Big data in the lab – How do consumers fare when predictive algorithms work against them?

SEBASTIAN DENGLER*

Draft version: May 8, 2016

VERY PRELIMINARY – Please do not distribute or quote!

Abstract

This paper investigates how consumers fare when predictive algorithms are not (only) working for, but (also) against them: when they predict not only which product a consumer is interested in, but also how much he or she is willing to pay for it. Do consumers anticipate that this can happen and do they take into account that such an algorithm might even outsmart them? Will people make use of costly privacy protective alternatives in such markets and can their behavior be explained by the limited strategic sophistication model of level-k thinking?

Introduction

Intelligent personal assistants (such as Google Now (Google), Siri (Apple), Cortana (Microsoft), and Echo (Amazon)) are becoming vital parts of many people's lives and some can already provide information before they have been asked for it. These and similar services are mostly driven by the recent technological development summarized as the rise of big data (Mayer-Schönberger and Cukier 2013), which encompasses the sheer volume and variety of data available as well as the velocity with which it can be analyzed and put to valuable use. As storing data has become relatively inexpensive and as more data is generated by the constant increase of web-based or -aided transactions, sellers can make more tailored offers. Such tailoring is not limited to consumer's preferences alone, but can be based on location data in real time and manifest as so-called "mobile targeting" (Luo et al. 2014).

This development supports Odlyzko's prediction that "in the Internet environment, the incentives towards price discrimination and the ability to price discriminate will be growing." (Odlyzko 2003, p. 365). The more data is generated by consumers and provided to sellers, the closer they will get to first-degree (or perfect) price discrimination, where the seller has complete information about every specific consumer's willingness to pay for any given product (and at any given time and/or location) (Pigou 1920).

An exemplary case of this trend is the shopping app Shopkick that rewards its users not only for visiting particular stores, but even guides them to scan specific products, take them to the dressing room or similar actions that make purchases more likely. According to a 2012 study by Nielsen, Shopkick (founded in 2009) already had a reach of 6 million users spending more than three hours per month using the app rendering it the most-used shopping app (The Nielsen Company 2012). Shopkick itself reported last year that it had created a total of more than USD 1 billion in revenues for partner stores by generating more than 50 million walk-ins and over 100 million product scans since its foundation (Shopkick 2014).

* PhD Student, Department of Economics, Tilburg University, PO Box 90153 5000 LE Tilburg, Netherlands. e-mail: s.dengler@tilburguniversity.edu

Related Literature

It is often pointed out (e.g. Solove 2006) that it is not precisely clear what is meant by "privacy" and hence different strands of literature have evolved: The early theoretical literature about the economics of privacy has underlined the negative welfare effects of consumers hiding information from sellers (Posner 1978; Stigler 1980; Posner 1981; Varian 1997). Later, the focus shifted towards the choices of firms that own some type of personal information about consumers (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 are the welfare consequences of keeping that information private or disclosing it to another firm, 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 of the papers.

Models of so-called "behavior-based price discrimination" have become prominent in addressing consumer-firm interactions (for an overview see Fudenberg, et al. 2006). Most of this literature focuses on second-degree price discrimination by assuming that a seller 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. The possibility of perfect price discrimination, though, has not received a lot of attention in the economics literature and has mostly been dismissed as a mere theoretical construct.[1]

Empirical studies regarding consumer's choices have shown that consumers' stated and revealed valuations of their own personal data differ highly and depend on the framing of the survey questions (Acquisti, et al. 2009; John, et al. 2011; Jentzsch, et al. 2012), which has become known as the "privacy paradox".

While shifts in privacy concerns can be observed empirically (Goldfarb and Tucker 2012) and although the performance of predictive algorithms for personal assistance is visible to users, they might be unaware of the potential for personalized pricing that comes along with this type predictive power and hence act non-optimally when making privacy-related choices. Regarding such choices Acquisti and Grossklags (2007, p. 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."

This paper investigates how consumers fare when predictive algorithms are not (only) working for, but (also) against them: when they predict not only which product a consumer is interested in, but also how much he or she is willing to pay for it. Do consumers anticipate that this can happen and do they take into account that such an algorithm might even outsmart them? What are the implications of privacy protective alternatives in such markets and will people make use of them if that means at the same time losing out on convenience or even having to pay money for an otherwise "free" product?

I experimentally test the theoretical results developed in Dengler and Prüfer (2016) in a laboratory experiment. The experimental setting closely followed the setup of the model of a monopolistic firm facing consumers with different valuations for a certain good. In an induced-value experiment consumers can choose to approach a monopolistic seller via two channels: a costly but privacy-preserving channel that hides their valuation of the good or the costless direct channel in which perfect price discrimination is possible. For all consumers in the privacy-preserving channel the seller then sets

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

a uniform price, absent other information about consumers, whereas he can charge perfectly personalized prices to all other consumers. To best reflect both, reality with ever more precise algorithms and the theoretical model, the seller is a computer-player employing an optimizing algorithm rather than a human subject.

This setup allows to test the underpinnings of the model of limited strategic thinking – so-called level-k thinking – among consumers employed in Dengler and Prüfer (2016). The idea of level-k thinking was introduced by Stahl and Wilson (1994; 1995) and Nagel (1995) and has been explored in a sizeable literature theoretically and empirically (Ho et al. (1998), Costa-Gomes et al. (2001), Crawford (2003), Camerer et al. (2004), Costa-Gomes and Crawford (2006), Crawford and Iriberri (2007b), Goldfarb and Yang (2009), and Binswanger and Prüfer (2012), among others). The interpretation in the theoretical model is that the consumers are only able to perform a limited number of iterations in their strategic assessment of the market situation (indicated by their k).[2] 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 on average (Camerer, et al. 2004; Crawford and Iriberri 2007a), and level-k models can be seen as a generalization to the usual assumption of unlimited strategic sophistication, which would be reflected by setting k equal to infinity. As the experiment is designed to test both, assumptions and results of our model, I will first provide an abridged version of the model before delving into specifics of the proposed experiment.

Theoretical Model and Sophistication-k Equilibrium Behavior

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 $c \geq 0$ and consumers have a heterogeneous valuation v_i for the good, where $v_i \sim \mathcal{U}[0,1]$.

We assume that consumers can buy the good from the seller in two different ways: using the usual direct channel (channel D) or making use of an anonymization technique (referred to as the anonymous channel A). The seller can neither directly influence consumers' channel choice nor close down one channel.[3] If a consumer chooses direct channel D, the seller knows her valuation (i.e. her maximum willingness to pay) perfectly through his large customer information database. However, if a consumer chooses to make use of the anonymization technique, her valuation is hidden from the seller, but using the anonymization technique comes at cost $s > 0$ (reflecting disutility stemming from a monetary payment for using the technique, a decrease in product search quality, lower internet connection speed or even increased search or transaction costs if the anonymization technique is "shopping offline").

Further, we assume that all agents are solely interested in their own material payoff (i.e. net monetary profit or consumer surplus). 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 the supply function) as well as the timing (see next paragraph) of the game are common knowledge among all agents. To account for the cognitive constraints that consumers face we assume that all consumers have a certain level of strategic sophistication, denoted by k. As a reflection of the seller's higher analytical capacities (due to his computing power, information databases and forecasting

[2] Loosely speaking, k can be understood as the number of steps a player can "think through / ahead".

[3] The seller might simply not be allowed to not serve consumers who make use of anonymization tools, if consumers have access to such tools. Also, the monopolist could face competitive pressure by potential entrants that credibly promise not to make use of consumer data and hence offer to serve anonymous consumers rather than to lose market share.

algorithms) we assume that the seller's level of strategic sophistication is always higher than the level of the consumers. Suppose the level of strategic sophistication of the consumers is given by k, then the seller will have a level of at least $k+1$. 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 her valuation v_i and her price expectations for the two channels, each consumer decides whether to use channel D or channel A, and incurs 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 consumer in channel D, 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.

Due to the limited strategic sophistication of consumers we cannot solve the game for sub-game perfect Bayesian Nash equilibria (which is nested in our analysis, though, if $k = \infty$) but for an equilibrium inspired by the level-k literature. While most of this literature deals with (fairly) symmetric decisions (e.g. the Beauty Contest game) we have to adapt the concept to the less symmetric situation or our model. We therefore employ the closely related concept of sophistication-k equilibrium as defined in Binswanger and Prüfer (2012) and tailor it to our needs:

Definition 1 (Modified Sophistication-k Equilibrium)
A sophistication-k equilibrium is a strategy combination and a set of beliefs 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.*

The belief structure we assume for our model is therefore as follows:

Assumption 1 (Belief structure)

- *All agents take into account that* $v_i \in [0,1]$ *and hence that* $p_i, p_A \in [0,1]$ *(eliminating beliefs higher than the prohibitive valuation).*
- *Consumers form a belief about the prices the seller will set in channel D and channel A, denoted by* $\mathbb{E}(p|D)$ *and* $\mathbb{E}(p|A)$ *respectively.*
- *As consumers' level of sophistication is lower than the seller's it does not need to be the case that* $\mathbb{E}(p|D) = (p|D)$ *and* $\mathbb{E}(p|A) = (p|A)$ *(as it would have to be in a Perfect Bayesian Nash equilibrium).*
- *The seller forms a belief about which consumers have chosen channel D and channel A and correctly anticipates the beliefs of the consumers.*

As part 1 of this equilibrium concept requires that strategies be Nash equilibria given the beliefs and strategies of the other player, the game is solved by using backward induction. Leaving the detailed analysis to the other paper (Dengler and Prüfer (2016)), I will provide a very reduced exposition around the most important intermediate results building towards that paper’s central proposition. Lemma 1 follows relatively straightforward from the usage of backward induction to solve the game and is hence stated without explanation.

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

The optimal response of the seller in channel D given his ability to perfectly discriminate is trivial, but in channel A the seller's optimal strategy crucially depends on the location of the anonymization threshold $\hat{v}$ relative to the regular optimal monopoly price p_M if price discrimination was impossible.

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

The optimal pricing strategy implies that the seller sets a higher price than consumers had expected whenever $p_A^* = \hat{v} = \mathbb{E}(p_A) + s > \mathbb{E}(p_A)$. This is due to the fact that s is going to be a sunk cost for consumers at Stage 3, which the seller can anticipate and exploit via increasing the price by exactly s compared to their expectation. Consumers, on the other hand, do only anticipate this best response in parts, depending on the degree of their limited strategic sophistication that influences their expectation formation.

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, which is why we will start with the case of $k = 0$. Consumers with $k = 0$ are called naïve consumers and naïvely expect the monopolist to engage in regular monopoly pricing in channel A, ignoring the fact that the very choice of channel A might be a signal to the seller.[4] For channel D, we assume that even the most naïve (but still rational) consumer foresees the perfect price discrimination as this does not require iterative thinking. The equilibrium behavior if consumers are completely naïve is as follows:

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

- *Consumers form the 0-beliefs* $\mathbb{E}_0(p_D) = p_i^*$ *and* $\mathbb{E}_0(p_D) = p_M = \frac{1+c}{2}$.
- *Consumers anonymize if and only if* $v_i > \hat{v}_0 = p_M + s$, *separating into the sets* $\mathcal{C}_D = [0, \hat{v}_0]$ *and* $\mathcal{C}_A = (\hat{v}_0, 1]$.
- *The seller correctly infers sets* $\mathcal{C}_D = [0, \hat{v}_0]$ *and* $\mathcal{C}_A = (\hat{v}_0, 1]$.
- *The seller sets the optimal prices* $p_i^* = max\{v_i, c\}$ *and* $p_{A_0}^* = \hat{v}_0 = p_M + s$.
- *All consumers in* $\mathcal{C}_D$ *with* $v_i \geq c$ *buy the product at the price offered to them.*
- *All consumers in* $\mathcal{C}_A$ *buy the product at the price offered.*

We see that there is a difference of exactly s between consumers' expectations of the price in channel A and the optimal price the seller eventually sets for this channel. Due to their limited capabilities in strategic reasoning the consumers do not foresee that once they reach Stage 3 the anonymization expenses of s will have turned into sunk costs and that the seller can exploit this fact. This in turn informs us about the way in which consumers form their price expectation for higher level of strategic sophistication $k > 0$.

---

[4] Starting with the more typical assumption of a level-0 player uniformly randomizing would not change the major outcomes, but make the point less clear. Since either assumption is supportable, we choose in favor of expositional simplicity.

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

$$\mathbb{E}_k\,(p_A) = p_M + k\,s$$
$$p^*_{A_k} = p_M + (k+1)s \;=\; \hat{v}_{k,s}$$

which implies the equivalence $\mathbb{E}_k\,(p_A) = p^*_{A_{k-1}}$. Thus, at every additional level of strategic sophistication the consumers will incorporate the sunk costs once more than at the previous level which induces the seller to raise his price once more. The more the population of consumers is strategically sophisticated, the fewer consumers will choose channel A. This iterative process comes to an end when $\mathbb{E}_{\overline{k}}(p_A) \;>\; 1 - s$ as this $\overline{k}$-belief causes the anonymization threshold $\hat{v}_{\overline{k},s} = p_M + (\overline{k}+1)s$ to exceed the maximum valuation $v_i^{max} = 1$ leading to $\mathcal{C}_A = \emptyset$. Then, channel A remains unused and the market for anonymization completely breaks down. We find that this level is reached at the finite number $\overline{k} = \left\lfloor \frac{1-c}{2s} \right\rfloor$ and hence we have also implicitly solved for the limit case of $k = \infty$ representing an approach with unlimited strategic sophistication of consumers. Thus, while unlimited strategic sophistication is a sufficient condition for a breakdown of channel A, it is not a necessary condition. This leads to the central proposition of the model and some additional insights.

Proposition 1 (Sophistication-k equilibrium)

For any non-prohibitively high cost of anonymization s and for any non-prohibitively high cost of production c it holds that:

1. *If consumers have a level of strategic sophistication of $\overline{k} < \left\lfloor \frac{1-c}{2s} \right\rfloor$, there is a unique sophistication-k equilibrium with the following characteristics:*
   - *Consumers form the k-beliefs $\mathbb{E}_k\,(p_D) = p_i^*$ and $\mathbb{E}_k\,(p_A) = p_M + k\,s = \frac{1+c}{2} + ks$.*
   - *Consumers anonymize if and only if $v_i > \hat{v}_k = p_M(k+1)s$, separating into the sets $\mathcal{C}_D = [0, \hat{v}_k]$ and $\mathcal{C}_A = (\hat{v}_k, 1]$.*
   - *The seller correctly infers $\mathcal{C}_D = [0, \hat{v}_k]$ and $\mathcal{C}_A = (\hat{v}_k, 1]$.*
   - *The seller sets the optimal prices $p_i^* = max\{v_i, c\}$ and $p^*_{A_k} = \hat{v}_k = p_M + (k+1)s$.*
   - *All consumers in $\mathcal{C}_D$ with $v_i \geq c$ buy the product at the price offered to them.*
   - *All consumers in $\mathcal{C}_A$ buy the product at the price offered.*

2. *If consumers have a level of strategic sophistication of $\overline{k} \geq \left\lfloor \frac{1-c}{2s} \right\rfloor$, there are multiple sophistication-k equilibria with the following characteristics:*
   - *Consumers form the k-beliefs $\mathbb{E}_k\,(p_D) = p_i^*$ and $\mathbb{E}_k\,(p_A) \in [1-s, 1]$, i.e. they expect the price in channel A to be prohibitively high.*
   - *No consumer anonymizes as $v_i \leq \hat{v}_k = p_M(k+1)s$ for all i, leading to the sets $\mathcal{C}_D = [0,1]$ and $\mathcal{C}_A = \emptyset$.*
   - *The seller correctly infers $\mathcal{C}_D = [0,1]$ and $\mathcal{C}_A = \emptyset$.*
   - *The seller sets the optimal prices $p_i^* = max\{v_i, c\}$ and any subgame perfection preserving off-equilibrium price $p_A^* \in [1-s, 1]$.*
   - *All consumers in $\mathcal{C}_D$ with $v_i \geq c$ buy the product at the price offered to them.*
   - *No consumer buys the product via channel A.*

Experimental Setting

So far, six experimental sessions have been conducted in Tilburg University's own laboratory for economic experiments (CentERlab) with student subjects. The data from the fifth session, however, had to be disregarded due to an error in the computer program, such that for now there are five properly conducted sessions, each with exactly 20 subjects, i.e. a total of 100 subjects.

Subjects were randomly allocated to one of the workstations in the laboratory and played a series of ten games. In sessions 1-4 subjects first played seven periods of the central market game, followed by three level-k elicitation games, whereas in session 5 & 6 (5 was the dropped session) subjects first played the level-k elicitation games to be then followed by the seven periods of the central market game. To encourage subject attendance and to cover for the potential loss from choosing the anonymous channel in the market game, subjects were paid a show-up fee of 6 €.

To avoid income effects, one period was selected at random using the following procedure: Before the experiment was started, the experimenter asked one subject to pick an envelope from a pile of sealed envelopes, containing cards with the numbers 1-10 on them. The subject subsequently signed the envelope which was then put in a corner underneath a box to be retrieved at the end of the experiment. At the end of the experiment the same subject verified the signature, opened the envelope, the selected period was announced to all subjects and the subject verified the correct entry by the experimenter into the computer program.

Each session lasted about one hour in total which includes time for individually recording payment information of subjects, which was then completed by bank transfer. The sessions resulted in average earnings of 8.12 €. The following paragraphs will explain the four different games before preliminary results are presented.[5]

Core Experiment: Market Game

For the central part of the experiment (the market treatment) subjects were told that they are in the role of a consumer facing a single seller who is a computer player offering one unit of a good to every consumer. Subjects were informed that they could be assigned one of 20 valuations ranging from 4.00 € to 13.50 € in increments of 0.50 €, that assignment was random and with equal chance, and that each valuation would be assigned to exactly one subject in the room, leading to the induced demand schedule depicted below (where the "optimum" marks the profit maximizing price if everyone were to hide their valuation).

For each valuation subjects had the choice two hide their valuation from the seller should this particular valuation be assigned to them ("hide" and "not hide" correspond to channels A and D of the theoretical model, respectively). This came at the cost of 1.00 €.

While the seller's price optimization algorithm was intentionally unknown to subjects, they were informed that the seller takes into account all available information, i.e. the demand schedule as well as any information that was not hidden by the subjects. Subjects were also informed that the seller would have to charge the same price to all subjects who chose to hide their valuation and that otherwise prices will be individualized.

[5] The instructions that were used during the experiment can be found in the appendix.

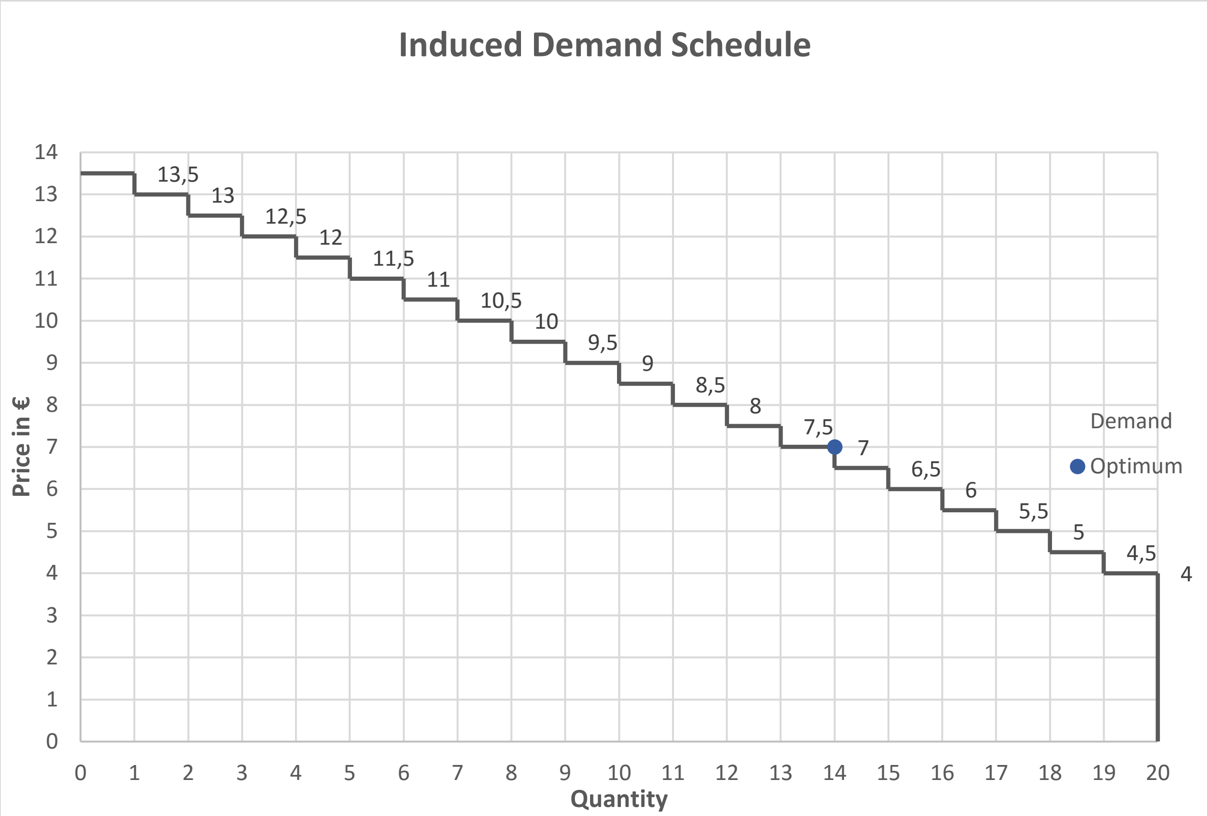

Assuming zero cost of production for the seller, the price for all subjects that chose to not hide their valuation was simply set to their respective valuation, leaving them with zero surplus from the transaction. For the subjects that chose to hide their valuation, though, the price was always the profit maximizing price given the realized demand schedule for all hidden valuations. If there was a tie between two different prices, the lower price (with higher demand and hence bigger consumer surplus) was picked – which would work against the model prediction of the market successively breaking down by lowering the uniform price rather than increasing it in tied cases. Subjects were unaware of this algorithm, though.

If the price the seller set for a subject was lower than her valuation, the subject received the difference between her valuation and the price as payoff. If the price was equal to or higher than her valuation, the payoff was zero. From this the cost of 1.00 € was subtracted if the subject had chosen to hide the assigned valuation. This procedure automated the third stage of the theoretical model to buy if and only if the price is lower than the respective valuation. Thus, while subjects could not make a loss from the comparison of the price to your valuation, their payoff could be negative (but limited to a loss of 1.00 €). When that happened the payoff would be deducted from the show-up fee if this particular period was selected for payment.

This game was repeated for seven periods, where subjects received feedback after every period consisting of their assigned valuation, the respective choice they made, the price the seller set for them and the price that the seller set for those subjects that chose to hide their valuation.[6] In each period, subjects were assigned a new valuation at random and submitted choices for all 20 possible valuations. No restrictions other than a complete schedule were imposed (especially it was not enforced that they would have to submit a schedule satisfying transitivity).

[6] This made sure that at least within sessions all subjects had the same information to reduce heterogeneity in information.

Given the model at hand the first hypothesis to be tested stems from the central assumption of the model, namely that subjects indeed exhibit limited strategic sophistication in this setting.

Hypothesis 1: *A significant part of the subject population chooses the costly channel A despite it not being part of an equilibrium supported by unlimited strategic sophistication.*

The case for the central assumption can be made even stronger, if this behavior persists over time.

Hypothesis 2: *A significant number of subjects does not converge to the equilibrium supported by unlimited strategic sophistication even after substantial learning opportunity.*

If this is the case, the result that a privacy-protecting channel can exist as long as consumers are not too sophisticated receives some empirical backing. However, the theoretical model also assumes that all consumers have the same level of strategic sophistication in order to be more tractable as this avoids the introduction of beliefs of some consumers about other less strategically sophisticated consumers.

To test whether this is overly simplistic, subjects played three additional games that giving representations of the different effects at play: the cognitive limitation of each subject and their estimate of limitations in others. The level-k literature has spawned different games that deal with these different aspects and the three chosen games were: a version of the Game of 21 (Dufwenberg, et al. 2010), a version of the Beauty Contest (Nagel 1995) and a version of the 11-20 Money Request Game (Arad and Rubinstein 2012).

Level-k Elicitations

Irrespective of the overall order of the experiment (Market Game *before* or *after* the level-k elicitations), the order within the level-k elicitation games was kept constant.

Level-k Elicitation Game 1: Beauty Contest

> "A large number of players have to state simultaneously a number in the closed interval $[0, 100]$. The winner is the person whose chosen number is closest to the mean of all chosen numbers multiplied by a parameter p, where p is a predetermined positive parameter of the game; p is common knowledge. The payoff to the winner is a fixed amount, which is independent of the stated number and p. If there is a tie, the prize is divided equally among the winners. The other players whose chosen numbers are further away receive nothing."(Nagel 1995, p. 1313)

The Beauty Contest game is very suitable to elicit what subjects think about how everyone else forms beliefs (about everyone else and their beliefs, etc.). If the parameter p is chosen such that $p < 1$ the only Nash equilibrium is for everyone to choose 0. However, this only holds true if everyone is perfectly strategically sophisticated. If there are players who do not achieve this level of strategic reasoning, the optimal response shifts upwards. As the response from the Beauty Contest takes the beliefs about capabilities (and beliefs) of all other players into account, it can be used to check whether this assessment of others' beliefs is a good predictor of the behavior in the market treatment and hence whether the simplifying assumption of a uniform level of strategic sophistication in the population is problematic. In order to avoid too much familiarity with the game if subjects have encountered it previously, either in other experiments or by having read about it somewhere, the game was slightly modified: The possible range was now from 1 to 200 rather than 0 to 100, with a maximum precision of 0.01, all with the parameter $p = 3/4$ rather than the more colloquial "guess 2/3 of the average" version.

Level-k Elicitation Game 2: Money Request Game

> "You and another player are playing a game in which each player requests an amount of money. The amount must be (an integer) between 11 and 20 shekels. Each player will receive the amount he requests. A player will receive an additional amount of 20 shekels if he asks for exactly one shekel less than the other player. What amount of money would you request?" (Arad and Rubinstein 2012, p. 3562)

This relatively new game is very unlikely to be known to subjects which makes for a good chance of getting a less noisy assessment of any subjects' initial reaction when encountering the game. Also, it triggers level-k thinking much more directly than other games.[7] As best-responding to any belief about the other player is very straightforward the results of this game can be "used to evaluate the plausibility of a distribution of types that best fits the data in [the] other games" as explicitly stated by (Arad and Rubinstein 2012, p. 3571). The only modification made was to rescale the game to the interval 1 € to 10 € and offering a bonus payment of 10 € for requesting exactly 1 € less than the other player.

Level-k Elicitation Game 3: Adding Game

> "Two players, call them *White* and *Green*, take turns. *White* begins. To start off, he can choose either 1 or 2. *Green* observes this choice, then increments the "count" by adding one or two. That is, if *White* chooses 1 *Green* can follow up with 2 or 3; if *White* chooses 2 *Green* can follow up with 3 or 4. *White* then observes *Green's* choice, and again increments the count by adding one or two. The game continues with the players taking turns, each player incrementing the count by one or two. The player who reaches 21 wins." (Dufwenberg, et al. 2010, p. 132–133)

As this game has a unique solution by backward induction but requires several iterations to reach this conclusion it is very well suited to elicit the cognitive constraints that are included in the theoretical model. Moreover, the optimal strategy is fully independent of beliefs about the other player irrespective of his or her level of strategic sophistication (i.e. the best response to a 0-belief is identical to the best response to an ∞-belief). The confounding expectations about other's limitation can be ruled out and hence this game can give a good isolated estimate of the strategic capabilities. In order to not lose any observation, however, subjects played this game against a random playing computer player, and again the game was changed to another interval to make it less similar to the original version. The implemented version allowed to pick from 1, 2, 3, 4, 5, or 6 and the goal was set to reach 50 or more (where the last part was necessary as the computer should also play randomly in the last round). This increase in the choice set also means that the chance for either the subject or the computer randomly picking the best response drastically decreased compared to the original version. If the subject won against the computer, she received 10 € as payoff, otherwise her payoff from this game would be 0 €.

These additional games allow the formulation of additional hypotheses to test whether or not subjects' behavior follows the pattern described by the theoretical model.

Hypothesis 3: *The behavior of a majority of subjects can be explained by levels of strategic sophistication of 0, 1, or 2.*

This would be in line with the empirical results of previous literature in which it has been shown that behavior consistent with a higher strategic sophistication than $k = 2$ is very rarely observed. Of course, it is expected that higher levels of strategic sophistication are also associated with different behavior in the market treatment:

[7] Several other features – such as the lack of dominated strategies and inexistence of a pure-strategy Nash equilibrium – are discussed in the paper by Arad and Rubinstein (2012).

Hypothesis 4: *Subjects for which a higher level of strategic sophistication is elicited in the level-k elicitation games also choose to stay in the free direct channel for lower induced values for the good than subjects with lower strategic sophistication.*

Using the variety of the different level-k elicitation methods allows to test separately whether the difference in behavior is more likely to be attributable to subjects' own cognitive limitations or to their differing beliefs about the cognitive limitations of others.

Hypothesis 5: *The level-k elicitation of the Adding Game is a better predictor of behavior in the market treatment than both other level-k elicitations.*

Not rejecting this hypothesis would back the simplifying assumption of disregarding beliefs about other consumers' behavior. If, however, the two other elicitation methods lead to a better prediction of behavior in the market treatment, the theoretical model should likely be extended to include beliefs of consumers about other consumers.

Preliminary Results: Market Game

The focus of these preliminary results lies on summary statistics of the games played, starting with the core market game. Let a "sensible" choice schedule be defined as having at most one switching point from "hiding" to "not hiding", which is further located such that "hiding" occurs for valuations above the switching point and "not hiding" below the switching point. Observing that a majority of subjects (52 %) submitted a not sensible choice schedule in at least in one out of the seven periods of this game already shows that subjects had difficulties to fully grasp what was going on in the game. This already gives some first support to the assumption of limited strategic sophistication.

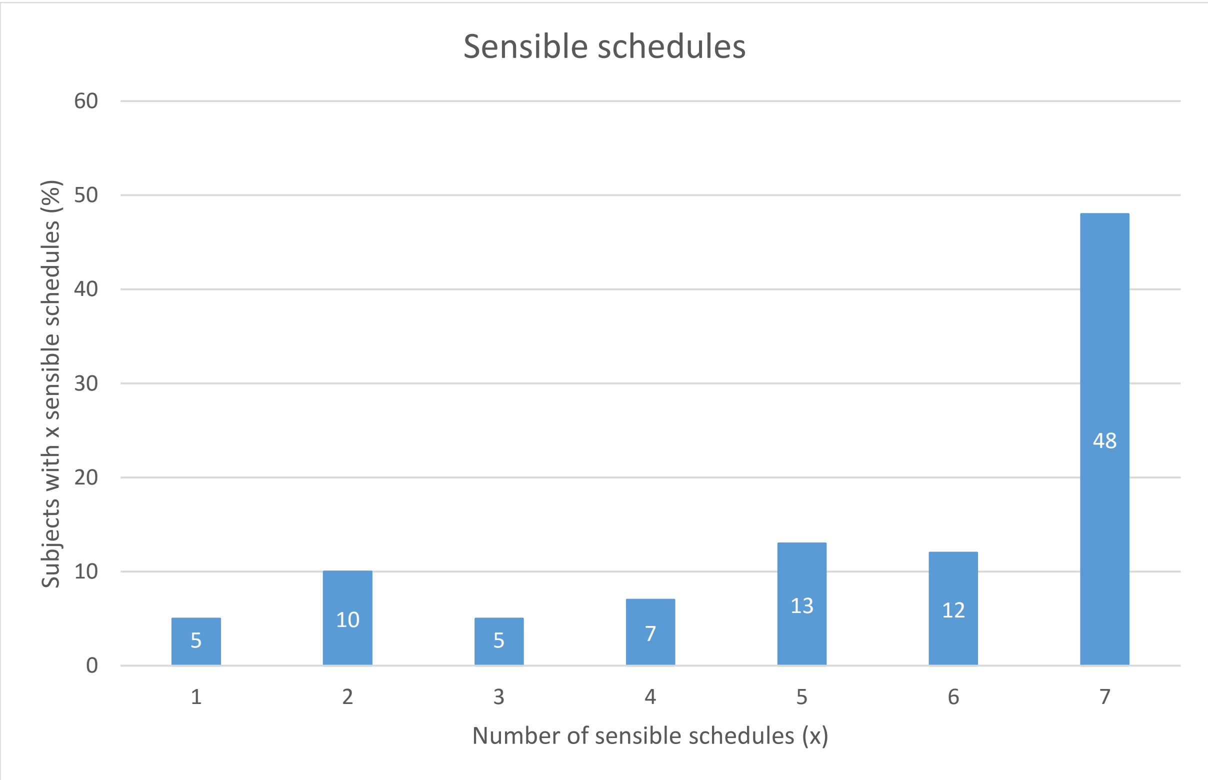

This difficulty notwithstanding, a closer look at the data allows testing of Hypothesis 1 and 2 in two different ways.

Hypothesis 1: *A significant part of the subject population chooses the costly channel A despite it not being part of an equilibrium supported by unlimited strategic sophistication.*

Hypothesis 2: *A significant number of subjects does not converge to the equilibrium supported by unlimited strategic sophistication even after substantial learning opportunity.*

The first, most literal interpretation is to look at which share of the subject population chose to hide their valuation for at least one of the twenty valuations in any given period of the game. This is depicted by the "HideCount > 0" graphic below. As can be seen clearly, a large majority of subjects does so and even more so after the first two periods, supporting both, hypothesis 1 and 2. Further, it does not make a difference whether only subjects who submit sensible schedules in all 7 periods or whether all subjects are taken into account.

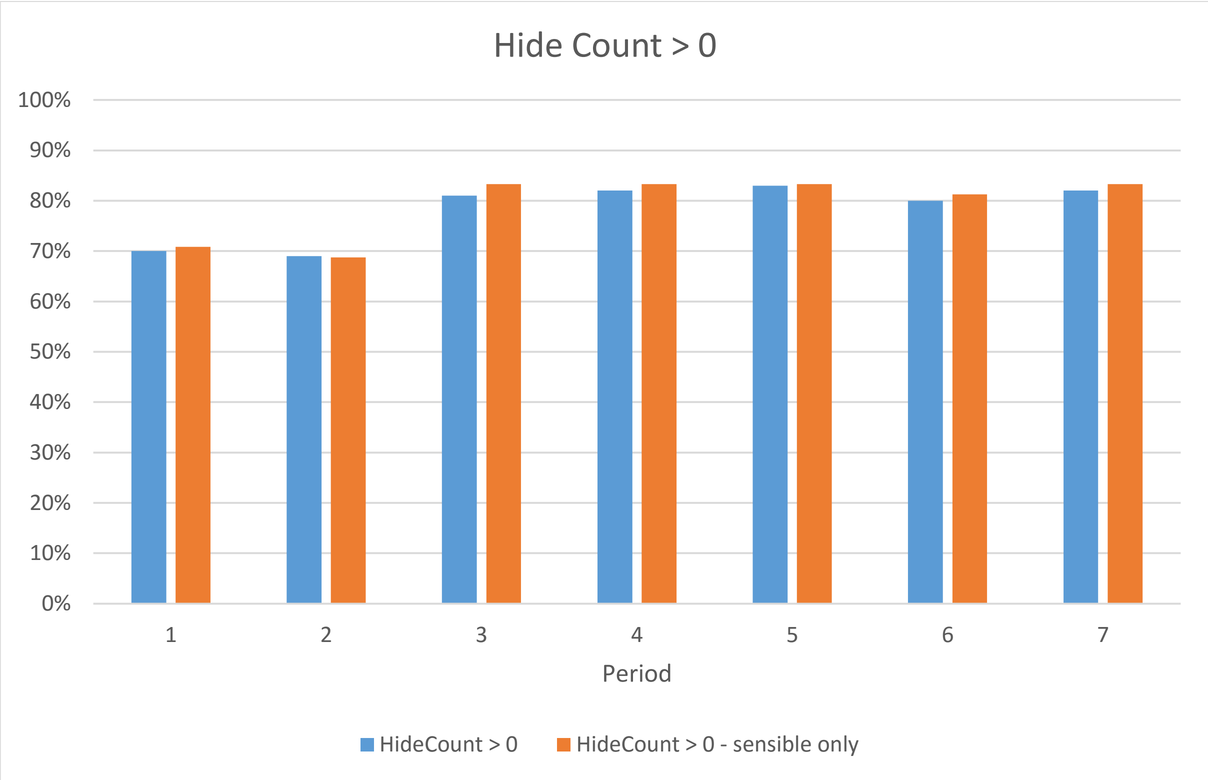

As the question whether a subject chooses "hide" for at least one valuation is a very crude measure, though, a more fine graned analysis can shed light on whether there is a more detailed pattern at work than the simple "black and white" picture so far.

A natural candidate is to look at the extent to which subjects chose to hide their valuation, shown in the graphic "Hide Choices" below. For every period, the graphic depicts the percentage of all choices that were made to hide rather than to not hide the respective valuation. Here too, it can be seen that subjects are chosing to hide more in later periods than in the first period, counter to the prediction that the market for anonymization would collapse. Again, there is no substantial difference in the overall pattern (though a slight difference in levels exists) between those that always submit a sensible choice schedule as opposed to the whole sample.

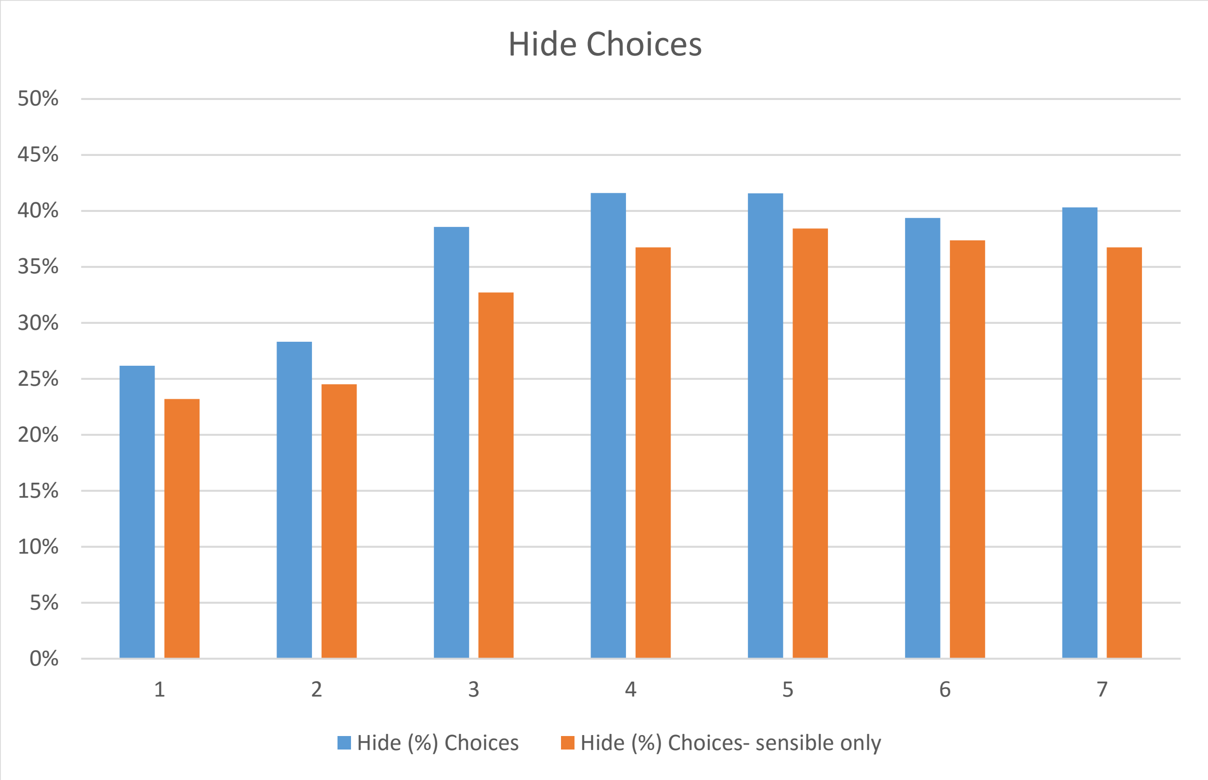

However, the increasing share of hiding choices need not be a complete rejection of learning behavior that might eventually lead to a breakdown of the anonymous channel. As all subjects learn the realized price, some of them might learn that it can be beneficial to hide higher valuations rather than to let them be known to the seller. The development of the realized monopoly prices across time for all sessions (but the dropped session 5) shows, that despite the positive linear trend, the realized prices are well below the maximum valuation of 13.50 € for a majority of the sessions.

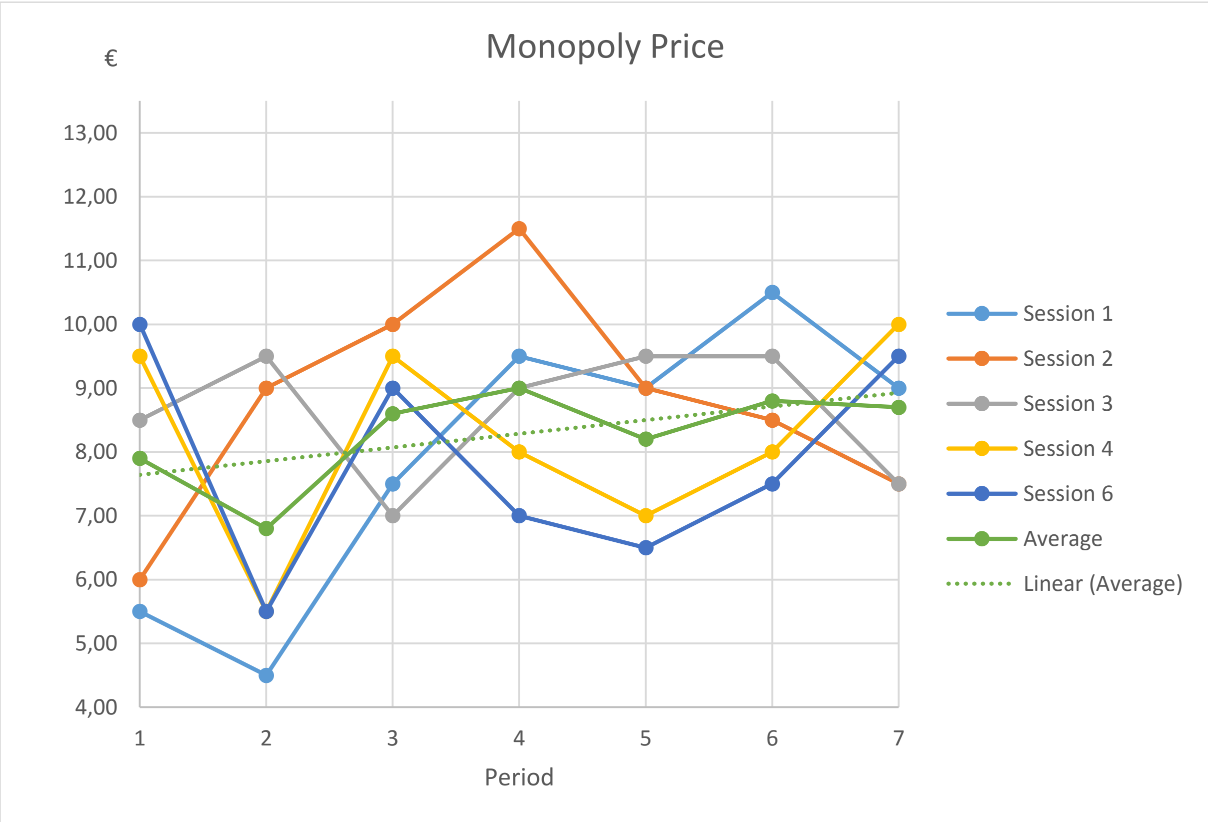

As a matter of fact, the monopoly price only exceeds 10.00 € twice in all 35 recorded periods, such that incurring the cost of 1.00 € to hide a high valuation is readily understandable to subjects. Thus, if subjects are indeed learning, the expected valuation conditional on it being hidden from the seller should increase over time. For the following graphic, the schedule of each subject was analyzed and the average valuation for which the subject chose to hide the valuation was calculated. Subsequently these average individual valuations were assigend equal weights to calculate the average hiding valuation in each period.

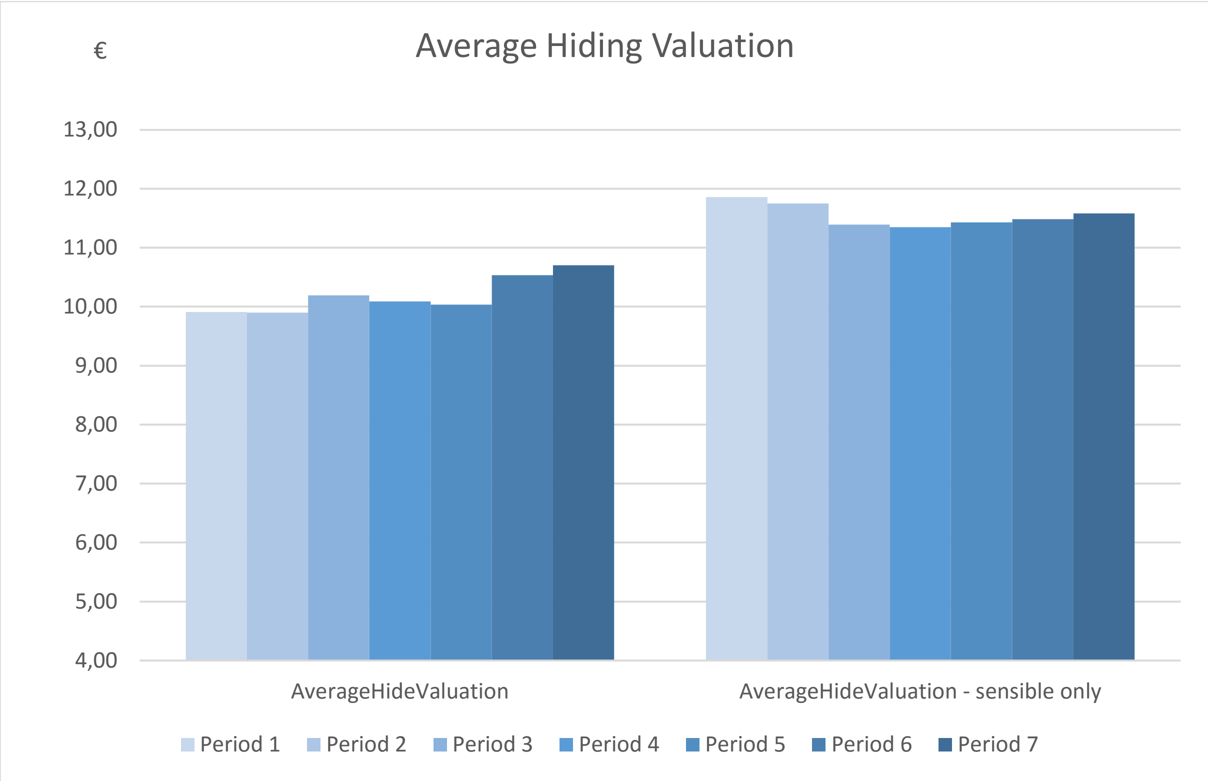

The graph for all valuations of all subjects does seem to show some increasing trend, whereas the average valuation for those who submitted a sensible schedule in all periods seems to exhibit a u-shape rather than an increasing trend. However, it has to be noted that this analysis neglects all sensible schedules that were submitted in which the subject chose to not hide any valuation whatsoever. Hence, it is not possible to include them in this calculation as arbitrarily assigning them any value might not represent their counterfactual choices. Since there are 153 of such schedules among the 700 submitted schedules in total (i.e. 21.86 %) and as 70 of these schedules make up a share of 20.83 % among the 336 schedules submitted by subjects who always submit a sensible schedule, further analysis accounting for this censoring will have to be done.

Apart from this, a look at the different distributions of the average hiding valuation across periods. In the graphic below all periods are shown, but the first (full black) and last (circled black) period are made more prominent to highlight the difference between them. Two major shifts can be noticed:

First, in Period 1 the distribution of average valuations is much more homogeneous than in Period 7, where there is a rapid increase in subject mass around the center of the valuation spectrum. This is most notably due to the fact that almost no subject has an average hiding valuation below 8.00 € whereas in Period 1 almost 20% of the subject mass are already covered.

Secondly, in Period 1 only 70% of all subjects have an average hiding valuation, which implies that 30% chose to not hide their valuation at any of the valuations available. In Period 7, however, only 18% of subjects chose not to hide their valuation at all.

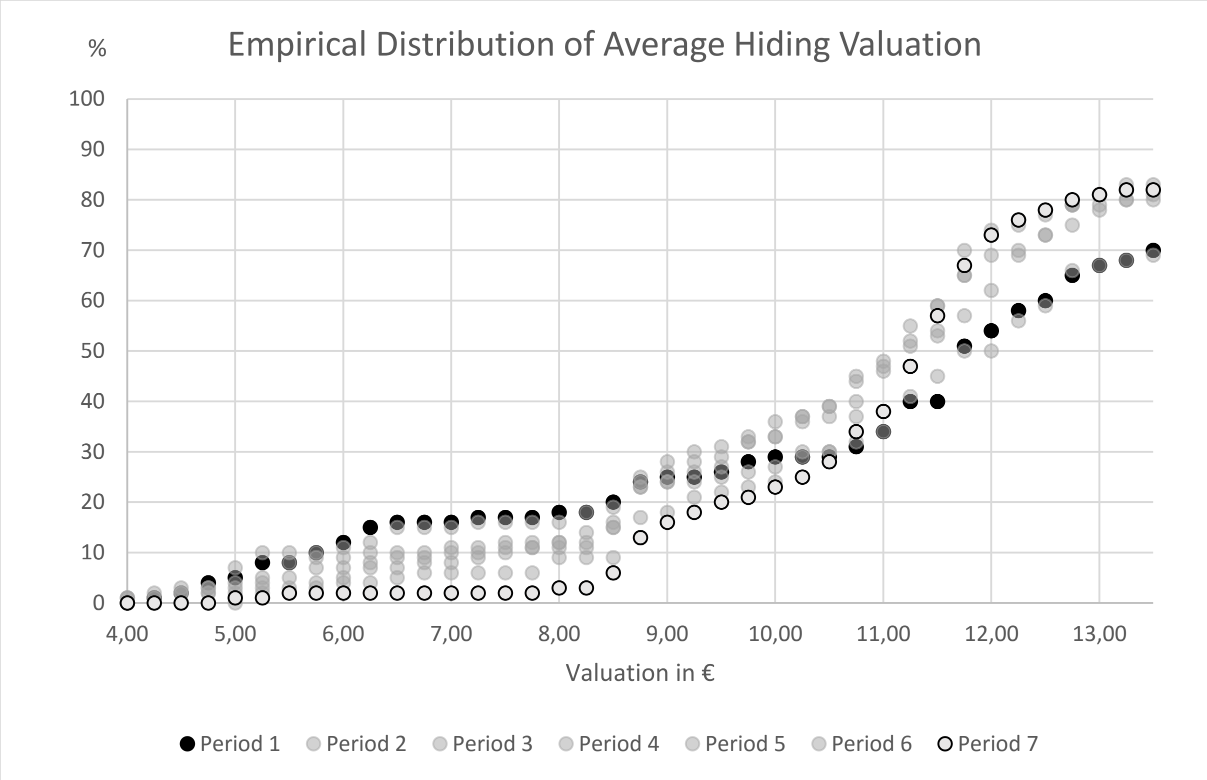

These two results seem to suggest that subjects on both ends of the spectrum indeed learn over time. Subjects who hide for very low valuations learn that it is indeed not optimal to do so (especially since a very low average hiding valuation implies that they do not hide for high valuations), whereas subjects who do not hide their valuation under any circumstance in the beginning realize that they could make profits by hiding for the highest valuations. Further robustness checks (such as controlling for the realized uniform price of the previous period) are necessary, though, to render this conclusion less speculative.

Very Preliminary Results: Level-k Elicitations

For two level-k elicitation games (Money Request & Adding Game) the level of strategic sophistication corresponding to a specific choice (or choice pattern) are readily available and straightforward. In the Money Request game each subject was simply associated with a level of sophistication by the difference between 10 and their request, e.g. a subject that requested 8 € was assigned k=2 in this game. In the Adding Game the number of subsequent perfect plays before the game ended was taken as the level of sophistication, e.g. a subject that added the current total up to 36, then after the computer turn to 43 and then to 50 or more would be assigned k=2 in this game.

For the Beauty Contest, two different measures were used so far. First, one can start by assuming that a naïve player plays randomly on the interval 1 to 200 and hence in expectation choses 100.5 to which a k=1 player would respond by picking 75.37 (or 75.38), a k=2 player by picking 56.53 and so forth. Secondly, one can start by assuming that levels correspond to elimination by iterative dominance, where a k=1 player understands that 3/4 of the average can never be larger than 150 (neglecting his own influence on the average) and hence pick 112.5, to which a k=2 player were to respond by picking 84.73 (or 84.74) and so forth. Usually, one should construct bins around these respective picks, but for now they have been taken as sharp cutoffs, assigning every subject who picked a number higher than the respective level-k number the next lowest k. While changing this to two-sided rather than one-sided bins, can shift some people from one level to another, the overall relative order is roughly preserved.

All these four measures are shown in the following graphic.

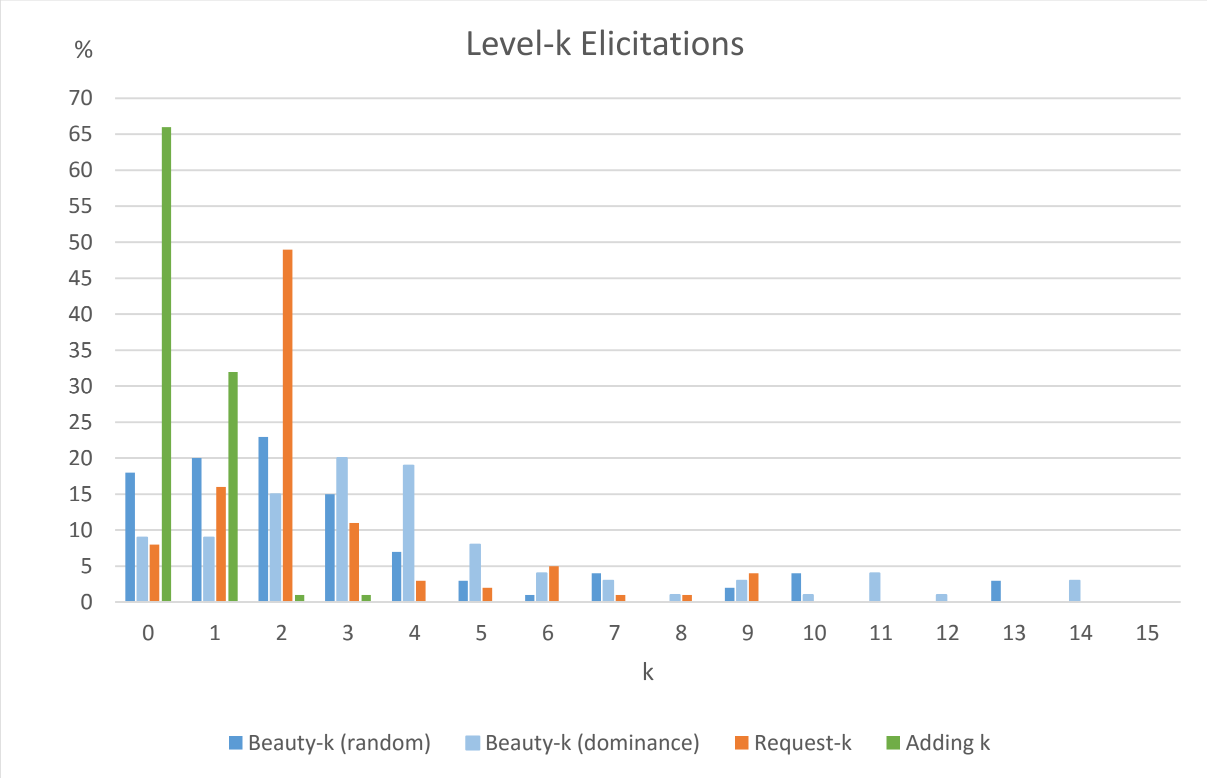

While the two measures elicited from the Beauty Contest game are fairly spread out, they are still single-peaked and attribute a level of 0, 1, or 2 to 61 % (random) or 33 % (dominance) of subjects, being roughly in line with previous research. The shift to higher levels in the dominance elicitation is of course due to the different starting point.

The measure from the Money Request game shows that 49 % of subjects requested 8.00 €, corresponding to a level k=2, and in total 73 % of all subjects were assigned a level of 0, 1, or 2.

The last elicitation method draws a much more pessimistic picture of subjects' ability to solve by backward induction as 66% of subjects did not manage to sum up to 43 and hence guarantee their own victory. 32% managed to reach this point and subsequently secure their victory, whereas only 2% of subjects played perfectly from a point earlier than 43. Being matched against a randomly playing computer of course provides less of an incentive to guarantee victory as one can rely to some extent on probabilities still being in one's favor.

A simpel correlational analysis, though, seems to indicate that the three measures indeed pick up different aspects of strategic sophistication as they are not very correlated to each other (except for the two Beauty Contest measures, of course) the following table shows:

Correlation Coefficients for Level-k Measures

	Beauty-k (random)	**Beauty-k** (dominance)	**Request-k**	**Adding-k**
Beauty-k (random)	1	0.985176282	-0.061647725	0.056184926
Beauty-k (dominance)		1	-0.085821753	0.043662847
Request-k			1	-0.093707765
Adding-k				1

Due to the missing average hiding valuation for subjects that always chose not to hide their valuation, a correlational analysis with these measures is not as straightforward and has not been conducted, yet. An alternative measure, the switching point of a subject from hiding to not hiding, is equally difficult as this is not readily determined for choice schedules that do not qualify as “sensible” due to multiple or reverse switching behavior.

As a first conclusion, though, it seems safe to say, that the usual theoretical assumption of unlimited strategic sophistication, does not fit the data very well, even in Period 7.

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.

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. 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.

Arad, Ayala, and Ariel Rubinstein. 2012. “The 11–20 Money Request Game: A Level-K Reasoning Study.” *The American Economic Review* 102 (7): 3561–73. doi:10.1257/aer.102.7.3561.

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.

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.” SSRN Scholarly Paper ID 2273006. Rochester, NY: Social Science Research Network. http://papers.ssrn.com/abstract=2273006.

Costa-Gomes, Miguel A., and Vincent P. Crawford. 2006. “Cognition and Behavior in Two-Person Guessing Games: An Experimental Study.” *The 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.

Crawford, Vincent P. 2003. “Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions.” *The American Economic Review* 93 (1): 133–49. doi:10.1257/000282803321455197.

Crawford, Vincent P., and Nagore Iriberri. 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): 1721–70. doi:10.1111/j.1468-0262.2007.00810.x.

———. 2007b. “Fatal Attraction: Salience, Naïveté, and Sophistication in Experimental ‘Hide-and-Seek’ Games.” *The American Economic Review* 97 (5): 1731–50. doi:10.1257/aer.97.5.1731.

Dengler, Sebastian, and Jens Prüfer. 2016. “Consumers’ Privacy Choice in the Big Data Era.” *In Progress*.

Dufwenberg, Martin, Ramya Sundaram, and David J. Butler. 2010. “Epiphany in the Game of 21.” *Journal of Economic Behavior & Organization* 75 (2): 132–43. doi:10.1016/j.jebo.2010.03.025.

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.” *The 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 K., Alessandro Acquisti, and George Loewenstein. 2011. "Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information." *Journal of Consumer Research* 37 (5): 858–73. doi:10.1086/656423.

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.

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 *Proceedings 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. http://www.econlib.org/library/NPDBooks/Pigou/pgEW.html.

Posner, Richard A. 1978. "The Right of Privacy." *Georgia Law Review* 12 (3): 393–422.

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

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.

Solove, Daniel J. 2006. "A Taxonomy of Privacy." *University of Pennsylvania Law Review*, 477–564.

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.

———. 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.

Stigler, George J. 1980. "An Introduction to Privacy in Economics and Politics." *The Journal of Legal Studies* 9 (4): 623–44.

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 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.

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. http://www.ntia.doc.gov/page/chapter-1-theory-markets-and-privacy.