Chapter 1

WHAT PRICE PRIVACY?

(and why identity theft is about neither identity nor theft)

Adam Shostack
adam@homeport.org

Paul Syverson
Naval Research Laboratory
syverson@itd.nrl.navy.mil

It is commonplace to note that in surveys people claim to place a high value on privacy while they paradoxically throw away their privacy in exchange for a free hamburger or a two dollar discount on groceries. The usual conclusion is that people do not really value their privacy as they claim to or that they are irrational about the risks they are taking. Similarly it is generally claimed that people will not pay for privacy; the failure of various ventures focused on selling privacy is offered as evidence of this. In this chapter we will debunk these myths. Another myth we will debunk is that identity theft is a privacy problem. In fact it is an authentication problem and a problem of misplaced liability and cost. When these are allocated to those who create them, the problem does not exist. Finally we consider the oft asked question of how much privacy should be given up for security. We find this to be the wrong question. Security of institutions may decrease and infrastructure costs may be increased by a reduction in privacy.

† 'This chapter is an expanded and revised version of both "The Paradoxical Value of Privacy" by Paul Syverson and " 'People Won't Pay For Privacy,' Reconsidered" by Adam Shostack. Both of these were presented at the 2nd Annual Workshop on Economics and Information Security, College Park MD, USA, May 2003.

1. The Meanings of Privacy

The word 'privacy' is heavily loaded with hard-to-disentangle meanings. It can mean anything from email confidentiality (PGP), to controlling who emails you (SPAM), to who sees your credit report (identity theft) to the ability of a woman to have an abortion (Roe v. Wade). The many meanings of 'privacy' contribute to the confusion which surrounds it, and some of the apparent contradictions may be resolved simply by paying close attention to them. Other work that has examined privacy and economics has chosen to focus on a single definition (Varian, 1997). By pointing out a rational way to behave in the context of a single definition of privacy, these analyses may actually contribute to the idea that people are acting irrationally.

Therefore, for clarity, we will try to use following words in place of 'privacy':

Unobservability is when you can not be observed. For example, shutting the door to the bathroom offers unobservability.

To Be Left Alone is a classic definition from Justice Brandeis. There is some subtlety in his writing, which we ignore, because the phrase is so powerful.

Untraceability is when you can not be traced from one identity to another. For example, "John, who we play softball with, but don't know his last name" is untraceable; you can't track down a phone number for him.

Informational self-determination is when you are confident that information you provide will be used only in ways you understand and approve. Giving your mother your new phone number probably qualifies.

Anonymity is when you are without any identifiers.

Many of these terms are based on other uses within the technical and legal privacy literature, and we believe that their uses here are very close to their understood meanings.

Each of these terms captures a meaningful aspect of privacy, and each of them is a goal which people pursue. There is also a measure of how important privacy is to people, which Westin breaks down into the "fundamentalists," "pragmatists" and "Don't cares" (Westin, 2001). It is the last group often cited as willing to trade away their privacy for a free hamburger.

Given these meanings, we will examine how people pay for them. From there, we will examine a number of areas where people don't pay for privacy. We will then explore what lessons can be learned from this.

2. Privacy People Pay For

The most obvious way people pay for privacy is in banking services, paying for informational self-determination, in the form of a guarantee that information about them won't be provided to some set of tax authorities, family members, or others. This is a business estimated at many billions of dollars per year.

In the realm of unobservability, privacy is one component of what drives purchases on curtains and drapes, as well as large shrubbery and fences. This statement is based on the easily observed fact that privacy is listed in most advertising for "window treatments" in home decoration magazines. We use advertising as a proxy for what people value because advertisers won't include things which they don't believe will sell their product, and they won't put in things which they expect will cause their audience to shake their heads. Drapery and curtains, whose sales are motivated not only by unobservability, but also by aesthetics and economics of insulation, were approximately 1.8 billion dollars in 1997 (US Census Bureau, 1999). We do not attempt to break down these numbers as to which motivator leads. We do note that see-through or lace curtains seem relatively rare. (Speaking of homes, privacy or distance from neighbors is often a reason to move to the suburbs or country.) On January 27, 2003, the New York Times published a story on college dormitories and private rooms. The story teaser on the web site was "With more students demanding – and paying for – privacy, the roommate is no longer the staple of college life it once was." Students at Boston University are paying an extra $1,400 per year, or about 4% extra for a private room.

Unobservability also drives mailboxes, private mail boxes, and mail receiving services in two ways. Some of this is unobservability with respect to the sender: one's real physical location is not revealed. Some of it is unobservability with respect to one's house-mates or family, who don't know what mail a person is getting. The post office rents more than 18 million post office boxes, for nearly 500 million dollars per year. It is unclear how many of these are personal or small business/sole proprietor sorts of rentals (USPS, 2001). Privacy is explicitly listed by both the US Postal Service and Mailboxes, Etc. as a motivator for renting of post-office boxes (USPS, 1998, and Mailboxes, Etc. web page).

Another area where the right to be left alone matters to people is their telephones. Some people find unwanted calls to be enormously annoying and intrusive. To address this concern, there are caller-ID, caller-ID blocking, voice mail services, and unlisted numbers. We consider both caller ID and the blocking service to be privacy driven. Caller ID is a desire to be left alone by unknown callers, a function of which is also served by answering machines with a call-screening function. Caller ID blocking is an untraceability feature, where

the caller desires privacy. Voice mail services regularly advertise themselves as a unobservability services, where roommates and others don't know what calls you're receiving. Unlisted numbers reflect a desire to be left alone; in California over half of all home phone lines are unlisted. (A perhaps interesting aside is that one of the authors no longer calls directory assistance to try to find people, only businesses, because he assumes that all of his friends have unlisted numbers.)

It might be interesting to add up the numbers above, but that presents several substantial difficulties. First, and most easily solved, the numbers are not for the same years. Secondly, many of the products are "tied", where privacy is bundled into a complex product, rather than a feature for which one can choose to pay. Some of this might be separable; for example, in the curtain example, We could pursue average sizing of curtains per dwelling, find the lowest cost option to block the view, and assign that as the privacy component. However, this strikes us as potentially misleading: Are all curtains purchased for privacy? If privacy were the only concern, would people re-use more older curtains? Similarly, with a post-office box, some portion of the rental may be to obtain a "professional appearance" or to avoid mail-theft issues. How to separate that out is not clear. Thirdly, we have not attempted to assemble a comprehensive list of markets where privacy is a factor. Lastly, and most importantly, its unclear what such numbers would mean, and thus they could not be used correctly. Therefore, we make no effort to add up these numbers. We simply point out that privacy is an important component of what people are paying for, refuting the claim that "people won't pay for privacy."

3. The Irrational Privacy Consumer: Selling your virtual self for a hamburger

Austin Hill has observed that people will tell you that privacy is very important to them, but then give you a DNA sample in exchange for a Big Mac. While there is clearly a bit of bemused (or frustrated) hyperbole in this statement, the thrust appears correct. But is it really so irrational to exchange private information for something of relatively little economic value?

We claim that there need be no inconsistency inherent in such behavior. Suppose a hamburger is worth two dollars, a full blown identity theft costs an average of 100K dollars, and the probability of such identity theft from giving name, address, and phone number to the hamburger vendor is 10^{-10}. In this case, the rational action is to trade the information for the hamburger. Expected value of such a transaction is still effectively two dollars.

But even assuming these numbers are reasonable, this example reflects a short-sighted consumer. Suppose the incremental probability given a previous history of such transactions is on average slightly higher, say 10^{-9}. A thousand

such transactions reduce the long term average expected value to a dollar. Thus even in the relatively long run, the consumer made no mistakes.

This is a very simplistic example. It overlooks the cost of discomfort the individual feels from her information being held by the vendor, the inconvenience from receiving resulting unwanted junk mail or the positive value if the consumer actually desires, e.g., the resulting coupons she receives, etc.

The cost of the discomfort felt at the collection of information is especially difficult to quantify. But, it may be reasonable to completely remove it from any analysis. For it is the expectation of how that information will be used that is significant. If such data were collected such that the individual felt genuinely sure that it would simply be filed away and never accessed, never correlated with any other actions of hers, never used in any way, it is unclear that she would care. Of course there is always some expectation that if an effort is made to collect the data, then someone intends to use it in some way. In any case, even adding such costs as the increase in junkmail, the expectation of unpleasant inferences about her by marketers, financial institutions, etc. it is at best unclear that the expected cost exceeds the value of the hamburger.

This is not to say that people are more *or less* rational with respect to privacy than any other aspect of their lives. They still understate or ignore risks in the temptation of immediate gratification, and there have even been some economic models of this in the privacy context (Acquisti, 2004). While quite insightful, such analysis can at best be hypothetical at this point, as we shall see presently. However, the point of the example is that while the consumer may violate some ideal rationality of an economic model, it is indeed hyperbole to claim obvious and extreme irrationality in such actions.

So, what is going on? Are privacy advocates just fanatics, themselves irrational about such things? Some have concluded as much with less justification. But there are other aspects to this issue.

First, the above numbers, however plausible, are made up. A shift of a few orders of magnitude could change things drastically. Second, real numbers are very difficult to come by and virtually impossible to justify. It might be possible to collect data on occurrence of identity theft correlated with consumer behavior so that probabilities of at least such clear privacy problems could be assigned to some actions. However, this is at best unclear and has not been done yet. And even this would ignore the other types of privacy cost, a few of which we have mentioned. Also, limiting ourselves to identity theft for the moment, any data collected would be of limited predictive value. According to the US FTC, the rate of identity theft is doubling every year. Obviously if true, that cannot continue for long. The situation is just too dynamic right now for there to be any empirically accurate analysis of current trends. Plus, the market typically needs to learn from experience, so consumer behavior is likely to lag behind any current reality. So one answer is that the expected cost

of privacy compromise, both large and small, is increasing. Privacy advocates (along with economic privacy modelers) are just ahead of their time.

Third, the example we have been considering is one involving the assessment of low probability but high value events. This is difficult enough for those who have good numbers and good understanding. Individuals may be somewhat polar in response to these circumstances. Horror stories of lost livelihood are met with sympathy but an expectation that it won't happen to me. And historically that has been statistically accurate. But, there may come a tipping threshold that will make this a major issue not just in polls but in individual behavior and in individual demands of government and business. Alternatively, the right sort of individual soundbite may resonate through society. A recent story in MSNBC recounts the plight of Malcolm Byrd who besides economic suffering, job loss, etc. has been arrested many times and spent time in jail more than once as the result of an identity theft (Sullivan, 2003). These stories may also desensitize people or leave them feeling helpless, since they have no meaningful way to respond. We will return to the advice people are currently given below. For now we note with trepidation that, while identity theft in general continues to rise exponentially, so-called *criminal* identity theft (as in this story) has increased as a percentage of the total, from 1.7% in 2001 to 2.1% in 2003 (FTC, 2004). On the other hand, the Anonymizer (the self-proclaimed "Kleenex" brand name in Internet privacy) claims a 500% increase in subscriber base from 2002 to 2003. Perhaps a tipping point is being approached.

Another indicator of privacy attitudes frequently cited is that people don't click through to read privacy policies. This is often cited in support of the assertion that people don't actually care about their privacy. We believe that it is more accurate to state that privacy policies rarely reveal anything in comprehensible language, and even more rarely give meaningful choices. Additionally, companies rarely distinguish themselves in their actual privacy commitments, so it is hard to choose a company for its privacy policies. (Initiatives such as the World Wide Web Consortium's P3P may help to change that, but it is still early to see.) Finally, most companies reserve the right to change their privacy policies at any time, and many exercise that right, meaning that even if a consumer chooses a company for its current privacy policies, he is unlikely to feel that he will have informational self-determination, or control over how information about him is used. As such, consumer decisions to not waste time with them reflect more on their utility than on consumer's privacy desires. Consumers failure to read, understand, and respond to bank privacy notices required under recent US laws may be understood the same way. However, in the case of those laws, the presence of the weasel word "affiliate" make it hard to determine if one would actually be left alone if one did bother to fill out the card.

Again, what appears to be insensitivity to privacy is actually a rational decision about the effect of investing time and energy in understanding a policy, and the expected value of that investment.

4. Analysis

Privacy is often a component of some other sale—home decoration or convenience. This makes it hard to place solid numbers on "The privacy market," although those would be quite interesting.

Consumers seem to spend money when there is a comprehensible threat, with an understandable solution, for example, with curtains. The concern of people looking in through windows is easily understood, and the solution is easily comprehended. In newer, or less transparent situations, understanding may be harder to come by. An example would be http cookies. It is not trivial to understand what an http cookie is, as this requires some understanding of the idea of a protocol, a server, and statefulness. Understanding the interaction of cookies with traceability and linkability is even more complicated, as it requires understanding of web page construction, cookie regeneration, and non-cookie tracking mechanisms. So, understanding the technical nature of the threat has a high threshold. From there, understanding the impact of the threat is complicated. What does it matter if all my browsing can be linked together to my real identity? What impressions or notes may be made when one goes to a pharmaceutical (or illegal) drug site, a gay rights site, or the web site of an accused terrorist organization? In contrast, understanding that anyone driving by can see in your windows if you don't have curtains is trivial. Protecting against threats too difficult for the average current consumer to grasp is a hard sell. A potential example is iPrivacy, a company that began five years ago offering comprehensive protection of consumer name, credit card information, and even address for physical delivery of goods. But it has never taken off. It does not help that vendors may try to convince consumers that it is in their best interest to provide personal information, whether or not this is true.

Businesses spend time and energy to present their activity in the best possible light, sometimes to the point of misdirection. For example, warranty cards which state they must be filled out completely to "ensure the best possible service" also ask for demographic information. Understanding what will be done with the information may take more effort than the result is worth.

Even if one does take the time to learn about and understand how different organizations will handle one's personal information, there may be little difference between them. In the financial services world the difference in actual policy may be very slim. In addition, information important for understanding what privacy an offer really entails may be lacking. Alternately, a choice may

appear to be a marketing ploy, not one based on real distinctions. As such, there may not be a real choice that can be made on privacy.

A recurring feature of the privacy world is that new issues are raised. New ways of invading privacy are suggested, people are outraged, studies are written, and the new technology succeeds or fails without apparent correlation to privacy issues. This is a phenomenon worth exploring. Those new technologies which succeed do so in one of two ways: First, they succeed in the marketplace. The benefits that they offer are so substantial that people are willing to give up their privacy for the benefits gained. It is worth asking in this instance, is this an informed choice? Will they regret it later? However, it is a choice which is sometimes freely made; for example, the capability to track cell phones deters very few people from carrying them. Concerns are raised more regularly about cancer risks. The second way new technologies succeed is that they are mandated. For example, cell phones will soon come with new and enhanced tracking technology, courtesy of the so called "Enhanced-911" mandate from the FCC. In this instance, the new privacy invasion is mandated, paid for, and only later will it be discovered what secondary uses are made of it. Then there are the technologies which fail. These generally have their failures attributed to non-privacy factors.

5. Default States

In making a purchase, sometimes there is an exchange of information that the buyer sees as needed for the transaction. A good example of this is the provision of credit card information online. It obviously needs to happen to make the purchase happen (absent such services as iPrivacy, which would make this relation more subtle), but what happens to the data afterwards? The consumer, if s/he has considered the issue at all, often believes that nothing should be done other than what needs to be done. The merchant, having considered things at great length, would like to be able to monetize the data in every way possible. As we discuss in the analysis section above, there is currently no easy way to find a merchant who will offer this choice, or to confirm that they are offering the choice that is want. (Again widespread adoption of P3P or related initiatives could change this.)

Informally, consumers feel strongly that they should not have to pay extra for their privacy to be protected. They feel taken advantage of if the basic transaction as they see it is not respected.

The only time we know of that this has been tested in a vote, the people of North Dakota voted to require banks to get permission to re-sell data, rather than offer them the choice of opting-out. This vote demonstrates that when offered the choice about their privacy (in the form of the right to be left alone),

those voters chose to make the default that information be used for the purpose for which it was provided.

Of course, this was a small vote: 128,206 ballots were cast, of which 119,028 voted on the question—the most votes cast on any question, compared to 113,182 on the other constitutional ballot question, or 108,747 votes cast in the US Congressional race. It would be incorrect to draw too many conclusions from the vote, as only twenty six percent of voters turned out. However, it is useful to note that the voters acted in a manner consistent with what they have told pollsters, that is, that their privacy matters to them, and to note that more voters who voted voted on this issue than on any other.

6. Why Identity Theft is Not About Identity or Theft

Why have we focused so much of this chapter identity theft? In addition to the above points, it illustrates how the allocation of the costs in protecting privacy do not currently reflect the value and incentives of those with control over its protection.

Malcolm Byrd, introduced in section 1.3 above, ended up in jail because the primary cost of misidentifying him was not born by the criminal who used his name, nor by the police who misidentified the criminal as Byrd, nor by any of the police, prosecutors, employers, credit issuers or others who continue to misattribute crimes to Byrd and act accordingly. The cost has been primarily born by Byrd. In our society while individuals are primarily legally responsible for their reputation, the actions of others (government entities, businesses, etc.) are increasingly causally responsible for how that reputation is constituted. This absurdity has absurd implications.

Current advice to protect oneself against identity theft includes checking one's credit record twice a year (up from once a year only a few years ago). Though prudent in the current US socio-economic environment, making individuals responsible for protecting their identity and reputation by such means is akin to requiring them to leave their homes unlocked while suggesting they check with the local pawn shop to see if any of their things are fenced as stolen. It is not a tremendous comfort that the 'pawn shop' in identity theft is larger, more centralized, and has in recent years made some efforts to return goods to their owners, i.e., correct credit records. Worse, as the far from unique case of Malcolm Byrd illustrates, it may only be a short time before one is well advised to check one's criminal record twice a year as well. In fact, Privacy Rights recommends that you "periodically obtain a copy of your driver's license record from your local DMV" for just such reasons (Privacy Rights Clearinghouse, 2002). For criminal identity theft, there is currently no centralized place to clear your record.

A longterm solution better than prosecuting the identity "thief" while leaving the victim to clean up the mess would be to structure the incentives in collecting, attributing, and dissemination information to accurately reflect costs. We have been looking at criminal records, but the same applies to other areas. If the sending of preapproved credit offers required that the senders bear the expected cost not just of duly reported fraudulent charges but of the resultant reputation damage, such offers might not be worth sending. Similarly if the expected damage caused by sharing of personal financial data were figured into the value of such sharing, there would be no need to push for legislation to allow people to opt in rather than opt out of such sharing. It would not be worthwhile for institutions to share; indeed the amount of data that is even worth collecting would probably greatly diminish as the responsibility not just the benefit for the correct value of that data were accounted.

How might this more accurate accounting be instituted? This is hard to say. Litigation is an easy answer. Another possibility is government reform of standards of evidence, not just for criminal trial but also for arrest, for attributions in best practice business accounting, etc. Many activities such as misdemeanor crimes and small value economic transactions might better be handled without affecting reputation at all. Any suggestion here would be very speculative; however, that some such change may be coming is reflected both in recent legislation in North America and Europe, and more importantly in corporate practice. Companies both large (IBM) and less large (Zero Knowledge) have made a substantial commitment to providing enterprise policy management service to corporations that would attempt to properly manage the data they have. If the proposal we suggest is followed, the potential exists to simplify the problem since less data are likely to be held. And certain types of data currently viewed as private might no longer need to be treated as such.

So far, this proposal still somewhat reflects the squishy worldview that treats Social Security numbers, credit card numbers, and such as quasi-private. This view relies on a notion that these are somehow secrets known only to the bearer of those numbers and those s/he trusts—as if one could have meaningful personal trust with thousands, indeed millions, of others. It also runs together these artificially private numbers with the actually private information associated with them: employment history, purchase history, etc. It is in their capacity to authorize transactions that such data acquire their need for privacy. If one could not use them to gain access to truly personal information, if one could not use them to create attributions of properties or behavior to the person assigned to them, then there would be no need to view them as private.

Much of the modern consumer economy is built on the offering of credit with minimal authentication. While the direct costs for bad authentication of credit card transaction may be primarily born by the credit card industry (assuming the consumer notices them in a timely manner *and* follows all the mea-

sures necessary to remove false charges from their accounts), indirect costs of cleaning credit records, jobs lost or not offered, loans lost, time, psychological effects, etc. are primarily born by the party whose identity is spoofed. If these costs were born by the parties that authenticate improperly *and* by any party that propagates such information it would be financially infeasible for them to continue the cavalier authorization of transactions that have been a hallmark of current practice.

This could be taken to mean that every action we take should be scrutinized and properly bound to us. However, the costs of such an approach, both literal and intangible are astronomical. Alternatively, our responsibility for any action could (at minimum) be proportional to the degree of authentication associating us with that action. Criminal and other personal records are currently reputation management systems with no probabilities (in compiling the entries). However, building such probabilities in is a daunting, perhaps hopeless, task especially given the dynamics of how reliable identifications of various types are.

An even more direct approach may be more viable. For example, suppose that a loan is denied or a job application turned down due to errors in a credit report. Currently the reporting agency is obligated to correct errors documented as such, but it is not liable for any effects of the denied loan, particularly if it is simply passing on information that it acquired in good faith; thus it may be appropriate for the agency to similarly pass on its liability. However, if the agency were responsible for any such losses and required to cover any losses it could not pass on, then it would be much more careful about the data it stores, the supporting documentation of it, the reputation and indemnification of the source of the information, but also it would be more cautious in its sharing of such information with others. This has the advantage that, e.g., preapproved credit cards are not themselves a liability for issuers in this sense. But, pursuing someone to pay for charges on such a credit card might be. The burden of proof that a charge is legitimate would of course be on the card issuer and the merchant, but the cost of rectifying errors, the time and any expense of the consumer in rectifying the errors should also be on the issuer (more strictly the authenticator of the transaction). The same approach applies as well in criminal cases. If someone is arrested based on a misauthentication of outstanding charges, this should count as the false arrest that it is rather than as an unfortunate side effect of due process. And associated liability may propagate from the arresting law enforcement agency through the source of its information.

Identity theft as a privacy problem simply goes away on this approach, to be replaced by the problem of properly authenticating transactions that affect the reputation and/or economic and social freedoms of individuals. This is not without large social and infrastructural costs. For example, it may become much harder to obtain unsecured loans in such a society, and the trend away

from cash transactions may reverse. However, it moves costs and incentives to those responsible for authorizations rather than those on whose behalf such authorizations are usurped.

7. Infrastructure Cost

We have already noted how accurate reflection of the costs of assigning, storing, and disseminating reputation would affect the incentives and behavior of infrastructure elements such as businesses and the components of the justice system. However, even without such reallocation, a more accurate assessment of infrastructure costs might lead to an increased emphasis on privacy.

Spam is a large privacy issue. (This is more from the right-to-be-let-alone aspect of privacy than the self-determination of reputation aspect we discussed in the last section.) But, it is not just an issue of personal inconvenience. Recent estimates of spam put it at approaching half of all email traffic in the US (Krim, 2003). This is a tremendous overhead born by business, government, and individuals. And, part of it comes from the distribution of email addresses without the consent of those who hold the addresses. Recent focus on SPAM as a major public thrust of some of the largest ISPs and software vendors, in addition to recent legislation, is evidence that this invasion of privacy is also recognized as a major cost to business.

Note that kneejerk 'solutions' to such problems, for example, wholesale automatic identification of mail senders, especially by the communications infrastructure, may actually cause more harm than good. This is not likely to deter spammers who have access to large numbers of zombie machines that will count as legitimate senders by protocols and who have easy access to jurisdictions from which sending spam is not a problem. In fact, such solutions provide an incentive for spammers to break into systems or otherwise steal accounts to send spam. Schneier notes, " anti-spam security that relies on positive identification isn't likely to work. It'll mean that more spam will rely on stolen accounts. It'll change the tactics of spammers, but not the amount of spam" (Schneier, 2004).

Thus, such approaches are likely to primarily hamper the private actions of the honest while at the same time making it more likely that they will be attacked and framed for sending spam rather than merely receiving it. For spam and for more directed communications, criminals already know how to communicate anonymously and privately. Another example, they can just steal cell phones for brief use, then toss them and steal more. Still another technique noted in the general press is to compromise a web host and leave files there for others to retrieve. Thus, monitoring communication primarily eavesdrops only on the law abiding.

One counterargument to this is that such activity by criminals involves transactional risk (Schechter and Smith, 2003). Thus, providing general private and anonymous Internet communication removes a disincentive to crime. True enough, but the analysis by Schechter and Smith does not account for the cost of privacy loss. If incorporated, an anonymous communications infrastructure may be more cost effective for the infrastructure providers.

Reduction in privacy also has a cost to security. A commonplace in recent polls is to ask how much privacy people would exchange for increased security. However, it is assumed rather than argued that decreasing privacy increases security. Just the opposite may be true. Law enforcement has made use of anonymous tips for years with the recognition that much of the information so gathered would not have been given without a plausible expectation of anonymity. Very shortly after September 11th, the Anonymizer set up an Web interface "providing anonymous access to the FBI's Terrorism Activity tip page to over 26,000 individuals around the world" (CNN, 2001). They have since added anonymous interface to the Utility Consumer's Action Network. Similarly, the Witness Protection Program relies on the ability to assign people a new identity. In an environment in which all commercial and public actions by individuals is monitored, this possibility becomes far less plausible. To effectively monitor to the degree necessary for effective authentication as discussed in section 1.6, the creation of a new identity would likely be noticed in a commercial database (whose entries would be shared without disincentives to do so). The person who recently turned in Khalid Sheikh Mohammed and received a new identity might not have risked doing so without a plausible new identity possible.

8. Conclusion

We have argued in this chapter that assumptions about privacy are not empirically justified. Contra that people will not pay for privacy we have found that when privacy is offered in a clear and comprehensible fashion, it sells well. Complex technologies offered for sale in response to nebulous threats don't sell well, even when those threats are against important targets. We also found that people are not wildly irrational in their dealings with privacy, especially when the cost of examining and understanding privacy policies and practices themselves is taken into account. Privacy is often a complex topic. Different people use the word to mean different things. What one person considers their deepest secret, another may announce to the world. For example, HIV-positive status is something that many people consider to be very private, but there are activists who make it the core of their public personas. That it is difficult to create products that address these complex needs should come as no surprise.

Finally, we observed that the cost of protecting privacy is not allocated in an accurate way. Reallocating appropriately and properly placing costs with those whose actions create them would remove identity theft as a privacy issue. A correct reallocation would also provide government and business with incentives to increase rather than decrease protection of individual privacy.

References

Acquisti, Alessandro (2004). Privacy in electronic commerce and the economics of immediate gratification. In Feigenbaum, Joan, editor, *ACM Conference on Electronic Commerce (EC'04)*. ACM Press.

CNN (Sept. 14, 2001). Web site passes anonymous tips to FBI. `www.cnn.com/2001/TECH/internet/09/14/anonymous.tips/`. Cf. also Secure Tips Online Program. `www.anonymizer.com/tips/`.

Mailboxes, Etc. Mailbox services. `www.mbe.com/ps/ms.html`.

FTC (2004). National and State Trends in Fraud & Identity Theft: January - December 2003. Federal Trade Commission Report. `www.consumer.gov/sentinel/pubs/Top10Fraud2003.pdf`.

Krim, Jonathan (March 13, 2003). Spam's Cost to Business Escalates. *The Washington Post*, page A01.

Privacy Rights Clearinghouse (2002). Fact Sheet 17(g): Criminal Identity Theft. `www.privacyrights.org/fs/fs17g-CrimIdTheft.htm`.

Schechter, Stuart E. and Smith, Michael D. (2003). How much security is enough to stop a thief? In Wright, Rebecca N., editor, *Financial Cryptography*, pages 122–137. Springer-Verlag, LNCS 2742.

Schneier, Bruce (Feb. 15, 2004). The economics of spam. *Cryptogram Newsletter* `www.schneier.com/crypto-gram-0402.html`

Sullivan, Bob (March 9, 2003). The darkest side of ID theft. `www.msnbc.com/news/877978.asp?0si=-&cp1=1`.

US Census Bureau (1999). Curtain and Drapery Mills, Economic Census Manufacturing Industry Series. Report EC97M-3141B, US Census Bureau.

US Postal Service (1998). Pub. 201 - Consumer's Guide to Postal Services & Products. `www.usps.com/cpim/ftp/pubs/pub201/pub201.htm#H1`.

US Postal Service (2001). Comprehensive Statement on Postal Operations. `www.usps.com/financials/`

Varian, Hal R. (1997). Economic Aspects of Personal Privacy. In *PRIVACY AND SELF-REGULATION IN THE INFORMATION AGE*. National Telecommunications and Information Administration, Washington, DC.

Westin, Alan (May 8, 2001). Opinion surveys: What consumers have to say about information privacy. Prepared Witness Testimony, The House Committee on Energy and Commerce.