A Franc Less for a Pound More: (Price) Discrimination and the Value of Privacy

Irina Gemmo* Wanda Mimra† Anastasia Sycheva‡

January 9, 2020

VERY PRELIMINARY VERSION - PLEASE DO NOT CIRCULATE

Abstract

Price discrimination based on consumers' personal data has become common practice in many markets. We analyze the willingness to share personal data when this data is used for price discrimination in subsequent markets. In a laboratory experiment, participants can sell a bundle of personal data. Participants are categorized based on the content of their personal data and receive category-dependent payoffs in a subsequent stage. The experimental variations modify the category-dependent payoff structure. We find no effect of subsequent financial discrimination on the general willingness to sell personal data. A significant change in the data reservation price is only observed under strong negative discrimination. Furthermore, we observe important gender differences in the reservation price for private information and the role of underlying privacy concerns.

JEL classification: C81, D12, D82, D83, D91
Keywords: personal data, price discrimination, privacy

*Irina Gemmo, Department of Management, Technology and Economics, ETH Zurich, Zürichbergstrasse 18, 8092 Zurich, Switzerland. Phone: +41 44 633 80 36; e-mail: igemmo@ethz.ch.

†Wanda Mimra, IÉSEG School of Management, LEM (UMR-CNRS 9221), 1 Parvis de La Défense, 92044, Paris, France. e-mail: w.mimra@ieseg.fr.

‡Anastasia Sycheva, Department of Management, Technology and Economics, ETH Zurich, Zürichbergstrasse 18, 8092 Zurich, Switzerland. Phone: +41 44 633 28 75; e-mail: asycheva@ethz.ch.

1 Introduction

Using customer information to boost sales is an old and time-tested strategy. Private doctors traveling between cities in Ancient Greece charged more to the rich than to the poor. With easier access to private information, price discrimination based on consumers' detailed personal data as well as individual behavior is increasingly becoming common practice in many markets.[1] Airlines and other companies, such as Home Depot, offer individual prices for different customers on their websites based on factors such as the time and day of the online activity as well as the customers' zip codes.[2] While price discrimination can be beneficial for consumers, these practices are often intransparent to consumers and violate their privacy.

The use of Big Data enables firms to pursue these practices with greater precision and therefore may lead to negative outcomes for consumers. Ezrachi and Stucke (2016) argue that online behavioral discrimination through big data will likely differ from the brick-and-mortar type of price discrimination in three ways: (1) a shift from third-degree, imperfect price discrimination to near-perfect or first-degree price discrimination by segmenting consumers into smaller groups and identifying their reservation prices more precisely (2) an increase in overall consumption through marketing strategies that target consumers' emotions more effectively and (3) a stronger durability of discrimination stemming from personalization and data-driven network effects. The authors point out that the increase in personalized product offers and individual pricing makes it harder for consumers to evaluate all options and assess general market prices.

A very controversial issue is risk-based pricing in insurance. In various lines of insurance business, such as health insurance, life insurance and automotive insurance, policies are priced based on policyholders' personal data that is used to predict their risk type. Risk-based pricing can incentivize policyholders to behave less risky and create desirable spill-over effects, such as an improvement in road safety and better general health. Risk-based pricing can become problematic though when potential insurance buyers are priced out of the market based on factors beyond their control, such as genetic conditions. A pricing scheme that makes health insurance prohibitively expensive for an individual with

[1]Price discrimination can be beneficial for consumers and retailers. The seller can often generate more revenue by offering services and products at lower costs to groups that tend to be more price sensitive, for example in the case of student and senior discounts.

[2]Access to consumer information may intensify competition. E.g., Choe et al. (2019) show that the ability to gather information and use personalized prices reduces firms' total profits in a model of asymmetric collection of personal information of consumers (e.g. via cookies).

such a condition seems unfair and possibly illegal. Genetic tests may even allow insurers to uncover medical conditions that the affected customer might not yet know about.

The willingness to share personal information that influences the demand for insurance products with such a pricing scheme differs across individuals. In a data set for pay-as-you-drive contracts, Muermann et al. (2019) show that such a car insurance policy that involves information sharing is more likely to be chosen by younger, female consumers who live in urban and/or wealthier areas.

Besides the potential direct economic consequences, privacy concerns play an important role for consumers' decision to purchase such products. The value of privacy has been subject to a public debate that has become increasingly relevant with public scandals, such as the Facebook–Cambridge Analytica data scandal.[3]

In this article, we make use of a laboratory experiment to elicit individuals' willingness to share personal data when this data is subsequently used to price discriminate. Thus, we analyze participants' privacy concerns as well as their response to payoff discrimination based on the content of their personal data. The personal data that the participants can sell in the experiment consists of the bundle of their height, weight, bank account balance information as well as a photo of their face. To implement price discrimination in the lab, participants are then categorized based on whether they sold their data to the experimenters, and importantly based on the content of their data, whereby the category-cutoffs depend in particular on a person's weight and bank account balance. These categories entail different payoffs in a subsequent stage, thus implementing data-based price discrimination in a reduced form.

Comparing treatments with and without data-dependent payoff differences, we find no effect of price discrimination per se on the general willingness to sell personal data. Thus, within the experiment, we don't find evidence of a disutility of financial discrimination attached to the content of personal data per se. With respect to the price of personal data demanded by participants, we find a significant change in the data reservation price under strong negative discrimination, i.e. when the subsequent payoff decreases strongly for one data category. Interestingly, this increase in the reservation price is observed also for participants that do not fall into the corresponding category. Furthermore, we find important gender differences in how general privacy concerns and trust related to the context of the experiment affect both the general willingness to sell the data as well as the reservation price of the data. The general privacy concerns and trust with respect to the decision context are thereby derived from answers of a comprehensive survey

[3]See for instance New York Times (2018).

about privacy-related attitudes and behavior. These findings are important in light of the consequences of personal data sharing for subsequent market interaction, not only with respect to price discrimination, but the usage of personal data more generally.

Related Literature

This article mainly relates to two areas of research: (1) Price discrimination based on consumers' personal data and (2) Consumers' valuation of private information.

Academic research on price discrimination based on consumers' personal data mainly focuses on the effects on market allocations, market efficiency, and social welfare. One common example in academic literature is the use of genetic information for the pricing of health insurance and life insurance contracts (Crocker and Snow (2013), Dionne and Rothschild (2014), Crainich (2017)).

Montes et al. (2018) analyze theoretically how price discrimination based on consumers' private information affects prices, profits, and consumer surplus in a consumption goods market. In their framework, firms can acquire consumer data for price discrimination from a third party intermediary and individual consumers can prevent the use of their private information by paying a *privacy cost*. The authors find that higher *privacy costs* decrease competing duopolists' profits and increase consumer surplus. In the monopoly case, the effect on consumer surplus and social welfare is ambiguous.

Belleflamme and Vergote (2016) assume that consumers can react to a monopolist's *tracking technology*, that identifies consumers' willingness to pay with a certain probability, by making use of a *hiding technology*. The authors show that while more accurate price discrimination by the use of such *tracking technologies* decreases consumer surplus, the availability of privacy protecting technologies may imply an even higher reduction of consumer surplus. The rationale is that the availability of privacy protecting technologies incentivizes the monopolist to limit the use of the *tracking technology* and to raise the regular market price of the product or service.

In an experimental setting, Richards et al. (2016) analyze perceptions of price fairness and *self-interested inequity aversion* in the context of price discrimination. Consumers with *self-interested inequity aversion* regard prices as unfair and tend to purchase less, if they perceive other consumers to pay a lower price. They tend to regard prices as more fair, if inequity is in their favor, which results in higher purchases. The authors find that the implications of such inequity aversion can be at least partially reversed if consumers are involved in the price formation.

The literature on consumers' valuation of private information is extensive and covers a wide range of fields. Acquisti et al. (2016) summarize and link various streams of theoretical and empirical economic research that investigates individual and societal trade-offs associated with protecting and disclosing personal information. The authors note that privacy related issues of economic relevance can be observed in diverse contexts and that situations can arise in which the protection of privacy can both enhance and reduce individual and social welfare. Further, they find that imperfect information about the purpose and the consequences of data collection severely hinders consumers' ability to make informed decisions about their privacy in digital economies.

Some experimental studies evaluate individuals' valuation of privacy by asking participants indirectly (Beresford et al. (2012), Regner and Riener (2017)) and directly (Benndorf and Normann (2017)) to sell private information.[4]

Beresford et al. (2012) conduct a field experiment to measure participants' willingness to pay for privacy. Participants are given the opportunity to purchase a DVD from one of two online stores, for which they have to provide personal information, such as last name, postal address, and e-mail address. In addition to those common data items, one store requires information about the date of birth and monthly income, whereas the other store asks for year of birth and favorite color. Except that the first store requires more sensitive personal information than the latter, both stores are identical. The authors find that when DVDs are offered for one Euro less at the store asking for more sensitive information, almost all participants choose to buy from the cheaper store. When prices are identical at both stores, however, participants buy equally often from either one.

In a Pay-What-You-Want (PWYW) online music store and a mimicking online experiment, Regner and Riener (2017) analyze the effect of revealing customer information, such as name and e-mail address, to the seller on consumers' purchasing behavior. While for donations and public goods, reduced anonymity can lead to higher PWYW revenues due to self-image motivations, Regner and Riener (2017) find that revealing customers' information in the online store context reduces the number of customers purchasing. Overall, lifting anonymity leads to a revenue loss of 25% (35%) in the online music store (in the online experiment). The authors conclude that the substantial reduction of customers might be explained by privacy concerns.

Benndorf and Normann (2017) conduct laboratory experiments to assess participants' willingness to sell personal data to a telecommunications company. Participants in the experiments can sell five different bundles of personal information that covered partici-

[4]For a detailed overview, see Kern et al. (2018).

pants' information on (1) preferences (2) contact data (3) both preferences and contact data (4) facebook profile and (5) facebook timeline. The authors find considerable heterogeneity in participants' willingness to sell personal information. About one sixth of the participants refuse to sell any personal data while a similar fraction sells for 2.50€ or less. The average price requested is 15€ for contact details and 19€ for Facebook data. The authors also find a gender effect: Female participants' valuation of personal data appears to be more sensitive to the type of data.

Various articles analyze the effect of external factors on the value of privacy. These include among others endowment effects (Acquisti et al. (2013)), pre-existing attitudes or dispositions, limited cognitive resources, and momentary affective states (Kehr et al. (2015)), data-breach notifications (Feri et al. (2016)), positive or negative information on companies' attitudes towards privacy (Marreiros et al. (2017)), the number of information recipients (Schudy and Utikal (2017)), and implicitly and explicitly stated prices, political orientation, income proxies and membership in loyalty programs (Plesch and Wolff (2018)).

Benndorf and Normann (2017) name three explanatory factors for differences in results between different studies on the value of privacy: incentivized decisions to share personal information, a salient focus on privacy issues, and transparent information with respect to the use of data shared.

In this experiment, we highlight in the experimental instructions for participants that personal information sold is not shared with third parties nor used for other purposes than for the data analysis in this experiment. Subjects are informed about the use of their personal information and data is sold explicitly. Hence, our experimental setting ensures salience, incentivization, and transparency. Apart from the pure value of privacy, we examine the effects of financial discrimination based on the personal information shared.

The remainder of this paper is organized as follows. In the next section, we present the experimental set-up and our hypotheses. Section 3 provides the data analyses and discusses the results. The final section concludes.

2 Experiment

2.1 Experimental Design

We apply a between-subject design to analyze individuals' valuation of personal data in light of potentially discriminatory use of this data. Specifically, we examine whether and how subjects' willingness to sell their personal data is affected by both inherent privacy concerns and financial discrimination based on their data. The mechanism by which personal data is bought/sold in the experiment is the Becker-DeGroot-Marschak (BDM) mechanism (Becker et al. (1964)), a standard incentive-compatible method for eliciting private values in laboratory experiments. Only if the participant wants to sell his or her personal data according to the BDM mechanism, the data is collected by the experimenters.

The experiment consists of three parts: In Part 1, participants practice the BDM procedure by having the option to sell back a 5 CHF coin which they receive as initial endowment to the experimenter. In Part 2, subjects can sell their personal data to the experimenters via the BDM mechanism. In the experiment, the personal data is the following bundle of personal information: The participant's height, weight, gender, bank account balance, and a picture of the participant's face. If participants sell their personal data via the BDM mechanism, their personal data is collected subsequently. In Part 3, participants first receive an additional payoff, and after that play a trust game, make decisions with respect to risky payoffs, and answer a post-experimental questionnaire. The payoff at the beginning of Part 3 represents the "payoff discrimination stage" in a reduced form: Depending on the experimental treatment, the payoff differs according to whether data was sold and the content of the data. There are five experimental treatments, which differ precisely in whether and how subjects are categorized based on their personal data, and the associated Part 3 payoffs. Participants have full information about data dependent categorization and associated payoffs prior to selling the data.

At the beginning of the experiment, control questions ensure that subjects have understood the instructions that describe the experiment. After all participants have successfully completed the control questions, and before Part I starts, subjects have the opportunity to self-verify their personal data in the absence of the experimenters. For this purpose, a measurement tape and a scale are located in the entry hall in front of the laboratory. Participants are informed that they can go to this entry hall in order to measure their height and weight and use their own cell phones to self-verify their bank account balance without experimenters observing it.

After the self-verification, participants begin with Part 1. Part 1 precedes the selling of personal data to familiarize participants with the BDM mechanism. Subjects receive a 5 CHF coin which they can sell back to the experimenters via the Becker-DeGroot-Marschak (BDM) mechanism (Becker et al. (1964)). The BDM mechanism by which the 5 CHF coin is bought/sold works as follows. Subjects have to state the minimum amount of money they would accept in exchange for an object they could sell to the experimenters. We refer to this amount as the reservation price for the respective object. The market price is then determined by a random draw. If this market price exceeds the reservation price stated by a participant, the object is sold and the subject receives the randomly determined market price as a payment. If the reservation price claimed exceeds the market price, the participant keeps the object but does not receive any money. Because the BDM procedure can be rather demanding, we make several arrangements to familiarize the participants with the mechanism. Following Grether and Plott (1979), we stress that subjects have an incentive to state their true valuation and that renegotiations are excluded. We also clarify that the random draw is independent of actual choices. Finally, prior to making their actual choices, subjects also have the possibility to conduct several tests with different prices and random draws using a payoff simulator that displays the hypothetical outcomes. Participants can use the simulator as long as they want to before continuing to Part 2.

In Part 2, participants can sell the bundle of personal data to the experimenters via the BDM mechanism. Participants are informed that their market price is randomly drawn from the range $[0\text{ CHF}, 60\text{ CHF}]$, where each 10 cent increment is equally likely. Prior to selling their data via the BDM mechanism, subjects receive information about what happens with the data sold to the experimenters, i.e. whether and how the data is used subsequently. The BDM ensures that participants only sell the data if they wish to do so. Their decision screen has two items: A field where they can put the minimum price (*reservation price*) at which they are willing to sell the data, as well as a box they can tick which reads "I don't want to sell the data in any case". Participants ticking this box will be classified as not agreeing to sell the data in the results section.[5]

In all parts, the markets are fully internal to the experiment such that no personal data is shared outside of the lab. Participants have this information and are asked a corresponding control question to ensure that there is no ambiguity about data sharing.

Subsequent to the decision to sell their personal data, all participants are brought one by one into a separate part of the laboratory, the *measurement room*, that other participants can neither enter, nor can they hear or see anything that is happening inside this room. If participants have decided to sell their data, their personal data is collected

[5]If participants tick the box but specify a reservation price lower than 60 CHF in the other field, they are asked whether they want to adjust their answer, as the box tick of not selling in any case will receive priority. Independent of that, all participants are asked whether they are certain about the decision or would like to adjust it before proceeding to the next screen.

by experimenters in this separate room.[6] If participants have decided to not sell their data, they are brought to the measurement room by the experimenters and are asked to wait there for 2-3 minutes before they are picked up and brought back to the main part of the laboratory again. This serves the purpose that all participants are brought to the separate room by the experimenters and other participants do not observe who sold their personal data. Participants know this process.

The experimental treatment variation is about whether and how personal data sold is used in Part 3 for discriminatory purposes. In the baseline Treatment (I), the personal data is not used, and all subjects receive the same payoff of 20 CHF from this part of the experiment.[7] I.e., the Part 3 payoff is independent of their personal data and of whether personal data is sold to the experimenters. In this treatment, with the BDM mechanism, we elicit the pure privacy value attached to the personal data.

In Treatments (II)-(V), participants are classified into one of three groups according to their data provided. Subjects that do not sell their personal data are classified as category A. For those that did sell their personal data, the classification is based on both Body-Mass-Index (BMI)[8] (gender-specific) as well as bank account balance. In particular, subjects with a bank account balance above or equal to 1000 CHF and a BMI below 22 (23.5) for female (male) participants are classified as category B.[9] The remaining subjects, i.e. participants who did sell their data and have a bank account balance below 1000 CHF or a BMI above 22 (23.5) for female (male) participants, are classified as category C. Table 1 summarizes how participants' categorization depends on their decision to sell their personal data and the content of the data.

In Treatment (II), subjects are categorized as described above but participants' payoff is independent of their category and of whether they sell their personal information to the experimenter. All subjects receive a payoff of 20 CHF from this part of the experiment as in Treatment (I). We use Treatment (II) additionally to Baseline Treatment (I) to have a second baseline with non-payoff relevant categorization based on personal data, and the corresponding willingness to sell the data.

[6] A detailed description of how data was collected in the measurement room is provided in the additional Online Appendix.

[7] Additional to this payoff, participants can earn money in the trust game and the lottery choice decision in Part 3. Besides providing information about social and risk preferences, these two decision situations serve the purpose of putting weight on Part 3.

[8] The BMI is calculated as $BMI = \frac{Body\ weight\ in\ kilogramm}{(Body\ height\ in\ meter)^2}$.

[9] We chose values in the interval $[18.5, 24.9]$ that defines a *normal* or *healthy* BMI according to the World Health Organization (2018a). Since in Switzerland, the mean BMI for males is slightly higher than for females World Health Organization (2018b), we chose a higher threshold for the BMI categorization for male participants than for female participants.

Table 1. Categorization based on BMI and Bank Account Balance

Sold Data	BMI	Bank Account (CHF)	Category
No	-	-	**A**
Yes	< 22 female (< 23.5 male)	≥ 1000	**B**
Yes	< 22 female (< 23.5 male)	< 1000	**C**
Yes	< 22 female (< 23.5 male)	≥ 1000	**C**
Yes	< 22 female (< 23.5 male)	< 1000	**C**

In Treatments (III)-(V), payoffs in Part 3 vary with the personal data based categorization. This is done in the following way: In Treatment (III), both categories, A and C, receive the baseline payoff of 20 CHF from this part of the experiment. Participants classified as category B, i.e., with a high bank account balance and a low BMI, receive CHF 30 instead. In both Treatments, (IV) and (V), categories A and B receive the baseline payoff of 20 CHF, whereas participants classified as C receive a lower payoff from this part of the experiment. In Treatment (IV), we employ the same payoff difference between category A and C as we have used in Treatment (III) between category A and B. Therefore, participants that are categorized as C receive 10 CHF as a payoff from this part of the experiment in Treatment (IV). In Treatment (V), we amplify the payoff difference and the payoff for subjects in category C is 0 CHF. Table 2 summarizes how the participants' payoffs depend on their data-based categorization in the respective Treatments. Participants receive full information in the experimental instructions about data dependent categorization and associated payoffs prior to selling the data, i.e. they have full knowledge about data-based payoff discrimination.

After payoffs in Part 3, we elicit each participant's tendency to trust others using the loss domain treatment of the trust game from Kvaløy et al. (2017). Further, we elicit participants' risk preferences using the standard (Holt and Laury (2002)) price list. One of the two games is randomly selected to be payoff-relevant. Subsequent to that, participants answer a comprehensive post-experimental questionnaire that collects information on privacy attitudes and behavior. At the end of the experiment, subjects observe a summary screen of their payoffs.

2.2 Experimental Procedure

We make use of a standard laboratory experiment at the ETH Decision Science Lab. The recruitment was performed by the Decision Science Lab using the joint subject pool of University of Zurich and ETH Zurich. Invitations were sent out via email to randomly

TABLE 2. PAYOFFS FOR DIFFERENT EXPERIMENTAL TREATMENTS

Treatment			Payoff		
Name	Abbreviation	Categorization	A	B	C
Baseline Treatment	(I)	no	20	20	20
Baseline Treatment (Categorization)	(II)	yes	20	20	20
Positive Discrimination	(III)	yes	20	**30**	20
Negative Discrimination	(IV)	yes	20	20	**10**
Strong Negative Discrimination	(V)	yes	20	20	**0**

selected German speaking subjects. The subjects signed up for a specific session, whereas the signing up procedure was limited in order to maintain a gender-balanced pool of subjects in each session. Participants are instructed upon arrival at the lab. This includes a short verbal introduction about the organizational procedure as well as written instructions that explain the detailed experimental procedure.[10] Participants receive a show-up fee of 5 CHF. The experiment is partly computerized using the standard z-Tree software (Fischbacher (2007)), and partly performed with pen and paper in the lab (collection of personal data).

The experimental sessions were conducted from August to October and in December 2018 at the ETH Decision Science Laboratory. 282 subjects participated in 29 experiment sessions, 5 sessions were conducted for Treatment (I) and (II) each, 6 sessions per treatment were conducted for Treatment (IV) and (V) and finally 7 sessions of Treatment (III). Participants were, on average, 22.29 years old; 50.0% of the participants were female. All participants were enrolled students. More than one third of the participants were enrolled for natural sciences (28.36%), roughly one fifth for engineering (22.69%), 9.21% for medicine, 6.38% for humanities, and 9.21% for economics. The remaining 24.11% of participants were enrolled in other subjects. Subjects participated in exactly one session. Sessions lasted on average about 75 minutes. Participants earned 55 CHF (including a 5 CHF show-up fee), on average, with some variance depending on participants' own decisions. A comprehensive set of control questions and the BDM practice procedure with the 5 CHF coin as initial endowment ensures that all participants understand the sequence of decisions in the experiment, the payoff consequences, and the BDM procedure.

After the main experiment and the trust game from Kvaløy et al. (2017) and standard (Holt and Laury (2002)) price list decisions, we also launch a comprehensive post-experimental questionnaire. Among other statements, participants can but do not have to specify whether they have self-verified their data, what category they would have been

[10]The German written instructions as well as the respective English translation for Treatment (I) and Treatment (IV) can be found in the additional Online Appendix.

in if they had sold their data (in case they have not sold it), whether they were allocated to the category that they had expected to be allocated to, and why they did not sell their data in case they had not done so. Furthermore, besides including the Falk et al. (2016) preference modules on social preferences and risk attitudes, participants are asked questions about their privacy concerns in other domains, their behavior with respect to private information in social networks and in the communication with various institutions, and about their self-assessment on trust in human beings as well as institutions.[11] In a last step, we ask participants about their rationale when taking decisions in the experiment and record the age, the gender and their field of study. Figure 1 illustrates the timeline of the experiment.

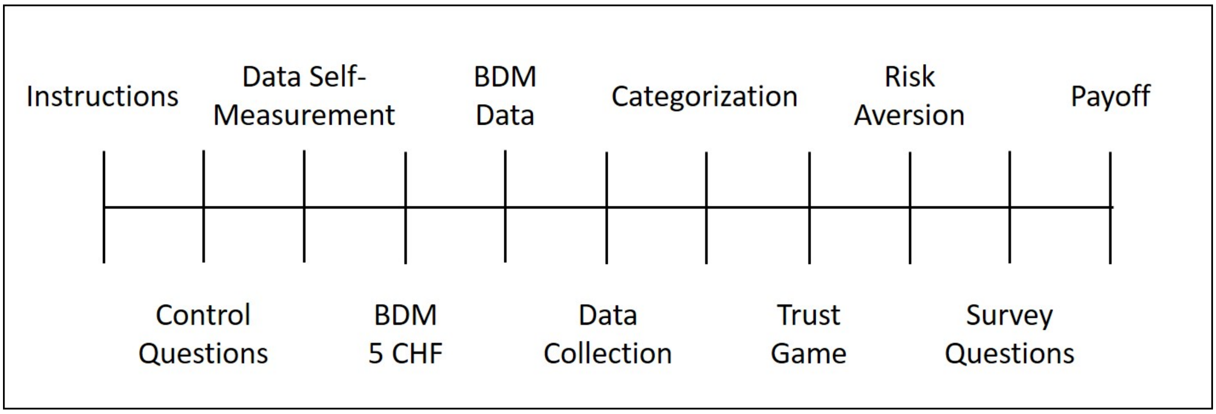

FIGURE 1. EXPERIMENTAL TIMELINE

At the end of each session, participants' payoff is calculated as the sum of the payoff from decisions in Parts 1-3. Participants are paid the total amount in private at the very end of the experiment.

Incentives to sell the data

Without selling personal data, with the show-up fee of 5 CHF, the 5 CHF coin in Part 1 that can result in a payoff exceeding 5 CHF, the baseline Part 3 payoff of 20 CHF as well as the two short games at the end of Part 3, participants receive a payoff that corresponds approximately to the average payoff of participating in a laboratory experiment of that length in Zurich. I.e., participants are remunerated for their time cost and according to standard expected payoffs even without selling the data. Thus, the stated reservation price for the personal data does not have the problem of needing to remunerate participation in the experiment per se, but captures the valuation for the personal data.

[11]The full German questionnaire as well as the respective English translation can be found in the additional Online Appendix.

2.3 Predictions

In the experiment, we elicit the valuation of personal data and analyze whether and how data-based financial discrimination is accounted for. In the experiment, subjects have full information about subsequent payoff discrimination. If a subject's utility depends only on final payoffs, then by backward induction the reservation price (RP) for personal data should fully adjust to payoff differences of the data-based categories.[12]

For the Baseline Treatments (I) and (II), the only difference is that participants are categorized based on their personal data, but there are no associated payoff differences. Unless there is a (dis)utility of categorization based on personal data per se, even without payoff consequences, the average RP should not differ. In the Discrimination Treatments (III-V), the RP adjustments should be upward for category C types in Treatments (IV) and (V) compared to both Treatment (III) and the Baseline Treatments. For category B types, the RP adjustments should be downward in Treatment (III) compared to all other treatments. Importantly, if utility depends only on final payoffs, then there should be no difference in category B types' RP between the Baseline Treatments and the Negative Discrimination Treatments (IV) and (V). Given the design, these RP adjustments by category/type are also underlying the overall RP adjustments by treatment. These predictions are summarized below. The average reservation price in treatment j is denoted by $\overline{RP}_j$ and the average reservation price for revealed category k in treatment j by $\overline{RP}_j^k$.

Hypothesis. ***Overall RP***
The average reservation price does not differ between Treatment (I) and (II).
Furthermore, $\overline{RP}_{III} \leq \overline{RP}_{I+II} \leq \overline{RP}_{IV} \leq \overline{RP}_V$.

Hypothesis. ***RP by Revealed Category***
For revealed category B, $\overline{RP}_{III}^B \leq \overline{RP}_{I+II}^B = \overline{RP}_{IV}^B = \overline{RP}_V^B$.
Furthermore, for revealed category C, $\overline{RP}_{III}^C = \overline{RP}_{I+II}^C \leq \overline{RP}_{IV}^C \leq \overline{RP}_V^C$.

3 Results

This section presents the main results of the experiment: Subsection 3.1 provides an overview of the results on the reservation price for personal data in light of subsequent

[12] As the personal data in our experiment pertains to personal attributes of the individual, i.e. type and not behavior, we cannot analyze the impact of financial discrimination in a within-subject design, as the personal attributes (contrary to behaviors) cannot be sold multiple times. In our between-subject design, we therefore need to analyze the impact of discrimination across groups of subjects.

financial discrimination. In section 3.2, we use the data from the post-experimental survey to uncover privacy attitudes. Section 3.3 uses these in regression analyses of the reservation price for personal data under financial discrimination.

3.1 The Willingness to Sell Personal Data

Before presenting the results on the valuation of personal data, we look at the results from Part 1—the selling of a 5 CHF coin using the BDM mechanism—to check the understanding of the BDM method among participants. In the following, we will refer to the minimum price for the good (either the 5 CHF coin or the personal data bundle) stated by the participant as the (participant's) 'reservation price' and to the random draw in the BDM as the 'market price'.

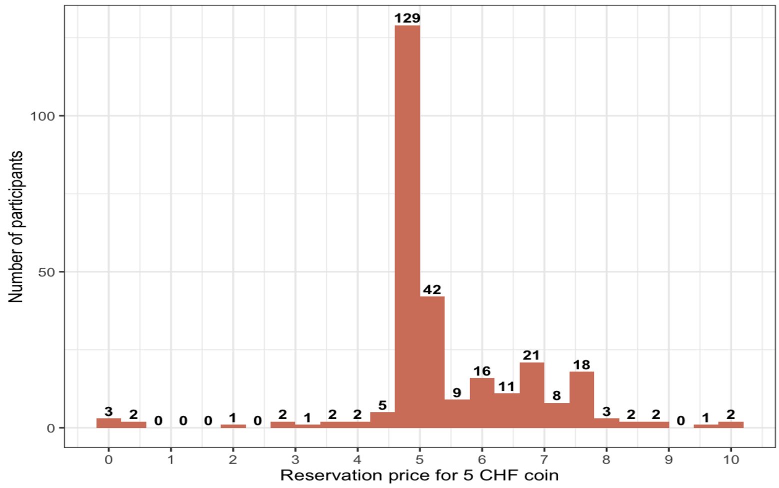

FIGURE 2. HISTOGRAM OF RESERVATION PRICES FOR THE 5 CHF COIN

The participants' reservation prices for the 5 CHF coins are reported in Figure 2. Values exceeding 5 CHF could be attributed to an endowment effect.[13] Reservation prices below 5 CHF are less easy to explain and suggest that the corresponding participant might not have understood the workings of the BDM or how to set their selling price. For this reason, we chose to exclude the 11 observations (participants) stating a reservation price below 4 CHF in the subsequent analysis of the willingness to sell personal data.[14]

A summary for the 271 remaining observations is provided in Table 3. Female and male participants are represented equally. Whether we observe the personal data (type) depends on the willingness to sell the data as well as on the market price, i.e. the results from the random number generator. There are three scenarios:

[13]Furthermore, a gambling motive might explain high stated reservation prices.
[14]These 11 observations are not excluded from the survey analyses.

1. Participant refuse to sell the personal data in any case and signal this by checking the box for "*I don't want to sell the data in any case*". In this case, no data is sold and consequently the data cannot be used in the analyses. In the following, these observations are classified as **"not agreeing to sell the personal data"**.

2. Participants are willing to sell the personal data for a certain price but the stated reservation price exceeds the market price. In this case, no data is sold and the participant's type is not observed. The reservation price is available.

3. Participants are willing to sell the personal data and the market price exceeds the reservation price. In this case, the personal data is collected and can be used in the analyses.

	Total no. of subjects			No. of subjects that did not agree to sell the data			No. of subjects that did not sell the data			No. of subjects that sold the data		
		female	male		female	male		female	male		female	male
Baseline Treatment (I)	46	21	25	5	2	3	22	13	9	24	8	16
Baseline Treatment with Categorization (II)	47	22	25	9	6	3	25	13	12	22	9	13
Positive Discrimination (III)	61	31	29	9	8	1	27	16	10	34	15	19
Negative Discrimination (IV)	58	30	28	5	2	3	29	14	15	29	16	13
Strong Negative Discrimination (V)	59	31	28	13	7	6	32	20	12	27	11	16
Total	271	135	135	41	25	16	135	76	58	136	59	77

Table 3. Number of observations, grouped by experimental treatment, gender and reservation prices

Columns (5) to (7) in Table 3 show the participants not agreeing to sell by treatment and gender, (8) to (10) in Table 3 summarize the numbers of observations for which no personal data was sold, that are all observations in the first two scenarios. Columns (11) to (13) in Table 3 show the numbers of observations for which the data was sold.[15]

Across all treatments, 15.1% (41/271) of the subjects refuse to sell the data at all, even though the data stay within the experiment. To compare, the rates of refusing to sell facebook timeline data and a combination of preference and contact data to a telecommunications company in Benndorf and Normann (2017) are at roughly 20%. Figure 3 displays

[15] In one Positive Discrimination Treatment (III) session, one of the participants did not indicate his or her gender in the post-experimental questionnaire. This participant agreed to sell private information for a certain price but did not sell the data (scenario 2). As a result, the observations for male and female subjects do not add up to the overall number of subjects for subjects who agreed to sell data (columns (11) to (13)), for subjects who did not sell the data (columns (8) to (10)) and for the total number of subjects (columns (2) to (4)).

the shares of participants not agreeing to sell the data by treatment. These shares do not differ significantly across treatments, in particular, the shares for the Treatments (III)-(V) do not differ significantly from the shares in the Baseline Treatments (I) and (II). Thus, a first observation is that financial discrimination based on personal data per se does not significantly impact the decision on agreeing to sell data overall, as measured by ticking the box "*I don't want to sell the data in any case*".

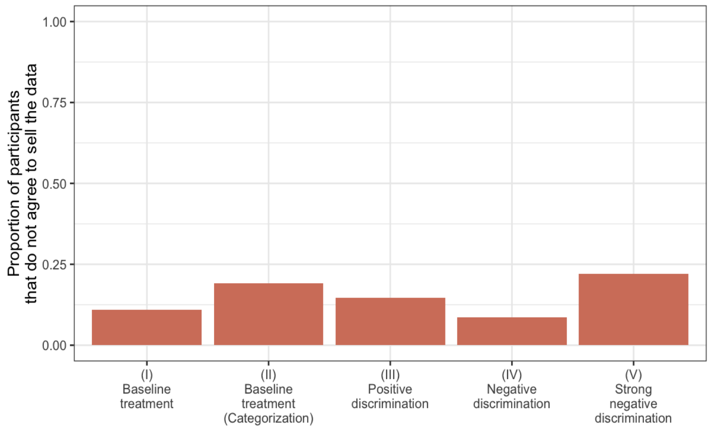

FIGURE 3. SHARE OF SUBJECTS THAT DO NOT AGREE TO SELL THE DATA BY EXPERIMENTAL TREATMENT

We now turn to the valuation of personal data as measured by the participants' stated reservation price for those who indicate one. Figure 4 displays the boxplots with quantiles, means and individual observations differentiated by whether the data was actually sold per treatment. The data presented is based on all subjects that indicate a RS, e.g. it is excluding the subjects that do not agree to sell the data. In both baseline treatments, in which the personal data has no impact on subsequent payoffs, the mean reservation price for participants who are willing to sell their data is roughly 24 CHF. This suggests that participants on average place a substantial monetary value on the privacy of their personal data, even when the data remains fully internal to the experiment. The share of subjects that indicate a reservation price of 0 is 8.7% and 8.5% in Treatments (I) and (II) respectively.

For a first analysis of the role of financial discrimination on the value of privacy, we group the non-discrimination treatments (I) and (II) together and compare these to the Discrimination Treatments (III)-(V).[16] In line with predictions, the mean reservation price is lower compared to the Baseline (I+II) when data-based discrimination is only positive in Treatment (III)(19.77 CHF), whereas it is higher when data-based discrimina-

[16]As the Baseline Treatments (I) and (II) do not differ in important ways in the distribution of reservation prices, they are pooled together for most parts of the analysis.

tion is negative (28.22 CHF in (IV) and 30.66 CHF in (V) respectively). For this average reservation price, the differences are statistically significant for (V) versus (I+II) (MWU: one-sided $p < 0.01$), as well as comparing the Negative Discrimination Treatments against the Positive Discrimination Treatments: (V) vs. (III) (MWU: one-sided $p < 0.01$) and (IV) vs. (III) (MWU: one-sided $p = 0.02$).[17]

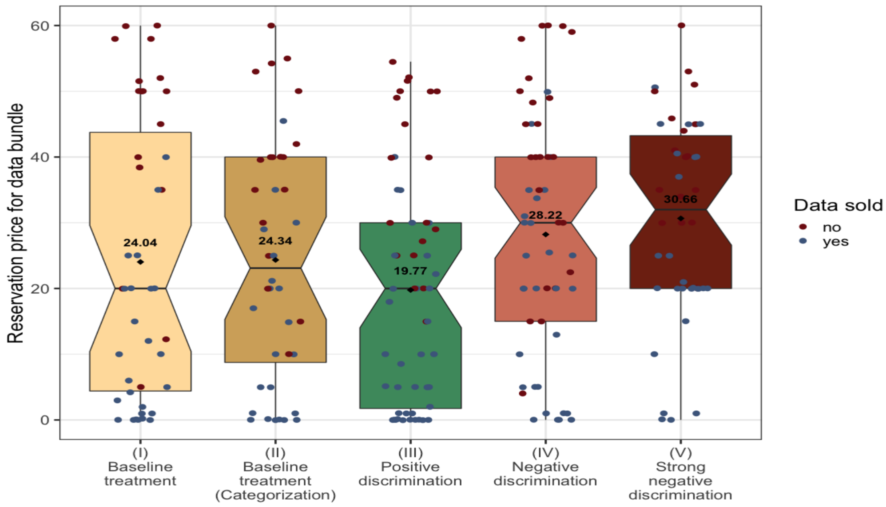

Figure 4. Reservation price of data by treatment. Points represent individual observations and their color indicates, whether the bundle was sold

Figure 5 displays the average reservation price stratified by gender. A first observation is that the mean reservation price of women is higher than that of men in each treatment. This difference in reservation prices between women and men is statistically significant in Baseline Treatment (I + II) (MWU: $p < 0.01$) and Treatment (V) (MWU: $p = 0.012$).

The higher mean reservation price also translates into lower shares of personal data revealed, as illustrated in Figure 6.[18] Interestingly, when stratifying by gender, we find that the differences in average reservation prices across experimental treatments for women are not statistically significant at the 5% level, except for (V) vs. (III). Thus, the above found treatment differences, in particular of Treatment (V) versus Baseline Treatment (I+II), are more strongly driven by changes in men's reservation prices.

Payoff discrimination is based on a subject's category. The experimental variation allows us to compare reservation price adjustments conditional on the category, and compare it to the payoff differences stemming from data-based discrimination. Table 4 displays the mean reservation prices by treatments, category, and gender as well as the corresponding numbers of observations. In the following, we will analyze reservation prices by category,

[17]Table 10 in Appendix A.1 shows nonparametric test results for all treatment comparisons and by gender.

[18]These differences are not statistically significant.

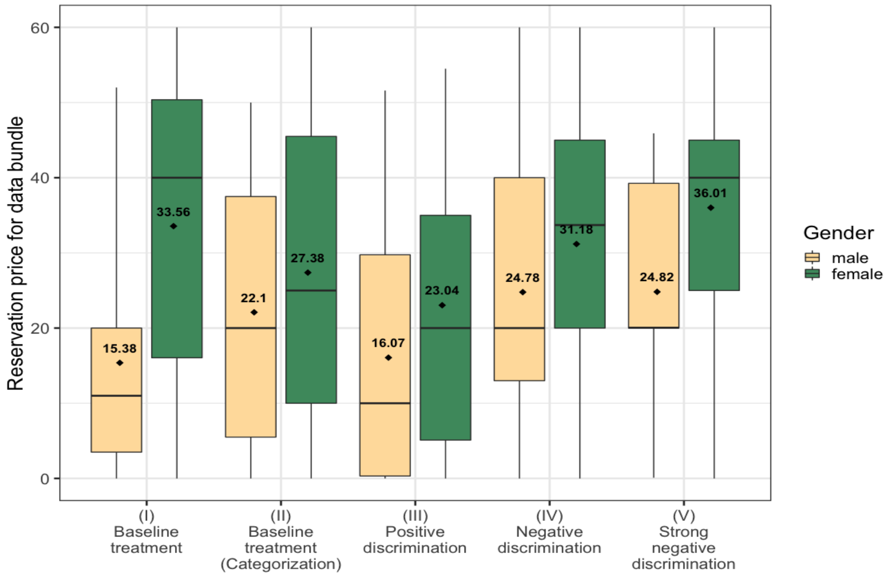

Figure 5. Mean reservation prices for data by treatment and gender

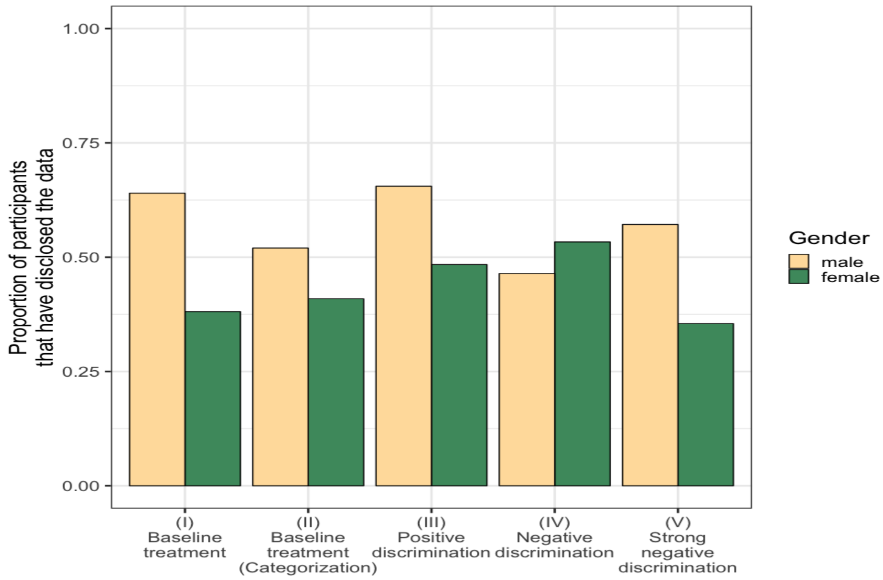

Figure 6. Proportion of participants selling personal data by gender

but refrain from a detailed discussion of gender differences by category due to our small sample sizes for gender-category groups.

Figure 7 displays the mean reservation prices by revealed category B and C. A first observation is that category C types have a higher average reservation price than category B types in Treatments (I) and (II), despite facing the same payoff consequences.

We find that the mean reservation price of category B types is significantly lower in (III) than in (V) (MWU: $p < 0.01$), as predicted based on simple payoff differences. The mean reservation price in (III) is also lower than that in the Baseline (I+II) as predicted,

		Overall all	Overall female	Overall male	Category B all	Category B female	Category B male	Category C all	Category C female	Category C male
(I)	mean	24.62	35.32	15.38	7.66	2.10	9.52	13.31	17.67	10.70
	(no. obs)	41	19	22	8	2	6	16	6	10
(II)	mean	23.69	26.81	21.28	11.74	9.15	14.33	14.89	27.90	9.31
	(no. obs)	38	16	22	12	6	6	10	3	7
(III)	mean	20.12	22.77	16.91	7.80	8.33	7.67	13.34	16.02	8.73
	(no. obs)	52	23	28	15	3	12	19	12	7
(IV)	mean	27.81	30.51	24.78	15.59	14.94	16.67	20.07	26.85	12.60
	(no. obs)	53	28	25	8	5	3	21	11	10
(V)	mean	31.29	36.01	26.35	22.11	32.23	14.23	28.42	32.62	26.01
	(no. obs)	46	24	22	16	7	9	11	4	7

Table 4. Mean reservation price of personal data, grouped by experimental treatment, gender and category

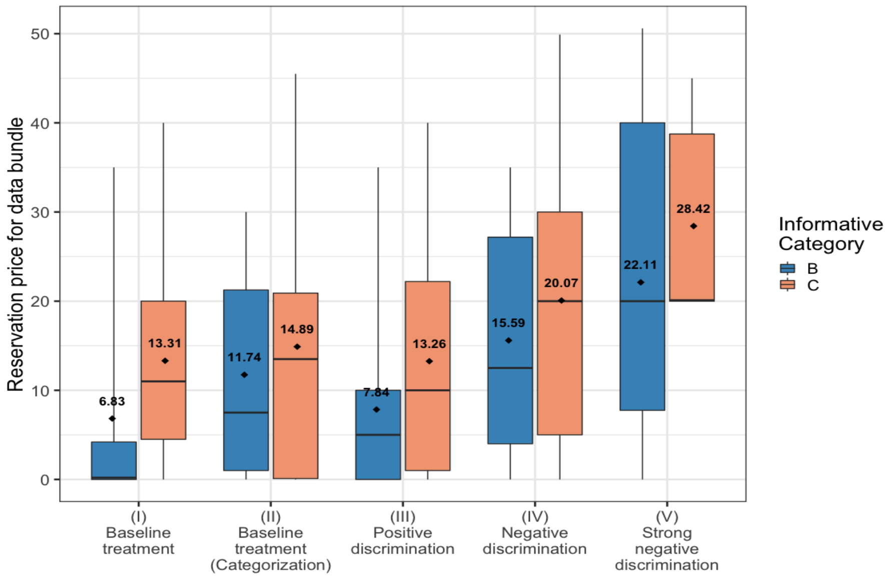

Figure 7. Mean reservation prices for data by treatment and revealed category

however the difference is not statistically significant. Furthermore, we find that mean reservation prices in Treatments (IV) and (V) are higher than in Treatments (I) and (II). Participants with revealed category B demand on average 7.4 CHF more in Treatment (IV) and 11.94 CHF more in Treatment (V) compared to the pooled Baseline Treatment (I + II) without financial consequences. The difference between Treatment (V) and (I+II) is statistically significant (MWU: p=0.04).[19] This finding is curious, as category B types face the same payoff consequences in Treatments (V), (I) and (II). One interpretation of this result is that, while there is full information on how the personal data is categorized, subjects perceive their data or in which category they fall as uncertain and therefore increase their reservation price. Another potential interpretation is that while category B

[19]Table 11 and Table 12 in Appendix A report the effect sizes alongside with p-values from several nonparametric tests by category.

types are not discriminated against, they experience a disutility similar to advantageous inequality aversion from the data-based financial discrimination and therefore increase the reservation price.

For category C, changes in subsequent payoff translate in statistically significant shifts of reservation prices in line with predictions: The mean reservation price is significantly lower in (III) vs. (V) (MWU: $p < 0.01$) and significantly higher in (V) vs. (I + II) (MWU: $p < 0.01$). The higher mean reservation price in (IV) compared to (III) is not statistically significant. Looking at the magnitudes, we observe that the average reservation price of category C types in (V) is 14.50 CHF higher than that in (I+II), compared to a payoff difference of 20 CHF.

3.2 Survey Analysis

Besides subsequent data-based financial discrimination, other characteristics may influence the decision to sell personal data. To minimize the risk of omitted variables in our analysis, in this section we will discern and quantify other relevant factors of the *privacy calculus* with the help of a post-experimental survey. In Section 3.2.1, we introduce the survey questions that participants are asked to answer, and in Section 3.2.2, we discuss the role of two latent variables that can affect participants' decision whether to sell their private information: trust towards the experimenters and general attitudes towards privacy related issues. Sections A.2 and A.3 in the Appendix describe the construction of these latent variables via exploratory and confirmatory factor analyses.

3.2.1 Survey creation

Table 5 presents parts of the post-experimental survey questions that are grouped into six different categories: (1) Risk aversion (2) Value of privacy (3) General privacy concerns (4) Privacy related behavior (5) Social network usage and (6) Lack of trust towards the experimenters. Henceforth, we refer to the individual questions using the corresponding abbreviations listed in the table.

Malhotra et al. (2004), Kehr et al. (2015) and Smith et al. (1996) have proposed and tested instruments to infer individuals' attitudes towards privacy. Following Kehr et al. (2015), we assess the general importance of privacy for the participants by slightly adapting 3 out of 5 questions from the construct *Global informational privacy concerns* of Malhotra et al. (2004). These questions are designated as VP.1 - VP.3.

Since, the motivation for our experiment partially comes from the secondary usage of personal information by companies, measuring attitudes towards privacy in this context could potentially be a strong predictor in our regression analyses. Therefore, we include questions GP.1 - GP.6 from Smith et al. (1996) in order to assess subjects' perception of information privacy practices in organizations.

Expressed general attitudes are not always reliable predictors for behavior. Recent research on online user behavior has discovered great discrepancies between users' stated attitudes and their actual behavior with respect to the value of their privacy. Specifically, users tend to claim a high level of concern about their privacy, while they engage in little action to protect their private information. This phenomenon is known as the privacy paradox.[20] Potential reasons behind this discrepancy are numerous: situation specific factors, affect, perceived behavioral control, or social desirability bias. To assess whether individual's concerns actually translate into concrete actions, we collect more information about participants' behavior in situations related to protecting personal data (PB.1 - PB.8).

Voluntary sharing of personal data in social networks is a frequently observed and studied phenomenon in the context of the privacy paradox. Active users of online social networks seem to be either unaware or unconcerned about the consequences of their data sharing. To capture this potentially informative behavior, we develop and add questions SN.1 - SN.10 to our survey.

In the context of our experiment, it is of great importance whether study participants believe in the integrity of the experimenters and hence whether subjects believe that all information stated in the instructions is true. To capture this, we include questions IT.1 - IT.3. Acquisti et al. (2016) consider control over information flow integral to the very definition of privacy. This further motivates question IT.3 that clarifies whether participants believe they can influence the subsequent usage of information provided during the experiment.

[20]For a literature review on the theories regarding the privacy paradox, see Barth and de Jong (2017).

Abbreviation	Question
Risk aversion	
RA	How do you see yourself: as a person who is generally willing to take risks, or as someone who prefers to avoid them?
Value of privacy	
VP.1	Compared to others, I am more sensitive/cautious with respect to how companies handle my personal information.
VP.2	To me, keeping my data private is of highest importance.
VP.3	Compared to others, I tend to be more concerned about threats to my data privacy.
General privacy concerns	
GP.1	It usually bothers me when companies ask me for personal information.
GP.2	Companies should not use personal information unless it has been authorized by the respective person.
GP.3	Companies should invest more time, effort, and costs in preventing unauthorized access to personal information.
GP.4	I (sometimes) think twice before sharing private information with companies.
GP.5	Companies should never share personal information with other companies without authorization by the respective person.
GP.6	I am concerned that companies collect too much personal information about me.
Privacy related behavior	
PB.1	Do you use a mobile app to execute transactions from your bank account or to check your account balance?
PB.2	Do you hide your bank card's PIN number when using an ATM or making purchases?
PB.3	Do you read the privacy policy before registering on a website?
PB.4	Do you remove cookies?
PB.5	Do you check your computer for spyware?
PB.6	Do you use the private browser mode?
PB.7	How often do you change your passwords?
PB.8	Do you use the same password for different websites and services?
Social networks usage	
SN.1	How often are you active or online on social networks (e.g. Facebook, Twitter, Instagram, ...)?
SN.2	I share my general contact information (name, hometown, age, occupation).
SN.3	I share my online contact information (email, Skype, MSN).
SN.4	I share my physical contact information (phone number, address).
SN.5	I use my own photograph in my profile.
SN.6.	I am honest with respect to the information about myself in my profile and my posts.
SN.7	I post information on my current mood.
SN.8	I share content and engage in activities that reveal my lifestyle.
SN.9	Do you use the privacy settings to control who can see which piece of your information in social networks?
SN.10	Do you delete anything that you have posted in the past?
Lack of trust towards experimenters	
IT.1	I am concerned that the information I share during this experiment could be misused.
IT.2	I am concerned about providing personal information during this experiment because it could be used in a way I did not foresee.
IT.3	I believe I am in control over how my personal information is used by the experimenters.

Table 5. Excerpt from the post experimental survey (english translation)

3.2.2 Latent variables

In order to identify latent characteristics of the participants, we first conduct an exploratory factor analysis to group the questions presented in the previous subsections into seven groups, of which each measures one latent characteristic. Via a confirmatory factor analysis, we identify two groups that satisfy our reliability criteria and that provide us with a good model fit. The detailed process of constructing and selecting the latent variables is described in Section A.2 and Section A.3 in the Appendix. In the current subsection, we introduce and interpret the two latent variables estimated by our measurement model (Figure 14 in Section A.3 of the Appendix). The plausibility of conclusions is checked by the means of data plots on the one hand and using correlation analysis on the other.

The first latent factor, denoted as PC, is measured through the responses to questions VP.1 - VP.3 using a 7-point Likert-scale anchored by the response options "*very untrue for me*" and "*very true for me*". The relationships between survey items and the estimated value of latent characteristics are depicted in Figure 8. This visualization suggests that subjects with higher values of PC put more value on their privacy and tend to be more concerned about it.

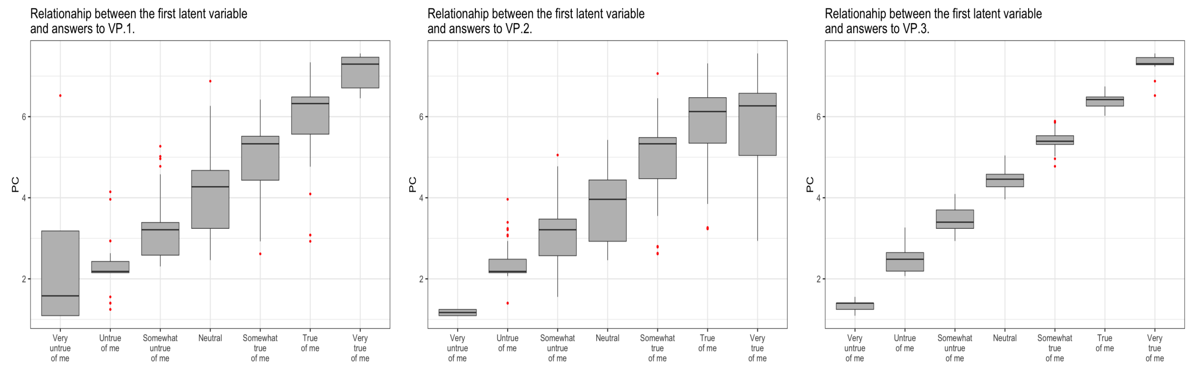

FIGURE 8. VALUES OF THE ESTIMATED VARIABLE PC GROUPED BY SURVEY RESPONSES

To cross-check our interpretation, we compute correlations between PC and various related characteristics. The results reported in Table 6 confirm our hypothesis that the latent variable PC captures general attitudes towards privacy related issues, as we find a significant correlation between PC and responses to the questions related to general attitudes towards privacy and privacy related behavior.[21]

[21]To measure the relation between the continuous latent variable PC and ordinal answers to post-experimental survey questions, we use Kendall's tau correlation coefficient. An association with the binary decision whether to sell private information is captured using point biserial correlation. For the relationship between PC and the reservation price for the personal data bundle, we use the Pearson correlation coefficient.

	General privacy concerns						Privacy related behavior							
	GP.1	GP.2	GP.3	GP.4	GP.5	GP.6	PB.1	PB.2	PB.3	PB.4	PB.5	PB.6	PB.7	PB.8
tau	0.346	0.106	0.219	0.372	0.147	0.454	-0.193	0.139	0.235	0.256	0.294	0.19	0.207	-0.274
pvalue	0.000	0.026	0.000	0.000	0.002	0.000	0.000	0.002	0.000	0.000	0.000	0.00	0.000	0.000

	Social Network usage									
	SN.1	SN.2	SN.3	SN.4	SN.5	SN.6	SN.7	SN.8	SN.9	SN.10
tau	0.138	0.155	0.125	0.081	0.162	0.159	0.020	-0.091	0.195	0.062
p-value	0.003	0.002	0.011	0.102	0.001	0.001	0.671	0.045	0.000	0.165

Table 6. Correlation between PC and other variables

We denote the second latent variable that is determined by questions IT.1 and IT.2 as *Distrust towards the experimenters* (DTE).[22] Figure 9 illustrates the relationship between the survey items and the estimated value of DTE. The boxplots show that high values of the latent variable DTE indicate that a participant is concerned about the protection and the use of private information that is sold to the experimenters.

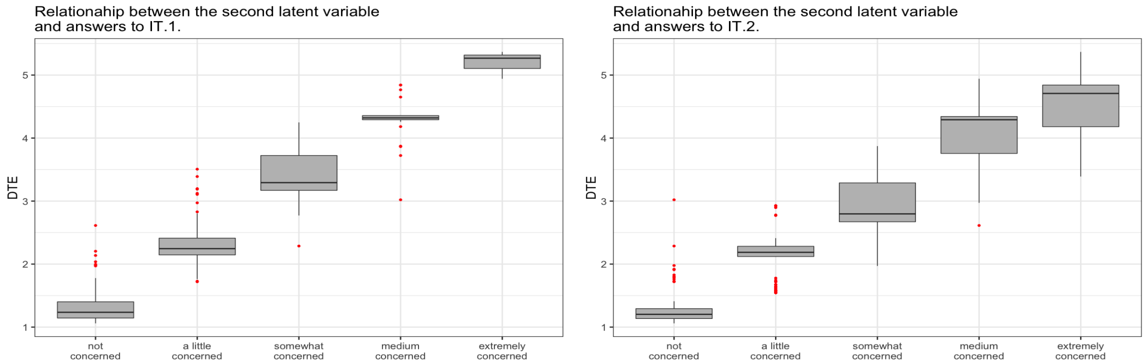

Figure 9. Values of the estimated variable DTE grouped by survey questions

Again, we cross-check our interpretation using other variables generated from the survey questions. DTE is not significantly correlated with answers to question I.16 and 10% correlated with answers to I.17. Furthermore, there is no association between the second latent variable DTE and the amount insured in the trust game. These discrepancies can be attributed, however, to the importance of context and situational factors. This suggests that DTE captures the specific trust towards the experiment with respect to the personal data rather than more general trust as it is captured in the trust game.

Our results from the previous subsection point at gender differences with respect to privacy related decision making. One question is whether the variables DTE and PC could shed light on the sources of gender differences. To address this, we first compare the

[22]The answer pattern suggests that responses to IT.1 and IT.2 are affected by a different unobserved factor as the answers to IT.3. Indeed, the first two questions assess trust towards the experimenters, whereas the last question refers more to the perceived control over subsequent usage. It is plausible that subjects do not trust the experimenters, but are still convinced they have control over the situation and vice versa. More details can be found in Section A.2.

variables' distribution by gender. Figure 10 presents population kernel density estimators of PC and DTE. Neither visual inspection nor statistical tests (Table 11) indicate that male or female subjects tend to be more concerned about their privacy or more distrustful towards experimenters.

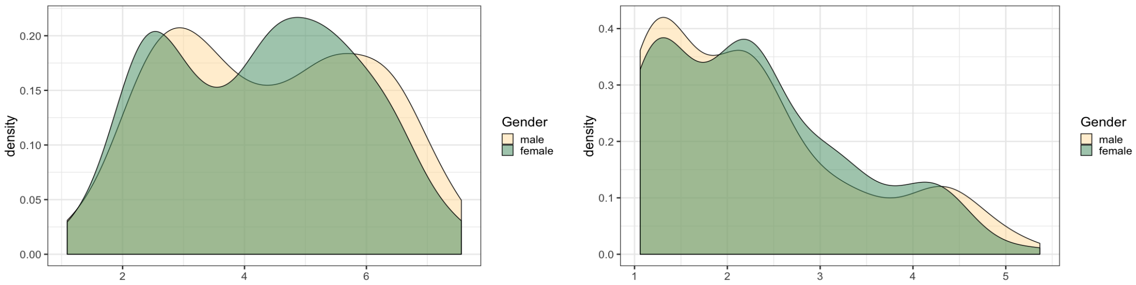

Figure 10. Distribution of created variables: PC (left) and DTE (right) grouped by gender

Variable	mean difference	effect size	MWU	KS test
PC	-0.087	-0.054	0.698	0.798
DTE	-0.083	-0.073	0.664	0.723

Figure 11. Comparison of distributions of latent characteristics by gender

For a first indication of whether these two latent characteristics are related to participants' willingness to sell personal data, we determine the latent variables' correlation with subjects' decision to agree to sell personal data, as well as with participants' reservation price for the data bundle. Table 7 presents these correlations separately for both gender and both latent variables. Since latent variables are also correlated, to account for potential confounding, we compute partial correlation coefficients. PC is significantly correlated with women's decision to agree to sell personal data and on the reservation price they demand in return. There is no statistically significant association between DTE and the main decision variables for female participants. For male participants, the opposite is the case: DTE exhibits a significant and positive correlation with the reservation price and a significant and negative correlation with the binary decision variable whether to agree to sell the data. There is no statistically significant association between PC and the main decision variables for male participants. Thus, both latent variables could increase the explanatory power of our regression models and are therefore included in the subsequent analyses.

	Female				Male			
	Reservation price for data		Agree to sell the data		Reservation price for data		Agree to sell the data	
	full	partial	full	partial	full	partial	full	partial
PC	0.365 (0.00***)	0.337 (0.00***)	-0.468 (0.00***)	-0.309 (0.00***)	0.141 (0.125)	-0.005 (0.961)	-0.165 (0.241)	0.001 (0.988)
DTE	0.157 (0.101)	-0.046 (0.636)	-0.173 (0.167)	0.068 (0.435)	0.308 (0.00***)	0.300 (0.00***)	-0.350 (0.012**)	-0.191 (0.027**)

Partial correlation: measures association between two variables controlling for the effect of the third.

Table 7. Correlation between latent variables and decision variables by gender.

3.3 Regression Analysis

We can now extend our analysis of the willingness to sell personal data under financial discrimination using the two latent variables for privacy concerns and trust derived from survey responses. We perform regression analyses to model two decisions made by the participants: whether to agree to sell the data at all and which reservation price (RP) to choose for the data.

As the first stage, we examine what factors determine whether an individual agrees to sell the data. We fit a probit regression model where the dependent variable equals 1 if the subject does not tick the box "I don't want to sell the data in any case", i.e. the subject indicates a reservation price. Since both treatments do not feature any financial discrimination, we pool together Treatment (I) and Treatment (II), as discussed in Section 3.1.

Besides controlling for underlying privacy concerns and trust, we include further covariates, which are either derived from responses in the survey or participants' decisions in other parts of the experiment. For the covariates from survey responses, first, we control for e-banking usage, since it is related to effort costs of information disclosure. Furthermore, we include whether participants self-checked their personal data. The variable 'perceived control over data' is a binary variable derived from the answers to the survey question of whether the subjects agree that they control over how the personal data is used in the experiment.[23] The data nonsensitivity score measures how sensitive the personal data bundle is for the participant. For each personal data item, weight, height and bank account statement[24], the participants are asked how sensitive the item is on a scale of 1 (most sensitive) to 5 (least sensitive). The data nonsensitivity score corresponds to the sum of the answers, i.e. a higher value indicates a lower overall sensitivity for the data bundle. Figure 12 plots the nonsensitivity by gender. From Figure 12, it is apparent

[23]The variable takes the value one if participants agree strongly, moderately or rather agree that they have control over how the personal data is used in the experiment.

[24]These are the three personal data items which are used for the categorization.

that the sensitivity of the data bundle does not differ importantly across gender.

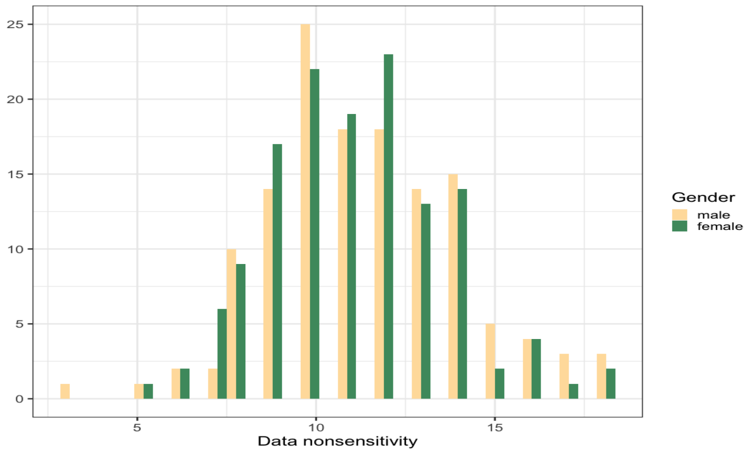

FIGURE 12. NONSENSITIVITY SCORE BY GENDER.

We furthermore control for risk aversion[25], age (self-reported), family income[26], 5 CHF coin stage income and the net reservation price (RP) for the 5 CHF coin, which is calculated as the participant's reservation price for the 5 CHF coin - 5 CHF.

Table 8 shows the estimation results of probit models for the propensity to agreeing to sell the data. The regression results, showing no significant effect of the Treatments (III)-(V), confirm the result from the previous sections that the general willingness to sell the data is not affected by data-based payoff discrimination. Model (2) shows a significant negative impact of the privacy concerns PC on the probability of agreeing to sell the data. However, when including gender and the interaction effects of gender with PC and DTE (Model 5), we can observe that the effect is moderated by gender: while PC is not statistically significant any more when including the interaction effects, we observe a significant negative interaction term of female and PC, suggesting that higher privacy concerns for females lead to a willingness to agree to sell the data. Interestingly, regarding trust towards the experimenters, the interaction term with female is positive, suggesting that for men a higher distrust translates into a lower willingness to agree to sell the data. We also find a significant positive impact of the perceived control over data and the data nonsensitivity, consistent with intuition. Furthermore, self-checking the data significantly increases the probability to agree to sell. This can be interpreted as an effect of risk/ambiguity about the content of the personal data on agreeing to sell. Interestingly, general risk aversion, as measured by the response to the Falk et al. (2016) preference module question on general risk preferences, is not significantly associated to

[25]Risk aversion is measured by the answer to the corresponding Falk et al. (2016) preference module.
[26]The family income is measured by the self-reported allocation into one of four income brackets in the post-experimental questionnaire.

Table 8. Probability of Agreeing to Sell the Personal Data (Probit) (Average marginal effects)

	Dependent variable: Agree to sell the data						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(III)	0.003	0.007	−0.011	−0.008	−0.024	−0.018	−0.020
	(0.253)	(0.256)	(0.260)	(0.263)	(0.268)	(0.286)	(0.294)
(IV)	0.069	0.070	0.067	0.068	0.060	0.035	0.024
	(0.283)	(0.284)	(0.297)	(0.298)	(0.305)	(0.334)	(0.342)
(V)	−0.065	−0.061	−0.059	−0.057	−0.074	−0.078	−0.048
	(0.242)	(0.243)	(0.253)	(0.254)	(0.260)	(0.284)	(0.304)
PC			−0.043***	−0.044***	−0.005	0.0001	−0.014
			(0.071)	(0.072)	(0.102)	(0.107)	(0.114)
DTE			−0.012	−0.012	−0.048	−0.057**	−0.042
			(0.101)	(0.102)	(0.142)	(0.153)	(0.160)
Female		−0.066		−0.065	0.110	0.111	0.086
		(0.190)		(0.197)	(0.669)	(0.708)	(0.759)
Female * PC					−0.080***	−0.064**	−0.043
					(0.151)	(0.160)	(0.169)
Female * DTE					0.083**	0.070*	0.038
					(0.205)	(0.220)	(0.234)
Self-check: any						0.125***	0.114***
						(0.259)	(0.264)
Perceived control over data						0.131**	0.153**
						(0.251)	(0.266)
Data nonsensitivity						0.019**	0.019**
						(0.044)	(0.047)
E-banking usage: false						−0.048	−0.055
						(0.223)	(0.231)
Risk Aversion							−0.001
							(0.058)
Net RP 5 CHF							0.006
							(0.107)
First stage (5 CHF coin) income							−0.011
							(0.061)
Age							−0.006
							(0.037)
Family income							0.018
							(0.112)
McFadden Pseudo R2	0.0182	0.0282	0.0760	0.0873	0.1193	0.2193	0.2352
Observations	271	270	271	270	270	267	262
Log Likelihood	−113.068	−111.756	−106.407	−104.959	−101.272	−89.392	−84.304
Akaike Inf. Crit.	234.136	233.512	224.815	223.918	220.543	204.784	204.609
Note:							*p<0.1; **p<0.05; ***p<0.01

the probability of agreeing to sell the personal data. Furthermore, we do not find an effect of prior experimental income (5 CHF coin income).

We now turn to the second-stage regression. Table 9 shows the OLS estimation results for the reservation price. Model (2), which includes the treatments and gender, confirms our results from Section 3.1: The coefficient of female is positive and statistically significant. We also find that the RP is weakly significantly lower in the Strong Negative Discrimination treatment (V) compared to the nondiscrimination treatments (I+II). As seen in Section 3.1, this is due not only to the higher RP of revealed and unrevealed category C types, but also due to a higher RP of revealed category B types. Furthermore, we find a significant positive effect of DTE on the RP for men. General privacy concerns as captured by PC are the driver of an increase of the RP for women. Thus, the gender-specific impact of privacy concerns and distrust towards the experimenters not only affects the general willingness to agree to sell the data, as seen in Table 8, but also the reservation price for the data. Interestingly, both perceived control and data nonsensitivity, which had a significant association with agreeing to sell the data at all, are not significant determinants of the reservation price for the data. The net RP 5 CHF, but not the first stage income, however is significantly positively related to the data reservation price. One interpretation of the net RP for the 5 CHF is that of a gambling motive: Although it is a dominant strategy to state the true valuation, participants in the 5 CHF selling and the data selling might try to state a high price to receive a higher selling gain. Importantly, all the effects discussed above, most importantly that of Treatment (V), are unaffected by controlling for this possible gambling motive.

The results from the regression analyses including privacy-related attitudes derived from the survey answers confirm the overview results from Section 3.1. In particular, subsequent data-based payoff discrimination does not affect the general willingness to agree to sell personal data, and a (weak) significant change in the reservation price is only observed when there is a strong negative payoff adjustment for one category. The regression analysis using the survey answers furthermore highlights the important gender-differentiated effect of privacy concerns—which are relevant for females— and distrust towards the experimenters—more relevant for males—for both the willingness to agree to sell the data as well as the reservation price.

Table 9. Reservation Price for Personal Data (OLS)

	Dependent variable: Reservation price of Personal Data				
	(1)	(2)	(3)	(4)	(5)
(III)	−3.971	−4.611	−4.326	−3.936	−4.333
	(3.253)	(3.157)	(3.136)	(2.992)	(3.005)
(IV)	3.714	2.916	4.267	3.735	2.885
	(3.234)	(3.126)	(3.117)	(2.951)	(3.001)
(V)	6.567*	5.831*	5.269	6.129*	5.421*
	(3.379)	(3.264)	(3.295)	(3.145)	(3.195)
PC			1.566*	−0.030	−0.172
			(0.847)	(1.047)	(1.048)
DTE			3.443**	5.405***	5.593***
			(1.368)	(1.702)	(1.703)
Female		9.359***		3.382	2.096
		(2.331)		(6.567)	(6.576)
Female * PC				4.929***	4.308***
				(1.650)	(1.632)
Female * DTE				−6.393**	−5.332**
				(2.593)	(2.593)
Self-check: any					−1.247
					(2.331)
E-banking usage: false					−3.003
					(2.385)
Perceived control over data					−0.367
					(2.857)
Data nonsensitivity					−0.129
					(0.434)
Risk aversion					−0.213
					(0.577)
Net RP 5 CHF					3.855***
					(1.090)
First stage (5 CHF coin) income					−0.473
					(0.644)
Age					0.308
					(0.384)
Constant	24.094***	19.947***	10.107***	7.918*	7.468
	(2.050)	(2.231)	(3.805)	(4.709)	(12.333)
Observations	230	229	230	229	226
R^2	0.040	0.109	0.118	0.223	0.284
Adjusted R^2	0.028	0.093	0.099	0.194	0.229

Note: *p<0.1; **p<0.05; ***p<0.01

4 Conclusion

We provide the results of an experiment on the value of personal data in light of data-based price discrimination. We find that in our experiment the general willingness to sell personal data is not significantly affected by subsequent data-based payoff discrimination. Furthermore, we only find a significant change in the reservation price of the data when one data-based category implies a strong decrease of the subsequent payoff. Interestingly, the change in the reservation price is not only driven by participants who fall into this category, but a general increase of the reservation price under strong negative price discrimination. We observe this effect even though the data-based categorization is fully known and risk/ambiguity-free.

The bundle of personal data that the participants can sell consists of their height, weight, gender, information on their bank account balance and a photo of their face. A comparison of the self-reported sensitivity of this personal data bundle shows that there are no important differences across gender how sensitive this data is perceived to be. Nevertheless, we find important gender differences in how general privacy concerns and trust related to the context of the experiment affect both the general willingness to sell the data as well as the reservation price of the data. These findings are important in light of the consequences of personal data sharing for subsequent market interaction, not only with respect to price discrimination, but the usage of personal data more generally.

In our setting, the rules of how the personal data is used were fully known and transparent: Both the data-based categorization was known to participants and exogenous, as well as the attached payoff consequences. When sharing personal data in online markets, however both how this data is interpreted, e.g. by algorithms, as well as the (financial) consequences are much less transparent, and often fully ambiguous. It is an interest avenue for future research to analyze how ambiguity about the informational content of personal data and the attached consequences affect the willingness to share this data.

References

Acquisti, A., L. K. John, and G. Loewenstein (2013): "What is privacy worth?" *The Journal of Legal Studies*, 42, 249–274.

Acquisti, A., C. Taylor, and L. Wagman (2016): "The economics of privacy," *Journal of Economic Literature*, 54, 442–92.

Barth, S. and M. D. de Jong (2017): "The privacy paradox–Investigating discrepancies between expressed privacy concerns and actual online behavior–A systematic literature review," *Telematics and Informatics*.

Becker, G. M., M. H. DeGroot, and J. Marschak (1964): "Measuring utility by a single-response sequential method," *Behavioral Science*, 9, 226–232.

Belleflamme, P. and W. Vergote (2016): "Monopoly price discrimination and privacy: The hidden cost of hiding," *Economics Letters*, 149, 141–144.

Benndorf, V. and H.-T. Normann (2017): "The willingness to sell personal data," *The Scandinavian Journal of Economics*.

Beresford, A. R., D. Kübler, and S. Preibusch (2012): "Unwillingness to pay for privacy: A field experiment," *Economics letters*, 117, 25–27.

Choe, C., S. King, and N. Matsushima (2019): "Pricing with Cookies: Behavior-Based Price Discrimination and Spatial Competition," *Management Science*, 64, 5669–5687.

Costello, A. B. and J. W. Osborne (2005): "Best practices in exploratory factor analysis: Four recommendations for getting the most from your analysis," *Practical assessment, research & evaluation*, 10, 1–9.

Crainich, D. (2017): "Self-Insurance With Genetic Testing Tools," *Journal of Risk and Insurance*, 84, 73–94.

Crocker, K. J. and A. Snow (2013): "The theory of risk classification," in *Handbook of insurance*, Springer, 281–313.

Dionne, G. and C. Rothschild (2014): "Economic effects of risk classification bans," *The Geneva Risk and Insurance Review*, 39, 184–221.

Ezrachi, A. and M. E. Stucke (2016): "The rise of behavioural discrimination," *European Competition Law Review*, 37, 484–491.

Falk, A., A. Becker, T. Dohmen, D. Huffman, and U. Sunde (2016): "The preference survey module: A validated instrument for measuring risk, time, and social preferences," .

Feri, F., C. Giannetti, and N. Jentzsch (2016): "Disclosure of personal information under risk of privacy shocks," *Journal of Economic Behavior & Organization*, 123, 138–148.

Fischbacher, U. (2007): "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental economics*, 10, 171–178.

Grether, D. M. and C. R. Plott (1979): "Economic theory of choice and the preference reversal phenomenon," *The American Economic Review*, 69, 623–638.

Holt, C. A. and S. K. Laury (2002): "Risk aversion and incentive effects," *American economic review*, 92, 1644–1655.

Honaker, J., G. King, M. Blackwell, et al. (2011): "Amelia II: A program for missing data," *Journal of statistical software*, 45, 1–47.

Kehr, F., T. Kowatsch, D. Wentzel, and E. Fleisch (2015): "Blissfully ignorant: the effects of general privacy concerns, general institutional trust, and affect in the privacy calculus," *Information Systems Journal*, 25, 607–635.

Kern, J., B. Fabian, and T. Ermakova (2018): "Experimental Privacy Studies-An Initial Review of the Literature," .

Kvaløy, O., M. Luzuriaga, and T. E. Olsen (2017): "A trust game in loss domain," *Experimental Economics*, 20, 860–877.

Malhotra, N. K., S. S. Kim, and J. Agarwal (2004): "Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model," *Information systems research*, 15, 336–355.

Marreiros, H., M. Tonin, M. Vlassopoulos, and M. Schraefel (2017): ""Now that you mention it": A survey experiment on information, inattention and online privacy," *Journal of Economic Behavior & Organization*, 140, 1–17.

Montes, R., W. Sand-Zantman, and T. Valletti (2018): "The Value of Personal Information in Online Markets with Endogenous Privacy," *Management Science*.

Muermann, A., A. Geyer, and D. Kremslehner (2019): "Asymmetric Information in Automobile Insurance: Evidence from Driving Behavior," *Journal of Risk and Insurance*, 1–27.

New York Times (2018): "How Trump consultants exploited the Facebook data of millions," https://www.nytimes.com/2018/03/17/us/politics/cambridge-analytica-trump-campaign.html, last accessed on Sep 30, 2018.

Plesch, J. and I. Wolff (2018): "Personal-Data Disclosure in a Field Experiment: Evidence on Explicit Prices, Political Attitudes, and Privacy Preferences," *Games*, 9.

Regner, T. and G. Riener (2017): "Privacy is precious: On the attempt to lift anonymity on the internet to increase revenue," *Journal of Economics & Management Strategy*, 26, 318–336.

Richards, T. J., J. Liaukonyte, and N. A. Streletskaya (2016): "Personalized pricing and price fairness," *International Journal of Industrial Organization*, 44, 138–153.

Schudy, S. and V. Utikal (2017): "'You must not know about me'—On the willingness to share personal data," *Journal of Economic Behavior & Organization*, 141, 1–13.

Smith, H. J., S. J. Milberg, and S. J. Burke (1996): "Information privacy: measuring individuals' concerns about organizational practices," *MIS quarterly*, 167–196.

World Health Organization (2018a): "Global Health Observatory (GHO) data," http://www.who.int/gho/ncd/risk_factors/bmi_text/en/, last accessed on Oct 24, 2018.

——— (2018b): "Mean body mass index trends among adults, crude Estimates by country," http://apps.who.int/gho/data/view.main.BMIMEANADULTCv?lang=en, last accessed on Oct 24, 2018.

A Additional Figures and Tables

A.1 Results of directional nonparametric tests

Table 10. Results of directional nonparametric tests. Dependent variable: price of informational bundle (not adjusted by 60)

Hypothesis		all observations		
			female	male
(V) ≥ (I + II)	mean difference	7.12	4.71	8.02
	effect size	0.36	0.22	0.49
	MWU	0.01***	0.27	0.02**
	KS test	0.01***	0.13	0.06*
	Permutation Test	0.02**	0.17	0.03**
	Bootstrap	0.01***	0.15	0.04**
(IV) ≥ (I + II)	mean difference	3.64	-0.79	6.45
	effect size	0.19	-0.04	0.40
	MWU	0.13	0.58	0.07*
	KS test	0.24	0.79	0.17
	Permutation Test	0.14	0.56	0.07*
	Bootstrap	0.14	0.57	0.07*
(I + II) ≥ (III)	mean difference	4.04	8.53	1.42
	effect size	0.23	0.47	0.08
	MWU	0.11	0.06*	0.26
	KS test	0.44	0.22	0.66
	Permutation Test	0.12	0.06*	0.36
	Bootstrap	0.11	0.06*	0.38
(V) ≥ (III)	mean difference	11.16	13.24	9.44
	effect size	0.62	0.74	0.55
	MWU	0.00***	0.00***	0.02***
	KS test	0.00***	0.02**	0.02**
	Permutation Test	0.00***	0.00***	0.03**
	Bootstrap	0.00***	0.00***	0.03**
(IV) ≥ (III)	mean difference	7.68	7.74	7.87
	effect size	0.43	0.43	0.45
	MWU	0.02**	0.07*	0.05**
	KS test	0.13	0.11	0.10*
	Permutation Test	0.02**	0.07*	0.05**
	Bootstrap	0.02**	0.07*	0.05**
(V) ≥ (IV)	mean difference	3.48	5.50	1.57
	effect size	0.19	0.29	0.09
	MWU	0.13	0.13	0.35
	KS test	0.30	0.33	0.76
	Permutation Test	0.16	0.12	0.37
	Bootstrap	0.16	0.13	0.38

Table 11. Results of directional nonparametric tests. Dependent variable: price of informational bundle. Category B

Hypothesis		Category B		
			female	male
(III) $\leq$ (I + II)	mean difference	2.31	-0.95	4.26
	effect size	0.21	-0.09	0.38
	MWU	0.28	NA	0.23
	KS test	0.68	NA	0.47
	Permutation Test	0.28	NA	0.20
	Bootstrap	0.24	NA	0.18
(I + II) vs (IV)	mean difference	5.48	7.55	4.74
	effect size	0.46	0.76	0.36
	MWU	0.28	0.55	NA
	KS test	0.87	0.71	NA
	Permutation Test	0.16	0.17	NA
	Bootstrap	0.14	0.15	NA
(I + II) vs (V)	mean difference	12.00	24.84	2.31
	effect size	1.00	2.49	0.17
	MWU	0.04**	0.03**	0.54
	KS test	0.14	0.04**	0.90
	Permutation Test	0.01***	0.00***	0.34
	Bootstrap	0.01***	0.00***	0.34
(III) $\leq$ (V)	mean difference	14.31	23.90	6.57
	effect size	1.33	2.30	0.58
	MWU	0.01***	NA	0.06*
	KS test	0.03**	NA	0.17
	Permutation Test	0.01***	NA	0.11
	Bootstrap	0.01***	NA	0.10*
(III) $\leq$ (IV)	mean difference	7.79	6.61	9.00
	effect size	0.72	0.63	0.80
	MWU	0.09*	NA	NA
	KS test	0.39	NA	NA
	Permutation Test	0.09*	NA	NA
	Bootstrap	0.08*	NA	NA
(IV) vs (V)	mean difference	6.52	17.29	-2.43
	effect size	0.45	0.97	-0.23
	MWU	0.46	0.17	NA
	KS test	0.67	0.30	NA
	Permutation Test	0.18	0.07	NA
	Bootstrap	0.19	0.06	NA

Table 12. Results of directional nonparametric tests. Category C

Hypothesis		Category C		
			female	male
(I + II) $\leq$ (V)	mean difference	14.50	11.55	15.88
	effect size	1.12	0.76	1.57
	MWU	0.00***	NA	0.00***
	KS test	0.00***	NA	0.00***
	Permutation Test	0.00***	NA	0.00***
	Bootstrap	0.00***	NA	0.00***
(I + II) $\leq$ (IV)	mean difference	6.15	5.78	2.47
	effect size	0.48	0.38	0.24
	MWU	0.09*	0.17	0.36
	KS test	0.19	0.36	0.59
	Permutation Test	0.07*	0.21	0.29
	Bootstrap	0.0*7	0.22	0.26
(III) vs (I + II)	mean difference	0.58	5.05	1.40
	effect size	0.04	0.35	0.13
	MWU	NA	NA	NA
	KS test	NA	NA	NA
	Permutation Test	0.44	0.22	0.40
	Bootstrap	0.42	0.21	0.40
(III) $\leq$ (IV)	mean difference	6.73	10.83	3.87
	effect size	0.50	0.75	0.35
	MWU	0.10*	0.05**	0.29
	KS test	0.29	0.18	0.58
	Permutation Test	0.08*	0.05**	0.26
	Bootstrap	0.07*	0.04**	0.23
(III) $\leq$ (V)	mean difference	15.08	16.60	17.29
	effect size	1.12	1.15	1.58
	MWU	0.00***	NA	0.01***
	KS test	0.00***	NA	0.01***
	Permutation Test	0.00***	NA	0.01***
	Bootstrap	0.00***	NA	0.02**
(IV) $\leq$ (V)	mean difference	8.35	5.77	13.41
	effect size	0.54	0.37	1.10
	MWU	0.08*	NA	0.02**
	KS test	0.12	NA	0.05**
	Permutation Test	0.06*	NA	0.02**
	Bootstrap	0.06*	NA	0.02**

A.2 Exploratory factor analysis

The latent traits that we try to measure are inextricably linked and boundaries between them can be quite elusive. For many items in our survey it is difficult to provide a definite answer what latent trait it is related to. For instance, positive answers to question GP.6[27] might be affected by both, lack of trust towards the companies and/or high level of privacy concerns. Instead of entering a lengthy discourse on the nature and relevance of various survey items for individual's latent characteristics, we take a more pragmatic approach and perform an exploratory factor analysis whereby letting the data speak for itself. Our general knowledge can be used later to check the validity of the conclusions and select the most promising measurement model. To perform an exploratory factor analysis, we need to provide a correlation matrix, select the factor analysis method and determine the number of factors to extract. The following paragraphs elaborate on each of these aspects.

Almost all survey items are measured using different types of Likert scales. These scales are conventional practice in social sciences and psychology, however they are prone to several behavioral biases that should be kept in mind.[28] We treat answers from the Likert scale as ordinary variables[29] and compute Spearman's rho and Kendall's tau correlation coefficients.

We extract common factors using the *weighted least squares* (WLS) procedure, since it relies on few distributional assumptions and has good numerical convergence properties.[30] The last required input parameter is the number of factors to be extracted. Costello and Osborne (2005) list various procedures, that could facilitate the choice.[31] Recommenda-

[27]"I'm concerned that companies collect too much information about me"

[28]The composition of a Likert scale makes the answers susceptible to an acquiescence bias, when participants tend to agree with most of the statements. To counter this effect, several reversed items are included. Furthermore, responses might exhibit a central tendency bias with which individuals avoid using extreme answers in fear of not confirming to the general norm. This is a subcase of the social desirability bias. Social desirability biases could further manifest themselves with subjects providing answers that rather reflect society's expectations than their own views.

[29]There is no general consensus on whether answers to a Likert scale should be treated as ordinal or interval variables, however the majority view tends to favor the former option.

[30]Sometimes the estimation algorithm might yield item commonalities equal or exceeding 1. This is the proportion of each variable's variance that can be explained by the factors (e.g., the underlying latent continua). These situations are referred to as Heywood and ultra-Heywood cases respectively and indicate that the resulting factor solution is invalid. Such anomalies could be caused by low sample size, too many factors extracted or other inadequacies in the common factor model. After trying different factoring method options available in the *psych* package in R, we have decided to extract factors using the weighted least squares approach, since it did not produce ultra-Heywood cases despite our low sample size. Note that the results of the inference will differ depending on the type of correlation coefficient used.

[31]An important pragmatic consideration is, that due to a limited sample size, the number of independent variables in the regression equations should not be too high, furthermore having too few observed variables per latent factor might lead to unstable solutions. Due to these considerations we have decided to restrict our search to models with no more than 10 latent variables.

tions based on various criteria are summarized in Table 13.

correlation coefficient	parallel	vss1	vss2	eBIC	SABIC
Spearman	9	2	9	9	10
Kendall	9	2	7	6	8

Table 13. Optimal number of factors based on different criteria for Spearman and Kendall correlation matrices.

The optimal number of factors varies significantly depending on the type of correlation coefficient and the measure chosen. Instead of selecting a single approach, we have decided to perform an exploratory factor analysis using all suggested parameters and aggregate the results. Our objective is to see, which survey items are more likely to assess the same latent trait. This is achieved using the following steps: for every model we compute a binary matrix, where rows and columns correspond to observed variables. An entry equals 1 if both variables have an absolute factor loading exceeding 0.3 for the same latent variable. All matrices are summed together, rows are divided by the size of the diagonal element, representing for how many latent variables the questions is used, and then substracted from a matrix of ones. The resulting object could be viewed as a distance matrix, where variables measuring similar characteristics are closer to one another.

Visual representation of this information is obtained using multidimensional scaling. The results are depicted in Figure 13 from which we see that certain groups of variables, for instance, PB.4 - PB.8, SN.1, SN.8, and SN.7, are almost always used together to assess a latent factor. Some of the variable clusters confirm our earlier intuition: items VP.1 - VP.3 measure the same latent characteristics - how much an individual values their privacy. Other suggested combinations, such as SN.9, SN.10, I.16, and I.17 are not anticipated. Post factum, the logic behind grouping them together becomes more apparent. Certain questions seem to be only loosely related to the common factors in the data set, this however does not render them invaluable for the subsequent analysis. Question PB.1, for instance, is directly related to the effort cost of providing information about the bank account.

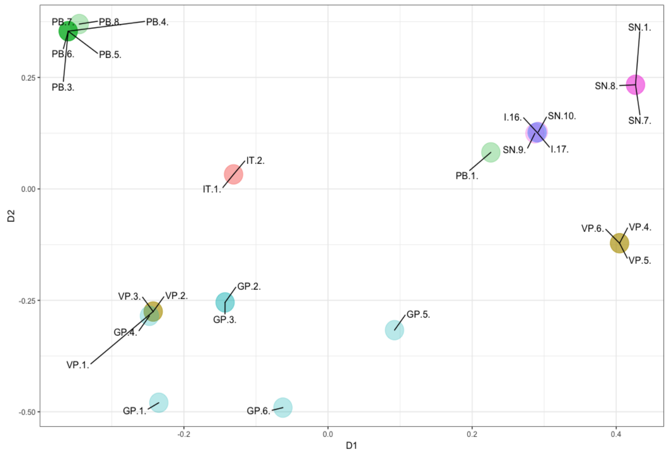

Figure 13. Multidimensional scaling of observed variables. Based on 224 observations

A.3 Confirmatory factor analysis

Based on the insights from Figure 13 and a literature review, we can start developing our measurement model and study its psychometric properties. Table 14 lists various groupings of variables suggested by the exploratory factor analysis, whereby each group of questions measures one latent characteristic. For each instrument proposed, we first ascertain its reliability by computing Cronbach's alpha and split half reliability. Values exceeding 0.7 are considered acceptable.

Group	Variables	Reliability		
		alpha	best alpha	split half
group 1	VP.1, VP.2, VP.3, GP.4	0.838	0.871	0.850
group 2	IT.1, IT.2	0.877	0.787	0.880
group 3	VP.4, VP.5, VP.6	0.542	0.565	0.532
group 4	PB.3, PB.4, PB.5, PB.6, PB.7, PB.8	0.743	0.736	0.567
group 5	SN.9, SN.10, I.16, I.17	0.419	0.429	0.453
group 6	SN.1, SN.7, SN.8	0.538	0.549	-0.285
group 7	GP.1, GP.2, GP.3, GP.5, GP.6	0.749	0.736	0.768
group 7 (full)	GP.1, GP.2, GP.3, GP.4, GP.5, GP.6	0.785	0.777	0.803

Table 14. Grouping of the survey questions

It reveals that group 3, group 5 and group 6 do not possess desirable psychometric properties and consequently won't be considered in further analyses. The reliability of group 1 can be slightly improved from 0.835 to 0.863 if the variable GP.4 is dropped. The confirmatory factor analysis is based on a correlation matrix with Spearmann's correlation coefficients. To ensure that the model is identified, we fix the variances of latent variables to 1. Fitting a model with 3 latent variables results in a poor model fit, with Chi = 101.054 on 39 df suggesting significant differences between model-based and observed correlation matrices. To achieve a better model fit, we drop the third latent variable. This decision should not have a strong negative impact on the explanatory power of our regression models since group 1 and group 4 measure similar latent characteristics (which is further confirmed by an estimated covariance of 0.536). The model with two latent variables has a good fit, reflected by a non-significant Chi-square p-value of 0.132, and value of RMSEA = 0.053. Path coefficients and relevant statistics are reported in Table 15. Figure 14 visualizes the final measurement model.

	Model			
	Estimate	Std. Err.	z	p
		Factor Loadings		
PC				
VP.1.	0.779	0.047	16.676	0.000
VP.2.	0.787	0.047	16.703	0.000
VP.3.	0.923	0.044	20.903	0.000
TE				
IT.1.	0.89	0.06	15.96	0.00
IT.2.	0.87	0.06	15.48	0.00
		Residual Variances		
VP.1.	0.26	0.03	9.27	0.00
VP.2.	0.27	0.03	9.25	0.00
VP.3.	0.06	0.02	2.63	0.01
IT.1.	0.14	0.06	2.27	0.02
IT.2.	0.19	0.06	3.07	0.00
		Latent Variances		
PC	1.00^{+}			
DTE	1.00^{+}			
		Latent Covariances		
PC w/TE	0.47	0.05	9.10	0.00
		Fit Indices		
χ^2	1.04(df=4)			0.132
CFI	0.997			
TLI	0.992			
RMSEA	0.053			

$^{+}$Fixed parameter

TABLE 15. COEFFICIENTS OF FITTED CFA MODEL

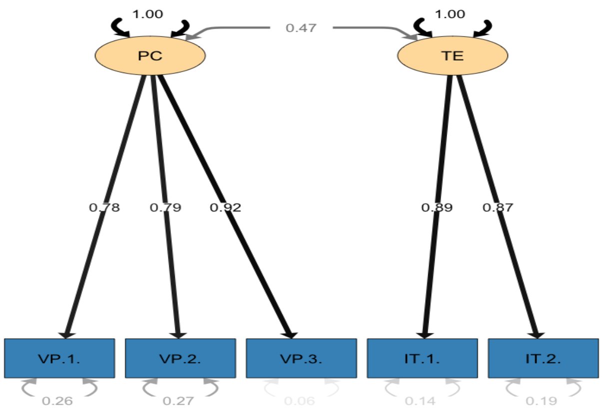

FIGURE 14. DIAGRAM OF FITTED CFA MODEL

Based on the model fit, we can check other psychometric properties of our scales. Table 16 contains all necessary information to check the composite reliability and convergent/discriminant validity of survey items. Both latent factors have a composite reliability exceeding 0.9. Convergent validity can be established if the average variance extracted

	CR	AVE	PC	DTE
PC	0.912	0.778	0.882	0.472
DTE	0.905	0.827	0.472	0.909

Table 16. Check for composite reliability, discriminant validity and Convergent Validity

(AVE) is higher than 0.5. Discriminant validity is checked using the Fornell–Larcker criterium that compares the square root of the average variance extracted with correlations between the latent factors. If the former is higher, discriminant validity can be established.

The objective of our analysis is to estimate the values of latent characteristics for all study participants. However, several subjects did not provide answers to all survey questions as this part of the experiment is neither mandatory nor financially compensated. Out of 282 people, 4 did not provide answers to one of the questions VP.1 - VP.3 and 2 other participants did not answer one of the questions IT.1 - IT.2. Due to the small sample size, we want to avoid losing observations and impute missing answers based on the information from other questions in the post-experimental survey. This task is accomplished using a combination of bootstrap and the EM algorithm (Honaker et al. (2011)) implemented in the R package *Amelia*.