Protecting Privacy with Economics: Economic Incentives for Preventive Technologies in Ubiquitous Computing Environments

Alessandro Acquisti*
UC Berkeley

Workshop on Socially-informed Design of Privacy-enhancing Solutions in Ubiquitous Computing
Ubicomp 2002

Abstract

Ubiquitous computing environments make the economic analysis of privacy more difficult as they exacerbate information asymmetries and uncertainties. This paper discusses why the actual marketization of privacy is more difficult than its technical protection in these environments. It then focuses on the economic incentives that can justify the adoption of preventive privacy enhancing technologies.

1 Privacy is Easy, Economics is difficult?

Surveys have repeatedly identified privacy as one of the most pressing concerns of those using new information technology.[1] Only in terms of Internet sales, billions of dollars are said to be lost every year because of privacy fears.[2] At the same time, academic research and industry efforts have developed protocols and technologies to protect individuals' privacy in almost any conceivable scenario - from browsing the Internet to purchasing on- and off-line. There is a demand, and there is an offer. So, why is there no market clearing?

This paper takes an economic approach to the study of why privacy enhancing technologies have failed to gain widespread adoption, while privacy and security of personal information have remained a concern for many.

It is clear that the economic incentives have failed to generate alone workable solutions: it seems like privacy is more difficult to *sell* than to *protect*. It could be that economic incentives find so much difficulty in causing technology adoption because privacy itself is a difficult concept to define in economic terms. One might delimit the privacy conundrum by referring to it as the relation between a *subject*; some *information* related to that subject (and possibly a *transaction* that subject might be participating to); and a set of *other parties* (that might or might not interact directly with that subject) that might have an interest in (or access to, or use of, or some other relation with) that information. As already highlighted in the economic literature (see e.g. Varian [1996]), the subject and the other parties are often in positions of information asymmetry with respect to what use will be made of that information. For example, a customer might not know how the merchant will use the information that she has just provided to him on his website. This creates problems to the economic analysis of privacy scenarios and to the design of appropriate economic incentives.

The subject might not know if, when, and how often the information she is providing will be used. In addition, she might not know what damage she will incur because of that information becoming known, she might not know how much profit others will make thanks to that information, or she might not know the benefits she will forego if her privacy will be violated. In extreme cases the subject might not even be aware of the fact that she is revealing information. One might picture price discrimination as a privacy issue, where economic agents might be revealing their "type," or preferences. "Myopic"

*Mail: SIMS, 102 South Hall, Berkeley, CA, 94720. Email: acquisti@sims.berkeley.edu.

[1] Privacy surveys are too numerous to be cited here. The Electronic Privacy Information Center (http://www.epic.org/) maintains archives of privacy related news and links to privacy surveys.

[2] See, for example, Federal Trade Commission (2000).

customers might not even know that their actions reveal their preferences to the merchant (see Acquisti and Varian [2001]). Deeper analysis highlights additional uncertainties: the subject and the parties she is interacting with often evaluate differently the same piece of information. In fact, the specific environmental conditions will affect in many unpredictable ways the value of information.

Ubiquitous computing environments in highly networked societies exacerbate these problems. There, the subject is more often unaware of what information she is revealing, to whom, and when. Ubiquitous computing environments also exacerbate the privacy issue because of the opportunities for linkage they provide between several activities and behaviors of a same individual. Privacy metrics (such as that proposed by Jiang, Hong, and Landay [2002]) are useful formalizations of the domain space. Economics can use these formalizations to assist, in turn, in the design process of privacy enhancing technologies. The rest of this paper therefore will discuss some economic issues related to preventive, avoiding, and detecting technologies. It will then focus on the economics of preventive privacy enhancing technologies.

2 Economic Issues in Privacy Enhancing Technologies

Starting with Posner (1981), economists have shown great interest in the study of privacy. It is clear that privacy is about trade–offs. The most basic one is between the incentive the subject has to *share* information, and the incentives she has to *hide* information.

We can express this trade-off in the most general way as:

$$v_a - c_a \lesseqgtr v_n - c_n \tag{1}$$

where the payoff from using a certain privacy enhancing technology, v_a, minus the cost of using it, c_a, is compared to the payoff and cost of *not* using it. For example, a customer might use an anonymous form of payment, like ECash (see Chaum [1983]), incurring some costs c_a but reducing the risks of bearing future losses, v_n. These losses might be related to having revealed private information during the transaction. Then again, what exactly could those costs, benefits, losses, and risks actually be?

Here lies the crux of the problem: as we move from an abstract representation - as the one proposed above - to practical implementations, we face an intricate web of trade-offs dominated by subjective evaluations and uncertainties. Table 1 gives an example of these uncertainties as it analyses more in depth the customer/merchant example.

Revealing personal identity		Pros	Cons
Privacy protected	Customer	No price discrimination, Sense of security/protection	No targeted services, No discounts in exchange for personal information, Cost of protecting privacy
	Merchant	More 'customer friendly' reputation	Less customer information
Privacy non protected	Customer	Targeted offers, Discounts in exchange for personal information	Price discrimination, Risk of encurring in future costs related to the revelation of personal information
	Merchant	Ability to price discriminate, Better customer relations because knows customer better	Worse customer relations because of customer fears

Table 1.

The above one is only a partial list of items of importance in the sharing and hiding of information between a merchant and its customers. Still, it shows the existence of many and conflicting trade-offs. It also suggests that many items have subjective evaluations or are related to stochastic factors. In fact, we could represent the general relations described in inequality 1 in a different and more complete form:

$$u_t = \delta \left[v_E(a), p^d(a) \right] + \gamma \left[v_E(t), p^d(t) \right] - c_t^d \tag{2}$$

where the utility u of completing a certain transaction t (which involves revealing certain personal information) is equal to some function of the *expected* payoff $v_E(t)$ from completing the transaction,

times the probability of completing the transaction with a certain technology d, $p^d(t)$; plus some function of the *expected* payoff $v_E(a)$ from maintaining certain information private during that transaction, times the probability of maintaining that information private when using technology d, $p^d(a)$; minus the cost of using the technology, c_t^d. Note that d could be a privacy enhancing technology or a non privacy enhancing technology used to complete the transaction. This notation is simply a compact way to compare the costs and benefits of completing a certain transaction trying or not to preserve one's privacy.

Note also that the costs and benefits from completing the transaction might be distinct from the costs and benefits from keeping the associated *information* private. For example, when Alice anonymously purchases a book, she gains a benefit equal to the difference between her valuation of the book and its price. But if her privacy is compromised during the process, she incurs losses which might be completely independent from the price of the book or her valuation of it. Note, finally, that the representation proposed in equation 2 takes into consideration that the payoff could be made of profits gained but also costs incurred. The payoff function u_t above allows to represent the duality implicit in all privacy issues, as well as the distinction between the value of completing a transaction and the value of keeping its associated personal information private: see Table 2.

Privacy	*Reliability*
Benefits from keeping information private / costs avoided keeping information private, or	Benefits from completing a transaction / costs avoided completing a transaction, or
Costs due to not keeping information private / profits missed because of not keeping information private	Costs due to not having completed a transaction / profits missed because of not having completed a transaction

Table 2.

For example, if a certain transaction is completed to gain some benefit, but privacy must be protected in order to avoid losses, then $v_E(t)$ will be positive while $v_E(a)$ will be negative and $p^d(a)$ will enter the payoff function as $\left(1-p^d(a)\right)$.[3] On the other side, if the agent must complete a certain transaction to avoid some losses but privacy ensures her some benefits, then $v_E(t)$ will be negative and $p^d(t)$ will enter the payoff function as $\left(1-p^d(t)\right)$, while $v_E(a)$ will be positive. Within this framework we can compare the utility the individual will gain from various possible actions, and in particular from using or not using a privacy enhancing technology.

This framework also lets us discuss the several factors that affect the decision of the individual whether to use or not a privacy enhancing technology.

The adoption of privacy enhancing technologies (whether they are used to prevent sharing information, to agree on what information should be shared, or to detect when certain information has become public) might be driven by the individual herself or by other parties that individual is interacting with. For example. imagine a customer using an anonymous independent payment service like PrivateBuy, that does not require cooperation on the side of the merchant. Or imagine the merchant providing an anonymous service, for example by accepting ECash payments. Either way, there will be both costs and benefits related to both using and not using preventive systems for all parties also indirectly involved.

Of course, both the individual and the other parties will also have costs when using *other*, non-privacy enhancing technologies to complete their transactions. For example, in equation 2, c_t^d could be the cost of sending an email through an anonymous mix-net system (see Chaum [1981]) or through a conventional, non anonymous channel. In addition, those who want to take advantage of an individual's personal information in many cases will have to incur costs to do so.

Here lies an important economic difference between privacy and anonymity. *Anonymity* requires noise - other similar pieces of information where to blend, mix, and hide one's personal data.[4] But no noise can stop an attacker with enough information, time, and resources. Anonymity then becomes impossible to defend. What instead we are discussing here is the attempt to defend the *privacy* of certain personal data by making the work of the "attacker" (whoever wants to access information revealed during a certain transaction that the individual wants to keep private) costly enough. In equation 2 the probability of maintaining certain information private will therefore depend both on the technical characteristics of the technology being used to complete the transaction but also on the incentives, resources, and costs of those who want to take advantage of the information revealed during that transaction.

[3] Being certain of mantaining privacy would therefore eliminate the risk of $v_E(a)$, while being certain of losing privacy would impose on the agent the full cost $v_E(a)$.

[4] See Acquisti, Dinglesine, and Syverson (2002).

For what relates to the costs c_t^d, a distinction should be drawn between fixed costs of adoption and variable costs of usage. Costs of adoption of anonymous systems might be high, because of the competition from legacy applications and existing patterns of behavior. However, usage costs might be comparable or even less than those associated to non-anonymous legacy technologies. This is relevant both for the individual subjects as well as the parties they interact with.

Individuals might have direct financial costs when they use a privacy enhancing technology (e.g., when they are asked to pay for the service). They might also have indirect costs (such as the adoption costs related to changes in their habits or the additional time spent to complete a certain transaction).

Recent surveys, anecdotal evidence, and experiments (cf. Spiekermann et al. [2001]), have shown that very few individuals are willing to pay for their own privacy. In fact, individuals are actually less concerned about privacy than what they claim to be. Many are willing to provide very personal information in exchange for small rewards. These individuals might be discounting the potential losses from losing control of their personal information[5] with the uncertain probability that such an outcome will take place. Then, they might be comparing the resulting value with the implicit or explicit costs of using an anonymizing technologies, eventually deciding against using them.

From an economic perspective, the subtlety lies in the fact that the individual loses control of her personal information. That information multiplies, propagates, and persists for an unpredictable span of time. Hence, the individual is in a position of information asymmetry with respect to the party she is completing transaction with. The negative utility coming from future potential misuses of somebody's personal information is a random shock whose probability and scope are both almost impossible to calculate. Because of identity theft, for example, an individual might be denied a small loan, a lucrative job, or a crucial mortgage. In addition, even if the expected negative utility could be estimated, when it comes to security of personal information, individuals might to look for immediate gratification, discounting hyperbolically the future risks (for example of being subject to identity theft), and choosing to ignore the danger. Hence, because of ignorance, self-gratification problems, over-confidence, or various other forms of misrepresentation studied in the behavioral economic literature, they might be acting *myopically* when it comes to protecting their privacy even when they might be acting *strategically* (as rational agents) when bargaining for short-term advantages such a discounted price for a certain good they want to purchase.

Other parties (for example, merchants) will have similar trade-offs. The economic analysis for these parties tries to examine whether the usage benefits - if any - of the anonymous system might overcome the usage costs thus justifying the adoption costs. For example, offering privacy services might be costs-saving in ways which are not directly related to the privacy they provide. Certain anonymous payment systems might have authentication devices that decrease the share of frauds or charge-backs compared to on-line credit card payments. In addition, offering to protect privacy might raise customer trust and therefore act as a marketing tool. However, if individuals are myopic about the future potential risks related to the personal information they reveal, also the parties they interact with tend to discount the incentives to take the burden of protecting the personal data of other individuals. In addition, without legal or contractual intervention, the parties which come to control an individual's information have fewer concerns about the use of that information. The database of a merchant, for example, might be hacked and the credit card numbers stored there might be stolen and then illegally used. The customers owning those credit cards will be often unable to identify the party where the "leak" took place. The merchant will not have to pay for its inability to protect its customers' financial information. This implies that without liability for misuse, abuse, or negligence in handling personal information, moral hazard ensues on the side of the other parties.

It must be noted that the above analysis does not distinguish different forms of privacy enhancing technologies. Preventive, avoiding, and detecting technologies (see Jiang, Hong, and Landay [2002]) can be all studied through the above framework. However, there are important differences between the three approaches from an economic perspective. Avoidance places a large burden of responsibility and rationality on the individual, and in particular on her correct evaluation of the trade-offs she is facing. We have seen however that individuals might be myopic with respect to their privacy risks. In fact, avoiding and detecting technologies can even bring additional uncertainties to the economic scenarios depicted above. Avoidance technologies based on agreements between parties with different contractual and bargaining power tend to fail under pressure without appropriate liabilities. Detection can be very difficult, expensive, and limited. It can also come too late, when the damage has already been done. Finally, for both detection and avoidance, it can be very difficult to incorporate

[5]Which are perceived as being small. For example, customers might rely excessively on credit card anti-fraud protections, and assume away several other risks. On the estimation of the social and private costs of privacy, see Gellman (2002).

the behavior of third parties in the contract designed for the subject and the party that subject is directly interacting with. On the other side, preventive technologies can in some cases decrease those uncertainties - in ways that will be described in the next section.

3 The Economics of Preventive Privacy Enhancing Technologies[6]

Preventive privacy enhancing technologies can be described as those that hide certain pieces of information while letting individuals share other pieces. They often do so by providing pseudonyms or pseudo identities.

Information technology can be used to track, analyze and link vast amounts of data. When registering on Amazon.com with an Hotmail.com email address, for example, an individual is only revealing what might be called an "on-line identity," which is not necessarily linkable to the legal identity of the individual. But when the same individual is completing a purchase at Amazon.com with her credit card, then she is also revealing her legal identity (let's call it her "off-line identity") and offering a linkage between that and her on-line identity represented by her email address.

But information technology can also be used to split those data and keep separate information and identities of the same individual in ways that are both effective (in the sense that linking back the information becomes either impossible or just complicated enough to be no longer profitable trying, since the laws of economics work not only for the defenders but also for the attackers of an individual's privacy) and efficient (in the sense that the transaction can be regularly completed with no additional costs for the parties involved). A purchase history at a merchant site, for example, can be associated to an on-line, non personally identifiable account. Then, the balance for that account can paid through one of many anonymous payment technologies (avoiding linkages to off-line identities and preserving anonymity), or with a personally identifiable method of payment, such as a credit card with the owner's name on it (providing linkages between the "off-line", legal identity, and the "on-line" identity of that individual). Similarly, information sharing between merchants can be realized through coupons and referrals that do not reveal the actual identity of the customer, or they might be based on the actual name of the customer.

Some recent economic studies (Acquisti and Varian [2001]; Calzolari and Pavan [2001], Taylor [2002]) have shown something interesting about the economics of privacy in relation to on-line and off-line identities during purchase transactions: when information about customers' tastes and purchase history is available and can be shared among sellers, market laws alone can produce Pareto-optimal outcomes. In Calzolari and Pavan (2001), sharing information between sellers reduces the distortions associated to asymmetric information between buyer and seller. In Taylor (2002), when the seller is facing strategic customers, she will autonomously tend to adopt a policy that protects the privacy of her customers. In a more abstract framework, Friedman and Resnick (2001) have found that "the distrust of newcomers is an inherent social cost of easy identity changes," but persistent pseudonyms can help both the society and the individual.

More specifically, in Acquisti and Varian (2001), under general conditions allowing firms to use cookies makes society better off, because the buyer can benefit from the seller knowing her better and providing her targeted services. In their model, merchants will not benefit from price-discriminating rational customers. That means that they will not benefit from using against them the information that customers have revealed in previous transactions, unless merchants are able to offer something in return. In this latter case also rational customers might find it useful to share certain personal information with merchants - for example their financial information, or their tastes - to facilitate one-click check-outs or personalized recommendations.

There is something important to note about the "identities" discussed in the above papers. These papers all deal with individuals as (economic) agents whose profiles might include information on taste, purchase histories, price sensitivity or risk aversion - all information that might be safely associated to an on-line identity. However, these profiles do not necessarily carry information about those individuals' legal, "off-line" identities. These papers show that for several types of transactions market laws tend towards fair use of "on-line" information, such as on-line accounts, or cookies. In other words there might be economic benefits from sharing and increasing the use of certain "on-line" information. However, these benefits would not be harmed by the protection of more personal and private information, such as the real, legal identity of the customer. Table 3, as compared to Table

[6]Some of the issues in this section are also discussed in Acquisti (2002).

1, shows that even when the "on-line" information is not protected, the merchant and the customer can still mutually benefit from their exchanges.

Revealing only on-line identity		Pros	Cons
On-line information protected	Customer	No price discrimination	No targeted services, No discounts in exchange for personal information/profile, Cost of protecting on-line information
	Merchant		Less customer information
On-line information not protected	Customer	Targeted offers, Discounts in exchange for personal information	Price discrimination
	Merchant	Ability to price discriminate, Better customer relations because knows customer better	

Table 3.

Existing information systems, however, are built in ways that link the various identities of their users. This is of particular concern in ubiquitous computing environments, where individual behaviors under various circumstances and scenarios can be studied and associated each to the other. This ability to profile users has caused concerns but it has not triggered vast adoption of preventive technologies.

The analysis of the incentives to use privacy enhancing technologies proposed in the previous section can help explain why the many technical solutions already proposed have not gained widespread adoption. Since the market of privacy conscious individuals willing to *pay* for their protection is small, it ends up not being satisfied. As the only economic interest in protecting personal information seems to belong to the owner of that information, who is also subject to "immediate gratification," the profit margins in this area of business are low. Low margins and small demand make it very hard for any technology to succeed - except in niche (and possibly disagreeable) markets. While actual usage costs of privacy enhancing technologies are low once adopted, their adoption fees become significant entry-barrier because they involve significant psychological and switching costs. Hence, as merchants decide against offering anonymizing technologies to their customers, the privacy concerned customers choose not to purchase on-line, or to purchase less. A latent, potentially large market demand remains therefore unsatisfied. The analysis of this section therefore shows that there are solid economic incentives to share certain pieces of personal information, but the market alone does not achieve the proper balance of sharing and hiding.

4 So, What can Economics Do?

Possibly, two things.

Firstly, in specific instances, economics can be used to define mechanisms which are privacy enhancing. For example, in anonymous protocols based on the interaction of many agents (see, e.g. Acquisti, Dingledine, and Syverson [2002]), economics can assist in the design process of mechanisms to solve the impasse when no party alone would have the incentive to perform certain actions (for example, sending dummy traffic to other parties in order to increase the level of anonymity in the system). Under an appropriate incentive compatible contract, different parties might be induced to support each other and therefore the anonymity of the system.

Secondly, and more generally, in the framework of socially-informed design of privacy technologies economics can be used to define what information should be shared, and what protected, as in the scenario analyzed in Acquisti and Varian (2001) and as discussed in the previous two sections. Thereafter economics will need to be assisted by law and technology to actually achieve the balances it proposes. Market forces might ensure fair use of data connected to some pseudo identities of individuals. However, because of the adoption costs and trade-offs analyzed in the previous section, they do not guarantee optimal use and appropriate protection of her legal identity. In these cases, legal intervention, on the model of the EU directive on data protection, or as proposed in Samuelson (2000), should place constraints and liabilities on the side of the parties receiving private information, calibrating them in order to compensate the moral hazard and asymmetric information in the market of personal data, and combining them with information technology as a "commitment" device in the system.

5 References

Acquisti, A., (2002). "Privacy and Security of Personal Information: Economic Incentives and Technological Solutions", SIMS Workshop on Economics and Information Security, May 2002.

Acquisti, A. and H. Varian, (2001). "Conditioning Prices on Purchase History", mimeo, UC Berkeley.

Acquisti, A., Dingledine, R., and P. Syverson (2002). "Open Issues in the Economics of Anonymity", mimeo.

Calzolari, G. and A. Pavan, (2001). "Optimal Design of Privacy Policies", mimeo, MIT and University of Toulouse.

Chaum, D., (1981). "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms." *Communications of the ACM*, 24 (2), pp. 84-88.

Chaum, D., (1983). "Blind signatures for untraceable payments." In *Advances in Cryptology. Proc. Crypto '82,* pp. 199–203. Plenum Press: New York.

Federal Trade Commision, (2000). "Privacy Online: Fair Information Practices In The Electronic Marketplace", May 2000.

Friedman, E. and P. Resnick, (2001). "The Social Cost of Cheap Pseudonyms", Journal of Economics and Management Strategy, 10 (2), 173-199.

Gellman, R., (2002). "Privacy, Consumers, and Costs", http://www.epic.org/reports/dmf privacy.html, March 2002.

Jiang, X., J. I. Hong, and J. A. Landay, (2002). "Approximate Information Flows: Socially-based Modeling of Privacy in Ubiquitous Computing", UBICOMP 2002.

Posner, R. A. (1981). "The Economics of Privacy", American Economic Review, 71 (2), 405-49.

Samuelson, P., (2000). "Privacy as Intellectual Property", Stanford Law Review, 52, 1125.

Spiekermann, S., Grossklags, J. and B. Berendt, (2001). "E-privacy in 2nd generation E-Commerce: privacy preferences versus actual behavior", Proceedings of EC'01: Third ACM Conference on Electronic Commerce, ACM, Tampa, Florida, 38-47.

Taylor, C. R., (2002). "Private Demands and Demands For Privacy: Dynamic Pricing and the Market for Customer Information", mimeo, Duke University.

Varian, H. R., (1996). "Economic Aspects of Personal Privacy", mimeo, UC Berkeley.