks by a corporation. Requiring multiple signatures ensures that someone else in the corporation is in a position to monitor large expenditures. More generally, corporations perform extensive audits to protect against negligence and wrongdoing by employees. The audits are typically conducted by "outsiders" -- often by persons from outside the corporation (an accounting firm), or at least from persons in a distinct part of the company (the audit division, reporting directly to top management). Securities laws and lenders often require such audits, but sensible managers perform audits even where they are not required.

Once we focus on the crucial role of monitoring in reducing agency costs, the implications on anonymous payments are substantial. It would run contrary to basic precepts of corporate governance to allow agents to act unmonitored on tasks crucial to the corporation. The corporation's keys cannot be entrusted solely to the Chief Financial Officer or Director of Management and Information Systems. Even a CFO might lose the keys, steal money, or be hit by a bus. At a minimum, corporations will need to have other persons in the firm be able to check the key management practices of whoever is managing cryptography. For at least two reasons, firms may also wish to have outside auditing of their key management systems: (1) outside parties may be needed for their expertise in double-checking the key management practices in the corporation; and (2) outside parties reduce the risk of collusion among the small number of agents in a corporation who have access to the keys. In short, the keys must be accessible at least to multiple parties within the corporation, and often to someone outside of the corporation.

In sum, corporations face many of the same problems as individuals in running a key management program that both protects secrets and dependably allows the secrets to be accessed by those who are authorized to do so. Indeed, the challenge of both protecting and accessing secrets can be so great that it sometimes will not be worthwhile for a corporation to use cryptography -- the benefits of cryptography for corporate security must be weighed against the costs of monitoring those who control the cryptography. Finally, for the reasons already stated, the need to monitor agents will often result in the use of outside parties to audit cryptographic practices, that is, for the corporation to give up the anonymity of its transactions.

II.C. Market Acceptance as a Limit on the Use of Strong Cryptography.

The discussion to this point has shown reasons why many users will avoid anonymity in financial transactions, even in a world where strong cryptography is permitted. Little or no lending will be made to anonymous borrowers. The risks associated with key management will give individuals and firms strong incentives to entrust their keys to other parties, quite possibly including their bank. A third important

limit on the use of strong cryptography is the somewhat amorphous concept of "market acceptance." The political outcry surrounding key escrow can obscure the many reasons that ordinary users are not likely to seek or use strong cryptography to gain effective anonymity.

Installed base. One reason to expect slow market acceptance is simple inertia. There is a tremendous installed base of non-anonymous financial transactions. Consider the usual sorts of checking accounts, credit cards, and debit cards, all of which contemplate the bank knowing the identity of the customer. A huge range of other, well-established business transactions also depend on the parties knowing each other's identities.

It is no small thing for a new technology to overcome the inertia of established patterns. The history of the ATM -- a new banking technology -- provides an instructive guide to what we might expect for anonymous banking transactions. It has taken a full generation for users to accommodate themselves to ATMs. Even today, use of ATMs varies widely by age. A 1996 survey showed that 36% of those above age 64 have an ATM card, compared with 75% of those aged 18 to 34. [2] I find these usage figures striking because of the seemingly obvious advantages of the ATM, such as the speed of the transaction and the ability to receive cash and do other transactions on weekends and after the historically-short hours at the local branch. Despite these advantages, millions of users have been reluctant to learn how to use ATMs and to trust them to work accurately.

Lack of compelling advantages. The history of the ATM shows that, as a simple business matter, it takes a compelling argument to change ingrained patterns of doing business. It is far from clear that strong cryptography offers any compelling advantages over the traditional business relationships in which the bank knows the customer's identity. The costs of strong cryptography are in some ways more apparent than the benefits. One cost of remaining anonymous is giving up the ability to borrow money. Another set of costs is associated with key management. And current cryptographic products are often clumsy or otherwise unattractive to use, such as when they slow the completion of a task. Widespread consumer use will depend on the development of more user-friendly cryptographic products.

In contrast to these significant costs, the benefits of strong cryptography to provide privacy are not entirely clear. As a preliminary matter, notice that we can expect a great deal of cryptography to protect the security of financial transactions. Media horror stories, for instance, have convinced many consumers that it is risky to send unencrypted credit card numbers over the Internet. Consumers may be willing, therefore, to cooperate with security measures in their financial transactions. Banks also have strong incentives to get their consumers to cooperate in the use of cryptography for security. When system security is breached, the banks who design and operate the system are likely to suffer losses. Banks and consumers thus share an interest in using cryptography to protect against attacks by malicious third parties.

From a business standpoint, banks have much less reason to encourage their customers to use cryptography to protect customer privacy. All other things being equal, banks would prefer to know the identity of their customers, in order to reduce fraud and to sell customers additional products. Widespread customer anonymity, then, will require strong consumer demand. Within the cryptographic community, various reasons have been offered for promoting use of anonymous payment mechanisms.[4] Without taking any position on the desirability of widespread anonymous payments, I limit my comments to a descriptive observation -- it is not yet clear that ordinary consumers have been offered a compelling reason to use anonymous payments. In the absence of such a compelling reason, a vast portion of users will not establish anonymous accounts with their banks, even if strong cryptography is permitted. The availability of mathematical solutions does not mean that users will avail themselves of the solutions.

Who will be trusted? An additional obstacle to the market acceptance of strong cryptography and anonymity concerns the question of whom the customers will trust. For Alice, the system for doing financial transactions involves enormous risk, the risk that accident or maliciousness will drain her accounts and wipe out her net worth. Aware of these customer concerns, banks and payment systems go

to great lengths to reassure their customers. Bank buildings historically are built of sober gray stone designed to convey a sense of permanence and invulnerability. Many institutions place the word "trust" in their name, and the Security First Network Bank bragged in Congressional testimony that "security is not our middle name, it's our first." [9]

How do crypto designers stack up in this trust game? Perhaps not so well. Imagine if Alice, an ordinary consumer, wandered into a cypherpunks convention. Does she want to trust her life savings to a product designed by these people? Aren't some of them "hackers"? And didn't she read in the papers that hackers do all sorts of criminal and anti-social things?

To give Alice a bit more credit, she might have other reasons to worry about whether the cryptographic product will operate as promised. She may have heard about how programmers can create back doors and Trojan horses. She might suspect her private financial information is being shipped directly to the programmer. She might never have heard of the individuals or companies who designed the anonymous payments system. Not being an omnicompetent cryptographer, she almost certainly lacks any personal ability to evaluate how well the product actually protects her anonymity. She might be reluctant to take a great deal of time to learn how to use the cryptographic product, and skeptical about her own ability to run it well enough that she will actually gain any privacy protection.

In short, customers have a series of potentially valid concerns about adopting cryptography, due to the risk that the product was badly or maliciously designed, or the reality that most customers will not be cryptographic adepts. If customers decide to use it at all, they will prefer to use cryptography that has been cleared by institutions they trust. Some people (likely not in the cryptographic community) will trust the government, and will feel comforted by some sort of government seal of approval. Other people will look to institutions that specialize in having customers trust them with money. These institutions, roughly speaking, are banks and other established financial institutions.

Consider what follows from the fact that consumers trust banks more than other purveyors of cryptographic systems. Much of the debate within the cryptographic community has assumed that government support of key escrow is the crucial barrier to widespread use of strong cryptography. The analysis here proceeds on the assumption that governments will allow strong cryptography. If so, then we reach the (perhaps unsurprising) conclusion that banks would stand as gatekeepers to widespread use of financial cryptography. As already discussed, however, banks generally prefer knowing their customers to having their customers remain anonymous. Banks will make it less expensive and easier to have accounts where they know the user's identity. Once again, the proponents of widespread anonymous transactions face the burden of explaining how and why consumers will take on the costs and risks of operating through anonymous accounts.

III. Can Technology, Rather than Law, Protect Privacy?

Let's recap what we have learned so far. We began with the Cypherpunk Credo, seeking "privacy through technology, not legislation." We assumed, for discussion purposes, that strong cryptography would be legal and readily available. In Part I, we saw that the use of cryptography is likely to spread rapidly in coming years, especially to protect the security of transactions. In Part II, we saw a number of important reasons for believing that cryptography would not become nearly so widespread to protect users' anonymity. Note that the difficulties of anonymous lending, key management, and market acceptance are not primarily technical; instead, the analysis here has emphasized economic, psychological, and other sorts of reasons that individuals and corporations would choose not to use anonymous transactions.

This Part of the paper will provide a more complete assessment of the likely use of anonymous transactions. We will relax the assumption that strong cryptography will be legal and readily available. We will also relax the implicit assumption that digital signatures are widely available and impose no penalty on anonymous users. Once the full nature of these constraints is appreciated, the discussion will

turn to what I call "overcoming legalphobia." My point here is that the cryptography community, traumatized by the Clipper Chip controversy, has become excessively allergic to the possible advantages of legal regulation for financial privacy.

III.A. A More Complete Assessment of the Constraints on Anonymous Transactions.

To this point, we have assumed that strong cryptography is legal and readily available. Nonetheless we have seen that *many* transactions will not have strong privacy protection -- another party will often have a database linking the user and the transaction, for lending, key management, and market acceptance reasons. The next step is to relax two assumptions: that strong cryptography is legal and readily available; and that anonymous digital signatures are widely accepted. The result of relaxing these assumptions would be that even fewer transactions would have mathematically-based privacy protection.

One possibility is that strong cryptography will become illegal for some or all financial transactions. The expected result of such a law would be to reduce the use of strong cryptography. Even if some cypherpunks manage to persist in their use of illegal strong cryptography, many individuals and established corporations would probably not risk criminal sanctions.

A related possibility is that strong cryptography would remain legal, but there would be government or other pressure not to use it. For instance, companies that contract with the government might be required to use key escrow systems for purposes of their government contract work. If so, it may be more convenient or politic for those companies to use key escrow for all of their cryptography needs. A related possibility is that leading software firms, hardware firms, and banks would decide not to support strong cryptography, perhaps under government pressure. More broadly, there could be a contest for the "hearts and minds" of the general public. Law enforcement and national security agencies would predictably encourage law-abiding citizens to use key escrow. Other parts of society, including cypherpunks and privacy advocates, would predictably encourage as many users as possible to take advantage of strong cryptography. Under any of these scenarios, the net effect would be less widespread use of anonymous financial transactions.

Turning to anonymous digital signatures, an implicit assumption of the discussion to this point is that they will become widely available at a reasonable cost. The chief purpose of digital signatures, as discussed extensively elsewhere, is to provide a way for the user to provide authentication over an electronic network. [6] For instance, a digital signature might allow an anonymous Alice to prove that she made a particular payment at a particular date and time. Such signatures are obviously important to many sorts of financial transactions, such as when Alice wishes to return a purchase as defective and wants her money back.

For our discussion, the relevant question is the extent to which a user can accomplish everything through an anonymous digital signature that he or she can accomplish by other means of authentication. As a society, we are in the early stages of learning how to establish digital signatures. We are debating how to structure the certification authorities that will help manage them. Depending on how the laws and practices develop in the area, anonymity may prove to have additional costs for consumers. If anonymous purchasers lose important consumer rights, then we would expect even fewer anonymous purchases.

B. Overcoming Legalphobia: A Positive Role for Legal Regulation of Privacy.

The discussion to this point has emphasized the reasons that many users will not use strong cryptography to conduct anonymous transactions. Faced with the many constraints on the power of mathematics to protect users' privacy, the next question is whether legal rules for protecting privacy would be desirable. In my discussions with cryptographers, there has often been a visceral reaction against the use of legal rules to protect privacy. After all, the Cypherpunk Credo says: "Privacy through technology, not legislation."

I would like to suggest that this "legalphobia" is both understandable and misguided. It is understandable because of the salience of the key escrow controversy within the cryptography community. The opposition to legal rules can be understood as based on this flawed syllogism: (1) Clipper was an example of government regulation of cryptography; (2) Clipper would result in reducing people's privacy; therefore (3) government regulation reduces people's privacy, which is a bad thing.

Table 1 is designed to show the flaw in this logic. The basic point is that legal rules act differently in cryptographic and non-cryptographic settings. The cryptographers' syllogism is based on the assumption that users will employ strong cryptography. If users indeed employ strong cryptography, then by definition privacy is at the maximum. That is Box I in Table 1. Legal rules in such circumstances can only have the effect of reducing privacy. Clipper and other mandatory key escrow schemes belong in Box II, and result in the wider availability of previously-private information.

	Crypto Transactions	"Ordinary" (Non-Crypto) Transactions
No Regulation	**I: SECRECY** Strong cryptography allows anonymity.	**III: DISCLOSURE** Market result is wide availability of information.
Regulation	**II: DISCLOSURE** Regulation leads to wider availability of information, as in Clipper.	**IV: "DISTRIBUTED PRIVACY"** Legal rules protect information -- information available only to selected 3rd parties.

Table 1

My argument in this paper, however, is that users very often will *not* employ strong cryptography. In such circumstances, the role of legal rules shifts 180 degrees. In an *unregulated* setting we will expect widespread disclosure of private information, as shown in Box III. In today's marketplace, the operators of databases often seek to profit from their information about customer transactions. They "data mine," in current parlance.

Against this backdrop of widespread disclosure, legal rules can actually promote privacy. It is of course true that legal rules cannot offer the same guarantee of anonymity as does strong cryptography. Legal rules, at best, can promote what might be called "distributed privacy." In a world where a great many transactions will not be anonymous, we face a choice between Box III -- disclosure through the marketplace -- or Box IV -- reduced disclosure on the basis of legal rules. In Box IV, transactional information is distributed from the consumer to other parties such as the bank. The bank, however, is under legal obligations not to reveal that information except where legally permitted.

The image here is one of contract, of delegation from the principal (the consumer) to the agent (the bank). For all of the reasons already discussed, it often makes sense for the consumer to wish to delegate key management and other tasks to an expert and trusted party, such as the bank. As is true more generally of principal/agent relationships, the principal Alice has broad powers to stipulate how the agent should act on her behalf. If the agent violates that agreement, such as by disclosing information in

ways that Alice would not desire, then the agent should pay contract damages. The next Part of this paper sketches how this contractual approach to privacy might operate.

I submit that having the distributed privacy in Box IV is imperfect, but is far more desirable than the widespread disclosure of Box III. I will not attempt here to explain all the reasons why people might care about their privacy, even if one institution such as a bank has a database listing their transactions. I will, however, cite one recent anecdote that I found telling. In the winter of 1997, the Cypherpunks listserv went through a lengthy process of shutting down. A proposal was then made to shift the discussion to a Usenet group. At that point, numerous complaints were made that participants in a Usenet group would lose their privacy, such as by being subjected to lots of spam. Lengthy discussions were held about how one might post to a Usenet group without losing one's privacy.

I suggest that this episode supports the importance of having distributed privacy. After all, the listserv participants had already revealed a good deal about themselves when they posted to the hundreds (perhaps thousands) of people who subscribed to the listserv. Much private information, such as e-mail headers, was already widely distributed. Yet these same people objected strongly to becoming a bit more public on the Usenet.

Turning to financial transactions, Alice does lose anonymity and privacy if information on her transactions is distributed to the bank. The problem is worse to the extent that the government can subpoena information. But Alice nonetheless retains an important measure of privacy compared, say, with the bank posting all of her transactions to a World Wide Web site. Where mathematics does not protect an individual's privacy, for all the reasons already discussed, then legal rules can play a crucial, supplementary role.

IV. Outline of a Contracts Approach to Privacy.

The paper to this point has discussed a series of ways in which privacy, to the extent it is protected, will not be protected solely by mathematics. Other institutional and legal developments will be needed if transactional information is not to be widely disseminated. The discussion thus serves as the introduction to a broader project, on "the role of law in protecting financial privacy."

Within the broader project, this paper explores the uses and limits of cryptography in protecting personal financial information. Where technology does not protect privacy, then law may provide supplemental protections. The second part of my project is tentatively entitled "Cyberbanking and Privacy: The Contracts Model." [12] That paper first claims that the use of customer information by banks can and should be conceptualized under contract principles. A second claim is that the law-and-economics approach to contracts -- an approach often thought to be hostile to protection of privacy -- supports having important rights in the information held by the customer rather than by the bank.

The third part of the larger project focuses on the question of which institutions should implement the rules for protecting personal information. This part was written for a report by the U.S. National Telecommunications and Information Administration. My paper is entitled "Markets, Self-Regulation, and Government Enforcement in the Protection of Personal Information." [13] As the title suggests, the analysis tries to specify the conditions under which we should expect markets, self-regulatory organizations, or government enforcement to operate most effectively. The emphasis is on understanding when self-regulation may be preferred to either the market (no regulation) or to government regulation. In general, I conclude that self-regulation will be more similar to the market than it will be to mandatory legal enforcement. The key choice will typically be between self-regulation and a more intrusive government role. Often the best government role may be to provide courts to enforce contracts between banks and customers. Governmental agency enforcement may also sometimes be appropriate.

In the remainder of this article, I will sketch the main points of the contracts approach to privacy. The goal of this brief description is to suggest the principle ways that a legal and contractual approach might provide significant protection of personal financial information. The discussion here will be necessarily incomplete -- it will not address all the possible objections or define all the points in economically-rigorous ways. The intention, however, is to make the main points of the contracts approach available to a non-legal audience.

IV.A. The Contracts Approach and Specifying Default Rules.

In understanding the contracts approach to privacy, the first job is to specify the "default rules" for the contract -- the rules that apply when the contract does not explicitly address a subject. Once the default rules are defined, the next task is to know what the parties must do in order to be governed by a different rule. Then, we must know the consequences of violating the contract.

Majoritarian default rules. A first approximation of the correct default rule is often to estimate what most people would have wanted if they had bargained on the subject. *If* a customer were well-informed about privacy issues, and *if* the customer were able to bargain with the bank about the uses of customer information, what uses would be approved under the contract? This majoritarian, "would have wanted" approach can further important goals such as autonomy and efficiency. Autonomy is furthered because of the focus on the wishes of the parties -- on what the parties would have agreed to had the bargaining taken place. Efficiency is furthered because we are looking to the wishes of the majority -- in the absence of an agreement between these two parties, we will provide them with the contract that most parties would prefer.

At the present time, information technology is leading to the creation of many new databases, linked in many new ways. Under such changing circumstances, it is hard to make conclusive estimates of what most parties would agree to if they bargained on privacy issues. Indeed, we are in a society-wide process of learning about the uses and abuses of private information. Despite these uncertainties, a few observations are possible about what most consumers would agree to if they bargained with their bank. We have the results of major polls, which consistently show a large and growing proportion of people reporting that they are concerned about threats to their personal privacy.[5] [4] We also have our own intuitions, which let us know that most people would not knowingly agree to give their bank unlimited power to publicize their transaction records. Any default rule on financial privacy should take account of this consumer objection to widespread dissemination of their transactional information.

Default rules with asymmetric information. A second important approach to default rules emphasizes information asymmetries, which are situations in which one party has significant information that the other party lacks. [1] Where such asymmetries exist, there is often a strong argument that, in order to get an efficient contract and an accurate meeting of the minds, the default rule should favor the party that lacks the information. Not to reveal the information can constitute a form of fraud -- the party lacking the information (often, the consumer) has never really agreed to the contract the way the well-informed party plans to use it. In order to force the formation of a contract to which both parties agree, the party having the information should not be permitted to profit from it, and the default rule should favor the party lacking the information.

There are two important information asymmetries that affect banks' use of their customers' personal information. The first asymmetry arises from the fact that customers in fact do not realize the extent to which banks and other organizations use personal information. My claim is that banks and other organizations accumulate and use databases of transactional information in more ways than customers suspect. In casual empirical support for this claim, I refer to the Cyberia listserv to which I subscribe. The listserv consists of persons interested in law and the Internet, and generally produces highly informed comments about a wide range of Internet issues. My experience in the past two years is that Cyberia subscribers are repeatedly surprised by particular ways that organizations gather and use data about

customers. By extension, if Cyberia subscribers are surprised, I suggest that most ordinary consumers would be far more surprised by the way their transactional data is used.

Under a contracts approach, the solution to this information asymmetry is straightforward. If the asymmetry is important enough, and I would argue that it is, then the bank would not have the contractual right to use the customer's information in ways that the customer does not reasonably contemplate. As an easy example, the customer would have contract rights against the bank if the bank suddenly posted the customer's transaction history on a Web site.

This asymmetry emphasizes the current lack of consumer knowledge about the ways that their information is used. Over time, one could imagine this asymmetry shrinking or disappearing. As we go deeper into the digital database era, consumers expectations may change. They may come to expect, for better or worse, that banks and other organizations will pervasively use their information.

Even in that instance, a second important asymmetry would exist. Even if customers accurately know how consumer information is generally used, it is costly or impossible for customers to tell how a particular company is using that information. Suppose that Alice strongly wishes to limit use of her personal information. Alice is disadvantaged because the cost and ineffectiveness of ymonitoring logically leads to over-disclosure of private information. Consider the incentives facing a company that acquires private information. That company gains the full benefit of using the information, notably in its own marketing efforts or in the fee it receives when it sells the information to third parties. The company, however, does not suffer the full losses from disclosure of private information. Because of imperfect monitoring, customers often will not learn of that use. They will not be able to discipline the company efficiently in the marketplace for its less-than-optimal privacy practices. Because the company internalizes the gains from using the information, but can externalize a significant share of the losses, it will have a systematic incentive to over-use private information. In terms of the contract approach, companies will have an incentive to use private information even where Alice would not have freely bargained for such use. In such a situation, the sensible legal approach is to deter the company from such uses of information. The default rule should once again be that companies cannot use information in ways that customers do not reasonably contemplate.

IV.B. Beyond Default Rules.

A full contractual regime provides more than the default rule. A default rule simply states what the rule shall be when the contract does not specifically address a topic. A contractual regime will also explain what the parties must do to contract around the default rule. The regime will provide for damages or other remedies in the event of contract breach. Finally, the regime will state which institutions will enforce and adjudicate the contract.

Contracting around the default rule. The above discussion suggested that the default rule should be that banks cannot use information in ways that customers do not reasonably contemplate. If banks wish to use customer information more broadly, one can then expect banks to insert favorable language in their deposit and other contracts with customers. For instance, the fine print of the contract might state: "The bank retains the right to use customer transactional information as it sees fit." The question for the legal regime will be whether to treat this fine print as an agreement binding on the customer. Resolving this question would take us far into the realm of consumer law. For present purposes it is enough to be aware that there will be tricky legal issues about whether customers have consented to uses of their personal information.

Contract damages and choice of institutions. Under current law, it is often difficult or impossible for a customer to prove damages if a bank or other company misuses personal information. If companies are to be deterred from misusing information, there will likely have to be new statutory or other legal remedies. The point of these remedies will be to shift the company's incentives, in order to make it unprofitable for the company to over-use the information and break the contract. As in the case of

other consumer statutes, it may be desirable to provide for class actions and attorney's fees, so that the company will not be able to profit from repeated, small violations of consumer contracts.

This discussion of damages assumes that individual consumers will go to court in order to remedy any breach of contract. Other institutional arrangements, however, are certainly possible. In my paper for the National Telecommunications and Information Administration, described above, I examine the conditions under which markets, self-regulatory approaches, or mandatory government rules would seem most appropriate.

V. Conclusion.

This paper addresses the uses and limits of financial cryptography. The analysis here suggests that cryptography should be deployed widely to assure security in the transmission of data, e.g., to prevent a malicious party from sniffing out credit card numbers on the Internet. Cryptography should also be deployed widely to assure the security of information in databases, e.g., to prevent employees from unauthorized snooping in the files.

The situation is more complicated when it comes to the use of cryptography to protect anonymity. If strong cryptography is legal, it will undoubtedly be used by highly motivated users to assure their anonymity. For less motivated users, strong cryptography may work, but it simply will not solve many of the problems that people have when they manage their finances. Nor will the situation necessarily change very much over time. Although new generations of users may be more comfortable with cryptography than current users, there are important economic, psychological, and other reasons why consumers may choose not to act anonymously. Consumers will need to reveal their names and credit ratings to borrow money, key management problems will persist, and banks may retain strong incentives to offer better terms to customers who agree to let the banks know their identity.

These limitations on the use of financial cryptography do not necessarily mean that anonymous electronic payments will be unavailable to ordinary users. At the Financial Cryptography '97 conference, noted cryptographer Ron Rivest speculated on the future of anonymous e-cash, and suggested that the most likely future widespread use would be in low-denomination stored-value cards. Stored-value cards are becoming more familiar today, such as for telephone calls or in the metro system in Washington, D.C. Purchases with stored-value cards can remain anonymous so long as no data link exists between the purchase of the card and the individual transactions that employ the card.

Stored-value cards avoid many of the problems discussed in this paper. They are debit cards, and so do not involve lending. Their small denominations mean that there is no need for fancy key management. And, as the growing popularity of telephone cards suggests, their ease-of-use leads to market acceptance. The widespread use of stored-value cards, including for purchases over electronic networks, could thus enable individuals to have anonymity for small purchases in daily life. Stored-value cards in low denominations might also be politically acceptable, because they might be more difficult to use for money laundering and other illegal purposes than would large anonymous bank accounts.

Whether or not stored-value cards proliferate in this way, they will not prove to be effective ways for individuals or firms to manage their ongoing financial affairs, especially for moderate- or large-sized transactions. For these transactions, the likely outcome will be non-anonymous accounts with banks or institutions that act like banks. Once banks can link transactions to the individual customer, technology alone cannot supply privacy. Then, as my ongoing project hopes to show, we will need good legal rules and other institutional ways to maintain privacy in a database world.

Acknowledgements.

Thanks to my colleagues, Ted Janger and Doug Whaley, for their help in developing this paper. I benefitted from comments from participants at the Financial Cryptography '97 conference in Anguilla, and the Computers, Freedom & Privacy '97 conference in San Francisco. Thanks also to my research assistants, Thomas Sinclair and Robert Wells.

Notes.

1. The discussion here implicitly assumes that Alice lives in the house that she is buying, so that simple surveillance will reveal that she is the resident and eliminate her anonymity. One can imagine instead that Alice owns the house but never visits it. In that event, it is indeed possible for Alice to maintain her anonymity, although she won't get any personal use out of the house.

It is even possible, in a limited sense, that the bank will make a loan without knowing the identity of the borrower. The bank can find it profitable to make such a loan without knowing anything about Alice's credit record, if two conditions are met: (1) the value of the house exceeds the loan amount; and (2) the expected revenue on the loan (up-front fees and interest rate) is high enough to pay for the cost of repossession. Notice, however, that these conditions transform the transaction once again into a debit transaction. The house becomes the equivalent of a deposit with the bank. As with the secured credit card, the bank only makes funds available to the extent that it keeps sufficient funds "on deposit" with the bank.

2. One caveat here. Because I am not myself a cryptographer, I cannot assess the extent to which a technical fix would be available. Perhaps there is some way to design the system so that the borrower loses anonymity upon nonpayment of a loan. To be effective, the penalty would need to be worse than mere imposition of a bad reputation on the pseudonym that is linked to the account. The penalty would need to include a way actually to identify and take action against the *person*, and not just the pseudonym. Otherwise, the deadbeat borrower simply establishes a new pseudonym, and begins borrowing again.

Put another way, I am challenging the cryptographic community to create a mechanism for revealing the identity of the borrower in the event of nonpayment, thus allowing enforcement against the deadbeat borrower.

3. An alternative exists for users of cryptography who do not wish to deliver their keys to some other party. Under a "secret-sharing protocol" [10, p. 71 & 182], Alice can divide up her key into a number of pieces. She then sends a different piece to each of several people. None of the pieces alone is the key, but the pieces all brought together can reconstruct the key. An obvious advantage of this secret-sharing protocol is that no one party acting alone can read Alice's encrypted messages.

Despite this advantage, I question whether the ordinary consumer would bother to create this sort of secret-sharing protocol. The protocol might be quite complicated to create and sustain. For instance, each of the other parties would need to be reliably available when their piece of the key were needed. For the typical consumer, it might be far easier to pick one trusted institution, such as the bank, and have that expert institution manage the keys.

4. Anonymous payments, for instance, might weaken the power of government surveillance and further the libertarian goals of the cypherpunks movement.

5. In 1995, 81% of Americans surveyed reported that they were "very concerned" or "somewhat concerned" about threats to their personal privacy. [4] As this volume was going to press, I heard a report that that figure had risen to a record 88% of Americans surveyed.

References.

1. Ian Ayres & Robert Gertner, "Filling Gaps in Incomplete Contracts: An Economic Theory of Default Rules," 99 Yale Law Journal 87 (1989).

2. Valerie Block, "ATM Cards Hit a Wall: The Next Breakthrough is Years Away, Bankers Say," American Banker, Jan. 2, 1997.

3. David Chaum, "Achieving Electronic Privacy," Scientific American, Aug. 1992, p. 96.

4. Equifax-Harris Mid-Decade Consumer Privacy Survey (1995).

5. A. Michael Froomkin, "The Metaphor Is the Key: Cryptography, the Clipper Chip, and the Constitution," 143 University of Pennsylvania Law Review 709 (1995).

6. A. Michael Froomkin, "The Essential Role of Trusted Third Parties in Electronic Commerce," 75 Oregon Law Review 49 (1996).

7. Furash & Co., Banking's Role in Tomorrow's Payments System: Ensuring a Role for Banks, A Study Prepared for the Banker's Roundtable (1994).

8. Michael C. Jensen & William H. Meckling, "Theory of the Firm, Managerial Behavior, Agency Costs and Ownership Structure," 3 Journal of Financial Economics 305 (1976).

9. Testimony of Michael S. Karlin, U.S. House of Representatives, Committee on Banking, hearing on "The Future of Money," Mar. 7, 1996.

10. Bruce Schneier, Applied Cryptography (John Wiley & Sons, 1996) (2d ed.).

11. Dashka Slater, "Secret Agents," Express: The East Bay's Free Weekly, Mar. 14, 1997, p. 1.

12. Peter P. Swire, "Cyberbanking and Privacy: The Contracts Model," www.osu.edu/units/law/swire.htm (preliminary draft available, April 1997).

13. Peter P. Swire, "Markets, Self-Regulation, and Government Enforcement in the Protection of Personal Information," www.osu.edu/units/law/swire.htm, scheduled to be released at www.ntia.gov (1997).