Unpacking Privacy: Willingness to pay to protect personal data

Anya Skatova^{1,2,3}, Rebecca McDonald^{4,2}, Sinong Ma³ and Carsten Maple³

¹School of Psychological Science, University of Bristol

²Warwick Business School, University of Warwick

³Warwick Manufacturing Group, University of Warwick

⁴School of Economics, University of Birmingham

Abstract

Data is key for the digital economy, underpinning business models and service provision, and a lot of these valuable datasets are personal in nature. Information about individual behaviour is collected regularly by organisations. This information has value to businesses, the government and third parties. It is not clear what value this personal data has to consumers themselves. Much of the digital economy is predicated on people sharing personal data, however if individuals value their privacy, they may choose to withhold this data unless the perceived benefits of sharing outweigh the perceived value of keeping the data private. Further, they might be willing to pay for an otherwise free service if paying allowed them to avoid sharing personal data. We used five evaluation techniques to study preferences for protecting personal data online and found that consumers assign a positive value to keeping a variety of types of personal data private. We show that participants are prepared to pay different amounts to protect different types of data, suggesting there is no simple function to assign monetary value that can be identified for individual privacy in the digital economy. The majority of participants displayed remarkable consistency in their rankings of the importance of different types of data, a finding that indicates the existence of stable individual privacy preferences in protecting personal data. We discuss our findings in the context of research on the value of privacy and privacy preferences, and in terms of implications for future business models and consumer protection.

Introduction

With the rising prominence of digital technology in the modern economy, the amount of data that is recorded about consumers' behaviour and that is retained in various virtual and physical places is enormous. This personal data can reveal private details about consumers' everyday behaviour. For example, records of individuals' daily journeys are reflected in

mobile phone or car GPS data; household composition can be deduced from electricity use in the home; loyalty card data can reflect personal lifestyle choices; and bank transaction data can reveal details of personal financial circumstances. These opportunities are increasingly being exploited in new and innovative ways: personal data has economic value for government organisations and businesses (e.g., in advertising, Shen and Villas-Boas 2018). It enables a diverse collection of digital economy services to be provided, serving a foundation for many successful business models (HM Treasury 2018, see Acquisti et al. 2016 for a review of the role of personal data in a variety of business models).

A fundamental mechanism underpinning the functioning of the digital economy is data sharing, whereby individual consumers are assumed to allow their data to be collected, so that it can be processed and analysed to support the provision of services. When data is shared, it can be used for the benefit of consumers. For example, Google Maps collects information about people's journeys, and this is used to create accurate travel maps and provide better estimates for journey times, reflecting current traffic conditions. If consumers stop sharing their data, many of these digital economy services may no longer be viable. However, the data individuals share can also be used to consumers' detriment (e.g. through unwanted personalised advertising or, in extreme cases, data breaches). In recognition of this possibility, the General Data Protection Regulation (GDPR) has been developed and enshrined in law in the EU, including the UK. Whilst effectiveness of GDPR is yet to be assessed (see Goldfarb and Tucker 2011 on effectiveness of privacy laws on online advertising, and Adjerid et al 2015, Miller and Tucker 2017 on health data sharing), this legislation puts greater emphasis on the transparency with which companies ought to deal with consumer personal data, aiming to give control to individuals over their personal data, and creating additional mechanisms to already existing ones (e.g., in the UK: Information Commissioner Office) to be able to financially penalise those companies that do not comply with protection of personal data. However, since little is known about what drives individuals'

decisions about whether to share their data online, little is known about the extent to which GDPR will influence the functionality of the digital economy.

Our research extends previous work on factors affecting individuals' decisions about whether to share personal data. We take an experimental approach and focus on whether consumers place a positive value on protecting their data by examining their willingness to share personal data in a variety of data sharing environments, including their Willingness To Pay (WTP) to protect it. Using five evaluation techniques, we systematically investigate whether members of the public value keeping their personal data private. Specifically, we determine (1) whether stated preferences for keeping personal data private are stable within individuals; and (2) whether stated preferences for keeping personal data private systematically vary between data sharing environments (e.g., electricity consumption vs physical activity monitoring vs spending patterns).

We find that 96% of individuals are willing to pay to avoid sharing their personal data in at least one of data sharing environments. We also find that individuals are consistent in their rankings of the importance of different types of personal data when these rankings are elicited in different ways. This suggests that the techniques we employ can reveal stable underlying preferences for keeping personal data private. We find consistent differences across our sample in their stated preferences for keeping different types of personal data private: financial and medical data were typically rated as the most important to protect, while supermarket loyalty cards, electricity use and physical activity data were rated as the least important to protect. We conclude that stability in stated preferences for keeping personal data private across a variety of elicitation methods is evidence that privacy preferences are well defined. If so, the stated preferences we elicit may help to explain how individuals may decide whether to share their personal data in the context of the digital economy. Finally, we draw conclusions on the links between our findings on willingness to pay to protect personal data in different environments with the literature on the value of privacy.

Literature

Economists have been long interested in how personal information is shared and, broadly speaking, how it is valued. Sharing personal information can help or hinder the efficiency of a variety of markets, not just digital economy markets. For instance, Stigler (1980) and Posner (1981) set out models in which individuals can hide personal characteristics about themselves, adding friction to insurance markets. A comprehensive review of the various economic theory perspectives on privacy is provided by Acquisti et al. (2015) and by Acquisti et al. (2016). This work brings into focus consumers' views on the consequences of their privacy-related behaviours and measures their preferences regarding privacy. The authors also explore the extent to which privacy preferences are malleable, calling into question whether they exist at all. Acquisti et al. (2016) demonstrate that there is no clear consensus about the welfare effects of information being withheld in an economic setting.

In contrast to the literature exploring the consumer perspective discussed so far, economic models of privacy have taken the perspective of the firm (e.g., see Casadesus-Masanell and Hervas-Drane 2015), by discussing the price discrimination that can arise when information about consumers is disclosed to, or otherwise obtained by, firms. This literature is concerned with modelling firms' behaviour in the market for goods or services, given different degrees of access to personal data. It particularly focuses on price discrimination, targeted advertising, or customised products (Montes et al. 2017). This approach usually involves the assumption that consumers' utility depends on the price and quantity of the goods that they buy and (sometimes) in the cost of foregone privacy (e.g. Montes et al. 2018), which is defined as a parameter that is typically exogenous to the model. Consumers' behaviour is assumed to maximise their utility and the welfare effects in these markets are typically defined in terms of consumers' utility and firms' profits at equilibrium.

More broadly, scholars in a range of disciplines have investigated the value of personal data (e.g., Spiekermann and Korunovska 2017, Benndorf and Normann 2017, Beresford et al. 2012, Acquisti et al. 2013), which we discuss in more details below. An individual deciding to

share data online may be compromising their privacy, since this data can unknowingly reveal private information about them (e.g., their health status). Therefore, an important factor influencing an individual's preference for protecting their personal data is how much they value their privacy, captured by their preference for privacy. We note that privacy preferences are a feature of some theoretical models of online privacy including Montes et al. (2017). Relatedly, the well-established idea of the "privacy calculus" (Laufer and Wolfe 1977) is that individuals assess the costs and benefits of protecting (or sharing) their information - essentially asking themselves whether it is "worth it" to reveal a certain type of information. To engage in such a calculus, individuals must have defined preferences over their privacy. However, some studies raise doubts about the existence of such preferences (e.g. Tucker 2014). We review the literature regarding the existence and nature of preferences for privacy in what follows.

An early empirical investigation into the value of privacy was undertaken by Hann et al. (2002), who conducted a Discrete Choice Experiment to investigate the trade-offs between privacy and benefit (financial and time saving) in the choice between two websites. The authors demonstrated that their sample of students were indeed willing to trade off their privacy for sufficient economic benefit. Huberman et al. (2005) explored individuals' willingness to accept compensation for revealing sensitive personal information, specifically one's bodyweight. They showed that people were willing to engage in such a trade, but that those who perceived their weight to be different from the norm demanded greater compensation for sharing this data. This again suggests that individuals are entering into a form of privacy calculus. Rose (2005) conducted an investigation into willingness to pay higher taxes for "stronger data protection laws" in New Zealand, concluding that people were willing to pay to for this type of general increased privacy.

Since these early investigations, there has been ongoing interest in attempting to value the protection of privacy online. A systematic review of the literature regarding "putting a price tag on personal information" was conducted by Wagner et al. (2018). They found that the

type of information being disclosed (or protected) was an important feature in determining the amount that individuals were willing to pay (or accept). The second feature that they discovered to be important was the privacy impact of personal information disclosure. Specifically, they found that:

“the more sensitive the data and the more identifiable people are, the higher has been the price people attach to their data as they perceive higher risks” (Wagner et al. 2018, p 3766).

As we have seen, there is substantial variation in the evidence about willingness to pay for protection of personal data. This may reflect the lack of any clear understanding amongst the participants - and arguably amongst the researchers - about exactly what they mean by privacy. Many of the articles refer to the value of privacy, but actually test willingness to pay not to share (or, willingness to accept compensation for sharing) a particular type of personal data (e.g. Jentzsch 2014 for the results of a logic test; Schreiner and Hess 2015 for social media details). This means that the value attributed to privacy will be confounded by all the other reasons that someone may wish to avoid sharing that particular type of personal data. Our view is that privacy relates directly to what sharing a piece of data can reveal about the person's behaviour. This will differ on many dimensions, including context, who is requesting the data, and on the specific data type. As such, to understand how individuals value their privacy, a first step is to consider their willingness to share a variety of different types of personal data.

However, while the values of different data types might vary, there still could be underlying individual privacy preferences. Many of the studies to date that explore the value of privacy (through eliciting the value of personal data) focus on just one type of data (e.g., demographic information, Beresford et al. 2012). In this case the consistency and stability of preferences for protecting personal data cannot be investigated, and the robustness of any conclusions about how much individuals value personal data in general cannot be ascertained. Exceptions to this mainly investigate the comparison between willingness to

accept and willingness to pay (e.g. Acquisti et al. 2013) or between different types of willingness to pay (e.g. Benndorf and Normann, 2014) but still with the same pieces of personal information. However, different types of data - for example energy use at home versus bank transactions - can reveal different information about an individual, and these different revelations may be sensitive to different extents. There is evidence that people care about some types of data more than others, and that people differ with respect to which types of data they care about (Skatova et al. 2013, Goldfarb and Tucker 2012). Arguably, then, the value of personal data is not a unified concept with a single value but is context-dependent and will vary by different types of data (Acquisti et al. 2016). Further, consumers with the same underlying privacy preferences might make different decisions while sharing different types of data because the assessment of risks and benefits will be different in each data sharing environment.

To summarise this literature, it appears that individuals are often willing to give up privacy for benefits defined in terms of money and/or convenience, and that their willingness to engage in such trade-offs may differ between contexts where the risks and benefits of the trade-offs vary, consistent with the idea that privacy preferences exist, that they depend on the context, and that they are inputs to the decision whether to share personal data online. However, an omission in the current literature is evidence about the stability and consistency with which individuals can express these preferences. If individuals' stated preferences regarding their personal data are not consistent, stable and well-defined, then behaviourally-relevant models, in which online behaviour is governed in part by privacy preferences, will be intractable. We next describe in more detail the one paper that currently exists that looks at the stability of privacy preferences and draw on related literature in other contexts to ascertain how best to capture the stability - or otherwise - of preferences.

Preference elicitation across contexts

To understand whether people can consistently report their preferences, it is necessary to elicit these preferences in different ways. To our knowledge just one paper has compared

different methods of preference elicitation in the context of online privacy. Benndorf and Normann (2014) investigated the effect of two different methods in eliciting the value of personal information, comparing a Becker-DeGroot-Marshak style second price auction (Becker et al. 1964) with a Take it or Leave it approach. They found a significant difference in values between these approaches, even though both aimed to elicit the same underlying construct. Hence, they provided preliminary evidence that preferences to keep personal data private - at least in terms of their monetary valuations - are subject to elicitation method effects. To our knowledge no research has yet considered the (in)consistency with which individuals express their preference ranking over the importance of protecting different types of personal data, and whether those rankings depend on elicitation methods. However, evidence about the stability and reliability of preferences expressed through different elicitation procedures has been considered in detail in other fields, most notably in experimental economics and in the non-market valuation literature applied to health outcomes.

Methods that exist to categorise and compare the strength of preference over different outcomes come in many forms, from simple choice tasks through to complex trade-off tasks between the outcome and some numeraire of its value (for reviews comparing methods in the health context, see Froberg and Kane 1989 or Nord 1991). One important way to categorise these methods is by their relative or absolute nature. Relative tasks allow a researcher to elicit the order - and sometimes the strength of preference - for one outcome compared to another. Examples include direct ranking tasks, where a participant must rank outcomes in order of preference, and binary choice tasks where one option is compared to another and the preferred option is chosen. Absolute tasks require participants to reveal the strength of preference for an outcome by itself without reference to any alternatives. For example, with rating scales an outcome is rated in terms of its quality, likeability or importance; while in valuation tasks an outcome is judged according to its monetary value.

According to standard economic theory, as long as individuals hold well-defined and consistent preferences over the outcomes in question, the type of task used to elicit these preferences should not influence the outcome of the preference elicitation. Put simply, if I like an apple more than an orange, then I would choose an apple over an orange in direct choice, and I would be willing to pay more for the apple than I would be willing to pay for the orange. However, a long and rich literature in experimental economics and psychology has revealed that choice and judgement tasks follow different psychological processes (e.g. Rosen and Rosenkoetter 1976) and prioritise different attributes of a comparison (e.g. Tversky et al. 1988). It is therefore an open question whether individuals can consistently reveal their preference ordering with respect to the importance of protecting different types of their personal data. Arguably, the more well-defined their underlying preferences, the more stable the preference ordering will be when compared across different elicitation tasks.

The health economics literature provides further evidence to support the notion that preference elicitation procedures matter. Discrepancies between values of a health outcome have been found depending on the method used to elicit them (e.g. Froberg and Kane 1989, Krabbe et al. 1997, Hornberger et al. 1992, Salomon and Murray 2004). There is a wide array of reasons given for these discrepancies, ranging from anchoring (e.g. Sutherland et al. 1982) to reference point effects and coherent arbitrariness (Ariely et al. 2003). Taken to the extreme, some authors have argued that there may not even be such a thing as underlying, stable and coherent preferences (see for example Fischhoff 1991, Slovic 1995, and Loomes 2006, for various discussions along these lines). There are striking similarities between valuing and ranking the importance of different health states, and valuing and ranking the importance of protecting different types of personal data. Both health states and personal data scenarios involve making judgements and choices that involve negative emotions like fear and distress. Both scenarios are relatively unfamiliar, in the sense that few people are regularly asked to directly evaluate different scenarios regarding their health, nor their privacy. Finally, both scenarios are difficult, arguably impossible, to incentivise, yet have

important implications for regulation and policymaking. The UK Treasury's Green Book guidance for policy appraisals states that

“If robust revealed preference data is not available, surveys that use willingness to pay and willingness to accept are an established alternative method known as stated preference techniques” (HM Treasury 2018, p. 42).

It seems wise, then, to draw on the lessons from health economics when considering the value of personal data and preferences to keep the data private. A key goal of this research is to put the methods used in the health economics literature to work in the context of privacy valuation, checking for the implicit ranking of types of data that they imply, and inferring the plausibility of using these elicitation techniques to establish the value of personal data to members of the public. We focus on the comparison between rankings of types of data that are based on absolute judgment (willingness to pay) tasks, and rankings of types of data based on more directly comparative relative judgement methods, to understand the extent to which preference orderings for personal data are stable across these frames.

In this study, we aimed to understand preferences for protecting personal data. Instead of assuming preferences for privacy to be captured by an exogenously determined “privacy value” constant in the utility function in all contexts, we explored how preferences differ across different data sharing environments. We did not assume that preferences are well-defined and stable but used a multitude of elicitation techniques to explore the consistency with which different types of data are ranked and valued.

We took a neutral approach when defining the context in which participants were asked to report the strength of preferences for keeping their personal data private, and we do not define the specific intended use of the information they were asked to share. Our study mimics real world online environments where individuals need to make their own assessment of costs and benefits associated with sharing of a particular type of personal data, and judge for themselves whether the trade-off is worth making.

Our approach does not intend to give a detailed exploration of the reasons behind the decision whether or not to share personal data, but instead focuses on examining two main issues: how individuals value keeping personal data private in a variety of contexts and how consistently do they report these preferences. We explored the reliability and stability of these preferences taking a holistic view of the prospect of sharing personal data in the context of the digital economy and compared individuals' overall willingness to share their data across personal data environments and elicited through different methods.

Experiment

This paper investigates the stability of participants' preferences for keeping their personal data private. In an online experiment, we asked whether participants place a positive monetary value on protecting different items of their personal data and explored how their willingness to pay differed between different data sharing environments. We drew inspiration from non-market valuation economic methodologies (e.g., contingent valuation, with its origins in Dupuit 1844; rating scales as explored in Sutherland et al. 1982; paired comparison, which draws on insights from Thurstone 1927), designing tasks in which participants expressed their absolute and relative judgments about the importance of protecting different items of their personal data. Through these tasks we could establish whether and how people value the privacy of their online data, whether they value different types of data differently and whether there is consistency between the rankings revealed by different elicitation methods.

The goals of our research were twofold. First, we aimed to ascertain the relative importance of each data type, as well as preliminary evidence about the levels of willingness to pay (WTP) for keeping each of these data types private. Second, we aimed to explore the consistency of the rankings of the data types across elicitation methods. To address these questions, we elicited the relative importance of protecting different personal data types using five different elicitation mechanisms, and gathered demographic information including age, gender, education and income to investigate whether WTP is affected by any of those

factors. Two elicitation methods focused on respondents' WTP for keeping their personal data private, one on the perceived absolute importance of personal data types and two on respondents' ranking of the importance of protecting different types of personal data - one pairwise ranking and one direct overall ranking. This is, to our knowledge, the first study that directly addresses the effect and consistency of different elicitation mechanisms in the content of data protection.

Participants & Procedure

A total of 265 participants took part in the online study run on the Qualtrics platform. The study took up to 30 minutes and participants received a fixed payment of £3 in return. The study was approved by Humanities and Social Sciences Research Ethics Sub-Committee, University of Warwick.

The survey was distributed through Prolific Academic (Prolific.ac). Three participants failed to answer all questions, and 262 responses were used in all of the following analyses. Of the respondents, 69.85% were female. Participants reported an average age of 37.28 (SD = 11.42), ranging from 18 to 72. Regarding education level, 53.44% reported no undergraduate degree or equivalent, 35.11% reported an undergraduate degree or equivalent, and 11.45% reported a degree higher than undergraduate. Our sample was more educated and contained a marginally higher proportion of female participants than the general population, since the UK population is 50.89% female (ONS, 2011a), and only 27.22% of the population have an undergraduate degree or higher qualification (ONS, 2011b). Income was measured on an interval scale with ten intervals: "less than £10,000"; "£10,000 to £19,999"; "£20,000 to £29,999"; "£30,000 to £39,999"; "£40,000 to £49,999"; "£50,000 to £74,999"; "£75,000 to £99,999"; "£100,000 to £149,999"; "£150,000 to £199,999"; and "£200,000 or more". There was also an option not to respond. The midpoint of each income interval was used as a proxy for the participant's income. The median of the household's annual income before tax for our sample was £25,000, equivalent to around \$33,700 (SD = £23,600 or \$32,000), which is lower than that UK population in general: the Office for National Statistics (2017) reported

median UK household pre-tax income of £33,705 for 2017. Income levels for people who preferred not to reveal this information (12 participants) were treated as missing values.

Experimental Design

The main preference elicitation section of the experimental design consisted of five conditions: two willingness to pay conditions, as well as three additional conditions to elicit individual preferences to protect different types of data. Every participant completed all five conditions, and the conditions were presented in a random order. The responses in the WTP conditions revealed whether the participant was willing to pay anything at all to protect the data types, and if so, how much they were willing to pay. The responses in the three non-WTP conditions revealed participants' ranking of data types in terms of their perceived importance of protecting the data. This allowed us to test the consistency of subjective rankings compared to the rankings implied by the WTP conditions.

All conditions had a common scenario. Participants were required to imagine that they had just bought a new smartphone. In addition to having all the usual functions such as camera or Internet access, it also had a special application which intended to assist them in their everyday life by tracking different types of their personal data. It was emphasised that the use of service was very desirable for a participant. The scenario specified that when they used the personal data application, they were given two choices; they could run the app in "free" mode, or they could run it in "paid" mode. In "free" mode, they did not have to pay anything but the company providing the app collected information about them when they used the app and used that information for various marketing and other purposes. In "paid" mode, they would pay a specified amount of money for the service and their data would remain private. In "paid" mode, the data would remain anonymous in a database and only used to maintain the functionality of the app. We did not specify the purposes for which the data might be used in "free" mode, allowing individuals to make up their own minds about the likely uses and value of their personal data for companies.

The personal data, defined here as “*D*”, could be one of the following: Banking Transactions, Browsing/Search/Click History, Electricity Use at Home, Loyalty Cards, Medical Records, Mobile Phone GPS, Physical Activity (exercise) Tracking and Social Media. These eight types of data were selected from a larger set of data types based on previous research (Skatova et al. 2013) as well as because they are commonly used in real life and are familiar to the general public.

In the *WTP Choice* condition, participants were first asked whether they were willing to pay anything, even a small amount, to keep the data *D* private, or alternatively whether they would prefer to use the app for free and share their data *D*. If they preferred to pay, we asked them whether they would be willing to pay Great British Pounds (£) 40 per month to keep the data private or else use the app for free and share the data. Based on their response, different amounts of money were presented, and respondents had to accept or reject paying that amount to keep the data private.

The amounts of money were iterated between the bounds 1 pence and £200 until the person was indifferent. The iteration procedure was designed to reveal bounds on the individual's indifference point, and the indifference in Willingness to Pay (WTP) was coded as the midpoint of these bounds. For instance, if an individual sequentially confirmed they would be willing to pay £40, £125 and £175 per month to keep data *D* private but they would not be willing to pay £200, then their highest acceptable amount is £175 per month while their lowest unacceptable amount is £200 per month. This individual's WTP would be recorded as £187.50 per month. If a participant indicated that they would be willing to pay £200, their indifference point could not be determined by the iteration. There were ten cases where participants were willing to pay more than £200 to protect some of their data types. In these cases, we asked them to state the maximum amount they would be willing to pay to keep this specific type of data private, which ranged from £210 to £300. Participants completed this procedure eight times, once per data type. The starting value of £40 and the maximum value of £200 was chosen based on pilot work with an open-ended WTP task.

In the *WTP Slider* condition, participants used a slider to directly assign the largest amount of money they would be prepared to pay to keep personal data D private. The slider ranged from £0 to £200 per month. The slider was set at a default of £40 per month to start with. This amount, and the £200 maximum, were chosen to match the £40 starting amount in WTP Choice condition. Participants answered eight of these WTP slider tasks, one per data type.

In the *Importance* condition participants used a five-point scale, ranging from “not important at all” to “very important”, to rate the importance of keeping each piece of personal data D private as opposed to sharing it. Participants answered eight of these tasks, one per data type.

In the *Direct Ranking* condition participants ranked the 8 types of personal data in order of the importance of keeping them private. They had to drag and drop types of data in a vertical ranking, where the top of the ranking was the data they were most willing to share, if they needed to use the app for free. At the bottom of the ranking was the data they were the most unwilling to share. Participants completed one ranking which included all eight data types.

In the *Forced Paired Choice* condition participants made a series of forced paired choices. Each choice was between two pieces of personal data, data D_1 and data D_2 . Each time, they needed to decide which data type they would prefer to give away if they needed to use the app for free. Participants completed 28 of these choices, one for every unique binary pairing of the data types.

Analytic approach

We first tested whether participants were consistent in their privacy preferences by calculating how consistent the rankings of data types were between the five conditions within each participant. We then used a probit and an OLS regression to analyse whether the demographic information explained variation in whether and how much individuals were WTP to keep their data private, whether WTP depended on the types of data, and whether it depended on within-person consistency scores.

Individual consistency in ranking different data types

To examine consistency in individual preferences for sharing different types of data, we converted people's responses in each condition into rankings. The *Direct Ranking* condition data was used without conversion. For the *WTP Choice* and *WTP Slider* conditions we assumed that a higher WTP for a data type always indicated a lower willingness to share. That is, if one's WTP for protecting banking transactions data is £10 and for protecting browsing history data is £1, the participant was assumed to be less willing to share banking transactions data than browsing history data. For the *Importance* condition we assumed a higher-rated importance of a data type indicates a lower willingness to share. For the *Forced Paired Choice* condition we counted the number of times each data type was chosen as the one to share, and assumed that a lower count implied lower willingness to share the data. This way, if banking data was never given away it would have the lowest willingness to share and thus the highest importance rank, and if loyalty cards data was always given away, it would have the highest willingness to share and the lowest importance rank. For any ties in rankings, standard competition ranking was employed to determine the Willingness to Share (WTS) rank. This was converted to an importance score where the number increases with importance of the data type. Table 1 demonstrates an example of one individual's scores. This individual was never willing to give away Mobile Phone GPS, was willing to give away Medical Records and Banking Transactions twice each and Browsing/Search/Click History and Social Media three times each. Mobile Phone GPS was ranked the most important (i.e., Willingness to Share (WTS) rank 1, importance rank 8), Medical Records and Banking Transactions was ranked second most important (i.e. WTS rank 2, importance rank 7), while Browsing/Search/Click History and Social Media will be 4th most importance (i.e., have WTS rank 4, importance rank 5).

Table 1. An example of ranking in Paired Choice condition.

Data type	Number of times preferred to share	Rank by importance
Banking Transactions	2	7
Browsing/Search/Click History	3	5
Electricity Use at Home	6	2
Loyalty Cards	5	3
Medical Records	2	7
Mobile Phone GPS	0	8
Physical Activity (Exercise) Tracking	7	1
Social Media	3	5

To measure consistency of each individual's ranking across different conditions we adopted Krippendorff's coefficient (Krippendorff 2004).

The alpha coefficient is defined as (Krippendorff 2004, page 222):

$$\alpha = (1 - D_o)/D_e$$

where D_o is a measure of the observed disagreement and D_e is a measure of the disagreement that can be expected by chance. In our context, the observed disagreement refers to the individual participants inconsistency in rankings for the same data type across different conditions. The coefficient ranges from -1 to 1: when individuals' sensitivity rankings are completely consistent across elicitation methods and inconsistent choices are completely

absent, $D_0 = 0$ and $a = 1$. When consistent choice and inconsistent choices are purely driven by chance, $D_0 = D_e$ and $a = 0$. Alpha less than zero implies systematic differences in ratings that are not due to random chance, for example if two methods elicited opposite ratings for the data types.

Regressions

To estimate whether demographic and control variables influenced the likelihood of being willing to pay something greater than zero to protect personal data, we used mixed-model probit regressions with a random effect to account for the fact that each participant provided eight answers in each of the two WTP conditions. The Banking Transactions data type was used as a reference category. The following model was specified:

$$Prob (y_i = 1) = F (\beta_0 + \beta_1 T_i + \beta_2 K_i + D_i' \gamma + \varepsilon_i)$$

where Y_i is a binary variable taking the value 1 if the individual chose to pay money to protect the data, and 0 otherwise; D_i' is a vector of demographic variables (age, gender, education and income); K_i is the individual Krippendorff's alpha coefficient (Krippendorff 2004) and T_i is the type of data. The function F is the cumulative distribution function of the standard normal distribution.

To estimate whether demographic and control variables influenced *the level of WTP to keep data private* we used mixed-model ordinary least squares (OLS) regressions, again with a random effect of a participant (e.g. Greene 1997; Cameron and Trivedi 2010):

$$Y_i = \beta_0 + \beta_1 T_i + \beta_2 K_i + D_i' \gamma + P_i' \delta + \varepsilon_i$$

where Y_i is an individual's WTP for keep their data private and is treated as continuous; D_i' , T_i , K_i are defined as in the probit model.

Results

Are people consistent in their judgements about different types of personal data between conditions?

We calculated average rankings for each data type in each condition, pooling across participants. The aggregate ranking of each data type between different conditions was similar, with the exception of Mobile Phone GPS: see Figure 1 for aggregate ranking scores and corresponding 95% Confidence Intervals, and Figure 2a, Figure 2b, and Figure 2c for density distributions of ranks for each data type and each condition.

Figure 1. Average ranking and corresponding 95% Confidence Intervals for all conditions.

Figure 1 demonstrates three clear “tiers” of data. Banking Transactions and Medical Records, which we determine to be “Tier 1” data, are ranked consistently as the most valuable between all conditions, and ranked on average as 7.52 and 6.98, out of 8, respectively. A second tier of data types includes Browsing History, Mobile Phone GPS and Social Media Data, which are consistently ranked just under 6, with the exception of a lower ranking implied by the Forced Paired Choice condition for all three data types, and higher ranking for Mobile Phone GPS data in the WTP Choice condition. Finally, the lowest importance category, “Tier 3”, of data includes Electricity Use, Loyalty Cards and Physical Activity data, rankings for which vary between 2 and 5 across all conditions. On average, the

ranking of different data types appears consistent between conditions, with the exception of Mobile Phone GPS which is ranked in the first tier in WTP Choice but not in other conditions.

Figure 2a. Density distribution of ranks for Tier 1 data: Banking and Medical Records Data.

Figure 2b. Density distribution of ranks for Tier 2 data: Browsing History, Mobile Phone GPS and Social Media Data.

Figure 2c. Density distribution of ranks for Tier 3 data: Electricity Use, Loyalty Cards and Physical Activity Data.

The Figures 2 show that Tier 1 data types are the most consistently ranked between different conditions, with density distributions that have similar shapes within the data type. Tier 2 data types have more variation and Tier 3 types have the least stability in ratings between conditions which is observable from different shapes of density distributions between conditions for the same data type. In other words, for Tier 1, and to a certain extent for Tier 2, it did not matter which elicitation procedure was used to elicit the ranks of the personal data types; the implied rankings were stable and consistent between conditions. However, for Tier 3, the shape of distribution seems to be mostly driven by the condition, not by data type: the distributions are more similar reading up the columns than along the rows for these data types. This can be interpreted to mean that the responses with respect to Loyalty Cards, Electricity Use and Physical Activity data were less well defined, and hence the responses were driven more by the elicitation procedure than by true underlying privacy preferences.

Individual level consistency

For each participant we calculated Krippendorff's alpha coefficient to reflect how consistent they were in ranking different data types. The average Krippendorff's alpha coefficient was 0.51 (with a median of 0.60, a standard deviation of 0.29), ranging from -0.17 to 0.92. To provide an intuition for the degree of consistency implied by these coefficients, we present the rankings provided by a randomly selected participant from the bottom 5th percentile of the consistency distribution, and by a randomly selected participant in the top 5th percentile of the consistency distribution, in Figure 3, which clearly demonstrates that Participant 49 who had a low consistency score lack of consistency of judging the same data types in different conditions, while Participant 149 who had a high consistency score is very consistent in their responses about data types across conditions.

Figure 3a. Randomly selected participant from the bottom 5% consistency score.

Figure 3b. Randomly selected participant from the top 5% consistency score.

To describe the whole sample in terms of consistency, we adopted the benchmark scales suggested by Landis and Koch (1977) (see also Krippendorff 2011) which allows participants to be classified in terms of their consistency scores. The share of participants who display poor, slight, fair, moderate, substantial and almost perfect consistency are shown in Table 2. The table demonstrates that only 20.6% of participants are below the threshold for “fair” consistency with the majority (70.61%) being moderately, substantially or almost perfectly consistent between conditions.

Table 2. Break down of Krippendorff alpha consistency scores by poor, slight, fair, moderate, substantial and almost perfectly consistent, Study 1.

Coefficient	Interpretation	Number of Participants	% of participants
<0	Poor	24	9.15%
[0.00 - 0.20)	Slight	30	11.45%
[0.21 - 0.40)	Fair	23	8.79%
[0.41 - 0.60)	Moderate	52	19.85%

[0.61 - 0.80)	Substantial	106	40.46%
[0.81 - 1.00)	Almost Perfect	27	10.30%

To pay or not to pay?

96% of participants stated for at least one data type that they would pay to protect it. There were differences in the likelihood with which individuals would be willing to pay for different data types. Models 1 and 2 in Table 3 summarise the results of Probit regression predicting whether to pay or not to pay for personal data from data types, consistency and demographic variables.

Table 3. Mixed-model Probit (predicting willingness to pay) and OLS (predicting the amount given people are willing to pay) for WTP Choice (Model 1 and Model 3) and WTP Slider (Model 2 and Model 4), N = 250. Banking transactions are used as a reference category.

The table presents beta coefficients and 95% Confidence Intervals for each coefficient in brackets, as well as random effects, Log Likelihood. Values highlighted in bold a significant at least on $p < .05$ level.

	Probit: Generalised linear mixed-effects		OLS: Linear mixed-effects	
	Model 1	Model 2	Model 3	Model 4
Intercept	1.72 [1.38; 2.07]	2.38 [1.88; 2.87]	28.09 [22.40; 33.78]	22.76 [17.77; 27.7]
Medical Records	-0.93 [-1.25; -0.61]	-0.87 [-1.28; -0.46]	-3.10 [-7.36; 1.17]	-0.27 [-3.08; 2.55]

Mobile Phone GPS	0.02 [-0.34; 0.38]	-1.91 [-2.22; 1.40]	0.32 [-3.95; 4.58]	-10.98 [-13.79; -8.17]
Browsing History	-1.90 [-2.23; -1.58]	-1.74 [-2.15; -1.34]	-15.97 [-20.23; -11.70]	-10.25 [-13.07; -7.44]
Social Media	-1.99 [-2.31; -1.66]	-1.78 [2.18; -1.37]	-18.05 [-22.31; -13.79]	-11.36 [-14.17; -8.55]
Electricity Use	-3.21 [-3.59; -2.83]	-3.05 [-3.48; -2.63]	-21.54 [-25.81; -17.28]	-15.54 [-18.35; -12.72]
Loyalty Cards	-3.37 [-3.48; -2.73]	-2.96 [-3.39; -2.54]	-21.57 [-25.83; -17.30]	-15.07 [-17.88; -12.25]
Physical Activity	-3.37 [-3.76; -2.97]	-3.25 [-3.69; -2.81]	-22.20 [-26.46; -17.93]	-15.64 [-18.46; -12.83]
Standardised Consistency	0.43 [0.32; 0.55]	0.18 [0.01; 0.35]	1.94 [-0.42; 4.30]	-1.38 [-3.60; 0.83]
Standardised Age	0.08 [-0.03; 0.20]	-0.29 [-0.46; -0.12]	-0.95 [-3.32; 1.42]	-2.27 [-4.50; -0.05]
Female	-0.20 [-0.43; 0.03]	-0.12 [-0.48; 0.23]	-8.06 [-13.11; -3.02]	-5.13 [-9.86; -0.38]
Degree:	0.06	0.02	-3.12	-1.85

Undergraduate	[-0.18; 0.30]	[-0.34; 0.38]	[-8.27; 2.03]	[-6.69; 2.98]
Degree: Postgraduate	0.48 [0.12; 0.84]	0.31 [-0.25; 0.88]	18.61 [10.84; 26.39]	10.06 [2.76; 17.35]
Standardised income	-0.08 [-0.19; 0.03]	-0.08 [-0.25; 0.09]	-1.22 [-3.61; 1.17]	-1.28 [-3.53; 0.96]
Random effects				
Intercept σ (participant)	0.38	1.34	273.0	273.3
Model Fit				
Deviance (-2LL)	-811.97	-936.93	18,894.57	17,413.04

Figure 4 demonstrates predicted effects for each data type, conditioned on all other variables. The data type that was considered the most important to protect was Banking Transactions, with an estimated likelihood of being willing to pay to protect at 95.8% in WTP Choice and 99.2% in WTP Slider conditions, respectively. For Medical Records, the estimated likelihood of being willing to pay to protect was 78.9% in WTP Choice and 93.6% in the WTP Slider conditions, respectively. The second tier of data are Browsing History (43% in WTP Choice and 74.1% in WTP Slider) and Social Media (39.8% in WTP Choice and 73.1% in WTP Slider). Mobile Phone GPS was closer to the first tier in the WTP Choice condition with 96% predicted probability to pay, while in WTP Slider condition it was less protected: 71.9%. Finally, the third tier included Loyalty Cards (6.9%, WTP Choice; 25.3%,

WTP Slider), Electricity Use (8.4%, WTP Choice; 28.4%, WTP Slider) and Physical Activity (5.1%, WTP Choice; 19.5% WTP Slider).

Figure 4. Predicted probability to pay for different data types, 95% Confidence Intervals.

Individual consistency between ranking orders predicted higher likelihood of being willing to pay to protect personal data: those participants who demonstrated higher consistency in their responses were more likely to decide to pay something in both WTP conditions - that is, a one standard deviation increase in people's consistency measurement is estimated to cause a 9% and a 4% increase in the likelihood to pay to protect their personal data in the WTP Choice and WTP Slider conditions, respectively. This might reflect more thorough consideration of consequences of sharing personal data by those who showed higher consistency.

Out of the demographic variables, only age was negatively related to the likelihood to pay to keep the data private, and this was the case only in the WTP Slider condition: younger people were more willing to pay to keep their data private, with a one standard deviation (~11.42 years) decrease in age leading to a significant increase of 6% in the likelihood of being willing to pay to keep data private. Further, more educated participants, specifically those who stated that they held a postgraduate degree, were 10% more likely to choose to pay in WTP Choice condition. Income level had no significant effect on this decision.

How much to pay?

We have considered the results of the Probit analysis which investigated the factors influencing whether or not participants were willing to pay a non-zero amount to protect their personal data. Next, we turn to the OLS regression (see Table 3, Model 3 and 4) which investigates the level of WTP, conditional on being willing to pay an amount greater than zero. The results of the OLS regression differed from (but did not seriously contradict) the results of the Probit analysis. We discuss the OLS results and differences from the Probit below.

Figure 5. Predicted amount individuals are willing to pay for different data types, out of those who were prepared to pay something, 95% Confidence Intervals.

The results related to the differences between data types were consistent with those from the Probit analysis, suggesting - reassuringly- that the data types that people are more likely to pay to protect are also the data types that they are willing the pay largest amounts to protect. Figure 5 demonstrates the estimated effects in pounds and the associated 95% Confidence Intervals for each data type. The estimated WTP to protect Bank Transaction data was

£28.10 per month in the Choice condition and £22.80 in the Slider condition. The estimated WTP for protecting Medical Records was £25.10 for the Choice condition and £22.50 for the Slider condition. The estimated WTP for Slider and Choice conditions for the Tier 2 data types were as follows: Mobile Phone GPS £28.50 and £11.80; Browsing History £12.12 and £12.60; and Social Media £10.10 and £11.40, respectively. Finally, our participants were willing to pay £6.61 and £7.27 for Electricity Use, £6.58 and £7.74 for Loyalty Cards and £5.95 and £7.17 for Physical Activity data in the WTP Choice the WTP Slider condition, respectively.

Individual consistency in ranking order within condition was not associated with the amount of money participants were willing to pay in either conditions. Similar to the decision about whether to pay anything at all, the amount that participants were willing to pay was associated negatively with age: one standard deviation decrease in years (~11.42 years) leads to an average payment to protect personal data that is £2.27 higher, but only in the WTP Slider condition. Female participants were willing to pay smaller amounts compared to men in both WTP conditions: £8.06 in Choice and £5.13 in Slider, on average. Finally, having a graduate degree was associated with being willing to pay higher amounts of money to protect data in both conditions: £18.71 more in the Choice condition and £10.06 more in the Slider condition. As with the Probit, income did not have a significant effect on WTP.

Discussion

This paper makes three main contributions. First, we showed that individuals state they are willing to pay non-zero amounts to protect their personal data. Second, individuals demonstrate consistent preferences for paying to protect different types of their personal data, with values and ranks being remarkably similar across different elicitation methods. We suggest that these elicited preferences are indicative of people's true preferences for protecting their personal data. Third, different types of personal data are valued differently, suggesting there is no unified "price tag for privacy". Our findings reveal consistency in

people's decisions about whether to share their personal data through a variety of elicitation procedures, which is an evidence for consistent preferences for privacy.

We now discuss how our findings relate to previous research exploring the value of personal data, explain how our methodology allowed us to unpack privacy and its value, and discuss how our results shed light on the possibilities for emerging business models and policy approaches in the domain of personal data protection. Different types of personal data might have different value for different stakeholders, and personal data may even have different value for the same consumer in various scenarios (e.g., a scenario in which a company provides services in exchange for consumer data vs a scenario in which consumers sell their data to company). Our study is the first to unpack the relative importance of protecting different types of personal data, ranging from very subjectively important, such as medical records and financial information, to - at present time - subjectively unimportant, such as loyalty cards and electricity data. Our study deliberately did not control for specific risks and benefits of sharing as we aimed to mimic the real-world choices that people often face, where risks and benefits of agreeing to share one's data are typically unclear, although future research could usefully contribute to the literature by exploring these types of effect.

Privacy preferences

Our work extends the growing literature on the monetary valuation of privacy (Huberman et al. 2005, Grossklags and Acquisti 2007). It also contributes to the body of knowledge concerning the factors that influence individuals' decisions about whether or not to share elements of their personal data online. To the best of our knowledge, this is the first report of a study that directly compares the value of protecting different types of personal data using five different elicitation techniques. We explored two main questions: whether individuals would state that they were willing to pay to protect different types of personal information, and whether their responses would imply rankings that were stable across different evaluation techniques, since economic theory suggests that if stable and well-defined privacy

preferences exist, individuals' stated WTP would be consistent across different evaluation techniques.

In line with some previous research (Wagner et al. 2018) and in contrast to other studies (Grossklags and Acquisti 2007), our findings show that individuals state positive willingness to pay to ensure their personal data remains private, at least for some types of data. We observed a surprising degree of consistency in participants' rankings of different data types elicited through different techniques, which indicates that stated preference evaluation techniques may be appropriate tools to reveal individuals' underlying preferences for privacy. In our study, all factors that could affect people's decisions about what value they assign to different types of personal data remained the same, with only data types and evaluation techniques varying between the questions. We suggest that while the stated preferences reported by the individuals in our study display remarkable consistency, there are still features of the decision context that could influence the value of privacy online. These features may include the context of data sharing, the type of data that is shared, what the purpose of sharing the data is, and who the data is shared with. We suggest that these features will interact with underlying privacy preferences and contribute to producing the observed values of personal data. By not specifying or deliberately manipulating these features of the data sharing environment in our scenarios, we examined the effects of data type and elicitation procedure in isolation.

Our methodological approach and results demonstrate the benefits of clearly specifying privacy decision-making scenarios as choices about protecting specific types of data. The clear differences in value between different data types suggests that the value of privacy in the digital world cannot be defined as a unified concept: it is not possible to assign a single price tag to individual privacy. Instead, our findings suggest that in order to understand the value of privacy, it is necessary to study how individual privacy preferences play out in different data sharing environments. For example, in our study there was a clear contrast between very high rates of willingness to pay to protect banking transactions data and very

low rates of willingness to pay to protect electricity use or physical activity data, demonstrating that public preferences for keeping the former type of data private are stronger than for the latter types. In addition, responses about low-valued data, referred to earlier in the paper as Tier 3, were noisier and appeared to be more strongly influenced by the elicitation techniques than for the high-valued Tier 1 and 2 data types. This result suggests that individuals have less well-defined preferences about protecting the Tier 3 data types than they do for the highly valued data types.

Understanding how much people value different types of personal data is an important first step in understanding individual privacy preferences and how they affect decisions in the real world. Everyday, people face scenarios where they are presented with requests to share different types of personal data, albeit the context might vary. Our results and the evaluation framework that we have introduced in this paper could be used to further investigate how those valuations change with different framing, or when the instructions manipulate the emphasis placed on the risks versus the benefits of sharing personal data online.

Future economic models of privacy

Currently, the digital economy does not typically provide a direct monetary benefit for consumers who create personal data. The benefits are instead provided through increased convenience of service delivery, or better opportunities for accessing certain services. Largely, this relies on the assumption that the actual value of one consumer's personal data is very hard to estimate and even if it is possible, the value is so small that consumers do not consider their personal data to be valuable enough to demand monetary compensation for sharing it. However, our findings, in line with previous research (see Acquisti et al. 2016) suggest that there may exist an imbalance in power between the consumer and the data collector/holder in terms of the allocation of the welfare gain, whereby the consumer places a positive value on protecting their personal data and yet the welfare gain from the use of that personal data is allocated solely to the data handler. We show that consumers place a positive monetary value on at least some types of their personal data, which implies that

business models that allow them to pay money to protect their personal data may be a viable option. Similarly, business models that financially compensate consumers for the use of their personal data may also be viable. At least some consumers might expect that if the data that they produced (over which they may or may not hold ownership, depending on the laws of the country), is used for commercial benefit, they should receive reasonable compensation in return. In our study we do not identify precisely what would this reasonable compensation might be, but, importantly, we demonstrate that it is not zero.

Limitations and future research

The choices in our experiments were hypothetical, following the tradition of the literature on the monetary valuation of non-market goods. We opted to take a hypothetical approach because it was not possible to identify real life choices (revealed preferences) in which different data types are traded in otherwise identical scenarios. Whilst incentivizing the sharing of personal data in experimental conditions is possible, and has been done before (e.g., Grossklags and Acquisti 2007), it was not appropriate for our research question because it is not possible to conduct an ethical experiment where individuals would be requested to pay money to avoid sharing their personal data in a variety of highly sensitive data sharing environments. Using hypothetical stated preference tasks avoids this necessity by allowing us full control of the data sharing and payment scenarios and making it possible to ask about sensitive data types.

Furthermore, it is not currently feasible to construct an incentivised study where these different types of personal data are traded. Based on our findings, which showed remarkable consistency in stated preferences for privacy, future research could be usefully conducted that focuses on a single type of data (e.g., retail loyalty cards) involving a field study where individuals could opt to pay for privacy or else opt to benefit from a free service. Building such a scenario would allow the researcher to manipulate different aspects of context of data sharing, to examine whether and how these aspects affect individual decisions. It would unlikely, however, to conduct such experiment with multiple different data types.

Decisions about sharing personal data are always contextual (Acquisti et al. 2016). In our study we deliberately did not define context and reasons for sharing personal data. Future research could manipulate whether the data is shared for commercial purposes or for the public good to investigate whether this affects people's decisions about whether to share their personal data, and whether different reasons for sharing affect people in a different way.

We might also consider willingness to accept instead of willingness to pay: the literature on the economics of privacy and security implicitly assumes that the market behaviour of individuals for both actions is identical, but related literature in psychology suggests otherwise. Some previous research (e.g., Beresford et al. 2012) elicits individuals' willingness to accept money in return for sharing their personal data, instead of willingness to pay to protect it. This approach can underpin business models in which individuals have the right to privacy but can choose to sell their personal data, for example through companies like datafairplay.com that give individuals the chance to directly choose how much data to share for financial reward. In contrast, we took a willingness to pay approach and found that individuals' willingness to pay to protect their data was positive. Our findings therefore support the possibility of a market in which individuals pay to protect their data. A final alternative, that might be explored in future research, is one in which consumers receive a specific, non-monetary benefit for sharing their data. Understanding the gap between willingness to pay and accept, as well as whether different scenarios affect values which are produced through different elicitation techniques, may have implications for future business models which evolve based on the use of personal data.

Finally, our results have implications for policy-making. Our results strongly suggest that privacy is not a unified concept and people make different decisions depending on what kind of data they are sharing. The types of detriment that are experienced when data are not protected can be diverse - including privacy loss *per se*, risks of adverse consequences for the individual, and a basic disapproval of situations where a private company makes money

from one's own personal data. It is important to understand the preferences of members of the public towards all of these elements when making decisions about how to create and regulate a fair environment in which individuals are able to make informed choices about protecting their privacy. Similar to health economics where resources are distributed based on public evaluations of different health outcomes (HM Treasury Green Book 2018), arguably government regulations with respect to personal data and digital privacy should take account of the value of personal data privacy to society. Our study is the first to look at the breadth of different data types in terms of how they are evaluated by members of public, and we provide reassuring evidence in favour of the stability of preferences for online privacy. Our results therefore indicate how public preferences might be included in policy making in this important context.

References

Acquisti A, Taylor C, Wagman L (2016) The economics of privacy. *J. Econ. Lit.* 54(2), 442-92.

Acquisti A, John LK, Loewenstein G (2013) What is privacy worth? *J. Legal Stud.* 42(2), 249-274.

Acquisti A, Brandimarte L, Loewenstein G (2015) Privacy and human behavior in the age of information. *Science* 347(6221), 509-514.

Acquisti A, Varian HR (2005) Conditioning prices on purchase history. *Marketing Science* 24(3), 367-381.

Adjerid I, Acquisti A, Telang R, Padman R, Adler-Milstein J (2015) The impact of privacy regulation and technology incentives: The case of health information exchanges. *Management Sci*, 62(4), 1042-1063.

Ariely D, Loewenstein G, Prelec D (2003) "Coherent arbitrariness": Stable demand curves without stable preferences. *Q. J. Econ.* 118(1), 73-106.

- Barnes SB (2006) A privacy paradox: Social networking in the United States. *First Monday* 11(9).
- Becker GM, DeGroot MH, Marschak J (1964) Measuring Utility By a Single-response Sequential Method. *Behav. Sci.* 226-32.
- Beresford AR, Kübler D, Preibusch S (2012) Unwillingness to pay for privacy: A field experiment. *Econ. Lett.* 117(1), 25-27.
- Benndorf V, Normann HT (2017) The willingness to sell personal data. *Scand. J. Econ.* 120(4), 1260–1278
- Blank G, Bolsover G, Dubois E (2014) A new privacy paradox: Young people and privacy on social network sites. *Computers and Security* 64(1), 122-134
- Brown B (2001) Studying the Internet experience. *HP Laboratories Technical Report* HPL 49.
- Calo R (2011) The boundaries of privacy harm. *Ind. L.J* 86, 1131.
- Cameron AC, Trivedi PK (2010) *Microeconometrics using Stata*, revised ed. (College Station, TX: Stata Press).
- Casadesus-Masanell R, Hervas-Drane A (2015) Competing with privacy. *Management Sci.* 61(1), 229-246.
- Dupuit J (1844) On the measurement of the utility of public works. *Internat. Econ. Papers*, 2, 83-110.
- Fischhoff B (1991) Value elicitation: Is there anything in there? *Am. Psychol.* 46(8), 835.
- Froberg DG, Kane RL (1989) Methodology for measuring health-state preferences—II: Scaling methods. *J. Clin. Epidemiol.* 42(5), 459-471.
- Goldfarb A, Tucker CE (2011). Privacy regulation and online advertising. *Management Sci.* 57(1), 57-71.
- Goldfarb A, Tucker CE (2012) Shifts in privacy concerns. *Am. Econ. Rev.* 102(3), 349-53.
- Greene WH (1997) *Econometric analysis*, 3rd edn (Upper Saddle River, NJ: Prentice Hall).

Grossklags J, Acquisti A (2007) When 25 Cents is Too Much: An Experiment on Willingness-To-Sell and Willingness-To-Protect Personal Information. *Workshop Econ. Information Security*.

Hann IH, Hui KL, Lee TS, Png IP (2002) The value of online information privacy: Evidence from the USA and Singapore. *Internat. Conf. Information Systems*.

HM Treasury Green Book (2018) The economic value of data: discussion paper. Available at:

https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/731349/20180730_HMT_Discussion_Paper_-_The_Economic_Value_of_Data.pdf

Hornberger JC, Redelmeier DA, Petersen J (1992) Variability among methods to assess patients' well-being and consequent effect on a cost-effectiveness analysis. *J. Clin. Epidemiol.* 45(5), 505-512.

Huberman BA, Adar E, Fine LR (2005) Valuating privacy. *IEEE Security & Privacy* 3(5), 22-25.

Jentzsch N (2014) Auctioning Privacy-Sensitive Goods: A Note on Incentive-Compatibility. *Privacy Technologies and Policy: Second Annual Privacy Forum, APF 2014, Athens, Greece, May 20-21, 2014, Proceedings* (Vol. 8450, p. 133). Springer.

Kokolakis S (2017) Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security* 64, 122-134.

Krabbe P, Essink-Bot ML, Bonsel G (1997) The comparability and reliability of five health-state valuation methods. *Soc. Sci. Med.* 45(11), 1641-1652.

Krippendorff K (2004) *Content analysis: An introduction to its methodology*, 2nd edn. (Thousand Oaks, CA: Sage).

Krippendorff K (2011) Computing Krippendorff's alpha-reliability. Working Paper, University of Pennsylvania, Pennsylvania.

Landis JR, Koch GG (1977) The measurement of observer agreement for categorical data. *Biometrics*. pp.159-174.

Laufer RS, Wolfe M (1977) Privacy as a concept and a social issue: A multidimensional developmental theory. *J. Soc. Issues*. 33(3), 22-42.

Loomes G (2006) (How) Can we value health, safety and the environment? *J. Econ. Psychol.* 27(6), 713-736.

Miller AR, Tucker C (2017) Privacy protection, personalized medicine, and genetic testing. *Management Sci*, 64(10), 4648-4668.

Montes R, Sand-Zantman W, Valletti TM (2018). The value of personal information in online markets with endogenous privacy. *Management Sci*. Article in Advance, p 1-21.

Norberg PA, Horne DR, Horne DA (2007) The privacy paradox: Personal information disclosure intentions versus behaviors. *J. Consumer Aff.* 41(1), 100-126.

Nord E (1991) Methods for establishing quality weights for life years, National Centre for Health Program Evaluation.

Office for National Statistics (2011a) 2011 Census: Population Estimates for the United Kingdom, http://webarchive.nationalarchives.gov.uk/20160108132257/http://www.ons.gov.uk/ons/dcp171778_292378.pdf

Office for National Statistics (2011b) 2011 Census: Key Statistics for England and Wales, <https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/bulletins/2011censuskeystatisticsforenglandandwales/2012-12-11>

Office for National Statistics (2017) Household disposable income and inequality in the UK: Financial year ending 2016. London: Office for National Statistics. Available at: <https://www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/incomeandwealth/bulletins/householddisposableincomeandinequality/financialyearending2016>.

Posner RA (1978) Economic theory of privacy. *Regulation*, 2, 19.

- Rose EA (2006) An examination of the concern for information privacy in the New Zealand regulatory context. *Information & Management* 43(3), 322-335.
- Rosen LD, Rosenkoetter P (1976) An eye fixation analysis of choice and judgment with multiattribute stimuli. *Mem. Cogn.* 4(6), 747-752.
- Salomon JA, Murray CJ (2004) A multi-method approach to measuring health-state valuations. *Health Econ.* 13(3), 281-290.
- Schreiner M, Hess T (2015) Why Are Consumers Willing to Pay for Privacy? An Application of the Privacy-freemium Model to Media Companies. *European Conf. on Information Systems*. Paper 164.
- Shen Q, Miguel Villas-Boas J (2017) Behavior-based advertising. *Management Sci*, 64(5), 2047-2064.
- Spiekermann S, Korunovska J (2017) Towards a value theory for personal data. *J. Inf. Technol.* 32(1), 62-84.
- Skatova A, Johal J, Houghton R, Mortier R, Bhandari N, Lodge T, Wagner C, Goulding J, Crowcroft J, Madhavapeddy A (2013) Perceived risks of personal data sharing. *Proc. Digital Economy: Open Digital*.
- Slovic P (1995) The construction of preference. *Am. Psychol.* 50(5), 364.
- Stigler GJ (1980) An introduction to privacy in economics and politics. *J. Legal Stud.* 9(4), 623-644.
- Sutherland HJ, Dunn V, Boyd NF (1982) Measurement of values for states of health with linear analog scales. *Med. Decis. Mak.* 3(4), 477-487.
- Thurstone LL (1927) A law of comparative judgment. *Psychol. Rev.* 34, 273.
- Tversky A, Sattath S, Slovic P (1988) Contingent weighting in judgment and choice. *Psychol. Rev.* 95(3), 371.
- Tucker CE (2014) Social networks, personalized advertising, and privacy controls. *J. Mark. Res.* 51(5), 546-562

Wagner A, Wessels N, Buxmann P, Krasnova H (2018) Putting a Price Tag on Personal Information-A Literature Review. *Proc. Hawaii Internat. Conf. on System Sci.*