Surfing Incognito: Welfare Effects of Anonymous Shopping

Johan N. M. Lagerlöf*

University of Copenhagen

May 27, 2019

Abstract

This paper studies consumers' incentives to hide their purchase histories when the seller's prices depend on previous behavior. Through distinct channels, hiding both hinders and facilitates trade. Indeed, the social optimum involves hiding to some extent, yet not fully. Two opposing effects determine whether a consumer hides too much or too little: the first-period social gains are only partially internalized, and there is a private (socially irrelevant) second-period gain due to price differences. If time discounting is small, the second effect dominates and there is socially excessive hiding. This result is reversed if discounting is large.

Keywords: behavior-based price discrimination, dynamic pricing, consumer protection, customer recognition, privacy

JEL classification: D42, D80, L12, L40

*Address: Department of Economics, University of Copenhagen, Øster Farimagsgade 5, Building 26, DK-1353 Copenhagen K, Denmark; email: johan.lagerlof@econ.ku.dk. I thank Antonio D'Agata, Mogens Fosgerau, Christoph Schottmüller, Christian Schultz, Peter N. Sørensen, Curtis Taylor, and a seminar audience at the University of Copenhagen for helpful discussions. Financial support from *Det Frie Forskningsråd*, grant number DFF—1327-00079, is gratefully acknowledged. Future versions of the paper will be available at www.johanlagerlof.com, where also the Supplementary Material can be found.

1 Introduction

In the last couple of decades it has become increasingly common that consumers make retail purchases on the internet. While online shopping often is convenient for consumers, it also makes it relatively easy for sellers to keep track of individual customers' purchasing decisions and thereby learn about their willingness to pay for the good. Using that knowledge, the sellers can charge personalized prices that leave certain consumers with a smaller surplus than otherwise. In a Washington Post article, Lowrey (2010) vividly describes this phenomenon and how it upsets consumers:

> *Retailers read the cookies kept on your browser or glean information from your past purchase history when you are logged into a site. That gives them a sense of what you search for and buy, how much you paid for it, and whether you might be willing and able to spend more. They alter their prices or offers accordingly. Consumers [...] tend to go apoplectic. But the practice is perfectly legal, and increasingly common—pervasive, even, for some products.*

In the economics literature, the seller practice described in the quotation has been referred to as dynamic pricing, history-based pricing, or behavior-based price discrimination (BBPD). A seller who practices BBPD does not necessarily have to do this in the blunt manner suggested in the quotation, directly presenting different consumers with different price tags. Often more subtle approaches are available. For example, a seller can distribute a specific discount coupon only to certain consumers, thereby effectively offering a personalized price to them. Introductory offers that entitle new customers to pay a lower price than returning customers pay are another example of BBPD.

BBPD has been studied theoretically both in monopoly and oligopoly settings. One insight from this literature is that a firm's opportunity to practice BBPD is not necessarily harmful to consumer welfare, as price discrimination tends to lead to more trade than otherwise and this can benefit also the consumers.[1] However, in specific situations and if she can, a consumer clearly has an individual incentive to hide her purchase history from the seller. Indeed, if being a returning customer is interpreted as being a high-valuation customer, then you are better off pretending to be a new customer. In practice, there are several possibilities for consumers to hide their purchase histories by using various anonymizing technologies: they can refrain from joining loyalty programs or set their browsers to reject cookies; they can choose to end a newspaper subscription and then start a new one, instead of renewing the old subscription; or they "can use a variety of credit cards or more exotic anonymous payment technologies to make purchases anonymous or difficult to trace" (Acquisti and Varian, 2005, p. 367). These kinds of defense measures might come at a cost (a financial expense or nuisance) but they are certainly available.

The present paper is an attempt, within an equilibrium framework where a firm engages in BBPD, to study the incentives of consumers to hide their purchase histories with the help of

[1] Another reason why the consumers can gain from the firm's opportunity to practice BBPD is that, due to the firm's inability to precommit to second-period prices, it is forced to lower its first-period price in order to sell anything then (a "Coase-conjecture effect").

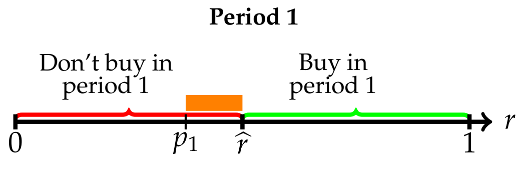

Figure 1: Equilibrium prices and cutoff value in a standard BBPD model. From a total surplus perspective, a consumer's incentive to hide is too weak in period 1 and too strong in period 2.

anonymizing technologies. Among other things I ask whether, at the equilibrium and given other imperfections in the market, consumers use these technologies too little or too much. That is, from a social welfare perspective, does the market tend to generate too strong or too weak incentives for consumers to hide their purchase histories?

To help us think about whether (or under what circumstances) the incentives are likely to be too strong or too weak, refer to Figure 1. This figure summarizes the results of a standard two-period monopoly BBPD model with a continuum of consumers, each having a valuation r (the same across the two periods and drawn from a uniform distribution on the unit interval). The firm has no production costs; hence, the social optimum involves all consumers purchasing the good. In period 1, the firm sets a single price p_1, knowing only the distribution of valuations; the consumers then choose whether to purchase. In period 2, the firm can distinguish between buying and non-buying consumers and thus set two prices, p_2^L and p_2^H. At the equilibrium, consumers with valuations above a cutoff $\widehat{r}$ purchase in period 1. Moreover, the period 2 equilibrium price for the returning customers is higher than that of the new customers ($p_2^H > p_2^L$). Finally, the period 1 equilibrium price is strictly lower than the cutoff ($p_1 < \widehat{r}$); this is because a consumer requires a positive first-period surplus in order to be willing to purchase in period 1 (for $p_2^H > p_2^L$).

Given the above framework, consider now the possibility that a consumer has the opportunity to hide her purchase decision in period 1 (thus being eligible to purchase at the price p_2^L in period 2, also if being a returning customer). Will her benefit from hiding be smaller or greater than society's benefit (measured as total surplus)? By inspecting Figure 1, we can identify three effects (or externalities) that determine whether the typical consumer has too weak or too strong incentives, and these effects point in different directions:

1. For consumers with $r \in (p_1, \widehat{r})$, hiding one's purchase history will lead to *more trade*, as the consumer's acquired status as a "new" customer makes it worthwhile for her to buy a good, in a situation where she would not have bought it if lacking that status. However, when evaluating the benefit from this trade, the consumer considers only the net surplus that is generated (valuation less price, $r - p_1$). In contrast, society cares about the full surplus r (as the price is a pure transfer from the consumer to the firm). This effect thus suggests that the consumer has a too weak incentive to hide her purchase history.

2. For consumers with $r \in (\hat{r}, 1]$, hiding one's purchase history enables the consumer, in period 2, to *buy the good at a lower price* than she otherwise would have paid, thus saving the amount $p_2^H - p_2^L$. The consumer cares about this benefit, while society does not (as the price is a pure transfer). Therefore, this effect suggests that the consumer has a too strong incentive to hide her purchase history.

3. Finally, an atomistic consumer's choice to hide her purchase history will not have any *impact on the equilibrium prices*. However, if many consumers in the market make that choice, the prices will indeed readjust to a new equilibrium. The direction of this effect is less clear than the direction of effects 1 and 2 above. But a plausible scenario would be that more hiding leads to a smaller difference between the second-period prices: the firm is less able price discriminate. This in turn might lead to less trade. If so, also this effect suggests that the consumer has a too strong incentive to hide her purchase history.

The first effect, which creates an incentive to hide too little, matters for the consumer in the period in which she makes her hiding decision. The second and third effects, which create an incentive to hide too much, matter only in the following period. This suggests that, if the consumer cares sufficiently much about the future, then the second and third effects might dominate and the consumer thus invests too much in anonymizing technologies. In the model that I set up and study, I show that this is indeed the case. I also show that, if the consumer instead is sufficiently myopic (i.e., she assigns a small weight to her second-period payoff), then we can reverse this result: the consumer invests too little in anonymizing technologies.

The formal framework that I develop is close to the standard BBPD model described above, except that I add the opportunity for the consumers to hide their purchase histories. Thus, a monopoly firm produces and sells a nondurable good in each of two time periods. Each consumer's valuation for the good is the same across the two periods and drawn from a uniform distribution. The valuation is initially the consumer's private information; the firm knows only the distribution of valuations in the market. However, by observing the first-period consumption choice, the firm can make a noisy inference about a consumer's valuation and then use this information when choosing the second-period price. In particular, a consumer's choice to purchase the good in the first period suggests that her valuation is relatively high, which creates an incentive for the firm to raise that consumer's second-period price. To protect herself from this, the consumer has the opportunity to hide her first-period purchase. This is modeled by letting the consumer choose a probability (i.e., any real number between zero and one) with which the information that she purchased the good is not available to the firm in the second period. Put differently, if the consumer chooses a particular hiding probability, then with that probability she will after her purchase look like a consumer who did not buy; with the complementary probability, she will indeed be identified by the firm as a consumer who purchased in period 1 and thus have to pay a higher price in period 2 (as in the standard setting). The choice of the hiding probability is made at an ex ante stage, before the consumer has learned her own valuation. This model feature captures the idea that the consumer adopts a long-term approach for dealing with privacy issues, for example, by choosing a setting on her computer that she sticks with through a large number of browsing sessions. I discuss this assumption at greater length in connection with the model description in Section 2. In the concluding section, I also briefly discuss the technical consequences of making the alternative assumption that the consumer's hiding decision is made *after* she has learned her valuation.

In Section 3, I begin the analysis by studying a version of the model in which the hiding probability is given exogenously and the same for all consumers. I first solve for the equilibrium (for given parameter values, this is unique). As in the standard model without the opportunity to hide one's purchase history, the equilibrium is characterized by a cutoff value of a consumer's valuation—above which she purchases in the first period, and below which she does not.[2] Given this model with an exogenous hiding probability, I take an initial look at social welfare. In particular, I show that if the firms and the consumers use the same discount factor (which is required for total surplus to be well defined), the welfare-maximizing value of the hiding probability is strictly interior. That is, if a total-surplus maximizing social planner could choose the hiding probability, the consumers would hide their first-period purchases to at least some extent, yet not fully. The reason why a strictly positive degree of hiding is socially optimal is that hiding generates gains from trade in the first period. Moreover, the social cost of hiding in terms of hindering second-period price discrimination is small, provided that the degree of hiding is sufficiently small.

In Section 4, I endogenize the hiding probability. I confine the analysis to equilibria where all consumers choose the same probability. I characterize and show existence of such an equilibrium. I then turn to the question whether the equilibrium value of the hiding probability is larger or smaller than the value that maximizes total surplus in the market. From here on in the analysis, I assume that the firm and consumers are equally patient (which means that total surplus is well-defined). I first show that for the case where the common discount factor equals one—so the weights assigned to the first- and second-period payoffs are the same—the equilibrium involves too much hiding. The intuition for this result can be understood in terms of the discussion in the beginning of this introduction. By hiding her purchase history, a consumer can in the second period buy the good at a lower price than otherwise; this individual gain does not enter total surplus, as it is a pure transfer (effect 2). Moreover, the act of hiding makes it harder for the seller to practice (trade-enhancing) price discrimination in the second period (effect 3). Both those effects suggest that the consumer hides too much. Another effect of hiding one's purchase history is that it makes it possible to exploit first-period gains from trade to a greater extent; this suggests that the consumer hides too little, as she does not internalize the full benefit of the extra trade (effect 1). When the discount factor equals one, the weight on the second-period payoff is large enough for effects 2 and 3 to dominate.

A more general analysis, for the case where the discount factor is strictly below unity, is harder. Nevertheless, with the help of a numerical analysis, I can first show that the result discussed above—that the consumer hides too much—holds also for a range of discount factor

[2] Interestingly, for a part of the parameter space, the equilibrium has the feature that some consumers purchase in the first period but in the second period they do not (a result which I believe is novel in the literature). However, for that to occur, the firm must use a smaller discount factor than the consumers; for the case where the firm's discount factor is equal to or larger than that of the consumers, any consumer who purchases in the first period also buys in the second (which is consistent with the results of the standard model).

values below one. Second, I provide examples where the discount factor is so low that the result above is reversed and, instead, the consumer hides too little. Intuitively, if the discount factor is sufficiently low, what matters for welfare is the first period and then the consumer does not internalize all the gains from trade that her hiding enables.

1.1 Literature Review

The present paper is, of course, closely related to the literature on behavior-based price discrimination. Examples of works in this literature include Chen (1997), who studies a two-period duopoly model with switching costs. At the equilibrium, each firm offers a relatively low introductory price, meaning that it effectively pays customers to switch from the rival. Fudenberg and Tirole (2000) set up a two-period Hotelling duopoly model where each firm, in the second period, can distinguish between its own returning customers and those of the rival and thus charge the two consumer groups different prices. Chen and Zhang (2009) study a problem similar to that of Fudenberg and Tirole but assume that the consumers have access to a richer set of instruments when they strategically try to avoid high prices; in particular, a consumer can, besides buying from the rival, instead choose to wait with purchasing. Villas-Boas (2004) develops a model of BBPD with overlapping generations of consumers (living for two periods) and an infinitely lived monopoly firm. He shows that the equilibrium involves price cycles; in particular, the price to the new customers alternates between a relatively low and a relatively high level. Gehrig, Shy, and Stenbacka (2011) investigate under what circumstances an incumbent monopoly firm's opportunity to do BBPD can act as an entry barrier. They formulate a simple two-period model where a potential rival, which does not have access to information about the purchase history of the incumbent's customers, chooses whether or not to enter the market. Esteves (2009a) studies the role of informative advertising in a duopoly market where firms can practice BBPD (although with naive consumers). For surveys of the literature on BBPD, see Armstrong (2006), Fudenberg and Villas-Boas (2006), and Esteves (2009b).

The work that is closest to the present one is Conitzer, Taylor, and Wagman (2012), who also model consumers' opportunity to hide their purchase histories in a BBPD context. Their model are in some regards similar to mine, although they assume a binary hiding choice. Most importantly, however, they do not address the welfare questions that are central in my analysis. Indeed, neither one of my two main results[3] appear in Conitzer et al.[4] The authors do, however, carry out interesting comparative statics with respect to the parameter that represents the cost of hiding, and they find that facilitating privacy (by lowering this cost) can either increase or decrease total surplus. Also Acquisti and Varian (2005) offer a very useful discus-

[3]Recall that those are: the social optimum involves hiding to some extent although not fully; and both socially (i) excessive and (ii) insufficient hiding are possible, and whether (i) or (ii) obtains depends on the discount factor.

[4]The closest Conitzer et al. get to any one of my two main results is when they in Section 4.6 (p. 288) compare consumer welfare under zero hiding cost with consumer welfare under a positive hiding cost. That comparison effectively shows that there can be excessive hiding, given a consumer surplus standard (Figure 4 in their paper shows that the same can hold also with a total surplus standard); the authors in that context also mention the negative externality that the consumers impose on each other. However, the social benefit of increased trade, due to hiding, is not mentioned and there are no results in their paper that suggest that there can be insufficient hiding. Nor can one find results, or any discussion, about the role of time preferences in this context.

sion of consumers' opportunity to hide their purchase histories. In Section 7 of their paper, the authors briefly consider a model where the firm cannot commit to second-period prices and where consumers have the option of hiding. Acquisti and Varian note that this "imposes a cost on the seller, in that it will not be able to implement a price-conditioning solution" (p. 378). However, the modeling choices (e.g., binary valuations and binary hiding choice) and the focus of the analysis mean that these authors fail to identify the tradeoff in the present paper or address the efficiency questions that are studied here.

On a more general level, the present paper adds to the literature on efficiency and the regulation of privacy, although in the specific setting of behavioral-based price discrimination in a monopoly market. University of Chicago scholars, notably Stigler (1980) and Posner (1981), have argued that privacy is harmful to efficiency. The reason is that privacy can, by hindering information flows, prevent gains from trade from being realized. For example, privacy can lead to informational asymmetries or discourage productive investments. The present paper contributes to this discussion by showing that, in an environment with BBPD, some degree of privacy can in fact generate more gains from trade (as it enables more consumers to purchase in the first period); moreover, I show that—because of this effect—privacy can indeed be under-provided from a social welfare point of view. Other more recent studies of the effects of privacy on allocative efficiency include Hermalin and Katz (2006), Taylor (2004), and Calzolari and Pavan (2006). See also the surveys of the economics of privacy provided by Hui and Png (2006) and Acquisti, Taylor, and Wagman (2016).

Finally, Johnson (2013) develops a model of a market where firms do (informative) advertising and consumers have the opportunity to block the advertisements. The feature of his model that consumers can choose the extent to which they protect their privacy is reminiscent of my framework. Moreover, like me, Johnson investigates under what circumstances the consumers over- and underinvest in privacy protection. However, the tradeoff that the consumers in Johnson's setting face is different from the one of the present model. In particular, protecting one's privacy cannot lead to gains from trade in his setting, whereas in mine it can.

2 Model

The following model of behavior-based price discrimination builds on the framework used by Armstrong (2006, Section 2) and Fudenberg and Villas-Boas (2006, Section 2.1), although I extend it by allowing each consumer to take a costly action that hides her purchase history from the seller.

There are two time periods, 1 and 2. In each period, a profit-maximizing and risk-neutral monopoly firm produces and sells a nondurable good. The production technology is characterized by constant returns to scale and the per-unit cost is normalized to zero. When making decisions in period 1, the firm discounts second-period profits with the discount factor $\beta \in [0, 1]$. The consumers form a continuum and differ from each other in terms of r, the gross utility from consuming one unit of the good. In particular, the per-period consumption utility, if p is the price of the good, equals $r - p$ if buying and zero if not buying. The r values are independent and uniformly distributed on the interval $[0, 1]$, and the total mass of consumers

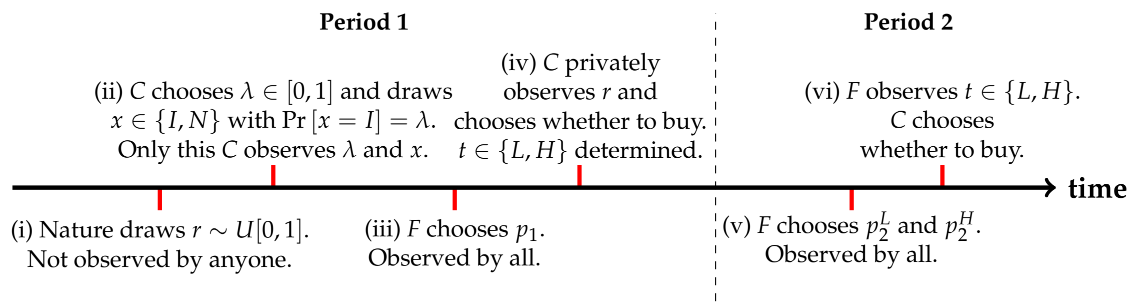

Figure 2: Sequence of events. Abbreviations: C = consumer, F = firm.

equals one. A given consumer's valuation r is the same across the two periods. Moreover, while the consumer knows her own r when making her first-period purchase decision, the valuation of an individual consumer cannot be observed by the firm. However, unless the consumer uses an anonymizing technology, the firm can keep track of individual consumers' purchase decisions. When making first-period decisions, consumers use the discount factor $\delta \in (0, 1]$.

How does the anonymizing technology work? For any given consumer, let $t \in \{L, H\}$ be an indicator variable that is determined as follows. If the consumer does not buy the good in period 1, then $t = L$ for sure. If the consumer indeed buys the good in period 1, then $t = L$ with probability λ and $t = H$ probability $1 - \lambda$, for some individual-specific $\lambda \in [0, 1]$. When interacting with the consumers in period 2, the firm can observe each individual consumer's value of t; it has no other information about whether that consumer actually bought in period 1 or not. That is, we can think of t as a marker that is initially attached to any consumer who purchases the good in period 1, but which is then removed with probability λ. In Section 3, I will assume that the "hiding" or "incognito" probability $\lambda \in [0, 1]$ is given exogenously (and is the same for all consumers). In Section 4, I will extend the analysis by letting λ be chosen by the consumer. There is a cost associated with choosing any $\lambda > 0$, which is denoted by $C(\lambda)$ and is subtracted from the consumption utility. The cost function is twice continuously differentiable and it satisfies $C(0) = C'(0) = 0$, $C'(\lambda) > 0$, and $C''(\lambda) > 0$ for all $\lambda \in (0, 1]$.

The informational assumptions stated above imply that the firm's second-period price can be made contingent on $t \in \{L, H\}$. The consumers understand that the firm may charge different second-period prices depending on if the consumer purchased in the first period or not and on the realization of t. They take this into account when deciding whether to purchase in period 1.

For the full model with endogenous λ, the timing of events is as follows—see also Figure 2. (i) Nature draws each consumer's valuation r. The realization of r is not, at this stage, observed by the firm or by any consumer. (ii) Each consumer chooses her own individual $\lambda \in [0, 1]$, at the cost $C(\lambda)$. Nature then with probability λ assigns an incognito status ($x = I$), and with probability $1 - \lambda$ a non-incognito status ($x = N$), to her. The consumer herself observes the realization of x. However, neither the choice of λ nor the realization of x is observed by the firm or by the other consumers. (iii) The firm chooses its first-period price $p_1 \geq 0$, which is observed by the consumers. (iv) Each consumer privately learns her own valuation r and then decides whether to make a first-period purchase or not. If she does buy and if $x = N$, then her indicator variable t equals H; otherwise, $t = L$. (v) We now move into period 2 and the firm chooses two second-period prices: $p_2^L \geq 0$ and $p_2^H \geq 0$. The price p_2^t must be paid by those consumers with $t \in \{L, H\}$. (vi) Consumers observe p_2^L and p_2^H and then choose whether to buy or not.

The above timing implies that the incognito probability λ is chosen by the consumer at an ex ante stage, before she has learned about her own valuation. This model feature captures the idea that the consumer adopts a long-term approach for dealing with certain kinds of privacy issues. For example, the consumer might today choose a particular setting on her computer and then, to save on hassle costs, stick with this for a long time and throughout many browsing and purchase situations. Alternatively, the choice of λ could represent the adoption of a simple behavioral rule or heuristic that the consumer uses in a wide range of situations, and which is updated only occasionally. If we nevertheless believe that consumers adjust their behavior in all new situations they face, we can think of the timing of the game as an analytical shortcut. That is, the timing is a simple way of capturing the broad tradeoffs we want to study, and there is no reason to believe that the simplification leads to systematically stronger or weaker incentives for the consumers to hide their purchase histories.[5]

I will impose some restrictions on the parameter space. The assumption below ensures that a pure strategy equilibrium of the game with an exogenous λ exists.

Assumption 1. *The parameters β, λ, and δ satisfy $\beta \geq \mathcal{B}(\lambda, \delta)$, where*

$$\mathcal{B}(\lambda,\delta) \stackrel{def}{=} \begin{cases} -\infty & \text{for } \lambda = 0 \\ \frac{\delta(1+\lambda)\left(1+\lambda+\delta\lambda^{\frac{3}{2}}\right)-2}{\sqrt{\lambda}\left(1+\sqrt{\lambda}\right)^2} & \text{for } \lambda \in (0,1]\,. \end{cases} \tag{1}$$

Assumption 1 requires that the firm does not discount too heavily. The assumption is necessarily satisfied if the firm is at least as patient as the consumers ($\beta \geq \delta$). But it holds also for $\beta < \delta$, if the difference between these discount factors is not too large.[6]

The solution concept that I employ is that of perfect Bayesian equilibrium. All players must make optimal choices at all information sets given their beliefs, and the beliefs are formed with the help of Bayes' rule when that is defined. This solution concept also requires that all consumers with access to the same information have the same beliefs.

[5]Suppose we instead assumed that each consumer knows her own r when choosing λ. Under such an assumption, only some consumers—namely, those who expect to purchase in the first period—would have an incentive to acquire an incognito status. However, we should expect that incentive to be stronger than the incentive of the ex ante consumer in the present framework. In aggregate and taking those two opposing effects (i.e., the number of incentivized consumers versus the strength of their incentives) into account, it is not clear whether the alternative setting would lead to more or less investment in λ.

[6]In particular, if $\delta \leq \sqrt{2} - 1 \approx .41$, then Assumption 1 is satisfied for all $(\beta, \lambda) \in [0,1]^2$.

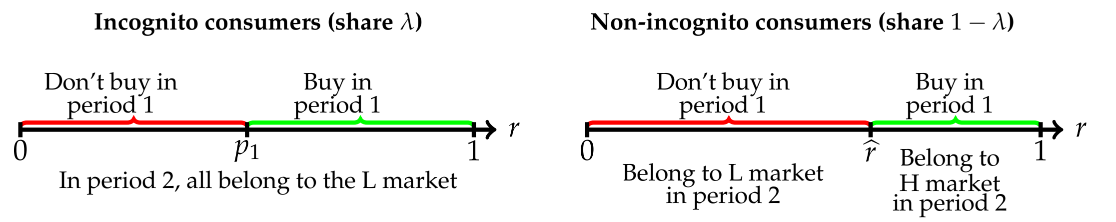

Figure 3: Period 1 behavior of incognito and non-incognito consumers.

3 Exogenous Fraction of Incognito Consumers

Before solving the full model as described in Section 2, it will be useful to study a setting where the incognito probability λ is exogenous and satisfies $\lambda \in [0,1)$.[7] This version of the model yields interesting insights in itself and it will also help us to later, in Section 4, solve the full model where the incognito probability is endogenous.

In the second period there are effectively two separate markets: a "low-valuation" market with consumers who pay p_2^L (as they either did not purchase in the first period or they did but have an incognito status) and a "high-valuation"market with consumers who pay p_2^H. Any equilibrium must be characterized by an endogenous threshold $\widehat{r} \in (0,1)$ with the property that, in the first period, a consumer without an incognito status (so with $x = N$) buys if $r > \widehat{r}$ and does not buy if $r < \widehat{r}$. In particular, any such consumer with valuation r has a (weak) incentive to buy in period 1 if, and only if,

$$r - p_1 + \delta \max\left\{0, r - p_2^H\right\} \geq \delta \max\left\{0, r - p_2^L\right\}. \tag{2}$$

The left-hand side of (2) is the consumer's utility if buying in period 1 (thus having to pay the second-period price p_2^H). The right-hand side is her utility if not buying in period 1 (which means that the second-period price is p_2^L). When solving for the equilibria of the model, we can exploit the fact that, for a consumer with $r = \widehat{r}$, inequality (2) must hold with equality.

The incognito consumers (i.e., those with $x = I$) always pay the second-period price p_2^L, regardless of whether they bought in the first period or not. They therefore optimally decide to purchase the good in the first period if, and only if, $r \geq p_1$. The first-period behavior of the incognito and the non-incognito consumers is summarized in Figure 3.

3.1 Equilibrium Behavior in Period 2

Consider, in turn, the firm's second-period profit-maximization problem in the high- and the low-valuation markets. Let q_2^H denote the demand that the firm faces in the high-valuation market. All consumers in this market lack an incognito status and they have valuations that

[7]The case $\lambda = 1$, which is excluded here, means that the firm cannot keep track of any consumers' first-period decision whether to purchase the good. This version of the model is straightforward to solve. In the second period there is effectively only one market, and at the equilibrium we have $p_1 = \widehat{r} = p_2^L = 1/2$.

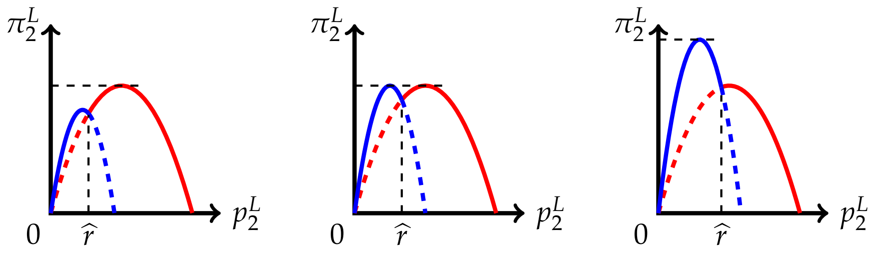

(a) $\widehat{r} < \frac{\sqrt{\lambda}}{1+\sqrt{\lambda}}$ **(b)** $\widehat{r} = \frac{\sqrt{\lambda}}{1+\sqrt{\lambda}}$ **(c)** $\widehat{r} > \frac{\sqrt{\lambda}}{1+\sqrt{\lambda}}$

Figure 4: Profit maximization in the second-period low-valuation market.

are uniformly distributed on $[\widehat{r}, 1]$; cf. Figure 3. We therefore have

$$q_2^H = \begin{cases} (1-\lambda)\left(1-p_2^H\right) & \text{if } p_2^H \in [\widehat{r}, 1] \\ (1-\lambda)\left(1-\widehat{r}\right) & \text{if } p_2^H \in [0, \widehat{r}]\,. \end{cases} \tag{3}$$

It is straightforward to see that the profits in the high-valuation market, $\pi_2^H = p_2^H q_2^H$, are maximized at $p_2^H = \max\left\{\frac{1}{2}, \widehat{r}\right\}$.

Next, consider the demand that the firm faces in the low-valuation market, denoted by q_2^L. This market consists of all consumers with an incognito status, uniformly distributed on $[0, 1]$, and of the consumers without an incognito status who did not buy in period 1, uniformly distributed on $[0, \widehat{r}]$; cf. again Figure 3. We thus get that

$$q_2^L = \begin{cases} \widehat{r} - p_2^L + \lambda\left(1-\widehat{r}\right) & \text{if } p_2^L \in [0, \widehat{r}] \\ \lambda\left(1-p_2^L\right) & \text{if } p_2^L \in [\widehat{r}, 1]\,. \end{cases} \tag{4}$$

That is, for relatively low values of p_2^L, there are both incognito and non-incognito consumers who find it worthwhile to purchase the good, while for higher values of p_2^L only incognito consumers do.

The firm's profits in the low-valuation market equal $\pi_2^L = q_2^L p_2^L$. This profit function is continuous in p_2^L (although with a kink at $p_2^L = \widehat{r}$). It is not, however, in general quasiconcave. Indeed, from the three panels of Figure 4 it is clear that the profit function may have two local optima: one where the price p_2^L is relatively low, which means that the firm sells to both incognito and non-incognito consumers; and one local optimum where the price is relatively high, which means that only (some) incognito consumer purchase the good. Which one of the local optima that is the global one (or if they both are global optima) depends on the relative magnitude of λ and $\widehat{r}$.[8]

[8]Lemma 1 below implies, in particular, that the firm's optimal price p_2^L (i.e., its best reply) makes a jump as $\widehat{r}$ increases. This is the technical reason why, if we had not imposed Assumption 1, an equilibrium in pure strategies would fail to exist for some parameter values.

Lemma 1. *The price p_2^L that maximizes the profits $\pi_2^L = q_2^L p_2^L$ is given by*

$$p_2^L = \begin{cases} \frac{\lambda+(1-\lambda)\widehat{r}}{2} & \text{if } \widehat{r} \in \left[\frac{\sqrt{\lambda}}{1+\sqrt{\lambda}}, 1\right] \\ \frac{1}{2} & \text{if } \widehat{r} \in \left[0, \frac{\sqrt{\lambda}}{1+\sqrt{\lambda}}\right] \end{cases}. \tag{5}$$

Proof. The proof of Lemma 1, as well as the results stated in the remainder of the paper, can be found in the Appendix.

3.2 Equilibrium Behavior in Period 1

In order to derive the optimal behavior in period 1 and to identify all possible equilibria of the game, we need to investigate three cases:

$$\text{(i) } \widehat{r} \geq \frac{1}{2}; \qquad \text{(ii) } \widehat{r} \in \left(\frac{\sqrt{\lambda}}{1+\sqrt{\lambda}}, \frac{1}{2}\right); \qquad \text{(iii) } \widehat{r} \leq \frac{\sqrt{\lambda}}{1+\sqrt{\lambda}}.$$

Refer to an equilibrium that arises under case (i) as a type (i) equilibrium, and analogously for cases (ii) and (iii). Below I will solve for type (i) and (ii) equilibria and specify under what conditions they exist. In the Appendix, I also show that a type (iii) equilibrium never exists (see the proof of Proposition 1). Overall, for given parameter values, there is a unique pure strategy equilibrium of the model.

3.2.1 Solving for a Type (i) Equilibrium

In a type (i) equilibrium, the second-period prices equal $p_2^H = \widehat{r}$ and

$$p_2^L = \frac{\lambda + (1-\lambda)\widehat{r}}{2} \tag{6}$$

(see subsection 3.1). We also know that the threshold $\widehat{r}$ must satisfy inequality (2) with equality, when (2) is evaluated at those two second-period prices:

$$\widehat{r} - p_1 + \delta(\widehat{r} - \widehat{r}) = \delta\left[\widehat{r} - \frac{\lambda + (1-\lambda)\widehat{r}}{2}\right] \Leftrightarrow p_1 = \widehat{r} - \frac{\delta\left[(1+\lambda)\widehat{r} - \lambda\right]}{2}. \tag{7}$$

Equation (7) gives us a relationship between the two endogenous variables $\widehat{r}$ and p_1. Anticipating this relationship and the optimal second-period prices, the firm chooses its first-period price p_1 so as to maximize the following overall profits:

$$\Pi = \pi_1 + \beta(\pi_2^L + \pi_2^H) = q_1 p_1 + \beta\frac{\left[\lambda + (1-\lambda)\widehat{r}\right]^2}{4} + \beta(1-\lambda)(1-\widehat{r})\widehat{r}, \tag{8}$$

where

$$q_1 = (1-\lambda)(1-\widehat{r}) + \lambda(1-p_1) \tag{9}$$

is the first-period demand. Rather than maximizing (8) with respect to p_1 (subject to (7) and (9)), we can equivalently maximize it with respect to $\widehat{r}$ (subject to (7) and (9)). Let $\widehat{\Pi}$ denote the reduced-form profit function that we obtain by eliminating p_1 from (8) with the help of (7) and (9). The function $\widehat{\Pi}$ is strictly concave in $\widehat{r}$. Therefore, the optimal $\widehat{r}$ satisfies the first-order condition $\partial\widehat{\Pi}/\partial\widehat{r} = 0$ or, equivalently,

$$\widehat{r} = \frac{2 - \delta(1+\lambda)\left(1+\lambda-\delta\lambda^2\right) + \beta(1-\lambda)(2+\lambda)}{4 - \delta(1+\lambda)^2(2-\delta\lambda) + \beta(1-\lambda)(3+\lambda)}. \tag{10}$$

For the above value of $\widehat{r}$ to indeed be part of an equilibrium, we must have $\widehat{r} \geq \frac{1}{2}$ and $\widehat{r} < 1$. The latter inequality can be shown to always hold, while $\widehat{r} \geq \frac{1}{2}$ is equivalent to $\beta \geq \delta^2\lambda$. Thus, if the firm is sufficiently patient ($\beta \geq \delta^2\lambda$), there is an equilibrium where $\widehat{r}$ is given by (10). The equilibrium values of the three prices are in turn obtained from (6), (7), and $p_2^H = \widehat{r}$.

3.2.2 Solving for a Type (ii) Equilibrium

In a type (ii) equilibrium, the second-period prices are given by eq. (6) and $p_2^H = \frac{1}{2}$ (see subsection 3.1). In particular, consumers with a valuation $r \in \left(\widehat{r}, \frac{1}{2}\right)$ and without an incognito status buy in the first period but do not buy in period 2. Thus, using (2), we have that the threshold $\widehat{r}$ satisfies $\widehat{r} - p_1 + 0 = \delta\left(\widehat{r} - p_2^L\right)$, which again yields (7). The firm's overall profits equal

$$\Pi = q_1 p_1 + \beta\frac{[\lambda + (1-\lambda)\widehat{r}]^2}{4} + \frac{\beta(1-\lambda)}{4}, \tag{11}$$

with p_1 and q_1 given by (7) and (9), respectively. Denote by $\widetilde{\Pi}(\widehat{r})$ the reduced-form profit function that we obtain by eliminating p_1 from (11) with the help of (7) and (9). The optimal $\widehat{r}$ satisfies the first-order condition $\partial\widetilde{\Pi}/\partial\widehat{r} = 0$ or, equivalently,

$$\widehat{r} = \frac{2 - \delta(1+\lambda)\left(1+\lambda-\delta\lambda^2\right) + \beta\lambda(1-\lambda)}{4 - \delta(1+\lambda)^2(2-\delta\lambda) - \beta(1-\lambda)^2}. \tag{12}$$

For this value of $\widehat{r}$ to indeed be part of an equilibrium, we must have $\widehat{r} \in \left(\frac{\sqrt{\lambda}}{1+\sqrt{\lambda}}, \frac{1}{2}\right)$ or equivalently $\beta \in \left(\mathcal{B}(\lambda,\delta), \delta^2\lambda\right)$, Thus if this condition holds, there is an equilibrium where $\widehat{r}$ is given by (12). The equilibrium values of the remaining prices are in turn obtained from (6) and (7).

3.3 Summing Up

Proposition 1. *Suppose Assumption 1 holds. Then there is a unique equilibrium of the model with exogenous λ. In this equilibrium, the relationship $\widehat{r} - p_1 = \delta\left(p_2^H - p_2^L\right)$ always holds. Moreover, the prices and the cutoff value in the equilibrium are as follows:*

(i) If $\beta \geq \delta^2\lambda$, then p_2^L, p_1, and $\widehat{r}$ are given by (6), (7), (10), respectively, and $p_2^H = \widehat{r}$. In particular, $p_2^L \leq p_1 \leq p_2^H = \widehat{r}$ and $\frac{1}{2} \leq \widehat{r}$.

(ii) If $\beta \in \left(\mathcal{B}(\lambda,\delta), \delta^2\lambda\right)$, then p_2^L, p_1, and $\widehat{r}$ are given by (6), (7), (12), respectively, and $p_2^H = \frac{1}{2}$. In particular, $p_2^L \leq p_1 \leq \widehat{r} < \frac{1}{2} = p_2^H$.

The results reported in Proposition 1 are illustrated in Figure 5. Panel (a) shows the (λ, β)-space (assuming $\delta = 0.6$) and indicates where a type (i) and a type (ii) equilibrium exists.

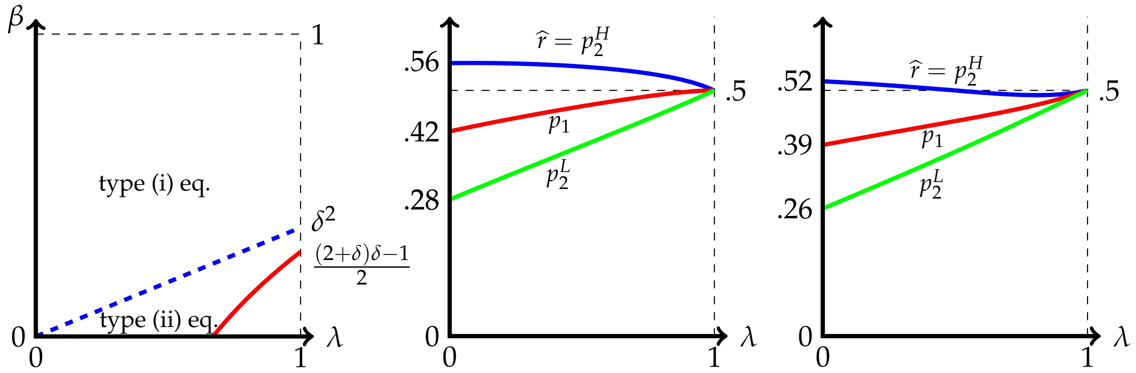

(a) The parameter space. **(b)** Example: $\beta = \delta = \frac{1}{2}$. **(c)** Example: $\beta = \frac{1}{8}$ and $\delta = \frac{1}{2}$.

Figure 5: Illustration of the results in Proposition 1.

South-east of the red (solid) curve, Assumption 1 is violated, so this part of the parameter space is left out of the analysis. Panel (b) of the figure graphs the equilibrium prices and the cutoff value in an example where the two discount factors are identical ($\beta = \delta = \frac{1}{2}$). Here the equilibrium is of type (i) for all values of λ. Panel (c), finally, shows an example with $\beta = \frac{1}{8}$ and $\delta = \frac{1}{2}$. Here the equilibrium is of type (i) for $\lambda \leq 0.5$ and of type (ii) for $\lambda > 0.5$. In particular, we see that the equilibrium cutoff value $\widehat{r}$ drops below one-half for $\lambda \in (0.5, 1)$, although the graph suggests that the drop is not very large.

It is useful to also consider an example in which neither the firm nor the consumers discount at all, i.e., where $\beta = \delta = 1$. As noted above, in this situation the equilibrium is of type (i). Moreover, the expressions become quite simple.

Example 1. *Suppose $\beta = \delta = 1$. Then the equilibrium is always of type (i) and the equilibrium prices and the cutoff value are given by*

$$p_2^L = p_1 = \frac{(3-\lambda)(1+\lambda)}{2(5-\lambda^2)} \quad \textit{and} \quad p_2^H = \widehat{r} = \frac{3-\lambda^2}{5-\lambda^2}.$$

One can check that the first expression (i.e., for p_2^L and p_1) is strictly increasing, and the second one (for p_2^H and $\widehat{r}$) is strictly decreasing. I will return to this example in Section 4.

3.4 Welfare

What is the effect of an exogenous change in λ on social welfare? To understand this, assume that the consumers and the firm are equally patient ($\delta = \beta$) and let our measure of social welfare be total surplus.[9] Proposition 1 tells us that, with $\delta = \beta$, the equilibrium is always of type (i) and hence $\widehat{r} \geq \frac{1}{2}$. Given that this kind of equilibrium is played, total surplus can be written as

$$W(\lambda) \stackrel{\text{def}}{=} \int_{\widehat{r}}^{1} r dr + \lambda \int_{p_1}^{\widehat{r}} r dr + \delta \int_{p_2^L}^{1} r dr, \qquad \text{for } \lambda \in [0,1]. \tag{13}$$

[9] With $\delta \neq \beta$, it is not clear how to define total surplus.

The two first terms in (13) represent the surplus generated in period 1. In that period, all consumers with valuation $r \geq \hat{r}$, regardless of their incognito status, purchase the good, which yields the surplus captured by the first term. Moreover, incognito consumers with valuations $r \in [p_1, \hat{r}]$ also buy the good in period 1, yielding the second term. In period 2, all consumers with valuations $r \in [p_2^L, 1]$ buy the good, which yields the surplus captured by the last term.

Let $\hat{\lambda}_W$ be the fraction of incognito consumers that maximizes the total surplus, as stated in (13).[10] To be able to say something about the value of $\hat{\lambda}_W$, first note that $W(0) > W(1)$ (this is shown in the proof of Proposition 2). That is, total surplus is strictly larger with no incognito consumers than with only incognito consumers. Intuitively, in the latter case price discrimination is not feasible and thus all consumers, in both periods, must pay the price one-half. In contrast, without any incognito consumers the firm can charge a separate second-period price (namely, $p_2^L < \frac{1}{2}$) for consumers with a relatively low valuation, which increases the amount of trade.

Next, differentiate the total surplus function with respect to λ:

$$\frac{\partial W}{\partial \lambda} = -\hat{r}\frac{\partial \hat{r}}{\partial \lambda} + \lambda \left[\hat{r}\frac{\partial \hat{r}}{\partial \lambda} - p_1 \frac{\partial p_1}{\partial \lambda} \right] + \int_{p_1}^{\hat{r}} r dr - \delta p_2^L \frac{\partial p_2^L}{\partial \lambda}. \tag{14}$$

The term with an integral sign in (14) represents the change in first-period surplus for the extra marginal incognito consumers who, thanks to their newly acquired incognito status, can consume the good in period 1. This term is clearly positive for all $\lambda < 1$ and in particular for $\lambda = 0$ (for λ close to one, however, the term must be close to zero, as then p_1 and $\hat{r}$ are close to each other). The preceding term (the one with square brackets) represents the change in first-period surplus for the infra marginal incognito consumers who already had an incognito status but now face adjustments in $\hat{r}$ and p_1. This is an indirect effect and it disappears for $\lambda = 0$. Also the two remaining terms capture indirect effects that are due to adjustments in one of the prices and in the cutoff value $\hat{r}$. These effects are harder to understand intuitively, but for $\lambda = 0$ one can show that they cannot overturn the positive effect coming from the term with the integral sign (see the proof of Proposition 2). All in all, this means that, at $\lambda = 0$, we have $\partial W / \partial \lambda > 0$. This result in combination with our observation from above that $W(0) > W(1)$ imply that the fraction of incognito consumers that maximizes total surplus must be strictly between zero and one.

Proposition 2. *Suppose $\delta = \beta$ and consider the model with an exogenous λ. The fraction of incognito surfers that maximizes total surplus lies strictly between zero and unity, $\hat{\lambda}_W \in (0, 1)$.*

Intuitively, to let all consumers have an incognito status is suboptimal as this hinders price discrimination, and price discrimination generates gains from trade. On the other hand, an increase in λ enables more consumers with valuations between $\hat{r}$ and p_1 to purchase the good in the first period, which also generates gains from trade; moreover, the difference between $\hat{r}$ and p_1 tends to be large for low values of λ, which makes this effect particularly important for $\lambda = 0$. Thus, increasing λ at least somewhat, starting from zero, always pays off.

[10] Formally, $\hat{\lambda}_W \in \arg\max_{\lambda \in [0,1]} W(\lambda)$.

4 Endogenous Fraction of Incognito Consumers

Now turn to the full model described in Section 2, where the incognito status is endogenous. I will assume that $\beta \geq \delta^2$, which implies that the equilibrium is of the first type reported in Proposition 1. At stage (ii) of the full model, each consumer chooses a level of λ that maximizes her expected utility (not yet knowing her valuation r) and expecting all other consumers to choose, say, $\lambda = \widetilde{\lambda}$. Because each consumer is infinitesimally small, her choice of λ has no impact on the prices or the cutoff value. Hence, only the direct effect on the consumer's utility, and of course the effect on the cost $C(\lambda)$, matter for the choice of λ.

The consumer's expected utility can be written as $EU(\lambda) = S(\lambda) - C(\lambda)$, where the gross consumer surplus $S(\lambda)$ is defined as

$$S(\lambda) \stackrel{\text{def}}{=} \int_{\widehat{r}}^{1} (r - p_1)\, dr + \lambda \int_{p_1}^{\widehat{r}} (r - p_1)\, dr + \delta \left[\int_{p_2^L}^{\widehat{r}} \left(r - p_2^L\right) dr + \int_{\widehat{r}}^{1} \left[r - (1-\lambda)\, p_2^H - \lambda p_2^L \right] dr \right] \quad (15)$$

and where all prices and the cutoff value are evaluated at $\lambda = \widetilde{\lambda}$. Differentiating $EU(\lambda)$ with respect to λ yields

$$\frac{\partial EU}{\partial \lambda} = \int_{p_1}^{\widehat{r}} (r - p_1)\, dr + \delta \int_{\widehat{r}}^{1} \left(p_2^H - p_2^L\right) dr - C'(\lambda)\,. \quad (16)$$

Thus, a consumer's marginal benefit from increasing λ has two components. In the first period the consumer increases her likelihood of earning a surplus whenever $r \in (p_1, \widehat{r})$; this effect is captured by the first term in (16). In the second period the consumer increases her likelihood of being eligible to pay p_2^L rather than p_2^H, when $r \in (\widehat{r}, 1]$; this is the second term in (16).

Given that the consumers are ex ante identical, it is natural to focus attention on symmetric equilibria, where all consumers choose $\lambda = \lambda^*$. In any such equilibrium $\partial EU / \partial \lambda = 0$ must hold, when (16) is evaluated at $\lambda = \widetilde{\lambda} = \lambda^*$.

Proposition 3. *Suppose $\beta \geq \delta^2$. Then there exists a symmetric equilibrium of the game with endogenous λ. The fraction of incognito surfers in this equilibrium, λ^*, satisfies $\lambda^* \in (0,1)$ and is implicitly defined by*

$$\int_{p_1^*}^{\widehat{r}^*} (1 - r) dr = C'(\lambda^*), \quad (17)$$

where p_1^* and $\widehat{r}^*$ are given by p_1 and $\widehat{r}$ as stated in Proposition 1 but evaluated at $\lambda = \lambda^*$.

What are the welfare properties of the equilibrium characterized in Proposition 3?[11] Is there too much or too little incognito surfing? Consider a total surplus standard; that is, let social welfare be defined as $W(\lambda) - C(\lambda)$, where $W(\lambda)$ is given by (13). To make sense of this standard, we must let the firm and the consumers use the same discount factor. I will thus,

[11]One can relatively easily verify that for the case $\beta = \delta = 1$, the equilibrium is guaranteed to be unique (because then the left-hand-side of (17) is downward-sloping in the incognito probability). I have no reason to believe that the equilibrium is not unique also for other values of β and δ. However, the algebra for the general case becomes quite intractable and I must therefore refrain from making such general uniqueness claims.

for the remainder of the paper, assume that $\beta = \delta$. The social marginal benefit, given the total surplus standard, is stated in (14). This marginal benefit consists of a direct welfare effect (namely, $\int_{p_1}^{\widehat{r}} r dr$) and an indirect welfare effect (the remaining terms in (14)).

Let the *direct external effect*, $\Delta_D(\lambda)$, be defined as the extent to which the atomistic consumer's direct private marginal benefit from a larger λ is greater than the direct welfare effect; that is,

$$\Delta_D(\lambda) \stackrel{\text{def}}{=} \int_{p_1}^{\widehat{r}} (1-r) dr - \int_{p_1}^{\widehat{r}} r dr. \tag{18}$$

Similarly, let the *indirect external effect*, $\Delta_I(\lambda)$, be defined as the extent to which the indirect private marginal benefit from a larger λ is greater than the indirect social welfare effect:

$$\Delta_I(\lambda) \stackrel{\text{def}}{=} \widehat{r}\frac{\partial \widehat{r}}{\partial \lambda} - \lambda \left[\widehat{r}\frac{\partial \widehat{r}}{\partial \lambda} - p_1 \frac{\partial p_1}{\partial \lambda} \right] + \delta p_2^L \frac{\partial p_2^L}{\partial \lambda} \tag{19}$$

(note that the indirect private marginal benefit is zero, as the consumer is atomistic). Finally, the *total external effect* is the sum of the direct and the indirect effects, $\Delta(\lambda) \stackrel{\text{def}}{=} \Delta_D(\lambda) + \Delta_I(\lambda)$. Given these definitions, I say that there is too much incognito surfing if $\Delta(\lambda^*) > 0$ and there is too little if $\Delta(\lambda^*) < 0$.

In order to understand under what circumstances we have too much or too little incognito surfing, consider first the direct external effect, $\Delta_D(\lambda)$. As already explained, this effect measures the extent to which the private (direct) marginal benefit from a larger λ is greater than the social direct marginal benefit from a larger λ. The latter marginal benefit equals the sum of the valuations of the additional consumers who, thanks to being able to surf incognito, find it worthwhile to purchase the good in the first period now when doing this has no impact on the second-period price; it is shown in Figure 6, panel (a), as the area $A_0 + A_1$. In contrast, the private marginal benefit corresponds to the area $A_0 + A_2$ in the same figure. The area A_0 is the sum of the *net* valuations of the additional consumers who purchase the good thanks to the incognito status (the consumers do not benefit from A_1 as this amount is paid to the firm). The area A_2 is the discounted value of the amount of money the consumer can save in period 2 thanks the incognito status, which entitles her to pay p_2^L instead of p_2^H for the good (these savings do not affect total surplus as they are just a transfer from the firm to the consumer). This discounted amount of money can be expressed as the area A_2 in the figure thanks to the equilibrium relationship $\widehat{r} - p_1 = \delta(p_2^H - p_2^L)$: at the equilibrium, $\widehat{r}$ is such that a consumer with the valuation $r = \widehat{r}$ is indifferent between between purchasing in the first period and not doing that—see eq. (2).

We thus have $\Delta_D(\lambda) = A_2 - A_1$. Moreover, it is clear from Figure 6, panel (a), that $A_2 > A_1$ if and only if $1 - \widehat{r} > p_1$. The latter inequality can be solved for δ, which yields the following result:

$$\Delta_D(\lambda) > 0 \Leftrightarrow \delta > \frac{2\lambda}{1+\lambda^2}. \tag{20}$$

That is, the direct external effect is positive if and only if the consumers care sufficiently much about the second period. (See Figure 6, panel (b), for an illustration.) In other words, if the consumers are sufficiently forward-looking, then the direct external effect, all else equal, works in

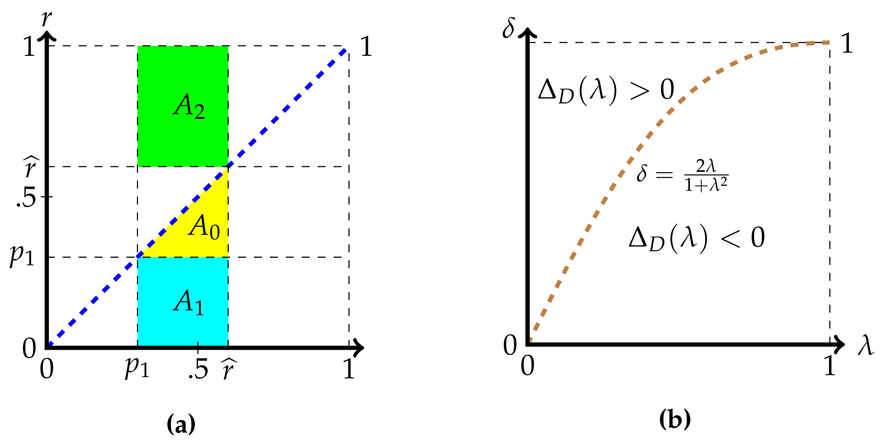

Figure 6: Social versus individual incentives to choose incognito: the direct effect. **(a)** Society cares about $A_0 + A_1$, but the atomistic consumer cares about $A_0 + A_2$. **(b)** The sign of the direct external effect, $\Delta_D(\lambda) = A_2 - A_1$.

the direction of too much incognito surfing. Intuitively, in the first period the consumers benefit too little from an incognito status,[12] whereas in the second period the consumers benefit too much.[13] Thus, if the second period matters sufficiently much relative to the first period, the consumers might have a too strong incentive to invest in λ.

Indeed, the graph in Figure 6, panel (b), shows that if the discount factor is sufficiently close to one, then the direct external effect is positive for all values of λ. Similarly, if the discount factor is sufficiently close to zero, then the direct external effect is negative for all values of λ. Of course, also the indirect external effect matters for whether there is too much or too little incognito surfing. However, it seems plausible that the indirect external effect might be of second-order importance (after all, this effect influences welfare only through the equilibrium values of the prices and the cutoff $\widehat{r}$, not directly through λ). For the case with a large discount factor, this reasoning turns out to be correct. That is, as stated in Proposition 4 below, if the discount factor equals unity, there is always too much incognito surfing.

Proposition 4. *Suppose $\beta = \delta = 1$. Then, relative to a total-surplus maximizing benchmark, the equilibrium yields too much incognito surfing:* $\Delta(\lambda^*) > 0$.

For the case where the discount factor is below unity, it is more difficult to obtain analytical results. However, let us assume the following functional form for the cost function:

$$C(\lambda) = \frac{c}{2}\lambda^2, \qquad c > 0. \tag{21}$$

We can then perform the welfare comparison using numerical methods. Also, let λ^*_W denote the value of λ that maximizes $W(\lambda) - C(\lambda)$. The numerical analysis reveals that—as hypothesized above—we can indeed obtain the result that there is too little incognito surfing

[12] The private benefit is A_0 but the social benefit is $A_0 + A_1$.
[13] The private benefit is A_2 but the social benefit is zero.

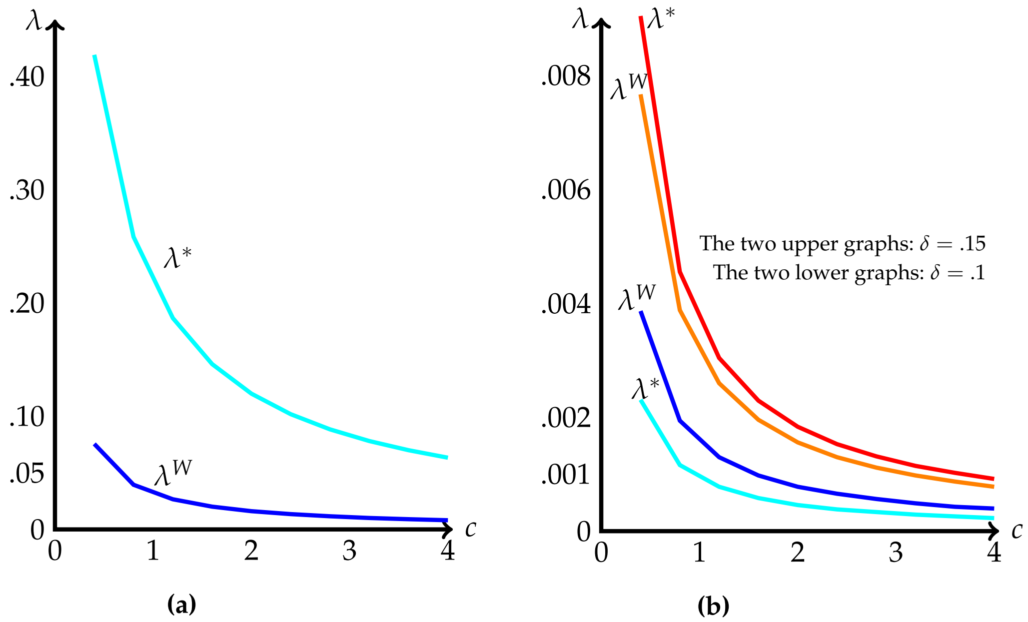

Figure 7: Socially optimal vs. equilibrium values of λ. **(a)** The case $\delta = 0.9$. There is too much incognito surfing at the equilibrium. **(b)** The cases $\delta \in \{0.1, 0.15\}$. Equilibrium undersupply of λ is possible for $\delta = 0.1$.

($\lambda^* < \lambda^*_W$), provided that we pick a value of the discount factor that is low enough. Results of the simulation exercise are shown in Figure 7. Panel (a) considers a case where the discount factor is relatively large, $\delta = 0.9$. We see that then, in keeping with Proposition 4, the equilibrium yields more incognito surfing than is socially desirable. Panel (b) considers examples with $\delta = 0.15$ and $\delta = 0.1$. For $\delta = 0.15$, we still obtain the result that there is too much incognito surfing, However, for $\delta = 0.1$, the graphs show that the equilibrium yields less incognito surfing than is socially desirable (for a range of different values of the cost parameter c).

We summarize the main insight from the simulation exercise as follows:[14]

Result 1. *Suppose $\beta = \delta$ and let the cost function $C(\lambda)$ be given by (21). Then, by letting δ be sufficiently small, we can construct numerical examples where, relative to a total surplus welfare benchmark, the market outcome yields too little incognito surfing: $\lambda^* < \lambda^*_W$.*

Summing up, the intuition for the results reported in Proposition 4 and Result 1 is that the consumers benefit from an incognito status both in period 1 and in period 2. However, relative to the social benefit, the consumers' benefit is too small in period 1 and too large in period 2. Therefore, if they discount heavily (low δ), they tend to have a socially too weak incentive to invest in an incognito status. In contrast, if they assign a sufficiently large weight to the second-period payoffs, the consumers tend to have a too strong incentive.

[14]Details about the simulation exercise can be found in the Supplementary Material, Lagerlöf (2018). This document and the Matlab code are available at www.johanlagerlof.com.

5 Concluding Remarks

I have studied the incentives of consumers to hide their purchases, in an environment in which a monopoly firm practices behavior-based price discrimination. The analysis yielded two main results. First, in the version of the model where the fraction of hiding consumers is exogenous (and hiding is costless), the total-surplus maximizing level of this fraction is strictly interior. The reason is that both (i) a fraction of zero and (ii) a fraction of one would fail to exploit gains from trade—in case (i) due to the fact that unnecessarily many consumers with valuations above the price do not purchase in period 1, and in case (ii) due to the firm's inability to practice price discrimination in period 2. Second, in the version of the model where the choice of hiding is endogenous (and comes at a cost), the market outcome yields, from a social welfare point of view, too much hiding if time discounting is small and too little if discounting is large. This is because the sign of the first-period externality differs from that of the second-period externality.

Throughout the analysis I have maintained the assumption that the hiding probability is chosen at an ex ante stage. I suggested in Section 2 that this model feature naturally captures the idea that a consumer adopts a simple rule or heuristic that she uses in a wide range of situations, updating it only occasionally. Moreover, I argued that—relative to an alternative model where the hiding choice is made with knowledge about the own valuation—nothing suggests that we systematically under- or overestimate the strength in the consumer's incentive by letting her choice be made at an ex ante stage (see f.n. 5). It would nevertheless be interesting to explore the alternative setting where the consumer makes her choice ex post. However, one analytical difficulty with such an extension is the fact that the endogenous hiding probability will depend on the consumer's valuation.[15] This, in turn, makes the firm's objective function at the stage where it chooses $\widehat{r}$ (or, equivalently, p_1) a polynomial of a higher degree than two (in the current model, it is quadratic). As a consequence, solving for the equilibrium value of $\widehat{r}$ analytically would become much less straightforward and in some cases impossible.

Another extension that might provide further insights would be to allow the firm to take some (costly) action that makes it harder for consumers to hide their purchasing history (cf. Johnson's (2013) model, in which firms can choose the number of advertisements that the consumers are exposed to and try to protect themselves from). Finally, it would be interesting to study the effects on the present paper's results of a change in the degree of competition in the market.

[15] More precisely, for consumers with $r \in (\widehat{r}, 1]$, the hiding probability will *not* depend on r, as the benefit from hiding equals the difference between the two second-period prices (cf. Figure 1). However, for consumers with $r \in (p_1, \widehat{r})$, the chosen hiding probability will indeed depend on the own valuation, as the benefit from hiding equals $r - p_1$.

Appendix: Proofs

Proof of Lemma 1. Using the information in the text, we can write total second-period profits in the low-valuation market as

$$\pi_2^L\left(p_2^L\right) = \begin{cases} [\lambda + (1-\lambda)\widehat{r} - p_2^L]p_2^L & \text{if } p_2^L \in [0, \widehat{r}] \\ \lambda\left(1 - p_2^L\right)p_2^L & \text{if } p_2^L \in [\widehat{r}, 1]. \end{cases}$$

This profit expression is continuous in p_2^L, but it is not necessarily concave or quasiconcave. However, the expression is clearly concave (and quadratic) in p_2^L in each of the two ranges $[0, \widehat{r}]$ and $[\widehat{r}, 1]$. The solution to the problem of maximizing π_2^L with respect to p_2^L subject to $p_2^L \in [0, \widehat{r}]$ can therefore easily be found with the help of a first-order condition. The solution is given by

$$\widehat{p}_2^L = \begin{cases} \frac{\lambda+(1-\lambda)\widehat{r}}{2} & \text{if } \widehat{r} \geq \frac{\lambda}{1+\lambda} \\ \widehat{r} & \text{if } \widehat{r} \leq \frac{\lambda}{1+\lambda}. \end{cases} \tag{A1}$$

Similarly, the solution to the problem of maximizing π_2^L with respect to p_2^L subject to $p_2^L \in [\widehat{r}, 1]$ is given by

$$\widetilde{p}_2^L = \begin{cases} \widehat{r} & \text{if } \widehat{r} \geq \frac{1}{2} \\ \frac{1}{2} & \text{if } \widehat{r} \leq \frac{1}{2}. \end{cases} \tag{A2}$$

Notice that the cutoff point in (A1) is strictly smaller than the one in (A2): $\frac{\lambda}{1+\lambda} < \frac{1}{2}$. Therefore, if $\widehat{r} \leq \frac{\lambda}{1+\lambda}$, we should compare the profits at $\widetilde{p}_2^L = \frac{1}{2}$ and at $\widehat{p}_2^L = \widehat{r}$ in order to find the global optimum for that region. With a bit of algebra one can verify that this optimum is at $\widetilde{p}_2^L = \frac{1}{2}$, which is indeed consistent with equation (5). Similarly, if $\widehat{r} > \frac{1}{2}$, we should compare the profits at $\widehat{p}_2^L = \frac{\lambda+(1-\lambda)\widehat{r}}{2}$ and at $\widetilde{p}_2^L = \widehat{r}$ in order to find the global optimum for that region. Again, with a bit of algebra one can check that this optimum is at $\widehat{p}_2^L = \frac{\lambda+(1-\lambda)\widehat{r}}{2}$, which is consistent with equation (5). The comparison that remains is the one for $\widehat{r} \in \left[\frac{\lambda}{1+\lambda}, \frac{1}{2}\right]$, in which we must compare the profits at $\widetilde{p}_2^L = \frac{1}{2}$ and at $\widehat{p}_2^L = \frac{\lambda+(1-\lambda)\widehat{r}}{2}$. The profits at the latter price equal $\pi_2^L\left(\widehat{p}_2^L\right) = \left[\lambda + (1-\lambda)\widehat{r}\right]^2/4$, whereas the profits at $\widetilde{p}_2^L = \frac{1}{2}$ equal $\pi_2^L\left(\widetilde{p}_2^L\right) = \frac{\lambda}{4}$. We can now compare these profit levels:

$$\pi_2^L\left(\widehat{p}_2^L\right) > \pi_2^L\left(\widetilde{p}_2^L\right) \Leftrightarrow \frac{\left[\lambda + (1-\lambda)\widehat{r}\right]^2}{4} > \frac{\lambda}{4} \Leftrightarrow \widehat{r} > \frac{\sqrt{\lambda}}{1+\sqrt{\lambda}}.$$

We can conclude that the optimal price is indeed as stated in Lemma 1. □

Proof of Proposition 1. We first must show that a type (iii) equilibrium cannot exist. To this end, note that in such an equilibrium, $p_2^H = p_2^L = \frac{1}{2}$ (this follows from $p_2^H = \max\left\{\frac{1}{2}, \widehat{r}\right\}$ and Lemma 1). The threshold $\widehat{r}$ must satisfy (2) with equality when evaluated at $p_2^H = p_2^L = \frac{1}{2}$, which implies that $\widehat{r} = p_1$. That is, the consumers choose to purchase in period 1 if, and only if, their valuation exceeds the first-period price. This means that there is effectively no interaction between the two periods and the firm's problem of choosing p_1 is tantamount to the problem of choosing p_1 in a one-period model. Hence, $p_1 = \frac{1}{2}$ $(= \widehat{r})$. However, $\widehat{r} = \frac{1}{2}$ contradicts the initial assumption that we have case (iii). It follows that a type (iii) equilibrium cannot exist.

The statements about existence, uniqueness, and characterization of the type (i) and type (ii) equilibria are proven in the main text. The claim that $\widehat{r} - p_1 = \delta\left(p_2^H - p_2^L\right)$ follows from (2) holding with an equality. It remains to verify the claims about the relationships between the cutoff value and the prices. First consider the type (i) equilibrium. The claim that $p_1 \leq \widehat{r}$ follows immediately from (7) and the fact that $\widehat{r} \geq \frac{1}{2}$. Next, to prove the claim that $p_2^L \leq p_1$, note that (2) holding with an equality yields

$$p_1 - \delta p_2^L = (1-\delta)\widehat{r} \Leftrightarrow p_1 - p_2^L = (1-\delta)\left(\widehat{r} - p_2^L\right) \geq 0, \tag{A3}$$

where the inequality follows from $\widehat{r} \geq \frac{1}{2}$ and $p_2^L \leq \frac{1}{2}$ (that the latter must hold can be seen from (6)). The other relationships for the type (i) equilibrium are proven in the main text or are straightforward.

Consider the type (ii) equilibrium. The claim that $p_1 \leq \widehat{r}$ follows immediately from (7) and the fact that

$\widehat{r} \geq \frac{\lambda}{1+\lambda}$.[16] Next, to prove the claim that $p_2^L \leq p_1$, note that for $\delta = 1$ we have $p_2^L = p_1$ (see (6) and (7)). Also note that for $\delta < 1$ we can write (again using (6) and (7))

$$p_1 \geq p_2^L \Leftrightarrow \widehat{r} - \frac{\delta\left[(1+\lambda)\widehat{r} - \lambda\right]}{2} \geq \frac{\lambda + (1-\lambda)\widehat{r}}{2} \Leftrightarrow \widehat{r} \geq \frac{\lambda}{1+\lambda}, \tag{A4}$$

which always holds (see footnote 16). The other relationships for the type (ii) equilibrium are proven in the main text or are straightforward. □

Proof of Proposition 2. In order to show that $\widehat{\lambda}_W$ is interior, first note the following relationship:

$$\frac{12 + 16\delta + 3\delta^2}{8(4+\delta)} = W(0) > W(1) = \frac{3(1+\delta)}{8}. \tag{A5}$$

That is, total surplus is strictly larger with no incognito surfers than with only incognito surfers. The expression for $W(0)$ in (A5) was obtained from (13) and by noting that, evaluated at $\lambda = 0$, $\widehat{r} = (2+\delta)/(4+\delta)$ and $p_2^L = (2+\delta)/[2(4+\delta)]$. Similarly, the expression for $W(1)$ was obtained from (13) and by noting that, evaluated at $\lambda = 1$, $\widehat{r} = p_1 = p_2^L = 1/2$.

Given (A5) and the arguments already provided in the main text, it remains to show that $\lim_{\lambda \to 0} \partial W(\lambda)/\partial\lambda > 0$ for all $\delta \in (0,1]$. To do that, first write the expression for $\widehat{r}$ stated in (10) as follows: $\widehat{r} = N(\lambda,\delta)/D(\lambda,\delta)$, where

$$N(\lambda,\delta) \stackrel{\text{def}}{=} 2 - \delta(1+\lambda)\left(1+\lambda-\delta\lambda^2\right) + \beta(1-\lambda)(2+\lambda),$$

$$D(\lambda,\delta) \stackrel{\text{def}}{=} 4 - \delta(1+\lambda)^2(2-\delta\lambda) + \beta(1-\lambda)(3+\lambda).$$

Differentiating yields

$$\frac{\partial N(\lambda,\delta)}{\partial\lambda} = -\delta\left[\left(1+\lambda-\delta\lambda^2\right) + (1+\lambda)(1-2\delta\lambda)\right] + \beta\left[-(2+\lambda)+(1-\lambda)\right],$$

$$\frac{\partial D(\lambda,\delta)}{\partial\lambda} = -\delta\left[2(1+\lambda)(2-\delta\lambda) - \delta(1+\lambda)^2\right] + \beta\left[-(3+\lambda)+(1-\lambda)\right].$$

Taking limits, we have

$$\lim_{\lambda\to 0} N(\lambda,\delta) = 2 - \delta + 2\beta, \qquad \lim_{\lambda\to 0} D(\lambda,\delta) = 4 - 2\delta + 3\beta,$$

$$\lim_{\lambda\to 0} \frac{\partial N(\lambda,\delta)}{\partial\lambda} = -2\delta - \beta, \qquad \lim_{\lambda\to 0} \frac{\partial D(\lambda,\delta)}{\partial\lambda} = -\delta(4-\delta) - 2\beta.$$

We can now write

$$\begin{aligned}
\lim_{\lambda\to 0} \frac{\partial\widehat{r}}{\partial\lambda} &= \lim_{\lambda\to 0} \frac{\frac{\partial N(\lambda)}{\partial\lambda} D(\lambda,\delta) - N(\lambda,\delta)\frac{\partial D(\lambda)}{\partial\lambda}}{[D(\lambda,\delta)]^2} \\
&= \frac{(-2\delta-\beta)(4-2\delta+3\beta) - (2-\delta+2\beta)\left[-\delta(4-\delta)-2\beta\right]}{(4-2\delta+3\beta)^2} \\
&= \frac{-3\delta(4+\delta) + \delta(2+\delta)(6-\delta)}{(4+\delta)^2} = \frac{\delta^2(1-\delta)}{(4+\delta)^2},
\end{aligned}$$

where the last line uses $\beta = \delta$. We also have

$$\lim_{\lambda\to 0} \widehat{r} = \frac{2-\delta+2\beta}{4-2\delta+3\beta} = \frac{2+\delta}{4+\delta},$$

where again the last equality uses $\beta = \delta$. Next, from (6) we have $p_2^L = \frac{1}{2}\left[\lambda + (1-\lambda)\widehat{r}\right]$, which yields

$$\frac{\partial p_2^L}{\partial\lambda} = \frac{1}{2}\left[1 - \widehat{r} + (1-\lambda)\frac{\partial\widehat{r}}{\partial\lambda}\right].$$

[16] We know that $\widehat{r} \geq \frac{\lambda}{1+\lambda}$, because $\beta > \mathcal{B}(\lambda,\delta) \Leftrightarrow \widehat{r} > \frac{\sqrt{\lambda}}{1+\sqrt{\lambda}}$ and $\frac{\sqrt{\lambda}}{1+\sqrt{\lambda}} \geq \frac{\lambda}{1+\lambda}$.

Thus, $\lim_{\lambda\to 0} p_2^L = \frac{1}{2}\lim_{\lambda\to 0}\widehat{r}$. Moreover, we can write $\int_{p_1}^{\widehat{r}} rdr = \frac{1}{2}\left(\widehat{r}^2 - p_1^2\right)$. Hence,

$$\lim_{\lambda\to 0}\int_{p_1}^{\widehat{r}} rdr = \frac{\left[\lim_{\lambda\to 0}\widehat{r}\right]^2 - \left[\frac{2-\delta}{2}\lim_{\lambda\to 0}\widehat{r}\right]^2}{2} = \frac{\delta(4-\delta)}{8}\left[\lim_{\lambda\to 0}\widehat{r}\right]^2,$$

where (7) was used to obtain $\lim_{\lambda\to 0} p_1$. By using (14) and the above results, we can write

$$\begin{aligned}
\lim_{\lambda\to 0}\frac{\partial W}{\partial\lambda} &= -\lim_{\lambda\to 0}\widehat{r}\lim_{\lambda\to 0}\frac{\partial\widehat{r}}{\partial\lambda} + \lim_{\lambda\to 0}\int_{p_1}^{\widehat{r}} rdr - \delta\lim_{\lambda\to 0} p_2^L\lim_{\lambda\to 0}\frac{\partial p_2^L}{\partial\lambda} \\
&= \left[\lim_{\lambda\to 0}\widehat{r}\right]\left[-\lim_{\lambda\to 0}\frac{\partial\widehat{r}}{\partial\lambda} + \frac{\delta(4-\delta)}{8}\lim_{\lambda\to 0}\widehat{r} - \frac{\delta}{2}\lim_{\lambda\to 0}\frac{\partial p_2^L}{\partial\lambda}\right] \\
&= \left[\lim_{\lambda\to 0}\widehat{r}\right]\left[-\lim_{\lambda\to 0}\frac{\partial\widehat{r}}{\partial\lambda} + \frac{\delta(4-\delta)}{8}\lim_{\lambda\to 0}\widehat{r} - \frac{\delta}{4}\left(1 - \lim_{\lambda\to 0}\widehat{r} + \lim_{\lambda\to 0}\frac{\partial\widehat{r}}{\partial\lambda}\right)\right] \\
&= \left[\lim_{\lambda\to 0}\widehat{r}\right]\left[-\frac{4+\delta}{4}\lim_{\lambda\to 0}\frac{\partial\widehat{r}}{\partial\lambda} + \frac{\delta(6-\delta)}{8}\lim_{\lambda\to 0}\widehat{r} - \frac{\delta}{4}\right] \\
&= \left[\lim_{\lambda\to 0}\widehat{r}\right]\left[\frac{-2\delta^2(1-\delta) + \delta(6-\delta)(2+\delta) - 2\delta(4+\delta)}{8\,(4+\delta)}\right] = \left[\lim_{\lambda\to 0}\widehat{r}\right]\left[\frac{\delta(4+\delta^2)}{8\,(4+\delta)}\right] > 0.
\end{aligned}$$

The last inequality holds for all $\delta \in (0,1]$, as $\lim_{\lambda\to 0}\widehat{r} > 0$ and the numerator of the ratio is also positive. □

Proof of Proposition 3. First note that we cannot have a symmetric equilibrium with $\lambda^* = 0$ or $\lambda^* = 1$. For if $\lambda^* = 1$, then all the prices are the same and thus the two first terms in (16) vanish, whereas the third term (i.e., the marginal cost) is strictly positive; hence, $\partial EU/\partial\lambda < 0$ at $\lambda = 1$ and the consumer would have an incentive to choose $\lambda < 1$. Similarly, if $\lambda^* = 0$, then the sum of the two first terms in (16) is strictly positive, but the marginal cost is zero; as a consequence, $\partial EU/\partial\lambda > 0$ at $\lambda = 0$ and the consumer would want to choose $\lambda > 0$. In order to show existence of a symmetric equilibrium with $\lambda^* \in (0,1)$, it suffices to note that the equation $\int_{p_1}^{\widehat{r}}(r-p_1)\,dr + \delta\int_{\widehat{r}}^{1}\left(p_2^H - p_2^L\right)dr = C'(\lambda^*)$ must have at least one root $\lambda^* \in (0,1)$. For the right-hand side is, by assumption, increasing in λ^* and it equals zero at $\lambda^* = 0$; moreover, the left-hand side is strictly positive evaluated at $\lambda^* = 0$ and zero evaluated at $\lambda^* = 1$. A symmetric equilibrium is characterized by the equation just stated, and the left-hand side of this can be rewritten as

$$\begin{aligned}
\int_{p_1}^{\widehat{r}}(r-p_1)\,dr + \delta\int_{\widehat{r}}^{1}\left(p_2^H - p_2^L\right)dr &= \int_{p_1}^{\widehat{r}}(r-p_1)\,dr + (\widehat{r}-p_1)(1-\widehat{r}) \\
&= \frac{1}{2}(\widehat{r}-p_1)(2-\widehat{r}-p_1) = \int_{p_1}^{\widehat{r}}(1-r)\,dr,
\end{aligned}$$

where the first equality is due to the relationship $\widehat{r} - p_1 = \delta\left(p_2^H - p_2^L\right)$ stated in Proposition 1. □

Proof of Proposition 4. Under the assumption that $\beta = \delta = 1$, we have $p_2^L = p_1 = (3-\lambda)(1+\lambda)/\left[2\left(5-\lambda^2\right)\right]$ and $p_2^H = \widehat{r} = \left(3-\lambda^2\right)/\left(5-\lambda^2\right)$ (see Example 1). Differentiating these expressions yields

$$\frac{\partial p_1}{\partial\lambda} = \frac{5-2\lambda+\lambda^2}{\left(5-\lambda^2\right)^2} \quad\text{and}\quad \frac{\partial\widehat{r}}{\partial\lambda} = -\frac{4\lambda}{\left(5-\lambda^2\right)^2}.$$

This means that the direct welfare effect can be written as

$$\Delta_D(\lambda) = \int_{p_1}^{\widehat{r}}(1-2r)\,dr = r(1-r)\,|_{p_1}^{\widehat{r}} = \widehat{r}(1-\widehat{r}) - p_1(1-p_1) = (\widehat{r}-p_1)(1-\widehat{r}-p_1) = \frac{(3+\lambda)(1-\lambda)}{2\left(5-\lambda^2\right)}.$$

Similarly, the indirect welfare effect can be written as

$$\Delta_I(\lambda) = (1-\lambda)\,\widehat{r}\frac{\partial\widehat{r}}{\partial\lambda} + (1+\lambda)\,p_1\frac{\partial p_1}{\partial\lambda} = \frac{\left(5-\lambda^2\right)\left(3-\lambda+5\lambda^2+\lambda^3\right)}{2\left(5-\lambda^2\right)^3}.$$

From inspection, it is clear that $\Delta_I(\lambda) > 0$ for all $\lambda \in [0,1]$ and $\Delta_D(\lambda) \geq 0$ for all $\lambda \in [0,1]$ (with an equality only for $\lambda = 1$). □

References

Alessandro Acquisti and Hal R. Varian. Conditioning Prices on Purchase History. *Marketing Science*, 24(3):367–381, 2005.

Alessandro Acquisti, Curtis Taylor, and Liad Wagman. The Economics of Privacy. *Journal of Economic Literature*, 54(2):542–492, 2016.

Mark Armstrong. Recent Developments in the Economics of Price Discrimination. In Richard Blundell, Whitney K. Newey, and Torsten Persson, editors, *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, Volume II*, pages 97–141. Cambridge University Press, 2006.

Giacomo Calzolari and Alessandro Pavan. On the Optimality of Privacy in Sequential Contracting. *Journal of Economic Theory*, 130(1):168–204, 2006.

Yongmin Chen. Paying Customers to Switch. *Journal of Economics & Management Strategy*, 6 (4):877–897, 1997.

Yuxin Chen and Z. John Zhang. Dynamic Targeted Pricing with Strategic Consumers. *International Journal of Industrial Organization*, 27(1):43–50, 2009.

Vincent Conitzer, Curtis R. Taylor, and Liad Wagman. Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases. *Marketing Science*, 31(2):277–292, 2012.

Rosa Branca Esteves. Customer Poaching and Advertising. *The Journal of Industrial Economics*, 57(1):112–146, 2009a.

Rosa Branca Esteves. A Survey on the Economics of Behaviour-Based Price Discrimination. Technical report, NIPE—Universidade do Minho, 2009b.

Drew Fudenberg and Jean Tirole. Customer Poaching and Brand Switching. *The RAND Journal of Economics*, 31(4):634–657, 2000.

Drew Fudenberg and J. Miguel Villas-Boas. Chapter 7: Behavior-Based Price Discrimination and Customer Recognition. In Terrence Hendershott, editor, *Handbook of Economics and Information Systems, Volume 1*, pages 377–436. Emerald, 2006.

Thomas Gehrig, Oz Shy, and Rune Stenbacka. History-Based Price Discrimination and Entry in Markets with Switching Costs: a Welfare Analysis. *European Economic Review*, 55(5):732–739, 2011.

Benjamin E. Hermalin and Michael L. Katz. Privacy, Property Rights and Efficiency: The Economics of Privacy as Secrecy. *Quantitative Marketing and Economics*, 4(3):209–239, 2006.

Kai-Lung Hui and I. P. L. Png. Chapter 9: The Economics of Privacy. In Terrence Hendershott, editor, *Handbook of Economics and Information Systems, Volume 1*, pages 471–498. Emerald, 2006.

Justin P. Johnson. Targeted Advertising and Advertising Avoidance. *The RAND Journal of Economics*, 44(1):128–144, 2013.

Johan N. M. Lagerlöf. Supplementary Material to "Surfing Incognito: Welfare Effects of Anonymous Shopping". Mimeo, University of Copenhagen, May 2019.

Annie Lowrey. How Online Retailers Stay a Step Ahead of Comparison Shoppers. Washington Post, December 12, 2010. URL http://www.washingtonpost.com/wp-dyn/content/article/2010/12/11/AR2010121100143.html.

Richard A. Posner. The Economics of Privacy. *The American Economic Review*, 71(2):405–409, 1981.

George J. Stigler. An Introduction to Privacy in Economics and Politics. *The Journal of Legal Studies*, 9(4):623–644, 1980.

Curtis R. Taylor. Consumer Privacy and the Market for Customer Information. *The RAND Journal of Economics*, 35(4):631–650, 2004.

J. Miguel Villas-Boas. Price Cycles in Markets with Customer Recognition. *The RAND Journal of Economics*, 35(3):486–501, 2004.