Games and Economic Behavior 130 (2021) 499–520

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Consumers' privacy choices in the era of big data

Sebastian Dengler [1], Jens Prüfer [*]

Department of Economics, CentER, TILEC; Tilburg University, P.O. Box 90153, 5000 LE Tilburg, the Netherlands

ARTICLE INFO

Article history:
Received 6 March 2020
Available online 6 October 2021

JEL classification:
L11
D11
D83
D01
L86

Keywords:
Privacy
Big data
Perfect price discrimination
Level-k thinking

ABSTRACT

Recent progress in information technologies provides sellers with detailed knowledge about consumers' preferences, approaching perfect price discrimination in the limit. We construct a model where consumers with less strategic sophistication than the seller's pricing algorithm face a trade-off when buying. They choose between a direct, transaction cost-free sales channel and a privacy-protecting, but costly, anonymous channel. We show that the anonymous channel is used even in the absence of an explicit taste for privacy if consumers are not *too* strategically sophisticated. This provides a micro-foundation for consumers' privacy choices. Some consumers benefit but others suffer from their anonymization.

1. Introduction

Two recent technological developments are revolutionizing seller-buyer transactions. First, aided by information and communication technologies (ICTs), sellers have the capability to analyze huge datasets with very detailed information about individual consumers' characteristics and preferences. Second, such datasets are increasingly available, owing to the fact that more economic and social transactions take place supported by ICTs, which easily and inexpensively store the information they produce or transmit.[2] These concurrent developments constitute the rise of big data (Mayer-Schönberger and Cukier, 2013). They imply that sellers can make ever more tailored offers which fit consumers' individual preferences or consumption patterns, approaching first-degree (or perfect) price discrimination as the limit case (Einav and Levin, 2014).

Because first-degree price discrimination can deprive consumers of all surplus from the transaction, they may want to protect their privacy and hide their willingness-to-pay (WTP) from sellers with market power by employing anonymization techniques. But anonymization is costly, which puts consumers at a disadvantage: it can come at an explicit cost or at an opportunity cost. Consumers are at a second disadvantage, compared to sellers, because they "will often be overwhelmed

* Corresponding author.
E-mail addresses: Sebastian.Dengler@ABCecon.com (S. Dengler), j.prufer@uvt.nl (J. Prüfer).
1 Present address: ABC economics, Boyenstr. 40, 10115 Berlin, Germany.

2 Data analytics firms collect and analyze huge commercial databases on consumers, offering help to marketers. For instance, *Acxiom's* "database contains information about 500 million active consumers worldwide, with about 1,500 data points per person. That includes a majority of adults in the United States" (The New York and Times, 2012). Smartphone apps with millions of users, such as *Shopkick*, reward users for checking into stores, scanning products, visiting the dressing rooms, and so forth. *Amazon* recently was issued a patent on a novel *Method and System for Anticipatory Package Shipping* (Patent number US008615473 (December 24, 2013), http://pdfpiw.uspto.gov/.piw?docid=08615473). "So Amazon says it may box and ship products it expects customers in a specific area will want – based on previous orders and other factors – but haven't yet ordered" (Wall Street Journal Blog, 2014).

https://doi.org/10.1016/j.geb.2021.09.007
0899-8256/

with the task of identifying possible outcomes related to privacy threats and means of protection. [...] Especially in the presence of complex, ramified consequences associated with the protection or release of personal information, our innate bounded rationality limits our ability to acquire, memorize and process all relevant information, and it makes us rely on simplified mental models, approximate strategies, and heuristics" (Acquisti and Grossklags, 2007, p. 369).

In our model, we study the effects of perfect price discrimination on consumers' choices and welfare when anonymization is possible but costly. We explicitly account for the discrepancy between cognitively challenged consumers and a seller whose strategic capabilities outperform them and investigate how limited strategic sophistication affects outcomes.

Our main contribution is to show under which conditions a costly privacy-protective sales channel is used even if consumers do not have an explicit taste for privacy, and how this depends on consumers' sophistication. We thereby provide a micro-foundation for consumers' privacy choices when facing a seller with access to big data.[3]

We construct a model where a mass of consumers with heterogeneous WTP for a product faces a monopolistic seller. Consumers can decide between two channels to buy the product from the seller. The direct channel (D) makes use of all personal information that the seller has about every single consumer, which includes each consumer's WTP. We assume that perfect price discrimination is feasible for the seller in channel D and that this channel economizes on transaction costs, which we normalize to zero. The anonymous channel (A) protects consumers' privacy by hiding individual identities, but comes at a cost, which we denote by s. As a consequence, perfect price discrimination is infeasible for the seller, who responds best by setting a uniform price for channel A.

Our model therefore represents a situation *after* a long period where consumers disclose information (due to neglect or lack of suitable technologies). Hence, the seller could acquire data shedding light on individual consumers' preferences, be it via collecting such information in the past (e.g. Amazon) or via buying such information from an intermediary (e.g. Google, Acxiom). However, as consumers decide about anonymization, the seller can neither directly influence their channel choice nor close down the anonymous channel.[4]

In the model, consumers first choose between channel D and channel A. Second, the seller observes which channel a consumer has used and sets prices in both channels. Third, every consumer decides whether to buy for the price offered, or not. Because prices are set by the seller after consumers chose a channel, each consumer only observes the price offered in the chosen channel. Hence, before choosing a channel, each consumer has to form a belief about the price charged in each channel. The price actually charged by the seller is a best-response to those beliefs. The critical question is to which extent consumers understand this.

Our analysis is based on a model of limited strategic sophistication, called *level-k thinking*, which was introduced by Stahl and Wilson (1994, 1995) and Nagel (1995). Models with level-k thinking are defined recursively, starting with so-called "naïve" level-0 players which employ a "naïve" (often random) strategy. Level-1 players then best respond to the level-0 strategy, level-2 players to the level-1 strategy, and so forth.[5] A sizable literature has developed that explores level-k thinking theoretically and empirically.[6] The literature has found strong experimental support for level-k thinking both in static (Camerer et al., 2004; Crawford and Iriberri, 2007b) and in dynamic settings (Kawagoe and Takizawa, 2012; García-Pola et al., 2020) and suggests values for k of one, two, or three.

Compared to the behavior-based price discrimination literature, where typically either unlimited strategic sophistication or complete naïveté of consumers is assumed, we zoom in and provide an analysis of behavior when players have *some* strategic sophistication.[7] We model consumers' cognitive constraints by their ability to anticipate k strategic iterations. The seller is able to outperform them in strategic thinking (i.e. has a level of $k+1$) due to superior access to data and computing power. Whether k is relatively low, as suggested by the empirical behavioral literature, or rather high turns out to crucially matter for our results.

We show that the higher consumers' level of sophistication, the *higher* the optimal price will be in the anonymous channel A. The intuition for this result is that consumers anonymize, at Stage 1 of the game, if their individual valuation for the product exceeds the expected price plus the anonymization cost. But when consumers decide about buying, at Stage 3, those anonymization costs are sunk. As the seller understands this fact, s/he increases the price above the one expected by consumers, namely to the expected price plus the anonymization cost. If the level of sophistication rises in the population, consumers are smart enough to expect this price increase. Hence, some of them—those with medium but not high WTP—preempt losses by choosing channel D instead of choosing Channel A at Stage 1. Responding to this changed behavior, the seller increases the price in the anonymized market even more, because s/he infers that only consumers with high WTP have chosen channel A at Stage 1.[8]

[3] Daughety and Reinganum (2010) also develop a model of a demand for privacy without assuming a direct preference for it, but study a public goods context and not a consumer market.

[4] Even Amazon.com shows prices to consumers that are not signed in (or have signed out).

[5] While most of this literature analyzes games with symmetric decisions (e.g. the beauty contest game (Nagel, 1995)), we will adapt the concept slightly to the asymmetric situation of our model where the seller has a different set of actions than the consumers.

[6] See Ho et al. (1998), Costa-Gomes et al. (2001), Crawford (2003), Camerer et al. (2004), Costa-Gomes and Crawford (2006), Crawford and Iriberri (2007a), and Goldfarb and Yang (2009), among others.

[7] What we call "unlimited strategic sophistication", is often referred to as "perfect rationality". However, players with limited strategic sophistication still act rationally given their (potentially wrong) beliefs, which is why we avoid the terms of "perfect" and "imperfect" rationality.

[8] Our model can be seen as employing a variant of the Diamond Paradox (Diamond, 1971). There, consumers with private valuations for a good produced by a monopolist have to incur some cost $s > 0$ to visit the store and be able to purchase the good, ultimately leading to a similar kind of reasoning and

Repeating this process, we show that, with any positive cost of anonymization, the anonymized market completely unravels for all sophistication levels $k \geq \bar{k}$, where $\bar{k}$ is a finite number. Hence, unlimited strategic sophistication is not a necessary condition for market unravelling. However, if consumers' k is sufficiently low, only a part of the market unravels and the anonymized sales channel can persist, serving consumers with high WTP. Among those who use the anonymous sales channel, some consumers suffer from net losses because prices turn out to be higher than expected. Consumers with very high WTP still get some surplus, however, because the unexpected price increase is less than what they are willing to pay.

Thereby, this model offers a micro-foundation for consumers' privacy choices: some consumers rationally use costly anonymization techniques even without an exogenous taste for privacy. Because a share of the anonymization cost could be interpreted as a fee that an intermediary can appropriate, this model also suggests that running an anonymous sales channel competing with a channel that tracks individuals and uses all personal data can be a profitable business model when consumers have limited strategic sophistication.

The model has important welfare results and policy implications. We show that consumer surplus is largest in the extreme case of costless anonymization, whereas the seller's profits are maximal if anonymization is prohibitively costly. Total welfare, however, is maximal at either of these extremes. The distributional consequences of both cases differ crucially: if anonymizing is free, consumers with high valuation anonymize and receive positive surplus, whereas all others do not anonymize and get perfectly discriminated against, leaving them zero consumer surplus. If anonymization is prohibitively costly, all consumers choose not to anonymize. Then all consumers buy the product for a personalized price and the seller appropriates all surplus.

A consumer-oriented policy maker should, hence, try to eliminate anonymization costs. For example, s/he could require marketers and online platforms to set anonymous shopping technologies as *default*. This would then require consumers to *opt in* to non-anonymous shopping instead of today's standard, where full tracking of consumers' choices is the default and a few providers offer *opt out* technologies. Consumers willing to reveal their characteristics to sellers (our model suggests these are those with low valuations) would log in to some service and receive the product for a price equal to their WTP. Consumers with higher valuations would stay in the (now default) anonymous channel and pay a higher price, but still retain some surplus.

Related Literature: First-degree (or perfect) price discrimination requires complete information of a seller about a specific consumer's WTP and was introduced into the economics literature by Pigou (1920). However, due to the very high information demand and the rather straightforward allocative and distributional implications, perfect price discrimination has not received a lot of scholarly attention and has been dismissed as a mere theoretical construct.

More prominent are models of "behavior-based price discrimination." Most of this literature focuses on third-degree price discrimination assuming that a seller learns about the WTP of a (re-)identifiable consumer after the first purchase. The idea is that, if a consumer bought a good at a certain price, the seller learns that this consumer's WTP must have weakly exceeded that price and raise the price for him/her in the future. If consumers anticipate this, they may adjust behavior in early periods and postpone purchases to avoid future price increases, or wait for future price cuts (Villas-Boas, 2004). In such cases, firms benefit from stricter privacy regulations if they lack commitment power to not increase prices after initial purchases (Taylor, 2004).[9]

However, lending support to the early conclusion of Odlyzko, "that in the Internet environment, the incentives towards price discrimination and the ability to price discriminate will be growing" (Odlyzko, 2003, p.365), online vendors and other retailers have already gone much further (see Footnote 2) and can approximate fully personalized prices more than ever. It has been shown empirically that "targeted advertising" techniques increase purchases (Luo et al., 2014), prices (Mikians et al., 2012), and sellers' profits (Shiller, 2013). Some consumers, however, feel repelled by this development and want to have control over their personal data back.[10] Many place a value on their privacy (Tsai et al., 2011).

The early theoretical literature about the economics of privacy, being based on the Chicago school argument that more information available to market participants increases the efficiency of markets, has underlined the negative welfare effects of hiding information from sellers (Posner, 1978; Stigler, 1980; Posner, 1981).[11] A lot of progress in our understanding has been made since then. Hermalin and Katz (2006, p.229) made clear: "With so many people making extreme claims in discussions of privacy and related public policy, and with so little understanding of the underlying economics, it is important to identify the fundamental forces clearly. A central fact is that, contrary to the Chicago School argument, the flow of information from one trading partner to the other can reduce ex post trade efficiency when the increase in information does not lead to symmetrically or fully informed parties."

market unravelling. However, in contrast to Diamond (1971), we entertain the notion that consumers may not be fully strategically sophisticated, and also assume that the seller knows each consumer's valuation in the default situation.

[9] For an overview of this strand of literature, see Fudenberg et al. (2006).

[10] Goldfarb and Tucker (2012) study three million observations between 2001 and 2008 and find that refusals to reveal their income in an online survey have risen over time. Tucker (2014) finds in a field experiment that, when Facebook gave users more control over their personally identifiable information, users were twice as likely to click on personalized ads.

[11] Even earlier, Warren and Brandeis (1890) study privacy as "right to be let alone", a point later discussed by Varian (1997) in the context of annoyance from telemarketing.

A related issue are the choices of firms that own personal information about consumers and can decide to disclose it to another firm (Taylor, 2004; Acquisti and Varian, 2005; Casadesus-Masanell and Hervas-Drane, 2015). In interactions between an upstream and a downstream firm for whose products consumers' WTP is positively correlated, the upstream firm will maintain full privacy of its customers if conditions on the upstream firm's preferences about the downstream firm and on the downstream relationship itself are met (Calzolari and Pavan, 2006). However, if one condition is not met, the upstream firm can find it optimal to disclose its customer list to the downstream firm (sometimes even for free), which need not be negative to consumers but could still yield a Pareto improvement (Calzolari and Pavan, 2006).

Core questions studied in these papers are the welfare consequences of privacy or disclosure, and who should own the property rights of consumers' personal data (Hermalin and Katz, 2006). The answers given have been ambiguous and depend on the specific application. Recently, the focus has shifted more towards consumers' privacy choices (Conitzer et al., 2012; Montes et al., 2018; Ichihashi, 2020) and the role of platform intermediaries (de Cornière and de Nijs, 2016; Bergemann and Bonatti, 2015). Akin to our model, Belleflamme and Vergote (2016) also study price discrimination by a monopolist with an option for consumers to hide again, reaching results similar in spirit, but structurally different to the ones presented in Section 4. For a thorough overview of the growing literature on the economics of privacy, see Acquisti et al. (2016).

With few exceptions, however, cognitive constraints of consumers have not been incorporated in theoretical studies of markets driven by big data. Taylor (2004), Acquisti and Varian (2005), and Armstrong (2006) assume the existence of a group of unlimitedly sophisticated consumers and a group of naïve consumers. The latter do not foresee that they may want to trade in the future again and, because of this, ignore the negative effects of disclosing personal data. Hence, if consumers are naïve, a seller may oppose stricter regulations as no commitment device is needed (Taylor, 2004). In our model, we allow for a more nuanced analysis of consumer sophistication.[12]

Benndorf et al. (2015) conduct lab experiments studying the revelation of private information about workers' types. Their study is based on the introduction of certification as a solution to adverse selection problems by Viscusi (1978). They find unraveling effects, which were theoretically predicted by Hermalin and Katz (2006), but also identify a tendency of under-revelation, which is consistent with level-k models. A key difference to our setup is that theses papers assume private information as default, which can be revealed costly, e.g. by certification. Our setup is reversed. Due to technological progress, we assume that a seller has perfect information about consumers' WTP unless they engage in costly anonymization.

The remainder of the paper is organized as follows. In Section 2, we construct a model, which is analyzed in Section 3. Section 4 studies welfare and payoff consequences of changing the level of sophistication k and the anonymization cost s. Section 5 is dedicated to alternative model specifications, covering heterogeneous cost of anonymization, increasing competition, and the same level of sophistication for all players. Section 6 concludes. Proofs are in the Appendix. Further extensions and robustness checks are in an Online Appendix.

2. Model

We consider an economy where a monopolistic seller of a single consumption good faces a unit mass of atomistic consumers who can buy at most one unit of the good and cannot resell it to each other.[13] Abstracting from potential fixed costs, we assume that the monopolist can produce the good at constant marginal cost which we normalize to $c = 0$.[14] Consumers have heterogeneous valuations v for the good, where $v \sim \mathcal{U}[0, 1]$ and can approach the seller in two different ways: directly (referred to as channel D) or after making use of an anonymization technique (channel A). Consumers choosing channel D incur no cost and the seller perfectly knows their individual valuation. Consumers choosing channel A, on the other hand, incur cost $s \geq 0$ and their individual valuation is hidden from the seller.[15] We assume that consumers do not have any exogenous taste for privacy and that they choose direct channel D in case of indifference between both channels.

After each consumer has made a choice between the channels, the seller sets prices based on the available information.[16] In channel D the seller can set personalized prices $p_i(v)$ conditional on each consumer's valuation. However, in channel A, due to the anonymization technique consumers used, the seller can only set a uniform price p_A for all consumers.

Finally, consumers decide whether they want to buy the good at the price the seller posted for them.[17] In case of indifference, we assume they buy the product. Outside options yield zero payoff (except for costs incurred within the game before opting out). The timing of the model is summarized as follows:

[12] The need to include cognitive constraints into economic models of privacy is spurred by empirical findings about the so-called *privacy paradox* (Norberg et al., 2007), an apparent discrepancy between people's privacy attitudes and their actual behavior. For a summary of the privacy paradox literature, see Acquisti et al. (2016).

[13] We discuss the case of monopolistic competition in Section 5.2.

[14] Generalized results for $c \geq 0$ do not affect the main results but decrease legibility.

[15] Consumers may need to pay for or install privacy-protective software, experience lower connection speed due to encrypted transmissions, or otherwise increase transaction cost (e.g. by shopping offline with additional costs if done by cash payments).

[16] This timing of stages is natural, because the seller sets prices aided by an algorithm that can best-respond to all potential consumer decisions.

[17] Note that exploring the price in the anonymous channel requires the spending of s. Hence, even if consumers could go back to the direct channel D for free, they would not recover s.

- Stage 1 (*Anonymizing*): Consumers choose channel D or channel A and incur costs of 0 or s, respectively. Indifferent consumers choose channel D.
- Stage 2 (*Pricing*): The seller sets prices $p = \{p_i(v), p_A\}$, where $p_i(v)$ are personalized prices in channel D, and p_A is the uniform price in channel A.
- Stage 3 (*Buying*): Consumers decide whether to buy the good for the offered price. Indifferent consumers are assumed to choose buying the good.

The distribution of v (and hence the demand function), the monopolist's cost structure (and hence the supply function), the cost for anonymization s as well as the timing of the game are common knowledge among all players.

Explicitly modeling consumers' cognitive constraints, we assume that all consumers have the same limited level of strategic sophistication, denoted by $k \in \mathbb{Z}_0^+$.[18] The seller, however, outperforms them in terms of sophistication and has a level of $k+1$. As in Nagel (1995), players with a level of $k > 0$ will generally act as if they believe that all other players had a level of strategic sophistication exactly one level below their own level. However, Nagel (1995) considers a setting where all players are symmetric and have the same action sets. As our model has one player (the seller) whose action set differs from everyone else's, and whose best response is therefore different, we adapt the concept slightly.

While we maintain that consumers believe that all other consumers are one level less sophisticated, we deviate in assuming that consumers expect the seller to share their level of sophistication. More formally, consumers form the beliefs $\mathbb{E}_i(k_{j \neq i}) = k_i - 1 = k - 1$ for j being a consumer and $\mathbb{E}_i(k_S) = k_i = k$ for S being the seller. Thus, consumers implicitly think of the seller as responding optimally to their own belief about the sophistication of all other consumers. This assumption is in turn based on the atomistic nature and the resulting insignificance of any individual consumer for the seller's choice. The seller, whose $k_S = k + 1$, however, forms the belief $\mathbb{E}_S(k_{j \neq S}) = k_S - 1 = k$ for j being any consumer, in line with Nagel (1995).

Their cognitive limitations in belief formation notwithstanding, players still act rationally by pursuing strategies which maximize their utility *given* their beliefs. Hence, we can solve the game by backward induction, but our solution concept differs from a Perfect Bayesian Equilibrium (which would be appropriate if all players had unlimited strategic sophistication) because it does not need to be the case that all expectations about others' strategies are eventually confirmed. Notably, we allow consumers' beliefs about the prices in both channels to not coincide with the prices the seller eventually sets. Hence, we do not *impose* that $\mathbb{E}(p|D) = (p|D)$ and $\mathbb{E}(p|A) = (p|A)$, where $(p|D)$ and $(p|A)$ denote the price after having chosen channel D or channel A, respectively. However, we do impose that players restrict their beliefs about possible prices to the support of the distribution of v, i.e. $\mathbb{E}(p|D) \in [0, 1]$, and $\mathbb{E}(p|A) \in [0, 1]$.

A consumer's strategy is a mapping from their own valuation for the good v, their own level of strategic sophistication k, and the exogenous parameter s to the action space $C \times B$, where $C = \{Channel\ D, Channel\ A\}$ denotes the set of choices in the anonymizing stage (Stage 1) and $B = \{(Buying|p), (NotBuying|p)\}$ denotes the set of choices in the buying stage (Stage 3). $p = p_i(v)$ after having chosen channel D and $p = p_A$ after having chosen channel A.

The seller's strategy is a mapping from his/her level of sophistication $k+1$, and the exogenous parameter s to a set of prices $p = \{p_i(v), p_A\}$, where $p_i(v)$ are personalized prices s/he can condition on the available knowledge about individual consumers in channel D, and p_A is a uniform price for all consumers in channel A.

The game is solved by backward induction. As models with level-k thinking are best solved recursively, the analysis starts with the case where consumers have a strategic sophistication level of $k = 0$ and form a so-called naïve belief. The seller, having a level of $k = 1$, believes (correctly) that all consumers have a level of $k = 0$. Thus, his/her level-1 best response is also objectively optimal. In later parts of the analysis, consumers have a level of $k > 0$ and believe that all other consumers have a level of $k - 1$, but that the seller employs the level-k best response. The seller, with a level of $k + 1$, will again be the only one with an objectively correct belief and his/her $k + 1$-level response is again objectively optimal.

3. Analysis

Stage 3 – Buying: A utility-maximizing consumer buys the good if the price does not exceed his/her valuation, i.e. if, and only if,

$$v \geq p \in \{p_i(v), p_A\}. \tag{1}$$

After choosing channel D, the price will be an individualized price $p_i(v)$, and after choosing channel A, s/he will receive the same uniform price p_A as all other consumers who have chosen channel D.

Stage 2 – Pricing: A profit-maximizing seller sets individual prices p_i for all consumers in channel D (denoted by set $\mathcal{C}_D$) and one uniform price p_A for all anonymized consumers in channel A (denoted by set $\mathcal{C}_A$). Knowing v precisely for all consumers in $\mathcal{C}_D$, the seller sets

[18] In Appendix C, we analyze the case with two levels of strategic sophistication. We show that, if some consumers have $k = 0$ and some have $k = 1$, the qualitative results are robust compared to our main model specification.

$$p_i^*(v) = v \text{ for all } i \in \mathcal{C}_D. \tag{2}$$

Despite being uninformed about individual valuations v of consumers in $\mathcal{C}_A$, the seller can infer which consumers are in $\mathcal{C}_A$ due to his/her higher level of strategic sophistication and set p_A accordingly. We therefore analyze consumers' general Stage 1 behavior first in order to inform the seller's pricing decision in channel A.

Stage 1 – Anonymizing: Consumers use the anonymization technique of channel A if the expected utility of doing so exceeds the expected utility of the direct channel D, i.e. if, and only if, $\mathbb{E}(u_i(A)) > \mathbb{E}(u_i(D))$, where

$$\mathbb{E}(u_i(D)) = \max\{v - \mathbb{E}(p|D), 0\}, \tag{3}$$

$$\mathbb{E}(u_i(A)) = \max\{v - \mathbb{E}(p|A) - s, -s\}. \tag{4}$$

The first value in each set in (3) and (4) reflects the expected payoff a consumer receives if s/he buys the product at Stage 3. The second value reflects the payoff of subsequently choosing not to buy the product. Although consumers might be limited in their strategic sophistication, they nonetheless understand the nature of the two channels, i.e. they realize that the seller has no incentive to decrease the price below their valuation in channel D and that the seller can only ask for a uniform price in channel A. Hence, consumers, irrespective of their level of strategic sophistication, form the price expectation for channel D

$$\mathbb{E}(p|D) = p_i^*(v) = v, \tag{5}$$

correctly expecting to be left with no surplus in channel D. Their exact expectation for the uniform price in channel A, though, still depends on their level of strategic sophistication. Hence, we only substitute $\mathbb{E}(p|A) = \mathbb{E}(p_A)$ yielding:

$$\mathbb{E}(u_i(D)) = \max\{v - v, 0\} = 0, \tag{6}$$

$$\mathbb{E}(u_i(A)) = \max\{v - \mathbb{E}(p_A) - s, -s\}. \tag{7}$$

This implies that consumers choose channel A if, and only if,

$$\max\{v - \mathbb{E}(p_A) - s, -s\} > 0. \tag{8}$$

Because $s \geq 0$, this can only hold if $v > \mathbb{E}(p_A) + s \equiv \hat{v}$, where $\hat{v}$ denotes the endogenous threshold dividing the population of consumers into $\mathcal{C}_D$ and $\mathcal{C}_A$.

Lemma 1 *(Anonymization Threshold). There exists a threshold $\hat{v} = \mathbb{E}(p_A) + s$ denoting the valuation of a consumer who is indifferent between both channels at Stage 1. Consumers with $v > \hat{v}$ choose channel A; consumers with $v \leq \hat{v}$ choose channel D, i.e. $\mathcal{C}_D = [0, \hat{v}]$ and $\mathcal{C}_A = (\hat{v}, 1]$.*

Stage 2 – Pricing (revisited): Having a higher level of strategic sophistication than the consumers, the seller correctly infers $\hat{v}$ and hence knows that $\mathcal{C}_A = (\hat{v}, 1]$. As the seller further anticipates that consumers will buy the product at Stage 3 if, and only if, $v \geq p_A$, s/he can easily infer demand $q_A(p_A)$ in channel A:

$$q_A(p_A) = \begin{cases} 0 & \text{if } p_A > 1, \\ 1 - p_A & \text{if } 1 \geq p_A > \hat{v}, \\ 1 - \hat{v} & \text{if } \hat{v} \geq p_A. \end{cases} \tag{9}$$

Charging $p_A = \hat{v}$ dominates all prices $p'_A < \hat{v}$ because any price below $\hat{v}$ decreases profits per unit sold without an increase in quantity to counter the loss. Thus, by setting $p_A = \hat{v}$, the seller can guarantee profits from channel A of:

$$\pi_A(\hat{v}) = q_A(\hat{v})\hat{v} = (1 - \hat{v})\hat{v} = \hat{v} - \hat{v}^2. \tag{10}$$

However, the seller could also charge a price $p_A > \hat{v}$, depending on where $\hat{v}$ lies in relation to the globally profit-maximizing price $p_M = \frac{1}{2}$ that the seller would set if only the anonymous channel A existed and the seller was unable to engage in price discrimination at all. Fig. 1 shows the three different cases for the location of the anonymization threshold $\hat{v}$ compared to p_M:

(a) The anonymization threshold is below the monopoly price ($\hat{v} < p_M$).
(b) The anonymization threshold is equal to the monopoly price ($\hat{v} = p_M$).
(c) The anonymization threshold is above the monopoly price ($\hat{v} > p_M$).

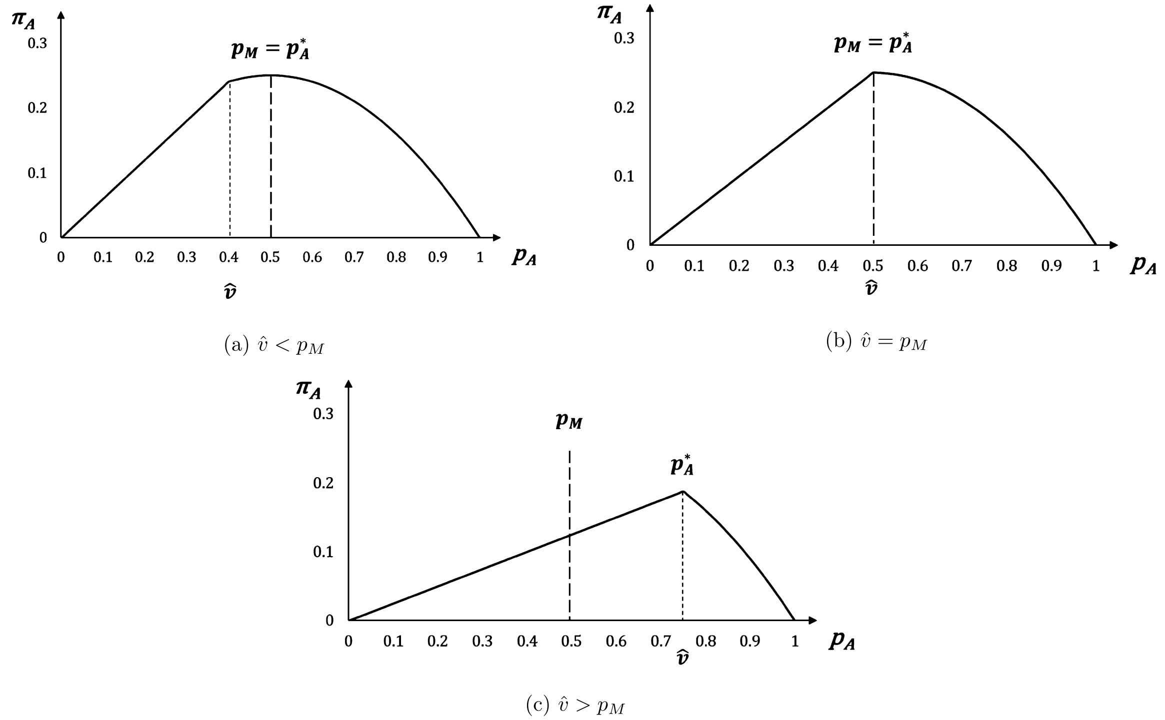

Fig. 1. Profits in channel A for different locations of $\hat{v}$.

In cases (a) and (b), the globally profit-maximizing price p_M is located within $\mathcal{C}_A$ and hence remains the optimal price to set. The only consumers that are not in $\mathcal{C}_A$ are those that the seller would not have served even if they had anonymized themselves. Only in case (c), where the globally profit-maximizing price p_M is not located within $\mathcal{C}_A$ anymore, any price below the anonymization threshold $\hat{v}$ is at least dominated by setting the price equal to $\hat{v}$. The seller also has no incentive to raise the price above $\hat{v}$ as profits are strictly decreasing to either side of the global maximum at p_M due to the strict concavity of the profit function. Hence, in this case the optimal price p_A^* is equal to $\hat{v}$.

Lemma 2 *(Optimal Pricing Strategy). The optimal strategy of the seller consists of a set of prices $\{p_i^*(v), p_A^*\}$ in channel D and A, respectively, where $p_i^*(v) = v$ and $p_A^* = \max\{\hat{v}, p_M = \frac{1}{2}\}$.*

This implies that the seller sets a higher price than consumers had expected (unless $s = 0$):

$$p_A^* \geq \hat{v} = \mathbb{E}(p_A) + s > \mathbb{E}(p_A). \tag{11}$$

With unlimited strategic sophistication, $\mathbb{E}(p_A) = p_A^*$ would be required in equilibrium for all values of s, leading to a contradiction. Because only beliefs about off-equilibrium paths can be wrong in a Perfect Bayesian Equilibrium, we conclude that if all players had unlimited strategic sophistication, channel A would remain unused for any $s > 0$.

With limited strategic sophistication, though, such discrepancy is possible. This is because s will be a sunk cost for consumers at Stage 3, which the seller can exploit by increasing the price by exactly s, compared to their expectations. Consumers, due to their limited strategic sophistication, cannot anticipate the seller's strategic response.

For $s = 0$, the difference between consumers' expectations and seller's optimal price disappears. Because anonymization is then available for free, failing to anticipate the seller's reaction to their anonymization decision does not lead consumers to incur an upfront sunk cost. While obviously having an effect on their payoffs, this special case does not lead to a change in consumers' or the seller's behavior, as analyzed so far.

Stage 1 – Anonymizing (revisited): The last missing piece to fully characterize the solution is the formation of consumers' price expectations in channel A, $\mathbb{E}(p_A)$, in Stage 1. We determine these recursively and will start with consumers with sophistication $k = 0$, which are referred to as "naïve" consumers: they naïvely expect the monopolist to engage in regular monopoly pricing in channel A, i.e. $\mathbb{E}_0(p_A) = p_M$.[19] Thereby, they ignore the fact that the very choice of channel A might

[19] The alternative starting point assumption of random play for naïve players, as it is used in some other applications of level-k thinking, is discussed in Appendix E. Depending on the specifics of the distribution, this can result in quantitative but no qualitative changes to the results of our main analysis.

be signaling a high valuation to the seller. For channel D, we have already assumed that naïve consumers foresee perfect price discrimination as it does not require iterative thinking about other consumers.

Lemma 3 *(Solution with level-0). For any non-prohibitively high anonymization cost $s \geq 0$ and with strategically "naïve" consumers ($k = 0$) there is a unique solution with the following characteristics:*

- *Consumers form the 0-beliefs $\mathbb{E}_0(p_D) = p_i^*(v)$ and $\mathbb{E}_0(p_A) = p_M = \frac{1}{2}$.*
- *Consumers anonymize if, and only if, $v > \hat{v}_0 = \mathbb{E}_0(p_A) + s = p_M + s$. Hence, $\mathcal{C}_D = [0, \hat{v}_0]$ and $\mathcal{C}_A = (\hat{v}_0, 1]$.*
- *The seller forms the 1-belief about the anonymization threshold $\mathbb{E}_1(\hat{v}_0) = \hat{v}_0$ and infers correctly that, therefore, $\mathcal{C}_D = [0, \hat{v}_0]$ and $\mathcal{C}_A = (\hat{v}_0, 1]$.*
- *The seller sets prices $p_i^*(v) = v$ and $p_{A_0}^* = \hat{v}_0 = p_M + s$.*
- *All consumers buy the product at the price offered.*

Lemma 3 shows that naïve consumers in channel A pay a premium of s as compared to their expectations, $p_{A_0}^* - \mathbb{E}_0(p_A) = s$.[20] Consumers do not anticipate that the seller can infer that only consumers with a valuation of at least $p_M + s$ choose the anonymous channel. Given this lower bound on the valuations in $\mathcal{C}_A$, the seller can ignore that anonymized consumers spent s on top, and extract the lower bound's full valuation. This divergence between expected price and realized price, in turn, informs us about the way in which consumers form their price expectations for higher levels of strategic sophistication, $k > 0$.

If instead of being naïve, all consumers are capable of one iteration of strategic reasoning ($k = 1$), they anticipate that the seller's best response to the 0-belief of naïve consumers is $p_{A_0}^* = p_M + s$. Recall, that consumers with $k > 0$ form the beliefs $\mathbb{E}_i(k_{j\neq i}) = k - 1$ for all other consumers and $\mathbb{E}_i(k_S) = k$ for the seller. Therefore, consumers with $k = 1$ assume that the seller responds optimally to a population of consumers with $k = 0$, and adjust their price expectation. As consumers are atomistic, their own anonymization choice is inconsequential for the seller's best response. Accordingly, they form the 1-belief $\mathbb{E}_1(p_A) = p_{A_0}^* = p_M + s$ leading to $\hat{v}_1 = p_M + 2s$, to which the seller's actual best response is $p_{A_1}^* = p_M + 2s$ (as reasoned above). This, in turn, would be the expected price in channel A by consumers with a level of $k = 2$, thus forming the 2-belief $\mathbb{E}_2(p_A) = p_{A_1}^* = p_M + 2s$, and so forth. More generally, we can write $\mathbb{E}_k(p_A) = p_{A_{k-1}}^*$ for all $k > 0$, which in combination with $\mathbb{E}_0(p_A) = p_M$ leads to:

$$\mathbb{E}_k(p_A) = p_M + ks, \tag{12}$$

$$p_{A_k}^* = p_M + (k+1)s = \hat{v}_k. \tag{13}$$

At every additional level of strategic sophistication, consumers incorporate the sunk cost once more than before, which induces the seller to incorporate s in the price once more, too. In the special case of free anonymization, i.e. with $s = 0$, this does not change the situation compared to the solution for naïve consumers ($ks = 0$ for all $k < \infty$). However, for all cases where $s > 0$, $\hat{v}_k$ increases and $\mathcal{C}_A$ shrinks in size as k increases. This implies, that the more strategically sophisticated the consumer population is, the fewer consumers anonymize. When $\hat{v}_k$ matches or exceeds the highest valuation, no consumer does so anymore. Channel A remains unused and the anonymous market breaks down completely. We denote the threshold level of strategic sophistication from which onwards this is the case by $\bar{k}$ and define:

$$\bar{k} \equiv \min\{k \in \mathbb{Z}_0^+ | \hat{v}_k \geq 1\}. \tag{14}$$

The inequality in (14) can hold with equality as indifferent consumers opt for channel D, by assumption. Using (13) in (14) and solving for $\bar{k}$ yields:

$$\bar{k} \geq \frac{1}{2s} - 1 \Rightarrow \bar{k} = \left\lceil \frac{1}{2s} - 1 \right\rceil \text{for } s > 0. \tag{15}$$

This shows that channel A breaks down at a *finite* level of strategic sophistication, in turn implying that unlimited strategic sophistication, while *sufficient*, is *not necessary* for this breakdown.[21]

Lemma 4 *(Usage of Channel A). For any strictly positive but non-prohibitively high anonymization cost $s > 0$ the anonymous channel is used if, and only if, consumers are not* too *strategically sophisticated, i.e. if $k < \bar{k} = \lceil \frac{1}{2s} - 1 \rceil$. If anonymization is costless, i.e. $s = 0$, the anonymous channel is used for all levels of sophistication.*

In many cases, it would lead to one additional reasoning step preceding the model as presented here. However, as this step is often non-linear, it would increase notational burden by adding more cases and likely distract from the main insights of the model. Therefore, we opted for expositional simplicity.

[20] Consequently, this premium is zero if anonymization is available for free ($s = 0$). In this special case, the resulting behavior of naïve consumers is the same that would ensue under unlimited strategic sophistication. This is purely driven by the assumption we made about the price expectation of naïve consumers: $\mathbb{E}_0(p_A) = p_M$. For alternative starting-point assumptions, see Appendix E.

[21] This insight is a contribution of the model at hand compared to the Diamond Paradox, which studies agents with unlimited strategic sophistication (see footnote 8).

Lemma 5 *(Solution with Level-k and Costless Anonymization). For anonymization cost $s = 0$ and any finite level of consumer sophistication k there is a unique solution with the following characteristics:*

- *Consumers form the k-beliefs $\mathbb{E}_k(p_D) = p_i^*(v)$ and $\mathbb{E}_k(p_A) = p_M = \frac{1}{2}$.*
- *Consumers anonymize if, and only if, $v > \hat{v}_k = p_M$. Hence, $\mathcal{C}_D = [0, p_M]$ and $\mathcal{C}_A = (p_M, 1]$.*
- *The seller forms the k+1-belief about the anonymization threshold $\mathbb{E}_{k+1}(\hat{v}_k) = p_M$ and infers correctly that, therefore, $\mathcal{C}_D = [0, p_M]$ and $\mathcal{C}_A = (p_M, 1]$.*
- *The seller sets $p_i^*(v) = v$ and $p_{A_0}^* = p_M$.*
- *All consumers buy the product at the price offered.*

That channel A breaks down at a finite level of sophistication $\bar{k}$ has consequences for the belief formation of consumers when $k > \bar{k}$. While for $k \le \bar{k}$ belief formation according to (12) does not violate that all players restrict price expectations to $p \in [0, 1]$, this is not the case for $k > \bar{k}$. Denoting any level of consumer sophistication $k > \bar{k}$ by $\bar{k}^+$, we specify beliefs $\mathbb{E}_{\bar{k}^+}(p_A)$ to meet this condition in (16). In line with Lemma 4, any belief $\mathbb{E}_{\bar{k}^+}(p_A)$ also has to render the choice of channel D optimal for consumers regardless of their valuation, leading to (17):

$$\mathbb{E}_{\bar{k}^+}(p_A) \in [0, 1] \qquad \Rightarrow \mathbb{E}_{\bar{k}^+}(p_A) \le 1, \tag{16}$$

$$\hat{v}_{\bar{k}^+} = \mathbb{E}_{\bar{k}^+}(p_A) + s \ge 1 \qquad \Rightarrow \mathbb{E}_{\bar{k}^+}(p_A) \ge 1 - s. \tag{17}$$

Both conditions are satisfied for any belief $\mathbb{E}_{\bar{k}^+}(p_A) \in [1 - s, 1]$. Hence, multiple beliefs are possible when $k > \bar{k}$, but all imply that channel A remains unused. For any level of consumer sophistication where channel A remains unused (including $k = \bar{k}$), the seller forms k+1-beliefs $\mathbb{E}_{k+1}(\mathcal{C}_D) = [0, 1]$ and $\mathbb{E}_{k+1}(\mathcal{C}_A) = \emptyset$. Therefore, setting p_A is an action on an unreached branch of the game tree and the seller can set any price $p^*_{A_{\bar{k}^+}} \in [0, 1]$ (restricted only by the support of the distribution of v). We summarize the analysis for general level-k and $s > 0$ in Proposition 1.

Proposition 1 *(Solution with Level-k). For any non-prohibitively high anonymization cost $s > 0$ and any finite level of consumer sophistication k it holds that:*

1. *If $k \le \bar{k} = \lceil \frac{1}{2s} - 1 \rceil$, there is a unique solution with the following characteristics:*
   - *Consumers form k-beliefs $\mathbb{E}_k(p_D) = p_i^*(v)$ and $\mathbb{E}_k(p_A) = p_M + ks = \frac{1}{2} + ks$.*
   - *Consumers anonymize if, and only if, $v > \hat{v}_k = p_M + (k + 1)s$. Hence, $\mathcal{C}_D = [0, \hat{v}_k]$ and $\mathcal{C}_A = (\hat{v}_k, 1]$ (where $\mathcal{C}_A = \emptyset$ if $k = \bar{k}$).*
   - *The seller forms the k+1-belief about the anonymization threshold $\mathbb{E}_{k+1}(\hat{v}_k) = \hat{v}_k$ and infers correctly that, therefore, $\mathcal{C}_D = [0, \hat{v}_k]$ and $\mathcal{C}_A = (\hat{v}_k, 1]$.*
   - *The seller sets $p_i^*(v) = v$ and* $\begin{cases} p^*_{A_k} = \hat{v}_k = p_M + (k+1)s & \text{if } k < \bar{k}, \\ p^*_{A_{\bar{k}}} \in [0, 1] & \text{if } k = \bar{k}. \end{cases}$
   - *All consumers buy the product at the price offered.*
2. *If $k > \bar{k} = \lceil \frac{1}{2s} - 1 \rceil$, there are multiple solutions with the following characteristics:*
   - *Consumers form k-beliefs $\mathbb{E}_{\bar{k}^+}(p_D) = p_i^*(v)$ and $\mathbb{E}_{\bar{k}^+}(p_A) \in [1 - s, 1]$.*
   - *No consumer anonymizes. Hence, $\mathcal{C}_D = [0, 1]$ and $\mathcal{C}_A = \emptyset$.*
   - *The seller forms the k+1-belief about the anonymization threshold $\mathbb{E}_{k+1}(\hat{v}_k) = 1$ and infers correctly that, therefore, $\mathcal{C}_D = [0, 1]$ and $\mathcal{C}_A = \emptyset$.*
   - *The seller sets $p_i^*(v) = v$ and any $p^*_{A_k} \in [0, 1]$.*
   - *All consumers buy the product at the price offered.*

In Proposition 1. 1, consumers with high valuations ($v > \hat{v}_k$) choose the anonymous channel A, consumers with low valuations ($v \le \hat{v}_k$) choose the direct channel D and are perfectly price discriminated against. Consumers with very low valuations ($v < p_M$) choose the direct channel D irrespectively of k and s as they cannot possibly hope to get a uniform price that is affordable to them via channel A. These are the consumers that are not served in monopolistic markets without possibility for perfect price discrimination. The multiplicity of solutions in Proposition 1. 2 depends only on the multiplicity of possible beliefs about unreached paths of the game. But all solutions feature the same behavior, where no consumer anonymizes and the seller charges individualized prices $p_i^*(v)$ to all.[22]

4. Welfare

Proposition 1 crucially depends on k and s. Hence, changes in both parameters have consequences for consumer surplus (CS), profits (π), and total welfare (W), for which we employ the customary definition $W = CS + \pi$. We focus on

[22] The iterative reasoning structure underlying the belief formation in our model also lends itself to a temporal interpretation. See Appendix F for the adjustments such an interpretation would imply.

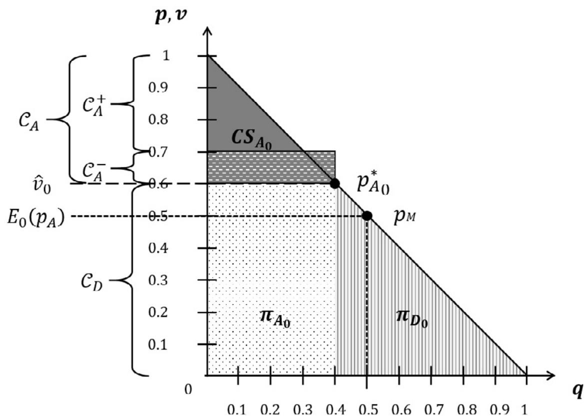

(a) $k = 0$

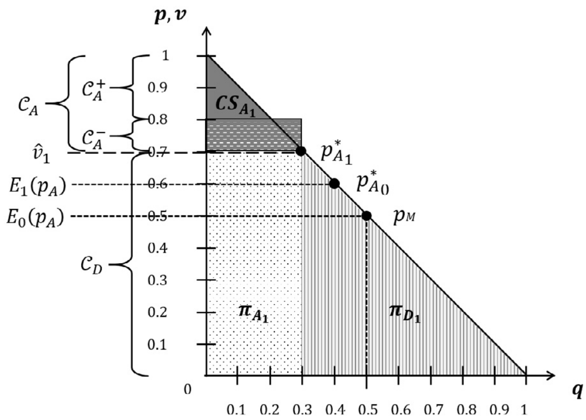

(b) $k = 1$

Fig. 2. Welfare analysis with $s = 0.1$.

the aggregate quantities and comparative statics in this section and derive consumer surplus and profits for each channel separately in Appendix A.1 and A.2.

Consumer surplus in channel A (CS_{A_k}) consists of two parts: the consumption utility at Stage 3 ($CS^+_{A_k}$) and the anonymization cost at Stage 1 ($CS^-_{A_k}$). In Fig. 2, $CS^+_{A_k}$ corresponds to the solid grey (upper) triangle, whereas the dashed rectangle that partially overlaps this triangle represents $CS^-_{A_k}$. Only some consumers in channel A end up with positive net surplus (those in $\mathcal{C}^+_A = [\hat{v}_k + s, 1]$), whereas others end up with negative net surplus (those in $\mathcal{C}^-_A = (\hat{v}_k, \hat{v}_k + s)$). Due to perfect price discrimination, consumer surplus in channel D (CS_{D_k}) is zero. The seller's profits in channel A (π_{A_k}) are represented by the dotted white rectangle, whereas profits in channel D (π_{D_k}) are depicted by the vertically striped (lower right) triangle. Finally, comparing Fig. 2a and Fig. 2b illustrates the effects of increasing k by one level.

4.1. Consumer surplus, profits, and welfare

We define $CS_k \equiv CS_{D_k} + CS_{A_k}$, $\pi_k \equiv \pi_{D_k} + \pi_{A_k}$, and $W_k \equiv CS_k + \pi_k$ (see Appendix A.1 and A.2 for details). This yields the following results:

$$CS_k = \frac{1}{2}(1-\hat{v}_k)^2 - (1-\hat{v}_k)s = \frac{1}{8} - \frac{1}{2}(k+2)s + \frac{1}{2}(k^2+4k+3)s^2, \tag{18}$$

$$\pi_k = \frac{1}{2}\hat{v}_k^2 + (1-\hat{v}_k)\hat{v}_k = \frac{3}{8} + \frac{1}{2}(k+1)s - \frac{1}{2}(k^2+2k+1)s^2, \tag{19}$$

$$W_k = \frac{1}{2} - (1-\hat{v}_k)s = \frac{1}{2} - \frac{1}{2}s + (k+1)s^2. \tag{20}$$

Like total consumer surplus and profits, total welfare can be identified graphically in Fig. 2. The first term in (20), $\frac{1}{2}$, corresponds to the whole area between the demand curve and the marginal cost curve (recall that $c = 0$), while the second term, $(1 - \hat{v}_k)s$, corresponds to the dashed rectangle. Although the market outcome of Stage 3 is efficient, because every consumer buys the product, total welfare is reduced by the cost of consumers' anonymization behavior as long as $\hat{v}_k < 1$, i.e., $k < \bar{k}$, and as long as $s > 0$. For $s > 0$ and $k \geq \bar{k}$, or when $s = 0$, a fully efficient outcome ensues. We prove Proposition 2 in Appendix A.3.

Proposition 2 *(Welfare). For any strictly positive but non-prohibitively high anonymization cost $s > 0$, and any finite level of consumer sophistication k, aggregated consumer surplus (CS_k), profits (π_k), and welfare (W_k) exhibit the following characteristics:*

- *$CS_k > 0$ for $k < \bar{k}_{CS}$, $CS_k \leq 0$ for $\bar{k}_{CS} \leq k < \bar{k}$, and $CS_k = 0$ for $k \geq \bar{k}$, where $\bar{k}_{CS} = \lceil \frac{1}{2s} - 3 \rceil$ and $\bar{k} - \bar{k}_{CS} = 2$.*
- *$\pi_k > 0$ for $k < \bar{k}$, and $\pi_k = W_k$ for $k \geq \bar{k}$.*
- *$W_k > 0$ for $k < \bar{k}$, and $W_k = W^*$ for $k \geq \bar{k}$, where $W^* = \frac{1}{2}$ is the first-best outcome.*

If anonymization is costless, i.e. $s = 0$, consumer surplus (CS_k), profits (π_k), and welfare (W_k) are given by: $CS_k(s=0) = \frac{1}{8}$, $\pi_k(s=0) = \frac{3}{8}$, and $W_k(s=0) = W^* = \frac{1}{2}$ for any finite level of consumer sophistication k and $\bar{k}_{CS}$ is undefined.

4.2. Comparative statics for k

In the comparative statics analysis the discreteness of k is taken into account by calculating changes as differences rather than derivatives. Given that behavior is unaffected by changes in k when $s = 0$, we focus this analysis on $s > 0$. Additionally, due to the potential non-linearity when increasing k from $\bar{k}-1$ to $\bar{k}$, these differences only hold for $k < \bar{k}-1$.[23]

For $k < \bar{k}-1$:

$$\Delta CS_k \equiv CS_{k+1} - CS_k = -\left(1-\hat{v}_{k+1}\right)s + \frac{s^2}{2} = -\frac{s}{2} + \frac{2k+5}{2}s^2, \tag{21}$$

$$\Delta \pi_k \equiv \pi_{k+1} - \pi_k = \left(1-\hat{v}_k\right)s + \frac{s^2}{2} = \frac{s}{2} + \frac{2k+3}{2}s^2, \tag{22}$$

$$\Delta W_k \equiv W_{k+1} - W_k = s^2 = s^2. \tag{23}$$

While the first term in (21) stems from the reduction of consumer surplus at Stage 3, the second term captures the gain from fewer consumers incurring the anonymization cost. When moving from Fig. 2a to Fig. 2b, the first effect is represented by the shrinking dark grey triangle, and the second effect by the shrinking dashed rectangle.[24] We prove Proposition 3 in Appendix A.4.

Proposition 3 *(Effects of Changing Consumer Sophistication). Raising the level of strategic sophistication of consumers from k to $k+1$ has the following effects on consumer surplus, profits, and welfare (ceteris paribus):*

- $\Delta CS_k < 0$ *for* $k < \bar{k}_{\Delta CS}$, $\Delta CS_k \geq 0$ *for* $\bar{k}_{\Delta CS} \leq k < \bar{k}$, *and* $\Delta CS_k = 0$ *for* $k \geq \bar{k}$, *where* $\bar{k}_{\Delta CS} = \left\lceil \frac{1}{2s} - \frac{5}{2} \right\rceil$ *and* $\bar{k} - \bar{k}_{\Delta CS} \in \{1, 2\}$.
- $\Delta \pi_k > 0$ *for* $k < \bar{k}$, *and* $\Delta \pi_k = 0$ *for* $k \geq \bar{k}$.
- $\Delta W_k > 0$ *for* $k < \bar{k}$, *and* $\Delta W_k = 0$ *for* $k \geq \bar{k}$.

Corollary 1 *(Individual Surplus). As long as consumers are not* too *sophisticated ($k < \bar{k}_{\mathcal{C}_A^+} = \left\lceil \frac{1}{2s} - 2 \right\rceil = \bar{k}-1$), some consumers who anonymize themselves (those in $\mathcal{C}_A^+$) end up with positive net surplus, whereas others (those in $\mathcal{C}_A^-$) end up with negative net surplus.*

Proposition 3 shows that (except for boundary cases) the strategic sophistication of consumers works to their disadvantage at an aggregated level and can break down the market for anonymous shopping. By contrast, the seller benefits from this stepwise breakdown, a development that would also be appreciated by a total welfare maximizer. The reason for this preference is, interestingly, not based on allocation: Due to perfect price discrimination in channel D, all consumers with a valuation above the marginal cost of production get the product independent of the existence of channel A. If such big data technologies are already in place, it is inefficient to incur the cost of anonymization. The less channel A is used, the smaller this inefficiency.

Corollary 1 zooms into anonymizing consumers. See Fig. 2, where three groups of consumers are distinguished: $\mathcal{C}_D$, $\mathcal{C}_A^-$, and $\mathcal{C}_A^+$. While the first denotes those consumers who chose channel D, $\mathcal{C}_A^-$ contains those consumers in channel A who make a net loss (because they do not anticipate the seller's incentive to increase the price by s fully), whereas those in $\mathcal{C}_A^+$ end up with a net benefit.

4.3. Comparative statics for s

The following analysis of the effects of changes in the cost of anonymization may inform whether policy efforts to reduce the cost of anonymizing techniques are desirable. Before delving into the analysis, we first include the case of $s = 0$ as the lower bound and identify the threshold where anonymization becomes prohibitively costly for channel A to be used at all (the equivalent of $\bar{k}$) as the upper bound for our analysis. Denoting this threshold by $\bar{s}$, we define:

$$\bar{s} \equiv \min\{s \in \mathbb{R}_0^+ | \mathcal{C}_A = \emptyset\}. \tag{24}$$

Recalling that $\mathcal{C}_A = (\hat{v}_k; 1]$ and $\hat{v}_k = \frac{1}{2} + (k+1)s$, yields:

$$\mathcal{C}_A = \emptyset \Leftrightarrow \hat{v}_k \geq 1 \Leftrightarrow s \geq \frac{1}{2(k+1)} \Rightarrow \bar{s} = \frac{1}{2(k+1)}. \tag{25}$$

[23] Recall that $\bar{k}$ is usually the result of rounding (unless $\frac{1}{2s} - 1 \in \mathbb{Z}_0^+$) and, hence, the last change in the composition of $\mathcal{C}_A$ and $\mathcal{C}_D$ is usually of different size than s. When increasing consumer sophistication from $\bar{k}-1$ to $\bar{k}$, the increase of $\mathcal{C}_D$ is bounded from above by s as *all remaining* consumers switch to channel D. Introducing separate cases in all difference equations is avoided for legibility, but addressed where necessary.

[24] The dark grey triangle shrinks by a trapezoid composed of a rectangle of area $(1-\hat{v}_{k+1})s$ and a triangle of area $\frac{s^2}{2}$, whereas the dashed rectangle has height s and shrinks in width by s, making for a decrease in area of s^2.

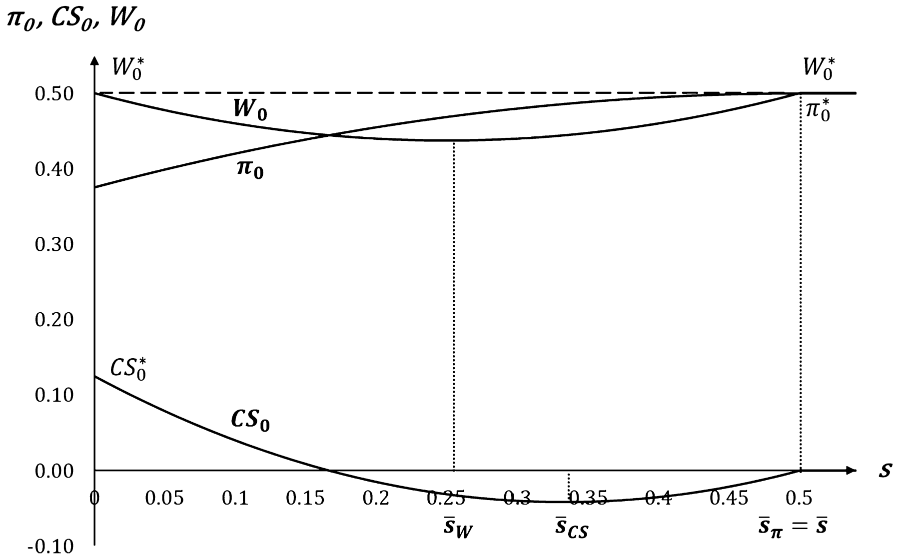

Fig. 3. Consumer surplus, profits and welfare as functions of s for $k = 0$.

Since s is a continuous variable, we do not need to take special cases into account and can use derivatives (instead of differences) of (18), (19), and (20).

For $0 \le s \le \bar{s}$:

$$\frac{\partial CS_k}{\partial s} = -\frac{1}{2}(k+2) + (k+1)(k+3)s = (k+1)s - (k+2)(1-\hat{v}_k), \tag{26}$$

$$\frac{\partial \pi_k}{\partial s} = \frac{1}{2}(k+1) - (k+1)^2 s \qquad = (k+1)(1-\hat{v}_k), \tag{27}$$

$$\frac{\partial W}{\partial s} = -\frac{1}{2} \qquad + 2(k+1)s \qquad = (k+1)s - (1-\hat{v}_k). \tag{28}$$

Defining thresholds similarly as above and limiting the analysis to $0 \le s \le \bar{s}$ yields:

$$\frac{\partial CS_k}{\partial s} \begin{cases} < 0 & \text{if } s < \frac{1}{2}\frac{k+2}{(k+1)(k+3)} \equiv \bar{s}_{CS}, \\ \ge 0 & \text{if } s \ge \frac{1}{2}\frac{k+2}{(k+1)(k+3)} \equiv \bar{s}_{CS}, \end{cases} \tag{29}$$

$$\frac{\partial \pi_k}{\partial s} \begin{cases} > 0 & \text{if } s < \frac{1}{2}\frac{1}{k+1} \equiv \bar{s}_{\pi}, \\ \le 0 & \text{if } s \ge \frac{1}{2}\frac{1}{k+1} \equiv \bar{s}_{\pi}, \end{cases} \tag{30}$$

$$\frac{\partial W_k}{\partial s} \begin{cases} < 0 & \text{if } s < \frac{1}{2}\frac{1}{2k+2} \equiv \bar{s}_{W}, \\ \ge 0 & \text{if } s \ge \frac{1}{2}\frac{1}{2k+2} \equiv \bar{s}_{W}. \end{cases} \tag{31}$$

It can further be shown that

$$\bar{s} = \bar{s}_\pi > \bar{s}_{CS} > \bar{s}_W, \tag{32}$$

which reveals that the seller's profits increase in s until the prohibitive level $\bar{s}$ is reached. With increasing anonymization cost, more consumers choose channel D instead of channel A, such that the seller appropriates their entire valuation. The effects on consumer surplus and total welfare, on the other hand, are ambiguous and depend on the initial level of s. This is due to the changing effects of raising s on the composition of $\mathcal{C}_A$: At first, $\mathcal{C}_A^-$ increases in size as s increases. But when there are no consumers in $\mathcal{C}_A^+$ anymore, a further increase will reduce the size of $\mathcal{C}_A^-$ again. The respective second derivatives provide further insights:

$$\frac{\partial^2 CS_k}{\partial s^2} = (k+1)(k+3) > 0, \tag{33}$$

$$\frac{\partial^2 W_k}{\partial s^2} = 2(k+1) \qquad > 0. \tag{34}$$

Both, consumer surplus and total welfare, are convex in s, implying that they reach a local maximum at $s = \bar{s}$. Recalling Proposition 2 for $s = 0$ and substituting $s = \bar{s}$ in (18), (19), and (20) allows us to compare consumer surplus, profits, and welfare at either extreme case:

$$CS_k(s = 0) = \frac{1}{8}, \qquad CS_k(s = \bar{s}) = 0, \tag{35}$$

$$\pi_k(s = 0) = \frac{3}{8}, \qquad \pi_k(s = \bar{s}) = \frac{1}{2}, \tag{36}$$

$$W_k(s = 0) = \frac{1}{2}, \qquad W_k(s = \bar{s}) = \frac{1}{2}. \tag{37}$$

Proposition 4 summarizes the comparative statics for s. Fig. 3 illustrates.

Proposition 4 *(Anonymization Cost and Welfare). For any non-prohibitively high level of consumer sophistication $k < \bar{k}$ anonymization is prohibitively costly for $s \geq \bar{s} = \frac{1}{2(k+1)}$. Then, channel A remains unused. As long as channel A is used, aggregated consumer surplus (CS_k), profits (π_k), and welfare (W_k) exhibit the following characteristics:*

- *CS_k is maximal at $s = 0$, decreases in s to its minimum (which is negative) at $s = \bar{s}_{CS}$, then increases in s back to 0 at $s = \bar{s}$.*
- *π_k is minimal at $s = 0$ and increases in s to its maximum at $s = \bar{s}_\pi = \bar{s}$.*
- *W_k is maximal at $s = 0$, decreases in s to its minimum at $s = \bar{s}_W$, then increases in s to another maximum at $s = \bar{s}$. Both maxima lead to the first-best outcome $W_k^* = \frac{1}{2}$.*

Proposition 4 shows that higher cost of anonymization is negative for consumers despite the fact that consumer surplus increases in s for relatively high values, which becomes apparent from the fact that consumer surplus is maximal when anonymization is costless. The seller, on the other hand, unambiguously benefits from higher cost of anonymization and, hence, prefers prohibitively high cost of $s = \bar{s}$. This maximizes the seller's profits as s/he can extract all consumer surplus via channel D. A total welfare maximizer, focusing on the welfare-deteriorating role of s, can choose either extreme to prevent consumers from incurring the cost: To achieve an efficient outcome, anonymization should be either costless ($s = 0$) or prohibitively costly ($s = \bar{s}$). Both options are welfare-maximizing, but lead to different surplus allocations. While the seller makes positive profits in either welfare-maximizing scenario, consumers only receive positive surplus if $s = 0$.

5. Extensions

5.1. Heterogeneous cost of anonymization

We have assumed that all consumers find it equally costly to use the anonymous channel A. However, it is easy to imagine that some people might find it less cumbersome to discover and make use of privacy-protecting technologies such as deleting cookies, activating "do not track" browser options, or installing various plugins.[25] Additionally, heterogeneity in s can stem from differing exogenous tastes for privacy in the consumer population, which would reduce the experienced disutility of using channel A. Heterogeneous values of s could be seen as the result of aggregating these and other effects.

If consumers have heterogeneous anonymization costs s_i, where $s_i \sim \mathcal{U}[\underline{s}, \bar{s}]$ (we redefine $\bar{s}$ for this section), they anonymize for $v > \mathbb{E}_k(p_A) + s_i$. As this expression now depends on two individually heterogeneous variables, there is no uniform cutoff $\hat{v}$ separating the sets $\mathcal{C}_D$ and $\mathcal{C}_A$. Rather, three segments of consumers' valuations v need to be distinguished (cf. Fig. 4):

(a) Consumers with $v > \mathbb{E}_k(p_A) + \bar{s}$,
(b) Consumers with $v \in \big(\mathbb{E}_k(p_A) + \underline{s}, \mathbb{E}_k(p_A) + \bar{s}\big]$,
(c) Consumers with $v \leq \mathbb{E}_k(p_A) + \underline{s}$.

Given their price expectation, consumers in segment (c) have a valuation v so low that they choose channel D even for the lowest possible cost of anonymization $\underline{s}$. Vice versa, consumers in segment (a) have a sufficiently high valuation v such that they choose channel A even if they face the highest possible cost of anonymization $\bar{s}$. For consumers in segment (b), however, the precise level of their anonymization cost s_i matters for their anonymization choice. Given any valuation $v \in \big(\mathbb{E}_k(p_A) + \underline{s}, \mathbb{E}_k(p_A) + \bar{s}\big)$ only those whose cost of anonymization s_i is sufficiently low choose channel A, while others with the same valuation v but a higher cost s_i choose channel D. Fig. 4 exemplifies this for consumers with a valuation of $v = 0.6$: among these consumers those with anonymization cost $s_i < 0.1$ choose channel A, but those with $s_i \geq 0.1$ choose channel D.

[25] For instance, TOR is a "free software and an open network that helps [users] defend against traffic analysis, a form of network surveillance that threatens personal freedom and privacy, confidential business activities and relationships, and state security" (https://www.torproject.org/). A more detailed list can be found in Sellenart's "A paranoid's toolbox for browsing the web": http://pierre.senellart.com/talks/cerre-20160915.pdf.

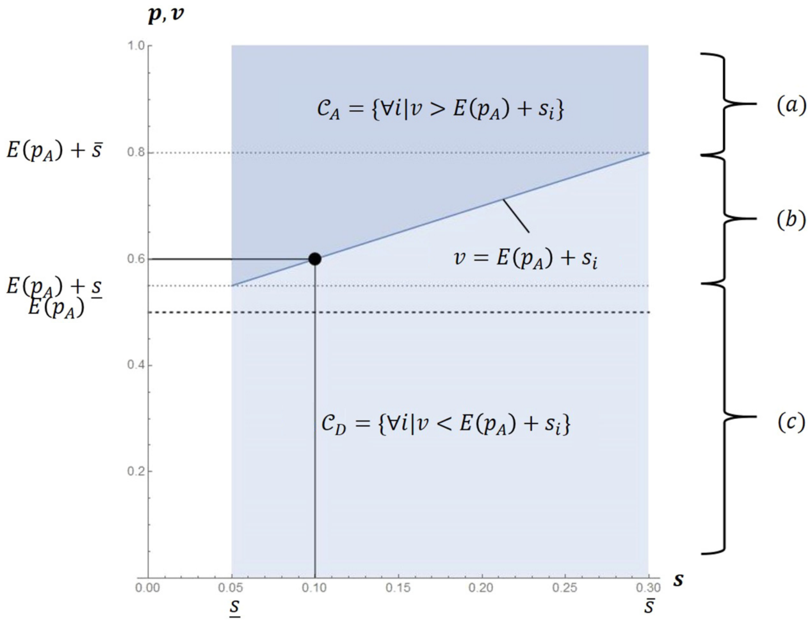

Fig. 4. Composition of sets $\mathcal{C}_D$ and $\mathcal{C}_A$ depending on v and s_i with parameters $v \sim \mathcal{U}[0,1]$, $s_i \sim \mathcal{U}[0.05, 0.3]$, $k = 0$.

The new composition of the sets $\mathcal{C}_D$ and $\mathcal{C}_A$ implies that demand in both channels is now defined differently for all three segments and hence becomes a piecewise (but still continuous) function of p. As the seller still perfectly price-discriminates in channel D, we focus on the implications of this change for channel A. There, demand is still linear for prices in segment (a) and constant for prices in segment (c). For price levels in segment (b), however, the uniformly distributed cost of anonymization leads to quadratic demand. We derive the demand function of channel A in general in Appendix B.1. For the specific parameters of Fig. 4, this leads to the following demand in channel A:

$$q_A(p_A) = \begin{cases} 0 & \text{if } p_A \geq 1, \\ 1 - p_A & \text{if } 0.8 < p_A < 1, \\ 1 - 0.625 - \frac{(p_A - 0.55)^2}{0.5} & \text{if } 0.55 < p_A \leq 0.8, \\ 1 - 0.625 & \text{if } p_A \leq 0.55. \end{cases} \tag{38}$$

The general derivation of demand in Appendix B.1 confirms that demand is always quadratic for $p \in \left(\mathbb{E}[p_A] + \underline{s}, \mathbb{E}[p_A] + \bar{s}\right]$. However, we show in Appendix B.2 that $p^*_{A_k} > \mathbb{E}_k(p_A) + \underline{s}$ as long as $\mathcal{C}_A \neq \emptyset$, i.e. the seller's optimal price is no longer equal to the valuation at the lower bound of $\mathcal{C}_A$. This in turn implies that there are some consumers in $\mathcal{C}_A$ that do not buy the product.

Continuing the example from above, this result is illustrated in Fig. 5. There, the left panel depicts the demand function $q_A(p_A)$ as well as the optimal price in channel A, given by $p^*_A = 0.6629$ with the chosen parameters, whereas the right panel replicates Fig. 4 to highlight the mapping from $\mathcal{C}_A$ to $q_A(p_A)$), but replacing the example point with the optimal price p^*_A. Both panels of Fig. 5 show that the optimal price p^*_A now exceeds the lowest valuation in $\mathcal{C}_A$. Consumers in the white area between $q_A(p^*_A)$ and $q_A(p_A)$ do not buy the product despite having a valuation of at least p_M (the unconditional monopoly price). Contrary to our baseline model, the seller is now willing to forgo profits from some consumers because the density of consumers in $\mathcal{C}_A$ across valuations is no longer uniform in the neighborhood of the lower bound of $\mathcal{C}_A$.

If we increase consumers' strategic sophistication from $k = 0$ to $k = 1$, consumers form the expectation $\mathbb{E}_1(p_A) = 0.6629$ and make their anonymization choice accordingly. The difficulty of finding an analytical closed-form solution for p^*_A for any k transmits to finding the cutoff level $\bar{k}$, from which onwards channel A remains unused. Fig. 6 illustrates that profit-maximizing prices $p^*_A(k)$ do not linearly increase in k, unlike in our baseline model (cf. (13)).

The qualitative result, that channel A is used less for higher levels of consumers' sophistication and eventually remains unused, however, is replicated: Once the optimal price falls within the interval $(0.95, 1]$, channel A is unused at the next higher k. Notably, this only holds if anonymization is costly for all consumers, i.e. if $\underline{s} > 0$. If $\underline{s} = 0$, we show in Appendix B.3 that a full breakdown of channel A is not achieved for finite k, because $p^*_{A_k} < 1$ for all finite k.[26]

[26] This reinstates unlimited sophistication as a *necessary* condition for a full breakdown of channel A if $\underline{s} = 0$.

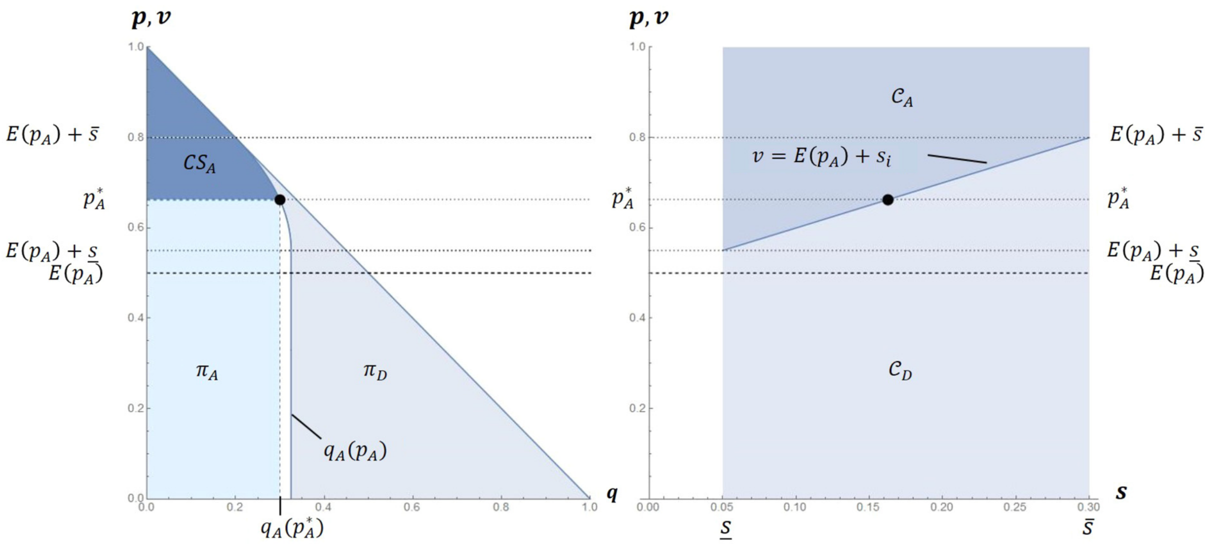

Fig. 5. Consumers' anonymization choice as a function of v and s_i with parameters $v \sim \mathcal{U}[0, 1]$, $s_i \sim \mathcal{U}[0.05, 0.3]$, $k = 0$.

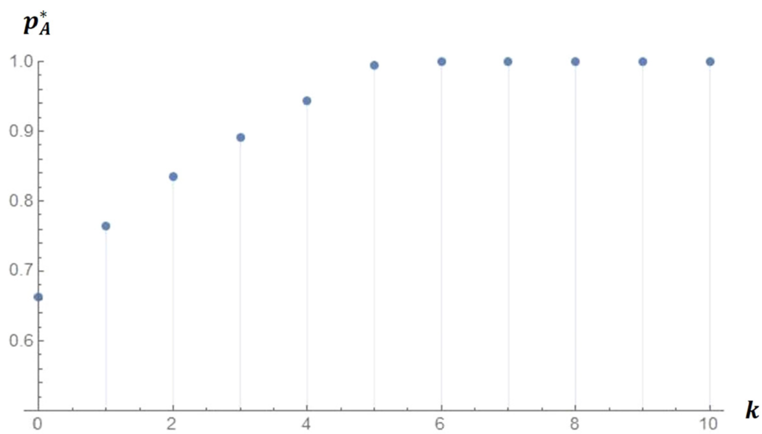

Fig. 6. Optimal price in channel A p_A^* for $k = 0, 1, 2, \ldots, 10$ with parameters $v \sim \mathcal{U}[0, 1]$, $s_i \sim \mathcal{U}[0.05, 0.3]$.

Summing up, we conclude that despite some quantitative changes, the general pattern of a gradual breakdown of the anonymous channel and the corresponding effects are not altered qualitatively if anonymization costs are heterogeneous.[27]

5.2. Increasing competition

Many markets where sellers have access to large amounts of data on consumers' preferences or characteristics, a prerequisite for perfect price discrimination, are dominated by one firm.[28] But to which extent would such a dominant firm, or a monopolist in a market niche, adapt behavior if consumers had access to an (imperfect) substitute product, thereby increasing competition? Assume a rival R offers a product competing with the monopolist's product. A consumer's net value of the rival's product is

[27] A related model variant is to consider $s < 0$, e.g. due to privacy gains outweighing the cost associated with anonymization. Here, consumers receive a positive utility from choosing the Channel A. Given that they still understand that they will be perfectly price-discriminated against in Channel D, they can at best realize a benefit of 0 by staying in channel D. However, in Channel A they can now realize at least a gain of size $|s|$ because they can always choose not to buy the product at Stage 3 of the game. Then, all consumers will pick the anonymous channel, and the seller cannot do any better than to set the profit-maximizing price of $p_M^* = \frac{1}{2}$, which in turn leads to the typical deadweight loss associated with a monopolist. Whether this loss in total welfare is compensated by the benefit consumers achieve from the "sunk benefit" of s, depends on the magnitude of s, which is no longer subject to any prohibitive levels where the market breaks down. Given that all consumers now have an incentive to choose channel A regardless of their eventual purchase, there is also no difference in behavior stemming from the level of strategic sophistication.

[28] Prüfer and Schottmüller (2021) explain this development in theoretical terms and cite empirical work to support the statement.

$$v^R \equiv \sigma v - p_R, \tag{39}$$

where $\sigma \in [0,1)$ denotes the degree of substitutability of the products and p_R denotes the price of the rival's product. Alternately, σ proxies the intensity of competition. As any $p_R > 0$ can be reflected by a lower intensity of competition σ, we assume $p_R = 0$ and focus on changes in σ to study the effects of increasing competition. In this scenario, consumers buy from the "monopolist" M if, and only if, $v - p \geq v^R$, i.e. if

$$v - v^R = (1-\sigma)v \geq p, \tag{40}$$

where $p \in \{p_i(v), p_A\}$. Knowing v precisely for all consumers in $\mathcal{C}_D$, M sets

$$p_i^* = (1-\sigma)v \text{ for all } i \in \mathcal{C}_D, \tag{41}$$

thus guaranteeing that no consumer in channel D buys the rival's product as long as M can earn a profit from this consumer. Consumers anonymize if $\mathbb{E}(u(A)) > \mathbb{E}(u(D))$, i.e. if, and only if,

$$v - \mathbb{E}_k(p_A) - s > v - p_i^* \Leftrightarrow \tag{42}$$

$$v > \frac{\mathbb{E}_k(p_A) + s}{1-\sigma} \equiv \tilde{v}_k, \tag{43}$$

where $\tilde{v}_k$ denotes the cutoff valuation akin to $\hat{v}_k = \mathbb{E}_k(p_A) + s$ in our main model.

Lemma 6 *(Anonymization and Monopolistic Competition). For any given price expectation of consumers for channel A, $\mathbb{E}_k(p_A)$, the presence of a rival selling a product of substitutability $\sigma \in [0,1)$ raises the anonymization threshold from $\hat{v}_k$ to $\tilde{v}_k = \frac{\hat{v}_k}{1-\sigma}$.*

Understanding this increase in the anonymization threshold, M might consider to also increase his/her price, as in the baseline model, and set $p_{A_k}^* = \tilde{v}_k$. But consumers in channel A might still buy from R. Thus, M faces the same participation constraint in channel A as in channel D: to leave every consumer with at least a net surplus of σv. It follows that pricing at $\tilde{v}_k$ is infeasible. M has to decrease the price to fulfill:

$$(1-\sigma)\tilde{v}_k \geq p_{A_k}^* \tag{44}$$

which yields, in combination with (43),

$$p_{A_k}^* = (1-\sigma)\frac{\hat{v}_k}{1-\sigma} = \hat{v}_k = \mathbb{E}_k(p_A) + s. \tag{45}$$

Thus, M cannot capitalize on the higher anonymization threshold as a consequence of increased competition. Even though all consumers in channel A now end up with a positive net surplus from the transaction (depending on how s compares to the guaranteed benefit of σv), some consumers are still worse off, having chosen channel A instead of channel D. Consumers with $k = 1$, however, do not adjust their expectation based on forgone surplus but simply update it to the price that would be M's best response to consumers with $k = 0$, as in the baseline model.

Summarizing, consumer surplus increases with competitive pressure. Therefore, there is less anonymization for any given price expectation. Prices in channel D decrease to account for consumers' improved outside option but the price in channel A is unaffected by competition.

5.3. Same level of sophistication for seller and consumers

Here we consider the case where the seller does not have an advantage in strategic sophistication and, in contrast to our main model specification, has the same level of strategic sophistication as consumers. Naturally, consumer behavior stays the same as in our main model specification. Before, the seller's level of sophistication was at least $k = 1$ and, thus, there was no need to also specify the seller's belief formation when fully naïve. Given that the seller and the consumers are now equally naïve, the seller is no longer able to infer exactly which set of consumers has chosen to anonymize. Rather, the seller has to follow the same iterative thinking process as the consumers. The seller, now likewise thinking that the market is a simple monopoly market, forms the naïve belief that all consumers with a valuation $v \geq p_M = \frac{1}{2}$ have anonymized. This means, the seller falls prey to a similar strategic misconception as the naïve consumers and does not properly take into account the strategic nature of the anonymization cost s. The seller therefore ends up setting the regular monopoly price as the price for channel A, incidentally proving the expectation of the consumers correct, which leads to the following resulting solution:

Lemma 7 *(Solution with joint level-0 of all agents). For any non-prohibitively high anonymization cost $s \geq 0$ and with both the seller and consumers being strategically "naïve" ($k = 0$) there is a unique solution with the following characteristics:*

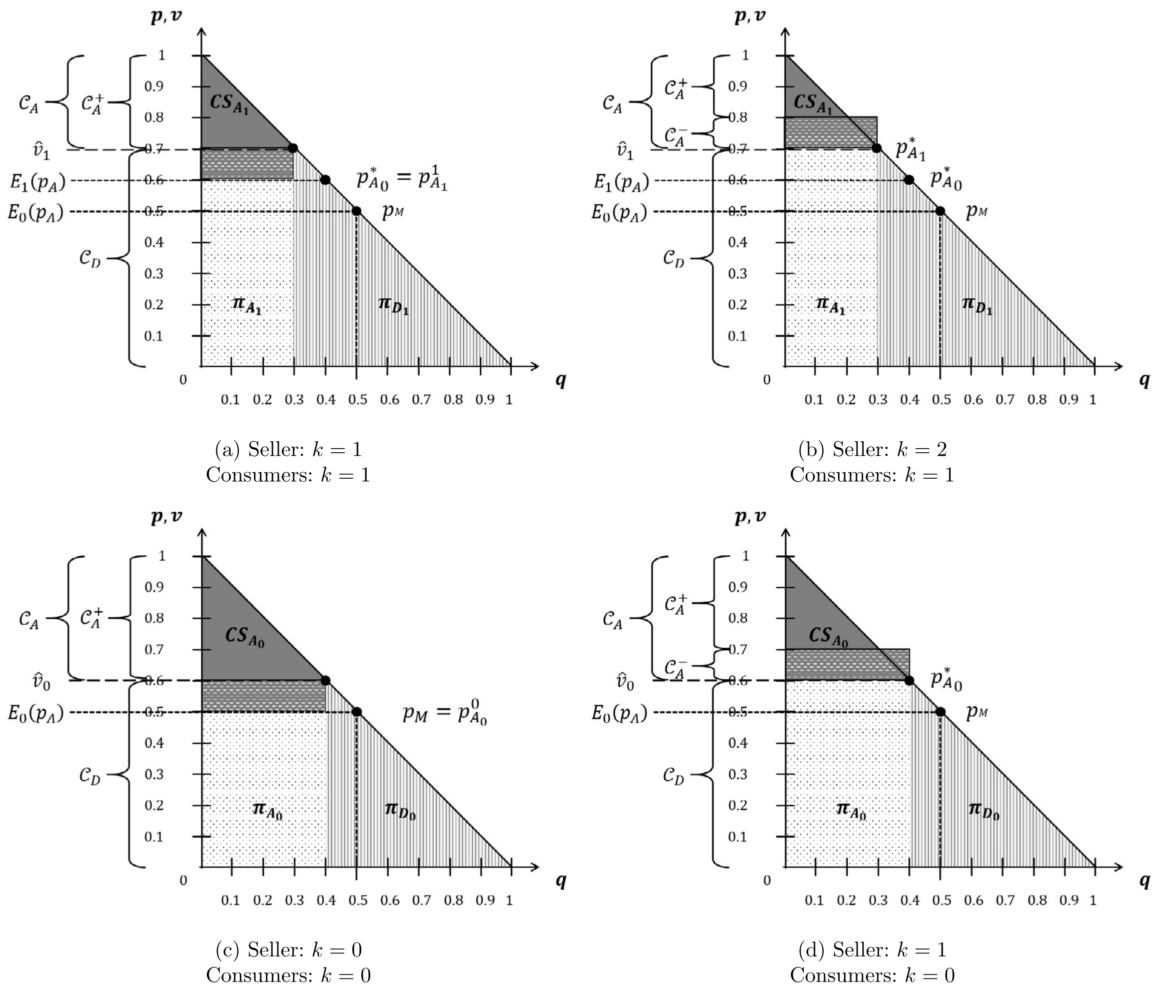

Fig. 7. Comparison of same and different levels of sophistication.

- *Consumers form the 0-beliefs* $\mathbb{E}_0(p_D) = p_i^*(v)$ *and* $\mathbb{E}_0(p_A) = p_M = \frac{1}{2}$.
- *Consumers anonymize if, and only if,* $v > \hat{v}_0 = \mathbb{E}_0(p_A) + s = p_M + s$. *Hence,* $\mathcal{C}_D = [0, \hat{v}_0]$ *and* $\mathcal{C}_A = (\hat{v}_0, 1]$.
- *The seller forms the 0-belief about the anonymization threshold* $\mathbb{E}_0(\hat{v}_0) = p_M$ *and infers correctly that, therefore,* $\mathcal{C}_D = [0, p_M]$ *and* $\mathcal{C}_A = (p_M, 1]$.
- *The seller sets prices* $p_i^*(v) = v$ *and* $p_{A_0}^* = p_M$.
- *All consumers buy the product at the offered price.*

From there on, both the consumers as well as the seller will always realize the same next step in their respective strategic reasoning when k increases by one. That is, they understand that the optimal price of the seller in response to a consumer population of one level below their current level of sophistication would have been to incorporate the anonymization cost into the respective belief about the cutoff valuation between the consumers in channel A and channel D. As a result, the seller will raise the price by s with every level of sophistication. However, because the consumers do the same, they now end up having a correct belief about the price the seller sets in the anonymous market. Consequently, no consumer ends up with a net loss from the transaction, anymore. The seller, on the other hand, ends up leaving some surplus on the table for consumers in channel A as compared to the situation where his/her level of strategic sophistication exceeds the consumers' level. Overall welfare, though, is unaffected by this because the fraction of consumers that incur the socially "wasteful" anonymization cost s stays the same and because all mutually beneficial trades still occur in Stage 3.

Fig. 7 depicts the resulting distribution of welfare in comparison to the respective situations of our main model specification. The two panels on the left (7a and 7c) show outcomes of same-level strategic sophistication and the two panels on the right (7b and 7d) show the respective outcome under different levels, where consumers' level of sophistication is the same as on the respective left panel.

6. Discussion and conclusion

This paper started from the empirical observation that the technological developments that are behind the "rise of big data" have led to asymmetries on markets for consumer goods (Mayer-Schönberger and Cukier, 2013). Sellers making use of datasets on choices of large masses of consumers can tailor prices to individual characteristics more and more and thereby appropriate a lion's share of the surplus created by market transactions. The specific model of perfect price discrimination used here is just the limit case that will be approximated when prediction methods, mostly using machine learning, improve in power. On top of the sheer amount of information that is available to sellers, consumers are at a second disadvantage. They face cognitive constraints regarding strategic sophistication (Acquisti and Grossklags, 2007), while the seller's data processing capabilities enable him to find best responses to consumers' behavior.

We have taken these developments seriously and constructed a model to study their implications on prices, consumption choices, and consumers' incentives to use anonymization technologies protecting their privacy. We have shown that under certain conditions, most notably under the assumption of imperfect strategic sophistication of consumers, a costly privacy-protective sales channel is used even if consumers do not have an explicit taste for privacy. In our model, consumers want to restore their privacy (i.e. choose channel A) solely based on their valuation of the good and their price expectation. We thereby provide a micro-foundation for consumers' privacy choices, to which the existence of a privacy-protective sales channel can cater.

Our model shows that unlimited strategic sophistication is a sufficient but not a necessary condition for the breakdown of the anonymous sales channel if anonymization is equally costly to all consumers. Allowing for heterogeneity in anonymization cost, sources of which can be different technological savviness but also differing preferences for privacy, can reinstate the necessity of unlimited strategic sophistication for a complete breakdown of the anonymous channel.

In general, the use of big data technologies by sellers improves total welfare by avoiding the deadweight loss usually associated with a monopoly: In contrast to uniform monopoly pricing, consumers with low valuations, $v < p_M$, can purchase the product now. This increases efficiency but not consumer surplus as the seller appropriates the entire surplus from these additional transactions. We have further shown that using the anonymous channel backfires and leads to a net loss for at least some (and, under certain conditions, for all) anonymized consumers (forming the set $\mathcal{C}_A^-$). We have shown that increasing consumer sophistication leads to a reduction in consumer surplus but to an increase in profits and total welfare. Although a policy-maker will unlikely be able to affect consumer sophistication directly, our comparative statics analysis on k has shown that any policy aimed at changing the anonymization cost s only matters if consumers are not too sophisticated ($k < \bar{k}$). Analyzing different anonymization cost levels, consumer surplus is largest in the extreme case of costless anonymization, $s = 0$, and profits are maximal in the extreme case of anonymization being prohibitively costly, $s = \bar{s}$.

Total welfare, however, is maximal at either extreme, $s = 0$ or $s = \bar{s}$. The two cases differ, however, in the way in which they ensure a first-best result. If $s = 0$, consumers with high valuation anonymize for free and receive positive surplus, whereas all others choose the direct channel, where they get perfectly discriminated against and are left with zero surplus. If $s = \bar{s}$, consumers choose the direct channel irrespective of their valuation for the good. Allocation is efficient because anonymization is too costly and the seller can appropriate all surplus.

Because the distribution of surplus is highly asymmetric, policy makers may want to have a second look. A consumer-oriented welfare maximizer should try to eliminate anonymization cost, whereas a seller-oriented welfare maximizer may seek to increase anonymization costs to a prohibitive level. A policy maker with a preference for consumer surplus could, for example, require marketers and online platforms serving as matchmakers between sellers and buyers of consumer goods to set anonymous shopping as *default*. This would then require consumers to *opt in* to non-anonymous shopping instead of today's standard, where full tracking of consumers' choices is the default and a few providers offer *opt out* technologies. This proposal is also discussed by Acquisti et al. (2016). Those consumers willing to reveal their characteristics to sellers (our model suggests these are those with low valuations) would log in to some service and receive the product for a price related to their WTP. Consumers with higher valuations would stay in the (now default) anonymous channel and pay a higher price, but still retain some surplus.

On the theory side, future research could shed light on the effects of heterogeneity in the level of strategic sophistication amongst consumers, relying on a more elaborate cognitive hierarchy model than this first attempt we undertook here.[29] This is a complex undertaking, however, because it is not only necessary to specify a distribution of k-levels across the population of consumers (and how it might be related to their WTP). It also requires to specify every consumer's belief about other consumers' level(s) of sophistication and the seller's response to them as a function of that consumer's own sophistication.

To test our theory empirically, we consider it most promising to conduct laboratory experiments where subjects could be assigned specific valuations and a perfectly price-discriminating algorithm could actually be implemented. Subjects could indicate their respective anonymization choice given a valuation and a known cost of anonymization. The implied thresholds for anonymization would correspond to a certain level of strategic sophistication according to our model, which in turn

[29] In a cognitive hierarchy model as introduced by Camerer et al. (2004) higher-level players would form beliefs about the distribution of all levels of sophistication lower than their own.

could be compared to other measures of strategic sophistication spawned from the level-k literature. Using measures that capture differences between belief interactions of subjects could inform whether the current model, which neglects more complex cognitive hierarchies, is a fair representation of subjects' approach to such a market or whether efforts to generalize our theory are needed.

Acknowledgments

We are grateful to seminar audiences at Tilburg, Mannheim (Governance of Platform markets in the Big Data era), Télécom ParisTech, ZEW Mannheim (14th ICT Conference), Carlos III Madrid (2016 ENTER Jamboree), ETH Zürich, Louvain (IODE workshop), Harvard (ISNIE conference), Tilburg (CLEEN meeting), Toulouse, and the ESNIE Spring School (2014). We are especially thankful to Alessandro Acquisti, Jan Boone, Eric van Damme, Clemens Fiedler, Ricard Gil, Víctor H. González-Jiménez, Heiko Karle, Christian König gen. Kersting, Marian W. Moszoro, Martin Peitz, Valerio Poti, Jan Potters, Renata Rabovič, Christoph Schottmüller, Florian Schütt, Katja Seim, Gyula Seres, Giancarlo Spagnolo, Vatsalya Srivastava, Birger Wernerfelt, the (associate) editor, and two anonymous reviewers, who provided valuable feedback on earlier versions of this paper. All errors are our own.

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Appendix A

A.1. Consumer surplus and profits in channel D

As the seller engages in perfect price discrimination for consumers in $\mathcal{C}_D$,

$$CS_{D_k} = 0, \tag{A.1}$$

whereas the seller appropriates the entire surplus in channel D:

$$\pi_{D_k} = \frac{1}{2}\hat{v}_k^2 = \frac{1}{8} + \frac{1}{2}(k+1)s + \frac{1}{2}(k+1)^2 s^2. \tag{A.2}$$

A.2. Consumer surplus and profits in channel A

In channel A, consumer surplus consists of two parts: the benefit from consumption of the good after the transaction at Stage 3 (denoted by $CS^+_{A_k}$) and the cost of anonymization incurred at Stage 1 (denoted by $CS^-_{A_k}$):

$$CS^+_{A_k} = \frac{1}{2}(1-\hat{v}_k)^2 = \frac{1}{8} - \frac{1}{2}(k+1)s + \frac{1}{2}(k+1)^2 s^2, \tag{A.3}$$

$$CS^-_{A_k} = (1-\hat{v}_k)s \quad = \frac{1}{2}s - (k+1)s^2. \tag{A.4}$$

Net consumer surplus in channel A, $CS_{A_k} \equiv CS^+_{A_k} - CS^-_{A_k}$, then amounts to:

$$CS_{A_k} = \frac{1}{2}(1-\hat{v}_k)^2 - (1-\hat{v}_k)s = \frac{1}{8} - \frac{1}{2}(k+2)s + \frac{1}{2}(k+1)(k+3)s^2. \tag{A.5}$$

Only some consumers in channel A end up with positive net surplus ($\mathcal{C}^+_A = [\hat{v}_k + s, 1]$), whereas others end up with negative net surplus ($\mathcal{C}^-_A = (\hat{v}_k, \hat{v}_k + s)$). The seller's profits in channel A are given by

$$\pi_{A_k} = \hat{v}_k - \hat{v}_k^2 = \frac{1}{4} - (k+1)^2 s^2. \tag{A.6}$$

A.3. Proof of Proposition 2

As (18) shows the cost of anonymization incurred by all consumers in $\mathcal{C}_A$ can outweigh the combined surplus from purchasing the good when $\hat{v}_k$ is close to 1 (i.e. k close to $\bar{k}$). We define the additional threshold $\bar{k}_{CS}$, where consumer surplus turns negative, and solve using (18):

$$\bar{k}_{CS} \equiv min\{k \in \mathbb{Z}^+_0 | CS_k \leq 0\} = \left\lceil \frac{1}{2s} - 3 \right\rceil. \tag{A.7}$$

Recalling $\bar{k} = \lceil \frac{1}{2s} - 1 \rceil$ reveals that

$$\bar{k} - \bar{k}_{CS} = \left\lceil \frac{1}{2s} - 1 \right\rceil - \left\lceil \frac{1}{2s} - 3 \right\rceil \qquad = \left\lceil \frac{1}{2s} \right\rceil - \left\lceil \frac{1}{2s} \right\rceil + 2 = 2, \tag{A.8}$$

i.e. consumer surplus turns negative at the penultimate level before channel A breaks down.

According to (19), profits are strictly positive as long as

$$\hat{v}_k - \frac{1}{2}\hat{v}_k^2 \geq 0 \Rightarrow 0 < \hat{v}_k < 2. \tag{A.9}$$

This holds for any k since $\hat{v}_k \in [\hat{v}_0 = p_M + s = \frac{1}{2} + s, 1]$.

According to (20), welfare is strictly positive as long as

$$\frac{1}{2} \geq (1 - \hat{v}_k)s. \tag{A.10}$$

This holds for any k since $(1 - \hat{v}_k) \in [0, 1 - \hat{v}_0 = 1 - (p_M + s) = \frac{1}{2} - s] < \frac{1}{2}$ because $s > 0$. □

A.4. Proof of Proposition 3

For any $\hat{v}_{k+1}$ sufficiently far from 1, consumer surplus decreases in k; see (21). Denoting the threshold level where consumer surplus stops decreasing by $\bar{k}_{\Delta CS}$, we define:

$$\bar{k}_{\Delta CS} \equiv min\{k \in \mathbb{Z}_0^+ | \Delta CS_k \geq 0\}. \tag{A.11}$$

Using (21) in (A.11), addressing the discreteness of k as before, and solving yields:

$$\bar{k}_{\Delta CS} \geq \frac{1}{2s} - \frac{5}{2} \Rightarrow \bar{k}_{\Delta CS} = \left\lceil \frac{1}{2s} - \frac{5}{2} \right\rceil. \tag{A.12}$$

We can pin down this threshold's location relative to $\bar{k} = \left\lceil \frac{1}{2s} - 1 \right\rceil$

$$\bar{k} - \bar{k}_{\Delta CS} = \left\lceil \frac{1}{2s} - 1 \right\rceil - \left\lceil \frac{1}{2s} - \frac{5}{2} \right\rceil = \left\lceil \frac{1}{2s} \right\rceil - \left\lceil \frac{1}{2s} - \frac{1}{2} \right\rceil + 1 \in \{1, 2\}, \tag{A.13}$$

which reveals that consumer surplus stops decreasing already one or two levels of sophistication before channel A breaks down. Recall that $\mathcal{C}_A = \mathcal{C}_A^- \cup \mathcal{C}_A^+$ and that $\mathcal{C}_A^-$ is situated below $\mathcal{C}_A^+$. Hence, as k increases, $\mathcal{C}_A^+$ seizes to contain consumers before $\mathcal{C}_A^-$ does. We define the additional threshold $\bar{k}_{\mathcal{C}_A^+}$, where no consumer in channel A gets net surplus from the transaction:

$$\bar{k}_{\mathcal{C}_A^+} \equiv min\{k \in \mathbb{Z}_0^+ | \mathcal{C}_A^+ = \emptyset\}. \tag{A.14}$$

Using the definition of $\mathcal{C}_A^+ = (\hat{v}_k + s, 1]$ in (A.14) and solving yields

$$\bar{k}_{\mathcal{C}_A^+} \geq \frac{1}{2s} - 2 \Rightarrow \bar{k}_{\mathcal{C}_A^+} = \left\lceil \frac{1}{2s} - 2 \right\rceil, \tag{A.15}$$

which implies for its location relative to $\bar{k}$:

$$\bar{k} - \bar{k}_{\mathcal{C}_A^+} = \left\lceil \frac{1}{2s} - 1 \right\rceil - \left\lceil \frac{1}{2s} - 2 \right\rceil = \left\lceil \frac{1}{2s} \right\rceil - \left\lceil \frac{1}{2s} \right\rceil + 1 = 1. \tag{A.16}$$

This shows that at the last level before channel A breaks down there are no consumers left in channel A making net surplus. Recalling that at $\bar{k}_{CS} = \bar{k} - 2$ the cost of anonymization incurred by all consumers in $\mathcal{C}_A$ outweighs the combined surplus from purchasing the good, i.e. consumer surplus is 0 at best. It can be seen that $\bar{k}_{\mathcal{C}_A^+}$ and $\bar{k}_{CS}$ jointly form the set in (A.13), the level at which consumer surplus stops decreasing. Due to the discreteness of k, the minimum can be attained at either level.

In any case, raising k from $\bar{k} - 1$ to $\bar{k}$ leads to an increase in consumer surplus because channel A remains unused and consumer surplus jumps to 0 as all consumers are being perfectly price discriminated in channel D.

The analysis for the seller's profits is much more straightforward. According to (22) they seem to be generally increasing in k. Conducting a similar analysis as before[30] delivers

$$\bar{k}_\pi \equiv \min\{k \in \mathbb{Z}_0^+ | \Delta\pi_k \leq 0\} \in \{\bar{k} - 1, \bar{k}\}. \tag{A.17}$$

However, all comparative statics difference equations only apply to $k < \bar{k} - 1$. Recalling that $\mathcal{C}_D$ increases until $k = \bar{k}$ and that the seller appropriates all surplus from consumers in channel D but only receives a share of the surplus in channel A, it follows that profits still increase when k changes from $\bar{k} - 1$ to $\bar{k}$.

[30] Substituting (22) in (A.17) yields $\bar{k}_\pi = \left\lceil \frac{1}{2s} - \frac{3}{2} \right\rceil$ resulting in $\bar{k} - \bar{k}_\pi = \left\lceil \frac{1}{2s} \right\rceil - \left\lceil \frac{1}{2s} - \frac{1}{2} \right\rceil \in \{0, 1\}$.

While increasing k has negative effects on consumer surplus and positive effects on profits, welfare generally increases in k as (23) shows (including the change from $\bar{k}-1$ to $\bar{k}$). This is driven by the fact that increasing the level of sophistication leads to fewer anonymized consumers, i.e. less cost of anonymization being incurred. Independently of k the surplus from transacting the good is always maximal due to perfect price discrimination in channel D (but surplus shifts from consumers to the seller as k rises). □

Appendix B. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.geb.2021.09.007.

References

Acquisti, Alessandro, Taylor, Curtis, Wagman, Liad, 2016. The economics of privacy. J. Econ. Lit. 54 (2), 442–492.

Acquisti, Alessandro, Grossklags, Jens, 2007. What can behavioral economics teach us about privacy. In: Acquisti, Alessandro, Gritzalis, Stefanos, Lambrinoudakis, Costas, De Capitani di Vimercati, Sabrina (Eds.), Digital Privacy: Theory, Technologies and Practices. Auerbach Publications, New York & London, pp. 363–377.

Acquisti, Alessandro, Varian, Hal R., 2005. Conditioning prices on purchase history. Mark. Sci. 24 (3), 367–381.

Armstrong, Mark, 2006. Recent developments in the economics of price discrimination. In: Blundell, R., Newey, W.K., Persson, T. (Eds.), Advances in Economics and Econometrics: Theory and Applications: Ninth World Congress, Vol. 2. Cambridge University Press, Cambridge, UK, pp. 97–141.

Belleflamme, Paul, Vergote, Wouter, 2016. Monopoly price discrimination and privacy: the hidden cost of hiding. Econ. Lett. 149, 141–144.

Bergemann, Dirk, Bonatti, Alessandro, 2015. Selling cookies. Am. Econ. J. Microecon. 7 (3), 259–294.

Benndorf, Volker, Kübler, Dorothea, Normann, Hans-Theo, 2015. Privacy concerns, voluntary disclosure of information, and unraveling: an experiment. Eur. Econ. Rev. 75, 43–59.

Calzolari, Giacomo, Pavan, Alessandro, 2006. On the optimality of privacy in sequential contracting. J. Econ. Theory 130 (1), 168–204.

Camerer, Colin F., Ho, Teck-Hua, Chong, Juin-Kuan, 2004. A cognitive hierarchy model of games. Q. J. Econ. 119 (3), 861–898.

Casadesus-Masanell, Ramon, Hervas-Drane, Andres, 2015. Competing with privacy. Manag. Sci. 61 (1), 229–246.

Conitzer, Vincent, Taylor, Curtis R., Wagman, Liad, 2012. Hide and seek: costly consumer privacy in a market with repeat purchases. Mark. Sci. 31 (2), 277–292.

Costa-Gomes, Miguel A., Crawford, Vincent P., 2006. Cognition and behavior in two-person guessing games: an experimental study. Am. Econ. Rev. 96 (5), 1737–1768.

Costa-Gomes, Miguel, Crawford, Vincent P., Broseta, Bruno, 2001. Cognition and behavior in normal-form games: an experimental study. Econometrica 69 (5), 1193–1235.

Crawford, Vincent P., 2003. Lying for strategic advantage: rational and boundedly rational misrepresentation of intentions. Am. Econ. Rev. 93 (1), 133–149.

Crawford, Vincent P., Iriberri, Nagore, 2007a. Level-k auctions: can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? Econometrica 75 (6), 1721–1770.

Crawford, Vincent P., Iriberri, Nagore, 2007b. Fatal attraction: salience, naïveté, and sophistication in experimental 'hide-and-seek' games. Am. Econ. Rev. 97 (5), 1731–1750.

Daughety, Andrew F., Reinganum, Jennifer F., 2010. Public goods, social pressure, and the choice between privacy and publicity. Am. Econ. J. Microecon. 2 (2), 191–221.

de Cornière, Alexandre, de Nijs, Romain, 2016. Online advertising and privacy. Rand J. Econ. 47 (1), 48–72.

Diamond, Peter A., 1971. A model of price adjustment. J. Econ. Theory 3 (2), 156–168.

Einav, Liran, Levin, Jonathan, 2014. The data revolution and economic analysis. Innov. Policy Econ. 14, 1–24.

Fudenberg, Drew, Villas-Boas, J. Miguel, Hendershott, T., 2006. Behavior-based price discrimination and customer recognition. In: Handbook on Economics and Information Systems, Vol. 1, pp. 377–436.

García-Pola, Bernardo, Iriberri, Nagore, Kovářík, Jaromír, 2020. Non-equilibrium play in centipede games. Games Econ. Behav. 120, 391–433.

Goldfarb, Avi, Tucker, Catherine, 2012. Shifts in privacy concerns. Am. Econ. Rev. 102 (3), 349–353.

Goldfarb, Avi, Yang, Botao, 2009. Are all managers created equal? J. Mark. Res. 46 (5), 612–622.

Hermalin, Benjamin E., Katz, Michael L., 2006. Privacy, property rights and efficiency: the economics of privacy as secrecy. Quant. Mark. Econ. 4 (3), 209–239.

Ho, Teck-Hua, Camerer, Colin, Weigelt, Keith, 1998. Iterated dominance and iterated best response in experimental 'P-beauty contests'. Am. Econ. Rev. 88 (4), 947–969.

Ichihashi, Shota, 2020. Online privacy and information disclosure by consumers. Am. Econ. Rev. 110 (2), 569–595.

Kawagoe, Toshiji, Takizawa, Hirokazu, 2012. Level-k analysis of experimental centipede games. J. Econ. Behav. Organ. 82 (2–3), 548–566.

Luo, Xueming, Andrews, Michelle, Fang, Zheng, Wei Phang, Chee, 2014. Mobile targeting. Manag. Sci. 60 (7), 1738–1756.

Mayer-Schönberger, Viktor, Cukier, Kenneth, 2013. Big Data: A Revolution That Will Transform How We Live, Work and Think. John Murray.

Mikians, Jakub, Gyarmati, László, Erramilli, Vijay, Laoutaris, Nikolaos, 2012. Detecting price and search discrimination on the Internet. In: Proceedings of the 11th ACM Workshop on Hot Topics in Networks. ACM, New York, NY, USA, pp. 79–84.

Montes, Rodrigo, Sand-Zantman, Wilfried, Valletti, Tommaso, 2018. Value of personal information in online markets with endogenous privacy. Manag. Sci. 65 (3), 1342–1362.

Nagel, Rosemarie, 1995. Unraveling in guessing games: an experimental study. Am. Econ. Rev. 85 (5), 1313–1326.

Norberg, Patricia A., Horne, Daniel R., Horne, David A., 2007. The privacy paradox: personal information disclosure intentions versus behaviors. J. Consum. Aff. 41 (1), 100–126.

Odlyzko, Andrew, 2003. Privacy, economics, and price discrimination on the Internet. In: Proceedings of the 5th International Conference on Electronic Commerce, ICEC 2003. ACM, New York, NY, USA, pp. 355–366.

Pigou, Arthur C. The Economics of Welfare. Macmillan and Co., limited, London.

Posner, Richard A., 1978. An economic theory of privacy. Regulation 2, 19–26.

Posner, Richard A., 1981. The economics of privacy. Am. Econ. Rev. 71 (2), 405–409.

Prüfer, Jens, Schottmüller, Christoph, 2021. Competing with big data. J. Ind. Econ. forthcoming.

Shiller, Reed Benjamin, 2013. First Degree Price Discrimination Using Big Data. Working Paper 58. Working Paper Series. Brandeis University, Department of Economics and International Businesss School.

Stahl, Dale O., Wilson, Paul W., 1994. Experimental evidence on players' models of other players. J. Econ. Behav. Organ. 25 (3), 309–327.

Stahl, Dale O., Wilson, Paul W., 1995. On players' models of other players: theory and experimental evidence. Games Econ. Behav. 10 (1), 218–254.

Stigler, George J., 1980. An introduction to privacy in economics and politics. J. Leg. Stud. 9 (4), 623–644.

Taylor, Curtis R., 2004. Consumer privacy and the market for customer information. Rand J. Econ. 35 (4), 631–650.
The New York Times, 2012. Acxiom, the quiet giant of consumer database marketing. N.Y. Times. June 16, sec. Technology.
Tsai, Janice Y., Egelman, Serge, Cranor, Lorrie, Acquisti, Alessandro, 2011. The effect of online privacy information on purchasing behavior: an experimental study. Inf. Syst. Res. 22 (2), 254–268.
Tucker, Catherine E., 2014. Social networks, personalized advertising, and privacy controls. J. Mark. Res. 51 (5), 546–562.
Varian, Hal R., 1997. Economic aspects of personal privacy. In: Privacy and Self-Regulation in the Information Age. U.S. Department of Commerce - National Telecommunications and Information Administration.
Villas-Boas, J. Miguel, 2004. Price cycles in markets with customer recognition. Rand J. Econ. 35 (3), 486–501.
Viscusi, W. Kip, 1978. A note on 'Lemons' markets with quality certification. Bell J. Econ. 9, 277–279.
Wall Street Journal Blog, 2014. Amazon wants to ship your package before you buy it. In: Digits – Tech News & Analysis from the WSJ. January 17. http://blogs.wsj.com/digits/2014/01/17/amazon-wants-to-ship-your-package-before-you-buy-it.
Warren, Samuel D., Brandeis, Louis D., 1890. The right to privacy. Harvard Law Rev. 4 (5), 193–203.