Choice Architecture, Privacy Valuations, and Data Market Efficiency

Tesary Lin* Avner Strulov-Shlain†

September 12, 2022

[*PRELIMINARY; COMMENTS ARE WELCOME*]

Abstract

How much do consumers value their personal data? How do their privacy valuations change under the influence of choice architecture? How does this influence impede the efficiency of a data market, not only in the equilibrium price and quantity of data traded, but also in the representativeness of data collected? To answer these questions, we run a large-scale choice experiment to elicit consumers' valuation for their Facebook data while randomizing two common choice frames: default and price anchor. Consumers' valuations for data vary in a systematic and coherent manner. An opt-out default decreases valuations by 14.4% compared to opt-in, while a \$0-\$50 price anchor decreases valuation by 51.1% compared to a \$50-\$100 anchor. Moreover, the susceptibility to frame influence is negatively correlated with consumers' valuation absent frames. Our results suggest that choice architecture can impede the efficiency of a data market beyond common mechanisms in a conventional market. In particular, a choice frame that maximizes the supply of consumer data may exacerbate the statistical bias in data collected by firms, thus decreasing the return of data for firms.

Keywords: privacy, choice architecture, market for data, experiment, social media

*Boston University Questrom School of Business; tesary@bu.edu

†University of Chicago Booth School of Business; avner.strulov-shlain@chicagobooth.edu. The authors thank Dan Bartels, Josh Dean, Berkeley Dietvorst, Kwabena Donkor, Andrey Fradkin, Ali Goli, Tanjim Hossain, Alex Imas, Jihye Jeon, Garrett Johnson, Yucheng Liang, Chiara Longoni, Nina Mazar, Sanjog Misra, Ilya Morozov, Olivia Natan, Omid Rafieian, Heather Sarsons, Ananya Sen, Oleg Urminsky, and Giorgos Zervas for their helpful comments. We thank Christy Kang, Paulina Koenig, and Kaushal Addanki for their excellent research assistance. This research is funded by the Becker Friedman Institute at the University of Chicago. It was approved by the Institutional Review Boards at the University of Chicago (IRB21-1376) and Boston University (IRB-6239X).

1 Introduction

How companies should collect and use consumers' personal data is at the center of recent policy debates. Yet, a missing piece in these debates is how much consumers value their own data. Although it is possible to observe consumers' privacy decisions in the field, these decisions are often influenced by choice architectures that companies deploy. As a result, their privacy choices may not accurately reflect their best interest, leading to failures in market-based mechanisms and privacy policies that rely on consumer choices. Further, different types of people may be more affected by choice architecture than others, creating distortions in who shares their data.

In this paper, we ask the following questions: How do choice frames influence consumers' privacy valuations? What is the heterogeneity of the choice frame effects, in particular, how much do they distort different parts of a data-supply curve? How do the choice frames change the composition of consumers willing to share their data beyond the distortive effects on the prices and the quantity of data traded? The returns of data to firms depend on both its quantity and its representativeness. Therefore, to assess how choice architecture affects the efficiency of the data market, we must account for its effect on both the volume of data traded at any given price point and the composition of data traded.

To answer these questions, we recruited a sample of 5,028 Facebook users and elicited their willingness-to-accept (WTA) to sell their data while randomizing the choice frames they faced. Within participants, we ask their valuation for sharing the following variable with researchers and advertisers: *about me* (their information on the profile page), *posts*, *likes*, *friends and followers*, and *responses to our survey*. For each variable, we elicited WTA using a multiple price list (Kahneman et al. 1990, Andersen et al. 2006), followed by a free-text entry. Across participants, we randomized the choice default and the price anchor. *Default* varies between opt-in, opt-out, and active choice. *Price anchor* is the range of prices in the multiple price list, which is either \$0-\$50 (low) or \$50-\$100 (high). We choose default and price anchor because they are either common or likely to be deployed if companies can buy data from consumers. We also collect consumer characteristics to explore heterogeneity in privacy valuations and choice frame effects. These variables include demographics, digital literacy, social media and Internet usage, and their belief about what data are already available to Facebook and the public.

Consumers' valuations for data are substantially different across both individuals and data. Across individuals, valuations for the same data range from \$0 all the way up to infinity, with 20% having their valuations above \$100. When we top-code the values at 119\$, the mean valuation across data is \$67. The difference in median WTA between the most valuable (friends and followers) and the second-most valuable data (posts) is \$11.5, while the difference between the most and the least valuable data (survey responses) is \$34.9.

We also find a significant influence of choice frames on consumers' valuations. Consumers decrease their valuations by 51.1% in the *low* versus *high* price anchor condition. Compared to an *opt-out* default, *active choice* increases the valuation of data by 6.9%, and *opt-in* increases the valuation by 14.4%. On average, the difference in valuation is \$14.7 lower due to the *low* price anchor, and is \$2.9 and \$5.9 higher in the *active* and *opt-in* defaults, respectively.

Although the qualitative effect of defaults is universal, the price anchor distorts data valuations in complex ways. Valuations bunch towards the endpoints of a price range. At the same time, the proportion of consumers who report extremely high values (e.g., "I do not want to share my data at any price") increases from 15.5% to 21.2% as a result of a low price anchor. This pattern suggests that a low anchor can cause "backlash" among some consumers while decreasing the valuation of others.

We deploy a set of causal forest models (Athey et al. 2019) to explore what consumer attributes correlate with the heterogeneity in data values and frame effects, as well as the correlation between the two. We find that consumers who value their data less across frames are overall younger, poorer, and less educated; they are also more likely to click on ads while on Facebook. Interestingly, these attributes also predict larger frame effects. As a result, consumers' valuations of their data and the choice frame effects are negatively correlated. Note that this negative correlation at the individual level is masked when we simply compare the aggregate supply curves with and without frames, as such comparison assumes that the frames do not change consumers' relative positions on the supply curve.

The negative correlation between privacy valuation and choice frame effects bears two implications on the efficiency of a data market. First, when a firm can choose the choice frame while buying data, it often wants to set a frame that maximizes the supply of data at any given price. The supply-maximizing frame is more effective when the equilibrium quantity of data sold is small absent frames. As a result, the loss of total surplus is larger when the market for data is not saturated. Second, such a supply-maximizing frame may exacerbate the statistical bias in the data collected by the firm. Consumer groups who already value their data less now give up their data even more willingly due to the frame. In this case, the collected data may oversample this group even more, while possibly undersampling other consumer groups if the firm also sets a lower price for data due to the supply expansion. A biased dataset compromises the quality of data-driven analytics and innovation. As such, this bias can counteract the benefits of collecting more data under this choice frame, and further reduce market efficiency and social welfare.

Despite its flexibility and power, a causal forest cannot pick up crucial features of the frame effects, such as bunching and backlash caused by the price anchor. Thus, we estimate a micro-founded model to capture these key features in the data. In our model, consumers have a prior on how much their data is worth. Instead of reporting this prior, their reported valuation is distorted by choice frames. An opt-in frame shifts up the reported valuation, while an opt-out frame does the opposite. The price anchor further attracts the reported valuation, which results from consumers

minimizing the joint distance from the reported valuation to the shifted prior and to the anchor. By allowing the second distance to have an L1-norm component, the model rationalizes bunching at endpoints. In addition, we allow consumers to ignore the anchor and report extremely high values when the anchor price is too far away from their shifted prior. In doing so, the model also accounts for backlash. We estimate the model using deep convolutional neural nets with influence function adjustments (Farrell et al. 2020), which allows us to accommodate heterogeneity efficiently under a structural model framework.

We plan to perform several counterfactual analyses to demonstrate how a firm's choice architecture design can change different aspects of data market efficiency. We assume that the firm can optimize the choice frame, uniform or personalized, while taking their demand curve for data as fixed. We examine how different objectives of the firm can lead to different choice architecture and market efficiency outcomes. We show that a volume maximizing (variance minimizing) objective and a bias minimizing objective typically lead to different choice architecture implementations, then characterize the bias-volume frontier. Depending on the objective of learning, the optimal data collection strategy involves a bias-variance tradeoff that lands the firm somewhere on this frontier. Finally, we discuss policy implications when firms can optimize the choice architecture.

Throughout the analysis, we assume that frames affect people's reported valuation but not their true value of privacy: in the language of Kahneman et al. (1997), choice frames can shift decision utility but not the consumption utility. We believe this assumption is reasonable, in that our choice frames are neither informative nor able to change consumers' subjective wellbeings in any meaningful way.

Our paper contributes to the existing literature on several fronts. The first is the literature on measuring consumers' value of privacy using a revealed preference approach (Goldfarb & Tucker 2012, Athey et al. 2017, Kummer & Schulte 2019). Although a wealth of papers have examined the value of consumer data to firms (Rossi et al. 1996, Trusov et al. 2016, Miller & Skiera 2017, Rafieian & Yoganarasimhan 2018, Bajari et al. 2019, Sun et al. 2021, Wernerfelt et al. 2022), empirical studies on the dollar value of data to consumers are nascent (Acquisti et al. 2013, Lin 2022, Tang 2019, Collis et al. 2020). One potential reason is that privacy preferences are considered context-specific (Martin & Nissenbaum 2016) and illusive. Lin (2022) highlights consumers' economic reasoning in different data usage scenarios as a contributor to this context effect, which is a normal part of market functioning. In contrast, we maintain the economic scenario as fixed, and focus on how the choice frame-induced specificity impedes the functioning of a market for data.

The second is the literature that examines the effects of choice architecture on privacy choices (Johnson et al. 2002, Acquisti et al. 2012, 2013, Athey et al. 2017, Adjerid et al. 2019, Kormylo & Adjerid 2021). Our goal is not to settle on documenting the choice architecture effects, but rather to relate the frame effects to the efficiency of the data market. We achieve this goal by tracing the effect heterogeneity along the data supply curve, and discussing the implication of such heterogeneity when the firms can set the price for data and the choice frame at the same time. As such, our

focus connects two disjoint threads of literature: behavioral biases in privacy valuations, and the efficiency of markets for data (Arrieta-Ibarra et al. 2018, Acemoglu et al. 2019, Bergemann et al. 2020).

Our project is also related to the broad literature in behavioral industrial organization that shows how consumers' behavioral biases create market inefficiency (see Heidhues & Kőszegi 2018 for a review); also related is the work on deceptive marketing (Peltzman 1981, Rao & Wang 2017). Our work is close to Taubinsky & Rees-Jones (2018), who argue that understanding the heterogeneity in behavioral biases is crucial for market efficiency assessments. In conventional markets, consumer biases decrease market efficiency by distorting the equilibrium trading quantity and price. In the context of data markets, we highlight the importance of considering a third dimension—the composition of consumers who trade their data.

A recent stream of literature characterizes externality in privacy choices when data shared by one consumer can be used to learn about others (Bergemann et al. 2020, Ichihashi 2021, Markovich & Yehezkel 2021). Although our starting point is consumers' internality (i.e., behavioral bias) in data sharing, the heterogeneity of this internality creates statistical biases in data shared and, in turn, the ability of the firm to leverage existing data to learn about other consumers (i.e., externality). How biases in the sample data complicate the externality pattern is not often discussed in the prior papers, which we think are worth more emphasis.

The rest of the paper is organized as follows. Section 2 introduces our conceptual model, and illustrates why accounting for the joint heterogeneity between privacy valuation and frame effect is crucial for market efficiency assessments. Section 3 illustrates the design of our experiment. Section 4 describes our data and reduced-form evidence, followed by a description of our structural model in Section 5. Section 6 introduces our intended counterfactuals and concludes.

2 Conceptual Framework

In this section, we lay out the framework that underpins our experiment design and subsequent analysis. We distinguish two types of data valuation. The first is a frame-neutral valuation v_0, which is the would-be valuation if all choice frames were absent; this valuation reflects consumers' prior knowledge about their consumption utility, that is, the willingness-to-accept for trading data that reflects their best interest. The second is the expressed valuation $\widetilde{v}$, which starts from the frame-neutral valuation but is subject to the influence of choice architecture, θ:

$$\widetilde{v} = f(v_0; \theta). \tag{1}$$

In other words, the choice architecture creates a gap between the expressed valuation and the consumers' prior judgment of what is best for them.[1]

We situate the two valuations in the context of a data market, where the firm has a demand for data and can choose the price and choice frame at the same time. Although $\widetilde{v}$ may be larger or smaller than v_0, the firm often wants to choose a frame that delivers the lowest $\widetilde{v}$ to maximize its gain per dollar offered. In essence, the firm chooses a choice frame that pushes the consumers' data supply curve towards the right (see Figure 1a). As a result, the equilibrium quantity of data traded is larger and the equilibrium price lower compared to the frame-neutral optimal. The welfare loss among consumers is represented by the sum of green and red regions. They represent consumers who sold their data under the frame-neutral supply and now sell it for a lower price, plus new consumers who now sell their data for a price below their frame-neutral valuation. On the other hand, the firm gains surplus equal to the green region. The total welfare loss is thus the net of the two, represented by the red triangle (ΔTW_1).

To understand how the choice frame distortion translates to market inefficiency and consumer welfare loss, we need to understand the heterogeneity of the choice frame effects, and in particular, how this effect is different for consumers situated at different parts of the data supply curve. It is the distortion at the margin where demand and supply intersect that matters. Example 2 of Figure 1 illustrates this point. Assume every consumer's valuation shifts down vertically and there is no change in their relative preference rankings. In this case, demand curve D_1 induces more total welfare loss than demand curve D_2, because the gap around the region where D_1 intersects with supply curves is larger.

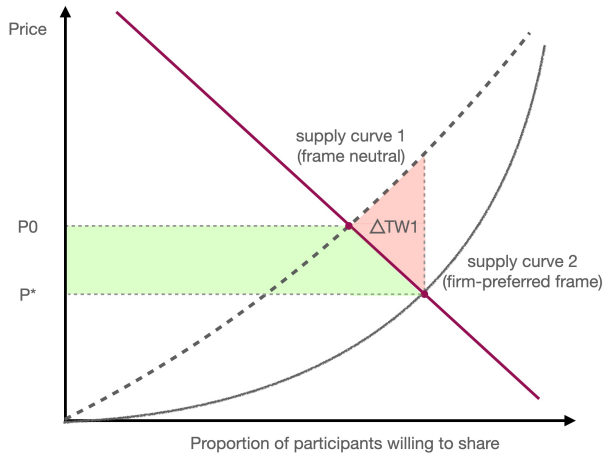
However, a simple comparison of supply curves is insufficient. To accurately calculate the welfare loss, we must also see the shifts in each consumer's valuation. As a simple example, imagine the scenario where some consumers decrease their valuation while others increase theirs when facing the same choice frame. In this case, even if the choice frame does not shift the aggregate supply curve, it still creates distortion because some consumers will over-share while others under-share their data under the frame. More generally, an accurate measure of individual-level utility distortion is crucial even when the choice frame distorts everyone's valuation in the same direction. To see this, we can look at Figure 1c. Although consumer A has a lower valuation than consumer B absent the frame, her valuation is less influenced by the frame. As a result, her reported valuation is still higher than the price P^* set by the firm, and she chooses not to share her data. In contrast, B's valuation under the frame is lower than P^* and thus chooses to share. The total welfare loss is now larger than what we would have estimated if we assumed the relative preferences among consumers did not change, represented by the new shaded region (ΔTW_3).

So far, we have illustrated how choice frames can induce behavioral distortions and impede market efficiency by treating data as a standard commodity. However, consumer data is not a

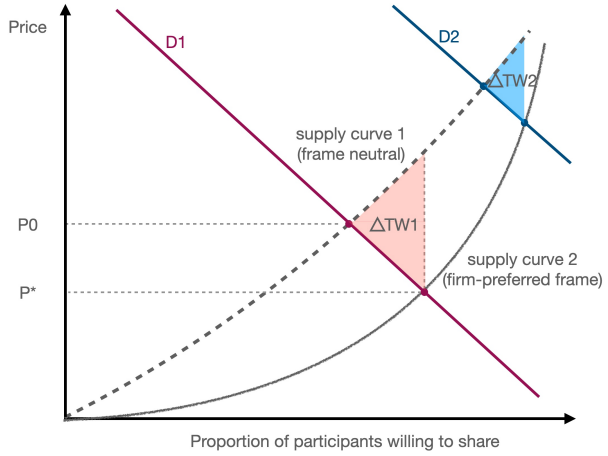
[1]We do not argue that the prior v_0 is the "true" value of the data. Rather, it is a potentially noisy representation of the true value.

Figure 1: Distortion of Data Market Due to Choice Frames: An Illustration

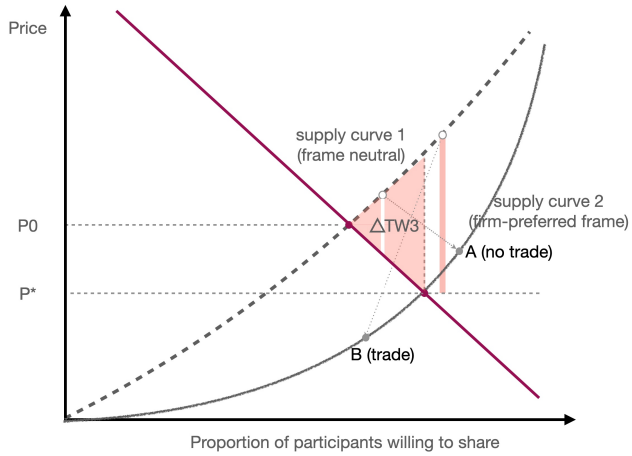
(a) Example 1



(b) Example 2



(c) Example 3



standard commodity. One unique feature of data is that its value to firms depends not only on its volume but also its representativeness. Therefore, when evaluating the impact of choice architecture, we care about not just its impact on the volume of data traded, but also which consumers are more likely to trade and how that affects the bias in the data sample collected by the firm.

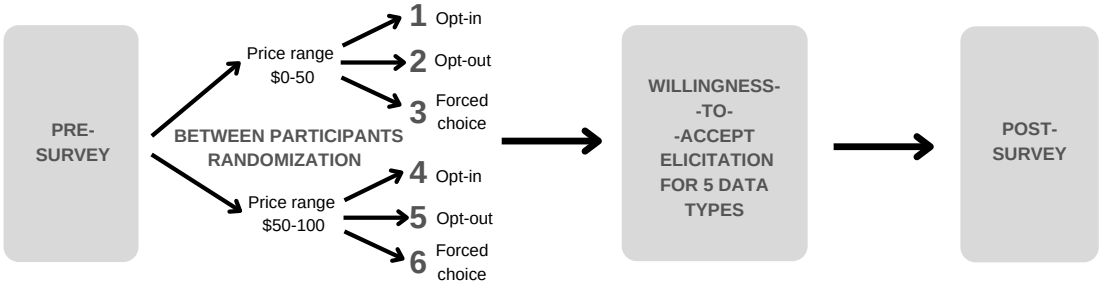
How will the consideration of sample bias affect our evaluation of frame effects? Imagine the following example: Absent the choice frame effects, low-income consumers value their privacy less compared with wealthier ones; however, their valuations are also more susceptible to the influence of choice architecture. Without the choice frame, for any given price the firm would have under-sampled wealthier consumers while over-sampling poorer ones. With a choice frame that pushes the data supply downwards, such selection bias becomes even more severe because poorer consumers push their valuations down even further compared to the wealthier ones. In this example, a volume maximizing choice frame would have helped the firm in a regular commodity market, but can end up harming the firm by inducing more bias in the collected data.

In what follows, we will first use an experiment to induce randomized changes in choice frames, while collecting consumer characteristics to explore the heterogeneity in the frame effects. Then we will put these data into a model, which allows us to show how the heterogeneity in frame effects correlates with their neutral valuation for data, as well as how the frames change the composition of consumers willing to share data at any given price point. In the counterfactual analysis, we will explore strategies that the firm may adopt to maximize the volume or to alleviate the bias in the data collected, as well as the tradeoff between the variance and bias. We then discuss how these various strategies may affect market efficiency and total surplus.

3 Experiment Design

The main component of our experiment is a multiple price list (MPL) that solicits consumers' valuation for their Facebook data in an incentive-compatible fashion. To test the effects of choice frames on reported valuation, we randomize the default choice and the range of price list when implementing the MPL. We also use baseline and endline questions to measure consumer characteristics, their Internet and social media usages, and their beliefs about which of their data are already available. Figure 2 summarizes the experiment flow. Below, we start with the sample recruiting procedure, then introduce the value solicitation components, the choice frame treatments, and end with the design of survey questions that measure consumer characteristics.

Figure 2: Experiment Flow



3.1 Sample Sources

We recruit our participants from two sources: Facebook Ads and Prolific. Using these sample sources confer two advantages. First, both sources allow us to screen participants by Facebook account availability.[2] Including only participants who have an active Facebook account ensures that their data sharing decision reflects only their privacy preferences and not the availability of their data. Second, including both samples facilitates external validity analysis. Existing studies that measure privacy preference often use survey panels or student populations, who may have lower valuations for privacy or respond to choice frames differently compared with the population. Having both the Facebook sample and a survey panel (Prolific) sample allows us to examine this possibility. We restrict our sample to English speakers living in the US, aged between 18 and 64.

To minimize selection into the experiment, we do not disclose the specific research topic in the recruiting ad (for the Facebook sample; see Figure A.1) or the study invitation email (for Prolific). A user who clicks on the survey invitation link from the recruiting ad or email is directed to our study introduction page. The introduction explains that we are university researchers who want to understand the public's social media usage and perceptions,[3] then ask for their consent to enter the study.

3.2 The Multiple Price List and Choice Frame Treatments

We measure participants' valuation for their Facebook data using a multiple price list (MPL). MPL resembles Becker-DeGroot-Marschack (Becker et al. 1964) in its usage of a lottery to ensure incentive compatibility. Its advantage is simplicity: since MPL uses simple yes/no questions at different price levels instead of a second-price auction, it is easier for participants to understand the procedure and why telling the truth is optimal. Standard MPLs are take-it-or-leave-it offers

[2]When recruiting the Facebook sample, we restrict our ad placement to positions only viewable by logged-in Facebook users, which excludes Facebook Audience Network and Instagram. For the Prolific sample, we use the internal selection tool provided by the platform to enlist only Facebook users.

[3]This information includes their data sharing behavior and privacy attitudes on social media, but we did not explicitly mention that on the introduction page to minimize selection bias.

repeated at discrete price points, giving us only the interval that includes a consumer's valuation for their data. To get more granular numbers, we follow each MPL with an open-text question asking participants their exact valuation of the data. Suppose a consumer's response is chosen by the lottery. In this case, the data-money exchange occurs if the number in their free-text entry is lower or equal to the number randomly generated by the computer.

After answering the baseline questions, participants see a message that asks their valuation for sharing data with advertisers. An example of the MPL procedure follows, showing them how their choices and the random price generated by the computer co-determine whether the data exchange will occur:

> Your answers to the survey questions, and other information, can help us understand browsing behavior better. It can also help companies and advertisers provide more products that they think you like, and show you fewer products that you are less likely to buy. **Would you be willing to share more data with us and advertisers?**
>
> **If you will, we will pay you a fair price.** When you started the survey, the computer already randomly selected if you will be asked to provide data at the end of the survey, and also randomly chose a price we will pay for it. If you are selected to participate, your payment and data shared with advertisers will base on what you choose. So you should answer carefully!
>
> For example: suppose you choose to share some data for a price of $Y. If the computer chose a price lower than $Y, you will not be asked to share the data and will not be paid. If the computer chose a price larger than $Y, you will be asked to download a copy of your data from Facebook and send them to us; you will then get the price the computer has chosen.

Participants then go through a practice question, where we ask them to imagine selling a gift card worth $14.5 by responding to a multiple price list. If a participant gives a value different from $14.5, we display an error prompt, asking them to think again and showing how a truthful response is optimal (see Figure A.2).

After the practice round, we show participants the actual MPL questions to get their valuation for different types of data. We inform participants that the price is randomly drawn between $0-$95 and does not depend on their response. Each participant receives five rounds of MPLs in random order, one for each personal variable. The following list shows the five variables and the definition we show to participants:

- *"About me" page: your Facebook information page;*
- *Posts: your Facebook posts and feed history;*
- *Likes: the posts and pages you liked on Facebook;*
- *Friends and followers: the people you befriended and followed, and people who followed you;*
- *Survey answers: the answers you gave earlier on your browsing behavior and demographics.*

We independently vary the following choice frames across participants. The first is the default choices in the multiple price list. A pre-selected "yes" for the question "will you share your posts for $50" is an *opt-out* default, while a pre-selected "no" is *opt-in*. For *active choice*, neither option is pre-selected; participants will have to click on one of them to answer the questions and proceed to the next screen. Second is the range of prices offered in MPL: the *low price* condition has prices between $0 and $50, while the *high price* condition ranges from $50 to $100. These choice frames are always the same within participant.

Apart from the choice frames, we also randomize the time range of behavioral data (*posts* and *likes*) between participants, which varies between *1 month*, *1 year*, and *since joining Facebook*. In doing so, we vary the value of data in a direction known to us. The goal is to see if consumers respond to the scope of data requested when valuing their personal data.

3.3 Multiple Price List for a Benchmark Product

We include a second practice question between the first practice question and the MPLs for personal data. In this question, we instruct consumers to imagine given either a $60 (if they are in the *low price* condition) or a $40 token (for the *high price* condition), and they need to choose the price at which they agree to sell it. Participants then go through our two-stage MPL procedure, with the same choice frame as the one they will face in the personal data MPL questions.

The second practice question fulfills two purposes. First, it demonstrates to the consumer that they can (and in this case, should) report values outside the MPL price range. Second, it allows us to control for the choice frame effects for standard products by including consumer responses to this question as additional covariates (Dizon-Ross & Jayachandran 2022).

3.4 Baseline and Endline Questions

The survey includes questions that capture consumer characteristics, their internet and social media usage, as well as their beliefs about personal data availability. In general, we put questions related to data sharing in the endline survey, so that participants are not primed to consider privacy before the MPL questions; otherwise, we include the questions in the baseline survey.

The baseline questions measure participants' social media consumption behavior and their demographics. We ask participants about their time spent on Facebook and online, when they started using Facebook, and their engagements with merchants and ads on Facebook. For demographics, we record their age, gender, ethnicity, income, and education.

The endline questions include measures of information-seeking behavior and participants' belief about data availability. To measure information seeking, we include the following question: "*Did you look up additional information when answering the questions that ask how much you value your*

Facebook data?" If they answer "yes", we ask what kind of information they looked up. To measure their belief about data already available to various parties in the market, we ask the following questions: (a) *Which of your personal information on Facebook do you think is available to the public?* (b) *What information do you think advertisers on Facebook already know about you?*

3.5 Discussion

In the experiment, we informed participants that the randomized price drawn from the computer was between $0 and $95. If a consumer's actual valuation is between this range, she has the incentive to report truthfully. If her valuation is above $95, reporting any value above $95 is optimal for her. In essence, it means that the reported values above $95 should be considered as "partially truthful": they truthfully reveal the fact that the underlying valuation is greater than a hundred, but are otherwise a stated preference.

Although many papers have used stated preference to measure privacy values (see Coyle & Manley 2022 for a review), debates abound over whether stated preferences are truthful (Spiekermann et al. 2001, Singleton & Harper 2002). One argument is that the gap between stated and revealed preferences can be alleviated once the context is controlled for (Prince & Wallsten 2020). Nevertheless, we adopt a variety of strategies in the reduced-form and structural model analysis to emphasize the valuations within the incentive-compatible range. The reduced-form analysis will focus on log valuation as our preferred specification. In the structural model estimation, we will allow truncation at different finite points above $95 and see whether the results are sensitive to different specifications.

4 Data and Reduced-Form Evidence

We recruited a total of 5,028 participants: 2,010 from Facebook during February 11-27, and 3,018 from Prolific during March 7-10, both in 2022. Table 1 provides summary statistics about our sample. Compared with the US representative sample,[4] ours includes more females and are overall better educated; the age, income, and percentage white distributions are similar to the national average. Compared to those recruited on Prolific, the Facebook ad samples include more females, are older, wealthier, better educated, more likely to be minorities, and spend more time on Facebook; they also click on ads and shop on Facebook more often. Table B.2 shows that our samples are balanced across the six treatment conditions.

Consumers' valuations for their data are heavily skewed to the right: in fact, 18.3% of consumers give their valuations at infinity across treatment conditions. The infinity values create challenges in reporting summary statistics and reduced-form analysis. We adopt several strategies

[4]https://www.census.gov/quickfacts/fact/table/US/PST045221; https://www.census.gov/library/visualizations/2022/comm/aging-nation-median-age.html.

Table 1: Summary Statistics of Sample Characteristics

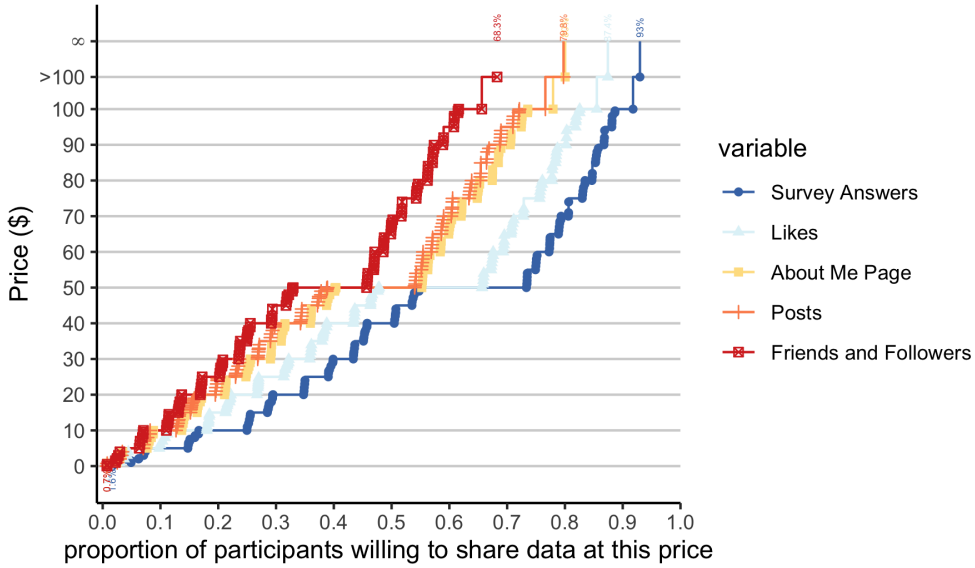
	Overall		Facebook		Prolific	
	Mean	SD	Mean	SD	Mean	SD
Number of participants						
N	5028		2010		3018	
Race (proportion)						
White	0.8	0.4	0.73	0.44	0.85	0.35
Black	0.07	0.25	0.08	0.27	0.06	0.23
Asian	0.12	0.32	0.16	0.37	0.08	0.28
Other	0.05	0.05	0.06	0.06	0.04	0.04
Gender (proportion)						
Female	0.59	0.49	0.76	0.43	0.48	0.5
Age						
Age	38.7	12.8	41	13.2	37.1	12.2
Income ($)						
Income	69238.9	48674.6	79954.9	51628.7	62102	45224.3
Education (proportion)						
College +	0.64	0.48	0.73	0.44	0.58	0.49
Some college	0.18	0.39	0.16	0.37	0.2	0.4
Less than college	0.17	0.38	0.11	0.31	0.22	0.41
Facebook and Internet Usage						
Average time spent on FB (h)	1.41	1.38	1.97	1.37	1.03	1.25
FB membership duration (y)	5.76	0.88	5.75	0.87	5.76	0.89
Average time spent on internet (h)	3.52	1.84	3.5	1.82	3.54	1.86
Active user (proportion)	0.31	0.46	0.47	0.5	0.21	0.41
Purchase from FB or Instagram (times/mo)	0.41	0.78	0.63	0.95	0.27	0.6
Click ads on FB or Instagram (times/mo)	1.66	1.81	2.47	1.92	1.12	1.51

to account for this challenge. In the model-free evidence, we report most results in the form of data supply curves, with the percentage of consumers with infinite valuations at the top of each curve. Since the data supply curves are essentially cumulative distribution functions of the valuation, they transparently visualize the distribution of consumer valuations and the impact of choice frames on different parts of the curve. In the reduced-form models, we use log valuation as the outcome variable in our main specification and indicate where we top-code the data. Focusing on log valuations makes the results robust to the impacts of extremely high values and our top-code choices. We note that the long-tail pattern in privacy valuations is universal (e.g., see Collis et al. 2020 and Lin 2022) and is not unique to our setting.

4.1 Valuations across Consumer Data

Our first result shows that consumers have systematically different valuations for different types of personal data (Figure 3). They value *friends and followers* data the most; next comes previous *posts* and *about me* information; *survey answers* is valued the least. When top-coded at \$119,[5] the difference between mean valuations across *friends* and *posts* is \$11.47, while the mean value difference between *friends* and *survey answers* is \$34.86.

Figure 3: Data Supply Curves by Data Type



The "coherent arbitrariness" theory (Ariely et al. 2003) argues that people can have coherent differences in valuation after picking an arbitrary starting point. To see if this hypothesis holds in our setting, we leverage the fact that the order of personal variables is also randomized across consumers, and examine the valuations using only the first question each consumer encounters. The differences in valuation across data persist with a similar magnitude, even when we focus on only the first valuation (see Appendix C.1). In other words, consumers' valuation for privacy is not arbitrary, but coherent both within and across people.

4.2 The Effects of Choice Architecture on Data Supply Curves

Despite the coherence, consumers' privacy valuations are prone to the influence of choice architecture. Table 2 summarizes the magnitude of average treatment effects using a set of Tobit regressions, with the valuations top-coded at \$119 as the outcome variable. An opt-out frame decreases the average valuation by 6.8%, while opt-in increases the valuation by 7.6% compared

[5]Why \$119? One of the authors insists on choosing a number that ends at 9 to represent the power of left-digit bias.

with active choice. The influence of a price anchor is more substantial. Consumer valuations for data decrease by 51.1% on average when priced in the low-price as opposed to the high-price condition.

Table 2: Average Treatment Effects and Valuation across Data: Tobit Regressions

	WTA	log(WTA)	WTA	log(WTA)
Intercept	67.116***	4.028***	49.281***	3.578***
	(1.014)	(0.027)	(1.030)	(0.029)
Price Anchor = Low	-14.717***	-0.510***	-14.771***	-0.511***
	(1.103)	(0.029)	(1.102)	(0.029)
Default = Active	2.936*	0.068+	2.930*	0.068+
	(1.324)	(0.036)	(1.323)	(0.036)
Default = Opt-in	5.948***	0.144***	5.934***	0.144***
	(1.333)	(0.036)	(1.332)	(0.036)
Likes			9.874***	0.286***
			(0.555)	(0.015)
About Me			21.248***	0.544***
			(0.646)	(0.016)
Posts			23.389***	0.597***
			(0.670)	(0.017)
Friends and Followers			34.861***	0.832***
			(0.791)	(0.019)
Num. Obs.	25140	25140	25140	25140

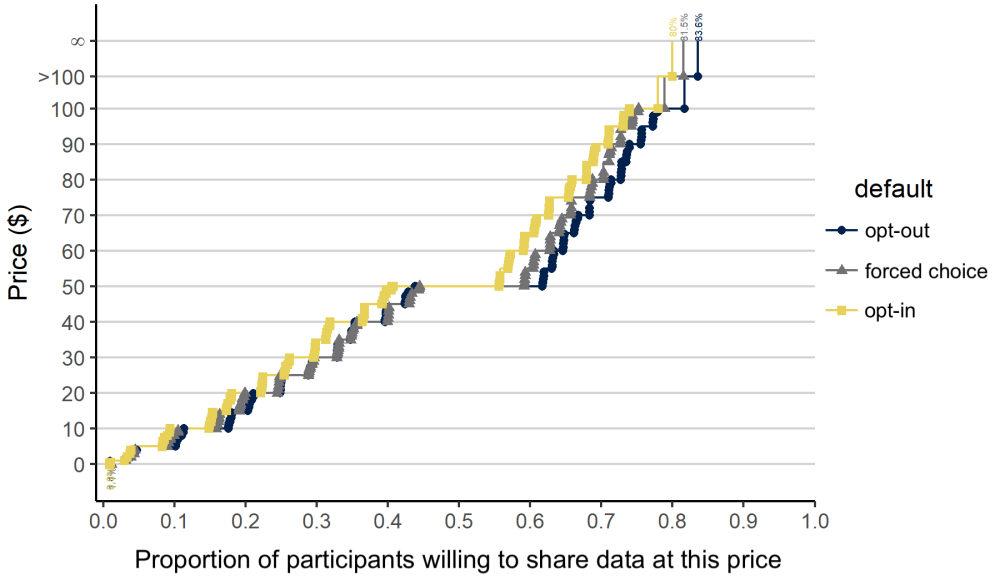
+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001. The outcomes are either valuations or log valuations, all top-coded at $119; standard errors are clustered at the participant level.

Figure 4 further compares the data supply curve across all six treatment combinations. The default conditions shift the supply curve uniformly. In contrast, the anchor price distorts different regions of the supply curve in different ways. In particular, the gap between supply curves is the largest in the middle region, due to valuations bunching towards the endpoints of the price range. Paradoxically, a low-price anchor also triggers more consumers to report extremely high (greater than $100) and infinite values. In the low-price condition, the average percentage of consumers reporting infinity is 21.2%, compared to 15.5% in the high-price condition. As a result, the supply curves from the two treatments intersect around the $95 price point, demonstrating the non-monotonicity of the price anchor effect.

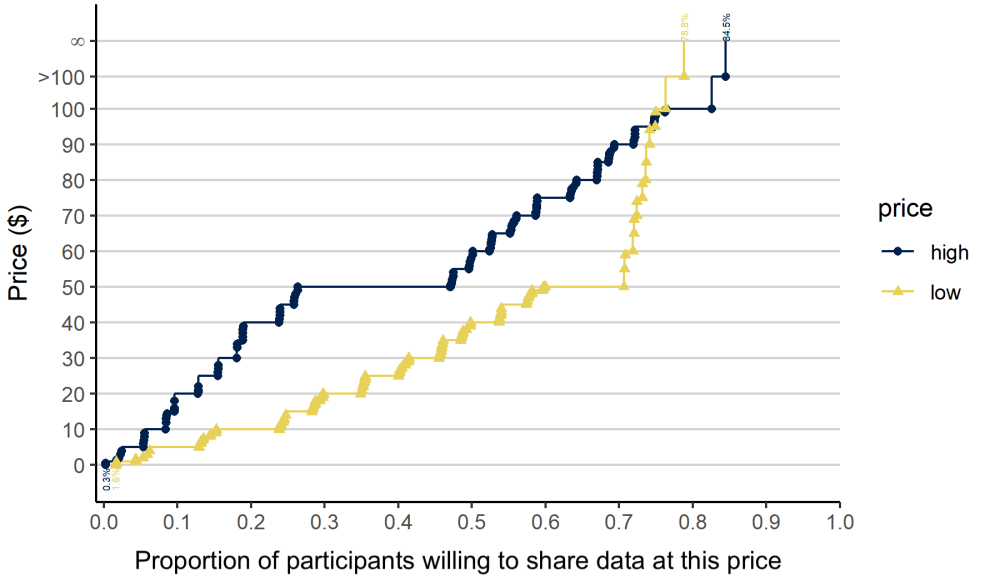
Appendix Table C.1 shows versions of Regression 4 with different top codes for the outcome. As the top code becomes less stringent, the coefficient representing the price anchor effect decreases in magnitude while all other coefficients increase. This pattern is consistent with the fact that price anchor creates "backlash" in a subset of consumers.

Figure 4: Data Supply Curves by Treatment Condition

(a) Default Option



(b) Price Range



4.3 Heterogeneous Effects of Choice Architecture

Although comparing data supply curves under different choice frames is informative, such a comparison does not tell us additional distortions caused by choice architecture due to misallocation. To capture such misallocation, we must go beyond the data supply curve comparison. Instead, we calculate each individual's heterogeneous responses to frames and relate them with their would-be valuation in a frame-neutral condition.

Without constructing a structural model, we first characterize heterogeneous frame effects using causal forest models (Athey et al. 2019). These models allow us to efficiently yet flexibly capture the treatment effects across consumer subgroups. We adopt two specifications. Our preferred specification is a survival forest, which takes care of censored values (i.e., our top codes for infinite values) in a Tobit-model fashion. However, survival forests only compare two treatments at a time; as such, it does not capture interaction effects between default and price anchor treatments. Our second specification is a multi-arm causal forest, which can capture the interaction effects between treatments but does not account for the censored values. The heterogeneity patterns in the two models are qualitatively similar, though the results from the multi-arm forests are smaller in magnitude due to censoring. In what follows, we use results from the survival forest when illustrating the main treatment effects, and results from the multi-arm forest when capturing the interaction effects is beneficial. We compare the estimation results from the two models in Appendix C.

Without a frame-neutral choice environment, it is hard to know what counts as a "frame-neutral" valuation. In the setting of forest models, we temporarily assume that the valuation marginalized over the six treatments is frame neutral. This is a simplifying assumption: we recognize the possibility that the neutral valuation may be closer to one of our frames than the others, rather than located exactly in the middle. We will replace this assumption in Section 5 with micro-founded economic structures. The results below use log valuation as the outcome variable, where the valuations are top-coded at $1,000 before taking the log. The covariates in the model include the demographics, Facebook and general internet usage, and consumer beliefs about what data are already available to Facebook as well as the public.

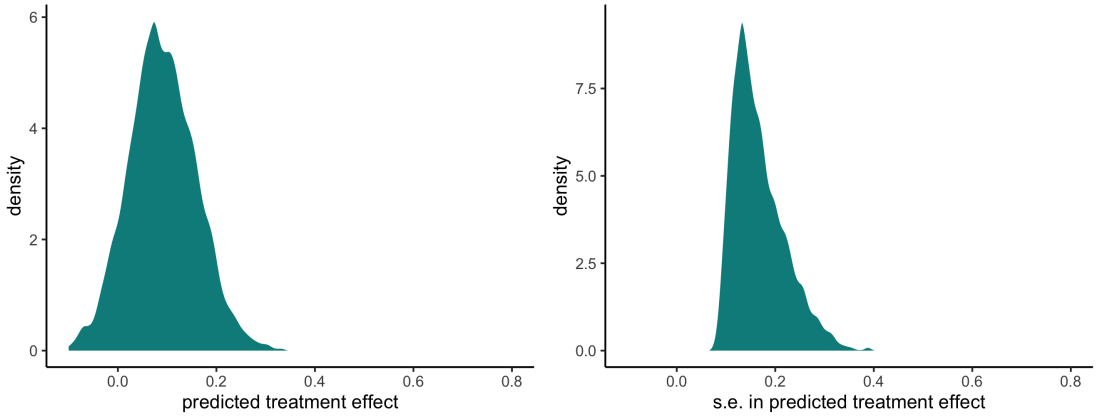
Figure 5 shows the distribution of the heterogeneous treatment effects for the two choice frames, with the clustered standard error distribution on the right. The average treatment effects for both frames are significantly above zero. Consistent with the raw data, the price anchor has a larger overall effect than the default.

Figure 6 shows the correlation between frame effects and the frame-neutral log valuation. The effect of default does not vary systematically as the valuation increases (correlation is -0.04). On the other hand, the effect of the choice anchor is strongly and negatively correlated with the neutral valuation (correlation is -0.46): consumers who have lower frame-neutral valuations are also more likely swayed by a price anchor. This pattern is consistent with Collis et al. (2020): They find that consumers who see an informative signal on the market value for data are more likely to revise their valuation when their prior is low, but are less likely to revise their prior the other way around.

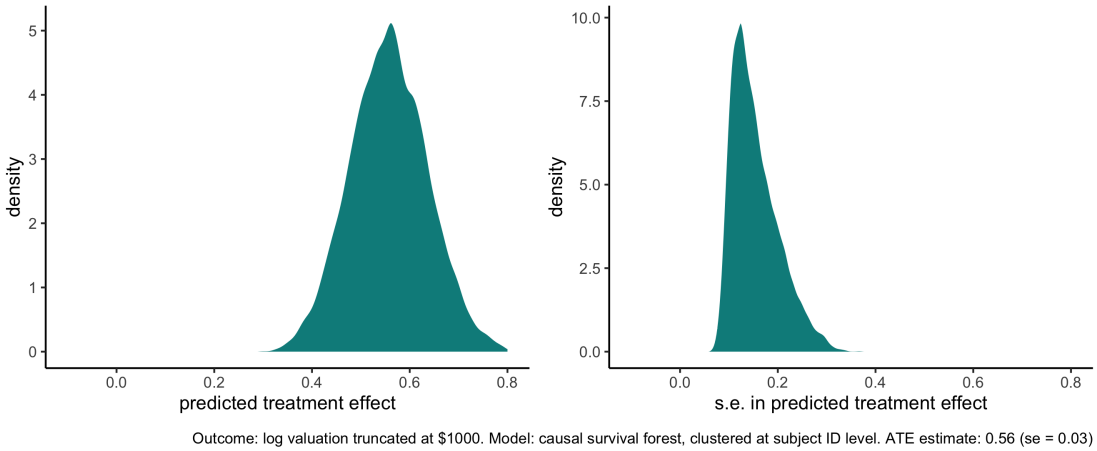
To further explore the overlap between consumers who experience larger frame effects and those with low privacy valuations, we project the heterogeneous treatment effects and the neutral valuation to the consumer attributes using a set of linear models. Tables 3 and 4 show the respective results. Consumers more easily influenced by choice frames are overall younger, poorer, less educated, and spend less time on Facebook; interestingly, they are more likely to click on ads while

Figure 5: Heterogeneous Treatment Effect Estimates and Standard Errors: Survival Forests

(a) Default Option



(b) Price Range

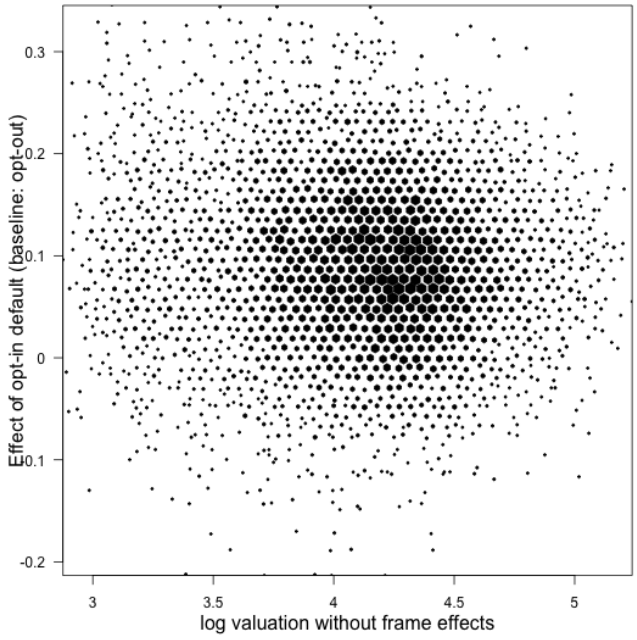


using Facebook. In comparison, consumers who value their personal data more are older, richer, better educated, more likely to be female and Asian, and spend less time on the Internet in general; they are also less likely to click ads on Facebook. Overall, we see that many attributes that predict high choice frame effects also predict lower valuations for personal data.

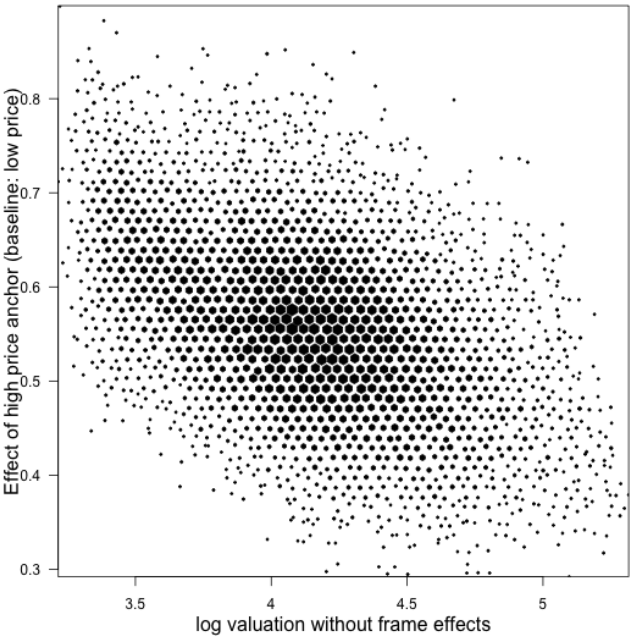
As a final note, we separately compare the neutral log valuation and choice frame effect sizes from our two sample sources, which is shown in Figure 7. The Facebook sample has a higher neutral valuation overall and is less influenced by frames. Although we do not know if the Facebook sample is more representative of the total population than the Prolific panel, our results suggest that studies using multiple sample sources have merit in assessing the robustness of the effects they aim to study.

Figure 6: Correlation Between Choice Frame Effects and Frame-Neutral Valuation

(a) Default Option



(b) Price Anchor



Note: In the figures above, treatment effects are estimated using survival forests, with the outcome as log valuation truncated at $1000. Frame-neutral valuations are generated alongside the model as the predicted outcome marginalized across the choice frames. The size of the points represents the number of observations in that region.

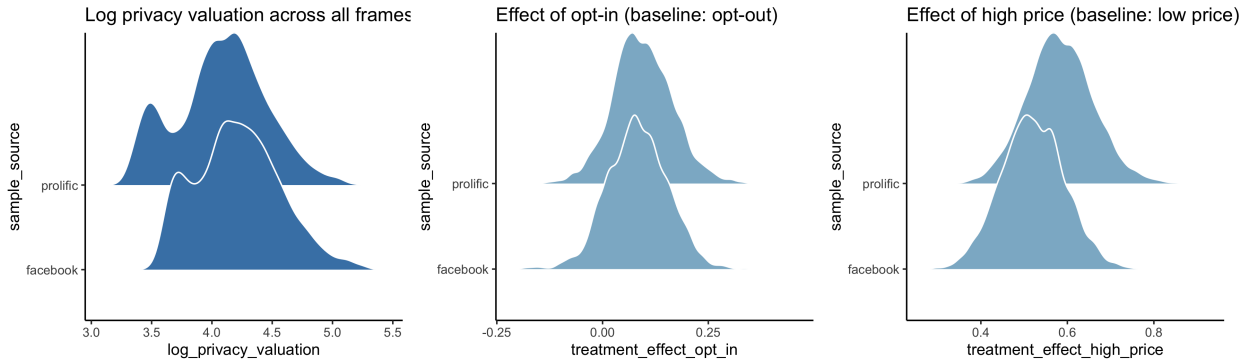
Table 3: Linear Projection of Maximal Frame Effect onto Consumer Characteristics

Data	Posts	About Me	Survey Answers	Likes	Friends	**All Data**
Intercept	0.602***	0.548***	0.877***	0.713***	0.385***	**0.632*** **
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	**(0.001)**
Age	-0.011***	-0.019***	-0.011***	-0.013***	-0.018***	**-0.027*** **
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	**(0.002)**
Income	-0.010***	-0.015***	-0.023***	-0.010***	-0.020***	**-0.026*** **
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	**(0.002)**
Education	-0.012***	-0.010***	-0.011***	-0.009***	-0.020***	**-0.021*** **
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	**(0.002)**
Female	0.000	0.006*	0.018***	-0.002	-0.027***	**-0.008** **
	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)	**(0.003)**
Race_black	0.008*	-0.001	0.004	-0.004	0.016***	**0.017** **
	(0.004)	(0.005)	(0.004)	(0.003)	(0.004)	**(0.006)**
Race_asian	-0.008**	-0.009*	-0.011***	0.002	0.001	**0.000**
	(0.003)	(0.004)	(0.003)	(0.003)	(0.003)	**(0.004)**
Time_on_FB	0.000	-0.002+	-0.003**	-0.008***	-0.012***	**-0.011*** **
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	**(0.002)**
Time_on_Internet	0.004***	-0.001	-0.009***	0.001	0.003*	**-0.005** **
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	**(0.002)**
Shop_on_FB	0.001	-0.002	0.002+	0.000	0.001	**0.001**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	**(0.001)**
Click_ad_on_FB	0.014***	0.008***	-0.003**	0.011***	0.002+	**0.004** **
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	**(0.001)**
Num. Obs.	5027	5027	5027	5027	5027	**25135**
R^2	0.160	0.111	0.188	0.143	0.275	**0.203**

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001.

Note: In each regression above, the outcome is the predicted *maximal treatment effect* for the variable specified in the column, represented as the difference between the highest and lowest valuation across the choice frames. Treatment effects are estimated using multi-arm causal forests, with log valuation censored at $1000 as the outcome.

Figure 7: Neutral Log Valuation and Choice Frame Effect Distributions by Sample Source



Note: In these figures, the treatment effects are estimated using survival forests, with the outcome as log valuation truncated at $1000.

Table 4: Linear Projection of Frame-Neutral Valuation onto Consumer Characteristics

Data	Posts	About Me	Survey Answers	Likes	Friends	**All Data**
Intercept	4.307***	4.242***	3.515***	3.887***	4.679***	**4.128*** **
	(0.004)	(0.003)	(0.003)	(0.002)	(0.004)	**(0.003)**
Age	0.050***	0.101***	0.081***	0.071***	0.255***	**0.115*** **
	(0.004)	(0.004)	(0.003)	(0.003)	(0.004)	**(0.003)**
Income	0.024***	0.000	0.009**	0.017***	0.003	**0.011*** **
	(0.004)	(0.004)	(0.003)	(0.003)	(0.004)	**(0.003)**
Education	0.088***	0.032***	0.019***	0.047***	0.057***	**0.050*** **
	(0.004)	(0.004)	(0.003)	(0.003)	(0.004)	**(0.003)**
Female	0.170***	0.119***	0.064***	0.162***	0.103***	**0.125*** **
	(0.008)	(0.007)	(0.005)	(0.005)	(0.008)	**(0.005)**
Race_black	-0.059***	0.041**	0.116***	0.008	-0.031*	**0.012**
	(0.015)	(0.014)	(0.011)	(0.010)	(0.015)	**(0.011)**
Race_asian	0.000	0.043***	0.063***	0.026***	0.046***	**0.040*** **
	(0.012)	(0.011)	(0.008)	(0.008)	(0.012)	**(0.008)**
Time_on_FB	-0.004	-0.006	0.044***	0.012***	-0.014***	**0.006***
	(0.004)	(0.004)	(0.003)	(0.003)	(0.004)	**(0.003)**
Time_on_Internet	-0.013***	-0.014***	-0.017***	-0.016***	-0.012**	**-0.014*** **
	(0.004)	(0.004)	(0.003)	(0.003)	(0.004)	**(0.003)**
Shop_on_FB	-0.003	0.012**	0.033***	0.010***	-0.005	**0.008***
	(0.004)	(0.004)	(0.003)	(0.003)	(0.004)	**(0.003)**
Click_ad_on_FB	-0.019***	-0.027***	0.004	-0.016***	-0.027***	**-0.017*** **
	(0.004)	(0.004)	(0.003)	(0.003)	(0.004)	**(0.003)**
Num. Obs.	5027	5027	5027	5027	5027	**25135**
R^2	0.241	0.208	0.304	0.362	0.523	**0.112**

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001.
Note: In each regression above, the outcome variable is a consumer's frame-neutral valuation for the variable specified in the column, represented as the predicted log valuation marginalized across the choice frame treatments. The frame-neutral valuations are predicted alongside the multi-arm causal forests, with log valuation censored at $1000 as the outcome.

5 The Model

Despite the rich insights from the reduced-form evidence, we need a micro-founded model for several reasons. First, a structural model allows us to accurately back out the choice frame effects. As we have seen in the data, the price anchor distorts valuation in a way that does not conform to a simple linear model formula. The causal forest model is unable to pick up the bunching effect caused by the price anchor, despite its ability to partition variables and allow for flexible interaction effects. Yet, our structural model allows us to use only three parameters to parsimoniously characterize the complex effects of price anchor, including bunching and backlash.

Second and more importantly, our model allows us to accurately represent consumers' would-be valuations if neither default nor price anchors were present. A price anchor is always present in empirical settings where we want to elicit privacy valuations under a revealed preference

framework (either take-it-or-leave-it or BDM/MPL), [6] yet the definition of a "neutral" price anchor is unclear. Our model allows us to characterize and separate the effect of the price anchor. In other words, we can tease out the effect of price anchor on privacy valuations for any anchoring point, without the need to know which anchor is neutral.

For the purpose of welfare calculation, we need to estimate not only the average frame effects and valuations but also their heterogeneity and how the two correlate. To estimate heterogeneity under a structural model framework, we use recent advances proposed by Farrell et al. (2020), who combine deep neural nets with automatic influence function adjustments to efficiently learn the heterogeneity of causal parameters.

5.1 Setup

We distinguish between the following valuations for personal data:

- v_{ik}^0: consumer i's prior valuation for data k before seeing any choice frame.
- $\widetilde{v_{ik}}$: valuation reported by the consumer, which is a function of the choice frames. This value is observable to us in the free-text entry.

Choice frames shift the reported valuation $\widetilde{v_{ik}}$ away from the prior valuation v_{ik}^0. The frames in our experiment include the following:

- $d_i^{in} = 1$ if the consumer is in the opt-in frame; $d_i^{in} = 0$ otherwise. d_i^{out} is similarly defined for the opt-out frame. $d_i^{in} = d_i^{out} = 0$ implies the active-choice condition.
- $[\underline{P_i}, \overline{P_i}]$ is the range of the offer price.

5.2 Frame Effects

The default distorts the reported valuation: instead of v_{ik}^0, now consumers' valuation is $v_{ik}^0 + m_1 \cdot d^{in} + m_0 \cdot d^{out} \equiv v_{ik}^d$. We choose the additive rather than multiplicative specification to account for the fact that the default effect does not seem to scale with the valuation in our non-parametric causal forest model.

The offer price serves as an anchor, but only when the consumer thinks the price range is reasonable. The consumer reports $\widetilde{v_{ik}} = \infty$ when the price is too low compared to the shifted valuation, i.e. $\overline{P_i} < v_{ik}^d - \delta$. Since this condition is more likely to be met when the offer price is low,

[6] One may think that a simple open-text question asking consumers how much they want to be paid to share their data is a simple and effective way to remove the effect of price anchoring. However, this is rarely done in practice, as consumers may not deliberate before giving their answers if only open-text questions are present. In our experiment, we include standard MPL questions before our open texts precisely to encourage deliberation.

consumers are more likely to report infinite values (i.e., "backlash") in the low price condition, which is consistent with our data.

If the "backlash" condition is not met, the offer price draws the reported valuation closer to the range. We assume that consumers choose the reported value to minimize the weighted sum of its distance from prior and from the price anchor. To accommodate the presence of bunching while also allowing for alternative anchoring rules, we specify an elastic-net flavored formula of anchoring, where the distance to the anchor is a combination of L-1 and L-2 norms:

$$\begin{aligned}\widetilde{v_{ik}} = \arg\min_{v} &\left(v - v_{ik}^{d}\right)^2 \\ &+2\lambda_1 \cdot \left[(\underline{P} - v) \cdot \mathbb{1}(v < \underline{P}) + (v - \overline{P}) \cdot \mathbb{1}(v > \overline{P})\right] \\ &+\lambda_2 \cdot \left[(v - \underline{P})^2 \cdot \mathbb{1}(v < \underline{P}) + (v - \overline{P})^2 \cdot \mathbb{1}(v > \overline{P})\right],\end{aligned} \tag{2}$$

Combining the solution of the minimization problem above and the backlash condition yields the following valuation reporting rule:

$$\widetilde{v_{ik}} = \begin{cases} \frac{v_{ik}^d + \lambda_2 \underline{P} + \lambda_1}{1+\lambda_2}, & \text{if } v_{ik}^d < \underline{P} - \lambda_1, \\ \underline{P}, & \text{if } \underline{P} - \lambda_1 \le v_{ik}^d < \underline{P}, \\ v_{ik}^d, & \text{if } \underline{P} \le v_{ik}^d < \overline{P}, \\ \overline{P}, & \text{if } \overline{P} \le v_{ik}^d < \overline{P} + \lambda_1, \\ \frac{v_{ik}^d + \lambda_2 \overline{P} - \lambda_1}{1+\lambda_2}, & \text{if } \overline{P} + \lambda_1 \le v_{ik}^d < \overline{P} + \delta, \\ \infty, & \text{if } v_{ik}^d \ge \overline{P} + \delta. \end{cases} \tag{3}$$

Having $\lambda_1 > 0$ explains bunching at either end of the price range, while $\lambda_2 > 0$ can rationalize any additional slope shifting effect caused by the price anchor. Later, we will allow both λ's to be functions of consumer characteristics. Therefore, the model can accommodate cases where some consumers only have $\lambda_1 > 0$ while others only have $\lambda_2 > 0$. We denote $\theta = \{m_1, m_0, \lambda_1, \lambda_2, \delta\}$ as the set of frame effect parameters.

5.3 Heterogeneity

Since there is only one observation for each (i, k), we let $v_{ik}^0 = v_{0k} + v(X_{ik}) + e_{ik}$, where X_{ik} is a set of consumer characteristics, some of which are also data-specific. X_{ik} includes the following variables: *demographics, belief about data availability, length of behavioral data, Internet and social media usage, and sample source*. We also allow frame effects θ to be functions of consumer-data characteristics in the form of $\theta(X_{ik})$. We let $e_{ik} \sim TIEV(0, \sigma_e)$ to represent the unobserved heterogeneity in prior valuation. The type-I extreme value specification accommodates the right tails in the privacy value distribution in the data without restricting all values to be strictly positive.

The setup above implies that v_{ik}^d is also TIEV distributed, with location parameter $u \equiv v_{0k} + v(X_{ik}) + e_{ik} + m_1 \cdot d^{in} + m_0 \cdot d^{out}$ and scale parameter σ_e. To simplify the notation, denote the cumulative distribution function of v_{ik}^d as $G_d(.)$ and its density function as $g_x(.)$. The log-likelihood takes the following form (we omit the function component of each choice frame parameter for notational simplicity):

$$\begin{aligned} LL = \sum_{i,k} & \mathbb{1}(\widetilde{v_{ik}} < \underline{P}) \cdot g_d((1+\lambda_2) \cdot \widetilde{v_{ik}} - \lambda_2 \cdot \underline{P} - \lambda_1) \\ & + \mathbb{1}(\widetilde{v_{ik}} = \underline{P}) \cdot \log\left[G_d(\underline{P}) - G_d(\underline{P} - \lambda_1)\right] \\ & + \mathbb{1}(\widetilde{v_{ik}} > \underline{P}) \cdot \mathbb{1}(\widetilde{v_{ik}} < \overline{P}) \cdot g_d(\widetilde{v_{ik}}) \\ & + \mathbb{1}(\widetilde{v_{ik}} = \overline{P}) \cdot \log\left[G_d(\overline{P} + \lambda_1) - G_d(\overline{P})\right] \\ & + \mathbb{1}(\widetilde{v_{ik}} > \overline{P}) \cdot \mathbb{1}(\widetilde{v_{ik}} < \infty) \cdot g_d((1+\lambda_2) \cdot \widetilde{v_{ik}} - \lambda_2 \cdot \overline{P} + \lambda_1) \\ & + \mathbb{1}(\widetilde{v_{ik}} = \infty) \cdot \log\left[1 - G_d(\underline{P} + \delta)\right]. \end{aligned} \tag{4}$$

5.4 Identification

In the model above, the key parameters include consumers' prior privacy valuation v^0 and the set of choice frame parameters θ. Since the reported valuation $\widetilde{v}$ is jointly determined by v^0 and θ, we identify consumer prior by separating the choice frame effects. Below, we illustrate how the main choice frame effects are identified.

The default effect parameters m_1 and m_0 are identified by the difference in valuations between default and active-choice treatments. By having a active-choice condition, we can allow opt-in and opt-out to have different effects rather than assuming the effects are symmetric. The finite valuations across price anchor treatments identify λ_1 and λ_2: In particular, λ_1 is pinned down by the degree of bunching at either end of the price range, while λ_2 explains the additional slope differences in data supply curves caused by the price anchor treatment. By identifying the choice frame parameters for people reporting finite $\widetilde{v}$, we can identify v^0 for this group. We let v^0 be a function of consumer characteristics to extrapolate v^0 from the finite-valuation sample to the infinite-valuation sample. Once v^0 is pinned down for all participants, we back out δ using the variation of reported infinite values.

6 Conclusion and Future Analysis

Throughout the paper, we have shown that choice architecture can substantially distort consumers' privacy valuations and thus the position of a data supply curve. The choice frames we examine—default and price anchor—shift the average consumer valuation for their Facebook data by 14.8% and 51.1%, respectively. Moreover, we find the susceptivity to frame influence is neg-

atively correlated with their valuation for data absent frames. Younger, less wealthy, and less educated consumers tend to respond more strongly to changes in the frame; they also value their personal data less when the frame effects are taken out.

Our results have several implications for the functioning of a data market. First, the negative correlation between choice frame effects and privacy valuations implies that the distortive effects of choice architecture are larger at the lower end of a data supply curve. As a result, the total welfare loss caused by choice architecture is larger when the quantity of data traded is smaller in equilibrium; that is, if the market for data is not saturated. Second, when a firm chooses a choice frame that naively maximizes the supply at each price point, the type of consumers who already value their data less will lower their valuations even more, whereas those who initially value their privacy strongly are less likely to be swayed. Though not conclusive, this correlational pattern suggests that a volume-maximizing frame may exacerbate the statistical bias in data collected by the firm.[7] Such bias can decrease the returns to data, counteracting the benefits of having more data to the firm, and bringing additional inefficiency from a collective welfare standpoint.

In future analysis, we will formally examine the bias in collected data and how it changes with different frame implementations. The structural model will allow us to estimate the frame-free valuation cleanly, overcoming the shortcomings of our reduced-form analysis. In the counterfactual analysis, we will study the regime where the firm can optimize the choice frames to maximize the volume of data collected at any given price point. Then, given a grid of market prices for data, we examine which consumers share their data and how they compare with the total participant pool. For example, we can compare the distribution of observed attributes between the shared and the full sample using Kullback-Leibler divergence, which measures the difference between the two distributions. We can then correlate the KL divergence and the volume of shared data across different frame implementations to examine the bias-variance tradeoff given by each frame choice. We will then discuss the impact of the firm's frame choice on consumer welfare and social surplus.

Our results contribute to the broad discussions about the efficiency of the data market. Several features unique to consumer data can lead to market failures, such as externality (Bergemann et al. 2020), incomplete information (Jin 2018), and non-rivalry (Jones & Tonetti 2020). Here, we have shown another feature that contributes to market inefficiency: a data market that allows consumers to make their own choices can often create bias in data collected and decrease the returns to data. In the context of choice architecture evaluation, we show how companies' efforts to maximize data collection can make this bias even more severe. We believe the bias aspect is worth emphasizing, as it represents a novel channel in the data market that impedes its functioning.

[7] It is inconclusive because the choice frames may differentially make vulnerable consumers even more likely to share their data without perturbing their value ranking, in which case the frame would not have an impact on the bias in collected data.

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?', *The 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.

Andersen, S., Harrison, G. W., Lau, M. I. & Rutström, E. E. (2006), 'Elicitation using multiple price list formats', *Experimental Economics* **9**(4), 383–405.

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

Arrieta-Ibarra, I., Goff, L., Jiménez-Hernández, D., Lanier, J. & Weyl, E. G. (2018), Should we treat data as labor? moving beyond 'free', *in* 'American Economic Association Papers and Proceedings', Vol. 108, pp. 38–42.

Athey, S., Catalini, C. & Tucker, C. (2017), 'The digital privacy paradox: Small money, small costs, small talk', *NBER Working Paper No. 23488* .

Athey, S., Tibshirani, J. & Wager, S. (2019), 'Generalized random forests', *The Annals of Statistics* **47**(2), 1148–1178.

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. M., DeGroot, M. H. & Marschak, J. (1964), 'Measuring utility by a single-response sequential method', *Behavioral science* **9**(3), 226–232.

Bergemann, D., Bonatti, A. & Gan, T. (2020), 'The economics of social data', *Cowles Foundation Working Paper* .

Carroll, G. D., Choi, J. J., Laibson, D., Madrian, B. C. & Metrick, A. (2009), 'Optimal defaults and active decisions', *The quarterly journal of economics* **124**(4), 1639–1674.

Collis, A., Moehring, A. & Sen, A. (2020), 'Economic value of data: Quantification using online experiments'.

Coyle, D. & Manley, A. (2022), 'What is the value of data? a review of empirical methods'.

Dizon-Ross, R. & Jayachandran, S. (2022), Improving willingness-to-pay elicitation by including a benchmark good, *in* 'AEA Papers and Proceedings', Vol. 112, pp. 551–55.

Farrell, M. H., Liang, T. & Misra, S. (2020), 'Deep learning for individual heterogeneity: an automatic inference framework', *arXiv preprint arXiv:2010.14694* .

Goldfarb, A. & Tucker, C. (2012), 'Shifts in privacy concerns', *American Economic Review* **102**(3), 349–53.

Heidhues, P. & Kőszegi, B. (2018), 'Behavioral industrial organization', *Handbook of Behavioral Economics: Applications and Foundations 1* **1**, 517–612.

Ichihashi, S. (2021), 'The economics of data externalities', *Journal of Economic Theory* **196**, 105316.

Jin, G. Z. (2018), 'Artificial intelligence and consumer privacy', *The Economics of Artificial Intelligence: An Agenda; Agrawal, A., Gans, J., Goldfarb, A., Eds* pp. 439–462.

Johnson, E. J., Bellman, S. & Lohse, G. L. (2002), 'Defaults, framing and privacy: Why opting in-opting out', *Marketing Letters* **13**(1), 5–15.

Jones, C. I. & Tonetti, C. (2020), 'Nonrivalry and the economics of data', *American Economic Review* **110**(9), 2819–58.

Kahneman, D., Knetsch, J. L. & Thaler, R. H. (1990), 'Experimental tests of the endowment effect and the coase theorem', *Journal of political Economy* **98**(6), 1325–1348.

Kahneman, Wakker & Sarin (1997), 'Back to bentham? explorations of experienced utility', *The quarterly journal of* .

Kormylo, C. & Adjerid, I. (2021), 'Reconsidering privacy choices: The impact of defaults, reversibility, and repetition', *Available at SSRN 3846620* .

Kummer, M. & Schulte, P. (2019), 'When private information settles the bill: Money and privacy in google's market for smartphone applications', *Management Science* .

Lin, T. (2022), 'Valuing intrinsic and instrumental preferences for privacy', *Marketing Science* .

Markovich, S. & Yehezkel, Y. (2021), "'for the public benefit": who should control our data?', *Available at SSRN 3945108* .

Martin, K. & Nissenbaum, H. (2016), 'Measuring privacy: an empirical test using context to expose confounding variables', *Colum. Sci. & Tech. L. Rev.* **18**, 176.

Miller, K. & Skiera, B. (2017), 'Economic damage of cookie lifetime restrictions', *Available at SSRN 3046254* .

Peltzman, S. (1981), 'The effects of ftc advertising regulation', *The Journal of Law and Economics* **24**(3), 403–448.

Prince, J. T. & Wallsten, S. (2020), 'How much is privacy worth around the world and across platforms?', *Journal of Economics & Management Strategy* .

Rafieian, O. & Yoganarasimhan, H. (2018), 'Targeting and privacy in mobile advertising'.

Rao, A. & Wang, E. (2017), 'Demand for "healthy" products: False claims and ftc regulation', *Journal of Marketing Research* **54**(6), 968–989.

Rossi, P. E., McCulloch, R. E. & Allenby, G. M. (1996), 'The value of purchase history data in target marketing', *Marketing Science* **15**(4), 321–340.

Singleton, S. M. & Harper, J. (2002), 'With a grain of salt: What consumer privacy surveys don't tell us', *Available at SSRN 299930* .

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', pp. 38–47.

Sun, T., Yuan, Z., Li, C., Zhang, K. & Xu, J. (2021), 'The value of personal data in internet commerce: A high-stake field experiment on data regulation policy', *Available at SSRN 3962157* .

Tang, H. (2019), 'The value of privacy: Evidence from online borrowers'.

Taubinsky, D. & Rees-Jones, A. (2018), 'Attention variation and welfare: theory and evidence from a tax salience experiment', *The Review of Economic Studies* **85**(4), 2462–2496.

Trusov, M., Ma, L. & Jamal, Z. (2016), 'Crumbs of the cookie: User profiling in customer-base analysis and behavioral targeting', *Marketing Science* **35**(3), 405–426.

Wernerfelt, N., Tuchman, A., Shapiro, B. & Moakler, R. (2022), 'Estimating the value of offsite data to advertisers on meta', *Available at SSRN* .

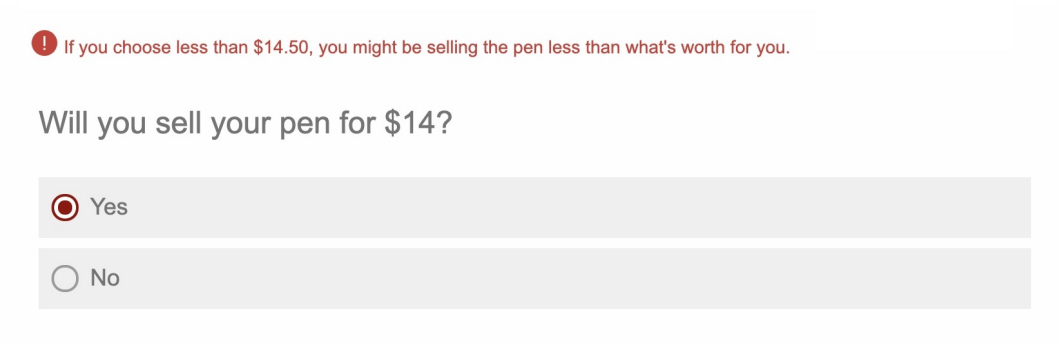
A Experiment Design Appendix

Figure A.1: Example Recruiting Ad

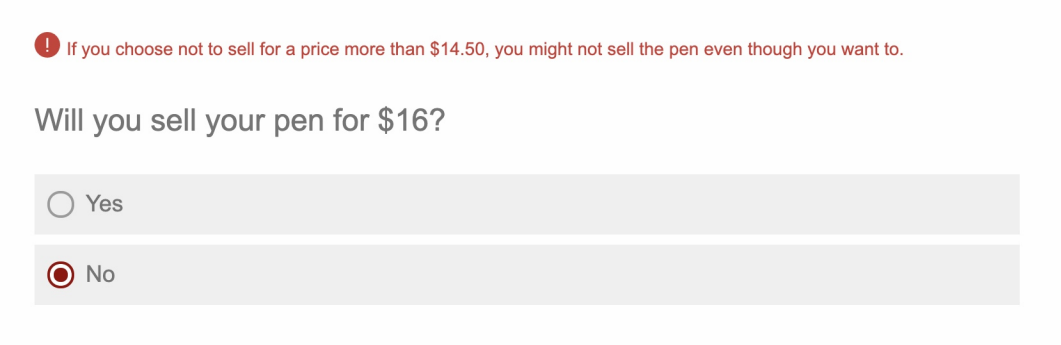


Figure A.2: Information Prompts in the MPL Practice Question

(a) Accepting an offer price too low



(b) Rejecting an offer price too high



B Data Appendix

Table B.1: Attrition throughout the Study

	Facebook	**Prolific**
Exposed to the ad	158453	-
Click on the ad	10135	-
Consent to participate	2348	3119
Complete all questions	2008	3018

Table B.2: Covariate Balance across Treatments

	Treament						T-test for equal mean
	low: opt-out	low: active	low: opt-in	high: opt-out	high: active	high: opt-in	p_value
Race							
White	0.827	0.792	0.813	0.795	0.793	0.809	0.398
Black	0.057	0.074	0.061	0.075	0.075	0.059	0.418
Asian	0.111	0.115	0.109	0.111	0.123	0.125	0.871
Other	0.037	0.047	0.049	0.057	0.052	0.05	0.558
Gender							
Is female	0.597	0.592	0.597	0.6	0.578	0.593	0.959
Age							
Age	39.318	37.915	38.35	38.803	39.097	38.534	0.245
Income ($)							
Income	66,508	67,676	72,021	70,248	72,875	66,077	0.011
Education							
College education and above	521	530	539	548	547	542	0.831
Some college	151	155	146	169	151	155	
Less than college	137	164	151	138	134	150	
Facebook Usage							
Avg time spent on FB	-0.054	0.064	0.003	-0.04	0.044	-0.019	0.102
Membership duration	0.002	-0.01	0.049	-0.014	-0.016	-0.011	0.772
Avg time spent on internet	0.002	0.048	-0.015	-0.021	0.002	-0.015	0.75
Purchase frequency/month	0.002	-0.004	0.021	-0.002	0.035	-0.051	0.602
Ad click frequency/month	-0.01	-0.025	0.008	-0.026	0.01	0.043	0.719
Use FB actively	0.302	0.313	0.301	0.329	0.332	0.294	0.448

C Empirical Results Appendix

Figure C.1: Data Supply Curves by Data Type (First-Round Answer Only)

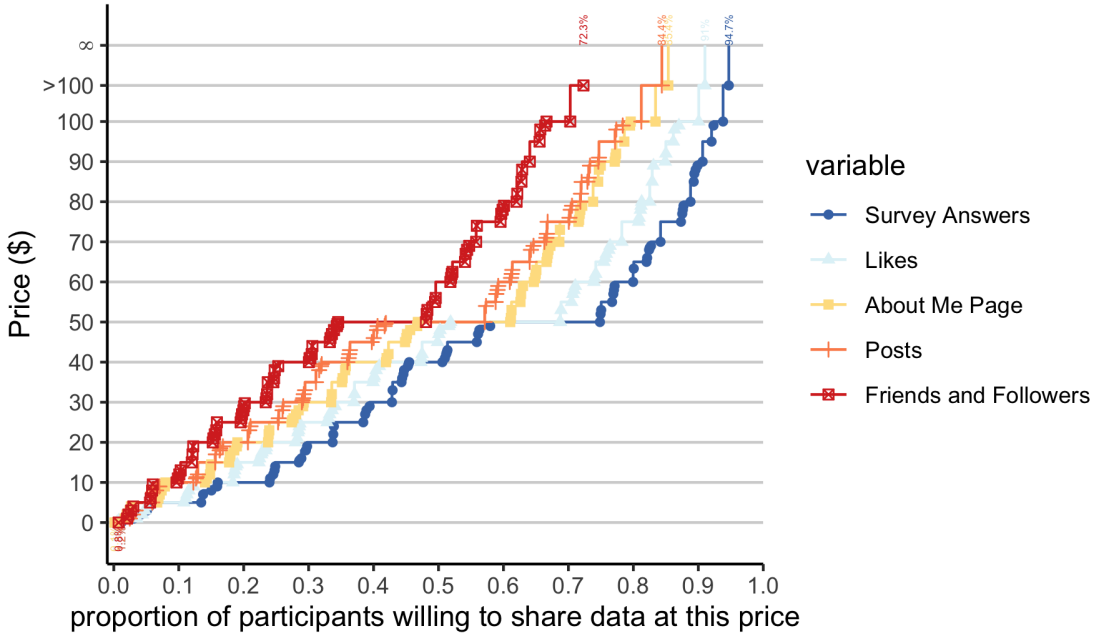


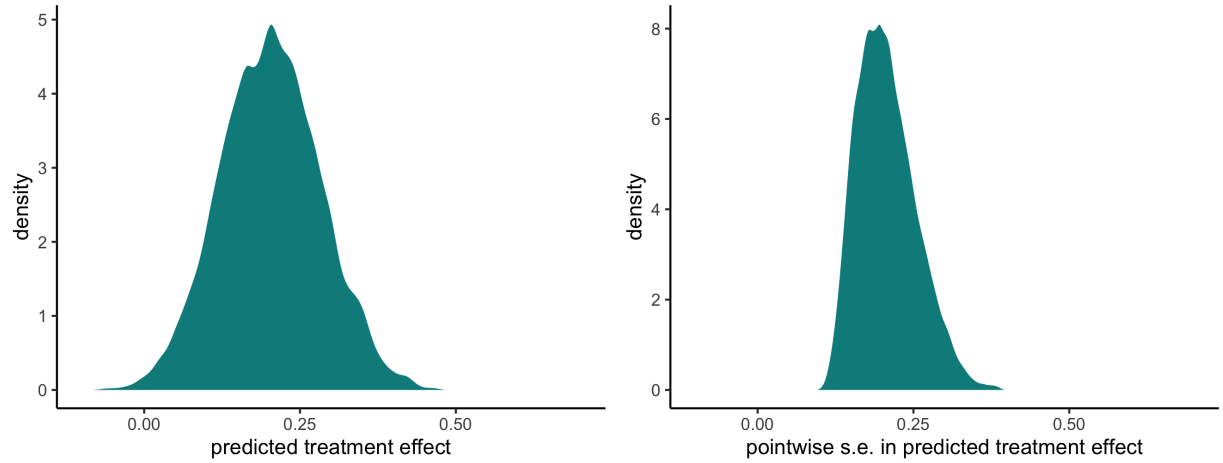
Table C.1: Average Treatment Effects and Data Valuation across Topcodes: Tobit Regressions

Top-code ($)	109	250	500	750	1000
Intercept	3.578***	3.597***	3.614***	3.625***	3.633***
	(0.029)	(0.033)	(0.037)	(0.040)	(0.042)
Price Range = Low	-0.511***	-0.462***	-0.417***	-0.392***	-0.374***
	(0.029)	(0.034)	(0.040)	(0.043)	(0.045)
Default = Active	0.068+	0.084*	0.102*	0.111*	0.118*
	(0.036)	(0.042)	(0.048)	(0.051)	(0.054)
Default = Opt-in	0.144***	0.169***	0.196***	0.210***	0.221***
	(0.036)	(0.042)	(0.048)	(0.052)	(0.055)
Likes	0.286***	0.338***	0.387***	0.413***	0.433***
	(0.015)	(0.017)	(0.020)	(0.022)	(0.023)
About Me	0.544***	0.654***	0.761***	0.822***	0.869***
	(0.016)	(0.020)	(0.023)	(0.026)	(0.027)
Posts	0.597***	0.716***	0.833***	0.895***	0.944***
	(0.017)	(0.021)	(0.024)	(0.027)	(0.028)
Friends and Followers	0.832***	1.043***	1.248***	1.365***	1.453***
	(0.019)	(0.024)	(0.029)	(0.032)	(0.034)
Num. Obs.	25140	25140	25140	25140	25140

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001. The outcomes are log valuations with top-codes indicated in the first row; standard errors are clustered at the participant level.

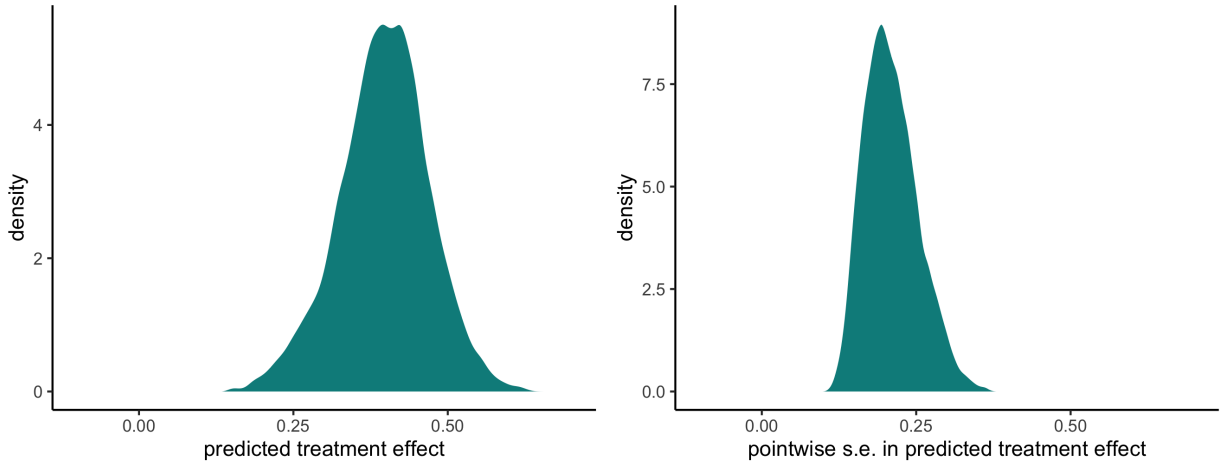
Figure C.2: Heterogeneous Treatment Effect Estimates and Standard Errors: Multi-Arm Causal Forests

(a) Default Option



Outcome: log valuation truncated at $1000. Model: multi-arm causal forest, clustered at subject ID level. ATE estimate: 0.19 (se = 0.07)

(b) Price Anchor



Outcome: log valuation truncated at $1000. Model: multi-arm causal forest, clustered at subject ID level. ATE estimate: 0.4 (se = 0.07)

Since the multi-arm forest does not account for the fact that the latent censored values are greater than the censoring point, the magnitudes in its estimates are potentially biased compared to the survival forest. Indeed, Figure C.2 shows that the default estimates from the multi-arm causal forest are slightly biased upwards while the price anchor effect is biased downwards.