On the Viability of Privacy-Enhancing Technologies in a Self-Regulated Business-to-Consumer Market: Will Privacy Remain a Luxury Good?

Rainer Böhme and Sven Koble
Technische Universität Dresden
Institute of Systems Architecture
01062 Dresden, Germany
{rainer.boehme,sven.koble}@inf.tu-dresden.de

WORKING PAPER

Abstract

The collection of personal data in business-to-consumer transactions and respecting the consumers' privacy preferences are fundamentally competing goals. This paper studies the effects of emerging user-controlled privacy-enhancing technologies on supply-side decision making with respect to technology adoption and pricing strategies. In particular, identity management systems that allow buyers to interact pseudonymously with online stores thwart the sellers' efforts to discriminate prices based on personal data. We present stylised micro-economic models (1) to compare self-regulated and government-enforced regimes towards the adoption of privacy-enhancing technologies, (2) to analyse the conditions under which it is profitable for sellers to support such technologies, and (3) to study implication on social welfare and consumer prices.

Keywords: economics of privacy, privacy-enhancing technologies, price discrimination, identity management, customer profiling

Contents

Note. The authors are grateful for helpful comments received from Gaurav Kataria, Karen Kreutel, and two anonymous reviewers of *Wirtschaftsinformatik*, a German journal for business information systems, in which an earlier version of this research has been published [7].

Revision 0.9: Workshop on the Economics of Information Security (WEIS)
Carnegie Mellon University, June 2007 Last update: June 9, 2007

1 Introduction

Collection and analysis of personal data is probably among the most far-reaching developments in retail and sales practices. Customer relationship management (CRM) and data warehousing solutions coupled with online analytical processing applications (OLAP) have become common keywords in corporate marketing divisions and academic business administration departments.

Although the advent of Internet technology has decreased the cost of storing and processing huge amounts substantially, the cost is still not negligible and businesses must see a concrete reason to justify allocation of financial and human resources to data warehousing tasks. There are three possible benefits: First, data about individual (potential) customers allows for targeted communication in the marketing mix and therefore materialises as a higher return on advertising investment. Second, albeit more difficult to quantify, insight in consumer preferences is valuable for the development of new products that are better targeted to the consumers' needs, and thus promises a competitive edge on the market. Third, information about an individual buyers' willingness to pay enables sellers with market power to impose pricing strategies that increase sales and revenues. For example, Acquisti and Varian [4] analyse how sellers with access to a technology for customer tracking can increase their profits by conditioning prices on past behaviour. Such endeavours, however, stand in clear contrast to the individual's right and desire for privacy and informational self-determination. Conversely, the motivation to plead for privacy protection might partly be driven in an attempt to escape price discrimination and thus retain consumer surpluses [34, 4].

This paper aims to shed light on the implications of privacy-enhancing technologies on pricing strategies and market equilibrium prices with stylised micro-economic models. The consensus in early economic research of privacy is that privacy and data protection impose a superfluous burden on flows of information, which are vital to the functioning of a modern economy. For example, Posner [26] concludes that privacy, by reducing the amount of information shared, leads to market-inefficiencies and misallocation of resources. Only recently, alternative views have suggested that privacy by itself has some value. As a result, the degree of privacy granted in a purchase act can be regarded as an additional quality attribute of the traded product. Therefore it is conceivable that privacy-respecting commerce is offered on the market if it is demanded [1, 31]. Existing proposals for privacy-enhancing technologies should ease the realisation of privacy goals particularly in electronic commerce [9, 19].

To capture the range of different aspects in a tractable economic model, we will focus on the interaction of a particular privacy-enhancing technology, namely *privacy-enhancing identity management* (PIM) with one of the aforementioned business reasons for personal data collection, namely pricing strategies. A privacy-enhancing identity management system supports its users to manage and protect their personal information. Following the idea of Chaum [10], it consists of mechanisms that allow pseudonymous interactions between business counterparts. By changing pseudonyms deliberately, users retain full control over which information can be linked to previous interactions. Additional functions support the user in keeping track of disclosed personal data and assure accountability, if desired, by means of cryptographic protocols. In the models presented below we solely regard the property of reducing the flow of personal information, i.e. realising gradual *unlinkability*, which limits the sellers' ability to impose pricing strategies conditional to collected personal data.

This paper is structured as follows. The next section contains a brief review of relevant literature and defines its relation to our research. Section 3 explains the basic relationship between the amount of information disclosed in a business transaction and the possibility to implement certain pricing strategies, in particular price discrimination. As PIM technology limits the flow of information, sellers must compensate this loss by selling to new market segments. By supporting PIM, vendors can attract new customers who value their privacy very much and thus would not purchase otherwise. We will study this basic economic trade-off for self-regulated (i.e. optional) PIM technology as well as for a scenario in which all buyers use PIM by default. Section 4 focuses on the economics of adopting such technology. It has been argued that the acceptance by a large user-base is a prerequisite for the success of privacy-enhancing technologies [2]. Unlike prior work, which addresses this issue mainly with technical means, such as calling for user-friendly systems [12, 21, 6], we believe that support from sellers and intermediaries is equally crucial. Therefore we will study the sellers' incentives to support PIM technology. Further, in section 5, we analyse the consequences of revenue-maximising price setting strategies for individual consumers with varying privacy preferences. Section 6 concludes with a discussion of possible implications and directions for future research.

2 Related work

Many survey results in the literature suggest that people actually do value privacy. The *Electronic Privacy Information Center* [14] cites a number of surveys supporting a general concern about privacy infringements in the U. S. population. According to a special report on data protection of *Eurobarometer* [13], a representative poll of EU citizens, a similar public opinion persists in Europe. More precisely, in 2003, 25 % of the respondents say they are "very concerned" about privacy protection. Some economic research organizations have tried to quantify the loss in sales revenue due to customers abstaining from transactions that do not comply with their privacy preferences. For instance, Jupiter Research predicts an annual loss amount of US$ 24.5 billion in 2006, for the U.S. market only [15, 24]. Westin's series of consumer surveys are not only a valuable source for the evolution of privacy preferences in the U. S. population over time (the earliest studies date back to the 1970) but also provide empirical evidence for heterogeneity in attitudes towards privacy between individuals [23]. This led to the definition of clusters, such as fundamentalists, pragmatists, and unconcerned. However, all self-reported data on topics related to privacy need to be interpreted with utmost care, since a number of studies find huge discrepancies between people's statements and their actual behaviour. This phenomenon is commonly referred to as *privacy paradoxon* [30, 3, 6].

Social implications of pseudonymous transactions have been subject to prior research. Friedman and Resnick [16] apply a game-theoretic framework by formulating a repeated prisoners' dilemma. In their model, changing pseudonyms frequently leads to a situation in which negative reputation does not persist over time. Therefore, mutual trust is reduced, especially in strangers without positive reputation, yielding to an overall decrease in welfare. This loss is characterised as *cost of cheap pseudonyms*. Zwick and Dholakia [35] compare free-market and government regulation approaches to address data protection and privacy concerns from a policy perspective. They argue that a self-regulated market solution is superior because the nature of privacy concerns differs between individuals, whereas any practical regulation would require the existence of a common "one-fits-all" understanding of privacy objectives. Bouckaert and Degryse [8] address the implications of different regulation approaches. They conclude that an *opt-out* policy, where personal information may be exchanged unless the affected individuals express their disagreement explicitly, is superior to *opt-in* (active consent is required before any transmission of personal data) or complete prohibition of personal data processing. Their results, however, are not directly comparable to our analyses due to some rigid model assumptions. Most importantly, Bouckaert and Degryse do not allow for heterogeneous privacy preferences in the population, which we consider as a core attribute of our models presented below. By contrast, Chellappa and Shivendu [11] do consider heterogeneous privacy preferences. They analyse how online service providers should adjust the degree of personalisation to satisfy privacy aware market segments.

As an alternative to user-controlled privacy-enhancing identity management, some marketing literature proposes unconstrained data sharing under the condition that customers are granted direct access to the data collecting organizations' databases to learn and possibly rectify information related to themselves [36, 28]. However, this approach suffers from a weak notion of control, as hidden action can neither be detected nor prevented. Consequently, it requires unlimited trust in all possible transaction counterparts. We will disregard this rather pessimistic scenario in our analysis.

Odlyzko [25], among others, has identified *price discrimination* as one of the main motivations for businesses to collect personal information about their customers. Price discrimination

Table 1: Overview of common assumptions made in all our models

• Market for homogeneous good • Linear demand function • Negligible marginal and transaction costs • Seller has market power • Absence of arbitrage • Heterogeneous (binary) privacy preferences • Privacy-aware buyers sanction seller if privacy-enhancing technology is not supported

occurs if sellers charge buyers different prices for the same product depending on the individual buyer's willingness to pay (see for example [32]). Sometimes referred to as *differential pricing*, this phenomenon has a long research tradition in micro-economics [27]. More recent work also deals with particularities of pricing strategies in electronic commerce (see [5] for a survey). In order to enforce price discrimination, the seller has to know (or infer) the willingness to pay of individual buyers, which can be done on the basis of collected personal data. This creates the link between privacy and pricing strategies.

3 Buyer privacy and price discrimination: A baseline model

Consider an ideal market for a homogeneous good with a monopolistic supplier and a linear demand function over Q_{T} consumers with reservation price p (see Fig. 1). A seller who faces negligible marginal costs—as for information goods and many industrial goods—would set a single price for the entire market to level $\frac{p}{2}$ in order to obtain a profit-maximising revenue of $r = \frac{p}{4} \cdot Q_{\mathrm{T}}$. This corresponds to the rectangular area in Fig. 1, following the theory of monopolistic pricing (see for example [32]; for the sake of brevity we refrain from reporting the formal derivation of profit-maximising conditions, which are analytically tractable solutions of a system of linear equations). If sellers can determine each individual's willingness to pay then they can implement perfect price discrimination and achieve revenue of $r = \frac{p}{2} \cdot Q_{\mathrm{T}}$, twice as much as before (the triangular area under the demand function). Note that there are some conditions to be fulfilled for price discrimination to appear, such as market power of the seller and buyers' inability to resell to third parties (absence of arbitrage).

Generalising the model, we assume two market segments of buyers with different attitudes towards privacy: buyers that are not at all concerned about their privacy (suffix $_{\overline{\mathrm{PA}}}$) and buyers with notable *privacy awareness* (suffix $_{\mathrm{PA}}$). We further assume that the latter group will only participate in a transaction if PIM is supported to protect their personal data. In other words, the utility function of privacy-aware individuals assigns an infinitely high weight to the quality dimension of a privacy-respecting transaction whereas individuals without privacy awareness assign zero weight. The distinction of heterogeneous privacy preferences can be justified against the backdrop of theoretical considerations (privacy needs differ between individuals [35]) as well as of empirical findings [6, 33]. Table 1 provides a summary of all assumptions with the intention to allow for a better assessment of the potential sensitivity of our results to the choice of assumptions.

Fig. 2 displays the demand functions for both market segments, where buyers without privacy awareness $Q_{\overline{\mathrm{PA}}}$ are depicted on the right-hand side, and buyers with high privacy awareness Q_{PA} on

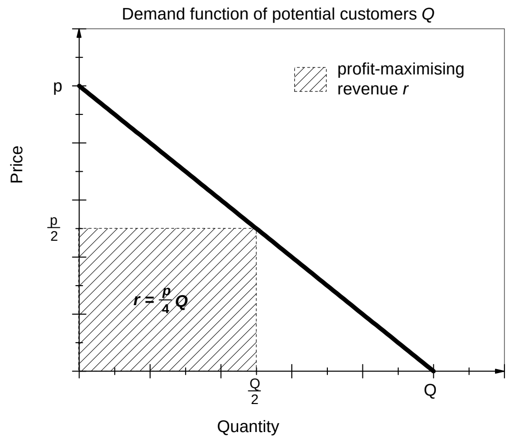

Figure 1: Demand function of Q buyers with reservation price p; no marginal costs. Price $\frac{p}{2}$ to maximises the seller's revenue.

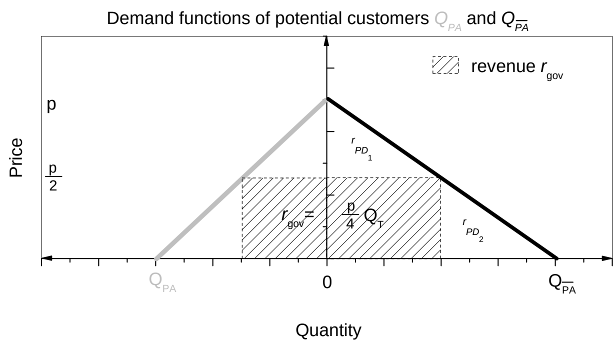

Figure 2: Demand function and profit-maximising revenue if the support of PIM is mandatory (e.g. as a result of government regulation).

the left-hand side (note the inverted quantity scale). Both halves together form the entire market Q_{T}. We define ratio λ in the domain [0,1] as the fraction of buyers *without* privacy awareness:

$$\lambda = \frac{Q_{\overline{\mathrm{PA}}}}{Q_{\overline{\mathrm{PA}}} + Q_{\mathrm{PA}}} = \frac{Q_{\overline{\mathrm{PA}}}}{Q_{\mathrm{T}}} \tag{1}$$

λ can be interpreted as a measure of privacy ignorance in the population. When no PIM is available, privacy-aware buyers will not purchase from the seller. Consequently, with only $Q_{\overline{\mathrm{PA}}}$ buyers left, the seller is able to implement perfect price discrimination within this market segment and achieves revenue of

$$r_{\mathrm{PD}} = \frac{p}{2} \cdot \lambda \cdot Q_{\mathrm{T}} \tag{2}$$

A completely different situation would be obtained if the usage of PIM technology was common practice for all kinds of business transactions, for example due to government regulation. As a result, price discrimination is impossible because all buyers are indistinguishable and the seller has to set one single price for the entire market. Then, as depicted in Fig. 2, the profit-maximising revenue changes to

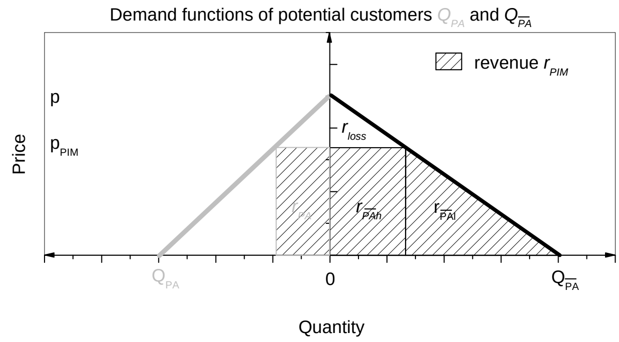

Figure 3: Demand function. r_{loss} is the revenue lost to "strategic buyers" if the seller decides to introduce PIM; however, extra revenue r_{PA} will compensate for this loss.

$$r_{\mathrm{gov}} = \frac{p}{4} \cdot (Q_{\overline{\mathrm{PA}}} + Q_{\mathrm{PA}}) = \frac{p}{4} \cdot Q_{\mathrm{T}} \tag{3}$$

This corresponds to the rectangular area spanning both market segments, since some privacy-aware buyers are willing to purchase now. The seller profits from this situation when the additional revenue in the market segment with high privacy awareness (r_{PA}) exceeds the revenue lost from the missing opportunity to apply perfect price discrimination (triangular areas r_{PD_1} and r_{PD_2}). This point concurs with the condition $\lambda < \frac{1}{2}$, which means that privacy-aware buyers constitute a majority in the population.

Welfare analysis sums up both consolidated supplier and consumer surplus to assess the overall effect of policy choices for the society at large. Interestingly, the same condition $\lambda < \frac{1}{2}$ has to be fulfilled to reach an outcome with higher social welfare than in the (privacy-unfriendly) perfect price discrimination scenario. Otherwise, the additional consumer surplus does not outweigh the losses in seller surplus caused by the inability to differentiate prices.

The scenario up to here is similar to the one described in an earlier workshop version of this research [22]. In this paper we augment the model with a self-regulation approach, where PIM technology is available, but its use is not mandatory. In a first step, the seller decides whether to support the technology or not. If so, each buyer can decide independently whether to use it or not and if all the requested personal data should be revealed. Sellers still implement price discrimination with those buyers of which they can obtain personal data whereas they set one single price p_{PIM} for all buyers that use PIM. This implies that buyers can act strategically: buyers without privacy awareness will choose using PIM not for privacy reasons but to extract surplus if the single price for PIM users is below the individual buyer's willingness to pay.

As illustrated in Fig. 3, the profit-maximising revenue is given by area r_{PIM}:

$$r_{\mathrm{PIM}} = \left[\frac{p}{4 - 2 \cdot \lambda} \right] \cdot Q_{\mathrm{T}} \tag{4}$$

In this scenario, all buyers without privacy awareness and some additional privacy-aware buyers do purchase from the seller. Note that some revenue in the right-hand market segment (upper triangle) is lost due to strategic buyers. However, this loss is over-compensated in all cases by the additional revenue from buyers with high privacy awareness. If no buyers are privacy-aware ($\lambda = 1$), r_{loss} becomes zero, and r_{PD} in (2) equals r_{PIM} in (4). Moreover, welfare effects are always positive compared to a situation without PIM technology (and strictly positive if at least one buyer is privacy aware). This is another indication supporting the view that a self-regulated market solution is superior to government regulation.

4 Viability of privacy-enhancing identity management

In the last section we have explained the basic relationship between customer data processing and pricing strategies. We have argued that a self-regulated approach is most likely superior to government-enforced usage of PIM in all business-to-consumer transactions. In the self-regulated regime, however, PIM technology will only succeed if its implementation is rational from a cost-benefit perspective. This section deals with the question when and under which conditions sellers would opt to support PIM technology. We do this by extending the baseline model of section 3, which mainly served as an introductory example, to meet more realistic scenarios. In the model to be developed in this section, PIM technology will be optional and we assume a binary market segmentation for both *privacy awareness* as well as for *willingness to pay*. Later in subsection 4.2 we will allow for linear dependence between the two attributes.

4.1 Optional PIM with binary market segmentation

The validity of the baseline model is limited by the assumption of perfect price discrimination. In reality, sellers often do not know each individual buyer's willingness to pay. This motivates an extension of the model to one in which the ability to price discriminate is constrained: sellers can only infer one bit of information about each buyers's willingness to pay. This means, they can tell for each buyer whether his or her willingness to pay is above (suffix $_h$ for *high*) or below (suffix $_l$ for *low*) some limit price p_{sep} (suffix $_{\text{sep}}$ for *separation*). The limit price is given exogenously, i.e. a single seller has no influence on it. One may think of a discrete criterion of civil status, such as *student* (low willingness to pay) or *employee* (high willingness to pay). As sellers can observe the demand function, they can estimate the numbers of buyers with high Q_h and low Q_l willingness to pay, respectively. Note that p_{sep} is equal for all buyers regardless of their privacy awareness. For the given linear demand function, p_{sep} is directly related to the fraction of buyers with high willingness to pay π.

$$\pi = \frac{Q_h}{Q_{\text{T}}} = 1 - \frac{p_{\text{sep}}}{p} \tag{5}$$

With *willingness to pay* and *privacy awareness* being two orthogonal dimensions, we have defined a model that segments buyers in four groups, as shown in Fig. 4. Given the parameters π, λ, and p (the reservation price of the demand function), sellers aim to maximise their revenue. In the absence of PIM, they do so by setting two prices, p_l and p_h, for buyers with low and high willingness to pay, respectively. Although we are dealing here with a multi-parameter optimisation problem, the specific setting in our model allows us to find the individual prices independently. Sellers apply the standard monopolistic pricing for the section of the demand function below p_{sep} to extract the maximum revenue from the market segment with low willingness to pay.

$$p_l = \frac{1}{2} \cdot p_{\text{sep}} = \frac{1}{2} \cdot p \cdot (1 - \pi) \tag{6}$$

where the second identity follows from (5). The choice of p_h depends on the size of the market segment with high willingness to pay (π). If buyers with high willingness to pay are in the majority ($\pi \geq \frac{1}{2}$), then sellers use monopolistic price setting as if it were for the entire market since all buyers with low willingness to pay are to be found in the section of the demand curve that would have been unsatisfied without price discrimination (see upper chart of Fig. 5 for illustration). In the opposite case ($\pi < \frac{1}{2}$) it is optimal to set $p_h = p_{\text{sep}}$, which is the closest possible solution to the unique monopoly price (Fig. 5, bottom). Therefore,

$$p_h = \begin{cases} \frac{1}{2} \cdot p & \text{for} \quad \pi \geq \frac{1}{2} \\ p \cdot (1 - \pi) & \text{for} \quad \pi < \frac{1}{2} \end{cases} \tag{7}$$

The corresponding revenues are given as follows:

$$r_{\overline{\text{PIM}}} = \frac{1}{4} \cdot \lambda \cdot p \cdot K(\pi) \cdot Q_{\text{T}} \tag{8}$$

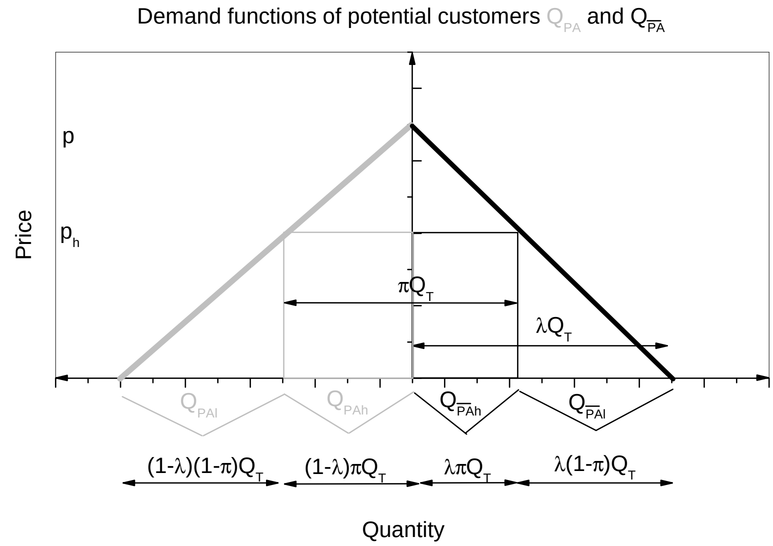

Willingness to pay	**Privacy awareness** high	low
high	$Q_{\mathrm{PAh}} = (1-\lambda)\cdot\pi\cdot Q_{\mathrm{T}}$	$Q_{\overline{\mathrm{PA}}\mathrm{h}} = \lambda\cdot\pi\cdot Q_{\mathrm{T}}$
low	$Q_{\mathrm{PAl}} = (1-\lambda)\cdot(1-\pi)\cdot Q_{\mathrm{T}}$	$Q_{\overline{\mathrm{PA}}\mathrm{l}} = \lambda\cdot(1-\pi)\cdot Q_{\mathrm{T}}$

Figure 4: Market segmentation by *willingness to pay* and *privacy awareness*. The formulas for segment sizes depend on parameters Q_{T}, λ and π. Graph of demand function (top) and crosstab representation (bottom)

$$\text{where} \qquad K(\pi) = \begin{cases} 1+(1-\pi)^2 & \text{for } \pi \geq \frac{1}{2} \\ 1+2\cdot\pi-3\cdot\pi^2 & \text{otherwise.} \end{cases} \tag{9}$$

Sellers who decide to support PIM have to find another price p_{PIM} for the users of PIM. Note that p_{PIM} imposes an upper bound for p_h because of strategic buyers. Consequently, p_{PIM} should never be set to p_l (or below) if the seller wishes to maintain the possibility to price discriminate. Abandoning price discrimination completely would mechanically imply a decreasing revenue. It turns out that the optimal setting of all three prices $(p_l, p_h, p_{\mathrm{PIM}})$ does not affect the choice of p_l. And p_{PIM} is set to the same level of p_h as follows:

$$p_{\mathrm{PIM}} = p_h = \begin{cases} \frac{1}{2}\cdot p & \text{for} \quad \pi \geq \frac{1}{2} \\ p\cdot(1-\pi) & \text{for} \quad \pi < \frac{1}{2} \text{ and } \lambda \geq \frac{1-2\pi}{1-\pi} \\ \frac{1}{2}\cdot p\cdot(1+\frac{\lambda\pi}{1-\lambda}) & \text{for} \quad \pi < \frac{1}{2} \text{ and } \lambda < \frac{1-2\pi}{1-\pi} \end{cases} \tag{10}$$

Now we will compare the optimal revenues when PIM is supported to the situation without PIM. For the comparison we regard both cases separately.

Comparison for $\pi \geq \frac{1}{2}$ (majority has high willingness to pay):

- Revenue *with* PIM:

$$r_{\mathrm{PIM}} = \frac{1}{4}\cdot p\cdot\big(\lambda\cdot(1-\pi)^2+1\big)\cdot Q_{\mathrm{T}} \tag{11}$$

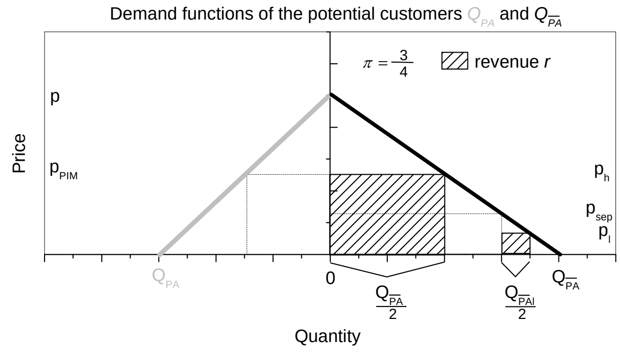

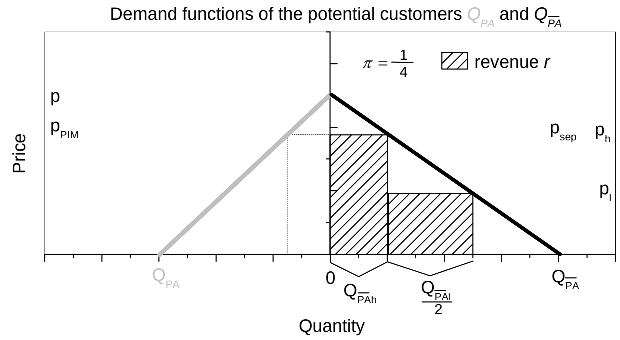

Figure 5: Demand function with parameters p_{sep} and π ; profit-maximising revenues without PIM when price discrimination is imposed to buyers with low privacy awareness. Satisfied demand differs for $\pi \geq \frac{1}{2}$ (top) and $\pi < \frac{1}{2}$ (bottom).

- The revenue *without* PIM follows from (8) after re-substitution of $K(\pi)$:

$$r_{\overline{\text{PIM}}} = \frac{1}{4} \cdot p \cdot \left(\lambda \cdot (1-\pi)^2 + \lambda\right) \cdot Q_{\text{T}} \tag{12}$$

As $\lambda \leq 1$ by definition, the revenue with PIM is always higher or equal to the revenue without PIM.

Comparison for $\pi < \frac{1}{2}$ (majority has low willingness to pay):

- Revenue *with* PIM for $\lambda \geq \frac{1-2\pi}{1-\pi}$:

$$r_{\text{PIM}} = \frac{1}{4} \cdot p \cdot \left[\lambda \cdot (1-\pi)^2 + 4 \cdot \pi \cdot (1-\pi)\right] \cdot Q_{\text{T}} \tag{13}$$

- Revenue *with* PIM for $\lambda < \frac{1-2\pi}{1-\pi}$:

$$r_{\text{PIM}} = \frac{1}{4} \cdot p \cdot \left[\frac{\lambda \cdot \pi^2}{1-\lambda} + 1\right] \cdot Q_{\text{T}} \tag{14}$$

- The revenue *without* PIM follows from (8) after re-substitution of $K(\pi)$:

$$r_{\overline{\text{PIM}}} = \frac{1}{4} \cdot p \cdot \left[\lambda \cdot (1-\pi)^2 + 4 \cdot \pi \cdot (1-\pi) \cdot \lambda\right] \cdot Q_{\text{T}} \tag{15}$$

It is easy to see that r_{PIM} in (13) is higher than $r_{\overline{\mathrm{PIM}}}$ in (15) as long as $\lambda < 1$. Subtracting (14) from (15) yields extra revenue r_e as "return on PIM", which is also strictly positive for $\lambda < 1$:

$$r_e \quad = \tfrac{1}{4} \cdot p \cdot \Big[\; \underbrace{4 \cdot \lambda \cdot \pi \cdot (1-\pi)}_{\geq 0} \; + \; \underbrace{\frac{\lambda \cdot \pi^2}{1-\lambda}}_{\geq 0} \; + \; \underbrace{1 - \lambda \cdot (1-\pi)^2}_{> 0 \text{ for } \lambda < 1} \; \Big] \cdot Q_{\mathrm{T}} > 0 \tag{16}$$

Therefore, in all cases, $\pi \geq \frac{1}{2}$ and $\pi < \frac{1}{2}$, the revenue *with* privacy-enhancing technology exceeds the benchmark level if at least some buyers are privacy-aware ($\lambda < 1$). The factor by which the revenue increases varies with the number of buyers that value privacy (related to λ) and the size of the market segments that can be separated with p_{sep} to implement price discrimination (related to π). As visualized in Fig. 6, the gains are comparatively lower when buyers with low willingness to pay constitute 60–80 % of the market. In these situations price discrimination is most effective, and sellers cannot sell to privacy-aware buyers with low willingness to pay because reducing p_{PIM} further would sacrifice the high margins from affluent buyers that would start using PIM for strategic reasons.

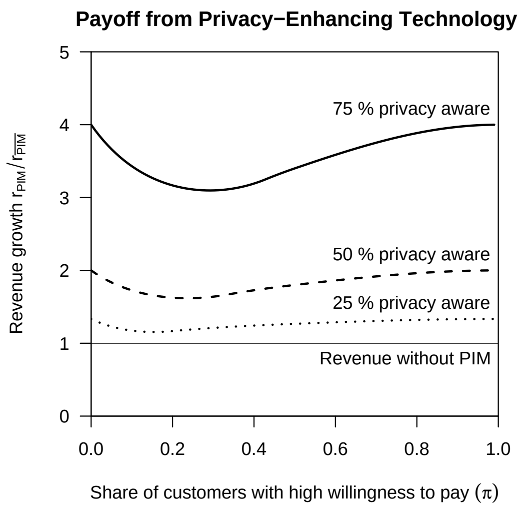

Figure 6: Potential increase in revenue after the introduction of PIM. A value of 1 corresponds to the revenue in a monopolistic price discrimination scenario without support for privacy-enhancing technologies. The graphs show different assumptions for the fraction of privacy-aware buyers in the population $(1-\lambda)$.

4.2 Linear dependence between buyer attributes

So far, the model is lacking an important property of reality as it assumes that the dimensions willingness to pay and privacy awareness are independently distributed in the population. Empirical evidence, however, suggests a positive correlation between willingness to pay and privacy awareness, i.e. that the wealthy are likely to be more privacy aware. Varian et al. [33] report this fact from an analysis of do-not-call lists in the United States. The positive correlation may be explained by factors such as affluent individuals being targeted by direct marketers more frequently and thus decide to subscribe to do-not-call lists, or alternatively, that wealthy people tend to value time more. A similar trend can be observed in representative survey data of EU citizens from Eurobarometer [13]. For example, in 2003, 13 % of managers as opposed to 3 % of house persons, 4 % of manual workers and just 2 % of the retired reported to use privacy-enhancing technologies—including encryption tools.

Table 2: Market segmentation with correlation: segment sizes depend on Q_{T}, λ, π and ρ.

Willingness to pay	**Privacy awareness** high	low
high	$Q_{\mathrm{PAh}} = Q_{\mathrm{T}} \cdot \Big[(1-\lambda)\cdot\pi + \rho\cdot\sqrt{\lambda(1-\lambda)\pi(1-\pi)}\Big]$	$Q_{\overline{\mathrm{PA}}\mathrm{h}} = Q_{\mathrm{T}} \cdot \Big[\lambda\cdot\pi - \rho\cdot\sqrt{\lambda(1-\lambda)\pi(1-\pi)}\Big]$
low	$Q_{\mathrm{PAl}} = Q_{\mathrm{T}} \cdot \Big[(1-\lambda)\cdot(1-\pi) - \rho\cdot\sqrt{\lambda(1-\lambda)\pi(1-\pi)}\Big]$	$Q_{\overline{\mathrm{PA}}\mathrm{l}} = Q_{\mathrm{T}} \cdot \Big[\lambda\cdot(1-\pi) + \rho\cdot\sqrt{\lambda(1-\lambda)\pi(1-\pi)}\Big]$

We use a measure of dependence between two variables based on Pearson's χ^2 statistic, which is the sum of squared differences between the actual size of the segments and the expected size if buyer attributes were independent. After adjusting for the total market size we obtain a measure ϕ^2 for the strength of the linear relation:

$$\phi^2 = \frac{1}{Q_{\mathrm{T}}^2}\cdot\Bigg[\begin{array}{ccc} \frac{(Q_{\mathrm{PAh}}-(1-\lambda)\cdot\pi\cdot Q_{\mathrm{T}})^2}{(1-\lambda)\cdot\pi} & + & \frac{\left(Q_{\overline{\mathrm{PA}}\mathrm{h}}-\lambda\cdot\pi\cdot Q_{\mathrm{T}}\right)^2}{\lambda\cdot\pi} \quad + \\ \frac{(Q_{\mathrm{PAl}}-(1-\lambda)\cdot(1-\pi)\cdot Q_{\mathrm{T}})^2}{(1-\lambda)\cdot(1-\pi)} & + & \frac{\left(Q_{\overline{\mathrm{PA}}\mathrm{l}}-\lambda\cdot(1-\pi)\cdot Q_{\mathrm{T}}\right)^2}{\lambda\cdot(1-\pi)} \end{array}\Bigg] \tag{17}$$

Note that by this definition, ϕ^2 quantifies dependence *strength*, but it lacks an indication of its *direction*. Therefore we define a correlation coefficient ρ as

$$\rho = \begin{cases} \sqrt{\phi^2} & \text{for} \quad \frac{Q_{\mathrm{PAh}}}{Q_{\mathrm{PA}}} > \frac{Q_{\overline{\mathrm{PA}}\mathrm{h}}}{Q_{\overline{\mathrm{PA}}}} \\ 0 & \text{for} \quad \frac{Q_{\mathrm{PAh}}}{Q_{\mathrm{PA}}} = \frac{Q_{\overline{\mathrm{PA}}\mathrm{h}}}{Q_{\overline{\mathrm{PA}}}} \\ -\sqrt{\phi^2} & \text{for} \quad \frac{Q_{\mathrm{PAh}}}{Q_{\mathrm{PA}}} < \frac{Q_{\overline{\mathrm{PA}}\mathrm{h}}}{Q_{\overline{\mathrm{PA}}}} \end{cases} \tag{18}$$

The domain of the correlation coefficient is $\rho \in [-1, 1]$, where values $\rho < 0$ denote that privacy awareness on average concurs with low willingness to pay whereas $\rho > 0$ indicate that privacy-aware buyers are more likely to have high willingness to pay.

Tab. 2 shows how the sizes of the four market segments $Q_{\overline{\mathrm{PA}}\mathrm{l}}$, $Q_{\overline{\mathrm{PA}}\mathrm{h}}$, Q_{PAl} and Q_{PAh} can be calculated from the exogenous parameters Q_{T}, λ, π and ρ. For $\rho = 0$, the second terms in the brackets become zero and the situation reduces to the independence model discussed above. The set of non-negativity constraints for the market segments ($Q_{\overline{\mathrm{PA}}\mathrm{l}} \geq 0 \wedge Q_{\overline{\mathrm{PA}}\mathrm{h}} \geq 0 \wedge Q_{\mathrm{PAl}} \geq 0 \wedge Q_{\mathrm{PAh}} \geq 0$) limits the domain of permissible combinations for parameters (λ, π, ρ) as follows:

$$\frac{\rho^2\cdot f(\lambda)}{\rho^2\cdot f(\lambda) + (1-f(\lambda))} \leq \pi \leq \frac{f(\lambda)}{\rho^2\cdot(1-f(\lambda)) + f(\lambda)} \quad , \tag{19}$$

$$\text{where} \quad f(\lambda) = \begin{cases} \lambda & \text{for} \quad \rho > 0 \\ \frac{1}{2} & \text{for} \quad \rho = 0 \\ 1-\lambda & \text{for} \quad \rho < 0 \end{cases}$$

For $\rho = 0$, condition (19) is true for any combination $(\lambda, \pi) \in [0, 1]^2$. This is in line with expectations as the model of the previous section is a special case of this more general model. Fig. 7 illustrates the domain of parameters (λ, π) for various values of ρ.

The correlation parameter has also an impact on the shape of the demand function, which is still assumed to be linear in both market segments along the dimension privacy awareness. While in the baseline model both demand curves intersect at the same reservation price p, now

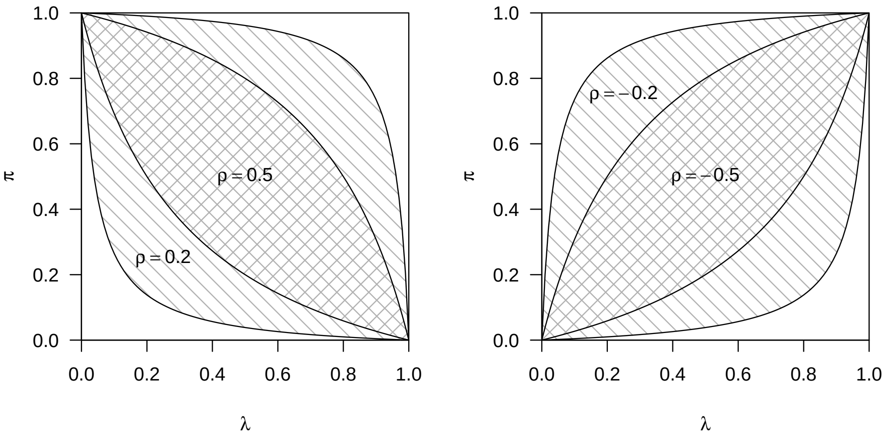

Figure 7: Domain of permissible parameters (λ, π) for different degree of correlation ρ between privacy awareness and willingness to pay (left: positive, right: negative)

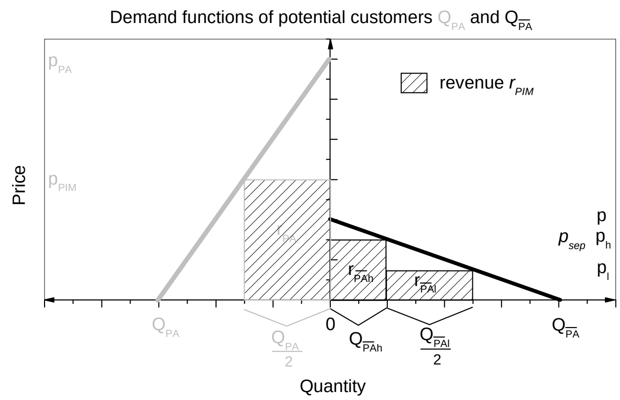

Figure 8: Demand functions for $p_{\mathrm{PIM}} > p_h$ and $\rho > 0$.

buyers with high and low privacy awareness may exhibit different reservation prices. We write the reservation price for privacy-aware buyers as p_{PA} as opposed to p. Fig. 8 shows a demand curve for positive correlation ($\rho > 0$), where p_{PA} is above p. Conversely, if $\rho < 0$ then $p_{\mathrm{PA}} < p$ (Fig. 9). Note that p_{PA} is fully defined by parameters λ, π, ρ, and p, although we omit the expression as it is rather bulky. The separation price p_{sep} is still assumed to be constant for both market segments and can be calculated as follows:

$$p_{\mathrm{sep}} = p \cdot \left(1 - \pi + \frac{\rho}{\lambda} \cdot \sqrt{\lambda(1-\lambda)\pi(1-\pi)}\right) \tag{20}$$

To assess the profitability of PIM, we will discuss the cases $\rho > 0$ and $\rho < 0$ separately. For positive correlation ($\rho > 0$), we find that the introduction of PIM is always rewarded with higher revenues. The intuition behind this proposition follows from the conclusion of the previous section, i.e. PIM is never disadvantageous when $\rho = 0$. As the reservation price for privacy-aware buyers p_{PA} is greater than p, a seller could always act as in the independent case by assuming $p = p_{\mathrm{PA}}$ (see Fig. 8). This means that the introduction of PIM is worthwhile despite leaving some extra consumer surplus to privacy-aware buyers. This is sufficient to back the proposition. Smart sellers would certainly employ a more appropriate price setting (which is complicated and not further detailed in this context) and thus increase revenues even further. As illustrated in the left chart

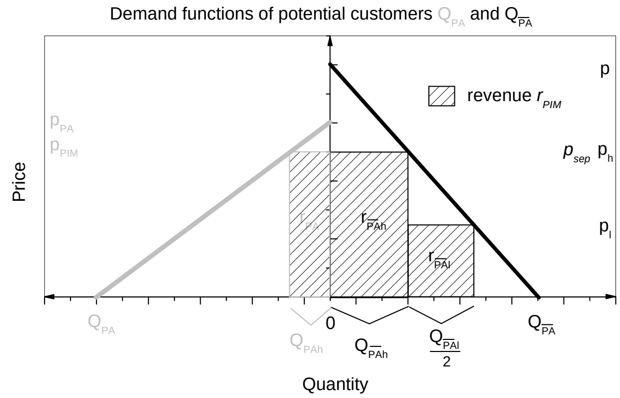

Figure 9: Demand functions for $p_{\mathrm{PIM}} = p_h$ and $\rho < 0$.

of Fig. 10, the "return on PIM" is much higher than in the independent case when the prices are set in a revenue-maximising way and π approaches 1.

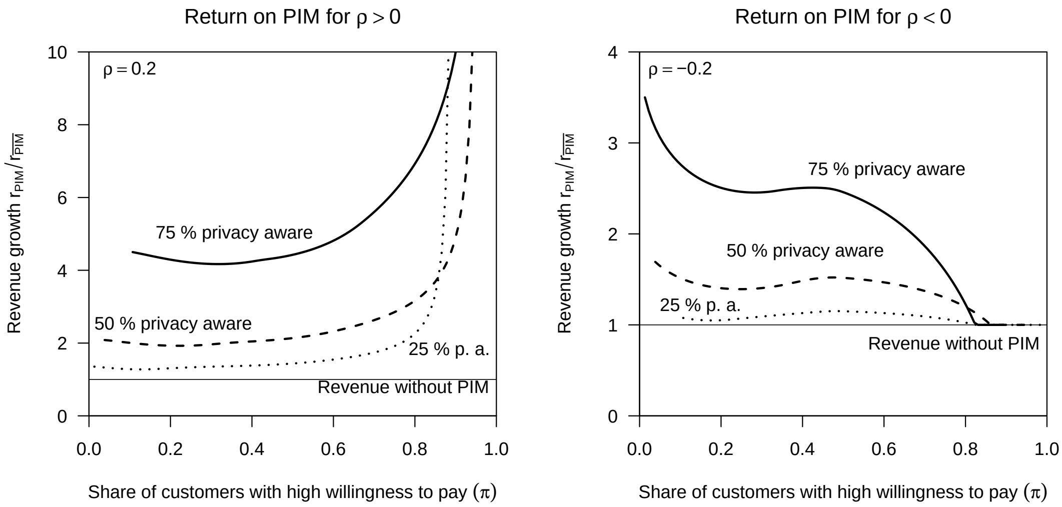

Figure 10: Potential increase in revenue after the introduction of PIM. A value of 1 corresponds to the revenue in a monopolistic price discrimination scenario without support of privacy-enhancing technologies. The graphs show different assumptions for correlation (left: $\rho > 0$, right: $\rho < 0$) and the fraction of privacy-aware buyers $(1 - \lambda)$. Compare with Fig. 6 for the independent case.

For $\rho < 0$, however, non-trivial cases exist in which the introduction of PIM is *not* profitable, as can be seen in the right chart of Fig. 10. This happens when the reservation price for privacy-aware buyers p_{PA} falls too far below p_h. Selling to the privacy-aware segment would then sacrifice large parts of the revenue from buyers with high willingness to pay and low privacy awareness so that sellers are better off if they do not support PIM at all. The situation occurs for high π, but the exact threshold depends also on λ. Fig. 11 shows the shape of the demand function in such a situation and Fig. 13 visualises those regions of combinations (λ, π) in which the introduction of PIM is profitable despite a negative correlation. It becomes apparent that for moderate negative correlation, PIM is still supported in large fractions of the joint domain of λ and π. However, we have to bear in mind that a strong negative correlation between privacy awareness and willingness

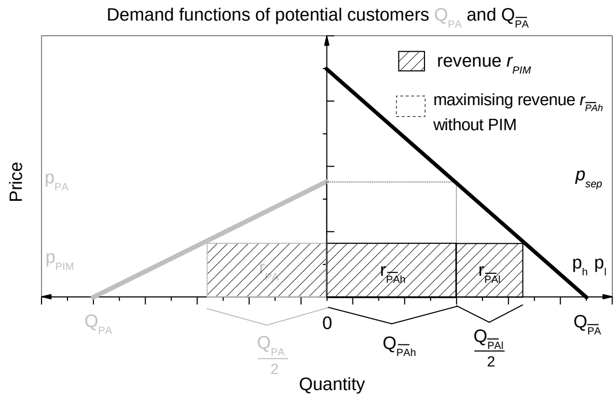

Figure 11: Demand functions for $\rho < 0$. Without supporting PIM the seller would extract higher revenue from buyers without privacy awareness and high willingness to pay $Q_{\overline{\mathrm{PA}}\mathrm{h}}$ (opaque rectangle).

to pay might form a serious market entry barrier for PIM technology, which intensifies even more if transaction costs would be taken into account.

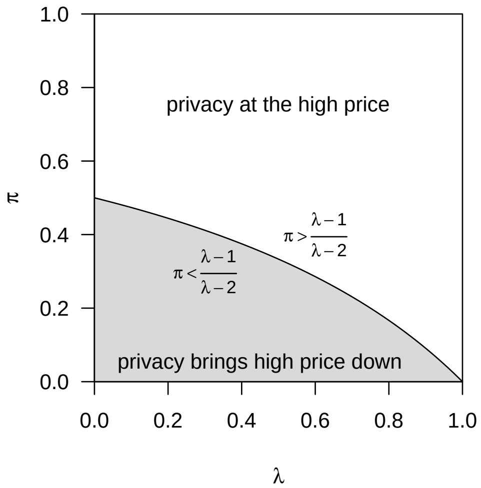

Figure 12: Prices p_{PIM} charged from privacy-aware buyers in comparison to the high prices in the default situation (price discrimination without the option to use PIM). Under certain conditions the optimal p_{PIM} is below the otherwise high price, but it will never come down to the low price.

5 Privacy at a premium

The model in the previous section has been set up with the intention to study the conditions under which rational sellers will decide to support privacy-enhancing technologies despite losing the opportunity to pursue price discrimination (at least in part of the market). This analysis was based on a comparison of expected revenues. In this section we are using the same market model; however we will focus on the prices buyers with and without privacy awareness will have to pay. Hence, after regarding the supply-side in the previous sections, we are now switching to a consumer

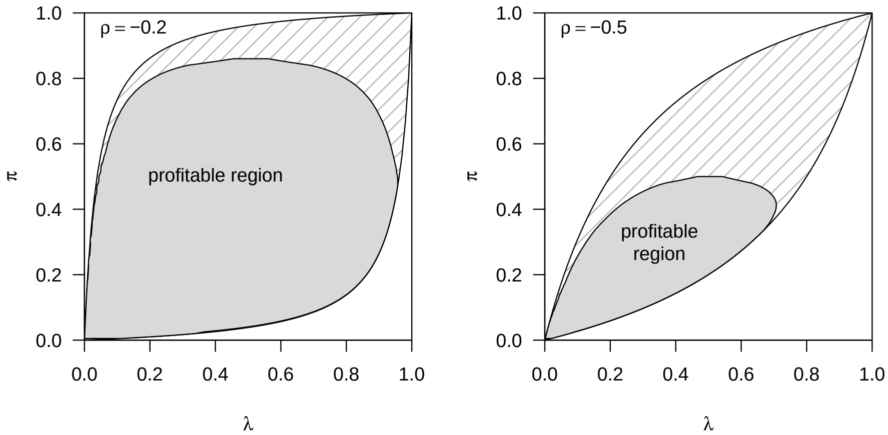

Figure 13: Combinations of ρ and λ where the introduction of PIM is profitable for the seller when willingness to pay and privacy awareness are negatively related ($\rho < 0$). Hatched regions show the domain of (λ, π) for given ρ (cf. Fig. 7, right).

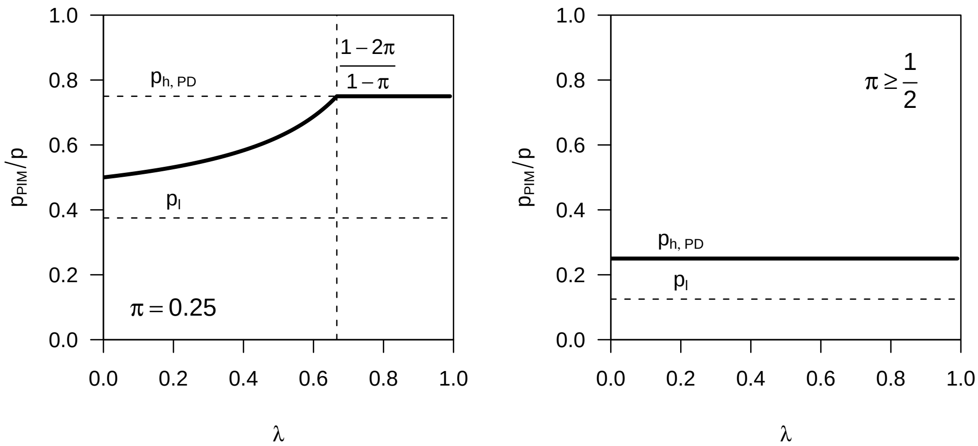

Figure 14: Prices charged from users of PIM compared to the high $p_{h,\mathrm{PD}}$ and low $p_{l,\mathrm{PD}}$ price in a situation with price discrimination in the absence PIM technology. Figures are scaled to a unit of $p = 1$, $\pi = \frac{1}{4}$ (left) and $\pi = \frac{3}{4}$ (right).

perspective. We will first discuss the implication for prices in the case of independent privacy awareness and willingness to pay before we advance to situations with positive and, respectively, negative correlation.

To assess the “price of privacy”, we have to set p_{PIM} into relation to the prices $p_{h,\text{PD}}$ and $p_{l,\text{PD}}$ if no PIM is supported (indicated by suffix $_{\text{PD}}$). These prices are calculated based on equations (6) and (7) above. As mentioned before, the low price p_l is not affected by the decision to support privacy-enhancing technology. However, the optimal high price as given in equation (10) may differ after the introduction of PIM, depending on both λ and π. The first two rows of (10) exactly correspond to (7), therefore p_{PIM} differs from $p_{h,\text{PD}}$ only if $\pi < \frac{\lambda-1}{\lambda-2}$. If this condition is fulfilled, the revenue-maximising prices p_{PIM} and p_h are *below* the original high price $p_{h,\text{PD}}$ when no PIM is supported (see Fig. 14, left chart). In other words, when buyers with low willingness to pay form a majority and the fraction of privacy-aware buyers is above a certain threshold, then the introduction of PIM not only increases the revenues for the seller but also slightly reduces the price for buyers with low privacy awareness and high willingness to pay (but will not alter the price for those with low willingness to pay). Fig. 12 shows regions in the (λ, π)-plane in which this happens.

In this model p_l is always a lower bound for p_{PIM}, as can be seen by subtracting (6) from the last equality of (10):

$$\left(1 + \frac{\lambda \cdot \pi}{1-\lambda}\right) - (1-\pi) = \frac{\pi}{1-\lambda} > 0 \quad \text{for } \pi > 0 \tag{21}$$

Obviously, p_{PIM} is equal to p_l only if $\pi = 0$, which means that there are no buyers with high willingness to pay and hence price discrimination would not be possible at all.

In contrast, when buyers with high willingness to pay are in the majority—think of a distinguishing criterion like student cards, where students are a minority in the population—, the optimal price is $p_{\text{PIM}} = \frac{1}{2} \cdot p$, which is greater than $p_{\text{sep}} = p \cdot (1-\pi)$. This means that people with low willingness to pay cannot afford to enforce their privacy preferences and thus privacy becomes a premium product. Note that this is an analytical result with the prior that privacy awareness and willingness to pay are independent. In the light of these findings, the common interpretation of an empirically reported positive correlation between privacy awareness and willingness to pay might need some reconsideration: it is well possible that the evidence for affluent consumers being on average more privacy-aware is not a natural or behavioural phenomenon by itself, but rather a consequence of market mechanisms that make privacy a premium product, which is not affordable by the entire population.

Finally, we turn to situations where $\rho \neq 0$, but for the sake of brevity we omit the analytical deviations and proof ideas. If privacy awareness and willingness to pay are *positively* correlated ($\rho > 0$), then p_{PIM} will be below $p_{h,\text{PD}}$ under exactly the same conditions as in the independent model, with the exception that the domain of (λ, π) is reduced as specified in equation (19). However, when π exceeds the threshold $\frac{\lambda-1}{\lambda-2}$, the seller can raise p_{PIM} above p_h because the reservation price of privacy-aware buyers allows for a higher equilibrium price. In other words, price discrimination by customer attributes is complemented with price discrimination by customer behaviour, as observed in the preference for privacy-friendly transactions. It is important to note that privacy-aware buyers will *not* act strategically because it would violate their privacy preference (though we acknowledge that this is a debatable assumption). We call this situation “privacy is expensive”: supply for opportunities to realise privacy objectives is made artificially scarce. This strengthens the arguments given above that privacy might be a premium product by its very nature. And it becomes even more expensive if privacy-aware buyers are known to be more affluent on average ($\rho > 0$). The top two charts of Fig. 15 show the regions for either case.

The situation becomes analytically more demanding if privacy awareness and willingness to pay are *negatively* correlated ($\rho < 0$). After accounting for the unprofitable combinations of (λ, π), we see that the relative proportion of situations in which the introduction of PIM is accompanied by lower prices $p_h = P_{\text{PIM}}$ increases (see bottom charts of Fig. 15). It is also noteworthy that p_{PIM} can drop to p_l in the marginal case where no privacy-aware buyers exhibit a high willingness to

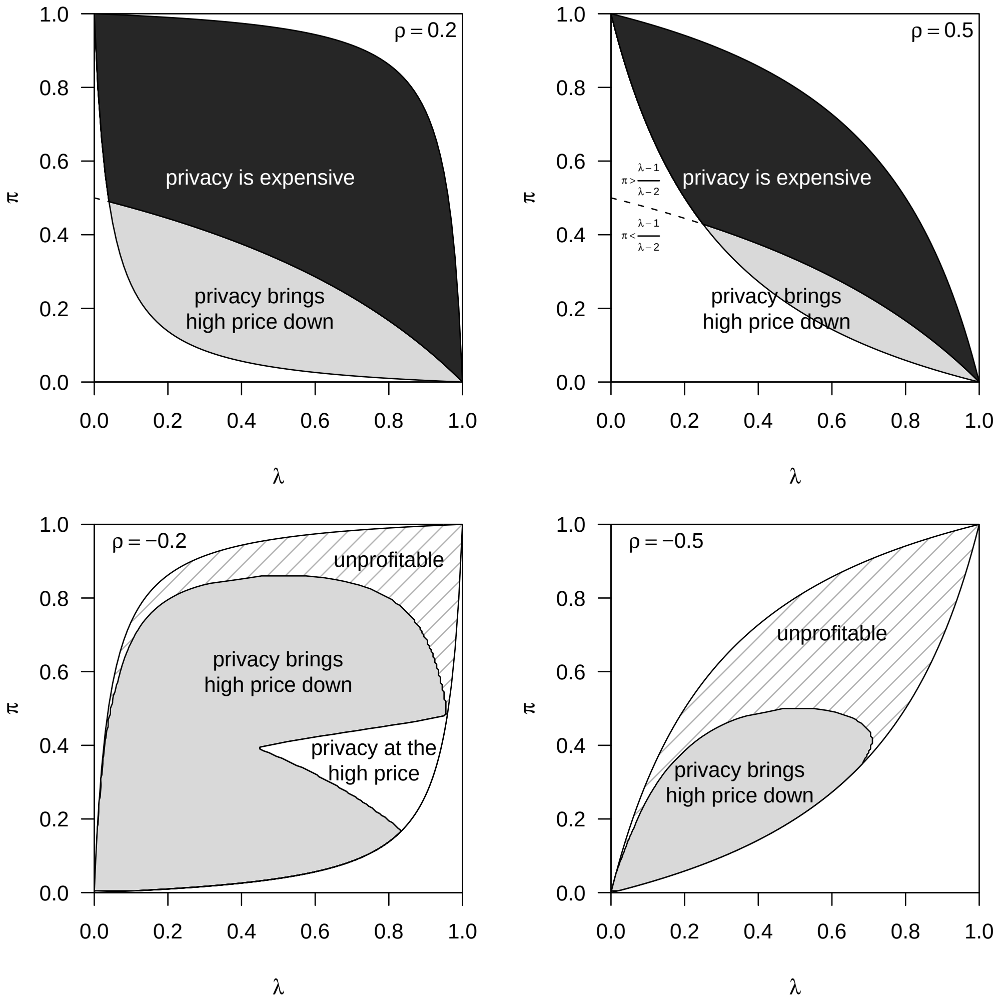

Figure 15: Prices p_{PIM} charged from privacy-aware buyers in comparison to the high prices in the default situation without PIM. For positive correlation $\rho > 0$ (top row), price discrimination by privacy preferences can be implemented and p_{PIM} raises above p_h. Negative correlation (bottom) increases the regions where a seller should lower the high price to sell to privacy-aware buyers.

pay (but certainly some buyers with low willingness to pay do have!). In such cases, the demand for privacy is strong enough to force the seller to set one single price $p_{\mathrm{PIM}} = p_h = p_l$ in all market segments. However, the precondition that there must not be one single privacy-aware buyer with high willingness to pay shows how unlikely such cheap privacy actually is. Even when the linearity constraint of the demand function would be replaced by a weaker assumption, Q_{PAh} must remain negligibly small. Therefore, we deem it justified to conclude this section as follows: in many situations users of privacy-emhamcing identity management systems will be charged an additional premium by sellers who otherwise would be able to price discriminate. It is possible that this "privacy tax", accumulated with acceptance problems of different nature, could hinder a wide deployment of such technologies by large parts of the population.

6 Summary and conclusion

As this paper has tried to show, new developments in the area of privacy-enhancing technologies in combination with persistence of privacy concerns in the population has tremendous implication on business-to-consumer relations. More precisely, sellers might have to give up one of the main advantages of electronic commerce: the power of processing personal information. Our analysis revealed that in most cases sellers can increase revenues—and thus profits—by voluntarily supporting interfaces for privacy-enhancing technologies even if this implies refraining from collecting customer information for the purpose of price discrimination. This proposition holds true for a variety of conditions, depending on the degree of price discrimination that could be realised through customer data processing: if sellers were able to implement perfect price discrimination (which rarely happens) and had to offer privacy-enhancing technologies to the entire market (which seems even less realistic), then to break even, the share of buyers that value privacy very much would have to exceed 50 %. In a more realistic scenario in which price discrimination is imperfect and based on a single binary attribute, the option to use privacy-enhancing technologies increases revenues as soon as there are a non-negligible number of privacy-aware buyers. This can be interpreted as the indication that privacy-enhancing technologies may thrive on the market, or—more prudently—that, at least in principle, no economic market entry barrier arises from the inability to employ price discrimination.

Assuming that optional support for privacy-enhancing technologies is commonplace, sellers are still able to implement some price discrimination, albeit on a meta-level, where the sole preference to use privacy-enhancing technology serves as new distinguishing criterion. As a result, rational sellers will opt to define a specific price for privacy-aware buyers. Our analyses show that this price will most likely be higher than the lowest price for buyers who accept to reveal personal information. This leads us to the notion that privacy is likely to remain a "luxury good", which consequently will not be affordable by the entire population. We acknowledge that this might be a controversial—and perhaps polarising—finding, the valuation of which we leave for others. However, it is somewhat surprising to see this result as a corollary of an analytic model as this fact is quite well supported by evidence in the literature [13, 33], where the premium status of privacy has previously been regarded as merely empirical phenomenon.

If privacy-enhancing technologies do not set out to conquer the market quickly, privacy activists might be tempted to lobby for government regulation to enforce the support of such technologies. Apart from anticipated difficulties in implementing such legislation, regulation by the government might also turn out to be a sub-optimal policy that could ceteris paribus lead to a decrease in social welfare. This finding concurs with related work in a competition policy context, where an abolition of price discrimination (by legal means) is reported to result in lower competitive pressure and hence a higher level of consumer prices [17].

This leads us to the main limitation of our analysis, namely the assumption of market power in a monopolistic modelling framework. This assumption is not completely ill-aligned since a number of real markets are structured as monopolistic competition (e.g. media) or artificially allow for market power through other imperfections, such as switching costs (e.g. software) [29]. But it does not cover all possible market structure in general, either. It is quite obvious that privacy-enhancing technologies will increase revenues in the case of perfect competition because here price discrimination is much more limited; if not impossible at all (the same rationale applies for the existence of arbitrage). The case of close oligopolies with strategic interdependencies between players remains a gap to be closed in future research. Another promising direction could be to replace the binary concept of privacy awareness with some sort of continuous elasticity measure. This would allow for substitution between privacy goals and monetary compensation and therefore provide a framework to better model the often-reported phenomenon that consumers are willing to give up privacy principles for fairly small rebates [20, 6, 18]. It is also conceivable to conduct a similar study on the trade-off for the two remaining benefits of customer data collection, viz. targeted advertising and market insight, as well as for additional properties of privacy-enhancing technologies, such as fewer customer defaults through better accountability (see for example [11]). Finally, research on economic aspects of privacy-enhancing technologies could also provide valuable feedback for the development of such technologies. For instance, cryptographic mechanisms, such as pseudonymous credentials, could be designed and implemented in a way that deliberately allows for certain price discrimination by authentically signalling information about the willingness to pay in well-defined attributes. This would ensure that no superfluous information is communicated, which is beneficial in terms of privacy, and at the same time reduce constraints for pricing strategies, which is beneficial for businesses and fair to consumers with low willingness to pay.

To sum it all up: the development of privacy-enhancing technologies is making considerable progress, but their ultimate success will depend on non-technical attributes, such as acceptance by users as well as economic rationality. If the technologies eventually break through, they will definitely change the shape of (electronic) commerce. The way this happens is a promising topic for interdisciplinary research.

References

[1] Alessandro Acquisti. Security of personal information and privacy: Technological solutions and economic incentives. In J. Camp and R. Lewis (eds), The Economics of Information Security, Kluwer, 2004. http://www.heinz.cmu.edu/~acquisti/papers/acquisti_eis_refs.pdf. 2

[2] Alessandro Acquisti, Roger Dingledine, and Paul Syverson. On the economics of anonymity. In R. N. Wright (Ed.): Financial Cryptography, LNCS 2742,84–102, 2003. 2

[3] Alessandro Acquisti and Jens Grossklags. Losses, gains, and hyperbolic discounting: An experimental approach to personal information security attitudes and behavior. In J. Camp and R. Lewis (eds), The Economics of Information Security, Kluwer, 2004. http://www.heinz.cmu.edu/~acquisti/papers/acquisti_grossklags_eis_refs.pdf. 3

[4] Alessandro Acquisti and Hal R. Varian. Conditioning prices on purchase history. *Marketing Science*, 24(3):1–15, 2005. http://www.heinz.cmu.edu/~acquisti/papers/privacy.pdf. 2

[5] Mark Armstrong. Recent developments in the economics of price discrimination. Forthcoming in: *Advances in Economics and Econometrics: Theory and Applications: Ninth World Congress*, eds. Blundell, Newey, Persson, Cambridge University Press, 2006. http://www.econ.ucl.ac.uk/downloads/armstrong/pd.pdf. 4

[6] Bettina Berendt, Oliver Günther, and Sarah Spiekermann. Privacy in e-commerce: Stated preferences vs. actual behavior. *Communications of the ACM*, 48(4):101–106, 2005. 3, 4, 18

[7] Rainer Böhme and Sven Koble. Pricing strategies in electronic marketplaces with privacy-enhancing technologies. *Wirtschaftsinformatik*, 49(1):16–25, 2007. 1

[8] Jan Bouckaert and Hans Degryse. Opt in versus opt out: A free-entry analysis of privacy policies. *Workshop on the Economics of Information Security (WEIS)*, Robinson College, University of Cambridge, UK, 2006. http://weis2006.econinfosec.org/docs/34.pdf. 3

[9] L. Jean Camp. *Trust & Risk in Internet Commerce*. MIT Press, Cambridge, MA, 2000. 2

[10] David Chaum. Security without identification: Transaction systems to make big brother obsolete. *Communications of the ACM*, 28:1030–1044, 1985. 2

[11] Ramnath K. Chellappa and Shivendu Shivendu. Incentive design for free but no free disposal services: The case of personalization under privacy concerns. *Workshop on the Economics of Information Security (WEIS)*, Carnegie Mellon University, Pittsburgh, PA, 2007. http://weis2007.econinfosec.org/papers/48.pdf. 3, 18

[12] Sebastian Clauß and Marit Köhntopp. Identity management and its support of multilateral security. *Computer Networks*, 37:204–219, 2001. 3

[13] Commission of the European Communities. Special Eurobarometer 196 - wave 60.0: Data protection. European Opinion Research Group, 2003. 3, 10, 18

[14] Electronic Privacy Information Center (EPIC). Public opinion on privacy. http://www.epic.org/privacy/survey/. 3

[15] Federal Trade Commission (FTC). Privacy online: Fair information practices in the electronic marketplace. http://www.ftc.gov/reports/privacy2000/privacy2000.pdf, 2000. 3

[16] Eric Friedman and Paul Resnick. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10(2):173–199, 2001. 3

[17] Thomas P. Gehrig and Rune Stenbacka. Price discrimination, competition and antitrust. The Pros and Cons of Price Disrimination, Konkurrensverket, Swedish Competition Authority, 2005. 18

[18] Jens Grossklags and Alessandro Acquisti. When 25 cents is too much: An experiment on willingness-to-sell and willingness-to-protect personal information. *Workshop on the Economics of Information Security (WEIS)*, Carnegie Mellon University, Pittsburgh, PA, 2007. http://weis2007.econinfosec.org/papers/66.pdf. 18

[19] Marit Hansen, Peter Berlich, Jan Camenisch, Sebastian Clauß, Andreas Pfitzmann, and Michael Waidner. Privacy-enhancing identity management. *Information Security Technical Report*, 9(1):35–44, 2004. 2

[20] Bernado A. Huberman, Eytan Adar, and Leslie R. Fine. Valuating privacy. *Workshop on the Economics of Information Security (WEIS)*, Kennedy School of Government, Harvard University, Cambridge, MA, 2005. http://infosecon.net/workshop/pdf/7.pdf. 18

[21] Junseok Hwang and Alexandre Repkine. The economic models of on-line digital identity, identity management systems and service interconnection policy, 2004. http://web.si.umich.edu/tprc/papers/2004/372/tprc2004-imsmodel-submission.pdf. 3

[22] Sven Koble and Rainer Böhme. Economics of identity management: A supply-side perspective. In G. Danezis and D. Martin (Eds.): PET 2005, LNCS 3856, 259–272, 2006. 6

[23] Ponnurangam Kumaraguru and Lorrie F. Cranor. Privacy indexes: A survey of Westin's studies. Tech. Report CMU-ISRI-5-138, Carnegie Mellon University, 2005. 3

[24] Robert Leathern. Online privacy — managing complexity to realize marketing benefits. Jupiter Research, 2002. http://www.jupiterresearch.com/bin/item.pl/research:vision/77/id=89141/. 3

[25] Andrew Odlyzko. Privacy, economics, and price discrimination on the internet. In N. Sadeh, editor, *ICEC2003: Fifth International Conference on Electronic Commerce*, pages 355–366, 2003. http://www.dtc.umn.edu/~odlyzko/doc/privacy.economics.pdf. 3

[26] Richard A. Posner. The economics of privacy. *American Economic Review*, 71(2):404–409, 1981. 2

[27] Joan Robinson. The economics of imperfect competition. Macmillan, London, 1933. 4

[28] Stefan Sackmann, Jens Strüker, and Rafael Accorsi. Personalization in privacy-aware highly dynamic systems. *Communications of the ACM*, 49(9):32–38, 2006. 3

[29] Carl Shapiro and Hal R. Varian. *Information Rules. A Strategic Guide to the Network Economy*. Harvard Business School Press, 1998. 18

[30] Paul Syverson. The paradoxical value of privacy. Workshop on Economics and Information Security (WEIS), University of Maryland, 2003. http://chacs.nrl.navy.mil/publications/CHACS/2003/2003syverson-privcost.pdf. 3

[31] Curtis R. Taylor. Private demands and demands for privacy: Dynamic pricing and the market for customer information. Duke Economics Working Paper 02-02, 2002. http://www.econ.duke.edu/Papers/Other/Taylor/private.pdf. 2

[32] Hal Varian. Intermediate microeconomics. W.W. Norton & Company, New York, 2003. 4

[33] Hal R. Varian, Fredrik Wallenberg, and Glenn Woroch. Who signed up for the do-not-call list? Workshop on Economics and Information Security, University of Minnesota, 2004. http://www.dtc.umn.edu/weis2004/varian.pdf. 4, 10, 18

[34] Luc Wathieu. Privacy, exposure, and price discrimination. *HBS Marketing Research Paper*, 02-03, 2003. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=347440. 2

[35] Detlev Zwick and Nikhilesh Dholakia. Models of privacy in the digital age: Implications for marketing and e-commerce. http://ritim.cba.uri.edu/Working%20Papers/Privacy-Models-Paper%5B1%5D.pdf, 1999. 3, 4

[36] Detlev Zwick and Nikhilesh Dholakia. Whose identity is it anyway? Consumer representation in the age of database marketing. *Journal of Macromarketing*, 24(1):31–43, 2004. 3