

“People Won’t Pay For Privacy,” Reconsidered

Adam Shostack
adam@homeport.org

March 14, 2003

Abstract

Privacy consistently tops polls regarding the concerns of Americans. However, after the change in focus or failure of several prominent companies selling privacy technology, it also seems that “people won’t pay for privacy.” Why these apparently contradictory stances? Why don’t people buy privacy technologies? Or do they? By examining these apparent contradictions in light of various definitions of privacy, we can discover where they do and do not hold. In general, privacy seems to work like any other product or feature: when people know they have a problem, an effective and understandable solution will sell. This also allows us to consider the likelihood of a “privacy Chernobyl” galvanizing people into taking broad action.

1 Introduction

The idea that “people won’t pay for privacy” is widespread, and worth examining. Privacy has for the last several years been a hot-button topic, near or at the top of polls about the concerns of Americans and Europeans. Many companies were founded to sell a variety of privacy technology, and none did well. So how can we reconcile this apparent contradiction? I examine a number of technological privacy initiatives and product offerings and consider why they have failed. We then ask if those failures were evidence that “people won’t pay for privacy” or evidence that people are making rational economic choices regarding privacy.

I am motivated by the general application of economic thinking to security issues, by my involvement in apparent privacy failures, and by a belief that people are generally pretty clever in their pursuit of their own self-interest. Given the ongoing, and apparently widely shared belief that privacy is important, I find it curious that getting people to pay for privacy seems so hard.

I also examine a related economic issue, namely people giving away their privacy for very little, or as Austin Hill has put it, people will tell you that privacy is very important to them, but then give you a DNA sample in exchange for a Big Mac. There is an apparent irrationality here, which many people find frustrating and inexplicable.

Through the course of this investigation, I find that privacy is paid for when it is offered as a comprehensible feature responding to a particular concern of the buyer, much like any other product. The apparent irrationality that “users want privacy, but won’t pay for it” will be shown to be largely a straw-man. When privacy features are sold properly, they sell well, and when they are confusing, intimidating, or over-priced, they don’t. The one exception to this is that a great many people believe that privacy should be “included” in the price; that the only use of data from a deal should be that needed to fulfill the contract. In the online world, this is often quite a bit of data, and many customers would like to see it used only for the reason it was provided. I examine this idea further in section 6.

1.1 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 which has examined privacy and economics has chosen to focus on a single definition (for example, [5]). By pointing out a logical way

to behave in the context of a single definition of privacy, these analyses may actually contribute to the idea that people are acting illogically.

Therefore, for clarity, I 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 I ignore, because the phrase is so powerful.

Untracability 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 a 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 I 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.” [6] It is the last group often cited as willing to trade the DNA samples for Big Macs.

Given these meanings, I’ll examine how people pay for them. From there, I will examine a number of areas where people don’t pay for privacy. Finally, we’ll see 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. I 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 esthetics and economics of insulation, were approximately 1.8 billion dollars in 1997 [1]. I do not attempt to break down these numbers as to which motivator leads. I 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: Your real physical location is not revealed. Some of it is unobservability with respect to your house-mates or family, who don’t know what mail you’re 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-comprehensive, USPS-financials. Privacy is explicitly listed by both the post office and Mailboxes Etc as a motivator for renting of post-office boxes [4, 2].

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. I 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 which is also served by answering machines with a call-screening function. Caller ID blocking is an untracability 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.¹

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, not a feature you can choose to pay for. Some of this might be separable; for example, in the curtain example, I 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 me 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, it's unclear what such numbers would mean, and thus they could not be used correctly. Therefore, I make no effort to add up these numbers. I 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 Privacy and the Big Mac

Having seen that people are willing to pay for aspects of privacy in a variety of purchases, it's important to consider people are willing to give up aspects of their privacy with little fight, and ask why that's so if people care about their privacy.

Returning to the example of a Big Mac for a DNA sample, I don't dispute the likelihood that it could happen, although to the best of my knowledge, it has not yet. Regardless, the example is uncomfortable because I feel it's likely that lots of people would take the Big Mac. It is worth examining why people might rationally do this. First, there is the question of linkability: As originally posed, the offer did not involve your name and DNA. As such, it's not clearly linkable. This analysis may well be weaseling out of it, and a Big Mac for your social security number and DNA sample would probably work, too. An alternate analysis is that an individual may not understand what her DNA can be used for, or more importantly, against her. It seems likely that in the United States, in the absence of legislation, DNA and DNA analysis is, or will soon be, stored and marketed to health insurers, employers, background checkers, and others. These uses of the DNA information may cause an employer to not hire her.² There is a trend in the US for new databases to start out with very poor consumer understanding, followed by their abuse, misinterpretation, and corruption with low quality data. This is followed by industry-specific legislation. I assume that after it becomes well understood that people are losing jobs, mortgages, and other opportunities to genetic databases with the typical flaws and abuses of new databases, there will be legislation in the US which controls the worst of these. I should note that due to the work of GeneForum means that Oregon has a genetic privacy information law, and for several years had the only information-as-property clause I am aware of, but it was never tested in court, as it was eliminated in 2001.^[3]

Another indicator 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 care about their privacy. I 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. 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" made it hard to determine if you'd actually be left alone if you did bother to fill out the card.

¹A perhaps interesting aside is that I no longer call directory assistance to try to find people, only businesses, because I assume that all my friends have unlisted numbers.

²There are also important questions of family privacy: If one family member agrees to the use of their genetic data, they expose a substantial amount about their siblings and parents. This is probably not obvious when you are asked for consent, and raises troubling questions of what consent means.

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. Understanding what an http cookie isn’t trivial, as it requires some understanding of the idea of a protocol, a server, and statefulness. Understanding the interaction of cookies with tracability 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 you go 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.

The effort required to understand a privacy issue, and its real impact, may be quite high. 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,” and then 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 your 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, important information to understand what the privacy on 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 then we will discover what secondary uses are made of it. Then there are the technologies which fail. These generally have their failures attributed to non-privacy factors.

The now nearly routine disclosures that hundreds or hundreds of thousands of people are put at risk of identity theft by accidental disclosure are also desensitizing people. From the lack of response, we see that people either don’t see how the issue affects them, or, if they do, how they can respond to it. From this we can see that a “Privacy Chernobyl” is highly unlikely. What aspects of privacy matter to different people sufficiently that it is unlikely that any revelation will cause enough people to become angry enough to mobilize for a given course of action.

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, but what happens to the data afterwards? The consumer, if they have 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 I discuss in the analysis section, there is often no easy way to find a merchant who will offer you this choice, or to confirm that they are offering the choice that you want.

Informally, (unless I can get a good reference, reviewers?) 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 I 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 26 voters turned out. However, it is useful to note that the voters acted in a manner consistent with what they tell pollsters, that is, that their privacy matters to them, and that more voters who voted voted on this issue than on any other.

6 Conclusions

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. However, 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.

7 About the Author

Adam Shostack is currently CTO of Informed Security, a startup building patch management tools. Previously, he was director of technology for Zero-Knowledge Systems, a leading maker of privacy tools.

References

- [1] US Census Bureau. Curtain and drapery mills, manufacturing industry series. Technical report, US Census Bureau, 1997.
- [2] Mailboxes Etc. Mailbox services. World-wide web page available at <http://www.mbe.com/ps/ms.html>.
- [3] Elizabeth Neus. DNA-Rights Defenders: Get Off My Genetic Property Money Creates Conflict of Research vs. Privacy, September 2000.
- [4] US Postal Service. Pub. 201 - consumer's guide to postal services & products. World-wide web page available at <http://www.usps.com/cpim/ftp/pubs/pub201/pub201.htm#H1>.
- [5] Hal R. Varian. Economic Aspects of Personal Privacy. In *PRIVACY AND SELF-REGULATION IN THE INFORMATION AGE*. National Telecommunications and Information Administration, Washington, DC, 1997.
- [6] Alan Westin. Opinion surveys: What consumers have to say about information privacy. Prepared Witness Testimony, The House Committee on Energy and Commerce, May 2001.