Consumers' Privacy Choice in the Big Data Era

SEBASTIAN DENGLER JENS PRÜFER
Tilburg University

Draft version: **May 10, 2015**
(Preliminary version, please do not distribute or quote!)

Abstract

Recent technological progress has led to the rise of big data. The availability of new datasets and search technologies enables sellers to approximate perfect price discrimination, charging every consumer individualized prices. Many consumers feel overwhelmed by the threats to their privacy and its economic consequences and face cognitive constraints when deciding about consumption and the disclosure of their personal data. We construct a model where consumers face the trade-off between using a sales channel that quickly leads them to their preferred product but collects and uses all personal data, leaving them no surplus from consumption. Alternatively, they can search anonymously, for a cost. Consumers do not fully anticipate a seller's strategic response to increase the price on the anonymous channel, which leads to some consumers ending up with negative, others with positive surplus, thereby endogenizing preferences for privacy. We demonstrate that the anonymous channel breaks down if consumers are *too* sophisticated and discuss different interpretations and resulting policy measures.

JEL Codes: L1, L5
Keywords: Privacy, Big Data, First Degree Price Discrimination

1 Introduction

> "Few consumers have ever heard of ***Acxiom***. But analysts say it has amassed the world's largest commercial database on consumers and that it wants to know much, much more. Its servers process more than 50 trillion data 'transactions' a year. Company executives have said its 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 Times 2012)*
>
> ***Shopkick*** offers a smartphone app that rewards users for checking into stores, scanning products, visiting the dressing rooms, and so forth. Founded in 2009, it has not only brought more than USD 1 billion in revenues for its corporate partners via generating more than 50 million walk-ins to partner stores and 100 million product scans (Shopkick 2014); with over 6 million users spending more than 3 hours per month it is also the most-used shopping app, according to Nielsen (2012).
>
> *(Shopkick 2014 and Nielsen 2012)*
>
> ***Amazon*** recently was issued a patent on a novel *Method and System for Anticipatory Package Shipping*.[1] "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. [...] [T]he patent demonstrates one way Amazon hopes to leverage its vast trove of customer data to edge out rivals. [...] Based on all the things they know about their customers they could predict demand based on a variety of factors."
>
> *(Wall Street Journal Blog 2014)*

These contemporary business cases exemplify two recent technological developments. On the one hand side, firms get better and better in drawing relevant information about certain groups of people or even individuals out of huge data sets. On the other side, such data sets are increasingly available, owing to the fact that more and more economic and social transactions take place aided by information and communication technologies (ICT), which easily and inexpensively store the information they produce or transmit. Taken together, these developments constitute the rise of big data (Mayer-Schönberger and Cukier 2013). They imply that sellers can make consumers ever more tailored contract offers, which fit their individual preferences or consumption patterns.[2]

One logical consequence of this process is that sellers can approximate first-degree price discrimination better and better. First-degree (or perfect) price discrimination is characterized by complete information of a seller about a specific consumer's willingness to pay for a certain

[1] Patent number US008615473 (December 24, 2013), http://pdfpiw.uspto.gov/.piw?docid=08615473

[2] Such offers can be made directly, for instance, in online retailing, or indirectly, via selling advertisers access to highly preselected consumer groups. Einav and Levin (2013) provide a list of illustrative examples how firms, public administration, and researchers can exploit such novel technological opportunities.

product (Pigou 1920). This information leads to the seller's ability to appropriate all surplus of the transaction, assuming that reselling is impossible or unprofitable, because she can set a price that just equals the consumer's valuation of the product. However, due to the very high information demand of the seller about consumers' preferences and due to the straightforward allocative and distributional implications, perfect price discrimination has not received a lot of attention in the economics literature and has mostly been dismissed as a mere theoretical construct.[3]

More prominent are models of so-called "behavior-based price discrimination." Most of this literature focuses on second-degree price discrimination by assuming that a seller usually learns about the willingness-to-pay of a re-identifiable or recognizable consumer after the first purchase of a good. The idea is that, if a consumer previously bought a product at a certain price, the seller would learn that this particular consumer's willingness to pay must have exceeded the price for which she bought the product.[4] However, as the introductory examples above illustrate, online vendors and other retailers have already gone much further and can approximate fully personalized prices more and more, which supports 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, 365).

It has been shown empirically that "targeted advertising" techniques, which make use of the vast amounts of information that sellers can get hold of about potential buyers today, increases purchases (Luo et al. 2014), prices (Mikians et al. 2012), and sellers' profits (Shiller 2013). Some consumers, however, feel repelled by this development, which assigns a passive role to them, facing apparently omniscient sellers who can exploit all their digital traces.[5] Many want to have control over their personal data back.[6] 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; Posner 1981; Varian 1997). A lot of progress in our understanding has been made since then.[7] Most research articles have focused on the choices of firms that own some type of personal

[3] For instance, the standard industrial organization textbook, Tirole (1988), spends three of its more than 1100 pages on perfect price discrimination.

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

[5] Shiller (2013, 5) reports: "Even if consumers did understand which behaviors result in low prices, they might prefer to ignore them rather than change potentially thousands of behaviors just to receive a lower quoted price for one product."

[6] 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. Brandimarte et al. (2012, 6) conclude: "Three experiments provide empirical evidence that perceived control over release plays a critical role in sharing/oversharing personal information, relative to the objective risks associated with information access and usage by others."

[7] Already Hermalin and Katz (2006, 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."

information about consumers and can decide to disclose it to another firm, or not (Taylor 2004; Acquisti and Varian 2005; Calzolari and Pavan 2006; Casadesus-Masanell and Hervas-Drane 2014). The core question studied in these papers is, what the welfare consequences of privacy or disclosure are, and who should own the property rights of consumers' personal data (Hermalin and Katz 2006).[8] The answers given have been ambiguous and depend on the specific application of the papers. Recently, the focus has shifted more towards the privacy choices of consumers (Conitzer et al. 2012) and to the role of platform intermediaries (de Corniere and De Nijs 2014).

Regarding these choices, however, Acquisti and Grossklags (2007, 369) note: "Consumers will often be overwhelmed with the task of identifying possible outcomes related to privacy threats and means of protection. [...] However, even if individuals had access to complete information, they would often be unable to process and act optimally on large amounts of data. 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. Bounded problem solving is usually neither unreasonable nor irrational, and it doesn't need to be inferior to rational utility maximization." The need to include cognitive constraints into economic models of privacy is spurred by empirical findings about the so-called privacy paradox.[9] With few exceptions, however, cognitive constraints of consumers have not been incorporated by theoretical studies of markets driven by big data.[10] In this paper, we account for the difference between sellers, whose data analysis capabilities outperform human capabilities by far, and consumers with limited cognitive sophistication. Combining both, the rising ability to perfectly price discriminate as well as cognitive limitations, we ask: What are the effects of perfect price discrimination on equilibrium behavior and welfare, when anonymization is possible but costly? And moreover, how are these effects affected by the strategic sophistication of consumers?

We therefore consider a market where a mass of consumers with heterogeneous willingness to pay for a product is facing a monopolistic seller who is equipped with two sales channels. Consumers can decide which channel to use to acquire product information and buy the product. The default channel (D) makes use of all personal information that the seller has about every single consumer, be it via collecting such information in the past (e.g. Amazon) or via buying such information from an intermediary (e.g. Google, Acxiom). We assume that the seller can perfectly discriminate prices when selling through D, but that D is transaction cost-economizing,

[8] On the Internet, for instance, the customer databases of sellers or intermediaries, such as search engines, tracing back the physical address of users on the basis of their IP address (or to clearly identify them as persons on the basis of their registration data or a unique identifier derived from a permanent cookie) was recently qualified as personal data (Opinion 1/2008 on Data Protection Issues Related to Search Engines, Advisory Working Party (adopted Apr. 4, 2008) (EC), Data Protection available at http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2008/wp148_en.pdf).

[9] A series of experimental research has shown that consumers stated and revealed valuations of their own personal data differ highly and depend on the framing of the survey questions. (Acquisti, John, and Loewenstein 2009; John, Acquisti, and Loewenstein 2009; Jentzsch, Preibusch, and Harasser 2012;).

[10] Taylor (2004), Acquisti and Varian (2005), and Armstrong (2006) assume the existence of a group of perfectly rational 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 negligence, ignore the negative effects of disclosing personal data. In our model, we allow for more nuanced, marginal analysis of consumers sophistication

in the sense that a consumer can shop without bearing additional cost for search, etc. Denote the individual price faced by consumer i in this channel by p_i.

The alternative channel (A) protects consumer privacy by anonymizing the search choices of individual users and does not track their behavior (e.g. Ixquick/Startpage, DuckDuckGo, or shopping offline with cash). As a consequence, perfect price discrimination is infeasible for the seller who consequently sets a uniform price for this channel, denoted by p_A. The disadvantage for using channel A is that it foregoes utilizing all information about consumers' choices, which decreases the precision of its search results compared to channel D, which results in extra time or effort that a consumer must spend to find her preferred product.[11] Moreover, using this channel might also require the consumer to first change privacy settings, install potentially costly privacy protective software, or use other privacy-enhancing techniques. Call the entirety of the additional costs associated with the anonymous channel s.

In our model, at Stage 1 each consumer decides about the sales channel to use. At Stage 2, the seller decides about the product price. At Stage 3, every consumer decides whether to buy for the price offered to her, or not. We solve this game by backward induction for a Sophistication-k equilibrium (defined in Binswanger and Prüfer 2012) since it turns out that the strategic sophistication of consumers plays a critical role for the solution. Our analysis is based on a model of limited strategic thinking, called *level k-thinking*. This model was introduced by Stahl and Wilson (1994; 1995) and Nagel (1995). Since then, a sizeable literature has developed that explores k-thinking theoretically and empirically, including Ho et al. (1998), Costa-Gomes et al. (2001), Crawford (2003), Camerer et al. (2004), Costa-Gomes and Crawford (2006), Crawford and Iriberri (2007a), Goldfarb and Yang (2009), and Binswanger and Prüfer (2012), among others. The literature has found strong experimental support for level k-thinking. Thus, we will assume that consumers face cognitive constraints expressed by their ability to anticipate k strategic iterations. The seller on the other hand – due to superior access to data and computing power – is able to outperform them in strategic thinking (i.e. has a level of at least $k+1$) and hence always employs the optimal response to the consumers' strategy.

The empirical evidence suggests values for k of one or two (Camerer, Ho, and Chong 2004; Crawford and Iriberri 2007b). This stands in sharp contrast to the usual assumptions in the behavior-based price discrimination literature of either unlimited strategic sophistication or complete naïveté of consumers.[12] Whether k is to be seen as a low number, as suggested by the empirical behavioral literature, or rather infinitely high turns out to crucially matter for our results. First, we find that the higher the consumers' degree of sophistication, the higher the equilibrium price will be on the anonymized market of channel A. If consumers are sophisticated, they anticipate the incentive of the seller to increase the price beyond the price expected by less sophisticated consumers. Therefore, if the level of sophistication rises in the population,

[11] See the literature cited in Argenton and Prüfer (2012) documenting the effect from access to more search log data on the quality of search engines as perceived by users.

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

more consumers (those with medium but not high willingness to pay) will expect to receive a negative payoff from using the anonymized market and choose sales channel D at Stage 1. The seller, in turn, has an incentive to increase the price in the anonymized market even more because he infers that only consumers with high willingness to pay, have chosen channel A at Stage 1.

We further show that, with any positive cost of anonymization s, the anonymized market completely unravels for all sophistication levels $k \geq \overline{k}$, where $\overline{k}$ is a finite number (and hence also covers the case of $k = \infty$ where consumers have unlimited strategic sophistication). If consumers have rather binding cognitive constraints, however, only a part of the market unravels and the anonymized sales channel can persist, serving consumers with relatively high willingness to pay. Among those who use the privacy-protecting sales channel, some consumers suffer from net losses because prices turn out to be higher than expected, but consumers with a very high willingness to pay get some surplus. Thereby, this paper delivers a micro-foundation of consumers' privacy preferences, as consumers can rationally want to use anonymization techniques (or their data erased) even without an explicit "taste for privacy". Thereby, we contribute to the discussion evolving around the recent ruling of the Google Spain case of the Court of Justice of the European Union (2014),[13] a topic that has been conversely debated on both sides of the Atlantic.

Because the anonymized sales channel exists in equilibrium and a share of the anonymization cost s could be interpreted as a (royalty) fee that the intermediary can appropriate, this model also suggests that running a consumer privacy-protecting sales channel that competes with a channel which tracks individuals and uses all personal data can be a sustainable business model in an economy populated by consumers with limited strategic sophistication.

2 Model

2.1 Setup

We consider an economy where a monopolistic seller of a single consumption good faces a continuum of consumers who can buy at most one unit of the good and cannot resell it to each other. Abstracting from potential fixed costs, we assume that the monopolist can produce the good at constant marginal cost $c > 0$. Consumers have a heterogeneous valuation v_i for the good, $v_i \sim \mathcal{U}[\underline{v}, \bar{v}]$, where $0 \leq \underline{v} \leq c < \bar{v}$. We assume that the seller knows the valuation, i.e. the maximum willingness to pay, of each consumer who chooses the default channel D to buy the good. However, each consumer can choose to use channel A, which provides a search anonymization technique, but requires additional search cost s, $0 < s$. If a consumer chooses to do so, her valuation will not be known to the seller. Further, we assume that all agents are solely interested in their own material payoff (i.e. monetary profit or consumption utility reduced by all eventually incurred costs). Specifically, consumers do not have any exogenous "taste for privacy". The distribution of v_i, the cost for anonymization s, the cost structure of the monopolist (and hence

[13] http://curia.europa.eu/juris/liste.jsf?num=C-131/12

the supply function) as well as the timing of the game are common knowledge among all agents. We further assume that levels of strategic sophistication, denoted by k, differ between the group of consumers and the seller but not within the group of consumers. Further we assume that if the level of strategic sophistication of the consumers is given by k the seller will have a level of at least $k+1$ as reflection of his large database of customer information and higher computing power.[14] Our model therefore describes the situation *after* a long period of behavior-based price discrimination interactions, during which the consumers have not employed anonymization techniques, reflecting the situation in today's big data-driven markets.
There are three stages in our model and the timing is as follows:

- Stage 1 (*Anonymization Decision*): Based on price expectations each consumer decides whether to use channel D or A, and incurs product search costs of 0 or s, respectively.
- Stage 2 (*Pricing Decision*): The seller sets prices $p = \{p_i, p_A\}$ where p_i are individual prices for each channel D consumer, and p_A is the uniform price p_A for all channel A consumers.
- Stage 3 (*Buying Decision*): Each consumers decides whether to buy the good for the price the seller has set for her.

Given the finite horizon of the game, sub-game perfect Bayesian Nash equilibria could be found by backward induction and application of Bayes' law. However, due to the limited strategic sophistication of consumers we will employ a modified version of the concept of sophistication-k equilibrium as defined in Binswanger and Prüfer (2012):[15]

Definition 1 (Sophistication-k Equilibrium) *A sophistication-k equilibrium is a strategy combination and a set of beliefs about the state of nature and about the behavior of the other player, such that at each node of the game between a level-k player (the consumer) and a level-k+1 player (the seller):*

1. *The strategies for the remainder of the game are Nash given the beliefs and strategies of the other player;*
2. *The level-k player holds a k-belief about the behavior of the other player;*
3. *The k+1 player anticipates the belief of the level-k player.*

As part 1 of this equilibrium concept requires that strategies be Nash equilibria given the beliefs and strategies of the other player, we will still use backward induction.By doing so, we will see about which components of the game the consumers and the seller will have to form expectations and which are therefore dependent on their respective level of k. Hence, we will start our analysis by looking at Stage 3.

[14] Since we have already assumed that the seller can use his database and computing power in the non-anonymized channel D to perfectly infer every consumer's valuation, this assumption about the seller's abilities with respect to strategic sophistication follows almost by definition.

[15] As our model does not include a draw of a state nature, we do not need the fourth part of their definition, which states: "The beliefs about the state of nature are rational and determined by Bayes' law."

2.2 Equilibrium Analysis

Stage 3 – Buying Decision: A utility maximizing consumer decides to buy the product if and only if her valuation of the good exceeds the price she has to pay for it, i.e. if and only if $v_i \geq p$.[16] Under our assumption of $v_i \sim U[\underline{v}, \bar{v}]$ the cumulative distribution function $F(p) = \frac{p-\underline{v}}{\bar{v}-\underline{v}}$ gives the mass of consumers for whom v_i is smaller than a particular price p. As only consumers whose v_i equals or exceeds p buy the product, demand is given by the complementary mass $q(p) = 1-F(p) = \frac{\bar{v}-p}{\bar{v}-\underline{v}}$ and inverse demand $p(q) = q^{-1}(1-F(p))$ accordingly by $p(q) = \bar{v}-(\bar{v}-\underline{v})q$. Having derived demand, we can now proceed to analyze the seller's pricing decision in Stage 2.

Stage 2 – Pricing Decision: A profit maximizing seller will set individual prices p_i for all consumers approaching him via channel D (denoted by $\mathcal{C}_D$) and one optimal uniform price p_A for all anonymized consumers in channel A (denoted by $\mathcal{C}_A$). Knowing v_i precisely for all consumers in $\mathcal{C}_D$ the seller trivially sets $p_i^* = \max\{v_i, c\}$ for all $i \in \mathcal{C}_D$, taking into consideration that $c \geq \underline{v}$ might imply not selling to some consumers (should $c > \underline{v}$ hold), as pricing below marginal cost would lead to losses. Being uninformed about the individual valuations v_i of all $i \in \mathcal{C}_A$, the seller can only infer which consumers form the set of anonymized consumers $\mathcal{C}_A$ by strategic reasoning about consumer behavior in Stage 1 and set p_A accordingly. We will therefore first have to have a look at consumers' general Stage 1 behavior to inform the seller's pricing decision for channel A.

Stage 1 – Anonymization Decision: Consumers will choose to use the anonymization technique of channel A if and only if the expected utility of doing so exceeds the expected utility of the default channel D, i.e. if and only if $\mathbb{E}(u_i(A)) \geq \mathbb{E}(u_i(D))$. The expected utility of each option is determined as follows:

$$\mathbb{E}(u_i(D)) = \max\{v_i - \mathbb{E}(p|D), 0\} \tag{1}$$

$$\mathbb{E}(u_i(A)) = \max\{v_i - \mathbb{E}(p|A) - s, -s\} \tag{2}$$

where the first value of each pair reflects the payoff the consumer were to receive in Stage 3 should she choose to buy the product and the second value accordingly reflects the payoff of subsequently choosing not to buy the product. Knowing the timing and structure of the game, consumers can always anticipate that it is optimal for the seller to set the price equal to their valuation when he has access to this information. Hence, consumers form the price expectation for channel D correctly $\mathbb{E}(p|D) = p_i^* = \max\{v_i, c\}$, irrespective of their level of strategic sophistication k. Therefore, they expect to be left with no surplus when choosing channel D. However, with respect to channel A consumers only know that the seller has to set a uniform price, but which price exactly they expect depends on their level of strategic sophistication. For the sake of a general result we replace the expectation $\mathbb{E}(p|A)$ only by $\mathbb{E}(p_A)$ rather than the specific price

[16] For the two possible histories $\{D; p_i\}$ and $\{A; p_A\}$ this implies that the consumer buys if and only if $v_i \geq p_i$, respectively if and only if $v_i \geq p_A$.

which consumers with a particular level k will expect:

$$\mathbb{E}(u_i(D)) = \max\{v_i - \max\{v_i, c\}, 0\} = 0 \tag{3}$$

$$\mathbb{E}(u_i(A)) = \max\{v_i - \mathbb{E}(p_A) - s, -s\} \tag{4}$$

Comparing these expected payoffs shows that consumers will choose channel A if and only if $\max\{v_i - \mathbb{E}(p_A) - s, -s\} \geq 0$.[17] This can only hold if $v_i \geq \mathbb{E}(p_A) + s$, because we assumed $s > 0$, which means there is a general threshold dividing the population of consumers into $\mathcal{C}_D$ and $\mathcal{C}_A$.

Lemma 1 (Anonymization Threshold) *There exists a threshold $\hat{v} = \mathbb{E}(p_A) + s$, above which consumers will prefer channel A to channel D, which shows $\mathcal{C}_D = [\underline{v}, \hat{v})$ and $\mathcal{C}_A = [\hat{v}, \bar{v}]$.*

Stage 2 – Pricing Decision (revisited): Given the utility maximizing behavior of consumers in Stage 1, we can now return to the seller's decision about which uniform price p_A to set for channel A. Having higher strategic sophistication than the consumers, the seller can infer $\hat{v}$ and hence knows that $\mathcal{C}_A = [\hat{v}, \bar{v}]$. This implies that the demand in channel A is given by:

$$q_A(p) = \begin{cases} \frac{\bar{v}-p}{\bar{v}-\underline{v}} & \textit{if } p \geq \hat{v} \\ \frac{\bar{v}-\hat{v}}{\bar{v}-\underline{v}} & \textit{otherwise} \end{cases} \tag{5}$$

It becomes immediately clear that charging a price below $\hat{v}$ in channel A will never be optimal for the seller as it would not change demand but only lead to foregone profits. However, depending on the level of $\hat{v}$, the opposite need not hold true.

Denote the price that would maximize the seller's profits if the entire consumer population was approaching him via channel A by $p_M = \frac{\bar{v}+c}{2}$, (which is identical to the regular monopoly price). This price p_M then provides another lower bound to the seller's optimal pricing strategy, which can be easily seen: If it is not optimal to charge below p_M for the entire population, then it can not be optimal to charge below p_M for the subset $\mathcal{C}_A$ of the population (which is only excluding people with a low valuation). Furthermore, the regular monopoly price p_M is also the optimal price as long as it is the binding lower bound (i.e. $p_M \geq \hat{v}$). To see this, remember that moving away from the monopoly price in either direction reduces profits when serving the entire population (if demand is concave or at least not *too* convex). Hence, increasing the price above p_M also reduces profits when the monopolist faces the subset $\mathcal{C}_A$ as long as doing so reduces demand in channel A.[18] Thus, as long as $\hat{v} \leq p_M$, the optimal price is given by $p_A^* \geq p_M$, whereas as soon as $\hat{v} > p_M$ we still only know that $p_A^* \geq \hat{v}$.

[17] We assume that consumers that are indifferent between channel A and D will choose A as a tie-breaking rule.

[18] Especially it is not the case that the increased lower bound of willingness to pay in the subset $\mathcal{C}_A$ compared to the entire population provides an incentive for the seller to increase the price, as it is irrelevant how many people he does not serve (i.e. how far below the optimal price p_M the demand function in channel A extends). To see this, recall that $p_M = \frac{\bar{v}+c}{2}$ does not depend on the lower bound of valuation for the good even with the entire population being in channel A, but on the marginal cost instead.

Combining the two insights that it is neither optimal to price below $\hat{v}$ as this only forgoes profits without affecting demand nor to moving away from p_M in either direction as long as this changes demand, we can easily deduce that for the case where $\hat{v} > p_M$ the optimal price is equal to the lower bound, i.e. $p_A^* = \hat{v}$.

The complete optimal pricing strategy of the seller for both channels is summarized in the following proposition:

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

Note that this optimal pricing strategy implies that the seller sets a potentially higher price than consumers had expected, when $p_A^* = \hat{v} = \mathbb{E}(p_A) + s > \mathbb{E}(p_A)$ which stems from the fact that s is going to be a sunk cost for consumers at Stage 3 and the seller's ability to anticipate this fact.

Stage 1: Anonymization Decision (revisited)

Having found the seller's optimal pricing strategy for any given $\hat{v} = \mathbb{E}(p_A) + s$, the last missing piece to fully characterize equilibrium behavior is the formation of consumers' expectations $\mathbb{E}(p_A)$ in Stage 1. As outlined earlier, we describe this by level-k thinking, which is best determined recursively and thus we will start with the case of $k = 0$.

Consumers with $k = 0$ are referred to as "naïve" consumers and in some environments it is suitable to treat them as if they were uniformly randomizing about the range of options. However, we will not employ this reasoning here as we already imposed that even consumers with $k = 0$ are able to understand the basic structure of the game (even anticipating the perfect price discrimination in channel D). Therefore, they especially know that they are dealing with a monopolistic seller and hence will "naïvely" expect the seller to engange in regular monopoly pricing and hence form $\mathbb{E}_0(p_A) = p_M$ as their price expectation ($\mathbb{E}_k(p_A)$ being the shorthand notation for $\mathbb{E}(p_A|k)$).[19] This leads to the next proposition:

Lemma 3 (Sophistication-0 Equilibrium) *For "naïve" consumers and anonymization cost s there is a sophistication-0 equilibrium with the following characteristics:*

1. *Consumers form the 0-beliefs $\mathbb{E}_0(p|D) = p_i^*$ and $\mathbb{E}_0(p|A) = p_M$ and choose channel A if and only if $v_i \geq p_M + s = \hat{v}_{0,s}$ and channel D otherwise.*
2. *The seller's best response to consumers' 0-beliefs is to charge prices $p_i^* = \max\{v_i, c\}$ and $p_{A_{0,s}}^* = \hat{v}_{0,s} = p_M + s$.*
3. *All consumers with $v_i \geq c$ buy the product (either via channel D or channel A)*

[19] Starting with uniform randomization would not change the major outcomes, but make the point less clear. Since either assumption is supportable, we choose in favor of expositional simplicity.

We see that there is a difference of exactly s between consumers' expectations of the price p_A in channel A, $\mathbb{E}_0(p_A)$, and the optimal price, $p^*_{A_{0,s}}$, the seller eventually sets given these expectations (note that we have included subscripts k and s for p^*_A and $\hat{v}$, and to remember that they depend on the parameters whose marginal effects will be studied in section 2.3). Due to their limited capabilities in strategic reasoning the consumers do not foresee that the seller will make use of the fact that once any consumer reaches Stage 3 the anonymization expenses of s will have turned into sunk costs. This result in turn informs us about the way in which consumers form their price expectation for any higher level of strategic sophistication $k > 0$.

If consumers have a strategic sophistication level of $k = 1$ instead of $k = 0$, they will be capable of one strategic iteration and hence anticipate that the optimal price in channel A given $k = 0$ is given by $p^*_{A_0} = p_M + s$. Thus, they form the 1-belief $\mathbb{E}_1(p_A) = p_M + s$ which will lead the seller to set the optimal price of $p^*_{A_1} = p_M + 2s$. This in turn leads to the 2-belief $\mathbb{E}_2(p_A) = p_M + 2s$ and so forth. Hence, we can write more generally for any level of k that

$$\mathbb{E}_k(p_A) = p_M + ks \tag{6}$$

$$p^*_{A_k} = p_M + (k+1)s = \hat{v}_{k,s} \tag{7}$$

which implies the equivalence $\mathbb{E}_k(p_A) = p^*_{A_{k-1}}$.

Proposition 1 (General sophistication-k equilibrium) *For consumers with limited strategic sophistication k and anonymization cost s there is a sophistication-k equilibrium with the following characteristics:*

1. *Consumers form the k-beliefs $\mathbb{E}_k(p|D) = p^*_i$ and $\mathbb{E}_k(p|A) = p_M + ks$ and choose channel A if and only if $v_i \geq \hat{v}_{k,s} = p_M + (k+1)s$ and channel D otherwise.*
2. *The seller's best response to consumers' k-beliefs is to charge prices $p^*_i = \max\{v_i, c\}$ and $p^*_{A_{k,s}} = \hat{v}_{k,s} = p_M + (k+1)s$.*
3. *All consumers with $v_i \geq c$ buy the product (either via channel D or channel A)*

Thus, while consumers with a relatively high consumption utility choose channel A and hence get anonymized, consumers with relatively low consumption utility ($v_i < p_M + (k+1)s$) will choose to stay in the default channel and be perfectly price discriminated against. It is worth noting that a (potentially large) share of consumers with low consumption utility ($v_i < p_M$) stays in the default channel D irrespectively of k and s as they cannot possibly hope to get a uniform price via channel D that is affordable for them.[20]

This iterative process comes to an end when $\mathbb{E}_{\overline{k}}(p_A) > \bar{v} - s$, as this $\overline{k}$-belief causes the anonymization threshold $\hat{v}_{\overline{k},s} = p_M + (\overline{k}+1)s$ to exceed $\bar{v}$ which in turn induces $\mathcal{C}_A = \emptyset$. Then, channel A remains unused and the market for anonymization completely breaks down. Solving

[20] Those consumers are the ones that account for the deadweight loss associated with monopolistic markets in which there is no possibility for price discrimination

$\hat{v}_{\overline{k},s} = p_M + (\overline{k}+1)s > \bar{v}$ for $\overline{k}$ and recalling that $0 < s$, we see that this point will be reached at a finite $\overline{k} > \frac{\bar{v}-p_M}{s} - 1$. As $\overline{k}$ has to be a natural number we have to round up for the inequality to hold. In the case where $\frac{\bar{v}-p_M}{s} - 1$ is a natural number itself, however, $\overline{k}$ is given by the next higher natural number:

$$\overline{k} = \begin{cases} \left\lceil \frac{\bar{v}-p_M}{s} - 1 \right\rceil & \text{if } \frac{\bar{v}-p_M}{s} - 1 \notin \mathbb{N} \\ \frac{\bar{v}-p_M}{s} & \text{if } \frac{\bar{v}-p_M}{s} - 1 \in \mathbb{N} \end{cases} \tag{8}$$

We thereby implicitly solve for the limit case of $k = \infty$ which represents an approach with unlimited strategic sophistication and delivers the same result as $k = \overline{k}$ with $\mathcal{C}_A = \emptyset$. As a corollary, we can note that a market for anonymization, although costly, can exist as long as k is sufficiently low, i.e. as long as consumers are not *too* strategically sophisticated (given that s is not prohibitively high even for $k = 0$). We summarize these arguments without formal proof: It can be shown that

Corollary 1 *$\mathcal{C}_A \neq \emptyset$ if and only if $k < \overline{k} < \infty$.*

2.3 Welfare analysis

Further investigating the effects of changes in k and s we continue with a short analysis of consumer surplus ($CS_{k,s}$), profits (π), and welfare (W) associated with the level-k sophistication equilibrium behavior. We employ the customary understanding of total welfare as $W = CS_{k,s} + \pi_{k,s}$ and thus abstract from any potential preferences by a social planner (or policy-maker) for either side of the market. Consumer surplus is given by $CS_{k,s} = CS_{D_{k,s}} + CS_{A_{k,s}}$ where the subscripts denote the respective channel chosen by the consumers. As the seller engages in perfect price discrimination for channel D we readily know that $CS_{D_{k,s}} = 0$. Recalling that consumers in $\mathcal{C}_A$ have incurred the cost s to anonymize, their surplus can then be expressed as $CS_{k,s} = CS^+_{k,s} - CS^-_{k,s}$ where $CS^+_{k,s}$ is the aggregate surplus from the transaction with the seller (i.e. the buy decision) and $CS^-_{k,s}$ represents the aggregate cost of anonymization. We then get

$$CS^+_{k,s} = \int_{q(\bar{v})}^{q(\hat{v}_{k,s})} p(q) - p^*_{A_{k,s}} \,\mathrm{d}q = \int_0^{\frac{\bar{v}-\hat{v}_{k,s}}{\bar{v}-\underline{v}}} p(q) - p^*_{A_{k,s}} \,\mathrm{d}q = \frac{1}{2}\frac{(\bar{v}-\hat{v}_{k,s})^2}{\bar{v}-\underline{v}} \tag{9}$$

$$CS^-_{k,s} = \int_{q(\bar{v})}^{q(\hat{v}_{k,s})} s \,\mathrm{d}q = \int_0^{\frac{\bar{v}-\hat{v}_{k,s}}{\bar{v}-\underline{v}}} s \,\mathrm{d}q = s \cdot \frac{\bar{v}-\hat{v}_{k,s}}{\bar{v}-\underline{v}} \tag{10}$$

and consequentially

$$CS_{k,s} = \frac{1}{2}\frac{(\bar{v}-\hat{v}_{k,s})^2}{\bar{v}-\underline{v}} - s \cdot \frac{(\bar{v}-\hat{v}_{k,s})}{\bar{v}-\underline{v}} \tag{11}$$

In Figure 1 (both panels) the first term, $CS^+_{k,s}$, is graphically represented by the solid grey (upper) triangle. Partially overlapping this triangle is the dashed rectangle that represents the

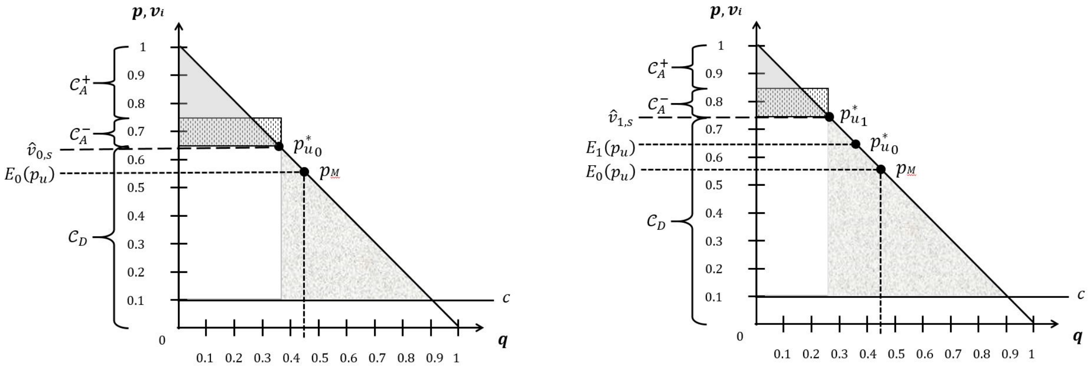

(a) $k = 0$ (b) $k = 1$

Figure 1: Welfare analysis with parameters $v_i \sim \mathcal{U}[0,1]$, $c = 0.1$, $s = 0.1$

second term of the equation, $CS^-_{k,s}$. The fact consumer surplus not only depends on s (which is immediately apparent), but via $\hat{v}_{k,s}$ as well on k can easily be seen from the difference between Figure 1a and Figure 1b. The precise effects [21] are given as follows:

$$\Delta CS_{k,s} = CS_{k+1,s} - CS_{k,s} = -\frac{s}{\bar{v} - \underline{v}} \cdot \left(\bar{v} - \left(\hat{v}_{k,s} + \frac{3}{2}s\right)\right) \tag{12}$$

$$\frac{\partial CS_{k,s}}{\partial s} = -\frac{k+2}{\bar{v} - \underline{v}} \cdot \left(\bar{v} - p_M - \frac{(k+3)(k+1)}{(k+2)} s\right) \tag{13}$$

Equation (12) shows that consumer surplus is decreasing in k as long as $\bar{v} > \hat{v}_{k,s} + \frac{3}{2}s$ which is closely related to the condition from which we derived $\overline{k}$. As long as channel A is used, consumer surplus is decreasing in k, with the notable exception that in the moment where an increase in k leads to complete market breakdown, consumer surplus increases from a negative value to zero. This is explained by the fact that at $k = \overline{k} - 1$ all consumers in $\mathcal{C}_A$ incur net losses, due to the unanticipated sunk cost – and those losses are countered by everyone getting perfectly discriminated at $k = \overline{k}$. Hence, $CS_{\overline{k}-1,s} < 0$ and $CS_{\overline{k},s} = 0$, which is why $\Delta CS_{\overline{k}-1,s} > 0$. More plainly, consumers lose more and more surplus the more strategically sophisticated they become as a group until there is nothing left to lose and everyone "gives in" to the price discrimination and choses channel D.

The second result seems more straightforward. From Equation (13) we can see that consumer surplus is decreasing s as long as $s < (\bar{v} - p_M)\frac{(k+2)}{(k+3)(k+1)}$. Noting that the rightmost factor in this equation, $\frac{(k+2)}{(k+3)(k+1)}$, is bounded from above by $\frac{2}{3}$ at $k = 0$, this condition simply tells us that consumer surplus is decreases in s, unless it is already prohibitively costly at the current level of k, which should come as no surprise.

[21] We only explicitly calculate the effects for any non-boundary k as the explicit mathematical analysis taking into account the special cases of $\overline{k}$ does not add anything beyond the discussion in the text.

Looking at the other side of the market, we get profits in each channel $\pi_{D_{k,s}}$ and $\pi_{A_{k,s}}$ as:

$$\pi_{D_{k,s}} = \int_{q(\hat{v}_{k,s})}^{q(c)} p(q) - c \, \mathrm{d}q \qquad = \int_{\frac{\bar{v}-\hat{v}_{k,s}}{\bar{v}-\underline{v}}}^{\frac{\bar{v}-c}{\bar{v}-\underline{v}}} p(q) - c \, \mathrm{d}q \qquad = \frac{1}{2}\frac{(\hat{v}_{k,s}-c)^2}{\bar{v}-\underline{v}} \tag{14}$$

$$\pi_{A_{k,s}} = \int_{q(\bar{v})}^{q(\hat{v}_{k,s})} p^*_{A_k} - c \, \mathrm{d}q \qquad = \int_{0}^{\frac{\bar{v}-\hat{v}_{k,s}}{\bar{v}-\underline{v}}} p^*_{A_k} - c \, \mathrm{d}q \qquad = \frac{(\bar{v}-\hat{v}_{k,s})(\hat{v}_{k,s}-c)}{\bar{v}-\underline{v}} \tag{15}$$

and consequentially total profits $\pi_{k,s} = \pi_{D_{k,s}} + \pi_{A_{k,s}}$ are given by:

$$\pi_{k,s} = \frac{1}{2}\frac{(\hat{v}_{k,s}-c)^2}{\bar{v}-\underline{v}} + \frac{(\bar{v}-\hat{v}_{k,s})(\hat{v}_{k,s}-c)}{\bar{v}-\underline{v}} \tag{16}$$

The first term, $\pi_{D_{k,s}}$, is depicted in Figure 1 by the white rectangle, whereas the second term, pi_A, is indicated by the grained grey (lower) triangle. The difference between the white rectangle in Figure 1a and Figure 1b nicely illustrates the earlier point about the effect that increasing p_A will have on the seller's profits as the relation of width to length of the rectangle moves further away from identity (which would render it the typical square of a monopolist's profits). Similar to consumer surplus profits also depend on k and s (hidden in $\hat{v}_{k,s}$ in Equation (16)) and are affected by changes in those variables as follows:[22]

$$\Delta\pi_{k,s} = \pi_{k+1,s} - \pi_{k,s} = \frac{s}{\bar{v}-\underline{v}}(\bar{v} - (\hat{v}_{k,s} + \frac{1}{2}s)) \tag{17}$$

$$\frac{\partial \pi_{k,s}}{\partial s} = \frac{k+1}{\bar{v}-\underline{v}}(\bar{v} - \hat{v}_{k,s}) \tag{18}$$

Equation (17) shows that profits are increasing in k as long as $\bar{v} > \hat{v}_{k,s} + \frac{1}{2}s$, which is exactly one level of k higher than the condition for consumer surplus in Equation (12) as profits still increase with the last group of consumers switching to channel A, whereas consumer surplus stops decreasing already when this change occurs.

Equation (18) reveals that profits are increasing in both s as long as $\bar{v} > \hat{v}_{k,s}$, which can be rewritten as $s < (\bar{v} - p_M)\frac{1}{(k+1)}$ to match the condition derived for consumer surplus. It follows that profits are increasing in s unless it is already at a prohibitive level given the sophistication of consumers (clearly the rightmost factor is bounded from above here as well, namely by 1 at $k = 0$). Since increasing either variable in general will lead to more consumers preferring channel D, which is strictly more profitable for the seller than channel A where he has to leave some surplus to the consumers due to his inability to perfectly discriminate anonymized consumers.

As both k and s have a negative marginal effect on consumer surplus, but a positive marginal effect on profits, it is important to check how total welfare is affected by changes in both variables.

[22] Again, we only explicitly calculate the effects for any non-boundary k and give a verbal analysis of the special cases in the text.

Recalling that we defined welfare to be given by $W = CS_{k,s} + \pi_{k,s}$, we get the following:

$$W = \frac{1}{2}\frac{(\bar{v}-c)^2}{\bar{v}-\underline{v}} - s \cdot \frac{(\bar{v}-\hat{v}_{k,s})}{\bar{v}-\underline{v}} \tag{19}$$

The first term of the last expression in Equation (19) corresponds to the whole area between the demand curve and the marginal cost curve in Figure 1 and represents the entire surplus generated from all buying transactions (via both channels). The second term, is again the dashed rectangle and denotes the loss generated by the anonymization behavior of consumers. Hence, it should not be surprising that welfare is generally increasing in k (see Equation (20)), since this always leads to fewer anonymized consumers (unless we cross $\overline{k}$, of course).

$$\Delta W_k = W_{k+1} - W_k = \frac{s^2}{(\bar{v}-\underline{v})} \tag{20}$$

$$\frac{\partial W}{\partial s} = -\frac{\bar{v}-\hat{v}_{k,s}-(k+1)s}{(\bar{v}-\underline{v})} \tag{21}$$

Equation (21) on the other hand shows that welfare is decreasing (i.e. that consumer surplus is decreasing faster than profits increase) in s as long as $s < (\bar{v} - p_M)\frac{1}{2(k+1)}$, showing that the negative effect on consumer surplus outweighs the gains in profits from driving some consumers to channel D instead of channel A. We summarize the welfare analysis in the following proposition.

Proposition 2 (Welfare effects) *As long as channel A is used in equilibrium, it holds generally that* $\Delta CS_{k,s} < 0$, $\frac{\partial CS_{k,s}}{\partial s} < 0$, $\Delta\pi_{k,s} > 0$, $\frac{\partial \pi_{k,s}}{\partial s} > 0$, $\Delta W_k > 0$, $\frac{\partial W}{\partial s} < 0$.

While increases in the anonymization costs have the expected negative effects on consumer surplus and welfare, it has been less expectable that a higher level of strategic sophistication k would beneficial for welfare, but detrimental to consumer surplus. As a concluding perspective on the welfare analysis, note that in both panels of Figure 1 three groups of consumers have been distinguished: $\mathcal{C}_D$, $\mathcal{C}_A^-$, and $\mathcal{C}_A^+$. While the first denotes those consumers who chose channel D, the superscript at the two remaining groups distinguishes those consumers in channel A who make a net loss from those ending up with a net benefit of the whole transaction.[23] The next section discusses these results further.

[23]Due to linearities in our model the group size of $\mathcal{C}_A^-$ is equal to s as long as $k < \overline{k} - 1$.

3 Discussion

We have shown that under certain conditions, most notably under the assumption of imperfect strategic sophistication of consumers, a costly privacy protective sales channel can exist even when consumers do not have an explicit "taste for privacy". In our model, consumers want to restore their privacy (i.e. they choose channel A) solely based on their valuation of the good and their price expectation. We thereby provide a micro-foundation as to why people might have such a "taste for privacy" to which the existence of an anonymization channel can cater and even potentially make a profit on (e.g. if a (royalty) fee is part of s).

We have demonstrated that such a market cannot exist if consumers have the typically assumed infinite level of strategic sophistication as they would then be able to foresee the seller's incentive to always price above their price expectation. Turning the focus on incremental changes of strategic sophistication rather than a binary comparison between unlimited strategic sophistication and complete strategic naïveté, we were further able to show that infinite strategic sophistication, while sufficient, is not necessary for the breakdown of the anonymization channel. A finite level of strategic sophistication, $\overline{k}$, is already sufficent for this.

Also, we have shown that increasing sophistication leads to a reduction in consumer surplus, but an increase in profits and total welfare. Thus, while the potential of sellers with access to big data and the relevant techniques improves welfare by avoiding the dead weight loss usually associated with a monopoly (consumers with $v_i < p_M$ can now get the product), policy interventions trying to increase awareness of costly alternatives, could backfire in monopoly markets for at least some consumers (those in $\mathcal{C}_A^-$). Shifting our focus towards the cost of anonymization s, we have shown that consumer surplus is strictly decreasing when this cost increases, but profits are increasing. So, policies aiming at decreasing anonymization costs to consumers, overall efficiency would decrease in favor of increasing consumer surplus. Whether that is a desirable goal, has to be determined by policy-makers or the societies for which the policy is suggested.

We conducted some preliminary robustness check by relaxing the assumption of s being homogeneous, and allowed for it to be uniformly distributed between some lower bound $\underline{s}$ and an upper bound $\bar{s}$. Although we have not found a closed form solution yet, our results seem to generalize to this case. Moerover, comparing the case of heterogeneous s with the respective heterogeneous case of $s = \frac{\bar{s}+\underline{s}}{2}$ being the mean of the distribution in the heterogeneous case, we found that the heterogeneous market will break down slightly faster with increases in k. Since we further allowed for distributions where $\underline{s} = 0$, we compared it to the formerly ruled out homogeneous case of $s = 0$ to assess differences in interpretations.

The heterogeneous case could reflect that some consumers have simply not generated a large data footprint yet, whereas the homogeneous case of $s = 0$ could be understood as making privacy enhancing techniques readily available or assigning them particular rights about their data such that consumers actually do not incur any cost when using channel A – an interpretation somewhat comparable to the Google Spain case.[24] In the homogeneous case, our model would

[24]The interpretation depends heavily on which costs specifically are covered by s and whether this generally rules

result in another fully efficient outcome: All consumers with $v_i \geq p_M$ anonymize for free and no consumer in $\mathcal{C}_A$ would make a net loss after Stage 3, whereas all consumers with $v_i < p_M$ would still choose channel D where they will get the good for their exact willingness-to-pay. (...)[25]

Future empirical research based on our model could be undertaken to test our results with respect to equilibrium behavior and moreover test the usually stated range of k between 1 and 2 by an induced-value laboratory or a field experiment. Theoretical research, on the other hand, could further advance the focus on different levels of strategic sophistication by allowing for heterogeneity in k within the consumer population, relying on a more complex cognitive hierarchy model than this first attempt we undertook here. In general, though, it will be crucial to figure out the empirical relationship between s, k and v for each relevant market as this will heavily influence which recommendations should be given to policy-makers, consumer groups or businesses, depending on a particular normative goal of surplus or welfare maximization.

out the personalized pricing in channel D...

[25] This section is still under construction due to the variety of possible conceptions for $s = 0$

References

Acquisti, Alessandro, and Jens Grossklags. 2007. "What Can Behavioral Economics Teach Us About Privacy?" In *Digital Privacy: Theory, Technologies and Practices*, edited by Alessandro Acquisti, Stefanos Gritzalis, Costas Lambrinoudakis, and Sabrina De Capitani di Vimercati, 363–77. New York & London: Auerbach Publications. Available at: `https://books.google.nl/books?id=Ye11dQQ5CbcC&lpg=PA363&ots=CWvOL7LOHj&dq=What%20Can%20Behavioral%20Economics%20Teach%20Us%20About%20Privacy%3F&pg=PA363#v=onepage&q&f=false`.

Acquisti, Alessandro, Leslie John, and George Loewenstein. 2009. "What Is Privacy Worth." In *Twenty First Workshop on Information Systems and Economics (WISE).* Arizona Biltmore Resort & Spa, Phoenix, AZ. Available at: `http://www.heinz.cmu.edu/~acquisti/papers/acquisti-privacy-worth.pdf`.

Acquisti, Alessandro, and Hal R. Varian. 2005. "Conditioning Prices on Purchase History." *Marketing Science* 24 (3): 367–81. doi:10.1287/mksc.1040.0103.

Argenton, Cédric, and Jens Prüfer. 2012. "Search Engine Competition with Network Externalities." *Journal of Competition Law and Economics* 8 (1): 73–105.

Binswanger, Johannes, and Jens Prüfer. 2012. "Democracy, Populism, and (un)bounded Rationality." *European Journal of Political Economy* 28 (3): 358–72. doi:10.1016/j.ejpoleco.2012.02.002.

Brandimarte, Laura, Alessandro Acquisti, and George Loewenstein. 2012. "Misplaced Confidences: Privacy and the Control Paradox." *Social Psychological and Personality Science* 4 (3): 340–47. doi:10.1177/1948550612455931.

Calzolari, Giacomo, and Alessandro Pavan. 2006. "On the Optimality of Privacy in Sequential Contracting." *Journal of Economic Theory* 130 (1): 168–204. doi:10.1016/j.jet.2005.04.007.

Camerer, Colin F., Ho, Teck-Hua Ho, and Juin-Kuan Chong. 2004. "A Cognitive Hierarchy Model of Games." *The Quarterly Journal of Economics* 119 (3): 861–98.

Casadesus-Masanell, Ramon, and Andres Hervas-Drane. 2014. *Competing with Privacy*. Management Science 61 (1): 229-246.

Conitzer, Vincent, Curtis R. Taylor, and Liad Wagman. 2012. "Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases." *Marketing Science* 31 (2): 277–92. doi:10.1287/mksc.1110.0691.

Costa-Gomes, Miguel A., and Vincent P. Crawford. 2006. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study." *American Economic Review* 96 (5): 1737–68. doi:10.1257/aer.96.5.1737.

Costa-Gomes, Miguel, Vincent P. Crawford, and Bruno Broseta. 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study." *Econometrica* 69 (5): 1193–1235. doi:10.1111/1468-0262.00239.

Court of Justice of the European Union. 2014. Google Spain SL and Google Inc. v Agencia Española de Protección de Datos (AEPD) and Mario Costeja González. Court of Justice of the European Union.

Crawford, Vincent P. 2003. "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions." *The American Economic Review* 93 (1): 133–49.

Crawford, Vincent P., and Nagore Iriberri. 2007a. "Fatal Attraction: Salience, Naïveté, and Sophistication in Experimental 'Hide-and-Seek' Games." *American Economic Review* 97 (5): 1731–50. doi:10.1257/aer.97.5.1731.

Crawford, Vincent P., and Nagore Iriberri. 2007b. "Level-K Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica* 75 (6): 1721–70. doi:10.1111/j.1468-0262.2007.00810.x.

De Corniere, Alexandre, and Romain De Nijs. 2014. "Online Advertising and Privacy." *RAND Journal of Economics* (accepted for publication). Available at: `http://papers.ssrn.com/abstract=2191124.`

Einav, Liran, and Jonathan D. Levin. 2013. *The Data Revolution and Economic Analysis*. Working Paper 19035. National Bureau of Economic Research. http://www.nber.org/papers/w19035.

European Commission. 2008. *Opinion 1/2008 on Data Protection Issues Related to Search Engines*. 00737/EN WP 148. Brussels: European Commission. Available at: `http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2008/wp148_en.pdf.`

Fudenberg, Drew, J. Miguel Villas-Boas, and T. Hendershott. 2006. "Behavior-Based Price Discrimination and Customer Recognition." *Handbook on Economics and Information Systems* 1: 377–436.

Goldfarb, Avi, and Catherine Tucker. 2012. "Shifts in Privacy Concerns." *American Economic Review* 102 (3): 349–53. doi:10.1257/aer.102.3.349.

Goldfarb, Avi, and Botao Yang. 2009. "Are All Managers Created Equal?" *Journal of Marketing Research* 46 (5): 612–22. doi:10.1509/jmkr.46.5.612.

Hermalin, Benjamin E., and Michael L. Katz. 2006. "Privacy, Property Rights and Efficiency: The Economics of Privacy as Secrecy." *Quantitative Marketing and Economics* 4 (3): 209–39.

Ho, Teck-Hua, Colin Camerer, and Keith Weigelt. 1998. "Iterated Dominance and Iterated Best Response in Experimental 'P-Beauty Contests.'" *The American Economic Review* 88 (4): 947–69.

Jentzsch, Nicola, Sören Preibusch, and Andreas Harasser. 2012. *Study on Monetising Privacy - An Economic Model for Pricing Personal Information*. European Network and Information

Security Agency (ENISA). http://www.enisa.europa.eu/activities/identity-and-trust/library/deliverables/monetising-privacy/at_download/fullReport.

John, Leslie, Alessandro Acquisti, and George Loewenstein. 2009. *The Best of Strangers: Context Dependent Willingness to Divulge Personal Information*. SSRN Scholarly Paper ID 1430482. Rochester, NY: Social Science Research Network.Available at: http://papers.ssrn.com/abstract=1430482.

Luo, Xueming, Michelle Andrews, Zheng Fang, and Chee Wei Phang. 2014. "Mobile Targeting." *Management Science* 60 (7): 1738–56. doi:10.1287/mnsc.2013.1836.

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

Mikians, Jakub, László Gyarmati, Vijay Erramilli, and Nikolaos Laoutaris. 2012. "Detecting Price and Search Discrimination on the Internet." In *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*, 79–84. Redmond, WA: ACM. doi:10.1145/2390231.2390245.

Nagel, Rosemarie. 1995. "Unraveling in Guessing Games: An Experimental Study." *The American Economic Review* 85 (5): 1313–26.

Odlyzko, Andrew. 2003. "Privacy, Economics, and Price Discrimination on the Internet." In P*roceedings of the 5th International Conference on Electronic Commerce - ICEC 2003*, edited by Norman M. Sadeh, Mary Jo Dively, Robert J. Kauffman, Yannis Labrou, Onn Shehory, Rahul Telang, and Lorrie Faith Cranor, 355–66. ICEC 2003. Pittsburgh, Pennsylvania, USA: ACM Press. doi:10.1145/948005.948051.

Pigou, Arthur C. 1920. *The Economics of Welfare*. London: Macmillan and Co., limited. Available at: http://www.econlib.org/library/NPDBooks/Pigou/pgEW.html.

Posner, Richard A. 1978. "An Economic Theory of Privacy." *Regulation* 2: 19.

Posner, Richard A. 1981. "The Economics of Privacy." *The American Economic Review* 71 (2): 405–9. doi:10.2307/1815754.

Shiller, Benjamin Reed. 2013. *First Degree Price Discrimination Using Big Data*. Working Paper 58. Working Paper Series. Brandeis University, Department of Economics and International Businesss School. http://ideas.repec.org/p/brd/wpaper/58.html.

Shopkick. 2014. "Shopkick Surpasses $1 Billion Generated for Partners." *PR Newswire*. October 15. http://www.prnewswire.com/news-releases/shopkick-surpasses-1-billion-generated-for-partners-279269882.html.

Spiegel, Joel R., Micheal T. McKenna, Girish S. Lakshman, and Paul G. Nordstrom. 2013. "Method and System for Anticipatory Package Shipping." http://pdfpiw.uspto.gov/.piw?docid=08615473.

Stahl, Dale O., and Paul W. Wilson. 1994. "Experimental Evidence on Players' Models of Other Players." *Journal of Economic Behavior & Organization* 25 (3): 309–27. doi:10.1016/0167-2681(94)90103-1.

Stahl, Dale O., and Paul W. Wilson. 1995. "On Players′ Models of Other Players: Theory and Experimental Evidence." *Games and Economic Behavior* 10 (1): 218–54. doi:10.1006/game.1995.1031.

Taylor, Curtis R. 2004. "Consumer Privacy and the Market for Customer Information." *The RAND Journal of Economics* 35 (4): 631–50. doi:10.2307/1593765.

The New York Times. 2012. "Acxiom, the Quiet Giant of Consumer Database Marketing." *The New York Times*, June 16, sec. Technology. http://www.nytimes.com/2012/06/17/technology/acxiom-the-quiet-giant-of-consumer-database-marketing.html.

The Nielsen Company. 2012. *Which Smartphone Apps Do Savvy Shoppers Use Most*. Newswire - Mobile. The Nielsen Company. http://www.nielsen.com/us/en/insights/news/2012/top-mobile-shopping-apps.html.

Tirole, Jean. 1988. *The Theory of Industrial Organization*. 1st ed. MIT Press.

Tsai, Janice Y., Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. 2011. "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study." *Information Systems Research* 22 (2): 254–68. doi:10.1287/isre.1090.0260.

Tucker, Catherine. 2014. "Social Networks, Personalized Advertising, and Privacy Controls". *Journal of Marketing Research* 51 (5): 546-562. doi: 10.1509/jmr.10.0355

Varian, Hal R. 1997. *Economic Aspects of Personal Privacy*. Privacy and Self-Regulation in the Information Age. U.S. Department of Commerce - National Telecommunications and Information Administration. Available at: http://www.ntia.doc.gov/page/chapter-1-theory-markets-and-privacy.

Wall Street Journal Blog. 2014. "Amazon Wants to Ship Your Package Before You Buy It." *Digits - Tech News & Analysis From the WSJ*. January 17. Available at: http://blogs.wsj.com/digits/2014/01/17/amazon-wants-to-ship-your-package-before-you-buy-it/tab/print/.