

THE UNIVERSITY OF CHICAGO

VALUING INTRINSIC AND INSTRUMENTAL PREFERENCES FOR PRIVACY

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
XIAOCAI LIN

CHICAGO, ILLINOIS

JUNE 2020

Copyright © 2020 by Xiaocai Lin
All Rights Reserved

To Alex, who encouraged me to pursue this path

TABLE OF CONTENTS

LIST OF FIGURES	vi
LIST OF TABLES	vii
ACKNOWLEDGMENTS	viii
ABSTRACT	ix
1 INTRODUCTION	1
2 THE CONCEPTUAL FRAMEWORK	7
2.1 Setup	7
2.2 Privacy Preferences	8
2.3 Who Chooses Not to Share Personal Data?	9
2.4 Extending the Stylized Framework	11
3 THE EXPERIMENT	13
3.1 Empirical Challenges and Solutions	13
3.2 Examination of the New Policy Regime	14
3.3 Experiment Design	15
3.4 Discussion	19
4 DATA AND DESCRIPTIVE EVIDENCE	21
4.1 Data Source and Cleaning	21
4.2 Sample Characteristics	22
4.3 Intrinsic Preferences	23
4.4 Instrumental Preferences	24
4.5 Dual Privacy Preferences and the Selection in Shared Data	26
5 THE STRUCTURAL MODEL	28
5.1 Setup	28
5.2 Identification	31
5.3 Estimation	32
6 ESTIMATION RESULTS	33
6.1 Model Comparison	33
6.2 Intrinsic Preferences	33
6.3 Beliefs that Generate Instrumental Preferences	36
6.4 Dual Privacy Preferences and Selection in Shared Data	37
7 REPLICATING THE EXPERIMENT IN THE FIELD	40

8	COUNTERFACTUALS	42
8.1	Setup	42
8.2	When and How to Buy Data from Consumers?	43
8.3	Learning about Selection: What Do Privacy Choices Reveal?	50
8.4	Summary	54
9	CONCLUSION	55
	REFERENCES	57
A	PROOF FOR PROPOSITION 1	62
B	SCREENSHOTS OF THE SURVEY	64
B.1	Conjoint and Demographic Questions	64
B.2	Compensation Schedules across Treatments	65
C	ATTRITION	67
D	CREDIBLE INTERVALS FOR INTRINSIC PREFERENCE ESTIMATES (WTA) . .	68
E	INTRINSIC WTA ESTIMATES WITH HETEROGENEOUS SENSITIVITY TO IN- COME	69
F	PSYCHOLOGICAL FACTORS	70
F.1	The Default Frame	70
F.2	Other Psychological Factors	76

LIST OF FIGURES

3.1	Treatment Design	18
4.1	Frequency of Data Sharing across Incentive Treatments	25
6.1	Posterior Predicted Density of WTA in Intrinsic Preference	35
6.2	Full vs. Firm Data across Different Magnitudes of Instrumental Incentive	38
8.1	Screenshot of the Conjoint Task and Focal Product Used for Price Optimization	43
8.2	Inferred Optimal Prices with or without Privacy Choice Indicator	51
8.3	Inferred Optimal Prices with or without Privacy Choices (by Consumer Subsets)	53
B.1	Example Questions in the First Stage	64
B.2	Displayed Compensation Schedule: Intrinsic Treatment	65
B.3	Displayed Compensation Schedule: Instrumental Treatment	66
C.1	Percentage of Participants Remained Throughout the Survey	67
F.1	Frequency of Data Sharing under Different Policy Regimes	70
F.2	Posterior Predicted Density of WTP in Intrinsic Preference	72
F.3	Sensitivity to Compensation with Rational Inattention	75

LIST OF TABLES

4.1	Demographics of Experiment Participants (Cleaned Sample)	23
4.2	Frequency of Data Sharing with Intrinsic Utility	24
4.3	t-Test for Equal Means ($H_1: E[D \mid \text{shared}] - E[D \mid \text{true}] \neq 0$)	26
6.1	Intrinsic and Instrumental Preference for Privacy: Estimation Results Comparison	34
6.2	Posterior Predicted Distribution of WTA in Intrinsic Preference	35
8.1	Profit Loss (\$/1,000 customers) with Different Prices for Data	46
8.2	Profit Loss When Using Firm Data (\$/1000 Consumers)	52
D.1	Posterior Estimates of Mean and Standard Deviation of the Intrinsic WTA	68
E.1	Posterior Distribution of WTA in Intrinsic Preference (with Heterogeneous Sensitivity to Income)	69
F.1	Privacy Preferences across Default Frames	71
F.2	Posterior Estimates of Parameters Associated with Intrinsic WTP Distribution .	73
F.3	Sensitivity to Economic Payoff across Default Frames (Pooled Estimation)	76

ACKNOWLEDGMENTS

I am deeply grateful to my advisors: Sanjog Misra and Pradeep Chintagunta, for their guidance and support throughout my PhD stage; Brad Shapiro and Oleg Urminsky, for showing me how to pursue research driven by intrinsic motivations while still succeed in a sea of instrumental incentives and constraints.

I feel fortunate to have participated in the NBER Economics of Digitization Boot Camp organized by Avi Goldfarb and Shane Greenstein. The Boot Camp helped me get a better sense of my own research identity by showing how researchers in the domain of digitization navigated their career; it has helped me a lot in organizing my research agenda and prioritizing projects since.

I would also like to thank Eric Anderson, Michael Dinerstein, Jean-Pierre Dubé, Avi Goldfarb, Wes Hartmann, Yufeng Huang, Garrett Johnson, Nitin Mehta, Kanishka Misra, Sarah Moshary, Harikesh Nair, Sridhar Narayanan, Aniko Öry, Anita Rao, Robbie Sanders, Avner Shlain, Valeria Stourm, Abigail Sussman, Anna Tuchman, Dana Turjeman, Kosuke Uetake, and Ken Wilbur. They provided valuable comments about the paper and encouraging words to me as a young researcher while I was on the job market. Meeting them and hearing their advice has made my experience truly enjoyable.

I would like to thank my friends, Huan Deng and Angela Denis, for their unconditional support in some of the most difficult times of my life; and my mother, Xiaoqin Zhang, who is always optimistic (though not without worry) for every step I have chosen to take. My journey wouldn't be the same without you.

ABSTRACT

In this paper, I separately measure two motives for consumers to protect privacy: an intrinsic motive, which is a “taste” for privacy; and an instrumental motive, which reflects the expected economic loss from revealing one’s private information to the firm. While the intrinsic preference is a utility primitive, the instrumental preference arises endogenously from a firm’s usage of consumer data. Combining a two-stage experiment and a structural model, I find that consumers’ intrinsic preferences for privacy range from 0 to 5 dollars per demographic variable, exhibiting substantial heterogeneity across consumers and categories of personal data. This rich heterogeneity in intrinsic preferences leads to a selection pattern that deviates from the “nothing-to-hide” argument predicted by a model with pure instrumental preferences. I then propose three strategies that firms and researchers can adopt to improve data-driven decisions when shared data are influenced by consumers’ dual privacy concerns. First, by using an experiment to measure the joint distribution of privacy preferences, firms can extrapolate selection patterns to cases where the data utilization method changes. Second, when the joint privacy preference distribution is unknown, data collection should focus on representativeness over quantity, especially when information externality is present. Lastly, firms can learn about the selection pattern in the shared data by leveraging information contained in consumers’ data-sharing decisions.

CHAPTER 1

INTRODUCTION

With the advent of privacy regulations across the globe, companies can no longer collect consumers' personal data without their explicit consent.¹ Across the EU, *General Data Protection Regulation* (GDPR) mandates firms to deliver transparent information and seek opt-in consent before data processing;² at the same time, it allows consumers to rectify, delete, or transfer their data to another firm.³ Outside Europe, the transparency and consent elements in GDPR have been emulated by many other countries and local states, including California in the US, Brazil, Chile, and India.⁴ They are also part of the "core principles" for the incoming US federal privacy legislation.⁵

When explicit consent is a prerequisite for personal data processing, consumers' **preferences for privacy** play a central role in determining what data and whose data firms are able to collect. For instance, in compliance with GDPR's consent seeking requirement, European websites lost 10%-12.5% of their recorded traffic and revenue as a result of increased consumer vigilance (Goldberg et al. 2019, Aridor et al. 2020). Prior literature shows that consumers exhibit heterogeneity in privacy choices when informed (Goldfarb & Tucker 2012b, Varian et al. 2005, Johnson et al. 2020). This heterogeneity leads to selection in the data shared by consumers. In the big-data era, selection in data has been a more pressing concern that affects the validity of data-driven technologies, including automatic resume sorting, voice and facial recognition, and public opinion polling. Firms need to understand the nature of consumers' privacy preferences in order to know how consumers' self-selection into data-sharing influences the quality of decisions based on

1. "Personal data" refers to any combination of data that can be used to identify a natural person. See <https://gdpr-info.eu/art-4-gdpr/>.

2. <https://gdpr-info.eu/recitals/no-39/>; <https://gdpr-info.eu/recitals/no-32/>.

3. <https://gdpr-info.eu/art-7-gdpr/>; <https://gdpr-info.eu/chapter-3/>.

4. <https://piwik.pro/blog/privacy-laws-around-globe/>.

5. See the "Individual consumer rights" section in the Privacy and Data Protection Framework proposed by the US Senate: <https://privacyblogfullservice.huntonwilliamsblogs.com/wp-content/uploads/sites/28/2019/11/Final-CMTE-Privacy-Principles-11.14.19.pdf>

these data.

To better understand the nature of consumers' privacy preferences and how they affect the insights developed from voluntarily-shared data, I empirically distinguish between two motives for protecting privacy. Privacy preference can emerge because privacy itself is valued as an **intrinsic** right (Warren & Brandeis 1890). It can also arise because of its **instrumental** value, namely, the economic payoff of preventing their private "type" from being revealed through data (Stigler 1980, Posner 1981). Consumers may hold both types of privacy preferences. Intrinsically, most people find it "creepy" to have smart thermostats tracking their activities at home, regardless of whether their behaviors are benign or objectionable (Pew Research Center 2015). Instrumentally, reckless drivers can be unwilling to install telematics offered by insurance companies that monitor their driving habits (Soleymanian et al. 2019, Jin & Vasserman 2018).

Empirically distinguishing intrinsic and instrumental motives is useful for two reasons. First and foremost, accounting for their coexistence is crucial for understanding how consumers self-select into sharing their data to firms, and the representativeness of the resulting dataset. The "if you've got nothing to hide, you've got nothing to fear" doctrine is only valid when people harbor a purely instrumental preference for privacy. On the other hand, assuming consumers value privacy only because of intrinsic motives may result in the misleading conclusion that people who value privacy more are no different from the rest of the population.

Moreover, empirically separating these two motives can help us extrapolate existing preference measurement results across economic situations. The endogenous nature of instrumental preference implies that it depends on how the firm uses consumer data. As such, privacy choices can shift with the (perceived) purpose of data collection, the performance of a firm's model used for processing the data, and what other data the firm already obtains. Even without actual changes, information that prompts consumers to update their beliefs about any aspect above will cause their data sharing decisions

to change accordingly. Separating the instrumental preference from the intrinsic allows us to see that these changes are due to incentives instead of behavioral anomalies, and model them accordingly. Accounting for the endogenous nature of instrumental preference is particularly crucial for calculating equilibrium outcomes of privacy regulations and managerial strategies. Together, understanding the selection in data resulted from heterogeneous privacy preferences and its endogenous nature is crucial for collecting and analyzing consumer data to gain better insights.

The goal of this paper is showing how to improve data collection and inference methods by accounting for the dual privacy preferences. To achieve this goal, I first introduce the conceptual framework using a model of consumer disclosure with two-dimensional preferences. This simple model clarifies the distinction between intrinsic and instrumental preferences, and shows how they jointly determine to what extent the “nothing to hide” argument (i.e. adverse selection) can be refuted. In particular, advantageous selection will occur if intrinsic preferences for privacy are sufficiently heterogeneous and if they are negatively correlated with instrumental preferences among consumers.

I design an experiment that measures revealed preferences for privacy in terms of dollar values, then characterizes the preference heterogeneity. Revealed preferences are solicited by requesting consumers to share sensitive data in a market research survey. To capture preference heterogeneity and its impact on selection in shared data, I use a novel two-stage design, which sequentially records consumers’ private types and their privacy choices. This design enables me to observe the contents of personal information even for consumers who choose not to share their personal data, and is the key to characterizing the relationship between privacy choices and data selection. The experiment then generates three layers of variation needed to identify the model: (a) The intensity of instrumental incentives allows me to separate the two preference components; (b) the amount of compensation enables me to calculate the dollar values of privacy preferences; (c) the default-choice condition permits comparison of privacy choices in different policy

regimes. The experiment also contains a conjoint survey, which allows me to calculate the information value of personal data for price optimization in the counterfactual analysis.

Data and structural estimation results reveal the following findings. First, intrinsic preferences are highly heterogeneous across both consumers and categories of data. Without instrumental incentives, the mean willingness to accept (WTA) to share personal data ranges from \$0.14 to \$2.37; however, at the 95% quantile, consumers value each personal variable from \$2.19 to \$5.08. To accommodate this heterogeneous WTA and obtain a representative set of data, a firm will need to pay as high as \$29.72 per consumer per demographic profile. Second, consumers' privacy choices respond to the economic consequences of revealing their personal information when information about data usage is plain and transparent, as is required in the GDPR regime. Moreover, consumers' beliefs and corresponding choices are approximately consistent with the actual data usage. This observation provides useful guidance for projecting changes in consumer privacy decisions when firms change their data utilization strategy. Third, intrinsic preference substantially moderates the direction and degree of selection due to its heterogeneity. In the experiment, the intrinsic motive plays a dominant role in determining the direction of consumer self-selection, even when intrinsic and instrumental preferences have the same magnitude on average. As a result, self-selection into data sharing is not always adverse, deviating from the "nothing to hide" doctrine and predictions offered by canonical economic models.

Based on these findings, I develop three approaches to improve the collection and analysis of consumer data that acknowledge the duality of privacy preferences. I begin by describing how the experiment can be replicated in the field, taking into account the institutional and informational constraints absent in my experiment setting. By measuring consumers' intrinsic and instrumental preferences, firms can learn about the nature of selection in shared data, which allows them to design more efficient data acquisition strategies. Then I use counterfactual analysis to demonstrate how data acquisition can

be improved when this joint distribution is unknown, in the context of price targeting. I show that firms can improve the efficiency of their data-acquisition strategy by sampling consumers where information externality is present. Lastly, I document that with dual privacy preferences, the information value of privacy choices (the decision of whether to share data) is different from its value predicted by a classical economic model. By interpreting privacy choices correctly, the firm can diagnose the nature of sample bias and improve aggregate-level inference.

Although the following analysis uses “firms” to refer to the data collector, the qualitative findings apply to any organization that wants to obtain insights from data voluntarily shared by consumers. These organizations can include research institutions, public opinion poll providers, and policy think tanks.

This paper contributes to several strands of the literature. First and foremost, it formalizes and extends Becker’s (1980) dual-privacy preference framework (also see Wathieu & Friedman 2009, Farrell 2012, and Choi et al. 2019 which adopts a similar treatment). Both Becker and Farrell distinguish the valuation for privacy as a “final good” and the valuation for privacy induced by the economic consequences of data sharing. Similarly, Tirole (2020) describes a model where citizens’ actions are influenced by a combination of intrinsic preference for dissent, and a reputation concern which is endogenous. Compared to these papers, my paper documents how the coexistence of intrinsic and instrumental preferences determines the selection pattern in data shared by consumers, and how it subsequently affects the quality of data-driven decisions based on shared data. In doing so, my paper builds the link between consumers’ privacy preferences and the quality of consumer data as firms’ input.

Second, my paper builds on existing work that measures revealed privacy preference as a whole, including Goldfarb & Tucker (2012*b*), Athey et al. (2017), Kummer & Schulte (2019), and (Acquisti et al. 2013, Benndorf & Normann 2018, Tang 2019) which also provide dollar values for privacy preferences. Compared with these papers, mine separately

measures intrinsic and instrumental preferences. Given the endogenous nature of instrumental preference, separating these two components is useful for extrapolating existing measurement results to different contexts. In doing so, my paper also develops a replicable method to measure heterogeneous privacy preferences.

My paper also contributes to the literature on context-dependent privacy preferences by highlighting how instrumental incentive induces responses to contexts, such as entities that access the data (Martin & Nissenbaum 2016) and information that changes the perceived usage of the data (John et al. 2010, Athey et al. 2017, Miller & Tucker 2017). As such, it complements the previous literature (Egelman et al. 2009, Acquisti et al. 2012, 2013, Adjerid et al. 2019, Lee 2019), which emphasizes psychological factors that generate context-dependence.

Last but not least, by discussing how consumers' privacy choices affect firms' inferences and resultant profits in the new policy regime, my paper adds to the research on how privacy regulations influence firms' managerial outcomes, including the effectiveness of advertising (Goldfarb & Tucker 2011, Tucker 2014), funds raised (Burtch et al. 2015), innovation activities (Goldfarb & Tucker 2012a, Adjerid et al. 2015, Jia et al. 2018, 2020), and profits (Johnson 2013, Johnson et al. 2020). My paper focuses on one mechanism, namely, how consumers' self-selection into data-sharing affects the quality of firms' data-driven decisions. In doing so, I also develop strategies that allow firms to address the impacts of selection.

The paper proceeds as follows. Chapter 2 introduces the conceptual framework, clarifies definitions, and illustrates the implications of the dual-privacy-preference framework. Chapter 3 describes the experiment design, and Chapter 4 provides descriptive evidence of intrinsic and instrumental preferences. Chapter 5 and 6 present the structural model and estimation results. Chapter 7 describes how the experiment can be replicated in the field. Chapter 8 describes the counterfactual analysis, and Chapter 9 concludes.

CHAPTER 2

THE CONCEPTUAL FRAMEWORK

This chapter uses a stylized model to clarify the distinction between intrinsic and instrumental preferences for privacy. It describes how the instrumental motive is endogenously derived from the economic context, how consumers self-select into sharing or protecting their personal data, and how this selection pattern differs from predictions generated by models that assume monolithic privacy preference.

To illustrate the key elements, for now I assume that consumers have rational expectation, and use the same notation to represent the actual economic payoff and the payoff perceived by consumers. I also assume that the firm does not have other information about consumers before requesting their data. Later in Section 2.4, I discuss the robustness of the results when these two assumptions are relaxed. In the empirical analysis, I directly measure consumers' beliefs and examine whether their degree of sophistication indeed generates the selection pattern predicted by the stylized model.

2.1 Setup

Consider a firm that sells a product or service to many consumers. Consumers have different types, which calls for customized offers; denote consumer i 's type as d_i . The firm requests personal data from consumers in order to know their types. At a later stage, the firm gives customized transfer $T(\cdot)$ to consumer i , which maximizes the firm's expected profits conditional on the firm's understanding of i 's type. For example, $T(\cdot)$ can be price discount while d is price sensitivity; or $T(\cdot)$ can be the annual limit in an insurance contract while d is risk type. To encourage data sharing, the firm may incentivize consumers using compensation, denoted as P . Examples of compensation include perks offered to consumers who sign up for a loyalty program, or gift cards for sharing email.

Consumer i owns personal data that can reveal their type. Without loss of generality,

assume the shared data perfectly reveal their type. Therefore, a one-to-one correspondence exists between the content of personal data and a consumer's type.¹ We can always construct consumer types such that the transfer $T(d)$ is monotonic in d . For example, suppose that d is age and that the middle-aged group has the lowest price sensitivity, followed by the older and then the youngest. Then we can label the middle-aged group as $d = 1$, the older group as $d = 2$, and the youngest group as $d = 3$. Without loss of generality, I define consumer types such that $T(\cdot)$ is increasing in d , and refer to consumers with higher d (who can obtain higher economic payoffs upon revealing their type) as the high type.

Consumers decide whether to share their data with the firm. $s_i \in \{0, 1\}$ indicates whether i decides to share d_i : $s_i = 1$ means the data is shared. For consumers who share no information, the firm forms beliefs about their types and chooses the amount of transfer accordingly: $T(s = 0) = T(F_d(d|s = 0))$, where $F_d(d|s = 0)$ is the distribution of consumer type conditional on the consumer choosing not to share his data. For consumers who share, the transfer can condition on the content of the data, written as $T(d_i)$.

2.2 Privacy Preferences

A consumer has an intrinsic motive for privacy c_i , which is a taste for protecting his data regardless of the economic consequences induced by revealing his type. He also has an instrumental motive for privacy, namely, the expected economic gain from not revealing his type: $\Delta T(d_i) \equiv T(F_d(d|s = 0)) - T(d_i)$. For example, suppose $T(\cdot)$ is price discount and d is age. If older consumers have higher price sensitivity, the firm will choose to give them

1. One way to understand the one-to-one correspondence assumption between data content and consumer type is the following. Suppose the data do not perfectly reveal the exact level at which the consumer values the product, but indicates a range for the consumer's valuation; in this case, we define the range of valuation as his type. In other cases where two income levels correspond to the same level of product valuation, we can code the two income levels as having the same value. Another way to understand the model is that some firms only care about the predictive performance of the pricing model but not consumers' price sensitivity per se; in this case, a direct mapping occurs from d_i to $P(d_i)$, and the intermediary "type" is unnecessary.

higher discounts upon learning their age. Anticipating this outcome, the instrumental preference for older consumers to protect their personal data will be low.

The key distinction between intrinsic and instrumental preferences is whether they are induced by the consequences of revealing one's private information to the firm.² The intrinsic preference is a utility primitive, which persists regardless of the market environment and the consumer's "type" relevant to this market. By contrast, the instrumental preference is endogenously derived from the economic environment; thus, it changes with the payoff function $T(\cdot)$ as well as his type in this particular market d_i . The intrinsic preference can also be correlated with a consumer's type. However, his instrumental preference for privacy changes with the (perceived) usage of data $T(\cdot)$, such as the purpose of data collection and the technology used for processing data; his intrinsic preference does not.³

Instrumental preference and the utility from compensation are also distinct constructs, even though both are derived from economic payoffs. The instrumental motive reflects the value of private information. It is a function of the hidden type that the firm cares about, that is, information about the consumer that can shift the level of optimal transfer between the firm and consumers. On the other hand, compensation does not necessarily hinge on a consumer's type; it is more properly viewed as the price for personal data.

2.3 Who Chooses Not to Share Personal Data?

A consumer shares data iff the privacy cost is offset by the compensation that the firm provides:

$$s_i = 1 \text{ iff } -c_i - \Delta T(d_i) + P > 0. \quad (2.1)$$

2. Under this taxonomy, costs that may at first glance appear instrumental still belong to the "intrinsic" bin (e.g. concern for identity theft, because this concern is common for all consumers).

3. This statement does not contradict the fact that intrinsic preference can be shifted by psychological factors. Rather, it distinguishes psychological and economic shifters, which have different impacts on the expression of the two preference components.

Firms often want to learn the characteristics of consumers choosing not to share their data, in order to optimize the transfer to consumers that maximizes profits. A model that assumes privacy preferences to be purely instrumental will generate the following prediction: Only low types choose to withhold their data in equilibrium, because these are the consumers who incur a larger loss upon sharing data (Grossman & Hart 1980, Milgrom 1981, Jovanovic 1982). This reasoning is the underpinning of the “nothing to hide” statement. Alternatively, a theory that assumes privacy preferences to be pure intrinsic may fail to capture the nuance of consumers’ self-selection into sharing.

The dual-preference framework paints a more nuanced picture of how consumers self-select into sharing personal data. The intrinsic preferences for privacy are likely to be heterogeneous among consumers. This heterogeneity should change the firm’s inference $F_d(d|s = 0)$ because nondisclosure no longer signals low-type customers unambiguously. The degree to which privacy choice reveals information about a consumer’s type depends on both the relative heterogeneity of the intrinsic preference and its correlation with the instrumental preference. This is formally characterized by the proposition below (see the proof in Appendix A).

PROPOSITION 1. *Denote the standard deviation of intrinsic and instrumental preferences respectively as σ_c and σ_t , and their correlation coefficient as ρ . The following conclusions hold:*

(A) *In data shared with the firm, sample selection goes in the same direction as predicted by a model with pure instrumental preference iff $\rho + \frac{\sigma_t}{\sigma_c} > 0$.*

(B) *Privacy choice is more indicative of a consumer’s type d_i when $\frac{\sigma_t}{\sigma_c}$ is higher.*

To illustrate this proposition, suppose older consumers (who would have obtained better discounts upon sharing their age) care more about privacy intrinsically, and that the intrinsic preference is highly heterogeneous compared with the instrumental. Then the intrinsic preference will play a dominant role in privacy decisions: On average, consumers who choose not to share their data are more senior and should receive more generous discounts. This pattern forms a stark contrast to the case with a pure instrumental

preference for privacy.

In sum, the dual presence of intrinsic and instrumental privacy preferences has two main implications. First, although the intrinsic preference is a utility primitive, the instrumental preference is endogenously determined by the market environment. This fact explains why preferences for privacy vary across the contexts of data used, who gets access to the data, and what data are requested. Second, when the intrinsic preference for privacy is heterogeneous, privacy choice no longer unambiguously signals a specific type of customer. The more the heterogeneous intrinsic preference is relative to the instrumental, the less we can assume a consumer's type based on his privacy decisions. Accounting for this fact is essential for analyses based on voluntarily contributed personal data.

2.4 Extending the Stylized Framework

Boundedly rational consumers. The selection pattern predicted above persists as long as consumers' belief about $T(d_i)$, that is, how the payoff depends on their type once this information is revealed to the firm, is correct. In particular, the prediction holds even when consumers do not have the correct belief about $T(F_d(d|s = 0))$, the payoff that the firm applies to consumers who refuse to share data. Intuitively, it is the difference in $T(d_i)$ among consumers that generates the adverse selection pattern caused by instrumental preference. We can think of correct belief about $T(d_i)$ as a first-order sophistication, and correct belief about $T(F_d(d|s = 0))$ as a higher-order sophistication which requires correct expectation about the firm's as well as other consumers' rationality and the latter's preference distribution. In the empirical analysis, I directly measure consumers' beliefs and focus on examining whether they have the correct first-order belief.

Firm's existing knowledge about consumers. In cases where the firm already has existing knowledge about consumers, the magnitude of instrumental preference depends on how much more the firm can learn about consumer types from the personal data, while the intrinsic preference stays still. In an extreme example, if a firm already learns perfectly

about consumers' purchase habits and still requests to record their activities at home, all privacy concerns will come from intrinsic motives.

Data that improves the horizontal match value. In the model above, consumers' types are thought of as purely vertical (i.e. "high" vs. "low"). Cases abound where the firm collects data to improve the horizontal match value, such as designing a better recommendation engine. Conceptually, the perceived benefits from improving product or service quality is a form of compensation, rather than a privacy preference component. In terms of its impact on the inference problem, if this perceived benefit is heterogeneous among consumers and unknown to the firm or researcher, then it plays the same role as the heterogeneous intrinsic preference does in affecting the selection in shared data.

CHAPTER 3

THE EXPERIMENT

The experiment serves three purposes. First, it generates variation crucial for separately identifying intrinsic and instrumental utility parameters. Second, it provides a setup that sequentially records consumers' private information and privacy choices, which enables me to characterize preference heterogeneity and selection in shared data. Lastly, it creates choice environments that match key features of the new privacy regulatory regime. In this chapter, I first introduce the empirical challenges and explain how my experiment addresses them, then describe key elements of the new policy regime that the experiment intends to examine. This is followed by the introduction of experiment design and implementation.

3.1 Empirical Challenges and Solutions

To understand how consumers self-select into different privacy choices, one needs to separately measure the intrinsic and instrumental preferences. However, doing so using observational data is difficult. First, the economic incentive is usually fixed in observational settings, making it infeasible to separate instrumental preferences from the intrinsic. Second, in most observational settings, the request for personal data is bundled with product provision. As a result, the preferences for privacy are confounded with the preferences for products concurrently offered. For example, consumers may keep using Gmail even after learning that Google analyzes all their email texts, due to either their low preferences for privacy or their high valuation of Gmail service. The last and the most difficult challenge is that both consumers' privacy choices and their private types need to be observed to identify my model, yet privacy choices are precisely the decision concerning whether to reveal these private types. If this challenge is not accounted for, the collected data will exhibit self-selection as long as variation in privacy decisions exists.

To circumvent these problems, I design an experiment that includes three main features. First, instrumental incentives are turned on or off across different treatment conditions. I can thereby measure the intrinsic preferences directly when the instrumental incentives are off, and use the difference between treatments to measure the instrumental preferences. Second, I exclude the confound from product preference by using monetary incentives (which have known values) to compensate for data sharing. Furthermore, the amount of compensation to encourage data sharing varies across treatments, allowing me to measure the dollar values of privacy preferences. To overcome the last challenge, I adopt a novel two-stage design, where the first stage collects participants' private information, and the second stage solicits revealed preferences for privacy.

3.2 Examination of the New Policy Regime

It is important to measure privacy preferences in a relevant choice environment given their context-dependent nature. To this end, my experiment specifies a choice environment that features key elements common in recent privacy regulations and principles, described below.

Transparency of data usage. Recent privacy regulations and principles require firms to deliver plain and accessible information about data collection and its purpose. For example, GDPR requires data controllers and processors to use “clear and plain language” to describe the purpose of data processing and consumer rights.¹ CCPA requires firms to give a “visible and accessible” notice at or before data collection, describing what data will be collected and the corresponding usage.² To match this element, my experiment explains clearly the usage and flow of the data, and explicitly notifies participants about their options related to data sharing.

Consumer control and consent. This is represented by various rights clauses in

1. <https://gdpr-info.eu/recitals/no-39/>.

2. <https://www.alston.com/en/insights/publications/2019/10/the-draft-ccpa-regulations>.

major regulations, such as *the right to know, the right to deletion, the right of access, and the right of data portability*.³ One key component in these rights clauses is the explicit consent requirement, which is implemented differently across regulations in terms of the default action. In particular, EU laws (GDPR and ePrivacy Regulation) requires opt-in consent, while practices in the US are mixed.⁴ Regardless of the regulation in force, requests effectively operate in an opt-in condition for data that are not generated by default, such as survey responses, tests, and membership sign-ups. In view of the ambiguity in regulatory focus and the empirical relevance of the opt-in regime, my experiment includes both opt-in and opt-out conditions, but the empirical analysis will focus on the opt-in condition. I compare privacy choices in different consent regimes and discuss the underlying mechanism in Appendix F.1.

3.3 Experiment Design

The experiment uses a survey as an instrument, but solicits revealed preference instead of stated attitude. This is achieved by including personal questions with varying degrees of sensitivity: A participant's decision to share the response to a question indicates his level of privacy cost associated with this personal variable. This technique has previously been deployed by Acquisti et al. (2012) and Goldfarb & Tucker (2012b). Research shows that in the domain of privacy preferences, attitude- and behavior-based measures often disagree (Harper & Singleton 2001, Spiekermann et al. 2001). I focus on revealed preference because it is not only incentive-compatible, but also more relevant than attitudes for managerial

3. See <https://www.bakerlaw.com/webfiles/Privacy/2018/Articles/CCPA-GDPR-Chart.pdf> and <https://www.eff.org/deeplinks/2019/12/sen-cantwell-leads-new-consumer-data-privacy-bill>.

4. Although many existing laws still adopt opt-out consent, opt-in has been used for regulating more sensitive data. Examples include Illinois' *Biometric Information Privacy Act*, the *Fair Credit Reporting Act* on medical data, CCPA on minor's data, the *Washington Privacy Act* for biometrics and geolocation data, and the *Consumer Online Privacy Rights Act*. Whether opt-in or opt-out should become the new regulation standard has been one of the debate topics in recent regulatory discussions. See <https://www.commerce.senate.gov/public/index.cfm/2019/2/policy-principles-for-a-federal-data-privacy-framework-in-the-united-states> and <https://fpf.org/2020/02/25/close-to-the-finish-line-observations-on-the-washington-privacy-act/>.

decisions and policy analysis. In addition, I avoid using a Becker-DeGroot-Marschak mechanism, which is shown to produce results closer to stated attitude than revealed-preference when used for measuring privacy preferences (Benndorf & Normann 2018).⁵

The experiment consists of two stages. In stage one, participants see themselves participating in a market research survey sent by the University of Chicago. The survey includes conjoint questions about smartwatch attributes, and their intent to purchase a digital device in the near future. These are followed by demographic questions, including gender, age, education level, income, relationship status, whether they have children, zip code, and ethnicity. Each personal question in the first stage includes a “prefer not to say” option; people who find the question too sensitive to answer are thus allowed not to respond rather than being forced to fabricate a response. Appendix B.1 shows examples of the conjoint and demographic questions.

Stage one serves two roles. The first is to record private information from consumers, including those choosing not to share data in the subsequent stage. This full information allows me to measure heterogeneity in privacy preferences, and characterize how the interplay between intrinsic and instrumental motives determine selection in shared data. Second, the conjoint questions provide inputs for calculating the value of data to firms in a pricing context, which becomes the basis for comparing the effectiveness of data collection and analyzing strategies in the counterfactual analysis. The conjoint questions also disguise the real purpose of the survey so that participants are not prompted to consider privacy.

Stage two solicits privacy choices. After finishing the survey, participants are directed to a new screen. Here, they are requested to share survey responses with a third party, which is a smartwatch manufacturer who wants to use the data to inform its product-design decision. Participants can choose whether to share each personal information

5. Although BDM is incentive-compatible in theory, its incentive structure can be hard to understand from participants’ perspectives in practice. Other potential sources of the measurement gap include differences in contextual cues, and forced attention in BDM caused by repeated testing participants’ understanding of the incentive structure.

variable separately via check boxes.⁶ Data sharing is encouraged by compensation in the form of a gift-card lottery. Participants are not told about the data-sharing step until they answer all questions in stage one; once consumers reach the second stage, the “return” button is disabled, preventing them from deliberately changing previous responses to facilitate sharing. These two features, together with the presence of the “prefer not to say” options in Stage one, are included to ensure responses in the first stage are truthful.

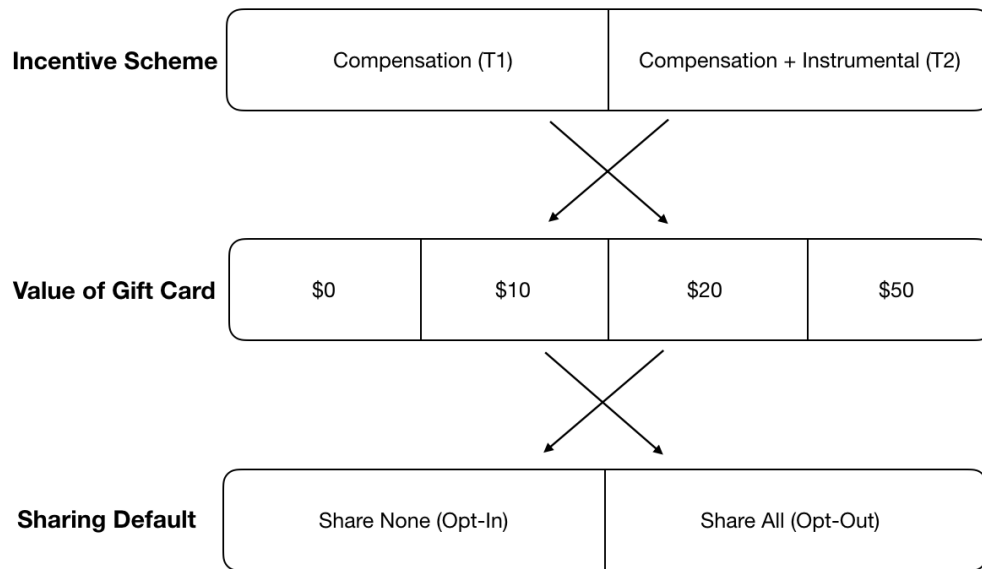
Stage two is also where all treatments take place. Figure 3.1 displays the three layers of treatments: the incentive scheme, the amount of compensation, and the sharing default. These treatments are orthogonal to each other.⁷ The first layer varies the incentive scheme across treatment groups:

- Treatment 1 (pure compensation): The amount of compensation increases proportionally to the amount of data shared and is common across all participants. In other words, the price for data is the same regardless of what the firm learns about the consumer. In particular, sharing one additional personal variable increases the probability of winning the gift card by one percentage point.
- Treatment 2 (compensation + instrumental incentive): A baseline level of compensation exists and takes the same form as in Treatment 1. The amount of compensation is subsequently adjusted based on whether the company perceives the participant to be a potential customer based on the data it obtains. Participants who are more likely to be their customers receive higher compensation than the baseline, whereas unlikely customers receive a lower amount. Participants are told the company’s target customers are high-income people who intend to buy a digital product, and therefore, they will receive more if the shared data indicate they fit this profile.

6. Only first-stage responses that are informative (responses other than “prefer not to say”) are allowed to be shared in stage two.

7. There is one exception. By design, participants who receive zero compensation do not receive any instrumental incentives.

Figure 3.1: Treatment Design



Note: The three layers of treatments are orthogonal to each other. Treatments are assigned with equal probability in each layer.

Appendix B.2 displays the information shown in each treatment. Overall, the incentive scheme is presented in a transparent and clear manner. The incentive scheme is displayed in two parts. The main page explains who collects the data and for what purpose, and how a participant's payoff will qualitatively depend on the data shared. The detailed screen shows quantitatively how the payment is calculated, and is accessible when a participant clicks the "see details" link. This design is similar to the format of most post-GDPR website banners.

In sum: Privacy choices in Treatment 1 alone identify intrinsic privacy preferences. Here, the stated purpose of data collection does not imply continuous tracking or any other future interactions with consumers. Moreover, this company is previously unknown to the firm; thus, it is unlikely that participants anticipate the instrumental consequences of sharing data from interacting with the firm in the future. By contrast, choices in Treatment 2 are motivated by both intrinsic and instrumental preferences. The instrumental preferences are induced by an incentive scheme that depends on a participant's income

and product-purchase intent. These two characteristics constitute a consumer's "type" in this experiment. Therefore, the differential responses between Treatments 1 and 2 identify instrumental preferences for privacy.

The other treatments are designed as follows. The second treatment layer changes the value of the gift card (essentially cash) across participants, creating variations for measuring the dollar values of privacy preferences. The third layer varies default choice, which is set to either sharing all data (opt-out) or sharing none (opt-in). Within each layer, treatments are assigned with equal probability.

To measure if participants understand and trust the validity of incentive treatments, a set of follow-up questions are prompted after participants make the data-sharing choices. These questions include the perceived purpose of the study, what determines the amount of expected compensation, the reasons they choose (not) to share the survey responses, and if they prefer a sure reward with the same expected value as the gift-card lottery.

3.4 Discussion

Using a controlled field experiment allows me to design a control group that measures intrinsic preference in a relatively clean manner. In real business settings where the firm is known to consumers, consumers' expectations are likely to fixate on how the firm usually uses their data, thus always having some instrumental preference. In Chapter 7, I discuss how researchers and firms can run a different version of my experiment in the field to separate intrinsic and instrumental preferences when combined with an assumption on consumers' belief stability.

The experiment uses type-dependent monetary compensation instead of personalized product prices to induce the instrumental incentive. Although the latter is more natural, it will not induce variations of instrumental preference in my setup. Given that participants have never interacted with the featured company (it is fictitious), they may not plan to engage in future transactions with this company. In this case, the firm's pricing practices

will not matter to them.

Using a lottery instead of sure rewards for compensation may bias preference measurement if participants predominantly have the same risk preference. If participants are risk-averse, their perceived gain from the gift-card lottery will be lower than its objective expected value, and the estimated dollar value of privacy preferences will be an upper bound of their true valuation; the opposite holds if participants are risk-seeking. In the follow-up survey question, 35% of the participants prefer the lottery, while the rest prefer the sure reward. This pattern suggests that risk aversion is not a dominating feature.

Consistent with the conceptual framework, the experiment focuses on the case where consumers cannot garble their personal information sent to the firm. Cases abound where consumers' personal data are truthfully recorded as long as they opt in to share, such as location tracking, browsing history tracking, and genetic testing. Garbling information is technically possible in some cases, but involves a high cost in practice and is usually adopted by only the most tech-savvy consumers. One can extend this framework by introducing heterogeneous costs of data fabrication as the third dimension in consumers' preferences. The measurement results in this paper serve as a useful building block for such extensions.

CHAPTER 4

DATA AND DESCRIPTIVE EVIDENCE

In what follows, I describe the data source and sample characteristics, and then present model-free patterns of intrinsic and instrumental preferences. The main analysis focuses on privacy choices in the opt-in regime. Data show how consumers purposefully share some data while protecting others, how the economic context changes the composition of consumers that share data, and how this compositional shift changes the quality of data shared. Section F.1 compares privacy choices in different default regimes, and discuss its implications for policy design and welfare analysis.

4.1 Data Source and Cleaning

Participants of the experiment come from Qualtrics Panels. To the extent that Qualtrics panel members are more willing to share personal information with others without anticipating any instrumental consequences, the measurement result provides a lower bound for the population-level intrinsic preferences. Nevertheless, existing work finds that the Qualtrics panel is more representative than alternative online panels in terms of demographics and political attitudes (Heen et al. 2014, Boas et al. 2018). To further reduce possible discrepancies, stratified sampling is applied so that the demographics of participants entering the survey resemble the distribution given by the 2018 US census. Qualtrics provides three demographic variables on the back end, including income, age, and ethnicity. I use these data to validate the truthfulness of responses in the first stage. Not all demographic variables I intend to collect are available through Qualtrics. Therefore, having the first stage is still necessary.

A total of 4,142 participants enter the survey; 3,406 of them proceed to the data-sharing-request stage. For people who leave the survey upon seeing the request, I code their choices as sharing nothing, regardless of the default condition. Figure C.1 shows

the participant attrition throughout the experiment. Among the 18.4% of participants who leave the survey before seeing the treatment, 91% exiting occurs before or during the conjoint survey. This pattern indicates that attrition is mainly caused by a lack of interest in the conjoint questions, rather than a reluctance to share personal data in the first stage.

To prevent treatment contamination, I deduplicate the respondents by IP address.¹ I also exclude respondents whose time spent on the survey, or time spent in responding to the data-sharing request is at the lowest decile. The cleaned data include 2,583 participants, comparable to other large-scale experiments that study consumers' utility from digital consumption, such as Brynjolfsson et al. (2019) and Allcott et al. (2020).

4.2 Sample Characteristics

Attrition and sample cleaning may change the characteristics of the final sample. Table 4.1 summarizes the demographics of survey participants in the cleaned sample, and compares them with the 2018 Current Population Survey (CPS) whenever similar statistics are available. Some discrepancies come from differences in counting. For example, the mean age provided by CPS includes juniors (ages 15–18), whereas my sample contains only adults; “black” in my sample includes mixed-race groups, while CPS’s definition excludes it. Another difference comes from the fact that some participants choose not to share all demographics during the first stage. As a result, the percentages of different income levels do not sum up to 1, whereas in the census, the disclosure is complete. Compared to the population, participants who finish the survey tend to be female, less educated, and have lower incomes.

Purchase intent is one of the consumer types in the instrumental-incentive treatment. It is calculated based on participants’ responses to two questions in the first stage: (A) “How likely will you buy a new smartwatch within the next 3 months?” (B) “How likely

1. For respondents using the same IP address, I keep the first response when the finishing time of the first respondent does not overlap with the starting time of the second respondent. If these times overlap, I discard both responses.

Table 4.1: Demographics of Experiment Participants (Cleaned Sample)

	Variables	Experiment Sample	2018 Census
	Female	65.31%	50.80%
	Married	47.39%	51.16%
	Have young kids	24.78%	–
	Mean age	47.60 (16.89)	45.9 (–)
Education	High school degree or less	47.00%	39.93%
	College degree	40.65%	48.67%
	Master’s degree or higher	11.39%	11.40%
Race	White	71.27%	76.60%
	Black	15.37%	13.40%
Annual Household Income	\$25,000 or less	21.99%	20.23%
	\$25,000 to \$50,000	29.54%	21.55%
	\$50,000 to \$100,000	30.12%	28.97%
	\$100,000 or more	13.55%	29.25%
No. Observations		2,583	–

Source of the census data: U.S. Census Bureau, Current Population Survey, 2018 Annual Social and Economic Supplement. “–” indicates that no corresponding statistics are available.

Note: For discrete variables, values in the survey are collapsed into larger groups to facilitate the exhibition. Numbers corresponding to the same category may not sum to 1, given that smaller groups are left out and that some participants choose not to respond in the first stage. For continuous variables, mean values are reported with standard deviation in parenthesis.

will you buy any other digital devices within the next 3 months?” Each question uses a 5-point Likert scale. Different answers are then given different scores. For example, “extremely likely” is scored 2, while “extremely unlikely” is scored -2. Purchase intent is then constructed by summing up these two scores; a higher value indicates higher purchase intent. Across participants, the mean purchase-intent score is -0.17, with a standard deviation of 1.72.

4.3 Intrinsic Preferences

Table 4.2 shows how the frequency of sharing varies with compensation and the category of personal data in Treatment 1 (pure compensation group). Consumers do not want to share personal data when not compensated: in the first column, the frequency of data

being shared are all at or below 50%, which is the indifference benchmark.

Table 4.2: Frequency of Data Sharing with Intrinsic Utility

Compensation	Category of Data								
	Gender	Age	Edu	Income	Relationship	Kids	Zip	Race	Purchase Intent
= 0	0.50	0.47	0.43	0.36	0.46	0.29	0.41	0.42	0.43
> 0	0.70	0.68	0.62	0.56	0.66	0.53	0.63	0.63	0.54

Note: “Relationship” corresponds to their responses about marital status. “Kids” corresponds to responses to the number of children they have. Among the compensated groups, the value of gift card is \$33 on average, with a 1% increase in the possibility of winning for each variable shared.

Compensation is effective in shifting privacy decisions. An average price of 33 cents per variable increases the probability of sharing by about 20% across variables. However, this average response among participants masks preference heterogeneity, which is crucial for understanding the impact of privacy decisions on the quality of shared data. I revisit preference heterogeneity in the estimation result chapter.

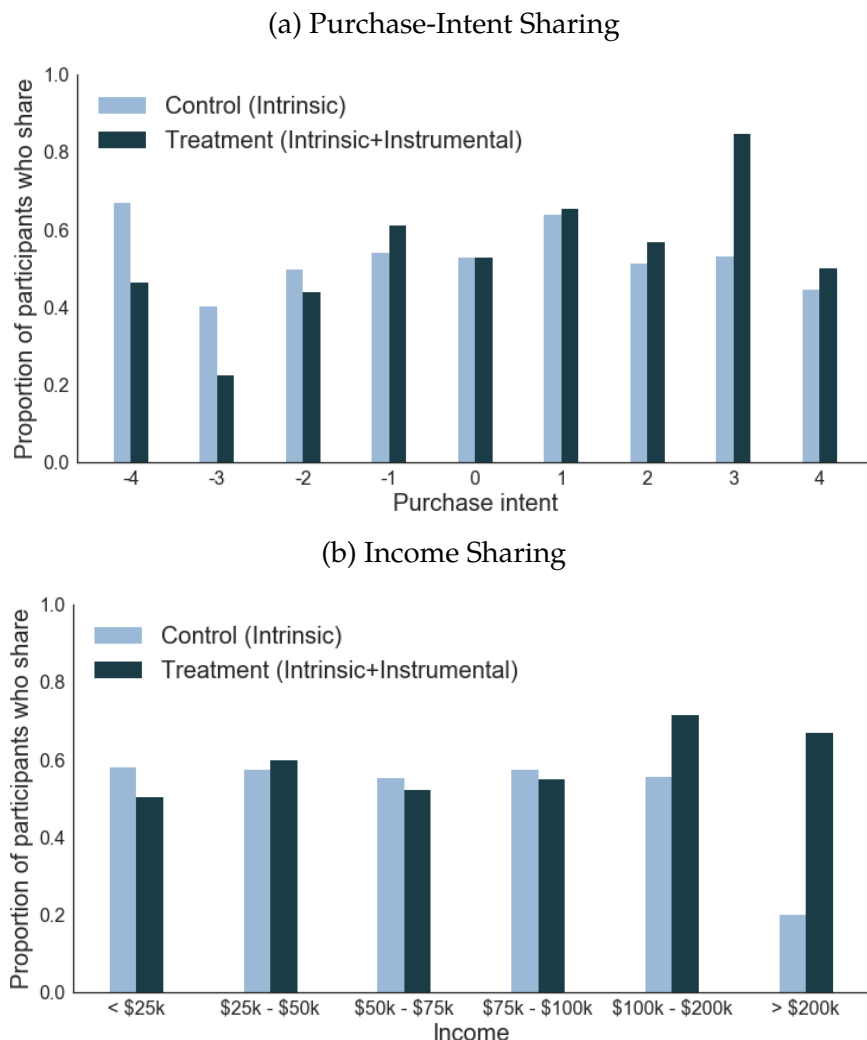
On the other hand, different data are valued quite differently, and the sensitivity ranking across personal variables remain largely unperturbed regardless of whether data sharing is compensated. Data about household income and about their children are valued the most, whereas gender is viewed as the least sensitive. Overall, the table shows that participants make attentive trade-offs in the experiment, and that different data are valued differently by consumers.

4.4 Instrumental Preferences

Treatment 2 introduces the instrumental incentive: Participants benefit more if they are perceived as wealthy or intend to buy digital products in the short term (hereafter “high types”). Figure 4.1 shows how instrumental incentives influence privacy choices and how this influence is moderated by intrinsic motives. Panel (a) plots the proportion of participants choosing to share their purchase intent data across purchase intent cohorts for each incentive treatment. High-type consumers are more willing to share in Treatment

2 than in Treatment 1, whereas for low-type consumers, the reverse pattern occurs. This pattern indicates that participants are attentive to the instrumental incentive when it is displayed in a plain and clear manner as required by GDPR-style regulations.

Figure 4.1: Frequency of Data Sharing across Incentive Treatments



Note: Intrinsic Utility = Treatment 1; Intrinsic + Instrumental Utility = Treatment 2. Frequency is calculated as the proportion of participants who share their income data within each income cohort (not across). The sum of bar heights can be greater than 1.

Purchase intent is calculated based on participants' responses to two questions in the first stage: (A) "How likely will you buy a new smartwatch within the next 3 months?" (B) "How likely will you buy any other digital devices within the next 3 months?" A higher value indicates higher purchase intent.

Panel (b) shows the same plots for the income sharing decision. Here, the behavioral differences between the treatment and control groups are overall insignificant. This lack

of behavioral difference may be caused by a greater heterogeneity in intrinsic preferences, which makes the utility variation caused by instrumental preference zoom smaller when translated to choice variation. Interestingly, wealthier participants have stronger intrinsic preferences for privacy than their low-income counterparts, which is opposite to the direction that instrumental preferences indicate.

4.5 Dual Privacy Preferences and the Selection in Shared Data

To further examine how the two privacy preferences affect the distribution of data shared, I compare the mean purchase intent and income between the shared (data reflecting Stage two sharing decisions) and the true data (all data collected in Stage one), separately for each treatment group. Table 4.3 displays the t-test statistics for this comparison. With purchase intent, the existence of instrumental incentive makes the shared data feature more high-types than the true data has (see column 2 of Panel (a)); the difference between the shared and true data is marginally significant at the 0.06 level. This selection pattern is consistent with prediction offered by the classical economic model, due to the fact that with purchase intent sharing, intrinsic preferences are largely homogeneous among different types.

Table 4.3: t-Test for Equal Means ($H_1: E[D \mid \text{shared}] - E[D \mid \text{true}] \neq 0$)

(a) Purchase Intent			(b) Income		
	Treatment 1	Treatment 2		Treatment 1	Treatment 2
Statistics	0.190	1.847	Statistics	-0.969	1.053
p-value	0.849	0.065	p-value	0.333	0.293

Note: Treatment 1 = Intrinsic Utility; Treatment 2 = Intrinsic + Instrumental Utility. Shared data are constructed based on consumers' decisions in the second stage as to whether to share their data with the firm; true data refers to all data collected from the first stage.

In comparison, Panel (b) shows that the instrumental preference does not cause a significant selection pattern among the shared data. This is because the intrinsic prefer-

ence for sharing income data is both more heterogeneous and negatively correlated with the instrumental incentive: Wealthier participants have stronger intrinsic preferences for privacy than their low-income counterparts. Taking the messages together, the joint distribution of the two preference components determines the final selection pattern in the shared data. It is therefore crucial for firms to either measure this joint distribution or adopt analysis tools that are agnostic about the privacy preference distribution, rather than imposing “nothing to hide” style assumptions.

CHAPTER 5

THE STRUCTURAL MODEL

The structural model serves three main purposes. First, by backing out the preference primitives, it estimates the dollar value of privacy preferences. The dollar value is an objective scale for utility measurement; furthermore, it facilitates the translation of consumers' privacy preferences to the costs of buying their personal data. Second, it clarifies how instrumental incentives shift privacy choices by changing consumers' beliefs about economic payoffs. While the instrumental incentive is endogenous, the ability of consumers in accounting for the economic consequences of revealing private information is the primitive for a given information environment. Lastly, the utility primitive estimates allow me to simulate privacy choices and evaluate the information value of shared data in counterfactual regimes where the firm's data utilization strategy becomes endogenous.

5.1 Setup

Consumer i is endowed with a vector of personal data $D_i = [d_{i1}, d_{i2}, \dots, d_{iK}]$; d_{i1} is income, and d_{i2} is purchase intent. His sharing decision is characterized by a vector with equal length S_i : Each entry is an indicator of whether the associated personal variable is shared. For example, $S_i = [0, 0, 1]$ means i shares d_{i3} but not d_{i1} or d_{i2} . Sharing decision S_i brings an intrinsic privacy cost, a type-induced payoff from sharing (if the consumer is in the instrumental treatment), baseline compensation, and a random utility shock:

$$U(S_i; C_i, D_i) = \sum_k - \underbrace{c_k(X) \cdot s_{ik}}_{\text{intrinsic cost}} + 1_{instr} \cdot 1_{k \in \{1,2\}} \cdot \underbrace{\beta \cdot p_i \cdot w_k \cdot \widehat{E}[d_{ik}|S_i, D_i]}_{\text{type-induced payoff}} + \underbrace{\beta \cdot p_i \cdot s_{ik}}_{\text{baseline compensation}} + \epsilon_{ik}. \quad (5.1)$$

$C_i = [c_1, c_2, \dots, c_K]$ is the intrinsic preference for privacy; each c_k can be expanded as a function of observables X (more details below). 1_{instr} is the instrumental-treatment indicator. $1_{k \in \{1,2\}}$ selects the data-sharing decisions that are subject to the influence

of instrumental incentives. β is the marginal utility of monetary rewards. p_i is the value of the gift card multiplied by 1%. w_k is the consumer's expected increases in the percentage winning probability for an adjacent, higher type. $\widehat{E}[\cdot]$ is "belief about belief": the consumer's expectation of the firm's expectation about his type. Hereafter, I refer to w_k as the first-order belief and $\widehat{E}[\cdot]$ as the higher-order belief. The first belief simply reflects consumers' understanding of how the payoff depends on the perceived type. On the other hand, the latter depends on a consumer's expectation about not only the firm but also potentially other consumers. The baseline compensation is proportional to the amount of data shared, represented by $p_i \cdot s_{ik}$. Lastly, ϵ_{ik} is the random utility shock associated with choice S ; $\epsilon_{i1}, \epsilon_{i2}, \dots, \epsilon_{iK} \stackrel{iid}{\sim} TIEV$.

Belief about a consumer's type depends on the contents of shared data, as well as the sharing decision itself: $\widehat{E}[d_{ik}|s_{ik} = 1, D_i] = d_{ik}$, $\widehat{E}[d_{ik}|s_{ik} = 0, D_i] = \widetilde{d}_k(p_i)$. I let $\widetilde{d}_k(p_i) = \delta_{k0} + \delta_{k1} \cdot p_i$ to allow for different levels of rationality. If both the firm and consumers are rational, the conjectured type of consumers not sharing their data can change with the stake of instrumental incentives. If instead agents form naive beliefs, δ_1 is zero. By treating consumer beliefs about type-related payoffs ($w_k, \delta_{k0}, \delta_{k1}$) as free parameters, the model imposes no assumption on consumer rationality. This flexibility is useful, given that consumer beliefs may fail to match actual usage of consumer data (Athey et al. 2017), due to either firms' purposeful obfuscation (Ben-Shahar & Chilton 2016) or the evolving data-utilization practices (Stutzman et al. 2013).

The belief parameters reflect the extent to which consumers understand the actual type-dependent payoff, that is, the degree of consumers' rationality. In the experiment, the targeting payoff is held fixed by the experiment; in the counterfactual where the firm starts changing the data utilization strategy, the degree of consumer rationality stays fixed, but the expected instrumental consequence will change, and so will consumers' instrumental preference.

Correctly estimating heterogeneity in intrinsic versus instrumental preferences is key

to understanding how consumers self-select into sharing. I characterize heterogeneity by allowing privacy preference parameters to be functions of observables X , including demographics, time entering the experiment, time spent on each question, browser used, and device specifications. In particular, in models that allow for heterogeneity in intrinsic preferences, $c_k(X) = c_{k0} + c_{kx} \cdot X$. $\delta_{k0}(X)$, $\delta_{k1}(X)$ and $\beta(X)$ are specified similarly, except that variables in δ_k 's do not include income or purchase intent so that the model can be identified. There is also a “built-in” heterogeneity in instrumental preference, coming from the fact that instrumental incentives vary with consumer types.

Psychological factors other than privacy preferences also affect choices. First is the default frame. The literature has proposed different mechanisms underlying the stickiness to default, which implies different ways that the default frame and utility parameters interact with each other (Bernheim et al. 2015, Goswami & Urminsky 2016, Goldin & Reck 2018). To be agnostic about the mechanism, I estimate models separately for each default frame. The estimated parameters represent behavioral preferences under each frame, which are the relevant objects for analyzing firm-side implications of privacy choices. Chapter 6 focuses on the opt-in regime given the current regulatory focus; a comparison between behaviors in the two regimes can be found in Section F.1. The model also includes a behavioral response term $m \cdot (p_i \geq 0) \cdot s_i$ to account for a combination of the mere-incentive effect and potential anchoring effects at the start of the survey. The estimation result and interpretation for this term can be found in Section F.2.

With the specification above, the log-likelihood can be written as the sum of log logit probabilities:

$$LL = \sum_{i=1}^N \sum_{k=1}^K s_{ik} \cdot (\Delta u_{ik}) - \ln(\exp(\Delta u_{ik}) + 1),$$

where Δu_{ik} is the difference in mean utilities between sharing and not sharing data k ,

experienced by consumer i (heterogeneity functions are omitted to simplify exposition):

$$\Delta u_{ik} = - \underbrace{c_k}_{\text{intrinsic cost}} - 1_{i,inst} \cdot 1_{k \in \{1,2\}} \cdot \underbrace{\beta \cdot p_i \cdot w_k \cdot [\delta_{k0} + \delta_{k1} \cdot p_i - d_{ik}]}_{\text{instrumental cost}} + \underbrace{\beta \cdot p_i}_{\text{util from compensation}} + m \cdot (p_i \geq 0). \quad (5.2)$$

5.2 Identification

Coefficients to be estimated include $c_k(X)$, w_k , $\delta_{k0}(X)$, $\delta_{k1}(X)$ for $k \in \{1,2\}$, β , and m . Parameters in $c_k(X)$ are identified as the utility intercept of the participants who enter the intrinsic treatment; since treatment is randomly assigned, these coefficients are the intrinsic preferences shared by all participants. Belief parameters are identified from the instrumental treatment. w_k is identified from how different types react differently to instrumental incentives. $\delta_{k0}(X)$ and $\delta_{k1}(X)$ are identified from responses to the instrumental incentives that are common across types. In particular, the identification of δ_{k1} comes from the interaction between the instrumental treatment and the amount of compensation. Parameter β is identified through the variation in gift-card values. Given that $\beta \cdot p_i$ is linear, and that there are multiple gift-card values across treatments, m is identified from the different responses to zero and non-zero incentives.

The key parameters in this model consist of the following: intrinsic preference, c_k ; first-order belief about the instrumental consequence, w_k ; and the sensitivity to income, β . Identification of these primitives allows me to construct consumers' privacy choices under different counterfactual scenarios. In particular, measuring the first-order belief is important, because it is this belief component that generates the adverse selection pattern created by instrumental incentives. To see this, note that w_k scales the type-dependent payoff when a consumer chooses to share his data $\widehat{E}[d_{ik}|s_{ik} = 1, D_i]$. In comparison, the higher-order belief $\widehat{E}[d_{ik}|s_{ik} = 0]$ does not affect the selection pattern, given that it is not a function of the consumer's private information. Other parameters in the model

are auxiliary: They provide flexibility so that the estimation of key parameters are not affected by confounding factors. For example, δ_{k0} and δ_{k1} may reflect not only consumers' higher-order belief, but also risk preferences that are common across types.

5.3 Estimation

I estimate the model under a Bayesian framework. Flat priors are placed for major parameters, while horseshoe priors are used for the heterogeneity parameters c_{kx} and δ_{kx} (Carvalho et al. 2009, 2010). The horseshoe is a form of continuous shrinkage prior; it accommodates the large number of parameters in the heterogeneity functions and avoids model over-fitting. Compared to other shrinkage priors such as Bayesian Lasso, horseshoe yields estimates that are the closest to results from the Bayes Optimal Classifier. Intercepts c_{k0}, δ_{k0} are left unregularized to obtain unbiased estimates for the mean of functions $c_k(X)$ and $\delta_k(X)$. Note that due to regularization, the estimated heterogeneity will be smaller than the heterogeneity displayed in raw data. This is a necessary trade-off to avoid model over-fitting.

I place non-negativity constraints on the sensitivity to compensation β , and bound constraints on δ such that they do not exceed the actual distribution support of consumer types. No sign constraints are placed on $c_k(X)$: This allows for the possibility that consumers have a “warm glow” in sharing insensitive data for improving research. In addition, the model directly estimates $\widetilde{\delta}_{ik} \equiv \beta \cdot w_k \cdot \delta_{ik}$ instead of δ_{ik} for numerical stability.¹ The distribution of δ_{ik} is then backed out from posterior draws.

1. That is, I estimate $p_i \cdot (\beta \cdot w_k \cdot d_{ik} - \widetilde{\delta}_{k0} - \widetilde{\delta}_{k1} \cdot p_i)$ instead of $\beta \cdot p_i \cdot w_k \cdot (d_{ik} - \delta_{k0} - \delta_{k1} \cdot p_i)$ as the instrumental-preference component.

CHAPTER 6

ESTIMATION RESULTS

6.1 Model Comparison

Table 6.1 compares estimation results from models with different heterogeneity specifications. To compare model performance, I calculate the expected log predictive density (elpd) using the Watanabe-Akaike information criterion (WAIC) approximation; a higher number indicates a better out-of-sample fit (Watanabe 2010). Preference estimates are very different between the model without heterogeneity (Model 1) and the models that allow for heterogeneity in intrinsic preferences (Models 2 to 4). The latter exhibit better fits, as is demonstrated by higher elpd values. On the other hand, allowing for heterogeneity in belief or sensitivity to income does not bring better out-of-sample fit: Estimation results are similar across Models 2, 3 and 4, and the elpd of Model 2 is the highest. Model 2 constitutes the basis for the main analysis.

6.2 Intrinsic Preferences

The willingness to accept (WTA) to give up one's privacy due to intrinsic motives is calculated as $\frac{c_k(X)}{\beta}$.¹ Figure 6.1 shows the predicted distribution of heterogeneous WTA for different data, and Table 6.2 summarizes the statistics corresponding to each distribution (see Table D.1 for credible intervals associated with these estimates). Consumers' WTA are highly heterogeneous. The mean intrinsic preference for sharing different personal variables range from \$0.14 for gender to \$2.37 for information about their children (in the follow-up survey question, many participants describe the request for information about their children as "irrelevant" and "improper"). In comparison, the 97.5% quantiles are more than twice as large as the mean valuations. The upper tails are worth more attention

1. Factors used for scaling p_i are multiplied back to get the correct dollar measure.

Table 6.1: Intrinsic and Instrumental Preference for Privacy: Estimation Results Comparison

Model		1. No Heterogeneity		2. Heterogeneous c		3. Heterogeneous c & δ		4. Heterogeneous c & β	
		mean	95% CI	mean	95% CI	mean	95% CI	mean	95% CI
intrinsic	c_{income}	0.57	[0.43, 0.70]	0.91	[0.59, 1.32]	0.93	[0.58, 1.51]	0.93	[0.60, 1.39]
	c_{intent}	0.55	[0.41, 0.70]	0.83	[0.42, 1.32]	0.84	[0.38, 1.38]	0.87	[0.41, 1.44]
	c_{gender}	0.02	[-0.12, 0.15]	0.19	[-0.16, 0.66]	0.24	[-0.16, 0.95]	0.20	[-0.20, 0.75]
	c_{age}	0.06	[-0.09, 0.20]	0.26	[-0.09, 0.73]	0.29	[-0.16, 0.91]	0.28	[-0.09, 0.82]
	c_{edu}	0.37	[0.23, 0.51]	0.62	[0.33, 1.05]	0.65	[0.29, 1.29]	0.65	[0.29, 1.24]
	$c_{relationship}$	0.20	[0.06, 0.33]	0.50	[0.12, 1.01]	0.55	[0.11, 1.23]	0.50	[0.16, 1.04]
	c_{kid}	0.74	[0.61, 0.88]	1.11	[0.79, 1.46]	1.09	[0.71, 1.51]	1.10	[0.75, 1.55]
	c_{zip}	0.29	[0.16, 0.43]	0.56	[0.23, 1.07]	0.60	[0.18, 1.22]	0.61	[0.19, 1.13]
instrumental	c_{race}	0.29	[0.16, 0.42]	0.60	[0.29, 1.10]	0.65	[0.26, 1.26]	0.65	[0.30, 1.33]
	w_{income}	2.00	[0.15, 3.87]	2.12	[0.11, 3.99]	2.02	[0.14, 3.92]	1.90	[0.04, 3.88]
	w_{intent}	2.63	[1.07, 3.88]	1.94	[0.38, 3.76]	1.97	[0.29, 3.77]	1.90	[0.35, 3.70]
	$\delta_{income,0}$	0.05	[-0.19, 0.29]	0.05	[-0.19, 0.28]	0.05	[-0.19, 0.28]	0.05	[-0.19, 0.29]
	$\delta_{income,1}$	0.05	[-0.19, 0.29]	0.04	[-0.19, 0.28]	0.05	[-0.19, 0.29]	0.04	[-0.19, 0.28]
	$\delta_{intent,0}$	0.08	[-0.35, 0.39]	0.06	[-0.35, 0.38]	0.07	[-0.36, 0.38]	0.06	[-0.34, 0.39]
sensitivity to compensation		-0.05	[-0.36, 0.31]	-0.05	[-0.36, 0.32]	-0.05	[-0.37, 0.31]	-0.04	[-0.34, 0.32]
β		0.13	[0.07, 0.21]	0.15	[0.07, 0.24]	0.15	[0.06, 0.24]	0.15	[0.07, 0.25]
log posterior		-8015	[-8022, -8010]	-7476	[-7540, -7407]	-7433	[-7501, -7352]	-7525	[-7588, -7450]
elpd _{WAIC}		6384	[6358, 6410]	6460	[6431, 6489]	6365	[6337, 6394]	6455	[6427, 6484]

Note: Variables are normalized using the Gelman method before estimation. Wherever heterogeneity is allowed, the table displays estimates on the intercept term only. The same seed is used for estimating different models.

from a data-acquisition perspective, since these are the prices that firms need to surpass to guarantee a representative set of data. For example, a data collector needs to pay \$3.82 per customer for 97.5% of them to share their income data, and \$5.08 per customer to get 97.5% of purchase-intent data. As a robustness check, Appendix E includes WTA estimates from Model 4, which allows for heterogeneity in both intrinsic utility and sensitivity to income. The WTA distribution is quantitatively and qualitatively similar to the main result produced by Model 2.

Are these privacy-preferences high or low? To interpret their magnitude, I compare these numbers with findings in the previous literature. In Hui (2007), consumers' willingness to answer one additional sensitive question in a market research survey amounts to 2.78 Singapore dollars (2.04 USD).² Acquisti et al. (2013) estimate consumers' WTA of

2. The original paper does not calculate WTA. However, according to their estimation, one additional Singapore dollar is associated with 0.39 util while answering one more sensitive question decreases utility

Figure 6.1: Posterior Predicted Density of WTA in Intrinsic Preference

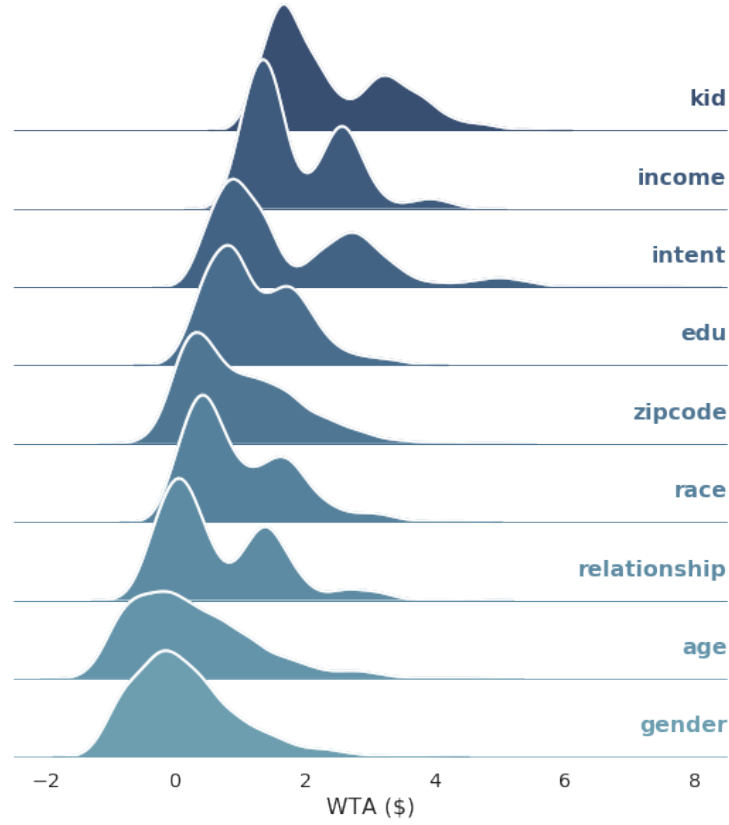


Table 6.2: Posterior Predicted Distribution of WTA in Intrinsic Preference

	mean	median	2.5%	97.5%
kid	2.367	2.069	1.220	4.311
income	1.870	1.546	0.944	3.823
intent	1.825	1.352	0.398	5.078
edu	1.228	1.051	0.228	2.845
zipcode	0.985	0.800	-0.157	2.916
race	0.980	0.737	-0.066	2.945
relationship	0.687	0.390	-0.448	2.894
age	0.260	0.084	-1.064	2.718
gender	0.142	0.006	-1.043	2.187

Note: Numbers in this table refer to statistics associated with the estimated WTA distribution among consumers; these are measures of preference heterogeneity.

attaching their names to transactions associated with a gift card to be \$1.04. Considering by 0.14. Thus, $0.39/0.14$ is the WTA for the marginal question that the consumer chooses to answer.

the variation caused by different economic contexts and categories of data requested, the magnitudes of WTA in these studies are similar to the mean intrinsic preference in my data.

Another way to gauge the magnitude of intrinsic preferences is calculating the WTA for a profile, which is essentially a bundle of different data. For example, if cookies used to identify online users are associated with different demographic tags examined above, the WTA for sharing the whole demographic profile will have a mean of \$10.34 and a 97.5% quantile of \$29.72. For categories of personal data that are highly granular and intimate, such as browsing and location histories, the WTAs are possibly even higher.

A third way of sensing the magnitude of privacy preferences is by comparing them with the firm's willingness to pay for these data. This comparison is further discussed in Section 8.2, where I calculate the firm's valuation of personal data under different data acquisition strategies.

6.3 Beliefs that Generate Instrumental Preferences

In the model, consumers' beliefs are utility primitives associated with the endogenous instrumental preference: They represent how consumers recognize the economic impact of sharing private information under various circumstances. One particular interest is the extent to which consumers' beliefs correspond to the actual data usage. In the experiment, a consumer whose type is one tier above receives an additional 2 percent probability of winning the gift card if his type is disclosed to the firm. This means that consumers' first-order beliefs are accurate if w equals 2. Column 2 of Table 6.1 shows that consumers' beliefs about w_{income} and w_{intent} are correct on average. As is mentioned in Section 2.4, the impact of instrumental incentives on selection in shared data will persist as long as these first-orders are accurate. Consumers' beliefs about the payoff of being anonymous are much noisier, as is reflected by the wide credible intervals for δ . This pattern is consistent with the fact that guessing the payoff from withholding data requires higher-

level thinking (need to form beliefs about the firm's as well as other consumers' reasonings) and information about other consumers' privacy preference distribution.

Overall, the belief estimates represent the level of consumer sophistication in making privacy choices when fully informed, as is required by GDPR and other similar regulations. My estimates suggest that with a transparent information environment, consumers are able to engage in strategic reasoning when making data sharing decisions, and their beliefs are accurate to the first order. In other policy regimes where firms are allowed to obfuscate information about how data will be used and accessed, consumers' beliefs are likely to be further away from actual practices.

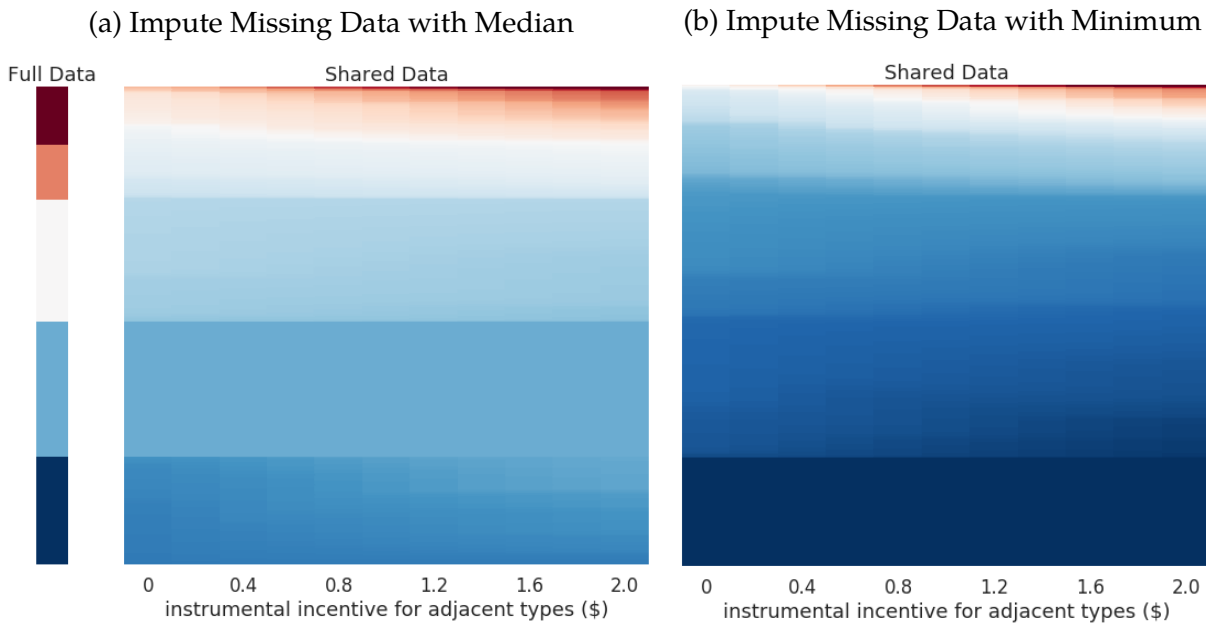
6.4 Dual Privacy Preferences and Selection in Shared Data

Given the degree of consumers' sophistication, the magnitude of instrumental preference is determined by the actual differential payoffs when their private information is revealed versus withheld. To examine how changes in the relative magnitudes of intrinsic and instrumental preferences change the selection pattern in data available to the firm, I vary the magnitude of the actual instrumental payoff, then simulate consumers' privacy choices and the data shared, taking intrinsic preference and belief parameters as fixed. This is implemented in the example of income sharing. To compute the firm's view about consumers, I assume two different imputation strategies, each consistent with a different type of firm's inference. The first strategy imputes the missing data using the median of observed data. This imputation strategy is consistent with a view that consumers care about privacy only intrinsically, and that people who share and who withhold their data have similar characteristics. The second strategy imputes missing data using the minimum of observed data, consistent with the view that privacy concerns are purely instrumental.

Figure 6.2 compares the distributions of full and firm data across the range of instrumental preferences and under different imputation strategies. Panel (a) shows that as instrumental incentive increases, the composition of consumers sharing their data tilts

more and more towards high-income cohorts, as is indicated by the expansion of red and shrinkage of blue regions from left to right. However, even when the mean instrumental incentive matches the mean intrinsic preference for income (2 dollars; the rightmost bar in the heatmap), the firm still ends up overestimating the proportion of low-income customers. When only motivated by intrinsic preference, it is the low-income consumers who are more willing to share their personal data. The rightmost vector of Panel (a) shows that the preference heterogeneity generated by intrinsic preference is not fully offset by instrumental incentive even when their magnitude matches on average. As a result, the median consumer type used for imputing missing values is still lower than average. In situations like these, adopting a “nothing to hide” argument will only end up exacerbating the bias in the firm’s view about consumers, as is shown in Panel (b).

Figure 6.2: Full vs. Firm Data across Different Magnitudes of Instrumental Incentive



In sum, taking a monolithic view about the nature of consumers’ privacy preferences will result in misleading inferences about consumers and managerial decisions. Instead, firms need to either learn about the joint distribution of privacy preferences of their consumers, preferably via experimentation; or adopt data collection and analysis strategies

that are agnostic about the joint distribution.

CHAPTER 7

REPLICATING THE EXPERIMENT IN THE FIELD

Privacy preferences are context-specific: This fact calls for repeated measurements of privacy preferences across scenarios. Only by doing so can firms and researchers better understand the nature of personal data shared by consumers and make valid inferences using these data.¹ In particular, measuring the joint distribution of intrinsic and instrumental preferences will allow firms and researchers to understand consumers' data sharing decisions, even when their data utilization strategy (and thus the instrumental preference) becomes endogenous. Below, I describe how firms and researchers can replicate my experiment to measure consumers' privacy preferences in the field.

To measure the selection in shared data, having a "ground truth" dataset is necessary. The ground truth data can be obtained by having a treatment group where consumers are given compensation high enough so that everyone chooses to share. Alternatively, distribution-level statistics about the relevant type (price sensitivity, risk type, etc.) may be available from a third-party intermediary or a government agency (e.g. the census bureau). The other treatments can then be designed to induce exogenous variations of instrumental preferences, for example, by changing the targeting rule and informing consumers about the change.

One challenge of implementing the experiment in the field is that instrumental preference is hard to be completely removed. Suppose Safeway asks its customers for data and promises not to use these data for business purposes. Without additional technological or legal guarantees, such promise will not have commitment power: Users may still expect Safeway to use these data to design customized coupons and promotions. More generally, consumers' belief about the consequences of revealing their personal information can anchor on the firm's routine practices of using the data.

Fortunately, we do not need a treatment that removes instrumental preferences com-

1. Advertising effect is also context-specific, and we haven't given up measuring it.

pletely. To separate the two preference components, it suffices to have variation in actual instrumental payoffs known to the customers. For example, if data are eventually used for designing customized coupons, the depths of coupons can be different across treatment arms. Consumers' privacy choices across treatments then allow us to back out the magnitude of changes in instrumental preferences and compare it with changes in actual instrumental payoffs. By comparing the two, we will be able to estimate consumers' degree of rationality, that is, how much they internalize the actual instrumental consequences when forming privacy preference. Assuming this degree of rationality is stable across treatments, we can then use it to calculate privacy preferences and data sharing choices in a hypothetical scenario where the instrumental preference is zero, thus backing out the intrinsic preference among consumers.

CHAPTER 8

COUNTERFACTUALS

In this chapter, I examine data acquisition and analysis strategies when the firm does not know about the distribution of consumers' privacy preferences. The following analysis focuses on two angles:

Ex ante: How should the firm allocate its resources for data collection?

Ex post: Can the firm learn about selection pattern in shared data from the data itself, instead of relying on faulty assumptions?

8.1 Setup

I investigate the questions above in the context of price targeting. Pricing is an area that has witnessed substantial efficiency improvement due to the recent influx of personal data (e.g. Dubé & Misra 2019). The focal firm is taken as the third-party company featured in the experiment's second stage. I take a choice scenario featured in the first-stage conjoint survey to serve as the market environment (Task 3) and the product that the firm sells (Option C); they are displayed in Figure 8.1. Consumers' valuation of product features and price sensitivity are calculated from their responses to the conjoint survey. The marginal cost of a smartwatch is assumed to be \$50.¹ All data sharing choices and their impact on firms are evaluated in a GDPR-like policy regime, where firms need to seek opt-in consent before collecting data.

The value of different data utilization strategies is reflected as the difference in profits with or without adopting this strategy. For data acquisition, this is represented by the value of different shared datasets associated with different levels of compensation; for data analysis, this is reflected by the value of different models, taking the shared data as

1. This amount is the average of the estimated production cost for Apple Watch (\$83.70) and the cost of Fitbit Flex (\$17.36). See <https://www.forbes.com/sites/aarontilley/2015/04/30/the-apple-watch-only-costs-83-70-to-make/#6e981e8d2f08>, and <https://electronics360.globalspec.com/article/3128/teardown-fitbit-flex>.

Figure 8.1: Screenshot of the Conjoint Task and Focal Product Used for Price Optimization

If you want to buy a smartwatch and these are the available options, which one will you choose? (Scenario 3/7)

Product	A	B	C	D
Fitness Tracking	Activity	Activity + heart rate	Activity + heart rate	Activity + heart rate
Voice Control	Yes	Yes	Yes	No
Mobile Payment	No	No	Yes	Yes
Encryption	TLS 1.0	SSL 3.0	SSL 3.0	TLS 1.0
Login Option	Pin, pattern	Pin, pattern	Pin, pattern	Pin, pattern, face
Price	\$299	\$149	\$199	\$249

☐ Product A
 ☐ Product B
 ☐ Product C
 ☐ Product D
 ☐ None of the above

Note: Highlights are added to illustrate the focal product used for the counterfactual. They were not present in the actual experiment.

fixed. To calculate the true profits, I estimate consumer demand based on the full data (from Stage one) and view this demand as the ground truth.

To construct firm data for each counterfactual compensation level, I first simulate 300 privacy choice draws, and then construct a shared dataset for each draw: If a consumer decides not to share data k , the value of variable k is left empty. Firm data also contain a “not sharing k ” indicator, which equals 1 when the consumer chooses not to share k , and 0 otherwise. I assume the firm imputes missing variables using mean values among the data available, and takes competitors’ prices as given when doing price optimization.²

8.2 When and How to Buy Data from Consumers?

The value of personal data to firms. Calculating the value of data is the logical prerequisite for assessing the value of data buying plans. Moreover, it is also a legal prerequisite for buying data from consumers in the most recent CCPA regulation. In particular, Section

² If a consumer chooses not to share the choice task responses, the outcome variable in the pricing model is missing. In this case, I assume the firm imputes the missing outcome using observed conjoint choices from consumers who are demographically similar to this observation. In practice, this assumption amounts to leaving the whole observation out of the pricing model and predicting the missing outcome afterwards.

999.336 of the *proposed CCPA Regulation* states,³

If a business is unable to calculate a good-faith estimate of the value of the consumer's data or cannot show that the financial incentive or price or service difference is reasonably related to the value of the consumer's data, that business shall not offer the financial incentive or price or service difference.

Nevertheless, the value of consumer data is hard to pin down, as it depends on what other data are already available, and is model- and question domain-specific.

Below, I provide one way to calculate the value of (additional) consumer data for a given model, application domain, and dataset already available to the firm. I then decompose this value by the role data plays. This decomposition will point to strategies that allow firms to design more efficient data buying plans. Each data buying strategy leads to a different dataset shared to the firm, denoted as d ; the full data is indicated as d_0 . The profit loss of not getting full data is

$$\Delta\pi_{total} = \pi\left(P_{d_0}(d_0)\right) - \pi\left(P_d(d)\right).$$

Note that in a new policy regime, a firm normally only observes the voluntarily shared data d but not the full data d_0 . It may, however, learn about the value of full data via a third-party data intermediary who has access to full data.

This valuation can be further decomposed into two parts. One part indicates the value of data in improving the pricing model; the other indicates its value in profiling consumers, taking the model as fixed:

$$\Delta\pi_{total} = \Delta\pi_{model} + \Delta\pi_{profile};$$
$$\Delta\pi_{model} = \pi\left(P_{d_0}(d)\right) - \pi\left(P_d(d)\right); \Delta\pi_{profile} = \pi\left(P_{d_0}(d_0)\right) - \pi\left(P_{d_0}(d)\right). \quad (8.1)$$

3. <https://www.oag.ca.gov/sites/all/files/agweb/pdfs/privacy/ccpa-text-of-mod-clean-020720.pdf>

Here, $P_{d_0}(d)$ is the firm's pricing model trained using d_0 and taking d as input. π is the true profit, which is a function of the pricing strategy.⁴ The reason for decomposing the information value is two-fold. First, information externality exists among consumer data through the model-estimation phase, but not the profiling phase. This information externality points to strategies that a firm can use to economize on data acquisition. Second, consumers' instrumental preferences are activated by profiling rather than modeling, since it is the former that assigns individualized prices based on consumers' private information. This aspect is useful for incentive design when instrumental incentives become the main obstacle for getting representative data. Both aspects are discussed extensively below.

As a starting point, I calculate the value of data when consumers only have intrinsic preferences. This may occur when consumers who receive requests for data provision do not directly experience the economic impacts from the firm's data analysis. For example, Nielson and ComScore maintain a panel of consumers and provide the data to other firms for analysis, but these firms' focal customers may not overlap with the panel (although they have similar demographics). Alternatively, a wedding vendor has one-off transactions with most of its customers, and those who already use its service will not expect direct economic consequences of sharing their data. The first two rows of Table 8.1 shows the posterior mean and credible intervals of the profit losses at different levels of compensation. Having to seek consent results in a profit loss of \$1,440 per thousand customers when no compensation is given; this amount is 3% of the total profits that could have been obtained using full data. Inaccuracy in profiling accounts for 38.6% of the total profit loss.

Information externality and firm's WTP for data. Suppose the firm obtains dataset d when consumers make informed data sharing choices under no compensation. The firm's WTP for obtaining d_0 given d is calculated as the profit difference divided by the unit

4. It is more natural to measure the information value on model building as $\pi(P_{d_0}(d)) - \pi(P_d(d))$ instead of $\pi(P_{d_0}(d_0)) - \pi(P_d(d_0))$, given that data at the profiling phase are often harder to obtain than are the data for modeling.

Table 8.1: Profit Loss (\$/1,000 customers) with Different Prices for Data

Privacy Preference	Role of Data	Price per Variable (\$)					
		0		1		2	
Intrinsic	Model+Profile	1,440	[657, 3,040]	1,126	[536, 2,400]	883	[382, 2,193]
	Profile	556	[379, 716]	465	[317, 636]	380	[218, 592]
Intrinsic+Instrumental	Profile	862	[799, 892]	857	[779, 892]	852	[748, 890]

Note: This table reports posterior mean estimates, with 95% credible intervals in brackets. Total profit loss is calculated as $\Delta\pi_{total} = \pi(P_{d_0}(d_0)) - \pi_d(d)$; profit loss associated with profiling is $\Delta\pi_{profile} = \pi(P_{d_0}(d_0)) - \pi(P_{d_0}(d))$.

difference between the two datasets:

$$WTP_{firm} = \frac{\Delta\pi}{N \cdot \bar{K}}.$$

Here, $\Delta\pi$ can be either $\Delta\pi_{total}$ or one of its subparts, depending on the particular data acquisition strategy. N is the number of consumers whom the firm wants to collect data from; with a mass-collection strategy, N equals to the market size. \bar{K} is the average number of variables withheld per consumer when they receive no compensation for sharing data. This is the break-even price that the firm is willing to give to each consumer and each personal variable. If the firm's WTP is lower than the consumers' WTA under a particular data acquisition strategy, then this strategy is not worth adopting, because matching the price for data to consumers' WTA will lead to a loss in profits. Consider the following two strategies:

(a) The firm buys data from all consumers, and uses the data for both modeling and profiling.

(b) The firm allocates resources to collect data that improves the model. In doing so, it randomly samples 1% of consumers, and only compensates them for sharing data.

When data is used for building the pricing model, there is information externality among consumers. Since data is used to learn a systematic relationship between optimal prices and personal characteristics, data coming from one consumer also improves

the inferred optimal prices for other consumers. On the other hand, such information externality does not exist when data is used for profiling: Knowing the characteristics of consumer A does not tell the firm about the characteristics of other consumers. Therefore, with strategy (a), the firm has to ask for data from every consumer; but with strategy (b), it can randomly sample, say 1% of total customers, and only buying data from the sample.

The calculation below shows the improvement from leveraging this information externality at the modeling stage. On average, a consumer withholds 5.31 variables without compensation. With strategy (a), $WTP_{firm} = \$1.440/5.31 = \0.27 . In comparison, consumers' mean WTA ranges from \$0.14 to \$2.37. This result indicates that collecting data from all consumers for both purposes is not a viable strategy. With strategy (b), however, $WTP_{firm} = (\$1.440 - \$0.556) \times 100/5.31 = \16.65 , four times as much as the 97.5% quantile of consumer WTA even for the most precious data.

The WTP calculated above represents the average value of data. Ideally, we want to know the marginal value of data. However, calculating this marginal value at the observational level is computationally demanding, as one will need to calculate the profit difference caused by deleting every single observation and then take the average. The computational time is proportional to the amount of time needed to refit the model for each new data times the number of observations in the dataset. Assuming the marginal value of additional data decreases with the volume of data already obtained, the average value calculated as above will serve as a lower bound for the marginal value evaluated at d .

Although the calculation above may be overly simplistic, the qualitative pattern is general. Recent work highlight the presence of information externality in privacy choices (Acemoglu et al. 2019, Bergemann et al. 2019, Choi, Jeon & Kim 2019). My paper further shows that to leverage information externality when collecting consumer data, firms need to know the stage at which it is present. In reality, the performance of the model will increase with the size of the estimation sample. For sampling to improve the efficiency of

data collection, a critical condition is that estimation data has decreasing returns to scale. This condition is empirically supported by recent literature, such as Bajari et al. (2019) and Claussen et al. (2019). The marginal return to data will diminish slower with more complex models and greater heterogeneity that the model intends to capture; in these cases, the sampling percentage should increase accordingly.

Data acquisition when instrumental preferences are present. In cases where a firm solicits data from its own consumers and applies its model to them, instrumental preferences will be present. In the context of price targeting, consumers' instrumental motives are derived from price differences that they expect to receive when sharing versus withholding their data. This is different from the instrumental incentive in the experiment. Below, I calculate the instrumental preferences based on the consumer belief estimates (which indicates their degree of sophisticated reasoning), and evaluate data acquisition strategies accordingly.

Previous estimation result shows that consumers have first-order rationality when making data sharing decisions. However, there is no sufficient evidence that they also conduct higher-level reasoning. The requirement to sustain higher-order rationality is high and unlikely to be satisfied: Consumers need to think that the firm is rational and knows the distribution of other consumers' privacy preferences. Therefore, in what follows, I assume that consumers have beliefs that are "approximately rational". To further simplify the analysis, I focus on the case where the firm has previously trained its pricing model using a set of representative data from other customers. That is, I only calculate the value of shared data for profiling. Taking the pricing model as given, consumer i expects to receive different prices when he withholds or shares data k :

$$E[P_i | s_{ik} = 0] = \bar{P}_{i'}; \text{ and } E[P_i | s_{ik} = 1, d_{ik}] = \bar{P}_{i', \forall d_{i'k} = d_{ik}}.$$

Here, i' denotes all other consumers in the market, $\bar{P}_{i'}$ is the mean price for all other

consumers, and $\bar{P}_{i', \forall d_{i'k}=d_{ik}}$ is the average price for all other consumers with the same attribute d_{ik} . Given that consumer i can always choose the outside option when the price is too high, his instrumental preference is the difference in log sums:

$$E[\Delta U] = \frac{1}{\beta_i} \left[\log(1 + \exp(v_i - \beta_i \bar{P}_{i'})) - \log(1 + \exp(v_i - \beta_i \bar{P}_{i', \forall d_{i'k}=d_{ik}})) \right], \quad (8.2)$$

where β_i is i 's price sensitivity and v_i is his valuation for the product.

The last row of Table 8.1 shows that when consumers harbor instrumental preference, the loss from not obtaining the full data is larger—in this case, around twice as large compared with when they only have intrinsic privacy concerns. This is driven by a more severe sample bias in the shared data. It also shows that compensation for data sharing is less effective in overcoming instrumental incentives. This is because the expected difference in log sum due to revealing private information ranges from \$20 to \$50 for each data-sharing decision.

To analyze the scenario where consumers have instrumental preferences both when sharing modeling data and profiling data, one needs to solve the full equilibrium, because the pricing model and the data shared now depend on each other. This is a computationally daunting task, as each iteration involves simulating R different dataset draws and computing the firm's pricing model for each of these datasets. This analysis is part of my future work. Based on the results from the intrinsic-only case in Table 6, I conjecture that when consumers have both intrinsic and instrumental preferences, the economic loss of having incomplete data for modeling is much larger than its impact solely on profiling. In this case, the firms may achieve efficiency gain by using a separate consumer panel for modeling. This way, the firm can “insulate” the modeling sample from instrumental concerns.

8.3 Learning about Selection: What Do Privacy Choices Reveal?

At the stage when available consumer data is a given, what can firms do to improve the quality of their inference? Incorporating consumers' privacy choices into the model may help. **Privacy choice** means the decision of whether to share personal data per se, apart from the contents of data. An example of the privacy choice variable is the do-not-track header: It is generated when a user declines to be tracked by third parties, and remains visible to websites.⁵ In a world where privacy preferences were purely instrumental, privacy choices would completely reveal a consumer's hidden type. If this were true, firms would be able to use privacy choices as additional targeting variables and substantially improve the targeting performance even without observing the actual contents of personal data for these consumers. However, the information value of privacy choices changes substantially when consumers have both types of privacy preferences.

In what follows, I examine what firms can and cannot learn by adding privacy choices to their model. This is again evaluated in the context of price targeting. Fixing the dataset shared by consumers, I compare the performances of two pricing models. In the without-indicator model, the firm sets prices based on the content of data provided by consumers, but not their privacy choices. In the with-indicator model, the firm sets prices based on both the content of available data and consumers' privacy choices. Note that the focus here is not evaluating the performance of data acquisition strategies. Therefore, I take d as actual data shared from the experiment's second stage, instead of simulating counterfactual datasets, which would add unnecessary sampling error. The metrics for evaluating pricing performances are

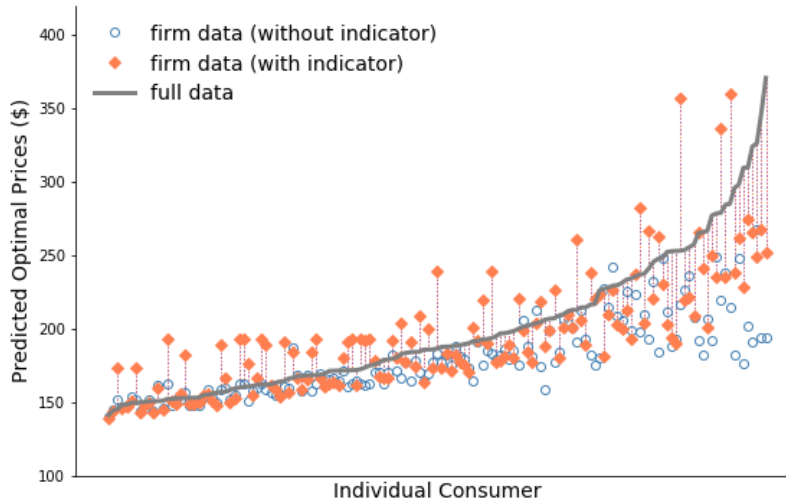
$$\Delta\pi_{without-indicator} = \pi(P_{d_0}(d_0)) - \pi(P_d(d)), \Delta\pi_{with-indicator} = \pi(P_{d_0}(d_0)) - \pi(P_{d+c}(d+h)),$$

where h refers to the privacy-choice indicator.

5. <https://www.eff.org/issues/do-not-track>.

Figure 8.2 compares the individual-level optimal prices predicted by these two models when applied to firm data, and the benchmark prices calculated based on full data. Prices predicted by the with-indicator model are less biased. The mean price that consumers receive under this model is \$194.26, close to the mean price \$199.05 when the firm has all data; in comparison, the mean prices under the without-indicator model is \$179.22. In other words, a model with privacy choice indicators can serve as a bias correction tool: By comparing average prices with and without the privacy choice indicator, the firm can learn about the direction and magnitude of selection in the shared personal data. This information is useful not only for interpreting insights obtained from the data, but also for designing and evaluating data acquisition schemes.

Figure 8.2: Inferred Optimal Prices with or without Privacy Choice Indicator



On the other hand, Figure 8.2 also shows that predictions generated by the with-indicator model are not necessarily more accurate. The with-indicator model surpasses the alternative model when predicting prices for consumers who have high valuations for the product, but performs worse at the opposite end of the spectrum. To quantify the impact of adding privacy choices on the performance of price targeting, Table 8.2 compares the profit losses from using each pricing model. Adding privacy choices improves overall targeting performance, but only to a small extent (see column 1). The performance gain

mainly comes from a better calibration of prices offered to privacy-sensitive consumers (column 3). On the other hand, the prediction accuracy for consumers who already share lots of data can suffer (column 2), because privacy choices add little additional explanatory power for the price sensitivity of these consumers.

Table 8.2: Profit Loss When Using Firm Data (\$/1000 Consumers)

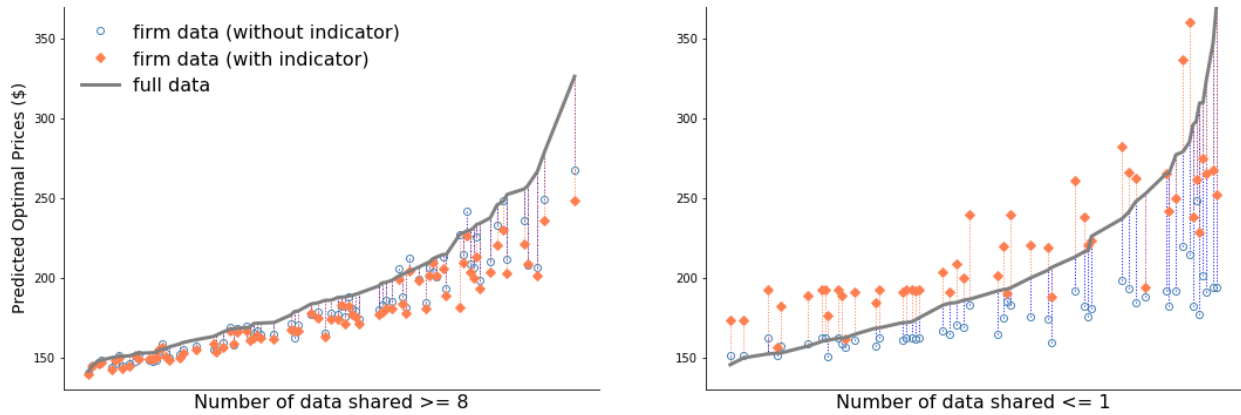
Model	Consumer Subset					
	All consumers		Share all data		Share no data	
without-indicator	2,441	[917, 5,113]	2,348	[926, 5,721]	2,492	[1,229, 4,186]
with-indicator	2,384	[957, 5,229]	2,405	[877, 5,193]	2,419	[1,138, 3,854]

Note: This table reports posterior mean estimates, with 95% credible intervals in brackets. Profit loss is calculated as $\Delta\pi = \pi(P_{d_0}(d_0)) - \pi(P_d(d))$; a lower number indicates a better performance.

To further understand how the with-indicator model affects consumers who make different privacy choices, Figure 8.3 separately displays prices for consumers who withhold at most one personal variable and prices for consumers who withhold most of their data. For consumers who already share a lot, prices set by the two models are not very different. For consumers who choose not to share most personal data, the actual optimal prices are higher than those for privacy-insensitive consumers on average, as can be seen by comparing the gray lines across the two panels. The privacy-choice indicators pick up this information, resulting in a rise in inferred optimal prices. Privacy choices reflect the average characteristics of non-sharing consumers, but not individual differences within this group. This is reflected by the fact that the new inferred prices are shifted up by almost the same amounts when compared with the original prices inferred.

Taken together, these results paint a nuanced picture of the information value of privacy choices. Incorporating consumers' privacy choices into a firm's decision model can reveal the direction and degree of sample selection, provided that the untruncated distribution of outcome (e.g. individual-level sales across all customers) is observed. On the other hand, the information value is limited when the goal is improving individual-

Figure 8.3: Inferred Optimal Prices with or without Privacy Choices (by Consumer Subsets)



level pricing. Intuitively, privacy choices capture systematic differences in price sensitivity between consumers who share and who withhold their data, but will not reflect the heterogeneity in sensitivity within the withholding consumers. The former reflects the impact of selection, while the latter is more useful for targeting. With greater heterogeneity in intrinsic preferences, consumers who decline sharing data are more likely to exhibit heterogeneous price sensitivity, and privacy choices become less useful in improving targeting as a result.

Several existing and proposed privacy regulations, such as CCPA, COPRA and the Washington Privacy Act, all include a non-discrimination requirement. This requirement prohibits firms from giving consumers differential prices based on their privacy choices, while still allow prices to vary with the contents of voluntarily shared data. In other words, under such a clause, firms are not allowed to base price targeting on privacy choices indicators. The analysis above suggests that such prohibition will not strongly affect targeting effectiveness or firm profits. On the other hand, Figure 8.3 suggests that allowing targeting to be a function of privacy choices has a substantive re-distributive effect on consumer welfare. That being said, the actual effect of these policies on firms and consumers largely depends on the joint distribution of privacy preferences and consumer price sensitivity. To evaluate the impact of these policies, researchers need to measure this

joint distribution among the consumers and markets that will be affected.

8.4 Summary

The counterfactual studies suggest that firms can improve their inference on consumers by sampling when collecting personal data used for estimation, and by using privacy choices to learn about self-selection bias when estimating models. Both strategies avoid the use of arbitrary assumptions on consumers' privacy motives when making inferences. The quantitative results are subject to the influences of assumptions on firm behavior and the application context. However, the strategies developed based on the qualitative findings are generic to inference problems. They can be applied not only when personal data are collected for pricing, but also when they are used for other managerial decisions such as targeted advertising and customized product recommendation. They can also be applied when personal data are requested to conduct general-interest research with a goal of inference.⁶

6. <https://socialscience.one/blog/first-grants-announced-independent-research-social-media%E2%80%99s-impact-democracy>.

CHAPTER 9

CONCLUSION

Privacy choices are motivated both by intrinsic preference—a taste for privacy, and by instrumental preference—the utility change from disclosing one’s type relevant to the specific market environment. While the intrinsic preference is a utility primitive, the instrumental preference is endogenous to how the firm uses consumer data to generate targeting outcomes. Separating these two preference components can help us understand how consumers self-select into sharing data, and how this selection pattern reacts to changes in the firm’s data utilization. Ultimately, understanding the selection in the voluntarily-shared consumer data is crucial for obtaining valid insights from these data and for designing effective data compensation strategies.

By separating intrinsic and instrumental motives using experimental variation, I establish the following findings. Consumers’ WTA corresponding to intrinsic preferences is highly heterogeneous and skewed to the right: The mean valuation for sharing a demographic profile is \$10, while the 97.5% quantile is \$30. When given GDPR-style transparent and plain information, consumers form their belief about the instrumental consequences in a manner that is first-order correct. The direction and magnitude of selection in shared consumer data are jointly determined by the heterogeneity and correlation of these two preference components. Firms and researchers can adopt the following strategies to improve their inferences based on shared consumer data. First, they can run an experiment to measure the joint privacy preference distribution among consumers, which can provide more information about selection in shared data. Alternatively, they can adopt strategies that are agnostic about the preference distribution. Ex ante, they can allocate resources to buying a more representative dataset rather than simply increasing its volume where the information gain is largest. Ex post, incorporating privacy choices into models can reveal the impact of consumers’ self-selection on inference.

Although this paper does not directly discuss welfare, measuring intrinsic and instru-

mental preferences separately is useful for welfare calculations. First, separating these two components can help us understand the extent to which privacy preferences change endogenously with firms' strategy to use shared consumer data. In addition, the relative magnitudes of these two preference components have distinct welfare implications. The intrinsic preference implies a pure loss of consumer welfare caused by data collection; the instrumental preference implies welfare transfer between consumers and firms and among consumers (subject to the influence of information externality).

Privacy preferences are known to be context-specific. Part of the context dependences comes from changes in the perceived instrumental consequences across scenarios, which my model formally characterizes. In addition, privacy choices can also be influenced by various psychological shifters. I show how the experiment can be replicated in the field to address the second type of context-dependence. Different versions of the experiment can be useful for unpacking how consumers' intrinsic preference respond to psychological shifters, and for examining consumers' beliefs about the instrumental consequences when firms obfuscate their data usage. The conceptual framework and the estimation can be extended to discuss cases where consumers can manipulate the contents of shared information at a cost, and when sharing data improves the quality of products offered by the firm.

Future analysis will enrich the model and further explore the implications of the dual preference framework. One direction is to investigate the substitution and complementarity in data-sharing decisions. Another direction is to develop better models to extract information from consumers' data-sharing decisions. A third extension is to explore optimal compensation for data procurement that incorporates instrumental incentives.

REFERENCES

- Acemoglu, D., Makhdoumi, A., Malekian, A. & Ozdaglar, A. (2019), 'Too much data: Prices and inefficiencies in data markets', *NBER Working Paper No. 26296* .
- Acquisti, A., John, L. K. & Loewenstein, G. (2012), 'The impact of relative standards on the propensity to disclose', *Journal of Marketing Research* **49**(2), 160–174.
- Acquisti, A., John, L. K. & Loewenstein, G. (2013), 'What is privacy worth?', *Journal of Legal Studies* **42**(2), 249–274.
- Adjerid, I., Acquisti, A. & Loewenstein, G. (2019), 'Choice architecture, framing, and cascaded privacy choices', *Management Science* **65**(5), 2267–2290.
- 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 Science* **62**(4), 1042–1063.
- Allcott, H., Braghieri, L., Eichmeyer, S. & Gentzkow, M. (2020), 'The welfare effects of social media', *American Economic Review* **110**(3), 629–676.
- Aridor, G., Che, Y.-K., Nelson, W. & Salz, T. (2020), 'The economic consequences of data privacy regulation: Empirical evidence from gdpr', *Available at SSRN 3522845* .
- Athey, S., Catalini, C. & Tucker, C. (2017), 'The digital privacy paradox: Small money, small costs, small talk', *NBER Working Paper No. 23488* .
- Bajari, P., Chernozhukov, V., Hortaçsu, A. & Suzuki, J. (2019), The impact of big data on firm performance: An empirical investigation, in 'AEA Papers and Proceedings', Vol. 109, pp. 33–37.
- Becker, G. S. (1980), 'Privacy and malfeasance: A comment', *Journal of Legal Studies* **9**(4), 823–826.
- Ben-Shahar, O. & Chilton, A. (2016), 'Simplification of privacy disclosures: An experimental test', *The Journal of Legal Studies* **45**(S2), S41–S67.
- Benndorf, V. & Normann, H.-T. (2018), 'The willingness to sell personal data', *The Scandinavian Journal of Economics* **120**(4), 1260–1278.
- Bergemann, D., Bonatti, A. & Gan, T. (2019), 'The economics of social data', *Cowles Foundation Discussion Paper* .
- Bernheim, B. D., Fradkin, A. & Popov, I. (2015), 'The welfare economics of default options in 401 (k) plans', *American Economic Review* **105**(9), 2798–2837.
- Bernheim, B. D. & Rangel, A. (2010), Behavioural public economics, in 'Behavioural and Experimental Economics', Springer, pp. 51–58.

- Boas, T. C., Christenson, D. P. & Glick, D. M. (2018), 'Recruiting large online samples in the united states and india: Facebook, mechanical turk, and qualtrics', *Political Science Research and Methods* pp. 1–19.
- Brynjolfsson, E., Collis, A. & Eggers, F. (2019), 'Using massive online choice experiments to measure changes in well-being', *Proceedings of the National Academy of Sciences* **116**(15), 7250–7255.
- Burtch, G., Ghose, A. & Wattal, S. (2015), 'The hidden cost of accommodating crowdfunder privacy preferences: A randomized field experiment', *Management Science* **61**(5), 949–962.
- Carvalho, C. M., Polson, N. G. & Scott, J. G. (2009), Handling sparsity via the horseshoe, in 'Artificial Intelligence and Statistics', pp. 73–80.
- Carvalho, C. M., Polson, N. G. & Scott, J. G. (2010), 'The horseshoe estimator for sparse signals', *Biometrika* **97**(2), 465–480.
- Choi, J. P., Jeon, D.-S. & Kim, B.-C. (2019), 'Privacy and personal data collection with information externalities', *Journal of Public Economics* **173**, 113–124.
- Choi, W. J., Jerath, K. & Sarvary, M. (2019), 'Customer purchase journey, privacy choices, and advertising strategies', *Working Paper* .
- Claussen, J., Peukert, C. & Sen, A. (2019), 'The editor vs. the algorithm: Economic returns to data and externalities in online news', *Available at SSRN* 3479854 .
- DellaVigna, S. (2009), 'Psychology and economics: Evidence from the field', *Journal of Economic Literature* **47**(2), 315–72.
- Dubé, J.-P. & Misra, S. (2019), 'Personalized pricing and customer welfare', *Working Paper* .
- Egelman, S., Tsai, J., Cranor, L. F. & Acquisti, A. (2009), Timing is everything? the effects of timing and placement of online privacy indicators, in 'Proceedings of the SIGCHI Conference on Human Factors in Computing Systems', ACM, pp. 319–328.
- Farrell, J. (2012), 'Can privacy be just another good', *Journal on Telecommunications and High Technology Law* **10**, 251.
- Goldberg, S., Johnson, G. & Shriver, S. (2019), 'Regulating privacy online: An early look at europe's gdpr', *Available at SSRN* 3421731 .
- Goldfarb, A. & Tucker, C. (2011), 'Online display advertising: Targeting and obtrusiveness', *Marketing Science* **30**(3), 389–404.
- Goldfarb, A. & Tucker, C. (2012a), 'Privacy and innovation', *Innovation policy and the economy* **12**(1), 65–90.
- Goldfarb, A. & Tucker, C. (2012b), 'Shifts in privacy concerns', *American Economic Review* **102**(3), 349–53.

- Goldin, J. & Reck, D. (2018), Optimal defaults with normative ambiguity, in 'AEA Papers and Proceedings', Vol. 108, pp. 98–102.
- Goswami, I. & Urminsky, O. (2016), 'When should the ask be a nudge? the effect of default amounts on charitable donations', *Journal of Marketing Research* **53**(5), 829–846.
- Grossman, S. J. & Hart, O. D. (1980), 'Disclosure laws and takeover bids', *The Journal of Finance* **35**(2), 323–334.
- Harper, J. & Singleton, S. (2001), 'With a grain of salt: What consumer privacy surveys don't tell us', *Available at SSRN* 299930 .
- Heen, M., Lieberman, J. D. & Miethe, T. D. (2014), 'A comparison of different online sampling approaches for generating national samples', *Center for Crime and Justice Policy* **1**, 1–8.
- Hsee, C. K. (1996), 'The evaluability hypothesis: An explanation for preference reversals between joint and separate evaluations of alternatives', *Organizational Behavior and Human Decision Processes* **67**(3), 247–257.
- Hsee, C. K. & Zhang, J. (2010), 'General evaluability theory', *Perspectives on Psychological Science* **5**(4), 343–355.
- Hui, K.-L. (2007), 'Consumer disclosure: The effects of company information presentation and question sequence', *Working Paper* .
- Jia, J., Jin, G. Z. & Wagman, L. (2018), 'The short-run effects of gdpr on technology venture investment', *NBER Working Paper No. 25248* .
- Jia, J., Jin, G. Z. & Wagman, L. (2020), 'Gdpr and the localness of venture investment', *Available at SSRN* 3436535 .
- Jin, Y. & Vasserman, S. (2018), 'Buying data from consumers: The impact of monitoring programs in u.s. auto insurance', *Working Paper* .
- John, L. K., Acquisti, A. & Loewenstein, G. (2010), 'Strangers on a plane: Context-dependent willingness to divulge sensitive information', *Journal of Consumer Research* **37**(5), 858–873.
- Johnson, E. J., Bellman, S. & Lohse, G. L. (2002), 'Defaults, framing and privacy: Why opting in-opting out', *Marketing Letters* **13**(1), 5–15.
- Johnson, G. (2013), 'The impact of privacy policy on the auction market for online display advertising', *Simon School Working Paper* (FR 13-26).
- Johnson, G. A., Shriver, S. & Du, S. (2020), 'Consumer privacy choice in online advertising: Who opts out and at what cost to industry?', *Marketing Science* **39**(1), 33–51.
- Jovanovic, B. (1982), 'Truthful disclosure of information', *The Bell Journal of Economics* pp. 36–44.

- Kahneman, D. (1979), 'Prospect theory: An analysis of decisions under risk', *Econometrica* **47**, 278.
- Karlan, D., McConnell, M., Mullainathan, S. & Zinman, J. (2016), 'Getting to the top of mind: How reminders increase saving', *Management Science* **62**(12), 3393–3411.
- Kummer, M. & Schulte, P. (2019), 'When private information settles the bill: Money and privacy in google's market for smartphone applications', *Management Science* .
- Lee, Y. (2019), 'Revealed privacy preferences: Are privacy choices rational?', *Working Paper* .
- Madrian, B. C. & Shea, D. F. (2001), 'The power of suggestion: Inertia in 401 (k) participation and savings behavior', *The Quarterly Journal of Economics* **116**(4), 1149–1187.
- Martin, K. & Nissenbaum, H. (2016), 'Measuring privacy: An empirical test using context to expose confounding variables', *Columbia Science and Technology Law Review* **18**, 176.
- Milgrom, P. R. (1981), 'Good news and bad news: Representation theorems and applications', *The Bell Journal of Economics* pp. 380–391.
- Miller, A. R. & Tucker, C. (2017), 'Privacy protection, personalized medicine, and genetic testing', *Management Science* pp. 1–21.
- Palmeira, M. M. & Srivastava, J. (2013), 'Free offer ≠ cheap product: A selective accessibility account on the valuation of free offers', *Journal of Consumer Research* **40**(4), 644–656.
- Posner, R. A. (1981), 'The economics of privacy', *American Economic Review* **71**(2), 405–409.
- Rainie, L. & Duggan, M. (2015), 'Privacy and information sharing', *Available at: <http://www.pewinternet.org/2016/01/14/2016/Privacy-and-Information-Sharing>* .
- Shampanier, K., Mazar, N. & Ariely, D. (2007), 'Zero as a special price: The true value of free products', *Marketing Science* **26**(6), 742–757.
- Soleymanian, M., Weinberg, C. & Zhu, T. (2019), 'Sensor data, privacy, and behavioral tracking: Does usage-based auto insurance benefit drivers?', *Marketing Science* **38**(1).
- Spiekermann, S., Grossklags, J. & Berendt, B. (2001), E-privacy in 2nd generation e-commerce: Privacy preferences versus actual behavior, in 'Proceedings of the 3rd ACM conference on Electronic Commerce', ACM, pp. 38–47.
- Stigler, G. J. (1980), 'An introduction to privacy in economics and politics', *Journal of Legal Studies* **9**(4), 623–644.
- Stutzman, F., Gross, R. & Acquisti, A. (2013), 'Silent listeners: The evolution of privacy and disclosure on facebook', *Journal of Privacy and Confidentiality* **4**(2).
- Tang, H. (2019), The value of privacy: Evidence from online borrowers, Technical report, Working Paper, HEC Paris.

- Tannenbaum, D. & Ditto, P. H. (2012), 'Information asymmetries in default options', *Working Paper* .
- Thaler, R. (1980), 'Toward a positive theory of consumer choice', *Journal of Economic Behavior & Organization* **1**(1), 39–60.
- Tirole, J. (2020), 'Digital dystopia', *13th Digital Economics Conference, Toulouse* .
- Tucker, C. E. (2014), 'Social networks, personalized advertising, and privacy controls', *Journal of marketing research* **51**(5), 546–562.
- Urminsky, O. & Kivetz, R. (2011), 'Scope insensitivity and the 'mere token' effect', *Journal of Marketing Research* **48**(2), 282–295.
- Varian, H., Wallenberg, F. & Woroch, G. (2005), 'The demographics of the do-not-call list', *IEEE Security and Privacy* **3**(1), 34–39.
- Warren, S. D. & Brandeis, L. D. (1890), 'Right to privacy', *Harvard Law Review* **4**, 193.
- Watanabe, S. (2010), 'Asymptotic equivalence of bayes cross validation and widely applicable information criterion in singular learning theory', *Journal of Machine Learning Research* **11**(Dec), 3571–3594.
- Wathieu, L. & Friedman, A. (2009), An empirical approach to understanding privacy concerns, Technical report, ESMT Working Paper.

APPENDIX A

PROOF FOR PROPOSITION 1

First, define the notation for the means and covariances of preference components: $E[c_i] = \mu_c$, $Var[c_i] = \sigma_c^2$; $E[-T(d_i)] = \mu_t$, $Var[-T(d_i)] = \sigma_t^2$; $Corr(c_i, -T(d_i)) = \rho$. Note that $\Delta T(d_i) = T(F_d(d|s = 0)) - T(d_i)$, where $T(F_d(d|s = 0))$ does not vary across consumers. Therefore, $Var[\Delta T(d_i)] = \sigma_t^2$ and $Corr(c_i, \Delta T(d_i)) = \rho$. σ_c^2 and σ_t^2 respectively represent the heterogeneity of the intrinsic and instrumental preference components.

Denote the total preference for privacy as g_i . Then,

$$Corr(g_i, \Delta T(d_i)) = Corr(c_i + \Delta T(d_i), \Delta T(d_i)) = \frac{Cov(c_i + \Delta T(d_i), \Delta T(d_i))}{\sqrt{Var[c_i + \Delta T(d_i)] \cdot Var[\Delta T(d_i)]}} = \frac{\rho\sigma_c + \sigma_t}{\sqrt{\sigma_c^2 + \sigma_t^2 + 2\rho\sigma_c\sigma_t}}. \quad (A.1)$$

$Corr(g_i, \Delta T(d_i))$ captures the degree to which privacy decisions can be explained by the instrumental preference $\Delta T(d_i)$. Because a one-to-one mapping exists between instrumental preference and a consumer's type (conditional on a fixed transfer to non-disclosing consumers $T(F_d(d|s = 0))$), $Corr(g_i, \Delta T(d_i))$ is a direct assessment of the information value of non-sharing decisions for inferring consumer types. The following observations hold:

1. $Corr(g_i, \Delta T(d_i)) > 0$ if and only if $\rho + \frac{\sigma_t}{\sigma_c} > 0$.
2. $Corr(g_i, \Delta T(d_i))$ increases with $\frac{\sigma_t}{\sigma_c}$, and strictly increases with $\frac{\sigma_t}{\sigma_c}$ if $|\rho| < 1$.
3. $Corr(g_i, \Delta T(d_i))$ increases with ρ iff $\sigma_c + \rho\sigma_t > 0$, and decreases with ρ if $\sigma_c + \rho\sigma_t < 0$.

Observation 3 reveals a more nuanced relationship between the explainability of instrumental preference and the correlation between the two preference components. In particular, if $\sigma_t > \sigma_c$, a regime $\rho \in [-1, -\frac{\sigma_c}{\sigma_t}]$ exists where an increase in ρ leads to a decrease in $Corr(g_i, \Delta T(d_i))$. The reason is that when ρ is close to -1, the variation in instrumental preference dominates intrinsic preference ($\sigma_t > \sigma_c$), leading to a perfect correlation between total preference g_i and instrumental preference $\Delta T(d_i)$. Once ρ deviates away from -1, this relationship is loosened. Note that when $\sigma_t < \sigma_c$, $Corr(g_i, \Delta T(d_i))$

always increases with ρ .

The proof goes through regardless of the level of $T(F_d(d|s = 0))$. In particular, consumers need not have rational expectations, such that their beliefs about $T(F_d(d|s = 0))$ are consistent with the actual transfer that the firm gives to consumers who withhold their data. By the same token, firms need not have correct inference about consumers who choose not to share data. In other words, the conclusions above are robust to scenarios where firms actively experiment or where information is inadequate for consumers or firms to form rational beliefs. The proof also remains valid when compensation for data sharing is present.

APPENDIX B

SCREENSHOTS OF THE SURVEY

B.1 Conjoint and Demographic Questions

Figure B.1: Example Questions in the First Stage

(a) Conjoint Question

If you want to buy a smartwatch and these are the available options, which one will you choose? (Scenario 3/7)

Product	A	B	C	D
Fitness Tracking	Activity	Activity + heart rate	Activity + heart rate	Activity + heart rate
Voice Control	Yes	Yes	Yes	No
Mobile Payment	No	No	Yes	Yes
Encryption	TLS 1.0	SSL 3.0	SSL 3.0	TLS 1.0
Login Option	Pin, pattern	Pin, pattern	Pin, pattern	Pin, pattern, face
Price	\$299	\$149	\$199	\$249

☐ Product A ☐ Product B ☐ Product C ☐ Product D

☐ None of the above

(b) Demographic Question

What is your total household income for the last calendar year, before taxes?

- ☐ Less than \$25,000
- ☐ \$25,000 to \$50,000
- ☐ \$50,000 to \$75,000
- ☐ \$75,000 to \$100,000
- ☐ \$100,000 to \$200,000
- ☐ \$200,000 or more
- ☐ Prefer not to say

B.2 Compensation Schedules across Treatments

Figure B.2: Displayed Compensation Schedule: Intrinsic Treatment

(a) Main Screen

You will have the chance to win a \$20 gift card if you choose to share your responses with Odde, our corporate partner. Odde is a high-end smart device manufacturer; it hopes to use the survey data to inform product development. To encourage participants to share their feedback, **it decides to increase the probability of winning for participants who share more information** ([see details](#)).

You can choose what information to share with Odde by selecting or deselecting the boxes below. Note that only the questions that you previously chose to give an informative answer to (i.e. not stating "prefer not to say") can potentially be shared. *Any information that you choose not to share with Odde will not be obtained by the company.*

- | | | |
|--|---|---------------------------------------|
| <input type="checkbox"/> Choice task responses | <input type="checkbox"/> Ethnicity | <input type="checkbox"/> Age |
| <input type="checkbox"/> Income | <input type="checkbox"/> Marital status | <input type="checkbox"/> Gender |
| <input type="checkbox"/> Zipcode | <input type="checkbox"/> Education | <input type="checkbox"/> Kids at home |

(b) Details Screen

A participant's winning probability is calculated by the following formula:

$$\text{Probability of winning} = \text{number of boxes checked} \times 1\%$$

For example, if you decide to share your responses to 5 questions that you previously gave, your probability of winning will be 5%.

Figure B.3: Displayed Compensation Schedule: Instrumental Treatment

(a) Main Screen

You will get the chance to win another \$50 reward if you choose to share your responses with Odde, our corporate partner. Odde is a high-end smart device manufacturer; it hopes to use the survey data to inform product development.

Your probability of getting the reward will increase with the amount of information that you share. Meanwhile, Odde is designing a new smartwatch geared towards tech-savvy, high-income consumers, and wants to get more feedback from this group of people. As a result, it chooses to assign higher winning probabilities to participants who fit into this profile. In particular, **if it infers you to be wealthy or likely to purchase a smartwatch in the near future, the probability of you winning the reward will increase substantially** ([see details](#)).

You can choose what information to share with Odde by selecting or deselecting the boxes below. Note that only the questions that you previously chose to answer (i.e. not stating "prefer not to say") can potentially be shared and therefore be displayed. *Any information that you choose not to share with Odde will not be obtained by the company, and therefore will not be used for determining the winning probability.*

- | | | |
|--|------------------------------------|---------------------------------------|
| <input type="checkbox"/> Choice task responses | <input type="checkbox"/> Education | <input type="checkbox"/> Ethnicity |
| <input type="checkbox"/> Marital status | <input type="checkbox"/> Gender | <input type="checkbox"/> Kids at home |
| <input type="checkbox"/> Income | <input type="checkbox"/> Age | |

(b) Details Screen

Your winning probability is determined both by the baseline probability and by the adjustment terms. The baseline winning probability is calculated as follows:

$$\text{Baseline probability of winning} = \text{Number of boxes checked} \times 1\%$$

This baseline probability is then adjusted to encourage response sharing from the customer group that Odde intends to serve, as shown in the following chart:

Income	< \$50,000	\$50,000 – \$75,000	> \$75,000
Adjustment	-2%	Unchanged	+2%
Plan to purchase a smartwatch in the next 3 months	Somewhat or extremely unlikely	Neither likely nor unlikely	Somewhat or extremely likely
Adjustment	-2%	Unchanged	+2%

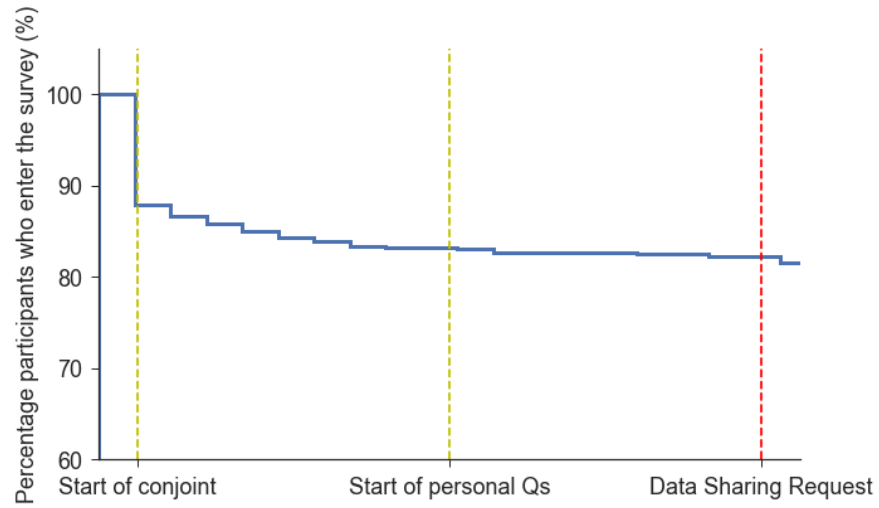
For example, if you have checked 5 boxes, then your baseline winning probability will be 5%. In addition, if the information you share indicates that your annual income is between \$75,000 and \$100,000, but you are unlikely to buy a smartwatch in the short run, then your final probability of winning will be $5 + 2 - 2\% = 5\%$. The final winning probability never goes below zero.

Any information that you choose not to share with Odde will not be accessed by the company, and therefore will not be used to adjust your winning probability. Meanwhile, *Odde might still be able to use the information that you choose to share (e.g. zipcode, age, education) to infer your income level and your willingness to purchase.*

APPENDIX C

ATTRITION

Figure C.1: Percentage of Participants Remained Throughout the Survey



APPENDIX D

CREDIBLE INTERVALS FOR INTRINSIC PREFERENCE ESTIMATES

(WTA)

Table D.1: Posterior Estimates of Mean and Standard Deviation of the Intrinsic WTA

(a) WTA Mean			(b) WTA Standard Deviation		
	mean	95% CI		mean	95% CI
income	1.870	[1.012, 3.518]	income	0.906	[0.379, 1.840]
intent	1.825	[0.981, 3.534]	intent	1.337	[0.702, 2.615]
gender	0.142	[-0.285, 0.709]	gender	0.929	[0.438, 1.965]
age	0.260	[-0.172, 0.805]	age	1.078	[0.536, 2.173]
edu	1.228	[0.619, 2.337]	edu	0.805	[0.330, 1.602]
relationship	0.687	[0.249, 1.454]	relationship	0.998	[0.477, 1.973]
kid	2.367	[1.337, 4.523]	kid	1.001	[0.465, 1.990]
zipcode	0.985	[0.450, 1.992]	zipcode	0.982	[0.455, 1.953]
race	0.980	[0.437, 2.008]	race	0.906	[0.406, 1.801]

APPENDIX E

INTRINSIC WTA ESTIMATES WITH HETEROGENEOUS SENSITIVITY TO INCOME

As a robustness check, I also calculate consumers' WTA distribution corresponding to Model 4, which allows consumers to have heterogeneous preferences in both the intrinsic value for privacy and monetary compensation. The estimated sensitivity to income is not very different among consumers. The median sensitivity is 0.15; for consumers at the bottom 2.5% quantile, $\beta = 0.13$, while for the top 2.5% quantile, $\beta = 0.18$. Table E.1 reports the posterior distribution of intrinsic WTA from Model 4. Compared to the main results in Table 6.2 and Figure 6.1, the estimated WTA distribution from Model 4 exhibits slightly larger heterogeneity among high-value variables and smaller heterogeneity among low-value ones. That being said, overall the two sets of estimates are similar both qualitatively and quantitatively.

Table E.1: Posterior Distribution of WTA in Intrinsic Preference (with Heterogeneous Sensitivity to Income)

	mean	median	2.5%	97.5%
kid	2.253	2.007	1.051	4.453
income	1.784	1.533	0.794	3.882
intent	1.742	1.261	0.341	5.097
edu	1.189	1.008	0.224	2.960
zipcode	0.959	0.740	-0.114	2.971
race	0.951	0.734	-0.059	2.919
relationship	0.691	0.404	-0.359	2.870
age	0.271	0.081	-0.927	2.647
gender	0.149	-0.010	-0.897	2.150

Note: Numbers in this table refer to statistics associated with the estimated WTA distribution among consumers; these are measures of preference heterogeneity.

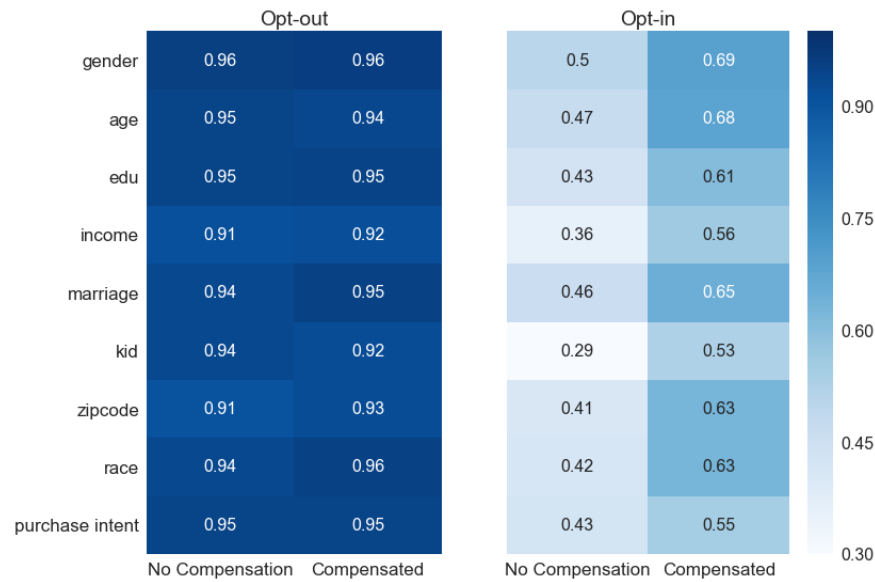
APPENDIX F

PSYCHOLOGICAL FACTORS

F.1 The Default Frame

Figure F.1 visualizes the data-sharing frequency in different default regimes. Under the opt-out regime, almost everyone shares everything, regardless of the amount and format of compensation. The lack of choice variation in the opt-out regime does not per se imply a weaker preference for privacy or economic incentives; it simply means the impact of a “share-all” frame is strong enough to dominate other components in utility.

Figure F.1: Frequency of Data Sharing under Different Policy Regimes



Interaction between the default regime and privacy preferences. The literature has widely acknowledged the fact that default frame influences choices (Kahneman 1979, Thaler 1980, Johnson et al. 2002, Acquisti et al. 2013). However, little consensus exists on how or how much default affects choices. To flexibly characterize how default influences privacy choices, I estimate separate models for each default frame. Table F.1 displays the estimated privacy preferences under opt-in and opt-out regimes. In the comparison below, I acknowledge the scaling differences across the models, and normalize parameters to the

same (dollar) scale when needed. The scaling does not affect the sign of parameters, nor the sensitivity ranking across categories of data within the same model. The comparison of belief parameters w and δ are not affected by the scaling either, since these parameters directly apply to the sensitivity to compensation parameter β .¹

Table F.1: Privacy Preferences across Default Frames

Default Frame		Opt-In		Opt-Out	
		mean	95% CI	mean	95% CI
intrinsic	c_{income}	0.906	[0.588, 1.323]	-1.903	[-2.705, -1.134]
	c_{intent}	0.826	[0.419, 1.322]	-2.704	[-3.653, -2.127]
	c_{gender}	0.189	[-0.162, 0.664]	-2.988	[-3.956, -2.184]
	c_{age}	0.262	[-0.088, 0.733]	-2.429	[-3.127, -1.729]
	c_{edu}	0.624	[0.329, 1.051]	-2.739	[-3.301, -2.161]
	$c_{relationship}$	0.497	[0.124, 1.010]	-2.734	[-3.331, -2.105]
	c_{kid}	1.109	[0.790, 1.461]	-2.143	[-2.692, -1.380]
	c_{zip}	0.560	[0.227, 1.066]	-2.093	[-3.448, -1.328]
	c_{race}	0.604	[0.285, 1.104]	-2.660	[-3.518, -1.805]
instrumental	w_{income}	2.118	[0.108, 3.989]	1.994	[0.136, 3.893]
	w_{intent}	1.942	[0.383, 3.762]	1.995	[0.109, 3.909]
	$\tilde{\delta}_{income,0}$	0.047	[-0.186, 0.282]	0.054	[-0.183, 0.280]
	$\tilde{\delta}_{income,1}$	0.037	[-0.192, 0.284]	0.052	[-0.185, 0.286]
	$\tilde{\delta}_{intent,0}$	0.059	[-0.352, 0.379]	-0.121	[-0.391, 0.350]
	$\tilde{\delta}_{intent,1}$	-0.049	[-0.362, 0.324]	-0.129	[-0.384, 0.281]
sensitivity to compensation	β	0.146	[0.070, 0.2359]	0.046	[0.001, 0.141]
log posterior		-7476	[-7540, -7407]	-2075	[-2166, -1981]

Note: The models are estimated separately for each default frame. Variables are normalized using the Gelman method before estimation. Both models allow for heterogeneity in intrinsic preferences.

To compare intrinsic-preference parameters across models, Figure F.2 displays the willingness to pay (WTP) of intrinsic preferences, which are heavily influenced by the default frame. The negative WTPs imply that once data are obtained by companies, consumers will not take back their control over personal data, unless they are incentivized

1. To see this point, note that if the instrumental utility is $w \cdot \beta \cdot \Delta E[d]$ in the utility space, then its dollar value is simply $w \cdot \Delta E[d]$.

by the amount indicated by the WTP. In my data, the gap between median WTA and median WTP amounts to \$69.18 (income) to \$88.06 (gender). In comparison, previous literature estimates dollar values of default in 401(k) plan enrollment decisions that range from \$37–\$54 (Bernheim et al. 2015) to \$1,200 (DellaVigna 2009). However, the WTP estimates are very noisy, due to the fact that the estimated sensitivity to compensation in the opt-out regime is close to zero (see Table F.2 for credible interval estimates).

Interestingly, Table F.1 shows that the default frame does not heavily influence consumer beliefs about the instrumental payoff. The differential impacts of default suggest that while subjective evaluations are more susceptible to the influence of the default condition, objective evaluations—beliefs about the instrumental payoff—are less so. In view of this fact, distinguishing between the intrinsic and instrumental preferences also reveals how default (and potentially other psychological factors) influences different privacy motives differently.

Figure F.2: Posterior Predicted Density of WTP in Intrinsic Preference

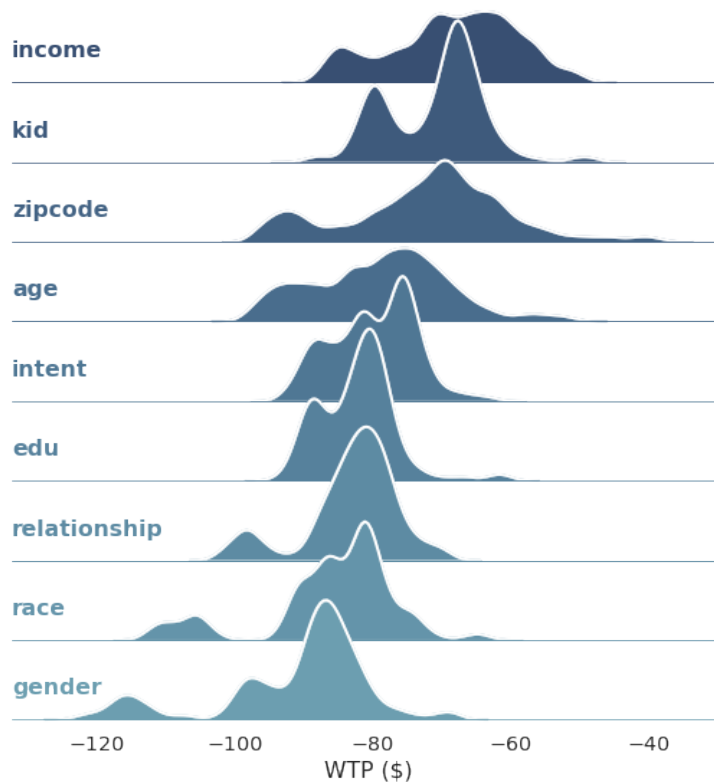


Table F.2: Posterior Estimates of Parameters Associated with Intrinsic WTP Distribution

(a) WTP Mean			(b) WTP Standard Deviation		
	mean	95% CI		mean	95% CI
income	-66.59	[-621.55, -6.92]	income	70.23	[2.41, 479.59]
intent	-78.87	[-733.52, -8.29]	intent	35.52	[0.96, 249.39]
gender	-89.84	[-866.19, -9.28]	gender	57.36	[1.92, 360.06]
age	-76.57	[-722.37, -8.03]	age	76.40	[2.70, 522.52]
edu	-81.10	[-767.11, -8.41]	edu	34.59	[1.22, 248.68]
relationship	-82.10	[-773.98, -8.63]	relationship	43.54	[1.22, 286.02]
kid	-70.15	[-634.81, -7.34]	kid	29.53	[1.37, 185.01]
zipcode	-70.52	[-653.87, -7.46]	zipcode	57.63	[3.04, 405.01]
race	-86.69	[-834.71, -8.97]	race	28.37	[1.30, 213.90]

The managerial implication is immediate. With a regulation that mandates opt-out consent, firms can still collect most customer data even if consumers are fully informed when making privacy choices. However, once the firm moves to an opt-in regime, it will incur substantial losses in the amount of data collected. The default paradigm is also useful for thinking about the real impact of data-portability rights.² Taking the incumbent as the default choice, consumers are less likely to opt out of incumbent tracking and transfer data to its competitors, unless the expected utility gain from switching is substantially large.

Mechanisms underlying the default effect. When the research goal is to estimate consumer welfare, distinguishing different mechanisms becomes necessary, because each mechanism implies a different way to calculate the welfare utility from the behavioral utility (i.e. revealed preference from choice). For example, if consumers stick to default due to inattention, the effect of default on choice is separate from welfare utility. On the other hand, if default shifts choices via implicit endorsement, that is, by providing information that changes consumers' evaluation of different options, it should be part of the welfare utility (see Bernheim & Rangel (2010), Goldin & Reck (2018) for a review).

² GDPR Article 20 and CCPA Title 1.81.5, Section 1798.100 (d).

Within my experiment context, three mechanisms are likely (broadly categorized):³

- **Inattention/cost of gathering information.** Both of them can be either rational or irrational. Rational means a decision maker endogenously determines whether it is worthwhile to pay attention or gather additional information, depending on the stakes of the decision. Irrational means no such ex-ante trade-off is involved (Karlan et al. 2016).
- **Anchoring (including endorsement effect).** It means that the default option serves as an attraction point. This can be caused by the fact that its advantages are more salient (Goswami & Urminsky 2016), or because it is viewed as an implicit recommendation (Madrian & Shea 2001).
- **Loss aversion (also called reference dependence).** It means that loss looms larger than gains (Kahneman 1979, Thaler 1980).

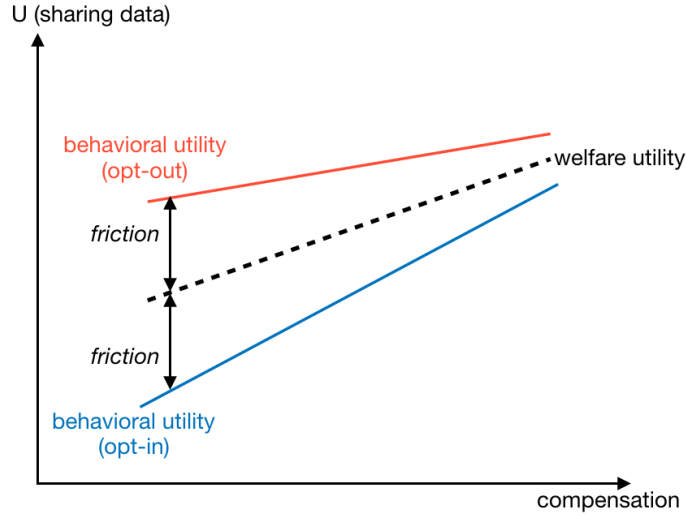
Below, I discuss how patterns in my data can help distinguish between different mechanisms:

Sign switch between WTA and WTP is inconsistent with pure loss aversion. Loss aversion merely scales gains and losses, which is sign preserving. The coexistence of positive WTA and negative WTP implies that other mechanisms must be at work.

Changes in sensitivity ranking across data suggest anchoring effect. Neither inattention (information gathering cost) nor loss aversion can generate a switch of rankings across different categories of data. Under inattention, the decision maker either sticks to default, or makes attentive choices that rank different options in the same way across frames. Loss aversion implies privacy costs for sharing different data are scaled by the same factor. Under anchoring, however, values that are more certain to the decision maker are less susceptible to the influence of defaults, which can generate changes in relative sensitivity

3. Other mechanisms do not apply to my setting. For example, choice procrastination is unlikely to be the driving factor because participants are not given a long time frame for making choices. Mechanical switching cost is also unlikely, since changing their decision takes only one click.

Figure F.3: Sensitivity to Compensation with Rational Inattention



among data. The fact that beliefs about the instrumental payoff are less influenced by the default frame is also consistent with this explanation.

Greater sensitivity to compensation in the opt-in frame is consistent with rational inattention but not loss aversion. The impact of rational inattention on sensitivity to economic payoffs can be best illustrated by Figure F.3. Since the utility from sharing data increases with the amount of compensation, the welfare utility is upward sloping (the dashed line). The behavioral utility in the opt-in frame is always closer to zero than the welfare utility (the blue line below), while behavioral utility in the opt-out frame exhibits the opposite pattern (the red line above). Under rational inattention, the impact of the default frame diminishes as the stake increases, generating a steeper utility response in the opt-in frame and a flatter response in the opt-out frame. By contrast, under loss-aversion theory, not sharing means “gaining” privacy and “losing” money in the opt-out frame; therefore, decision makers should be more sensitive to monetary payoff when the default is opt-out.

To directly compare the sensitivity to economic payoffs across frames, data from different default conditions should be contained in the same model to avoid mechanical differences caused by scaling. Table F.3 displays the estimation result corresponding to

β from the pooled regression. It shows that participants in the opt-in frame are indeed more sensitive to compensation than in the opt-out frame.

Table F.3: Sensitivity to Economic Payoff across Default Frames (Pooled Estimation)

Default Frame	Opt-In		Opt-Out	
	mean	95% CI	mean	95% CI
β	0.138	[0.059, 0.231]	0.041	[0.001, 0.138]

To conclude, my data suggest that inattention/information-acquisition costs and anchoring are the main drivers of the default effect in my experiment. The fact that anchoring is one of the drivers implies the impact of default is likely to be asymmetric: Previous literature shows that the welfare utility is often closer to the behavioral utility in the opt-in frame (Madrian & Shea 2001, Tannenbaum & Ditto 2012). Given this normative ambiguity, I refrain from trying to back out the welfare utility given the current data constraints.

F.2 Other Psychological Factors

The model includes a behavioral response term $m \cdot (p_i \geq 0) \cdot s_i$, to account for a combination of a mere-incentive effect and potential anchoring effects at the start of the survey. Behavioral response to the mere presence of incentives is well documented in the psychology literature (Shampanier et al. 2007, Urminsky & Kivetz 2011, Palmeira & Srivastava 2013), which can be explained by the theory that people are insensitive to scopes when evaluating options separately (Hsee 1996, Hsee & Zhang 2010). In treatment groups that distribute positive amounts of compensation, participants are told at the beginning that they can enter a gift-card lottery upon finishing the survey. This information may inadvertently create an additional anchoring effect, making all participants in these groups more inclined to share their data in order to get the anticipated gift-card rewards. The parameter m captures the combination of these two forces. Under the second mechanism, the additional anchoring effect will be stronger for participants in the opt-in group (be-

cause an opt-out condition per se also has a substitutive anchoring effect); this possibility is accounted for by having separate m 's for different default conditions.

In the opt-in frame, the point estimate for m is 0.76, with the 95% credible interval being $[0.65, 0.87]$. In the opt-out frame, the point estimate is 0.07, with the credible interval being $[-0.17, 0.30]$. The strong effect asymmetry and the fact that the effect is almost non-existent in the opt-out condition suggest anchoring is more likely to be the main driver of this effect.