WI – Schwerpunktaufsatz

Pricing Strategies in Electronic Marketplaces with Privacy-Enhancing Technologies

The Authors

Rainer Böhme
Sven Koble

Rainer Böhme, M.A.
Dipl.-Inf. Sven Koble
Technische Universität Dresden
Institut für Systemarchitektur
01062 Dresden
{rainer.boehme | sven.koble}@inf.tu-dresden.de

The authors greatly appreciate the comments and feedback on earlier versions of this manuscript provided by Gaurav Kataria and Karen Kreutel.

Eingereicht am 2006-05-01,
nach zwei Überarbeitungen
angenommen am 2006-09-22 durch
die Herausgeber des Schwerpunktheftes.

1 Introduction

Collection and analysis of personal data are 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. There are three reasons for companies to allocate financial and human resources to data warehousing. 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, knowledge about customer 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 the customers' willingness to pay enables vendors with market power to impose pricing strategies that increase sales and revenues. For example, Acquisti and Varian [AcVa05] show how vendors with access to a technology that allows tracking customers over time can increase sales revenues by conditioning prices on the past behaviour of their customers. Such endeavours, however, stand in clear contrast to the individual's right and desire for privacy and informational self-determination. Conversely, customers' motivation to plead for privacy protection might partly be driven by an attempt to escape price discrimination and so to retain consumer surpluses [Wath03; AcVa05].

This article aims to shed light on the implications of privacy-enhancing technologies on pricing strategies from an economic perspective. The common understanding in early economic research of privacy was 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 [Posn81] concludes that by reducing the amount of information shared, privacy leads to market-inefficiencies and misallocation of resources. Recently, alternative views suggest that privacy has a certain value in itself, and therefore it is conceivable that privacy-respecting commerce is offered on the market if demanded [Acqu04; Tayl02]. Existing proposals for privacy-enhancing technologies should ease the realisation of privacy goals particularly in electronic commerce [HBCC04].

To capture the range of different aspects in a tractable economic model, we will concentrate on the interaction of one particular privacy-enhancing technology, namely *privacy-enhancing identity management* (PIM), with one of the aforementioned business reasons for personal data collec-

Executive Summary

This article approaches the field of conflict between personalisation and privacy protection with micro-economic models. It discusses the compatibility of pricing strategies with privacy-enhancing technologies as well as anticipated consequences for revenues and consumer prices.

- Privacy-enhancing technologies limit the possibility to implement price discrimination.
- A free-market solution with optional use of privacy-enhancing technologies yields higher social welfare than mandatory use through government regulation.
- Privacy will likely remain a "luxury good".

Keywords Privacy-Enhancing Technologies, Identity Management, Price Discrimination, E-Commerce, Economics of Privacy

tion, namely pricing strategies. A privacy-enhancing identity management system supports its users in managing and protecting their personal information. Following the idea of Chaum [Chau85], it consists of mechanisms that allow pseudonymous interactions among business counterparts. By changing pseudonyms deliberately, users retain full control over which information can be combined from previous interactions. What is more, additional functions support users in keeping track of disclosed personal information and assure accountability, if desired, by means of cryptographic protocols.

Social implications of pseudonymous transactions have been subject to prior research. Friedman and Resnick [FrRe01] 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 [ZwDh99] compare free-market and government regulation approaches to deal with 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 [BoDe06] address the implications of different regulatory approaches. They conclude that an *opt-out* policy, where customer 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 customer information) or complete prohibition of customer data processing. Their results, however, are not directly comparable to our analyses due to a number of rigid model assumptions. Most importantly, Bouckaert and Degryse do not allow for heterogeneous privacy preferences in the population, which is a core attribute of our models presented below.

This article is structured as follows. Section 2 explains the basic relationship between the amount of information disclosed in a business transaction and the possibility to implement price discrimination. As PIM technology limits the flow of information, vendors have to compensate the lost revenue and sell to new market segments. We regard the basic economic trade-off for self-regulated (optional) PIM technology as well as for a scenario in which all customers use PIM by default. Section 3 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 [AcDS03]. Unlike prior work, which addresses this issue mainly with technical means such as calling for user-friendly systems [ClKö03; HwRe04], we believe that support from vendors and service providers is equally crucial. Therefore, we will study the vendors' incentives to support PIM technology. Further, in section 4, we discuss the optimal price setting strategies for vendors and their consequences for individual consumers. Section 5 concludes with a discussion of possible implications and directions for future research.

2 Privacy and Price Discrimination: A Baseline Model

Odlyzko [Odly03], among others, has identified *price discrimination* as one of the main motivations for companies to collect personal information about their customers. Price discrimination occurs if vendors charge customers different prices for the same product depending on the individual customer's willingness to pay [see for example Vari03]. Sometimes referred to as *differential pricing*, this phenomenon has a long research tradition in micro-economics [Robi33]. More recent work also deals with particularities of pricing strategies in electronic commerce [see Arms06 for a survey]. In order to enforce price discrimination, the vendor has to know (or infer) individual customers' willingness to pay, which can be done on the basis of collected personal data. This creates the link between consumer privacy and pricing strategies.

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 vendor 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 $r = \frac{p}{4} \cdot Q_T$. This corresponds to the rectangular area in Fig. 1, following the theory of monopolistic pricing [see for example Vari03; 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 vendors can determine each individual's willingness to pay, then they can implement perfect price discrimination and achieve revenue

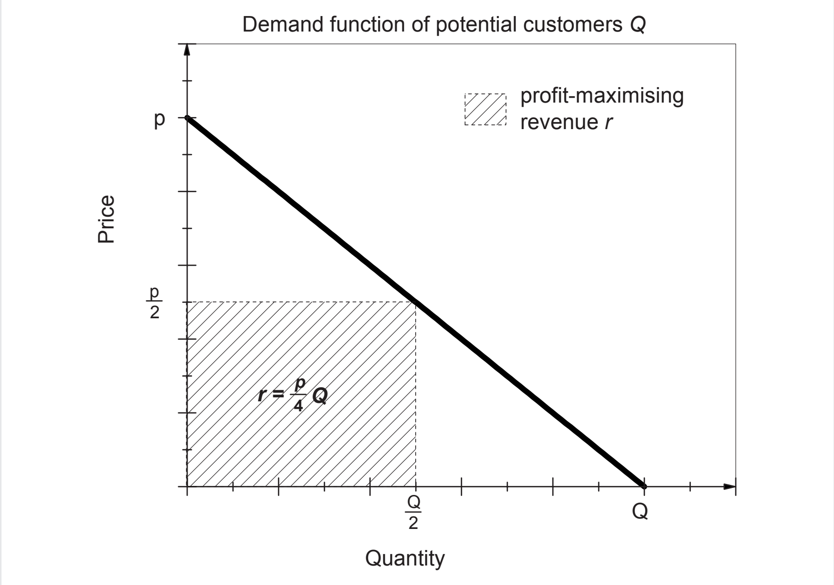

Fig. 1 Demand function of *Q* customers with reservation price *p;* no marginal costs. Price $\frac{p}{2}$ maximises the vendor's revenue

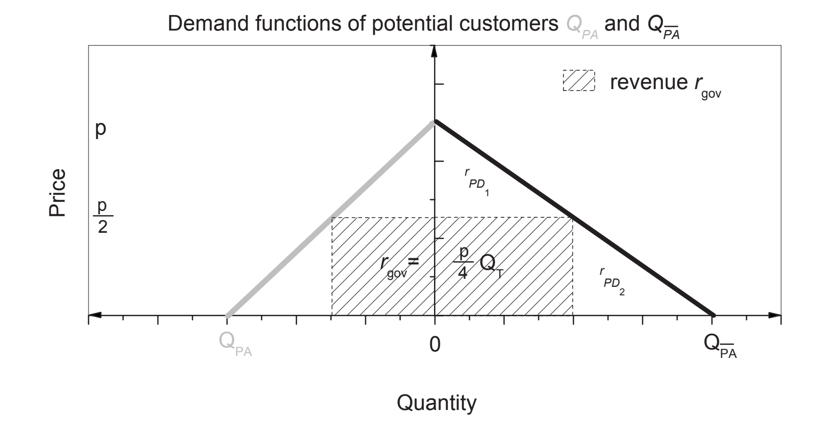

Fig. 2 Demand function and profit-maximising revenue if the introduction of PIM is mandatory (e.g. as a result of government regulation)

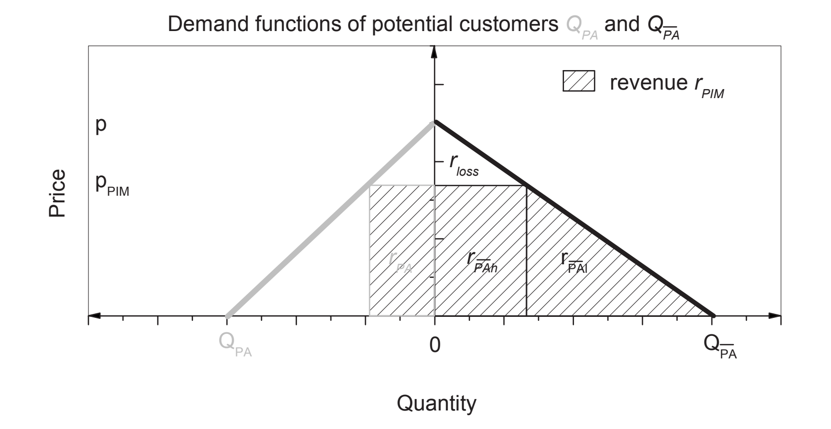

Fig. 3 Demand function of potential customers Q_{PA} and $Q_{\overline{PA}}$ · r_{loss} is the revenue lost to "strategic customers" if the vendor decides to introduce PIM; however additional revenue r_{PA} will compensate for this loss

$r = \frac{p}{2} \cdot Q_T$, twice as much as before (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 vendor and customers' inability to resell to other customers (absence of arbitrage).

Generalising the model, we assume two market segments of customers with different attitudes towards privacy: customers that are not at all concerned about their privacy (suffix $_{\overline{PA}}$) and customers with notable *privacy awareness* (suffix $_{PA}$). We further assume that the latter group will only participate in a transaction if PIM is supported to protect their personal data. This distinction of heterogeneous privacy preferences can be justified against the backdrop of theoretical considerations (privacy needs differ between individuals [ZwDh99]) as well as of empirical findings [BeGS05; VaWW04].

Fig. 2 displays the demand functions for both market segments, in which customers without privacy awareness $Q_{\overline{PA}}$ are depicted on the right-hand side, and customers with high privacy awareness Q_{PA} on the left-hand side (note the inverted quantity scale). We define ratio λ in the domain [0,1] as the fraction of customers *without* privacy awareness:

$$\lambda = \frac{Q_{\overline{PA}}}{Q_{\overline{PA}} + Q_{PA}} = \frac{Q_{\overline{PA}}}{Q_T} . \tag{1}$$

When no PIM is available, privacy-aware customers will not purchase from the vendor. Consequently, with only $Q_{\overline{PA}}$ customers left, the vendor is able to implement perfect price discrimination and achieves revenue of

$$r_{PD} = \frac{p}{2} \cdot \lambda \cdot Q_T . \tag{2}$$

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 customers are indistinguishable and the vendor has to set one single price for the entire market. Then, as depicted in Fig. 2, the profit-maximising revenue changes to

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

This corresponds to the rectangular area spanning both market segments as some privacy-aware customers are willing to purchase now. Vendors profit from this situation when the additional revenue in the market segment with high privacy awareness (r_{PA}) exceeds the lost revenue 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 customers constitute a majority in the population.

Welfare analysis adds up both consolidated supplier and consumer surplus to assess the overall effect of policy choices on 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 would not outweigh the losses in vendor surplus caused by the inability to differentiate prices. This scenario is similar to the one described in an earlier workshop version of this research [KoBö06].

In a self-regulated approach, PIM technology is available, but its use is not mandatory. Vendors support the technology, and each customer can decide whether to use it or not and if all the requested personal information should be revealed. Vendors still implement price discrimination with those customers of which they can obtain personal data whereas they set one single price p_{PIM} for all customers that use PIM. This implies that customers can act strategically: those without privacy awareness will choose PIM not for privacy reasons but to extract surplus if the price for PIM users is below the individual customer's willingness to pay.

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

$$r_{PIM} = \left[\frac{p}{4 - 2 \cdot \lambda}\right] \cdot Q_T \,. \tag{4}$$

In this scenario, all customers without privacy awareness and some additional privacy-aware customers do purchase from the vendor. Note that some revenue in the right-hand market segment (upper triangle) is lost due to strategic customers. However, this loss is over-compensated in all cases by the additional revenue from customers with high privacy awareness. Only if no customers 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 (and strictly positive if at least one customer is privacy-aware). This is another indication supporting the view that a self-regulated market solution is superior to government regulation.

3 Will Privacy-Enhancing Identity Management Thrive?

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.

3.1 Optional PIM with binary market segmentation

The validity of the baseline model is limited by the assumption of perfect price discrimination. Vendors often do not know exactly each customer's individual willingness to pay. This motivates an extension of the model to those cases in which the ability to price discriminate is constrained: vendors can only infer one bit of information about each customer's willingness to pay. This means vendors can tell for each customer whether his or her willingness to pay is above (suffix $_h$ for *high*) or below (suffix $_l$ for *low*) a certain limit price p_{sep} (suffix $_{sep}$ for *separation*). The limit price is given exogenously, i.e. a single vendor 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 vendors know the demand function, they can calculate the number of customers with high Q_h and low Q_l willingness to pay, respectively. For the given demand function, p_{sep} is directly related to the fraction of customers with high willingness to pay π.

$$\pi = \frac{Q_h}{Q_T} = 1 - \frac{p_{sep}}{p} \,. \tag{5}$$

With *willingness to pay* and *privacy awareness* being two orthogonal dimensions, we have defined a model that divides customers into four groups. Given the parameters π, λ, and p (the reservation price of the demand function), vendors aim to maximise their revenue. In the absence of PIM, they do so by setting two prices, p_l and p_h, for customers 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. To extract the maximum revenue from the market segment with low willingness to pay, vendors apply the standard monopolistic pricing for the section of the demand function below p_{sep}, hence

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

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 customers with high willingness to pay are in the majority ($\pi \geq \frac{1}{2}$), then vendors use monopolistic price setting as if it were for the entire market since all customers 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. 4 for illustration). In the opposite case ($\pi < \frac{1}{2}$), the optimal decision is to set $p_h = p_{sep}$, which is the closest possible so-

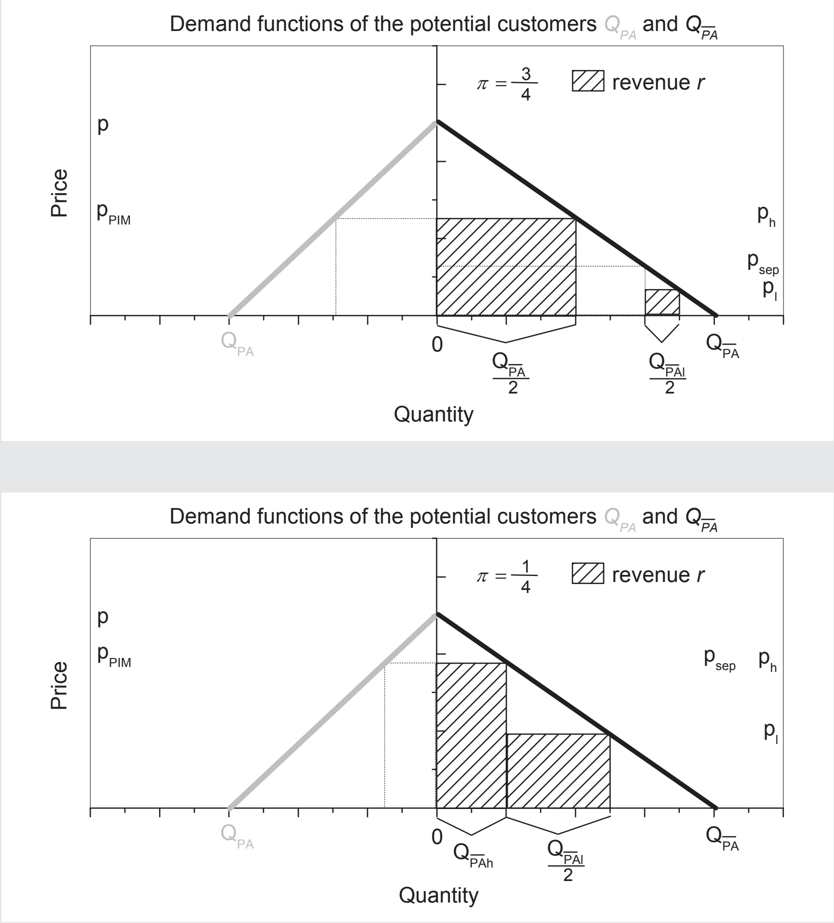

Fig. 4 Demand function with parameters p_{sep} and π: profit-maximising revenue without PIM when price discrimination is imposed to customers without privacy awareness. Satisfied demand differs for $\pi \geq \frac{1}{2}$ (top) and $\pi < \frac{1}{2}$ (bottom)

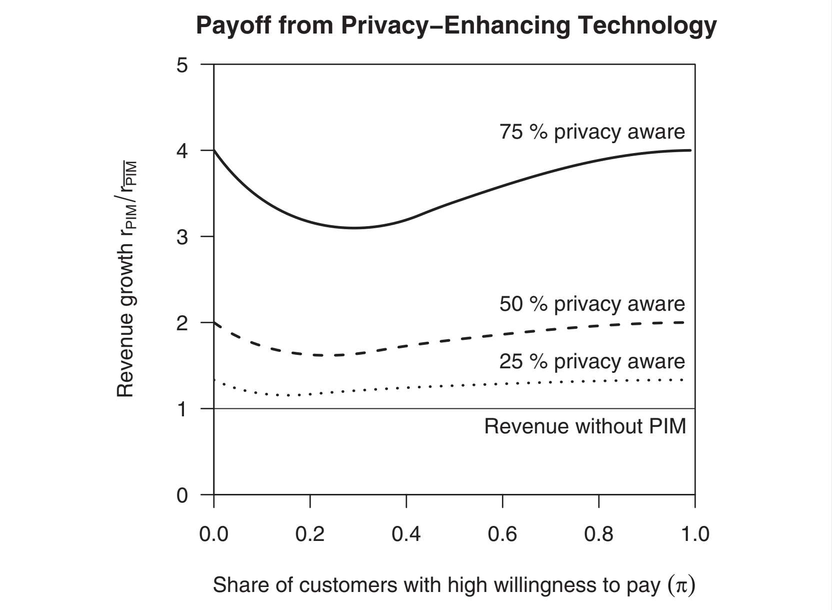

Fig. 5 Potential increase in revenue after the introduction of PIM. A value of 1 corresponds to the revenue in a monopolistic price discrimination scenario without support for privacy-enhancing technologies. The graphs show different assumptions for the fraction of privacy-aware customers $(1-\lambda)$

lution to the unique monopoly price (Fig. 4, 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:

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

where

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

If a vendor decides to support PIM, he or she has 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 customers. It turns out that the optimal setting of all three prices (p_l, p_h, p_{PIM}) does not affect the choice of $p_l \cdot p_{PIM}$ is set to the same level of p_h as follows:

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

Now we will compare the optimal revenues with PIM being 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:

$$r_{PIM} = \frac{1}{4}\cdot p\cdot(\lambda\cdot(1-\pi)^2+1)\cdot Q_T\,. \tag{11}$$

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

$$r_{\overline{PIM}} = \frac{1}{4}\cdot p\cdot(\lambda\cdot(1-\pi)^2+\lambda)\cdot Q_T\,. \tag{12}$$

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

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

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

$$r_{PIM} = \frac{1}{4}\cdot p\cdot[\lambda\cdot(1-\pi)^2+4\cdot\pi\cdot(1-\pi)] \times Q_T\,. \tag{13}$$

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

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

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

$$r_{\overline{PIM}} = \frac{1}{4}\cdot p\cdot[\lambda\cdot(1-\pi)^2+4\cdot\pi\cdot(1-\pi)\cdot\lambda] \times Q_T\,. \tag{15}$$

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

$$r_e = \frac{1}{4}\cdot p \times\left[\underbrace{4\cdot\lambda\cdot\pi\cdot(1-\pi)}_{\geq 0}+\underbrace{\frac{\lambda\cdot\pi^2}{1-\lambda}}_{\geq 0}+\underbrace{1-\lambda\cdot(1-\pi)^2}_{>0 \text{ for } \lambda<1}\right] \times Q_T > 0\,. \tag{16}$$

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 prospect customers are privacy-aware $(\lambda < 1)$. The factor by which the revenue increases varies with the number of customers 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. 5, the gains are comparatively lower when customers with low willingness to pay constitute 60–80 % of the market. In these situations, price discrimination is most effective, and vendors cannot sell to privacy-aware customers with low willingness to pay because reducing p_{PIM} further would sacrifice the high margins from affluent customers that would start using PIM for strategic reasons.

3.2 Relaxing the independence assumption

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. [VaWW04] report this fact based on an analysis of do-not-call lists in the U.S. The positive correlation may be explained by factors such as affluent individuals being

Tab. 1 Market segmentation with correlation: size of the customer groups depending on parameters Q_T, λ, π and ϱ

Willingness to pay		Privacy awareness: high	Privacy awareness: low
	high	$Q_{PA,h} = Q_T \cdot ((1-\lambda) \cdot \pi + \varrho \cdot \sqrt{\lambda(1-\lambda)\pi(1-\pi)})$	$Q_{\overline{PA},h} = Q_T \cdot (\lambda \cdot \pi - \varrho \cdot \sqrt{\lambda(1-\lambda)\pi(1-\pi)})$
	low	$Q_{PA,l} = Q_T \cdot ((1-\lambda) \cdot (1-\pi) - \varrho \cdot \sqrt{\lambda(1-\lambda)\pi(1-\pi)})$	$Q_{\overline{PA},l} = Q_T \cdot (\lambda \cdot (1-\pi) + \varrho \cdot \sqrt{\lambda(1-\lambda)\pi(1-\pi)})$

targeted by direct marketers more frequently and thus decide to subscribe to a do-not-call list, 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* [Comm03]. For example, in 2003, 13 % of the managers interviewed as opposed to 3 % of the house persons, 4 % of the manual workers and only 2 % of the retired reported to use privacy-enhancing technologies – including encryption tools.

We use a measure of dependence between two variables based on Pearson's χ^2 statistic, which is the sum of squared difference between the actual size of the segments and the expected size if customer attributes were independent. The domain of the correlation coefficient is $\varrho \in [-1, 1]$, where values $\varrho < 0$ denote that privacy awareness on average concurs with low willingness to pay, whereas $\varrho > 0$ indicate that privacy-aware people are more likely to have high willingness to pay. Tab. 1 shows how the sizes of the four different customer groups $Q_{\overline{PA}_l}$, $Q_{\overline{PA}_h}$, Q_{PAl} and Q_{PAh} can be calculated from the exogenous parameters Q_T, π, λ and ϱ. The set of non-negativity constraints for the market segments ($Q_{\overline{PA}_l} \geq 0 \wedge Q_{\overline{PA}_h} \geq 0 \wedge Q_{PAl} \geq 0 \wedge Q_{PAh} \geq 0$) limits the domain of reasonable combinations for parameters (λ, π, ϱ) as follows:

$$\frac{\varrho^2 \cdot f(\lambda)}{\varrho^2 \cdot f(\lambda) + (1 - f(\lambda))} \leq \pi \leq \frac{f(\lambda)}{\varrho^2 \cdot (1 - f(\lambda)) + f(\lambda)} , \qquad (17)$$

where

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

For $\varrho = 0$, condition (17) is true for any combination $(\lambda, \pi) \in [0, 1]^2$. This is consistent with our expectations as the model of the previous section is a special case of this more general model. The impact of the parameter constraints for $\varrho \neq 0$ can be seen in Fig. 8 and 9 below.

The correlation parameter also affects the shape of the demand function, which is still assumed to be linear in both market segments. While in the baseline model, both demand curves intersect at the same reservation price p, now customers with high and low privacy awareness may exhi-

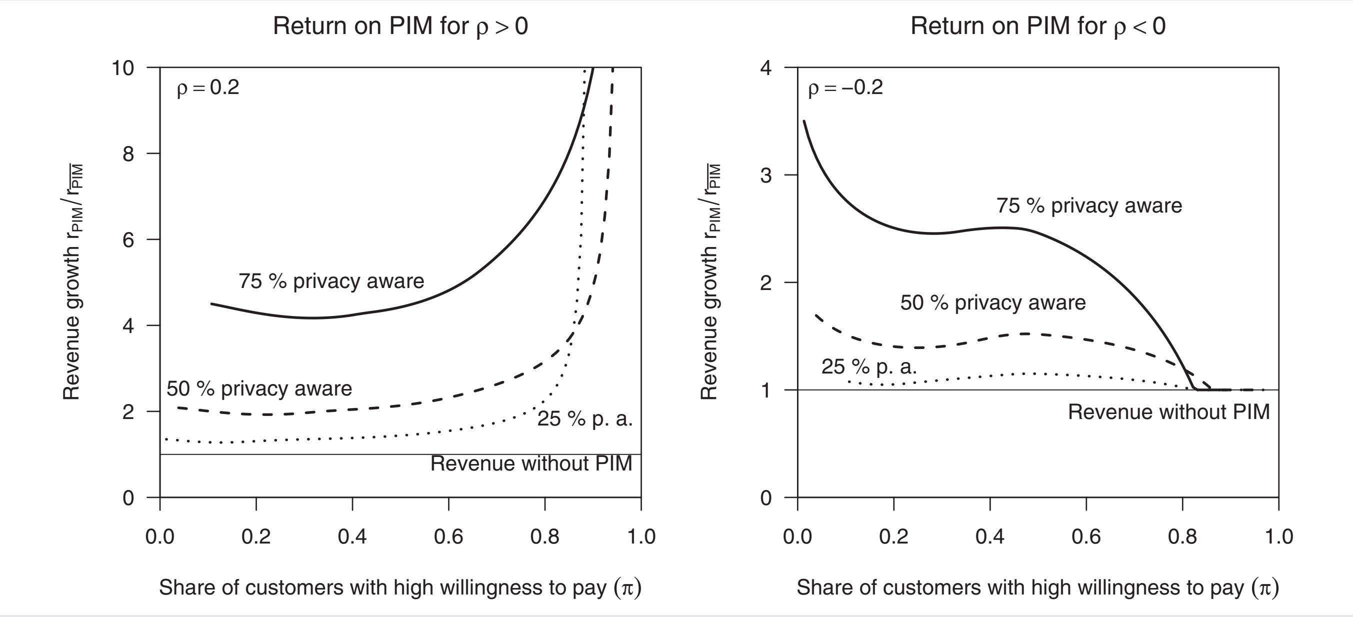

Fig. 6 Potential increase in revenue after the introduction of PIM. A value of 1 corresponds to the revenue in a monopolistic price discrimination scenario without support for privacy-enhancing technologies. The graphs show different assumptions for correlation (left: $\varrho > 0$, right: $\varrho < 0$) and the fraction of privacy-aware customers $(1 - \lambda)$. Compare with Fig. 5 for the independent case

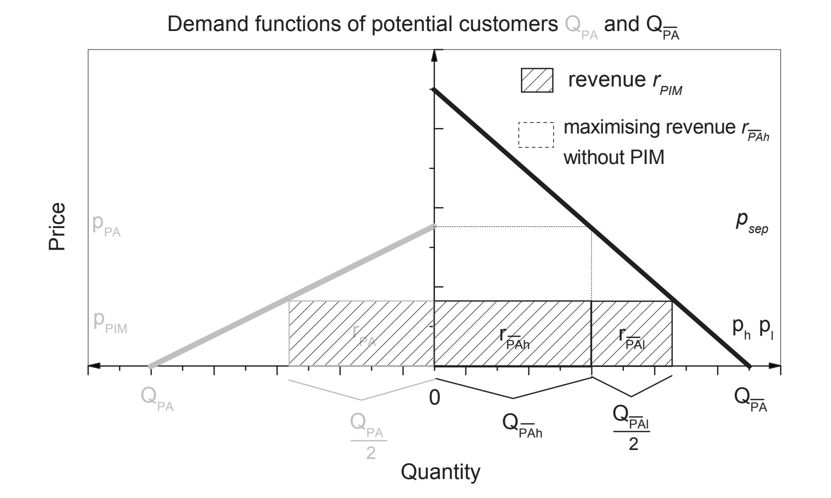

Fig. 7 Demand function of potential customers Q_{PA} and $Q_{\overline{PA}}$ with $\varrho < 0$. Without supporting PIM, the vendor would extract higher revenue from customers without privacy awareness and high willingness to pay (opaque rectangle)

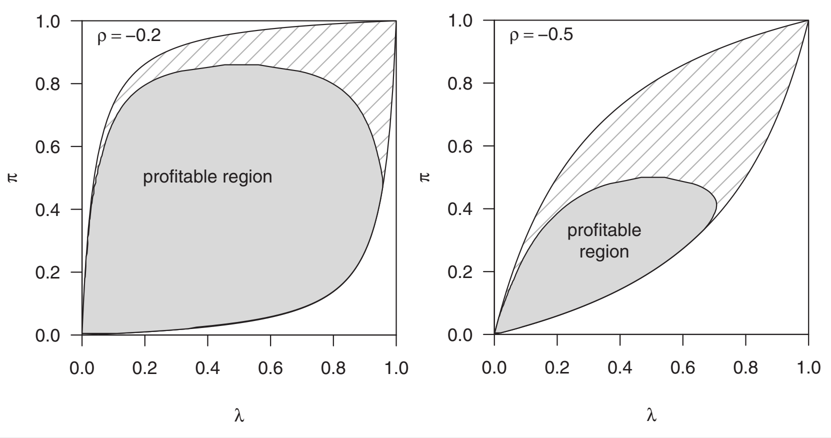

Fig. 8 Combinations of λ and π for which the introduction of PIM is profitable for the vendor even though willingness to pay and privacy awareness are negatively related ($\varrho < 0$). Hatched regions show the domain of (λ, π) for given ϱ (cf. equation (17))

bit different reservation prices. We write the reservation price for privacy-aware customers as p_{PA} as opposed to p.

To assess the profitability of PIM, we discuss the cases $\varrho > 0$ and $\varrho < 0$ separately. For positive correlation ($\varrho > 0$), 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 $\varrho = 0$. As the reservation price for privacy-aware customers p_{PA} is greater than p, a vendor could always act as in the independent case by assuming $p = p_{PA}$. This means that the introduction of PIM is already worthwhile despite leaving some extra consumer surplus to privacy-aware customers. This is sufficient to back the proposition. Smart vendors 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 of Fig. 6, 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.

For $\varrho < 0$, however, the situation becomes much more complicated. For the first time, non-trivial cases exist, in which the introduction of PIM is *not* profitable, as can be seen in the right chart of Fig. 6. This happens when the reservation price for privacy-aware customers p_{PA} falls too far below p_h. Selling to the privacy-aware segment would then sacrifice large parts of the revenue from customers with high willingness to pay and low privacy awareness so that vendors are better off if they do not support PIM at all. The situation occurs for high π, but the exact threshold also depends on λ. Fig. 7 shows the shape of the demand function in such situations while Fig. 8 visualises those regions of combinations (λ, π) in which the introduction of PIM technology 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 to pay might form a serious market entry barrier for PIM technology, which intensifies even more if transaction costs are taken into account.

4 Privacy at a Premium?

The model in the previous section has been set up in order to determine the conditions under which rational vendors will decide to support privacy-enhancing technologies despite losing the opportunity to pursue price discrimination (at least in parts 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 now focus on the prices customers with and without privacy awareness alike will have to pay. Hence, after regarding the supply-side in the previous sections, we are now switching to a consumer 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 $p_{h,PD}$ and $p_{l,PD}$ if no PIM is supported (indicated by suffix $_{PD}$). 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

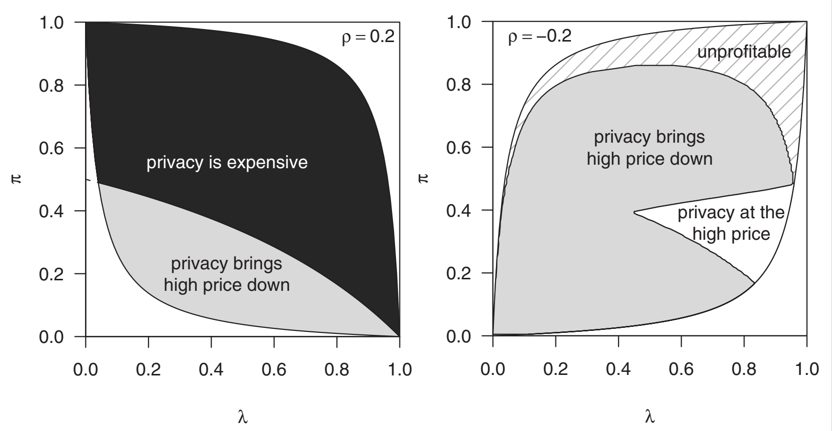

Fig. 9 Prices p_{PIM} charged from privacy-aware customers in comparison to the high prices in the default situation without PIM. For positive correlation (left), price discrimination by privacy preferences can be implemented and p_{PIM} climbs above p_h. Negative correlation (right) inflates the regions where a vendor should reduce the high price in order to sell more to privacy-aware customers

rows of (10) exactly correspond to (7), therefore p_{PIM} differs from $p_{h,PD}$ 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 $p_{h,PD}$ when no PIM is supported. In other words, when customers with low willingness to pay are in the majority and the fraction of privacy-aware customers is above a certain threshold, the introduction of PIM not only increases the revenues for the vendor but also slightly reduces the price for customers with low privacy awareness and high willingness to pay (but will not alter the price for customers with low willingness to pay).

In contrast, when customers with high willingness to pay are in the majority – think of a distinguishing criterion like student cards, and students are a minority in the population –, the optimal $p_{PIM} = \frac{1}{2} \cdot p$, which is greater than $p_{sep} = p \cdot (1 - \pi)$. 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 assumption 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 may require 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 $\varrho \neq 0$, but for the sake of brevity we omit the analytical deviations and proof ideas. If privacy awareness and willingness to pay are *positively* correlated ($\varrho > 0$), then p_{PIM} will be below $p_{h,PD}$ under exactly the same conditions as in the independent model, with the exception that the domain of (λ, π) is reduced as shown in equation (17). However, when π exceeds the threshold $\frac{\lambda-1}{\lambda-2}$, the vendor can rise p_{PIM} above p_h because the reservation price of privacy-aware customers 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-enhancing technology. It is important to note that privacy-aware customers 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 that it becomes even more expensive if privacy-aware customers are known to be more affluent on average ($\varrho > 0$). The left chart of Fig. 9 shows the regions for either case.

The situation becomes more difficult analytically if privacy awareness and willingness to pay are *negatively* correlated ($\varrho < 0$). After accounting for the unprofitable combinations (λ, π), we see that the relative proportion of situations in which the introduction of PIM is accompanied by lower prices $p_h = p_{PIM}$ increases (see Fig. 9, right chart). It is also noteworthy that p_{PIM} can drop to p_l in the marginal case where no privacy-aware customers have a high willingness to pay (but certainly some customers with low privacy awareness do have!). In such cases, the demand for privacy is strong enough to force the vendor to set one single price $p_{PIM} = p_h = p_l$ in all market segments. However, the condition that there must not be one single privacy-aware customers with high willingness to pay shows how unlikely such cheap privacy actually is. Even when the linearity constraint of the demand function is 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-enhancing identity management systems will be charged an additional premium by vendors who otherwise would be able to price discriminate. It is possible that this "privacy tax", coupled with acceptance problems of different nature, could hinder a wide deployment of such technologies by large parts of the population.

5 Summary and Conclusion

As this article has tried to show, new developments in the area of privacy-enhancing technologies in combination with persistence of privacy concerns in the population have tremendous implications for business-to-consumer relations. More precisely, vendors 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, vendors 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 vendors 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 customers that value privacy very much would have to exceed 50 %. In a more realistic scenario where 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 customers. 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 limitations to employ price discrimination.

Assuming that optional support for privacy-enhancing technologies is commonplace, vendors 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 vendors will opt to define a specific price for privacy-aware users. Our analyses show that this price will most likely be higher than the lowest price for customers 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 [Comm03; VaWW04], 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 may be tempted to demand 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 suboptimal 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 price level for consumers [GeSt05].

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) [ShVa98]. But it does not cover all possible market structures 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 [AcGr04; BeGS05]. It is also conceivable to conduct a similar tradeoff 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. 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.

References

[Acqu04] *Acquisti, Alessandro:* Security of Personal Information and Privacy: Technological Solutions and Economic Incentives. In: *Camp, J.; Lewis, R.* (eds.): The Economics of Information Security. Kluwer, 2004.

[AcDS03] *Acquisti, Alessandro; Dingledine, Roger; Syverson, Paul:* On the Economics of Anonymity. In: *Wright, R. N.* (ed.): Financial Cryptography. LNCS 2742, 2003, pp. 84–102.

[AcGr04] *Acquisti, Alessandro; Grossklage, Jens:* Losses, Gains, and Hyperbolic Discounting: An Experimental Approach to Personal Information Security Attitudes and Behavior. In: *Camp, J.; Lewis, R.* (eds.): The Economics of Information Security. Kluwer, 2004.

[AcVa05] *Acquisti, Alessandro; Varian, Hal:* Conditioning Prices on Purchase History. In: Marketing Science 24 (2005) 3, pp. 1–15.

[Arms06] *Armstrong, Mark:* Recent Developments in the Economics of Price Discrimination. De-

Abstract

Pricing Strategies in Electronic Marketplaces with Privacy-Enhancing Technologies

Collecting customer information in electronic commerce and respecting consumers' privacy preferences are fundamentally competing goals. This article studies the effects of emerging user-controlled privacy-enhancing technologies on pricing strategies pursued by vendors. In particular, identity management systems that allow users to interact pseudonymously with businesses thwart the vendors' efforts to set different prices based on user attributes or purchase histories. Applying micro-economic models, we will compare different possible regimes for the implementation of privacy-enhancing technologies, analyse the conditions under which it is profitable for vendors to support privacy-enhancing identity management systems and study respective welfare implications. Accordingly, we will address the basic questions of whether and how such technologies will become ready to be brought to the market.

Keywords: Privacy-Enhancing Technologies, Identity Management, Price Discrimination, E-Commerce, Economics of Privacy

partment of Economics, University College London, 2006 (forthcoming in: *Blundell; Newey; Persson* (eds.): Advances in Economics and Econometrics: Theory and Applications. Cambridge University Press).

[BeGS05] *Berendt, Bettina; Günther, Oliver; Spiekermann, Sarah:* Privacy in E-Commerce: Stated Preferences vs. Actual Behavior. In: Communications of the ACM 48 (2005) 4, pp. 101–106.

[BoDe06] *Bouckaert, Jan; Hans Degryse, Hans:* Opt In versus Opt Out: A Free-entry Analysis of Privacy Policies. In: Workshop on the Economics of Information Security (WEIS), University of Cambridge, UK, 2006. http://weis2006.econinfosec.org/docs/34.pdf.

[Chau85] *Chaum, David:* Security Without Identification: Transaction Systems to Make Big Brother Obsolete. In: Communications of the ACM 28 (1985), pp. 1030–1044.

[ClKö01] *Clauß, Sebastian; Köhntopp, Marit:* Identity Management and Its Support of Multilateral Security. In: Computer Networks 37 (2001), pp. 205–219.

[Comm03] Commission of the European Communities: Special Eurobarometer 196 – Wave 60.0: Data Protection. European Opinion Research Group, 2003.

[FrRe01] *Friedman, Eric; Resnick, Paul:* The Social Cost of Cheap Pseudonyms. In: Journal of Economics and Management Strategy 10 (2001), pp. 173–199.

[GeSt05] *Gehrig, Thomas P.; Stenbacka, Rune:* Price Discrimination, Competition and Antitrust. The Pros and Cons of Price Discrimination Konkurrensverket, Swedish Competition Authority, 2005.

[HBCC04] *Hansen, Marit; Berlich, Peter; Camenisch, Jan; Clauß, Sebastian; Pfitzmann, Andreas; Waidner, Michael:* Privacy-Enhancing Identity Management. In: Information Security Technical Report 9 (2004), pp. 35–44.

[HwRe04] *Hwang, Junseok; Repkine, Alexandre:* The Economic Models of Online Digital Identity, Identity Management Systems and Service Interconnection Policy, 2004. http://web.si.umich.edu/tprc/papers/2004/372/tprc2004-imsmodel-submission.pdf.

[KoBö06] *Koble, Sven; Böhme, Rainer:* Economics of Identity Management: A Supply-side Perspective. In: *Danezis, G; Martin, D.* (eds.): Privacy Enhancing Technologies (PET 2005). LNCS 3856, 2006, pp. 259–272.

[Odly03] *Odlyzko, Andrew:* Privacy, Economics, and Price Discrimination on the Internet. In: Fifth International Conference on Electronic Commerce. ACM, 2003, pp. 355–366.

[Posn81] *Posner, Richard A.:* The Economics of Privacy. In: American Economic Review 71 (1982) 2, pp. 404–409.

[Robi33] *Robinson, Joan:* The Economics of Imperfect Competition. Macmillan, London 1933.

[ShVa98] *Shapiro, Carl; Varian, Hal:* Information Rules: A Strategic Guide to the Network Economy. Harvard Business School Press 1998.

[Tayl02] *Taylor, Curtis R.:* 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.

[Vari03] *Varian, Hal:* Intermediate Microeconomics. W. W. Norton & Company, New York 2003.

[VaWW04] *Varian, Hal; Wallenberg, Fredrik; Woroch, Glenn:* Who Signed Up for the Do-Not-Call List? In: Workshop on Economics and Information Security (WEIS), University of Minnesota, 2004. http://www.dtc.umn.edu/ weis2004/varian.pdf

[Wath03] *Wathieu, Luc:* Privacy, Exposure, and Price Discrimination. HBS Marketing Research Paper, 02-03, 2003. http://papers.ssrn.com/sol3/papers.cfm?abstract id=347440.

[ZwDh99] *Zwick, Detlev; Dholakia, Nikhilesh:* Models of Privacy in the Digital Age: Implications for Marketing and E-Commerce (1999). http://ritim.cba.uri.edu/Working%20Papers/Privacy-Models-Paper%5B1%5D.pdf.