Identity Management, Privacy, and Price Discrimination

In economics, privacy is usually discussed in the context of consumer preferences and price discrimination. But what forms of personal data privacy are compatible with merchants' interests in knowing more about their consumers, and how can identity management systems protect information privacy while enabling personalization and price discrimination?

Alessandro Acquisti
Carnegie Mellon University

In the economics literature, privacy is usually discussed in the context of consumer preferences and reservation prices: merchants are interested in finding out a consumer's preferences because from those they can infer the consumer's maximum willingness to pay for a good (his *reservation price*). The ability to identify consumers and track their purchase histories, therefore, lets merchants charge prices that extract as much surplus as possible from the sale, which is what economists call *price discrimination*.[1–3] In this context, consumer privacy concerns reduce to individuals' issues with their personal preferences being known to merchants and exploited for profit.

However, economists also acknowledge that consumer privacy is not just about hiding the price paid for a good.[4,5] During any economic transaction, a consumer might rationally wish to share with a merchant certain types of personal data while keeping others private.[6] Consumers can incur several costs (and, also, gain several benefits) when revealing personal data during a purchase: costs associated with spam, profiling, or financial fraud when addresses, names, or financial data are revealed; and benefits associated with targeted offers or personalized recommendations when information about tastes and interests is shared. The personal data shared during a transaction does not need to be *personally identifiable* for those cost or benefits to occur. For instance, a merchant can infer a consumer's preferences without knowing the consumer's name or other public identifiers (such as his or her credit-card number); or, a consumer's status (as a student, a senior citizen, or member of the military) can be shared with a merchant without also disclosing his or her purchasing history.

Identity management systems can support such selective information revelation strategies by giving consumers greater control over which identities are established, which attributes are associated with them, and under what circumstances they're revealed to others. Therefore, such systems allow for transactions in which some level of information sharing is accompanied by some level of information hiding. At the same time, economic views of privacy that are more granular than the one formal micro models ordinarily focus on—that is, privacy as the protection of a consumer's set of preferences—show that there are both costs and benefits when information other than a consumer's preferences and reservation prices are protected or revealed. The issue becomes what dimensions of personal data privacy are compatible with merchants' interests in knowing more about their consumers and their valuations for goods or services. In this article, I'll examine several ways in which identity management systems can protect certain types of information privacy while simultaneously supporting various forms of personalization and price discrimination.

Identity management and privacy

Identity management systems make it possible for individuals and organizations to engage in selective information revelation strategies. [7–9] By offering consumers some control over how their identities and associated attributes are established and revealed to others, they become tools for privacy protection and for an efficient economic balancing of information

hiding and sharing.

Different types of personal information raise different privacy concerns. Merchants can use certain data (such as a consumer's preferences) for price discrimination. Other data (such as the consumer's credit-card number or personal address) can lead to financial fraud or spam. Accordingly, different types of privacy-enhancing identity management strategies protect different types of information. In the context I'm considering here, it's worthwhile to differentiate between privacy-enhancing strategies that aim to provide anonymity and those that aim to provide pseudonymity.

Pseudonymizing technologies can link various transactions (payments, emails, HTTP requests) by the same agent to the same pseudonym identity, although they aren't traceable to her permanent public identifiers (such as her name). *Anonymizing* technologies not only make any transaction from a certain agent untraceable to that agent's permanent, public identity, but also ensure that adversaries can't link together various transactions by the same agent. In the realm of privacy-enhancing electronic payments, for instance, David Chaum's eCash[10] is an example of an anonymizing technology, whereas Steven Low, Nicholas Maxemchuk, and Sanjoy Paul's credit-card approach[11] is better defined as a pseudonymizing technology.

A pseudonymizing technology can protect a purchaser's financial identity during a transaction. But an anonymizing technology, in addition to that, might also protect the purchaser from having her purchase history tracked or might offer the additional psychological comfort of complete anonymity. This technological distinction is important from an economic perspective because different combinations of transaction linkability and traceability allow different types of information to be shared and different types of price discrimination to be implemented.

Price discrimination, identity, and tracking

As noted earlier, price discrimination refers to a seller's ability to provide the same commodity or service at different prices to different consumers. This price is based on the seller's estimation of the price a buyer might be willing to pay for that good.

Price discrimination is very common in all types of markets: at the cinema, in airline booking systems, and, in fact, online, where increasingly sophisticated tracking technologies let merchants adjust prices based on the visitor's location (as revealed by his or her IP address), time spent on the site, cookie information, history of previous purchases, and so on.

Economists distinguish between three types of price discrimination, which they call "degrees" for technical reasons beyond this article's scope. In first-degree price discrimination, prices are based on individual preferences (in the extreme case, individual buyers could receive a customized price matching their maximum willingness to pay, or reservation price, for the good). In second-degree price discrimination, customers self-select into buying different versions or quantities of the good; in other words, the seller offers a menu of options for a product (for instance, standard and premium version), and consumers freely choose the option they desire (and the associated price). In third-degree price discrimination, differential prices are assigned to different consumer segments based on some observable group characteristics, such as age, student status, or geographical location.

Each degree of price discrimination relies on different types of personal information being available to the merchant, and therefore raises different privacy issues. It's generally believed that consumers don't accept price discrimination (the notorious "randomized" price experiment that Amazon.com attempted a few years ago provoked angry consumer reaction[12]). However, customers don't mind price discrimination when it implies lower prices than those charged to other customers—that is, when they benefit from it. Economists usually regard price discrimination favorably because it can be "welfare enhancing:"[1,3,13,14] under certain conditions, it can increase aggregate economic welfare—for instance, when a good wouldn't even be produced unless its producer could target a segment of consumers willing to pay high prices for it. Price discrimination can also increase the welfare of consumers with lower evaluations for a certain good, who otherwise might not have been offered the good at prices matching their willingness to pay.

Given that price discrimination often relies on consumer identification, it might seem incompatible with privacy protection. This, in turn, would imply that adopting privacy-enhancing technologies would come at the cost of the welfare enhancements that price discriminative strategies otherwise provide.[1] Andrew Odlyzko[14] observed that the current privacy debate in e-commerce is fueled by the clash between consumers and merchants around the use of personal information for price discrimination, and that the movement to reduce privacy online might be motivated by the incentives to price discriminate. Consumers want privacy, which implies freedom from being tracked, yet merchants want to track consumers, which implies their ability to charge prices that improve their profits (for a more recent view on this theme by the same author, see Odlyzko's other work[15,16]).

Although such opposing interests are readily observed in many transactions, consumers can use identity management systems for selective disclosure of identity and attribute information. Previous works have hinted at using privacy tech-

nologies in ways that protect certain types of consumer information but allow other individual data to be exchanged;[1,4,17–19] and others have analyzed some of the economic implications of such technologies.[1,20,21] Whether privacy-enhancing identity management systems are compatible with price-discriminative strategies depend on what type of information they're designed to shield, and therefore on the scope of their privacy protection.

Umpteen wart hogs grew up, but two mats tickled Paul. One wart hog grew up, however five dwarves auctioned tickets 4-line pull quote

Price discrimination with identity management

Different degrees of price discrimination are compatible with different privacy-preserving identity management strategies. I'll look at some specific examples of psuedonymity or anonymity in electronic purchases. I should note that using an anonymous payment technology isn't sufficient for protecting an online buyer's identity when other personal information could be revealed during the transaction: a shipping address, the name of the sender in a gift card, an IP address, or unique tastes, selections, or purchasing patterns can lead to the buyer's re-identification, even when he or she uses anonymizing technologies. In what follows, I'll assume a scenario in which the suite of technologies buyers use (from privacy-preserving payment technologies to anonymous email and anonymous browsing) hides all data that could otherwise reveal his or her public identity.

First degree

As noted earlier, first-degree price discrimination consists of differential prices being offered to individual consumers based on their willingness to pay for a good. Although this observation might suggest that in order to implement first-degree price discrimination, a seller must know the consumer's identity, this isn't necessarily the case.

First-degree price discrimination relies more on the seller's ability to estimate individual reservation prices than on its ability to recognize actual individual identities. Such reservation prices depend on consumers' preferences and tastes, and are predictable—for example, by observing consumers' behavior or purchase histories. Therefore, a merchant might estimate them even in presence of pseudonymizing technologies or, under uncertain conditions, anonymizing technologies as well.

First, the merchant can link consumer purchase histories or customer profiles that include a customer's preferences or tastes to individual pseudonymous identities, and even trace them to persistent identifiers (such as an email address used repeatedly to login to an account with a merchant or a loyalty card used for multiple purchases from a grocery-store chain). Such identifiers can be pseudonymous: in other words, they can separate information about the online persona (for instance, the purchase history in a certain account) from the owner of that account's public identity (name, credit-card number, and so on). Such separation can be enforced through pseudonymous payment technologies that link individual transactions by the same subject to each other through persistent pseudonymous identities, but that aren't traceable back to the purchaser's offline identity. An example of such technologies is Low and colleagues' credit-card protocol,[11] which a consumer would use in coordination with other anonymous browsing and messaging technologies. The linkages between transactions help merchants recognize a consumer and offer that person targeted services and prices, although the offline identity isn't revealed.[4] Of course, as noted earlier, given enough complementary information, traffic and trails analysis could let a determined adversary break pseudonymous and anonymous shields and re-identify the buyer's actual identity. A realistic goal of such protective technologies is simply to ensure that such attacks aren't cost effective.

Separating offline and online identities offers additional advantages—for example, it hides financial information about the consumer from the seller (as well as third parties); hence, it can help the consumer (as well as, possibly, the merchant) avoid additional (and possibly even significant) financial fraud and identity theft costs. It also protects against or decreases the risk of non-monetary forms of discrimination or the creation of an individual's digital dossier, shared with third parties over which the consumer has no control. In spite of these protections, consumers still receive targeted or customized offers and differential prices. However, for merchants, the risk is that consumers will engage in economic arbitrage subtracting revenues from them—particularly if pseudonymous accounts are cheap to create.

Second, a form of behavioral-based first-degree price discrimination can also be compatible with anonymizing technologies, as long as the individual reveals certain information: the online merchant observes the (anonymous, rather than pseudonymous) individual who arrives at the site after a certain search query, browses through different pages on the site, and spends time looking at particular products. Similarly to how brick-and-mortar sellers might try to asses the

consumer's willingness to pay for a commodity in face-to-face transactions, based on traits or information other than that consumer's name or social security number, an online merchant might, under certain conditions, estimate an anonymous visitor's willingness to pay for a certain item, and, accordingly, set prices dynamically.

There's obviously a third, intermediate, case in which a consumer can adopt an anonymous payment technology (which in theory makes transactions both untraceable to the originating public, permanent identity, as well as unlinkable to other transactions by the same consumer); however, the consumer actually renders that payment pseudonymous (by design or oversight) by providing information that lets the seller track his or her transactions. For instance, the use of persistent email addresses or PO boxes, online accounts, or cookies in combination with anonymizing technologies might shield the buyer's public identity (for instance, her actual name), while allowing repeated interactions with the seller.

Second degree

Second-degree price discrimination consists of customers voluntarily choosing differential prices for different versions or quantities of a good or service. This form of pricing therefore isn't based on personal information and doesn't rely on individual recognition. Neither linkability nor traceability across different transactions or identities needs to be exploited for this form of price discrimination to be enforced. In other words, a merchant can implement second-degree price discrimination even when customers adopt pseudonymous and anonymous payments strategies that shield personal information (including online information, such as email accounts or IP addresses). Price discrimination of this form is compatible with anonymous transactions—although, as mentioned earlier, the consumer's actions could provide ulterior information adversaries could use for re- identification.

Third degree

In third-degree price discrimination, merchants assign differential prices to different consumer segments based on some observable group characteristics. The combination of privacy technologies and prices targeted to customer segments is discussed by other researchers, such as Partha Das Chowdhury,[17] who suggested the use of ring signature protocols for anonymous purchases. However, even the combination of existing anonymous payments protocols and anonymous credentials (which can prove the ownership of certain group attributes, such as age, employment status, and so on[8]) meets the requirements of this form of price discrimination. The anonymous payment protocol makes the transaction untraceable to the consumer's permanent and public identity, while the anonymous credentials allow her and the merchant to converge on a price for a particular segment of the consumer population.

If both merchants and consumers adopted them, identity management systems could help them fine-tune personal information revelation to strike a balance between consumers' privacy protection and merchants' price discrimination. In this article, I didn't aim to establish which exact proportion of price discrimination and privacy protection may be optimal or simply desirable for society, or whether price discrimination might or might not be inherently privacy intrusive. I pointed out that many such proportions may be possible thanks to technology, with different types of privacy-preserving identity management strategies being compatible with various forms of differential pricing. The ensuing combination of price discrimination and privacy protection could be a desirable middle path whenever unlinkable transactions would end up decreasing social welfare, or when traceable transactions would expose consumers to potential costs by revealing their identities.

Customer reaction to price discrimination isn't always adversarial: Sarah Spiekermann[22] reports on survey results that confirm how consumer reactions to price discrimination depend on what degree of discrimination is imposed on them (the first degree being the least liked, and the second being the most accepted). For consumers, the advantages of adopting privacy technologies that allow for some level of individual tracking lie in the combination of sensitive information protection and the ability to receive personalized and targeted services (as well as, for some consumers, lower prices than other purchasers). For high-value customers, the disadvantages lie in adverse price discrimination.

For merchants, combining privacy-enhancing technologies and price discrimination strategies via identity management systems lets them attract and satisfy the needs of privacy-sensitive consumers without harming their ability to implement pricing and marketing strategies. With identity management systems, technologists have given us tools that let some information be shared while other data is protected; proper use of those tools can meet both merchants' and consumers' needs. □

Acknowledgments

I gratefully acknowledge research support from US National Science Foundation grant number 0713361 ("IPS: Evaluating and Enhancing Privacy and Information Sharing in Online Social Networks") from the Carnegie Mellon Berkman Fund, and from the US Army Research Office

under contract number DAAD190210389 ("Personal Information Security and Online Social Networks"). I also thank Ramnath K. Chellappa, Alfred Kobsa, Susan Landau, Deirdre Mulligan, Andrew Odlyzko, Sarah Spiekermann, Hal Varian, and three anonymous referees for their helpful comments.

References

1. A. Acquisti and H.R. Varian, "Conditioning Prices on Purchase History," *Marketing Science*, vol. 24, no 3, 2005, pp. 1–15.
2. C.R. Taylor, "Consumer Privacy and the Market for Customer Information," *RAND J. Economics*, vol. 35, no. 4, 2004, pp. 631–651.
3. G. Calzolari and A. Pavan, "On the Optimality of Privacy in Sequential Contracting," *J. Economic Theory*, vol. 130, no. 1,2006.
4. A. Acquisti, "Privacy and Security of Personal Information: Economic Incentives and Technological Solutions," *The Economics of Information Security*, J. Camp and S. Lewis, eds., Kluwer, 2004.
5. B. Hermalin and M. Katz, "Privacy, Property Rights and Efficiency: The Economics of Privacy as Secrecy," *Quantitative Marketing and Economics*, vol. 4, no. 3, 2006, pp. 209–239.
6. H. Varian, *Economic Aspects of Personal Privacy, in Privacy and Self-Regulation in the Information Age*, NTIA report, 1996; www.sims.berkeley.edu/hal/people/hal/papers.html.
7. D. Chaum, "Security without Identification: Transaction Systems to Make Big Brother Obsolete," *Comm. ACM*, vol. 28, no. 10, 1985, pp. 1030–1044.
8. S. Brands, *Rethinking Public Key Infrastructure and Digital Certificates -Building in Privacy*, MIT Press, 2000.
9. M. Hansen et al., "Privacy-Enhancing Identity Management," *Information Security Tech. Report*, vol. 11, no. 3, 2006, 119–128.
10. D. Chaum, "Blind Signatures for Untraceable Payments," *Advances in Cryptology* (CRYPTO 82), Plenum Press, 1983, pp. 199–203.
11. S. Low, N.F. Maxemchuk, and S. Paul, "Anonymous Credit Cards," *Proc. 2nd ACM Conf. Computer and Communications Security*, ACM Press, 1994, pp. 108–117.
12. "Amazon, The Software Company," *Economist*, 18 Dec. 2001; www.economist.com/displayStory.cfm?Story_ID=393096.
13. H.R. Varian, "Price Discrimination and Social Welfare," *Am. Economic Rev.*, vol. 75, no. 4, 1985, pp. 870–875.
14. A. Odlyzko, "Privacy, Economics, and Price Discrimination on the Internet," *Proc. 5th Int'l Conf. Electronic Commerce*, ACM Press, 2003, pp. 355–366.
15. A. Odlyzko, "The Evolution of Price Discrimination in Transportation and its Implications for the Internet," *Rev. Network Economics*, vol. 3, no. 3, 2004, pp. 323–346.
16. A. Odlyzko, "Privacy and the Clandestine Evolution of E-Commerce," *Proc. 9th Intl Conf. Electronic Commerce* (ICEC 07), ACM Press, 2007, pp. 3–6.
17. P. Das Chowdhury, B. Christianson, and J. Malcolm, "Privacy Systems with Incentives, *Proc. First Int'l Workshop Information Systems*, 2006.
18. P. Das Chowdhury, *Anonymity and Trust in the Electronic World*, PhD thesis, computer science dept., Univ. of Hertfordshire, 2005.
19. A. Acquisti, "Personalized Pricing, Privacy Technologies, and Consumer Acceptance," CHI Workshop on Personalization and Privacy, 2006; www.isr.uci.edu/pep06/papers/Proceedings_PEP06.pdf.
20. S. Koble and R. Böhme, "Economics of Identity Management: A Supply-side Perspective," *Privacy Enhancing Technologies Workshop* (PET 05), G. Danezis and D. Martin, eds., 2006, pp. 259–272; www.petworkshop.org/2005/workshop/call.html.
21. R. Böhme and S. Koble, "On the Viability of Privacy-Enhancing Technologies in a Self-regulated Business-to-Consumer Market: Will Privacy Remain a Luxury Good?," Proc. Workshop on Economics of Information Security (WEIS 07), 2007.
22. S. Spiekermann, "Individual Price Discrimination: An Impossibility?," CHI Workshop on Personalization and Privacy, 2006; www.isr.uci.edu/pep06/papers/Proceedings_PEP06.pdf.

Alessandro Acquisti *is an assistant professor of information technology and public policy at the H. John Heinz III School of Public Policy and Management, Carnegie Mellon University. His research interests include the economics of privacy and privacy in online social networks. Acquisti has a PhD in information systems from the University of California, Berkeley. He is a member of Carnegie Mellon CyLab and Carnegie Mellon Usable Privacy and Security Lab (CUPS). Contact him at acquisti@andrew.cmu.edu.*