



Marketing Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases

Vincent Conitzer, Curtis R. Taylor, Liad Wagman,

To cite this article:

Vincent Conitzer, Curtis R. Taylor, Liad Wagman, (2012) Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases. Marketing Science 31(2):277-292. <http://dx.doi.org/10.1287/mksc.1110.0691>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2012, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases

Vincent Conitzer

Levine Science Research Center, Departments of Computer Science and Economics, Duke University,
Durham, North Carolina 27708, conitzer@cs.duke.edu

Curtis R. Taylor

Department of Economics, Duke University, Durham, North Carolina 27708,
curtis.taylor@duke.edu

Liad Wagman

Stuart School of Business, Illinois Institute of Technology, Chicago, Illinois 60661,
lwagman@stuart.iit.edu

When a firm can recognize its previous customers, it may use information about their past purchases to price discriminate. We study a model with a monopolist and a continuum of heterogeneous consumers, where consumers have the ability to maintain their anonymity and avoid being identified as past customers, possibly at a cost. When consumers can freely maintain their anonymity, they all individually choose to do so, which results in the highest profit for the monopolist. Increasing the cost of anonymity can benefit consumers but only up to a point, after which the effect is reversed. We show that if the monopolist or an independent third party controls the cost of anonymity, it often works to the detriment of consumers.

Key words: anonymity; customer recognition; price discrimination; identity management

History: Received: April 11, 2011; accepted: October 11, 2011; Miguel Villas-Boas served as the guest editor-in-chief and Dimitri Kuksov served as associate editor for this article. Published online in *Articles in Advance* January 24, 2012.

1. Introduction

In the modern information economy, firms routinely use consumer data to target their marketing appeals and to price discriminate (Odlyzko 2003). This practice has, not surprisingly, generated privacy concerns among consumers (Goldfarb and Tucker 2011). The proliferation of online search and transactions on sites such as Google, Amazon, Groupon, and Netflix; social interactions on sites such as Facebook, Twitter, and Tumblr; and credit card transactions, the use of subscription services, and participation in retailer loyalty programs have brought an unprecedented amount of personal data within reach of marketers. Small bits of self-revelation via past transactions can increasingly be collected and reassembled by computers to help create a detailed picture of an individual's identity.

For example, in 2006, Netflix initiated a contest to analyze the movie rental history of 500,000 subscribers and improve the predictive accuracy of its recommendation software by at least 10% (see <http://www.netflixprize.com/faq>). Narayanan and Shmatikov (2008) then showed that the customer data released for that contest, despite being stripped of names and other direct identifying information, could often be “deanonymized” by pinning down an

individual's distinctive pattern of movie ratings and recommendations. In 2010, Netflix said that it was shelving plans for a second contest—bowing to concerns raised by the Federal Trade Commission (FTC) (Lohr 2010).

Similarly, Acquisti and Gross (2009) showed that information about an individual's place and date of birth can be exploited to predict his or her social security number (SSN). Using only publicly available data, they observed a correlation between individuals' SSNs and their birth data and found that, for younger cohorts, the correlation allows statistical inference of private SSNs. These inferences are made possible by the public availability of the Social Security Administration's Death Master File and the widespread accessibility of personal information from multiple sources, such as data brokers or profiles on social networking sites.

In this paper, we analyze the ability of firms to track individual purchasing patterns in order to practice behavior-based price discrimination (Armstrong 2006, Fudenberg and Villas-Boas 2006) and the ability of consumers to block data collection in order to avoid differential pricing. Behavior-based advertising and price discrimination are already ubiquitous

(Odlyzko 2003, Hann et al. 2007). Records containing the sequence of websites visited and the online purchases made by individuals are regularly used to target tailor-made offers to them (Chen 2006, Wathieu 2006, Pancras and Sudhir 2007, Chen and Zhang 2008). Similarly, retailer loyalty programs such as supermarket “MVP” cards allow sellers to observe individual purchase histories on which they can condition offers (Larsen and Voronovich 2004).

Chicago-based electronic coupon seller Groupon was recently accused of price discriminating against consumers who purchased a popular Valentine’s Day offering for a flower provider. When going through the link provided by the Groupon offering, customers allegedly faced higher prices compared to those who did not.¹ Another—now notorious—example of behavior-based price discrimination is the 2000 incident in which Amazon charged past customers higher prices for DVDs that their purchase histories suggested they would likely want to buy.²

There are numerous other cases. Examples include Netflix offering a month of free service to new subscribers, AOL offering special “new customer” accounts that are opened only by revealing credit card numbers that have not been used before for similar offers, and credit monitoring agencies offering promotions to new subscribers. The key to these examples is that sellers have difficulty committing to future prices and—in particular—committing not to use information about past purchases when formulating future offers to consumers.

Although technological advances have allowed sellers to track, store, and process individual purchases, consumers nevertheless do have some control over allowing sellers to record their behavior. For instance, they can exert effort to understand sellers’ privacy disclosures and take actions to circumvent being tracked. Such actions can include erasing or blocking browser cookies, using temporary e-mail addresses for creating new accounts, maintaining several online identities, paying with different or “virtual” credit cards, and spreading purchases among a variety of unrelated vendors (Low et al. 1994). In general, these practices, which can require time, effort, and even money, fall under the rubric of what is known as *identity management strategies* (Acquisti 2008). Hence, it is appropriate to think of firms and consumers engaged in a complex technological game of *hide and seek*, in which sellers attempt to identify individual consumers in order to price discriminate and in which buyers endeavor to conceal their personal data.

To investigate these issues, we construct a model featuring a monopolist who is able to track consumer

purchases. Consumers, however, are able to avoid being identified as past customers (or to “anonymize” their identity), possibly at a cost. In our framework, the firm charges past customers more than what it charges “new” customers because past purchases signal a higher willingness to pay.³ We find that when consumers can freely anonymize, they all individually choose to do so, which—somewhat paradoxically—results in the highest profit for the firm. We show that increasing the cost of obtaining anonymity can benefit consumers but only up to a point; at that point, the effect is reversed.

The intuition for this finding is closely related to the celebrated Coase conjecture (Coase 1972) and runs as follows. When the cost of maintaining anonymity is high, the seller is better able to recognize past customers and to price discriminate against them. Thus, consumers hesitate to make an initial purchase, knowing that this will cause them to pay a premium on future purchases. Anticipating this reluctance by consumers, the seller is forced to offer a lower initial price, and this effect actually dominates the increase in profits arising from price discrimination in future periods. In other words, the seller would prefer to commit itself to not price discriminate based on prior purchases. When the cost of maintaining anonymity is low, consumers, in effect, give the seller this commitment power when they each rationally choose to keep their purchases private.

This paper is related to work in the literatures on intertemporal price discrimination, consumer recognition, and online privacy. Research on intertemporal price discrimination and the “ratchet” effect, where the firm sets higher prices for consumers who signaled higher willingness to pay, dates back to the late 1970s. Stokey (1979) and Salant (1989) showed that intertemporal price discrimination is never optimal for a monopolist who can commit to future prices. This is analogous to the above-mentioned fact that in our model, the monopolist obtains its highest profit when anonymity is costless.⁴

A relatively small literature on consumer recognition and online privacy has begun to develop over the past several years.⁵ Contributions by Chen (1997), Fudenberg and Tirole (1998, 2000), Villas-Boas (1999), Shaffer and Zhang (2000), Taylor (2003), and Chen

¹ See Guy (2011).

² See Ramasastry (2005).

³ To focus on a worst-case scenario, and to provide a robust baseline, we leave to future work settings where the seller may want to give discounts to past customers, for instance, because of diminishing marginal utility.

⁴ Villas-Boas (2004) showed that committing to future prices can also help in a model with overlapping generations of consumers.

⁵ For a general discussion of price discrimination, see Stole (2007). For an analysis of privacy with respect to lawful search and seizure, see Mialon and Mialon (2008).

and Zhang (2008, 2009) developed the notion of consumer recognition and personalized pricing but did not explicitly consider privacy issues in online environments. Fudenberg and Tirole (1998) explored what happens when the ability to identify consumers varies across goods. They considered a model in which consumers can be anonymous or “semi-anonymous,” depending on the good bought. Villas-Boas (1999) and Fudenberg and Tirole (2000) analyzed a duopoly model in which consumers have a choice to remain loyal to a firm or defect to the competitor, a phenomenon they refer to as “consumer poaching.” They showed that in the second period, a firm always has the incentive to offer discounts to a rival firm’s customers who have revealed, through their prior purchases, their preference for the rival firm’s product. Such discounts tend to reduce consumer price sensitivity for a firm’s product in the first period, as consumers rationally anticipate them, and hence prices rise in the first period thanks to anticipated customer poaching. Taylor (2003) showed that firms can be worse off because of competitive targeted pricing. Chen and Zhang (2008) analyzed a “price for information” strategy, where firms price less aggressively to learn more about their customers. They extended their analysis in Chen and Zhang (2009) to show that price competition can be mitigated by firms vying to distinguish their loyal customers from price-sensitive shoppers, which can actually make firms better off.⁶ In contrast to these works, this paper aims to study the social, firm-level, and consumer-level impacts of the price of anonymity, when consumers have the ability to decide whether or not to remain anonymous, even after they make purchases.

Closest to our work is an emerging literature on optimal online privacy policies. These were first studied by Taylor (2004), Acquisti and Varian (2005), and Calzolari and Pavan (2006). Fudenberg and Villas-Boas (2006) and Esteves (2010) offered surveys of this literature. Taylor (2004) and Villas-Boas (2004) showed how strategic consumers could make firms worse off in the context of dynamic targeted pricing. The reason is that once a consumer anticipates future prices, he or she may choose to forgo a purchase today to avoid being identified as a past customer and thus be able to purchase at the lower price targeted at new consumers. This strategic “waiting” on the part of consumers can hurt a firm through both reducing sales and diminishing the benefit of price discrimination. Acquisti and Varian (2005) showed that it is never profitable for a monopolist to condition its pricing on purchase history unless a sufficient proportion of consumers are not sophisticated

enough to anticipate the firm’s pricing strategy or the firm can provide enhanced services to increase consumer valuation in subsequent purchases. Acquisti and Varian (2005) also began to study consumers’ use of anonymizing technologies (so as to circumvent identification by a firm as a past customer) but did not fully study the welfare implications. Calzolari and Pavan (2006) considered the case where two principals sequentially contract with a common agent and where the upstream principal can sell its information to the downstream principal. They assumed that the agent’s valuations with the two sellers were perfectly correlated (similar to our setup, where valuations are constant across time). In their setting, the second principal posted its contract after the consumer decided whether to accept the contract of the first firm. By selling information to the downstream principal, the upstream principal may get some rent. Calzolari and Pavan gave conditions under which, if the upstream principal can commit to privacy, it will choose to do so. They identified cases where, given disclosure among the principals, the increase in the rent that has to be assigned to the agent always offsets any potential benefit from the sale of information, as the agent becomes more protective of revealing information about his type. We obtain a similar result in our framework, particularly when we compare the benchmarks of publicly available purchase histories (full recognition) and no purchase histories (no recognition).

The papers above provide important insights regarding the fundamental tensions between consumer privacy and price discrimination. This paper considers a richer environment in which a firm’s strategic customers can choose to remain anonymous at some cost.⁷ We study how this cost affects equilibrium behavior and welfare. Our model allows us to gain additional insights into the effects of tightening privacy regulation. In particular, we show that although the firm obtains its highest profit when consumers can freely maintain their anonymity, consumers can, under some circumstances, be better off when maintaining anonymity is costly. The reason is that when identity management is very expensive, few consumers will do it. This leaves them with two viable alternatives: buy in the first period and face a high discriminatory price in the future, or maintain anonymity by refusing to buy in the first period and

⁶ Rossi et al. (1996) showed a similar result when consumers are passive recipients of a targeted price and do not react when a firm takes away their surplus.

⁷ Armstrong et al. (2009) studied a related model where consumers can choose to “opt out” of receiving sellers’ marketing. They similarly showed that each consumer has a private incentive to opt out of intrusive marketing but that when all consumers do this, price competition is relaxed and consumers are harmed. However, they did not study intermediate cases where opting out is costly but not prohibitively so; this is a main focus in this paper.

receive a low price offer in the future. To induce consumers to opt for the first alternative, the firm sets a low *introductory* price in the first period that can, on net, benefit consumers. However, we also identify a range on the cost of becoming anonymous, where no consumer maintains his or her anonymity in equilibrium, yet the firm increases prices. In this region of costs, consumer surplus suffers, because a greater cost of anonymity enables the firm to intensify its price discrimination while offering higher introductory prices. Consequently, we find that facilitating privacy over this region benefits consumers and increases overall welfare.

This result is in contrast to the prior literature because it actually agrees with the common intuition that more privacy can be better (even when consumers are strategic). Even more surprising is that this welfare behavior happens in a region of anonymizing costs where no consumer chooses to anonymize. In other words, we find that added privacy can benefit consumers and increase overall surplus, even when no consumer decides to take advantage of it.

We also study two extensions to our base model where the cost of anonymity is determined endogenously. In one, a third party—a privacy gatekeeper—controls the cost of anonymizing. Here, we show that the gatekeeper prefers to bargain with the firm and actually set this cost to zero, which works to the detriment of consumers.⁸ In the other extension, the firm itself has control over the cost of anonymizing. We show that when the firm can set this cost up front, this too can work to the detriment of consumers.

Our analysis also connects with the marketing literature on couponing, market segmentation, and addressability. In our model, the cost of anonymity is, in a sense, a repeat customer's cost of accessing a seller's introductory offers, which are targeted at new customers.⁹ Hence, anonymizing costs play an important role in segmenting the consumer population into (at least) three segments: repeat customers who are identified as such, repeat customers who anonymized, and new (potential) customers, where the latter two segments may overlap.¹⁰ When the firm has control over the cost of anonymizing, it would choose to segment the population optimally. We show that this optimal segmentation is obtained when the cost of anonymity is nil and no price discrimination takes place. Shaffer and Zhang (1995, 2002), Chen et al. (2001), and Chen and Iyer (2002) obtained

complementary results. They showed that price discrimination can lead to intense price competition where firms may have an incentive to (i) decrease the level of accuracy of targeted promotions, (ii) differentially invest in customer addressability, and (iii) seek commitment mechanisms not to price discriminate. In our context, the ability to anonymize at a cost of the firm's choosing facilitates such a commitment mechanism. In this regard, our analysis is also related to Hermalin and Katz (2006), who evaluated the efficiency of various privacy policies. Complementary to their work, we show that a firm's ability to set the cost of anonymity up front enables it to commit to prices before consumers make purchasing decisions. Hence, in setting this cost, the firm is able to control the extent of private information that purchasing consumers disclose.

The remainder of this paper is organized as follows. Section 2 describes the model. Section 3 gives two benchmarks: when there is no customer recognition and when consumers cannot anonymize their identities after purchasing. The equilibrium of the game with the ability to anonymize is derived in §4, along with comparative statics. Section 5 extends the base model to allow for an endogenous cost of anonymity, and §6 concludes. The proofs are relegated to the electronic companion, available as part of the online version that can be found at <http://mktsci.journal.informs.org/>.

2. Model

2.1. The Consumers

There is a continuum of consumers with total mass normalized to 1. All consumers are risk neutral, possess a common discount factor $\delta \in (0, 1]$, and maximize their present expected utilities.¹¹ Each consumer demands at most one unit of a nondurable, indivisible good in each of two periods. Consumer i 's valuation for the good is the same in each period and is determined by the realization of a random variable v_i with support normalized to be the unit interval. Consumer valuations are independently and identically distributed (i.i.d.) according to a cumulative distribution function $F(v)$ with density $f(v)$, which is strictly positive on $(0, 1)$. Consumer i 's valuation v_i is initially private information.

2.2. The Firm

There is a monopolist that produces and sells the good in each period. The firm's production cost is normalized to 0, and the firm possesses the same

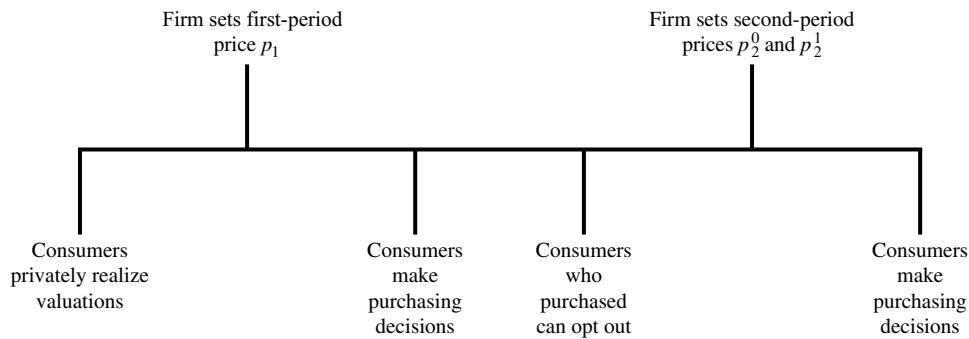
⁸ See Pancras and Sudhir (2007) for a study of an intermediary that stores and sells customer data rather than keep them private.

⁹ See Narasimhan (1984) for an analysis of a consumer's decision of whether to exert costly effort in using coupons.

¹⁰ See Gerstner and Holthausen (1986) for an early analysis of a firm's pricing strategies when market segments overlap.

¹¹ The role of the discount factor δ will be in breaking the indifference of nonmarginal types, i.e., tipping them in favor of buying or not buying in the first period.

Figure 1 Timeline of the Game



discount factor δ as consumers and maximizes its discounted expected profit. It does not observe consumer valuations directly but maintains a database containing purchasing histories. Each consumer is either *anonymous* or *identifiable*. If a consumer is anonymous, then there is no record of any prior purchases by her; i.e., she is not in the database. If she is identifiable, then in the second period the firm knows the purchasing decision that she made in the first period. We emphasize that the firm has no commitment power; i.e., the firm is unable to set and commit to second-period prices in the first period.¹² Because there is a continuum of consumers, each of them realizes that her first-period purchasing decision alone does not affect the prices charged by the firm in the next period.¹³

2.3. The Game

All aspects of the environment, including the distribution of valuations $F(v)$, are common knowledge. At the beginning of the game, all consumers are anonymous. Hence, the firm offers the same first-period price p_1 to all of them.¹⁴ Next, each consumer decides whether to buy the good in the first period, $q_1^i = 1$, or not to buy it, $q_1^i = 0$. Consumers who elect to buy the good also decide whether to let the firm keep a record of the transaction ($r^i = 1$) or to *anonymize* their identity with the firm by deleting the record of the sale ($r^i = 0$). The cost to any consumer who

anonymizes is $c \geq 0$,¹⁵ and we, without loss of generality, assume that this cost is expended in the second period (i.e., it is discounted by δ). This cost represents the time, effort, and any monetary expense of maintaining anonymity. We also allow consumers who purchase to randomize between anonymizing and staying identified. A consumer who does not purchase the good continues to be anonymous and is thus pooled with the buyers who anonymized, from the firm's perspective. At the beginning of the second period, the firm posts a price p_2^0 to the unidentified (anonymous) consumers and a price p_2^1 to the identified ones.¹⁶ Consumers can buy the good only at the price offered to them, $q_2^i \in \{0, 1\}$; i.e., no arbitrage is possible. Hence, a consumer i with valuation v_i who purchases in both periods has (present discounted) utility $v_i - p_1 + \delta(v_i - p_2^1)$ if he does not anonymize and utility $v_i - p_1 + \delta(v_i - p_2^0 - c)$ if he does. Figure 1 summarizes the timeline of the game.

The solution concept we use is perfect Bayesian Nash equilibrium (PBE). A PBE here consists of the firm's strategy (composed of first-period price p_1 and second-period prices p_2^0 and p_2^1 , corresponding to the firm's two information sets in the second period¹⁷),

¹² In a world where the monopolist can commit to its future prices, our analysis of privacy is uninteresting, because no user would pay the cost for remaining anonymous.

¹³ The results hold when there is a finite number of consumers, provided that we add the following assumption: the firm cannot update its beliefs over how many consumers anonymized based on an inventory count. Without this assumption, the firm could infer how many consumers anonymized based on the inventory count. With a continuum of (massless) consumers, inventory expectations on the equilibrium path are confirmed even if a single agent deviates.

¹⁴ We note that, given our setup, the firm cannot do any better by using mechanisms that are more sophisticated than simply posting prices.

¹⁵ The qualitative nature of the results still holds under certain conditions when the cost of maintaining anonymity varies across consumers or when it is correlated with their valuations. It is also possible to extend the analysis to the case where privacy is a continuous choice (e.g., with a convex cost and a concave benefit), but obtaining insightful comparative statics while doing so requires a parameterized cost function, bringing us to a similar, albeit more complicated, setup. To simplify the analysis, we model privacy as a binary choice (although we allow for randomization) and assume that consumers incur the same cost of maintaining their anonymity. See Taylor (2004) for a model with correlated valuations and Acquisti and Varian (2005) for a model with varying levels of consumer sophistication.

¹⁶ Subsection 4.1 considers the case where the firm sets second-period prices *before* consumers decide on whether or not to anonymize.

¹⁷ Technically, the strategy should also specify what the firm would have charged in the second period if it had made a different pricing decision in the first period, but we omit this for notational simplicity.

the consumers' strategies (composed of first-period purchasing decisions q_1^i and anonymizing decision $r^i \in \{0, 1\}$ as a function of p_1 and v_i , and second-period purchasing decision¹⁸ q_2^i as a function of v_i and p_2^i), and the firm's beliefs about consumers' valuations given their identification statuses (F^1 and F^0 for identified and anonymous consumers, respectively¹⁹). These constitute a PBE if all strategies are sequentially rational given the beliefs and the beliefs are consistent given the strategies.

We assume that $p(1 - F(p))$ is concave, i.e., the firm's marginal revenue in a single-period game is decreasing in p (or $-2f(p) - pf'(p) \leq 0$), and that marginal revenue is concave (or $-3f'(p) - f''(p)p \leq 0$). We denote the firm's optimal price in a one-shot version of the game by p^* ; i.e., $p^* = \arg \max_p p(1 - F(p))$.

3. Benchmarks with Exogenous Privacy

3.1. No Recognition

First, consider as a benchmark the case where there is no consumer recognition, so that the firm cannot price discriminate in the second period between consumers who bought and did not buy in the first period. In other words, data on past purchases are exogenously unavailable. Because the firm does not price discriminate based on purchasing history, a consumer will buy the good in each period in which his valuation exceeds the price. Thus, the firm sets the same price in each period, $p^* = \arg \max_p p(1 - F(p))$, generating a per-period profit of $p^*(1 - F(p^*))$. Consumer surplus in each period is given by $\int_{p^*}^1 (v - p^*) dF(v)$.

3.2. Full Recognition

Consider now the opposite extreme, in which data on past transactions are public information. The firm is then able to recognize its previous customers, and consumers are unable to maintain their anonymity at any cost (as in Hart and Tirole 1988, Schmidt 1993, Villas-Boas 2004, Taylor 2004, Fudenberg and Villas-Boas 2006). In this setting, the firm can discriminate between two different groups of consumers in the second period: identified consumers who purchased in the first period and unidentified consumers who did not. The firm consequently sets two different prices in the second period: p_2^1 to identified consumers and p_2^0 to unidentified consumers. (We emphasize again that the firm has no commitment power.)

¹⁸ In principle, the second-period decision can also directly depend on the first-period price, but in equilibrium, it will only depend on v_i and p_2^i .

¹⁹ The firm will also have beliefs about what actions an anonymous agent took in the first period (whether the agent purchased and anonymized or did not purchase at all), but this will not affect the analysis.

PROPOSITION 1 (FUDENBERG AND VILLAS-BOAS 2006). Assume that the history of purchases is publicly observable. Then,

(i) There exists essentially a unique equilibrium that is characterized by prices p_1 , p_2^0 , and p_2^1 and a threshold \tilde{v} such that consumers with valuations $v \in [\tilde{v}, 1]$ purchase in both periods; consumers with valuations $v \in [p_2^0, \tilde{v}]$ purchase only in the second period. The cutoff type \tilde{v} satisfies $\tilde{v} \geq p^*$.

(ii) The equilibrium prices and the threshold \tilde{v} satisfy $p_1 = (1 - \delta)\tilde{v} + \delta p_2^0$ and $p_2^1 = \tilde{v}$ with \tilde{v} and p_2^0 jointly determined from $F(p_2^0) + f(p_2^0)p_2^0 = F(\tilde{v})$ and $\tilde{v} = 1 + \delta(\partial p_2^0 / \partial \tilde{v})(1 - F(\tilde{v})) / f(\tilde{v})$. The first-period price p_1 is then determined from $p_1 = \tilde{v}(1 - \delta) + \delta p_2^0$.

Let us consider the monopolist's pricing strategy toward identified consumers in the second period. If the cutoff type for identified consumers (those who purchase in the first period) \tilde{v} satisfies $\tilde{v} \geq p^*$, then the monopolist sets $p_2^1 = \tilde{v}$. If, on the other hand, $\tilde{v} < p^*$, the monopolist sets $p_2^1 = p^*$. That is, $p_2^1 = \max\{\tilde{v}, p^*\}$. From Proposition 1, since $\tilde{v} \geq p^*$ holds on the path of play of the full-recognition equilibrium, $p_2^1 = \max\{\tilde{v}, p^*\} = \tilde{v}$. Hence, the marginal consumer who buys in the first period—the one with valuation \tilde{v} —gets no surplus in the second period. This is the *ratchet effect* of consumers who reveal their types (Freixas et al. 1985, Laffont and Tirole 1988). The proof of Proposition 1 is in the electronic companion.

Paradoxically, the full-recognition case can result in higher consumer surplus than the no-recognition case, because the firm will need to set p_1 lower to attract consumers in the first period. Correspondingly, in the model with the ability to anonymize, we will show that a low cost of anonymizing one's identity can lead to a prisoner's dilemma situation where consumers maintain anonymity but collectively suffer as a result, compared with the situation where anonymity is prohibitively costly and consumers face the ratchet effect.

By definition, a consumer with valuation \tilde{v} is indifferent between purchasing in both periods and purchasing only in the second period. It follows that the indifference condition that characterizes \tilde{v} is given by $\tilde{v} - p_1 = \delta(\tilde{v} - p_2^0)$. Hence, $p_1 = \tilde{v} - \delta(\tilde{v} - p_2^0)$. Using $p_1 = \tilde{v} - \delta(\tilde{v} - p_2^0)$ and $p_2^1 = \tilde{v}$, one can simplify the firm's present discounted profit to obtain

$$\tilde{v}(1 - F(\tilde{v})) + \delta p_2^0(1 - F(p_2^0)).$$

For $\delta > 0$, since p^* uniquely maximizes $p(1 - F(p))$ and $\tilde{v} \geq p^*$, we have

$$(1 + \delta)p^*(1 - F(p^*)) \geq \tilde{v}(1 - F(\tilde{v})) + \delta p_2^0(1 - F(p_2^0)).$$

Hence, the firm's profit under full recognition is lower than its profit under no recognition. The intuition is that some consumers refrain from purchasing in the

first period because they anticipate a lower price in the next as a result, and the firm is unable to fully recoup the loss in first-period profit by price discriminating in the second period. Hart and Tirole (1988) and Fudenberg and Villas-Boas (2006) showed that if the firm is able to commit to second-period prices in the first period, it would set $p_2^1 = p_2^0 = p^*$, a result that extends to our framework. Hence, the firm's profit under commitment coincides with its profit in the no-recognition equilibrium, i.e., when the firm does not have the ability to track consumers' purchases.

4. Equilibrium with Endogenous Privacy

We now consider the setting in which consumers who purchase in the first period can anonymize at a cost of c . Consumers who purchase in the first period and do not anonymize are identified by the firm in the second period (the firm recognizes that they purchased at price p_1 and did not anonymize) and will be offered price p_2^1 in period 2. All other consumers are offered p_2^0 in period 2. As above, let \tilde{v} denote the lowest consumer type to purchase in the first period. We note that, given that a consumer with valuation \tilde{v} prefers to buy in the first period (i.e., $\tilde{v} - p_1 + \delta \max\{\tilde{v} - p_2^1, \tilde{v} - p_2^0 - c, 0\} \geq \delta \max\{\tilde{v} - p_2^0, 0\}$), then all consumers with valuations $v \geq \tilde{v}$ do as well. Denote by $\alpha(v)$ the probability that a consumer of type $v \in [\tilde{v}, 1]$ anonymizes after purchasing. Then the distribution of valuations among anonymous consumers is

$$F^0(v) = \begin{cases} \frac{F(v)}{F(\tilde{v}) + \int_{\tilde{v}}^1 \alpha(x)f(x)dx} & \text{if } v \leq \tilde{v}, \\ \frac{F(\tilde{v}) + \int_{\tilde{v}}^v \alpha(x)f(x)dx}{F(\tilde{v}) + \int_{\tilde{v}}^1 \alpha(x)f(x)dx} & \text{if } v > \tilde{v}, \end{cases}$$

and the distribution of valuations among identifiable consumers (for $v \geq \tilde{v}$) is given by

$$F^1(v) = \frac{\int_{\tilde{v}}^v (1 - \alpha(x))f(x)dx}{\int_{\tilde{v}}^1 (1 - \alpha(x))f(x)dx},$$

where $F^1(v) = 0$ if $v < \tilde{v}$.

4.1. Partial Commitment

Let us briefly consider the case where the firm posts second-period prices *before* consumers choose whether or not to anonymize. In effect, the firm has some degree of commitment power, because it is able to commit to second-period prices after first-period purchases take place. As the Lemma 1 shows, a consequence of this timing is that, reminiscent of the full-recognition benchmark, no consumer anonymizes in equilibrium.

LEMMA 1. *For any $c > 0$, if the firm posts second-period prices before consumers choose whether to anonymize, then prices satisfy $p_2^1 - p_2^0 \leq c$ —that is, the benefit of anonymizing is smaller than its cost, so no consumer anonymizes in equilibrium.*

From Lemma 1, it follows that the firm's second-period problem here is to choose prices subject to $p_2^1 - p_2^0 \leq c$. Letting $p_2^{0,FR}$ and $p_2^{1,FR}$ denote second-period prices in the full-recognition benchmark, if $c \geq p_2^{1,FR} - p_2^{0,FR}$, it is straightforward to see that the constraint is nonbinding, and full-recognition outcome results. For $c < p_2^{1,FR} - p_2^{0,FR}$, however, the constraint is binding, which can work to the firm's advantage. In particular, the constraint $p_2^1 - p_2^0 \leq c$ enables the firm to commit to not price discriminate against identified customers by an amount larger than c .

In either case, given that the firm posts second-period prices prior to consumers' anonymity decisions, no consumer ends up anonymizing. In effect, the firm has the power to dissuade consumers from taking wasteful anonymizing actions. As will become clear from the following analysis, this limited commitment power benefits the firm relative to the case where prices are posted *after* anonymity decisions.

4.2. Costless Anonymity

Let us now return to the timing in our model, where consumers first choose whether or not to anonymize and the firm *then* posts second-period prices. As a starting point, we first consider the case where $c = 0$. Here, we show that the equilibrium is essentially unique and corresponds to the no-recognition benchmark.

PROPOSITION 2. *When anonymity is costless ($c = 0$), every²⁰ PBE satisfies (and a PBE exists that satisfies) the following:*

- (i) *The firm sets the one-shot monopoly prices $p_1 = p_2^0 = p^*$ and $p_2^1 \geq p^*$.*
- (ii) *Consumers with valuations $v \in [p^*, 1]$ purchase in both periods and anonymize.*
- (iii) *The no-recognition benchmark outcome is obtained.*

This result says that if the cost of maintaining anonymity is nil, then it is in the best interest of every individual who purchases the good in the first period to maintain her anonymity, effectively resulting in the no-recognition outcome from §3.1. However, as indicated in §3.2, this turns out to be exactly what the firm wants.

From the perspective of consumers, we show in §4.5 that this outcome is a prisoner's dilemma situation: individually, each consumer chooses to maintain her anonymity, but as a result, consumer

²⁰ This is excluding the possibility of a PBE in which a measure zero subset of the consumers uses a different strategy.

surplus ends up being lower because there is no price discrimination. In fact, relative to the case where anonymity is costly, every consumer ends up being (weakly) worse off overall when there is no cost associated with anonymizing. In other words, by maintaining anonymity, consumers impose a negative externality on other consumers. Below, we study how firm profit and consumer surplus are affected by the cost c of maintaining anonymity.

4.3. Costly Anonymity

We now move on to the general case in which there is some cost $c \geq 0$. We will restrict our attention to PBEs in which the following holds: all consumers who purchase the good in the first period anonymize with the same probability α (alternatively, α can be thought of as the proportion of consumers who anonymize, i.e., the fraction of purchasing consumers who take actions to maintain their anonymity). This restriction is motivated by the fact that all consumers who purchased in the first period face the same trade-off when deciding to anonymize: pay either p_2^1 or $p_2^0 + c$. (In equilibrium, all first-period buyers will buy again in the second period.) Hence, the benefits of anonymizing (or not) are a function of the prices and the cost of anonymizing. A consumer's level of v only determines whether the consumer is in the group that buys in both periods or only in the second period.²¹ We refer to such an equilibrium as a *pooling equilibrium*.^{22, 23}

²¹ We thank an anonymous referee for helping us clarify this argument.

²² The restriction to pooling equilibria can also be justified using a purification argument. Suppose that, instead of all agents facing the same cost of anonymizing, each agent's cost is drawn from a commonly known distribution. Furthermore, suppose that these costs are drawn i.i.d. across agents and are independent of the agent's valuation. Let d_i , $i \in \mathbb{N}$, denote a sequence of continuous distributions over the anonymizing cost that an individual agent faces such that $\lim_{i \rightarrow \infty} d_i$ is the degenerate distribution on c (where c is the cost in the original game G). Let G^{d_i} denote the cost-perturbed game where each consumer's cost of maintaining anonymity is realized, for simplicity, immediately after the first-period purchasing decisions according to d_i (e.g., because of uncertainties about which information the firm actually tracks, individual shocks regarding opportunity costs, and/or time delays prior to the availability of anonymizing actions). It can be shown that the pooling equilibrium we characterize is the unique equilibrium that results from taking the limit of the equilibria of G^{d_i} when $i \rightarrow \infty$. Appendix B in the electronic companion shows the possibility of an equilibrium in pure strategies.

²³ A pooling equilibrium exists for the following reason. By setting $p_1 \geq 1$, the firm reaps no profits in the first period; similarly, when setting $p_1 = 0$, the firm's first-period profit is 0. In either case, present-discounted second-period profit is given by $\delta p^*(1 - F(p^*))$. The firm can thus perform strictly better by setting $p_1 \in (0, 1)$ and still reap the same second-period profit by setting $p_2^1 = p_2^0 = p^*$. Because the firm's objective is continuous in p_1 , and because the firm's problem is defined on a compact set, it follows that a solution exists in the interior, $p_1 \in (0, 1)$.

The firm's second-period beliefs over valuations in a pooling equilibrium are given by

$$F^0(v) = \begin{cases} \frac{F(v)}{F(\tilde{v}) + \alpha(1 - F(\tilde{v}))} & \text{if } v \leq \tilde{v}, \\ \frac{F(\tilde{v}) + \alpha(F(v) - F(\tilde{v}))}{F(\tilde{v}) + \alpha(1 - F(\tilde{v}))} & \text{if } v > \tilde{v}; \end{cases} \quad (1)$$

and

$$F^1(v) = \begin{cases} 0 & \text{if } v \leq \tilde{v}, \\ \frac{F(v) - F(\tilde{v})}{1 - F(\tilde{v})} & \text{if } v > \tilde{v}. \end{cases} \quad (2)$$

In the second period, the firm chooses its prices to maximize profit according to

$$\max_{p_2^r} (1 - F^r(p_2^r)) p_2^r \quad \text{for } r = 0, 1. \quad (3)$$

The following lemma shows that when $c > 0$, there does not exist a PBE in which all of the consumers who purchased in the first period anonymize. The intuition is that when $c > 0$, it is optimal for the firm to lower the first-period price to attract more customers in the first period. For some of these customers, anonymizing is not a best response unless some other consumers purchase and stay identified.

LEMMA 2. *For $c > 0$, there does not exist a PBE in which all first-period customers anonymize.*

The intuition for Lemma 2 is rooted in the commitment problem of the seller. The seller is unable to directly commit not to price discriminate in the second period but is able to influence consumers' decisions to become anonymous in the second period by raising the first-period price. However, doing so results in a loss of profit, both because of the inability to capture the cost consumers expend on anonymizing and because of lower revenues in the first period. Hence, the seller ends up choosing to set a lower first-period price, and not all consumers anonymize.

The next lemma provides a useful ordering of the equilibrium prices and cutoff type, and it proves that anonymous consumers pay a discounted price in the second period. The key drivers for the result are the seller's inability to commit to second-period prices and consumers' strategic prediction of future prices in the first period.

LEMMA 3. *For $c > 0$, if $\tilde{v} \geq p^*$, then $p_2^0 \leq p_1 \leq \tilde{v} = p_2^1$ holds on the path of play of any pooling equation—that is, new and anonymous consumers pay discounted prices relative to identified consumers.*

The following lemma characterizes the pooling equilibrium for sufficiently small values of the cost of anonymizing, c . (Proposition 3, which follows, gives the relevant range on c .)

LEMMA 4 (POOLING EQUILIBRIUM). *For sufficiently small $c > 0$, every pooling equilibrium satisfies the following:*

(i) *Consumers with valuations $v \in [\tilde{v}, 1]$ purchase in both periods and anonymize with probability α , consumers with valuations $v \in [\tilde{v} - c, \tilde{v}]$ purchase only in the second period, and $\tilde{v} \geq p^*$.*

(ii) *Prices satisfy $p_1 = \tilde{v} - \delta c$, $p_2^1 = \tilde{v}$, and $p_2^0 = \tilde{v} - c$. The firm's beliefs about anonymous and identified consumers' valuations are given by (1) and (2).*

(iii) *For $h(v) = F(v) + vf(v)$ and $h'(v) = 2f(v) + vf'(v)$, the cutoff type \tilde{v} and anonymizing probability α are determined from*

$$\tilde{v} = \delta c + \frac{1 - F(\tilde{v})}{f(\tilde{v})} + \delta \frac{1 - h(\tilde{v} - c) - ch'(\tilde{v} - c)}{f(\tilde{v})}, \quad (4)$$

$$\alpha = \frac{h(\tilde{v} - c) - F(\tilde{v})}{1 - F(\tilde{v})}. \quad (5)$$

Excluding deviations of measure zero, Lemma 4 uniquely determines the behavior on the path of play. The resulting equilibrium, with prices $p_1 = \tilde{v} - \delta c$, $p_2^0 = \tilde{v} - c$, and $p_2^1 = \tilde{v}$, has the following properties. A consumer with valuation at least \tilde{v} will purchase in the first period as well as in the second period and be indifferent between anonymizing and staying identified. A consumer with valuation \tilde{v} will be indifferent among the options of only buying in the first period, only buying in the second period, and buying in both periods. A consumer with valuation at most \tilde{v} will not purchase in the first period and will purchase in the second period if and only if her valuation is at least $\tilde{v} - c$. In essence, the firm offers anonymous customers “introductory” prices in each period (although the introductory price is more attractive in the second period).

We now move on to general values of c (not necessarily small). Let $\alpha(c)$ denote the probability that a consumer who purchased in the first period maintains anonymity, when the cost of doing so is c . (In a pooling equilibrium, by definition, this probability is the same for all agents who purchase in the first period.) Also, let $p_2^{1,FR}$ and $p_2^{0,FR}$ denote the second-period full-recognition benchmark prices, and finally, let $\bar{c} = p_2^{1,FR} - p_2^{0,FR}$.

We proceed with several lemmas that address a higher level of costs, which we then integrate to give a more general characterization of equilibrium. The first lemma addresses the case where the cost of anonymizing is prohibitively high.

LEMMA 5. *For anonymizing costs that are prohibitively high, $c \geq \bar{c}$, the outcome from any pooling equilibrium coincides with the full-recognition benchmark outcome.*

The intuition for this result is the following. For prohibitively high costs of maintaining anonymity,

the full-recognition benchmark outcome is obtainable by the firm. In fact, the firm's problem is a constrained version of its counterpart in the full-recognition benchmark, where its optimal strategy and corresponding outcome in the latter is feasible. In essence, because the cost of anonymizing is high, the ability to do so does little to help the firm's profit in terms of encouraging more consumers to purchase in the first period; in fact, it does more to hurt the firm's profit by failing to capture a significant part of the cost incurred by anonymizing consumers. Consequently, the firm sets prices that encourage consumers to stay identified in equilibrium.

The next lemma addresses the region of cost where consumers no longer anonymize.

LEMMA 6. *Let \hat{c} denote the smallest cost such that $\alpha(\hat{c}) = 0$. Then, $\hat{c} < \bar{c}$.*

The intuition for Lemma 6 relates to that of Lemma 5. In particular, the firm works to mitigate its loss of potential profit from consumers anonymizing by reducing consumers' incentive to anonymize. According to Lemma 6, the firm begins to do so for costs smaller than \bar{c} .

The following lemma shows that, indeed, no consumer anonymizes in equilibrium for all $c \geq \hat{c}$.

LEMMA 7. *For all $c \geq \hat{c}$, $\alpha(c) = 0$ —that is, given anonymizing costs that exceed the threshold \hat{c} , consumers do not anonymize in equilibrium.*

The intuition for this result is that the firm maintains the status quo in terms of anonymizing behavior once it is able to eliminate consumers' incentives to do so at \hat{c} . The firm achieves this by setting a first-period price that is sufficiently low (with a corresponding cutoff type \tilde{v} that is also sufficiently low).

Proposition 3 characterizes the pooling equilibrium across different values of c .

PROPOSITION 3. *Let $p_2^{1,FR}$ and $p_2^{0,FR}$ denote second-period prices in the full-recognition benchmark, and let $\bar{c} = p_2^{1,FR} - p_2^{0,FR}$. A pooling equilibrium exists, and any pooling equilibrium satisfies the following properties on the path of play:*

1. *There exists a threshold anonymizing cost beyond which no consumer anonymizes, given by $\hat{c} \in (0, \bar{c})$, such that $\alpha(c) \in (0, 1)$ for all $c \in [0, \hat{c})$ and $\alpha(c) = 0$ for all $c \geq \hat{c}$.*

2. *For $c \in (0, \hat{c}]$, the unique pooling equilibrium outcome is characterized by Lemma 4.*

3. *For $c \in (\hat{c}, \bar{c})$, no consumer anonymizes. Let \tilde{v} be defined by $F(\tilde{v} - c) + (\tilde{v} - c)f(\tilde{v} - c) = F(\tilde{v})$. In equilibrium, \tilde{v} is nondecreasing in c , with $\tilde{v}(\hat{c}) < \tilde{v}(\bar{c})$; that is, fewer consumers purchase in the first period as c increases. Here, \tilde{v} is set to maximize $\tilde{v}(1 - F(\tilde{v})) \cdot (1 + \delta) + \delta p_2^0(F(\tilde{v}) - F(p_2^0))$, subject to $\tilde{v} \leq \bar{v}$, $F(p_2^0) + f(p_2^0)p_2^0 = F(\tilde{v})$, and $p_1 = (1 - \delta)\tilde{v} - \delta p_2^0$.*

4. For $c \geq \bar{c}$, the outcome from any pooling equilibrium coincides with the full-recognition outcome.


Figure 2 shows how the probability of maintaining anonymity is affected by the cost of anonymizing, c . The region $[\hat{c}, \bar{c}]$ is of particular interest: as we will show shortly, consumer and social surplus (weakly) decrease in this region and the firm's profit (weakly) increases, but the probability of maintaining anonymity is fixed at 0. The various regions can be explained as follows. First, the firm loses profit when consumers maintain their anonymity. When a consumer chooses to anonymize, because the firm's second-period price for anonymous consumers is c lower than for identified consumers (to keep consumers indifferent between anonymizing and staying identified), this effectively costs the firm c . That is, the cost of maintaining anonymity is passed on to the firm. Second, because some consumers anonymize, the second-period price to anonymous consumers, p_2^0 , targets both first-time buyers and repeat customers who anonymized. Hence, (anonymous) repeat customers are interfering with the firm's ability to capture more first-time buyers in the second period, thus lowering the firm's profit. On the other hand, because consumers can anonymize, more consumers decide to buy in the first period, which helps the firm's profit. This latter effect dominates when c is low but is overcome by the former two effects as c grows larger—to the point (at $c = \hat{c}$) where it pays off for the firm to lower the first price sufficiently so that no consumer anonymizes. Once the cost reaches \hat{c} , no consumer will anonymize. This allows the firm to more easily price discriminate as c increases.

When valuations are uniformly distributed, Proposition 3 can be used to more precisely characterize the equilibrium outcome. In particular, we have $\hat{c} = (1 + \delta)/(4 + 3\delta)$, and $\bar{c} = (2 + \delta)/(8 + 2\delta)$; consumers with $v \in [\tilde{v}, 1]$ purchase in the first period and anonymize with probability α , where

$$\tilde{v} = \begin{cases} \frac{1 + \delta - \delta c(1 + 2\delta)}{2(1 + \delta)} + \delta c & \text{if } c \leq \hat{c}, \\ \min\{2c, 2\bar{c}\} & \text{if } c > \hat{c}; \end{cases} \quad (6)$$

$$\alpha = \begin{cases} \frac{1 + \delta - (4 + 3\delta)c}{1 + \delta(1 - c)} & \text{if } c \leq \hat{c}, \\ 0 & \text{if } c > \hat{c}. \end{cases} \quad (7)$$

Figure 2 Equilibrium Probability α of Maintaining Anonymity as a Function of the Anonymizing Cost c on the Path of Play



The firm sets the first-period price

$$p_1 = \begin{cases} \frac{1 + \delta - \delta c(1 + 2\delta)}{2(1 + \delta)} & \text{if } c \leq \hat{c}, \\ \min\{(2 - \delta)c, (2 - \delta)\bar{c}\} & \text{if } c > \hat{c}. \end{cases} \quad (8)$$

Second-period prices satisfy $p_2^0 = \frac{1}{2}(\tilde{v} + \alpha(1 - \tilde{v}))$ and $p_2^1 = \tilde{v}$. The firm's beliefs about anonymous and identified consumers' valuations are given by (1) and (2).

4.4. Firm Profit

For $c \in [0, \hat{c}]$ and $c = \bar{c}$, the firm's present-discounted profit is given by

$$\Pi(c) = (\tilde{v}(c) - \delta\alpha(c)c)(1 - F(\tilde{v}(c))) + \delta p_2^0(1 - F(p_2^0)). \quad (9)$$

At $c = 0$, the firm obtains the no-recognition benchmark profit (where all consumers anonymize) given by $(1 + \delta)p^*(1 - F(p^*))$. This profit was shown to exceed the one in the full-recognition benchmark in §3.2. The following lemma shows that the firm's profit for $c \in (\hat{c}, \bar{c})$ is even lower than at $c = \bar{c}$.²⁴

LEMMA 8. The firm's profit is nondecreasing in c over $[\hat{c}, \bar{c}]$ and is strictly higher for $c \geq \bar{c}$ than \hat{c} .

Lemma 8 and the discussion preceding it indicate that the firm's profit is nonmonotonic in the cost of anonymity: it is higher at $c = 0$ than at either \hat{c} or \bar{c} , but it is lower at \hat{c} than at \bar{c} . Proposition 4 summarizes the above-mentioned findings. We include more precise profit characterization for the special case where valuations are uniformly distributed.

PROPOSITION 4 (FIRM PROFIT). (i) The firm's profits, $\Pi(c)$, are nonmonotone in c . They reach a global maximum at $c = 0$ and are nondecreasing over $[\hat{c}, \bar{c}]$, with $\Pi(c) = \Pi(\bar{c}) > \Pi(\hat{c})$ for all $c \geq \bar{c}$. (ii) In the case where valuations are uniformly distributed, there exists a threshold $c^\# < \hat{c}$ such that $\Pi(c)$ is strictly decreasing over $[0, c^\#]$, strictly increasing over $[c^\#, \hat{c}]$, and constant over $[\bar{c}, \infty)$.


4.5. Uniformly Distributed Valuations

When valuations are uniformly distributed, $p^* = \arg \max_p p(1 - F(p)) = 1/2$. The solution to the firm's second-period problem is given by $p_2^0 = \frac{1}{2}(\tilde{v} + \alpha(1 - \tilde{v}))$ and $p_2^1 = \tilde{v}$. When consumers anonymize (i.e., $\alpha > 0$), we also have $p_2^1 - p_2^0 = c$. Substituting for p_2^0 and p_2^1 , we can obtain $\alpha = \max\{(\tilde{v} - 2c)/(1 - \tilde{v}), 0\}$. From the indifference condition for the marginal type \tilde{v} , we have

$$\tilde{v} - p_1 + \alpha\delta(\tilde{v} - p_2^0 - c) + (1 - \alpha)\delta(\tilde{v} - p_2^1) = \delta(\tilde{v} - p_2^0).$$

²⁴ For all $c \geq \bar{c}$, the firm obtains the full-recognition equilibrium profit.

Figure 3 Comparative Statics in the Case with Uniform Valuations and $\delta = 1$



Substituting for second-period prices and rearranging gives us $\tilde{v} = p_1 + \delta c$ when $\alpha > 0$ and $\tilde{v} = 2p_1/(2 - \delta)$ when $\alpha = 0$. Thus, we have

$$\alpha = \max \left\{ \frac{p_1 - (2 - \delta)c}{1 - p_1 - \delta c}, 0 \right\} \quad (10)$$

and

$$\tilde{v}(p_1) = \begin{cases} p_1 + \delta c & \text{if } p_1 \geq (2 - \delta)c, \\ \frac{2p_1}{2 - \delta} & \text{if } p_1 < (2 - \delta)c. \end{cases} \quad (11)$$

Given $c \geq 0$, the firm's first-period problem is then obtained by solving

$$\max_{p_1(c)} (1 - \tilde{v})(p_1 + \delta(1 - \alpha)p_2^1 + \delta\alpha p_2^0) + \delta(\tilde{v} - p_2^0)p_2^0 \quad (12)$$

subject to $p_2^0 = \frac{1}{2}(\tilde{v} + \alpha(1 - \tilde{v}))$, $p_2^1 = \tilde{v}$, (10), and (11). Letting $\hat{c} = (1 + \delta)/(4 + 3\delta)$ and $\bar{c} = (2 + \delta)/(8 + 2\delta)$, the equilibrium first-period price and marginal type are characterized by the following:

$$p_1(c) = \begin{cases} \frac{1 + \delta - \delta c(1 + 2\delta)}{2(1 + \delta)} & \text{if } c \leq \hat{c}, \\ \min\{(2 - \delta)c, (2 - \delta)\bar{c}\} & \text{if } c > \hat{c}; \end{cases} \quad (13)$$

$$\tilde{v}(c) = \begin{cases} \frac{1 + \delta - \delta c(1 + 2\delta)}{2(1 + \delta)} + \delta c & \text{if } c \leq \hat{c}, \\ \min\{2c, 2\bar{c}\} & \text{if } c > \hat{c}. \end{cases} \quad (14)$$

Figures 3(a)–3(c) show the specific case of $\delta = 1$.²⁵ For $c = 0$, the outcome results in no recognition, with the firm setting the monopoly price in each period, 50% of consumers purchasing in each period, and all of them anonymizing. At the other extreme, if $c \geq \bar{c} = 0.3$, the full-recognition outcome obtains, with 40% of consumers purchasing in the first period, 70% purchasing in the second period, and no consumer anonymizing. Comparative statics over the range of costs slightly greater than $\hat{c} = 2/7$ are particularly interesting. As Figures 3(a) and 3(b) show, there is a steep increase in the first-period price

and a steep decline in first-period purchases. The reason can be gleaned from Figure 3(c): once no consumer anonymizes (at $c = \hat{c}$), as c increases, the firm raises its profit by intensifying price discrimination. It does so by raising the first-period price, which, as the next subsection showcases, is to the detriment of consumers.

4.6. Consumer Surplus

We now turn to consumer surplus. Proposition 4 shows that the firm obtains its highest profit when consumers can costlessly maintain their anonymity. However, this does not immediately imply that consumer surplus is at its lowest in this case, because the total surplus may vary depending on the cost of anonymizing. Specifically, because there is no cost to production, the efficient outcome in this model—the first best—would be for every consumer to obtain the good in each period. Hence, the efficient outcome is not obtained for any $c \geq 0$, because some consumers do not purchase.

For $c \in [0, \hat{c}]$ and $c \geq \bar{c}$, consumer surplus as a function of c , $CS(c)$, is given by²⁶

$$\begin{aligned} & \underbrace{\int_{\tilde{v}}^1 v f(v) dv - (1 - F(\tilde{v}))(\tilde{v} - \delta c)}_{(*)} \\ & + \underbrace{\delta \left(\int_{\tilde{v}-c}^1 v f(v) dv - (1 - F(\tilde{v}-c))\tilde{v} + (F(\tilde{v}) - F(\tilde{v}-c))c \right)}_{(**)}. \end{aligned} \quad (15)$$

In (15), (*) is the consumer surplus from first-period transactions, where consumers with valuations $v \in [\tilde{v}, 1]$ purchase the good and pay a price $p_1 = \tilde{v} - \delta c$; (**) is the consumer surplus from second-period transactions, where consumers with valuations $v \in [\tilde{v}, 1]$ are repeat customers and end up expending \tilde{v} (factoring in the cost c), and consumers with valuations $v \in [\tilde{v}-c, \tilde{v}]$ are first-time customers who receive a price discount of c .

²⁵ Given $\delta = 1$, we have $\hat{c} = 2/7$ and $\bar{c} = 3/10$. For $c \in [0, 2/7]$, we have $p_1 = \frac{1}{4}(2 - 3c)$, $\tilde{v} = (2 + c)/4$, and $\alpha = (2 - 7c)/(2 - c)$. For $c \in [2/7, 3/10]$, we have $p_1 = c$, $\tilde{v} = 2c$, and $\alpha = 0$.

²⁶ Proposition 5 shows that the expression in (15), evaluated at $c = \bar{c}$, gives a lower bound on consumer surplus for $c \in (\hat{c}, \bar{c})$. Consumer surplus for $c \geq \bar{c}$ equals consumer surplus at $c = \bar{c}$, because the equilibrium outcome is unchanged.

PROPOSITION 5 (CONSUMER SURPLUS). (i) *Consumer surplus, $CS(c)$, is nonincreasing over $[\hat{c}, \bar{c}]$, with $CS(c) = CS(\bar{c}) < CS(\hat{c})$ for all $c \geq \bar{c}$.* (ii) *In the case where valuations are uniformly distributed, CS is nonmonotone in c ; it reaches a global minimum at $c = 0$ and a global maximum at \hat{c} , is strictly increasing in c over $[0, \hat{c})$, and is strictly decreasing over $[\hat{c}, \bar{c}]$. Moreover, each consumer is individually (weakly) better off under $c > 0$ than at $c = 0$.*

Surprisingly, Proposition 5 shows that consumers can actually be worse off as the cost of anonymizing increases. In other words, facilitating privacy can actually improve consumers' welfare. The reason is the following. First, for costs $c \in [\hat{c}, \bar{c}]$, no consumer anonymizes. Second, as c rises, the fraction of the market that buys in the first period decreases as $\bar{v}(c)$ is nondecreasing; moreover, the firm's optimal prices, which are weakly increasing in c , also change to the detriment of consumers. Hence, the cost of anonymizing has a passive effect on the market outcome, even if this cost is never incurred in equilibrium.²⁷

This finding extends the previous literature, which finds that strategic consumers are better off under full recognition compared with no recognition. The above result brings forth the observation that although consumers may not anonymize in some region of cost (i.e., for $c \in [\hat{c}, \bar{c}]$), they could still be better off when privacy is more accessible, even if they choose not to anonymize. Said another way, there are different shades of "full recognition," and as far as consumers (and overall welfare) are concerned, some are better than others—namely, the ones where privacy is more accessible. In this region of costs, consumer surplus suffers, because a greater cost of anonymity enables the firm to intensify its price discrimination while offering higher introductory prices.

The last part of Proposition 5 clearly demonstrates the prisoner's dilemma nature of the outcome: when $c = 0$, individually, each consumer chooses to maintain her anonymity; however, every consumer ends up being (weakly) worse off as a result of there being no price discrimination. The reason is that by anonymizing, consumers impose a negative externality on other consumers; a nonzero cost of maintaining anonymity serves to at least partially alleviate this negative effect by diminishing the incentive to anonymize. We relegate intuition for the rest of Proposition 5 to the next subsection, where we also address social (total) surplus.

4.7. Social Surplus

Social welfare in our framework can be interpreted in two different ways, depending on whether the cost of

anonymizing is deadweight loss or collected as a fee by a third party (for example, one can rent an anonymous postal box for a fee). In the former case, social surplus equals firm profit plus consumer surplus; in the latter, social surplus is higher than this sum if consumers anonymize at positive costs.

Over the region of anonymizing costs $[\hat{c}, \bar{c}]$, because no consumers choose to maintain their anonymity, social surplus is simply a function of how many consumers purchase in equilibrium. We thus have the following result.

PROPOSITION 6. (i) *Social surplus, $SS(c)$, is nonincreasing over $[\hat{c}, \bar{c}]$, with $SS(c) = SS(\bar{c}) < SS(\hat{c})$ for all $c \geq \bar{c}$.* (ii) *With uniform valuations, if the cost c is deadweight loss, SS is nonmonotone in c ; it reaches a global minimum at $c^+ \in (0, \hat{c})$ and a global max at \hat{c} , is strictly decreasing over $[0, c^+)$, is strictly increasing over $[c^+, \hat{c}]$, and is strictly decreasing over $[\hat{c}, \bar{c}]$. (iib) If c is not wasted, SS differs from (ia) by having a global minimum at $c = 0$ and is strictly increasing over $[0, \hat{c}]$.*

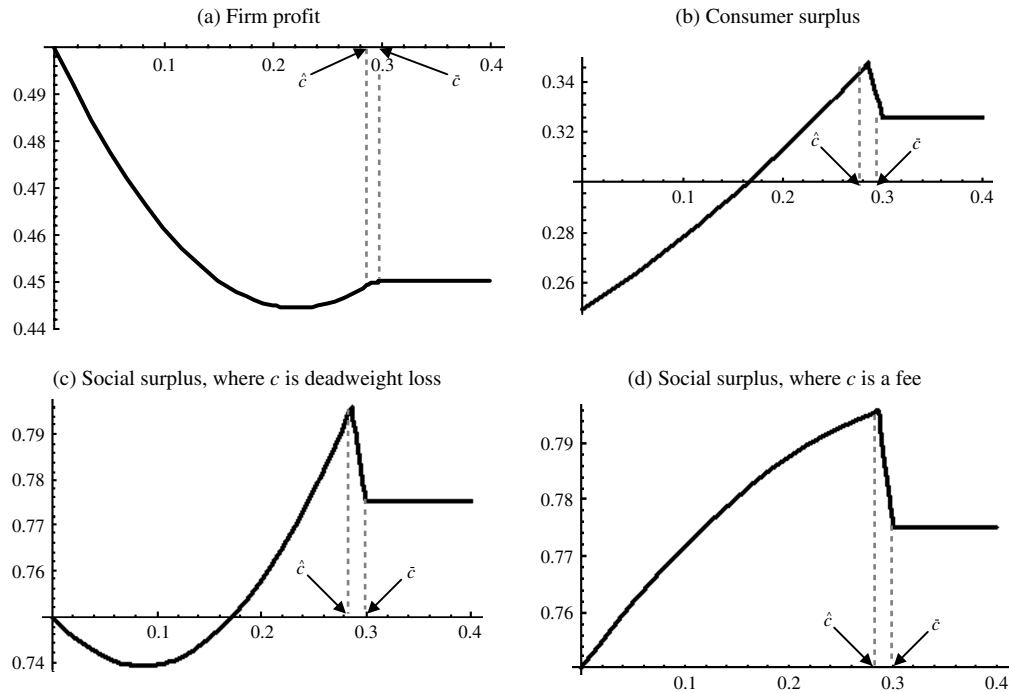
Figures 4(a)–4(d) show comparative statics for the uniform case when $\delta = 1$.

The intuition is as follows: when $c = 0$, all consumers who purchase in the first period choose to maintain their anonymity. As c begins to rise, consumers must pay a nontrivial cost to anonymize. The firm resorts to reducing the first-period price to counteract the negative effect on consumers' buying incentives: consumers are reluctant to purchase in the first period because of the cost they will incur in the second period either from anonymizing or from paying a high price. This results in a lower profit for the firm than it would have had with a lower c , but it also results in higher consumer surplus. The case where c is deadweight loss also results in lower social surplus because many consumers anonymize in this region of cost.

As c approaches \hat{c} , fewer consumers anonymize. This gives the firm more flexibility in setting second-period prices and allows it to better price discriminate, which leads to a slight increase in profit. The firm continues to depress the first-period price over this range; additionally, better price discrimination allows the firm to target more low-valuation consumers in the second period. This results in higher consumer surplus. Hence, as c approaches \hat{c} , both profit and consumer surplus are increasing so that social surplus increases when the cost c is wasted. When it is not wasted, social surplus is increasing but not as steeply because there is no surplus recovered directly from fewer consumers anonymizing (we recall that in this case, social surplus equals the sum of profit, consumer surplus, and the total cost of anonymizing). When c is in $[\hat{c}, \bar{c}]$, no consumer anonymizes. The firm increases prices in this range to better price discriminate in the second period, which

²⁷ We thank an anonymous referee for crystalizing this intuition.

Figure 4 Comparative Statics When Valuations Are Uniformly Distributed and $\delta = 1$ as a Function of c



results in fewer consumers purchasing and lower surplus overall.

5. An Endogenous Cost of Anonymity

In this section, we consider two simple extensions of our model. In the first extension, the firm itself is able to directly control the cost of maintaining anonymity, and it can possibly collect this cost as a fee. In the second, a third party—a privacy gatekeeper—enables consumers to shop anonymously for a fee of its choosing.

5.1. In-House Privacy

In the case where the firm sets the cost of anonymizing, c , there are several possibilities to consider. First, the firm may be able to set c at the beginning of the game. Alternatively, the firm may not be able to commit ex ante to a particular level of c , whereby c is only determined after first-period purchases take place. Second, the firm may also be able to collect c as a fee. The following result characterizes the firm's behavior.


PROPOSITION 7 (IN-HOUSE PRIVACY). *If the firm can set the cost of anonymizing c (i) prior to first-period purchases, it would set $c = 0$, and the no-recognition outcome would result; (ii) after first-period purchases, it would set $c \geq \bar{c}$, which would result in full recognition. The firm would make the same choices if it were to collect c as a fee.*

Figure 5 depicts the firm's profit when it does and does not collect the anonymizing cost c as a fee for

the case of uniform valuations and $\delta = 1$, where c is set prior to first-period purchases.

Although the firm's ability to set the cost of anonymizing up front gives it the power to commit to having the same prices in both periods, perhaps more interesting is that the firm can do no better than the no-recognition outcome even when it can collect anonymity payments as fees. We also observe that if the cost of maintaining anonymity were constrained to be above a relatively high threshold (e.g., a threshold $c \in [\hat{c}, \bar{c}]$ resulting from a minimum level of anonymizing effort), it is in the firm's best interest to make it costlier for consumers to anonymize. In fact, the firm's profit is maximized when it increases the cost of anonymizing to a prohibitively high level. Therefore, the firm's ability to

Figure 5 Firm's Profit as a Function of c



Note. The dashed line indicates that c is collected as a fee; the solid line indicates that c is not collected.

set the cost of anonymizing up front tends to work to the detriment of its customer base.²⁸

When the firm cannot commit to the cost of anonymizing up front, it would choose to set it at a prohibitively high level, which results in the full-recognition outcome. Hence, the firm's lack of commitment regarding anonymizing costs can work to the benefit of consumers (e.g., relative to the case where the cost of anonymizing is low). It is interesting that in this case as well, the firm effectively chooses not to collect any anonymity payments (when it has the ability to do so) by making these payments prohibitively large.

5.2. Privacy Gatekeeper

Let us now suppose that the firm does not control the cost of maintaining anonymity. Instead, there is a third party—a privacy gatekeeper of sorts—that is able to set and commit to a fee for anonymizing at the beginning of the game.^{29, 30} We also consider a second possibility where the gatekeeper can concurrently negotiate with the firm for setting the anonymizing fee at a particular level. This can be modeled most simply with a take-it-or-leave-it offer (e.g., the gatekeeper offers the firm terms given by a price schedule $x, y > 0$, where the firm must pay an amount x for setting $c = y$). We assume that the gatekeeper can commit to a price schedule up front.³¹ The following proposition characterizes the gatekeeper's behavior in equilibrium.

PROPOSITION 8 (PRIVACY GATEKEEPER). *If a privacy gatekeeper sets the cost of anonymizing ex ante and can (i) only charge consumers, it sets $c = c^* \in (0, \hat{c})$, where*

²⁸ In our base model, the firm essentially subsidizes anonymity costs by offering a discounted first-period price. However, this is not a "true" subsidy in the sense that when consumers face their anonymizing decisions, they take into account the full cost c . Indeed, in the case with uniformly distributed valuations, the firm has an incentive to subsidize anonymizing costs up front on a range of costs (specifically, lower costs), and in some cases (higher costs), it would prefer to raise the cost of anonymizing. We thank an anonymous referee for bringing this to our attention.

²⁹ For instance, a consumer could rent a postal box from UPS so as not to disclose her home address. However, UPS could make such an "anonymizing service" much simpler: it could provide customers with individualized one-time codes. Consumers would give these codes to sellers instead of their home addresses. When sellers ship purchases via UPS, they would print the corresponding codes on the labels, and UPS would use these codes to determine the customers' addresses. Google Checkout provides a related service to mask a consumer's e-mail address in online transactions (at no charge to consumers).

³⁰ If the fee for anonymizing is charged in the first period, the gatekeeper sets c/δ , which corresponds to setting c in the base model.

³¹ The results here readily extend to the more general case where the outcome of a negotiation between the gatekeeper and the firm is obtained using a Nash bargaining solution with arbitrary bargaining powers.

$c^* = \arg \max_c \alpha(c)(1 - F(\bar{v}(c)))c$ (that is, some consumers pay a fee to the gatekeeper in order to anonymize); or (ii) charge consumers and negotiate terms with the firm for setting c , it sets $c = 0$ (that is, all purchasing consumers anonymize at no cost and the gatekeeper is paid by the firm).

The intuition for this result is as follows. When anonymizing is costly, the firm offers a discounted price to consumers who purchase in the first period. These discounts are enjoyed by consumers who subsequently anonymize as well as by those who remain identified. The gatekeeper's revenues from anonymizing fees thus only partially capture profit that is lost to the firm. Hence, if negotiation with the firm is possible, there is always a positive surplus associated with an agreement between the gatekeeper and the firm to make anonymizing costless to consumers.

6. Conclusions

We studied a model in which a firm is able to recognize and price discriminate against its previous customers while consumers maintain their anonymity at some cost. We showed that the firm obtains its highest profit when consumers can costlessly maintain their anonymity, but consumers can be better off when maintaining anonymity is costly, although only up to a point. We then considered two extensions to the base model. In the first extension, the firm controlled the cost of anonymizing, which worked to the detriment of consumers when the firm could set this cost up front. In the second, a third party—a privacy gatekeeper—controlled this cost. Here, we showed that the gatekeeper preferred to bargain with the firm and actually set the cost of anonymizing to zero, which also worked to the detriment of consumers.

This paper suggests that certain aspects of consumer privacy may be misjudged by policy makers, firms, and consumer advocacy groups. In particular, facilitating privacy can work to reduce consumer and social surplus when the cost of maintaining anonymity is already low, although the opposite (or a neutral) effect takes place at higher costs. Of course, in practice, many other considerations need to be taken into account that are not within the scope of our model, such as the intrinsic value of privacy and the (possibly accidental) release of sensitive information and corresponding spillover effects—for example, the release of an individual's medical records to her employer.

Our model can, in principle, be extended to take certain other economic considerations into account. An obvious direction is to study a setting with competition. Indeed, we have some initial findings in a similar model with two firms selling differentiated products; these findings suggest that phenomena similar to those identified above continue to

occur. Another direction for further study is to consider settings where consumers obtain some benefit from being identified, such as smaller search costs or better technical support. Our preliminary finding here is that the results from the basic model carry through except when the relevant range of costs grows larger.³² One can also consider an opt-in policy. For instance, the firm could pay consumers to be identified (as in the case of membership programs that offer discounts).³³ Another direction is to increase consumer heterogeneity by, for instance, allowing each consumer to have a different, possibly correlated cost for anonymizing. One can also enrich the model by studying consumers with diminishing marginal utility for future units of the product. Finally, it would be interesting to study the steady-state equilibrium of an infinite-horizon version of our model with overlapping generations of consumers.

There are, in fact, a multitude of questions concerning issues of privacy that are both interesting and potentially important. A primary message of this paper is that the answers to such questions may not be as obvious as they first appear.

Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://mktsci.journal.informs.org/>.

Acknowledgments

The authors have benefited from helpful discussions with session participants at the 2009 Summer Meeting of the Econometric Society, the 2009 FTC-Northwestern Microeconomics Conference, the 2009 NET Institute Conference, the 2009 International Industrial Organization Conference, and workshops at Duke University. They especially thank Atila Abdulkadiroglu, Giuseppe Lopomo, Huseyin Yildirim, Rachel Kranton, Hanming Fang, David McAdams, Sasa Pekec, Tracy Lewis, Daniel Graham, Vasiliki Skreta, and Eric Rasmusen for useful feedback. V. Conitzer's research is supported by an Alfred P. Sloan Fellowship and by the National Science Foundation [Award IIS-0812113 and CAREER Grant 0953756]. C. R. Taylor's research is supported by a NET Institute grant. L. Wagman's research is supported by the Program for Advanced Research in the Social Sciences Fellowship and by an Illinois Institute of Technology Summer Research grant.

References

Acquisti, A. 2008. Identity management, privacy, and price discrimination. *IEEE Security Privacy* 6(2) 46–50.

³² It is worth nothing that, in such environments, consumers could benefit from having multiple accounts: one account to obtain these benefits and another to obtain potential access to lower prices.

³³ See Campbell et al. (2011) for a recent study of the effects consumers' costs of opting in may have on innovation.

- Acquisti, A., R. Gross. 2009. Predicting social security numbers from public data. *Proc. Natl. Acad. Sci. USA* 106(27) 10975–10980.
- Acquisti, A., H. Varian. 2005. Conditioning prices on purchase history. *Marketing Sci.* 24(3) 367–381.
- Armstrong, M. 2006. Recent developments in the economics of price discrimination. R. Blundell, W. K. Newey, T. Persson, eds. *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress of the Econometric Society*. Cambridge University Press, New York, 97–141.
- Armstrong, M., J. Vickers, J. Zhou. 2009. Consumer protection and the incentive to become informed. *J. Eur. Econom. Assoc.* 7(2–3) 399–410.
- Calzolari, G., A. Pavan. 2006. On the optimality of privacy in sequential contracting. *J. Econom. Theory* 130(1) 168–204.
- Campbell, J., A. Goldfarb, C. Tucker. 2011. Privacy regulation and market structure. Working paper, Rotman School of Management, University of Toronto, Toronto.
- Chen, Y. 1997. Paying customers to switch. *J. Econom. Management Strategy* 6(4) 877–897.
- Chen, Y. 2006. Marketing innovation. *J. Econom. Management Strategy* 15(1) 101–123.
- Chen, Y., G. Iyer. 2002. Consumer addressability and customized pricing. *Marketing Sci.* 22(2) 197–208.
- Chen, Y., Z. J. Zhang. 2008. Price-for-information effect and benefit of behavior-based targeted pricing. Working paper, Wharton School of Business, Philadelphia.
- Chen, Y., Z. J. Zhang. 2009. Dynamic targeted pricing with strategic consumers. *Internat. J. Indust. Organ.* 27(1) 43–50.
- Chen, Y., C. Narasimhan, Z. J. Zhang. 2001. Individual marketing with imperfect targetability. *Marketing Sci.* 20(1) 22–43.
- Coase, R. H. 1972. Durability and monopoly. *J. Law Econom.* 15(1) 143–149.
- Esteves, R. B. 2010. A survey on the economics of behaviour-based price discrimination. Working paper, University of Minho, Minho, Portugal.
- Freixas, X., R. Guesnerie, J. Tirole. 1985. Planning under incomplete information and the ratchet effect. *Rev. Econom. Stud.* 52(2) 173–191.
- Fudenberg, D., J. Tirole. 1998. Upgrades, trade-ins, and buybacks. *RAND J. Econom.* 29(2) 238–258.
- Fudenberg, D., J. Tirole. 2000. Customer poaching and brand switching. *RAND J. Econom.* 31(4) 634–657.
- Fudenberg, D., J. M. Villas-Boas. 2006. Behavior-based price discrimination and customer recognition. T. Hendershott, ed. *Handbook on Economics and Information Systems*, Vol. 1. North-Holland, Amsterdam, 377–436.
- Gerstner, E., D. Holthausen. 1986. Profitable pricing when market segments overlap. *Marketing Sci.* 5(1) 55–69.
- Goldfarb, A., C. Tucker. 2011. Privacy regulation and online advertising. *Management Sci.* 57(1) 57–71.
- Guy, S. 2011. Groupon offers refunds after complaints about flower deal. *Chicago Sun-Times* (February 17) <http://www.suntimes.com/business/3822291-420/more-groupon-woes.html>.
- Hann, I.-H., K.-L. Hui, S.-Y. T. Lee, I. P. L. Png. 2007. Overcoming information privacy concerns: An information processing theory approach. *J. Management Inform. Systems* 24(2) 13–42.
- Hart, O., J. Tirole. 1988. Contract renegotiation and Coasian dynamics. *Rev. Econom. Stud.* 55(4) 509–540.
- Hermalin, B., M. Katz. 2006. Privacy, property rights and efficiency: The economics of privacy as secrecy. *Quant. Marketing Econom.* 4(3) 209–239.
- Laffont, J.-J., J. Tirole. 1988. The dynamics of incentive contracts. *Econometrica* 56(5) 1153–1175.
- Larsen, K. R., Z. A. Voronovich, eds. 2004. *Technology and Privacy in the New Millennium*. Ethica Publishing, University of Colorado, Boulder.
- Lohr, S. 2010. How privacy vanishes online. *New York Times* (March 16) <http://www.nytimes.com/2010/03/17/technology/17privacy.html>.

- Low, S., N. F. Maxemchuk, S. Paul. 1994. Anonymous credit cards. *Proc. 2nd ACM Conf. Comput. Comm. Security*, ACM, New York, 108–117.
- Mialon, H. M., S. H. Mialon. 2008. The effects of the fourth amendment: An economic analysis. *J. Law, Econom., Organ.* **24**(1) 22–44.
- Narasimhan, C. 1984. A price discrimination theory of coupons. *Marketing Sci.* **3**(2) 128–147.
- Narayanan, A., V. Shmatikov. 2008. Robust de-anonymization of large sparse data sets (how to break anonymity of the Netflix prize data set). *IEEE Sympos. Security Privacy*, IEEE, New York, 111–125.
- Odlyzko, A. 2003. Privacy, economics, and price discrimination on the Internet. *ICEC 2003 5th Internat. Conf. Electronic Commerce*, ACM, New York, 355–366.
- Pancras, J., K. Sudhir. 2007. Optimal marketing strategies for a customer data intermediary. *J. Marketing Res.* **44**(4) 560–578.
- Ramasastri, A. 2005. Web sites change prices based on customers' habits. *CNN.com* (June 24) <http://www.cnn.com/2005/LAW/06/24/ramasastri.website.prices/>.
- Rossi, P., R. McCulloch, G. Allenby. 1996. The value of purchase history data in target marketing. *Marketing Sci.* **15**(4) 321–340.
- Salant, S. W. 1989. When is inducing self-selection suboptimal for a monopolist? *Quart. J. Econom.* **104**(2) 391–397.
- Schmidt, K. M. 1993. Commitment through incomplete information in a simple repeated bargaining game. *J. Econom. Theory* **60**(1) 114–139.
- Shaffer, G., Z. J. Zhang. 1995. Competitive coupon targeting. *Marketing Sci.* **14**(4) 395–417.
- Shaffer, G., Z. J. Zhang. 2000. Pay to switch or pay to stay: Preference-based price discrimination in markets with switching costs. *J. Econom. Management Strategy* **9**(3) 397–424.
- Shaffer, G., Z. J. Zhang. 2002. Competitive one-to-one promotions. *Management Sci.* **48**(9) 1143–1160.
- Stokey, N. 1979. Intertemporal price discrimination. *Quart. J. Econom.* **93**(3) 355–371.
- Stole, L. A. 2007. Price discrimination and competition. M. Armstrong, R. H. Porter, eds. *Handbook of Industrial Organization*. North-Holland, Amsterdam, 2221–2300.
- Taylor, C. 2003. Supplier surfing: Competition and consumer behavior in subscription markets. *RAND J. Econom.* **34**(2) 223–246.
- Taylor, C. 2004. Consumer privacy and the market for customer information. *RAND J. Econom.* **35**(4) 631–650.
- Villas-Boas, J. M. 1999. Dynamic competition with customer recognition. *RAND J. Econom.* **30**(4) 604–631.
- Villas-Boas, J. M. 2004. Price cycles in markets with customer recognition. *RAND J. Econom.* **35**(3) 486–501.
- Wathieu, L. 2006. Marketing and privacy concerns. Working paper, Harvard Business School, Boston.