

# Choice Architecture, Privacy Valuations, and Selection Bias in Consumer Data

Tesary Lin\*

Avner Strulov-Shlain<sup>†</sup>

April 21, 2023

*[PRELIMINARY; COMMENTS ARE WELCOME]*

## Abstract

How much do consumers' privacy valuations change under the influence of choice architecture? How does this influence affect the efficiency of data collection, by changing not only the quantity of data collected but also its representativeness? 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. An opt-out default decreases valuations by 14%-22% compared to opt-in, while a \$0–50 price anchor decreases valuation by 37%-53% compared to a \$50–100 anchor. Moreover, in some consumer segments, the susceptibility to frame influence negatively correlates with consumers' average valuation. Our results suggest that conventional frame optimization practices with a goal to maximize the supply of data can have opposite effects on its representativeness: A bias exacerbating effect emerges when consumers' privacy valuations and frame effects are negatively correlated. On the other hand, such a volume-maximizing frame may also mitigate the bias by getting a high percentage of consumers into the sample data, thereby improving its coverage.

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

---

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

<sup>†</sup>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, Avi Goldfarb, Michael Grubb, Tanjim Hossain, Yufeng Huang, Alex Imas, Jihye Jeon, Garrett Johnson, Yucheng Liang, Nina Mazar, Sanjog Misra, Ilya Morozov, Olivia Natan, Omid Rafieian, Heather Sarsons, Ananya Sen, Tim Simcoe, Oleg Urminsky, Giorgos Zervas, and Jinglong Zhao 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. Companies often deploy some form of “choice architecture” (Thaler & Sunstein 2008) when collecting consumer data, which are choice environments designed to nudge consumers towards sharing more data, all else equal. For instance, after the GDPR takes effect, websites commonly use a combination of default settings, salient options, and obstructions to nudge their users towards sharing all forms of cookie identifiers (Matte et al. 2020, Nouwens et al. 2020). Several major consent management platforms, such as OneTrust and Usercentrics, offer products and resources that help websites find the user interface design that will maximize user opt-ins.<sup>1</sup>

Existing choice architecture optimization practices emphasize maximizing the volume of data collection as their end goal. What is often neglected is a second dimension of data quality—the representativeness of the data collected. Biased input data often leads to biased insights and decision-making. For instance, Amazon had to abandon its AI hiring tool due to its gender bias; one leading cause was that the algorithm had insufficient female applicant resumes in its training sample (Dastin 2018). As another example, Cao et al. (2021) show that gender bias in Product Hunt’s product votes leads to severe bias in the predicted appeal of new products, leading to miscalibration in startups’ product launch decisions. Such bias in shared consumer data is likely to exist when different consumer groups have different valuations of their data (Lin 2022). However, an under-appreciated aspect is how choice architecture affects the representativeness of data shared when different types of consumers respond to choice architecture to different degrees.

In this paper, we ask the following questions: How do choice frames influence consumers’ privacy valuations, and what is the heterogeneity of the choice frame effects? How do the choice frames change the composition of consumers willing to share their data beyond its influence on the quantity of data shared? The economic returns of data to firms depend on both its quantity and its representativeness. Therefore, to assess how choice architecture affects the quality of data collected and the efficiency of data collection, we must account for its effect on both the volume and the composition of data shared.

To answer these questions, we recruited 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

---

<sup>1</sup><https://www.onetrust.com/blog/onetrust-launches-consent-rate-optimization-to-maximize-opt-ins/>; <https://usercentrics.com/resources/opt-in-optimization/>

choose default and price anchor because they are common or likely to be deployed if companies can buy data from consumers. We also collect consumer characteristics to explore heterogeneity in their privacy valuations and responses to choice frames. These variables include demographics, 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 \$100, the mean valuation across data is \$63.9. The difference in median WTA between the most valuable (friends and followers) and the second most valuable data (posts) is \$8.9, while the difference between the most and the least valuable data (survey responses) is \$28.4. As we relax the stringency of the top code, the valuations for each personal variable increase substantially, as well as the value differences.

We also find a significant influence of choice frames on consumers' valuations. Consumers decrease their valuations by 37.4%-52.6% in the *low* versus *high* price anchor condition. Compared to an *opt-out* default, *active choice* increases the valuation of data by 5.8%-11.8%, and *opt-in* increases the valuation by 13.6%-21.1%. On average, the difference in valuation is \$16.2 lower due to the *low* price anchor, and is \$2.2 and \$5.1 higher in the *active* and *opt-in* defaults when we use the \$100 topcode. Although the qualitative effect of defaults is uniform, 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.

To explore how choice frames affect the composition of the shared data, we deploy a set of causal forest models (Athey et al. 2019) to see what consumer attributes correlate with the heterogeneity in data values and frame effects. Overall, consumers' valuations of their data and their responses to choice frames are negatively correlated. Such a negative correlation is stronger across specific consumer segments. Consumers who value their data less across frames are overall younger, poorer, less educated, and more likely to click on ads while on Facebook. Interestingly, these attributes also predict larger frame effects.

The negative correlation between privacy valuation and choice frame effects poses a potential tension between volume-maximizing and bias-minimizing objectives during data collection. When a firm can choose the choice frame while buying data, it often wants to adopt a frame that maximizes the volume of data collected for a given price. Such a volume-maximizing frame has the potential to exacerbate the statistical bias in the data collected by the firm: Consumer groups who already value their data less absent frame effects 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, the potential bias induced by a volume-maximizing choice frame can counteract the benefits of collecting more data and decrease the value of shared data as a result.

We perform counterfactual analysis to explore how a firm may want to deviate from their current choice architecture design when both the volume and data composition matter. We start with the natural example of aiming for a representative sample. The prevalent consent optimization practices aim to find a choice frame that maximizes the volume of data collected. We show that such a volume-maximizing frame can have countervailing effects on the bias in collected data, compared to a benchmark in which we average across different frames in our experiment. With a negative correlation between privacy valuation and choice frame responsiveness, the volume-maximizing frame can potentially exacerbate the bias in sample data. In particular, holding fixed the sample size target, the volume-maximizing frame tends to collect data with more bias compared to the benchmark frame. On the other hand, increasing the volume of data can be useful for mitigating the bias: Since the volume-maximizing frame gathers more data for any given price point compared to the benchmark frame, the impact of the volume-maximizing frame on bias in sample data is ambiguous. Eventually, if the firm is able to gather 100% of consumer data, the resulting data has zero bias. Therefore the bias-mitigating effect will dominate as the percentage of consumers willing to share their data grows large.

We conclude by discussing a generalization of the above point. To the extent that the data collector cares about a specific population, there is a possible trade-off between the volume of data collected and data composition, or conversely, a frame can act as a screening or discrimination device. For example, if a firm’s clients are mostly young men, that might be the client type a company cares most about collecting data from. One alternative, if personalization is not practical, is to find the choice environment that *relatively* maximizes the share of young men sharing their data at a particular price. Another approach, especially lucrative if third-degree price discrimination is not allowed, is to personalize the choice environment such that young men face their volume-maximizing choice environment and other types of consumers face their volume-minimizing frame.

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, empirical studies on the 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 context-specific (Martin & Nissenbaum 2016) and hard to measure. Lin (2022) highlights consumers’ economic reasoning in different data usage scenarios as a contributor to this context effect. In contrast, here we focus on how choice frames influence privacy valuations and how this influence varies across consumer segments.

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, D’Assergio et al. 2022, Tomaino et al. 2022). Our goal is not just to document the choice architecture effects, but rather to relate the frame effects to how consumers self-select into sharing data and the quality of data available to firms. We achieve this goal by tracing the effect heterogeneity along the data supply curve, and showing how the effect heterogeneity translates to different selection patterns when the firms choose different choice frames. As such, our focus connects two disjoint threads of literature: behavioral biases in privacy valuations, and the efficiency of data markets (Arrieta-Ibarra et al. 2018, Acemoglu et al. 2022, Bergemann et al. 2022, Ichihashi 2021, Markovich & Yehezkel 2021). We also contribute to the empirical literature on behavioral industrial organization that takes consumers’ biases and fallibility into firms’ decision-making processes (Peltzman 1981, Rao & Wang 2017, Strulov-Shlain 2023, Miller et al. 2022).

Our project is closely related to the literature on the value of consumer data to firms (Rossi et al. 1996, Trusov et al. 2016, Miller & Skiera 2017, Rafieian & Yoganarasimhan 2021, Bajari et al. 2019, Peukert & Sen forthcoming, Sun et al. 2021, Wernerfelt et al. 2022). Iansiti (2021) theorizes that the marginal value of data is initially very high as firms try to overcome the model “cold-start” phase. However, diminishing marginal returns quickly set in, after which a firm’s competitive advantage of data ownership mainly comes from having unique data points. This theory is consistent with recent computer science literature (Hestness et al. 2017, Kaplan et al. 2020), which show that the generalization error of deep learning models follows a power-law distribution as the data size grows. In our framework, this theory implies that a firm may initially place more weight on increasing data volume, but will eventually shift to getting a more representative dataset with broader coverage to gain a competitive edge.

Our work is also related to the recent research on how bias in input data creates biased algorithms (Cao et al. 2021, Agan et al. 2023), variation of statistical accuracy in business analytics (Lin 2022, Neumann et al. 2022), and other market outcomes (Johnson et al. 2020). Although the sources of input data bias vary, individual differences in privacy concerns is often one of them. Our work explores this angle further by examining how a choice architecture chosen by the firm may exacerbate or alleviate this bias.

Lastly, our work is related to the literature on selection markets, where the firm’s cost of serving products depends on which consumers purchase the product (Einav & Finkelstein 2011, Veiga & Weyl 2016, Mahoney & Weyl 2017, Einav et al. 2021). In the context of data markets, it is the return to data, not the cost, that varies with consumer composition. In particular, a consumer is “good to include” if similar consumers are scant in the firm’s database. Similar to a selection market where prices are regulated, the firm can improve its profit by using non-price instruments to screen in low-cost or high-value consumers.

The rest of the paper is organized as follows. Section 2 uses a conceptual model to illustrate why accounting for the joint heterogeneity of privacy valuation and frame effect is crucial for

understanding how frames affect bias in the sample data. Section 3 illustrates the design of our experiment. Section 4 describes our data and reduced-form evidence, followed by heterogeneity analysis for the privacy valuations and choice frame effects. Section 5 introduces our counterfactuals to demonstrate when a volume-maximizing frame may create the bias-volume tradeoff, and Section 6 concludes.

## 2 Conceptual Framework

In this section, we lay out the framework that underpins our experimental design and subsequent analysis. We start by distinguishing 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' best guess about the true value of their privacy. We note that this true value is not observed by the researcher or the consumer. The observed valuation is what we call the expressed valuation  $\tilde{v}$ , which is influenced by the frame-neutral valuation but is subject to the influence of choice architecture,  $\theta$ :

$$\tilde{v} = f(v_0; \theta). \quad (1)$$

In other words, the choice architecture creates a gap between the expressed valuation and the consumers' prior judgment of what the value of their data is.

We situate the two valuations in the context of a data market, where firms directly interact with consumers to get their consent for sharing data. In such a market, consumers form their supply for data based on their privacy valuations; the firm has a demand for data, offers a price for buying data, and can choose the choice frame to influence the supply from consumers.<sup>2</sup> Although  $\tilde{v}$  may be larger or smaller than  $v_0$ , the firm often chooses a frame that delivers the lowest  $\tilde{v}$  to maximize its gain per dollar offered. In essence, the firm-chosen frame 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 is lower compared to the frame-neutral equilibrium.


So far, we have shown how choice frames can induce behavioral distortions and change data collection by treating data as a standard commodity. However, consumer data is not a standard commodity. One unique feature of data is that its value to firms depends on not only its volume but also its representativeness. When evaluating the impact of choice architecture, the firm should care about not just its impact on the volume of data shared, but also which consumers are more likely to trade and how that affects the bias in the data sample.

How does the consideration of sample bias affect our evaluation of frame effects? Consider the following example shown in Figure 1b. Absent the choice frame effects, low-income consumers value their privacy less compared with wealthier ones; however, their valuations are also more

---

<sup>2</sup>The price for consumer data can either be explicit or implicit. For example, a form of implicit price can take the form of personalization benefits upon data sharing.

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



susceptible to the influence of choice architecture. If a frame-neutral choice architecture exists, then 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 downward, such selection bias becomes more severe because poorer consumers push their valuations down even further compared to the wealthier ones. A trade-off emerges: The 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.

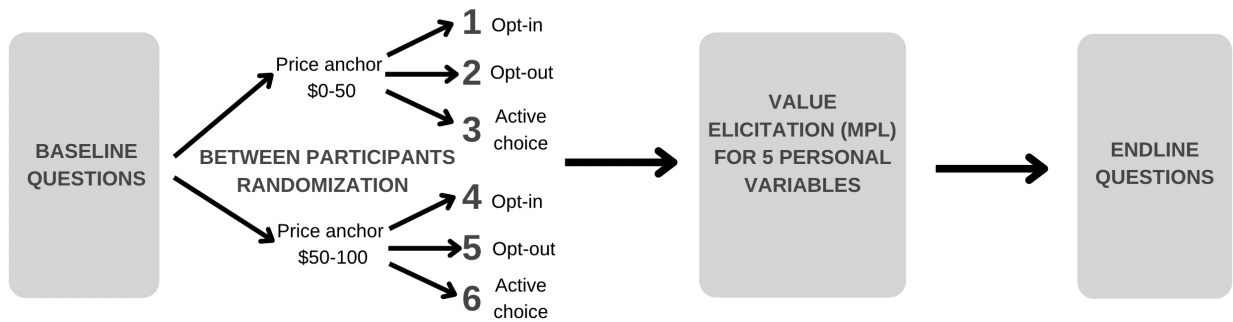
In our illustrative example above, deploying a volume-maximizing frame ends up exacerbating bias in the sample data due to the negative correlation between consumers’ privacy valuations and their frame responses. However, we note that in empirical settings, this need not always happen. For instance, suppose all consumers who share their data in the neutral frame belong to the low-income segments. In this case, deploying the volume-maximizing frame will never exacerbate the bias (it was maximally biased to begin with), and may even alleviate the bias by including more consumers in the sample. In other words, the volume maximizing frame can also have a bias-mitigation effect through supply expansion.

In what follows, we will first use an experiment to randomize choice frames between participants, 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 baseline 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 evaluate frames that firms commonly use to maximize the volume, and show when the trade-off between the variance and bias will emerge when the firm deploys different frames.

### 3 Experiment

The main component of our experiment is a multiple price list (MPL) that elicits consumers’ valuation of 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 between participants 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 flow of our experiment. Below, we start with the participant recruiting procedure, then introduce the value elicitation components, the choice frame treatments, and end with the design of survey questions that measure consumer characteristics.

Figure 2: Experiment Overview





### 3.1 Participant Sources

We recruit our participants from two sources: Facebook Ads and Prolific. Using these sources confers two advantages. First, both sources allow us to screen participants based on the availability of their Facebook accounts.<sup>3</sup> 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 participants from both sources 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 and survey panel (Prolific) participants allows us to examine this possibility. We restrict our participants to English speakers living in the US, with an age 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 participants; 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,<sup>4</sup> then ask for their consent to enter the study.

### 3.2 The Multiple Price List and Choice Frame Treatments

We measure participants’ valuation of their Facebook data using a multiple price list (MPL). MPL resembles Becker-DeGroot-Marschack (Becker et al. 1964) in its use of a lottery to ensure incentive compatibility. Its advantage is simplicity: Since MPL uses simple take-it-or-leave-it offers repeated at different price points instead of a second-price auction, it is easier for participants to understand the procedure and why telling the truth is optimal. Standard MPLs give us only the interval that includes a consumer’s valuation of their data (see Figure A.3). To get more granular numbers, we follow each MPL with an open-text question asking participants their exact valuation of the data. If a participant chooses not to share data in all the MPL questions displayed, we also give them the option to indicate “I do not want to share my data at any price” in the open-text prompt (see Figure A.4). 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 than or equal to the offer price randomly generated by the computer.

After answering the baseline questions, participants see a message asking for their valuation to share data with advertisers. An example of the MPL procedure follows, showing them how

---

<sup>3</sup>When recruiting the Facebook participants, we restrict our ad placement to positions only viewable by logged-in Facebook users, which excludes Facebook Audience Network and Instagram. For the Prolific participants, we use the internal selection tool provided by the platform to enlist only Facebook users.

<sup>4</sup>This information includes their data sharing behavior and privacy attitudes on social media, but to minimize selection bias, we did not explicitly mention that on the introduction page.

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 frame 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 (see Figure A.3). The second frame is the **price-range** 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.

We inform participants that their responses will be chosen to implement the data-price exchange with positive probability. For those who were chosen and whose valuations are lower than our randomized price, we sent them a step-by-step guide to download the Facebook variable. They receive the payment within 24 hours after sending their data to us.

### 3.3 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?*

### 3.4 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 within 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 \$95, 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 (Spiek-

ermann et al. 2001, Singleton & Harper 2002). One argument is that the gap between stated and revealed preferences diminish once the context is controlled for (Prince & Wallsten 2020). Nevertheless, we adopt a variety of strategies in the data analysis to emphasize the valuations within the incentive-compatible range. The reduced-form analysis will focus on log valuation as our preferred specification. In addition, we will allow truncation at different finite points above \$95 and see whether the results are sensitive to different specifications.

## 4 Data and Model 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 of our participants. Compared with the US representative demographics,<sup>5</sup> ours includes more females and are overall better educated; the distributions of age, income, and ethnic majority are similar to the national average. Compared to those recruited from Prolific, the Facebook ad participants 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. Thus, having participants from both sources allow us to cover a wider demographic range, which allows us to demonstrate how privacy valuations and frame effects differ across the demographic spectrum. Table B.2 and B.3 shows that our participants are balanced across the six treatment conditions.<sup>6</sup>

The second step in the WTA elicitation is unrestricted; thus participants can be inconsistent. For example, a participant may say that they are willing to sell their posts for \$40 but not for \$30 in MPL, and then ask for \$56 on the next screen. Appendix Figure B.1 shows responses from all participants who completed our study, with their free-text valuations on the Y-axis and the implied WTA from their MPL responses on the X-axis. We find that 93% of participants give consistent valuations throughout. Among participants who have given inconsistent answers, many deviate to a range close to where they were. More importantly, 83% consumers who give inconsistent valuations only do so occasionally, suggesting that their inconsistency is more a sign of regret than a byproduct of inattentiveness.

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. We allowed for infinity reporting because earlier work posits that some consumers can be "privacy fundamentalists", who would reject any benefits from data uses in exchange for privacy (Westin 2003, Woodruff et al. 2014) and we wanted the valuation measurement to allow for this possibility. Nevertheless, most

<sup>5</sup><https://www.census.gov/quickfacts/fact/table/US/PST045221;comm/aging-nation-median-age.html>.

<https://www.census.gov/library/visualizations/2022/>

<sup>6</sup>Here, we separate the covariate balance tests for the attributes measured in the baseline and endline surveys. One concern about the endline survey responses is that they may be influenced by the treatments, especially when it comes to beliefs about data usage. Table B.3 shows that the endline responses do not differ significantly across treatments, and can be included in our heterogeneity analysis.

Table 1: Summary Statistics of Participant Characteristics

	Overall		Facebook		Prolific		U.S. Census
	Mean	SD	Mean	SD	Mean	SD	Mean
<b>Number of participants</b>							
N	5028		2010		3018		
<b>Race (percentage)</b>							
White	0.8	0.4	0.73	0.44	0.85	0.35	0.76
Black	0.07	0.25	0.08	0.27	0.06	0.23	0.14
Asian	0.12	0.32	0.16	0.37	0.08	0.28	0.06
Other	0.05	0.05	0.06	0.06	0.04	0.04	0.05
<b>Gender (percentage)</b>							
Female	0.59	0.49	0.76	0.43	0.48	0.5	0.51
<b>Median age</b>							
Median age	39.5	12.77	39.5	13.23	39.5	12.22	38.8
<b>Median household income (\$)</b>							
Median household income	62500	48675	62500	51629	62500	45224	64994
<b>Education (percentage)</b>							
High school graduate or higher	0.99	0.1	0.99	0.09	0.99	0.11	0.89
Bachelor’s degree or higher	0.64	0.48	0.73	0.44	0.58	0.49	0.33
<b>Facebook Questions</b>							
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 (percentage)	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	
FB or Instagram ad click (times/mo)	1.66	1.81	2.47	1.92	1.12	1.51	


consumers in our experiment do not seem to adopt a fundamentalist attitude in that they are selective when reporting infinite values. For example, 7% of participants report infinity valuation for their survey answers compared to 31.7% for friends and followers. Only 3.7% participants report infinity values for all items.

Genuine or not, the infinity values create challenges in reporting summary statistics and reduced-form analysis. We adopt several strategies 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 *posts* and *about me* information; *survey answers* is valued the least. When top-coded at \$100, the mean valuation for *friends* data is \$77.04; the average valuation across all Facebook data is \$67.14, compared to an average of \$48.64 for sharing the survey responses. Given that 20% consumers report their data valuation at above \$100, these value differences serve as the lower bounds for the actual differences.

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 choosing an arbitrary starting point. One may expect such a phenomenon to be likely for private data valuations, as the pros and cons of sharing data can be uncertain. 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 B.2). 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 Tobit regressions, with the valuations top-coded at \$100 as the outcome variable. Our preferred specification is Model 4, which uses the log form of valuations to decrease the sensitivity of estimates on top code choices, and includes the types of personal variables as additional controls. Compared to active choice, an opt-out frame decreases the average valuation by 5.8%, while opt-in increases the valuation by 7.8%. The influence of a price anchor is more substantial. Consumer valuations for data decrease by 52.6% 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	WTA	log(WTA)	log(WTA)
Intercept	63.884 *** (0.876)	48.636 *** (0.881)	4.001 *** (0.026)	3.570 *** (0.027)
Price Anchor = Low	-16.112 *** (0.947)	-16.234 *** (0.921)	-0.523 *** (0.028)	-0.526 *** (0.028)
Default = Active	2.377 * (1.139)	2.205 * (1.108)	0.063 + (0.035)	0.058 + (0.034)
Default = Opt-in	5.178 *** (1.147)	5.101 *** (1.119)	0.138 *** (0.035)	0.136 *** (0.034)
Likes		8.269 *** (0.455)		0.268 *** (0.014)
About Me Page		17.875 *** (0.523)		0.509 *** (0.015)
Posts		19.459 *** (0.535)		0.550 *** (0.016)
Friends and Followers		28.399 *** (0.628)		0.764 *** (0.018)
Num. Obs.	25140	25140	25140	25140

+  $p < 0.1$ , \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . Outcomes variables are top-coded at \$100; 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, with the supply curve corresponding to active choice sitting squarely between opt-in and opt-out. 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.<sup>7</sup>

Figure 4: Data Supply Curves by Treatment Condition




Table 3 shows versions of the log outcome model with different top codes for the data valuation. As the top code becomes less stringent, the coefficient representing the price anchor effect decreases

<sup>7</sup>Appendix Table C.2 includes further controls and variables of interest. Consumers who ask for more than the token's value in the practice round also report higher WTA for their data. In addition, consumers who believe specific data are already available to Facebook and to the public value their data less when it comes to sharing data with advertisers.



in magnitude while all other coefficients increase. This pattern is consistent with the fact that price anchor creates “backlash” in a subset of consumers.

Table 3: Average Treatment Effects and Data Valuation Across Topcodes: Tobit Regressions

	DV: Free-Text valuation					
	100	250	385	500	750	1000
Intercept	3.556 *** (0.024)	3.584 *** (0.028)	3.596 *** (0.030)	3.603 *** (0.031)	3.614 *** (0.034)	3.622 *** (0.035)
Price Anchor = Low	-0.536 *** (0.023)	-0.479 *** (0.028)	-0.455 *** (0.031)	-0.440 *** (0.032)	-0.418 *** (0.035)	-0.402 *** (0.037)
Default = Active	0.042 (0.029)	0.065 + (0.034)	0.075 * (0.037)	0.081 * (0.039)	0.090 * (0.042)	0.096 * (0.045)
Default = Opt-in	0.105 *** (0.029)	0.138 *** (0.034)	0.152 *** (0.037)	0.161 *** (0.039)	0.174 *** (0.042)	0.183 *** (0.045)
Likes	0.226 *** (0.012)	0.286 *** (0.014)	0.313 *** (0.016)	0.329 *** (0.017)	0.353 *** (0.018)	0.370 *** (0.019)
About Me Page	0.409 *** (0.013)	0.540 *** (0.016)	0.598 *** (0.017)	0.633 *** (0.018)	0.688 *** (0.020)	0.726 *** (0.021)
Posts	0.446 *** (0.013)	0.590 *** (0.016)	0.653 *** (0.018)	0.691 *** (0.019)	0.749 *** (0.021)	0.791 *** (0.022)
Friends and Followers	0.562 *** (0.014)	0.810 *** (0.017)	0.919 *** (0.019)	0.986 *** (0.021)	1.089 *** (0.023)	1.163 *** (0.025)
Num. Obs.	25140	25140	25140	25140	25140	25140
R2	0.108	0.089	0.083	0.080	0.076	0.074

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

### 4.3 Heterogeneous Effects of Choice Architecture

In Section 2, we show that the correlation between consumers’ privacy valuations in the neutral benchmark condition and their responses to choice frames is the key driver that creates the tension between volume-maximizing and bias-minimizing effects across frames. Simply comparing data supply curves under different frames is insufficient for capturing such correlation, as the supply curves do not show how much the same consumer (segment) would value their data under different conditions. Instead, we must rely on heterogeneous effect models to characterize the joint distribution of consumers’ privacy valuations and the frame effects.

To achieve this goal, we estimate causal forest models proposed by 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. The caveat of the survival forest is that it only compares two treatments at a time; as such, it does not capture the potential interaction 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.

Note that for many choice frames including the price anchor, what counts as a “neutral” frame in practice is unclear. In view of this fact, we instead construct the benchmark privacy valuation as the average of a participant’s would-be valuations across the six treatments, predicted by our model. This average valuation underpins a consumer’s data sharing choice when the firm randomly chooses among the six possible frames. It is possible that a true frame-neutral valuation lies closer to one of the frames or even outside this range, but for now this approach will allow us to get a comparable valuation across participants. 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 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 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 range has a larger overall effect compared to the default. The estimated effects have similar standard errors, but since the effects of defaults are smaller in magnitude, their heterogeneous treatment effects are generally insignificant while the price-range effects are significantly different from zero.<sup>8</sup>

Figure 6 shows the correlation between frame effects and the average 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 price anchor is negatively correlated with the average valuation: consumers who have lower privacy valuations are also more likely swayed by a price anchor (correlation is -0.46). This pattern is consistent with Collis et al. (2020) who 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 have lower privacy valuations and those more responsive to choice frames, we project the average valuation and the heterogeneous “maximal frame effects” to the consumer attributes using linear models. Here, we define the “maximal” frame effect as the difference between the two frames that give the highest and lowest average treatment effects. In other words, we calculate for each consumer the difference in predicted valuation between the frame that maximizes the average reported valuation (opt-in, high price range) and the frame that minimizes it (opt-out, low price range).

---

<sup>8</sup>The standard errors for heterogeneous treatment effects are generally larger than for the average treatment effect. Therefore, it is natural for default to have a significant treatment effect in Table 2 while having insignificant effects here.

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

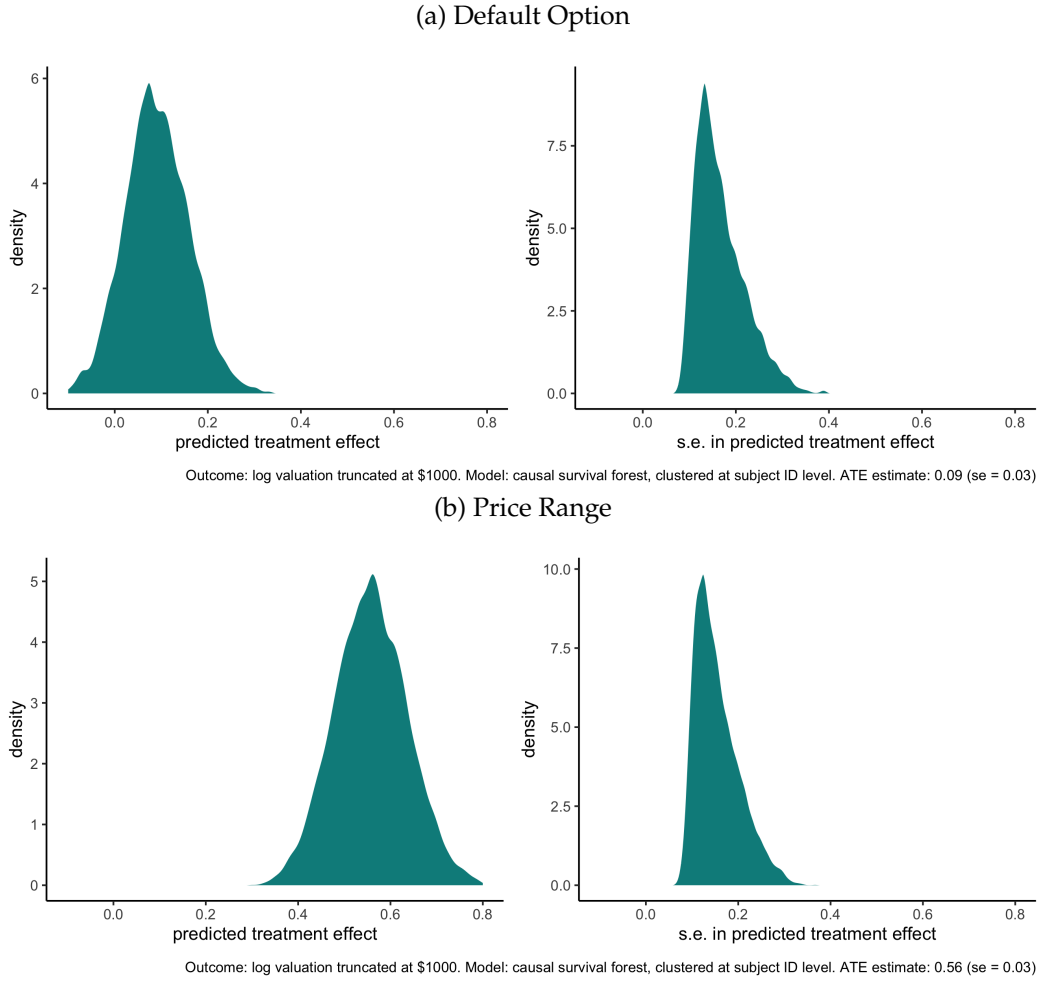



Figure 7 shows the linear projection estimates and 95% confidence intervals. The top panel shows the predictors of a high baseline valuation. Consumers who value their personal data more are older, richer, better educated, more likely to be female and Asian, spend less time on the Internet, and are less likely to click ads on Facebook. This is broadly in line with findings in the existing literature. For example, Goldfarb & Tucker (2012) find that older people and women value their privacy more, and Collis et al. (2020) show that high-income consumers and Asian communities value their Facebook data more. In comparison, those more easily influenced by our choice frames are overall younger, poorer, less educated, and more likely to click on ads while using Facebook. Overall, these are attributes that predict high choice frame effects while also predicting lower valuations for personal data.


As a final note, we separately compare the average log valuation and choice frame effect sizes from our two participant sources, which is shown in Figure 8. The Facebook participants have a higher privacy valuation overall and is less influenced by frames. Although we do not know if the Facebook participants are more representative of the total population than the Prolific panel,

Figure 6: Correlation Between Choice Frame Effects and Average Log Valuation



*Note:* In the figures above, treatment effects are estimated using survival forests, with the outcome as log valuation truncated at \$1000. The average 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.


Figure 7: Heterogeneous Privacy Valuation and Treatment Effects by Consumer Subgroups



Note: In the first projection model, the outcome is a consumer's predicted log valuation marginalized across the choice frame treatments. In the second projection model, the outcome is the predicted *maximal treatment effect*, represented as the difference between the choice frames with the highest and lowest ATE. Both the treatment effects and marginalized valuations are estimated using multi-arm causal forests, with log valuation censored at \$1000 as the outcome.

our results suggest that studies using multiple participant sources have merit in assessing the robustness of the effects they aim to study.

Figure 8: Average Log Valuation and Choice Frame Effect Distributions by Participant Source



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

## 5 Exploring the Volume-Bias Tradeoff

Our analysis has shown that choice frames can change the composition of consumers who share data since different consumers respond to frames differently—the question is how. With different frames, consumers have different willingness to share data. Does the frame chosen by the firm exacerbate or alleviate the existing selection bias? A company is often inclined to choose a choice architecture that maximizes the volume of data collected. Is this the optimal frame for data collection, or does volume comes at the cost of representativeness?

To answer these questions, we construct a set of counterfactual simulations based on the forest-model estimates. Our goal is to examine the volume and bias in data collected under different choice frames. In particular, we compare the performances of two frames: a *volume-maximizing* frame, which maximizes the supply of data at each price point; and an *average* benchmark frame, which we construct by averaging a consumer’s would-be valuations across all frames we have tested. Our results show that a volume-maximizing choice frame can have opposite impacts on the bias in sample data through two distinct mechanisms, illustrated below.

### 5.1 Setup

We use our causal forest estimates to construct the data supply curves under each choice frame. In our setting, the (uniform) volume-maximizing frame means opt-out default and low price anchor. To construct valuation and choice under the average frame, we average the valuations across the six frames in our experiment. These supply curves allow us to figure out who will share data for a given price under each frame, thereby constructing the *sample data* under different price points and frames. We repeat the sample data calculation for a grid of prices between \$30 and \$80, with \$5 increments.<sup>9</sup> For each sample data, we calculate its bias and volume:

- *Bias* is represented as the standardized difference in attributes between consumers who share data and all consumers in our experiment. For continuous attributes such as income, this is the standardized difference in means; for discrete attributes, this is the difference in the percentage of consumers with the attribute. We first focus on biases in individual attributes to illustrate the different mechanisms, then average over the standardized mean difference to form an aggregate bias measure.
- *Volume* is represented as the percentage of consumers captured in the sample data. Although alternative metrics for volume exist, using the percentage metric allows us to better characterizes how the volume-maximizing frame changes the bias in data in different regions.

---

<sup>9</sup>We choose the range of price grid such that the lowest price (\$30) still guarantees that a nonzero percentage of consumers will share data in the average frame. If the percentage of consumers sharing the data is zero, we will not be able to calculate the bias metric. The highest price can be more flexible, though due to the long-tail nature of consumer valuation, increasing prices in \$5 increments at the higher end does not move the amount of data shared as much, thus we stop at \$80.

Before we proceed, it is worth clarifying what kind of data bias the firm cares about and when such bias decreases the value of consumer data. The value of data-driven analytics comes from learning and prediction. In other words, there is a series of outcomes  $Y_i$  with distribution  $f(Y|X)$  that the firm wants to learn about using consumer data.<sup>10</sup> The firm observes  $X$  from all consumers at the point of data collection, but can only learn about  $Y$  among consumers who agree to share their data. A biased sample means the sample distribution  $f(X|\tilde{v}(\theta, X) > P)$  does not equal the target population distribution  $f(X)$ . Note that the target population can be the firm’s desired customer database and not necessarily representative of the general population. Such a bias can compromise the statistical accuracy of data-driven insights and thus the value of shared consumer data in the following scenarios:

1. The firm cares about the average outcome  $E[Y] = E[f(Y|X) \cdot f(X)]$ . However, the sample data gives  $f(Y|X) \cdot f(X|\tilde{v}(\theta, X) > P) = f(Y|\tilde{v}(\theta, X) > P)$ . When the firm does not know either  $f(X)$  or  $f(X|\tilde{v}(\theta, X) > P)$ , it is unable to recover  $E[Y]$  by reweighting the sample data to match the target population distribution. This is represented by the Product Hunt example (Cao et al. 2021), where startups knew their target customer base but had no idea about the gender (or other demographic) representation of the votes from the platform.
2. The firm cares about the full, heterogeneous distribution of the outcome  $f(Y|X)$  and is able to observe  $X$  from the sample data. That allows them to learn  $f(Y|X)$ ; however, some consumers  $X = x_1$  are underrepresented in the dataset, which means that the accuracy of estimate for  $f(Y|X = x_1)$  is low. This is typical in settings where the firm (or other data users) needs to know the heterogeneity of outcomes to design targeted campaigns and interventions.

In principle,  $f(Y|X)$  may not always be different across  $X$ ’s (i.e., no systematic relationships between the outcome and observable characteristics). However, the data buying firm often needs to use the same consumer data to learn about many different outcomes, making this possibility exceedingly low in practice. In a typical data market, the data buying firm is an intermediary (e.g., Meta, Google, Experian) that later resells the raw or derived data to different end-user firms.<sup>11</sup> Since different end-user firms have different customer bases and different analytic objectives, the data-buying firm has to integrate over all these use cases. Even in cases where the data buying and data using firms are one and the same, the firm often still needs to form an expectation over potential future uses of data, hence the phrase “data is a strategic asset”. In situations like these, the objective of minimizing biases in the learned outcome vector  $Y$  boils down to the objective of minimizing biases in observed attributes  $X$  in the sample data, while the firm remains agnostic about the specific data production in each use cases. Motivated by this observation, we focus on the representativeness of observables  $X$  in the sample data in our following analysis.

<sup>10</sup>Examples of  $Y$  include consumers’ willingness to pay for certain products, interest in certain political topics or product categories, etc.

<sup>11</sup>Here, selling derived data can mean selling impressions from certain consumer segments, where the advertisers specify what segments they want to reach based on demographics or inferred consumer interests.

## 5.2 The Bias and Volume of Data Across Frames

We start by looking at the biases in attributes where we see a negative correlation between the average valuation and frame responses across segments. These attributes include income, age, and education. Figure 9 shows the biases and volumes of the sample data collected under each frame conditional on the *sample size*. This is akin to data collection with a fixed sample size target, such as when the firm needs a certain sample size to run an experiment or to construct a targeting segment. Fixing the sample size target, we see that data collected under the volume-maximizing frame are further away from the representative benchmark (the dashed line), meaning they are more biased. This pattern confirms our intuition: when consumer segments initially with low valuations for data are influenced by the frame more, they become much more willing to share data under the frame, and thus the sample data over-represents these consumers.


However, another pattern shown in this figure is the positive relationship between volume and representativeness within each frame: as the percentage of consumers sharing data increases, the sample becomes less biased overall. Intuitively, as the sample data approaches a coverage of 100%, it will eventually gather all consumers and eliminate the bias. As another example, imagine the first 50% of consumers who share the data are all low-income consumers, while the next 50% are high-income ones. As the coverage goes beyond 50%, the sample eventually becomes less biased, because all the low-income consumers already choose to share data, and the new consumers at the margin serve to mitigate the sample bias. In other words, although the volume-maximizing frame can exacerbate the bias when there is a negative correlation between privacy values and frame effects (Mechanism 1), it can also mitigate the bias, since it gets a higher coverage of consumers for a given price (Mechanism 2).

Figure 10 shows the net effect of the two mechanisms from a different perspective by comparing sample data collected under the same *price*. The two mechanisms counteract each other: the bias exacerbating effect tends to dominate when the coverage of data is low, while the bias mitigating effect starts to dominate as the coverage of data increases. Since the bias exacerbating effect is only present when privacy valuation and frame effects are negatively correlated, the volume-maximizing frame generally alleviates bias in dimensions without such a negative correlation. Aggregating across demographic attributes, we see that in our setting the volume maximizing frame mitigates sample bias conditional on price but not conditional on sample size (see Figure 11).

Figures 9 and 10 represent two extreme cases of market equilibrium: the first case compares the data quality when the firm's demand for data is perfectly inelastic, while the second represents the case when the firm has a perfectly elastic demand. The more likely scenario is a downward sloping demand curve: if the marginal value of data decreases but is not zero. In this case, the firm decreases its price in response to expansion in data supply under the volume maximizing




Figure 9: Bias and Volume in Sample Data Across Frames: Fix Sample Size Comparison



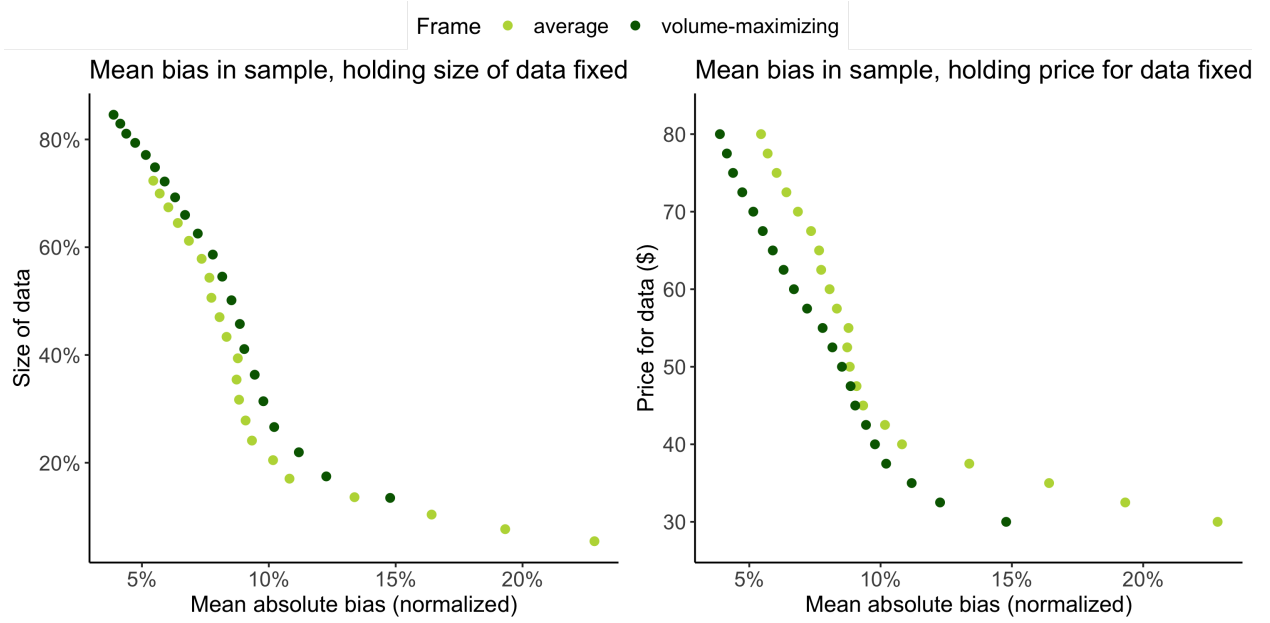
*Note:* In the figures above, each point represent a sample dataset collected under a choice frame  $\times$  price combination. The x-axis represents the average of the attribute values among the sample data, and the y-axis represents the size of the dataset. The dashline shows the average value of the attribute in the full data. The prices for data vary from \$30 to \$80 in \$5 increments within each frame, thus there are 11 sample data realizations accordingly.

Figure 10: Bias and Volume in Sample Data Across Frames: Fix Price Comparison



*Note:* In the figures above, each point represents a sample dataset collected under a choice frame  $\times$  price combination. The x-axis represents the average of the attribute values among the sample data, and the y-axis represents the price offered to consumers. The dashed line shows the average value of the attribute in the full data. The size of the dots represents the sample size as a percentage of the total consumers.

Figure 11: Bias and Volume in Sample Data Across Frames: All Demographics Data



frame. The equilibrium price, volume, and bias in data shared will be somewhere between the two extreme cases described above.

Should firms adopt the volume-maximizing frames when collecting consumer data? We show that the answer is yes if the firm can get a high percentage of consumers to share their data, either by providing a high price, or a product valuable to most consumers that justifies their privacy costs. If not, the effect of volume-maximizing choice framing is ambiguous: it depends on whether consumers who value their data less are also more easily swayed by frames. In our setting, such a negative correlation between privacy valuation and choice frame response exists among certain consumer segments. We note that the correlational pattern is likely to be context-specific, given that the privacy preference and its heterogeneity distribution can depend on the consequences of data sharing in different contexts (Lin 2022).

How can a firm choose a more efficient frame in its own setting? The key to selecting the optimal frame lies in knowing the joint heterogeneity of privacy preferences and choice frame responses. We propose a sequential optimization strategy, where the firm uses a model-enriched multi-armed bandit (MAB) to simultaneously learn about this joint heterogeneity and choose the optimal frame given its current knowledge. Imagine a typical scenario where consumers visit the firm in sequence. The firm can then allocate some traffic to see a randomized set of choice frames for learning, and the rest of the traffic to the expected optimal arm. Using a multi-armed bandit to search for the optimal design is an existing feature in several experiment-as-a-service platforms.<sup>12</sup>

<sup>12</sup>See: <https://vwo.com/blog/multi-armed-bandit-algorithm/>, and <https://www.optimizely.com/optimization-glossary/multi-armed-bandit/>.

Vanilla MAB is inefficient when the payoffs depend on multiple parameters (i.e., heterogeneity in privacy preferences and frame effects) and are correlated across different arms. However, previous work (Qiang & Bayati 2016, Schwartz et al. 2017, Misra et al. 2019) shows that enriching MAB with a model can substantially improve learning efficiency by reducing the dimension of the parameter space to learn about. We believe that such a model-enriched MAB has great potential to improve the framing choices and data collection efficiency over existing practices.

## 6 Conclusion and Future Analysis

Choice architecture can substantially distort consumers' privacy valuations and thus the supply of consumer data. The choice frames we examine—default and price anchor—shift the average consumer valuation for their Facebook data by 14% and 53%, respectively. Moreover, for some consumer segments, the susceptibility to frame influence is negatively correlated with their valuation for data absent frames. Younger, lower-income, 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.

The fact that different consumers respond to frames differently implies that choice frames can substantially change the composition of data that firms collect. We show that a conventional practice—choosing a frame that maximizes the supply of data—can have opposite effects on the bias in the data collected. In cases where consumers who value data less also respond more to choice architecture, a frame that aims to decrease valuations and maximize data supply can exacerbate the bias in data. However, a frame that maximizes the supply of data can also mitigate the bias, as increasing the volume also means improving the coverage of the sample data. We show that the bias mitigating effect tends to dominate when the firm can get a high percentage of consumers to share their data by providing a high price.

In future analysis, we plan to explore how a firm can improve the efficiency of data collection by using customized choice frames or screening contracts. The feasible customization rule depends on what information is available to the firm before gathering the consumer data. One scenario is that the firm knows the general privacy sentiments and the likely susceptibility to frame influences across consumer segments and has some segment tags available. For example, firms may know the demographic attributes of the consumers before asking them to share their behavioral data. In situations like this, a set of choice frames customized for each demographic segment may push the quality of the collected data further toward the bias-representativeness frontier. Another scenario is that the firm does not have any existing information about the customers, but can construct payment schemes that have a better appeal to consumers who have high privacy preferences and are hard to be influenced by frames. Imagine an airline company that knows high-income consumers have higher privacy preferences in general and are less likely to be swayed by choice

frames. In this case, they can design membership tier benefits that appeal more to the high-income segments to encourage more of them to sign-up for membership.

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. 2022), 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 when companies' efforts to maximize data collection can exacerbate or alleviate the bias in different market conditions. We believe the bias aspect is worth emphasizing, as it represents a novel channel in the data market that impedes its functioning.

## References

- Acemoglu, D., Makhdoumi, A., Malekian, A. & Ozdaglar, A. (2022), 'Too much data: Prices and inefficiencies in data markets', *American Economic Journal: Microeconomics* **14**(4), 218–256.
- 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.
- Agan, A. Y., Davenport, D., Ludwig, J. & Mullainathan, S. (2023), 'Automating automaticity: How the context of human choice affects the extent of algorithmic bias', *NBER Working Paper No. 30981*.
- 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', *Quarterly Journal of Economics* **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 'AEA 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. (2022), 'The economics of social data', *The RAND Journal of Economics* **53**(2), 263–296.
- Cao, R., Koning, R. M. & Nanda, R. (2021), 'Biased sampling of early users and the direction of startup innovation', *NBER Working Paper No. 28882*.
- 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'.
- D'Assergio, C., Manchanda, P., Montaguti, E. & Valentini, S. (2022), 'The race for data: Gaming or being gamed by the system?', *Available at SSRN 4250389*.
- Dastin, J. (2018), 'Amazon scraps secret ai recruiting tool that showed bias against women'.  
**URL:** <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>
- Einav, L. & Finkelstein, A. (2011), 'Selection in insurance markets: Theory and empirics in pictures', *Journal of Economic perspectives* **25**(1), 115–38.

- Einav, L., Finkelstein, A. & Mahoney, N. (2021), The role of selection markets, in 'Handbook of Industrial Organization', Vol. 5, Elsevier, pp. 389–426.
- Goldfarb, A. & Tucker, C. (2012), 'Shifts in privacy concerns', *American Economic Review* **102**(3), 349–53.
- Hestness, J., Narang, S., Ardalani, N., Diamos, G., Jun, H., Kianinejad, H., Patwary, M., Ali, M., Yang, Y. & Zhou, Y. (2017), 'Deep learning scaling is predictable, empirically', *arXiv preprint arXiv:1712.00409*.
- Iansiti, M. (2021), 'The value of data and its impact on competition', *Harvard Business School NOM Unit Working Paper* (22-002).
- 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.
- 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.
- 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.
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J. & Amodei, D. (2020), 'Scaling laws for neural language models', *arXiv preprint arXiv:2001.08361*.
- 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*.
- Mahoney, N. & Weyl, E. G. (2017), 'Imperfect competition in selection markets', *Review of Economics and Statistics* **99**(4), 637–651.
- 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.
- Matte, C., Bielova, N. & Santos, C. (2020), Do cookie banners respect my choice?: Measuring legal compliance of banners from iab europe's transparency and consent framework, in '2020 IEEE Symposium on Security and Privacy (SP)', IEEE, pp. 791–809.
- Miller, K., Sahni, N. S. & Strulov-Shlain, A. (2022), 'Sophisticated consumers with inertia: Long-term implications from a large-scale field experiment', *SSRN Electron. J.*.
- Miller, K. & Skiera, B. (2017), 'Economic damage of cookie lifetime restrictions', *Available at SSRN 3046254*.
- Misra, K., Schwartz, E. M. & Abernethy, J. (2019), 'Dynamic online pricing with incomplete information using multiarmed bandit experiments', *Marketing Science* **38**(2), 226–252.

- Neumann, N., Tucker, C. E., Kaplan, L., Mislove, A. & Sapiezynski, P. (2022), 'Data deserts and black box bias: The impact of socio-economic status on consumer profiling'.
- Nouwens, M., Liccardi, I., Veale, M., Karger, D. & Kagal, L. (2020), Dark patterns after the gdpr: Scraping consent pop-ups and demonstrating their influence, *in* 'Proceedings of the 2020 CHI conference on human factors in computing systems', pp. 1–13.
- Peltzman, S. (1981), 'The effects of ftc advertising regulation', *The Journal of Law and Economics* **24**(3), 403–448.
- Peukert, C. & Sen, A. (forthcoming), 'The editor and the algorithm: Recommendation technology in online news', *Management science* .
- Prince, J. T. & Wallsten, S. (2020), 'How much is privacy worth around the world and across platforms?', *Journal of Economics & Management Strategy* .
- Qiang, S. & Bayati, M. (2016), 'Dynamic pricing with demand covariates', *arXiv preprint arXiv:1604.07463* .
- Rafieian, O. & Yoganarasimhan, H. (2021), 'Targeting and privacy in mobile advertising', *Marketing Science* **40**(2), 193–218.
- 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.
- Schwartz, E. M., Bradlow, E. T. & Fader, P. S. (2017), 'Customer acquisition via display advertising using multi-armed bandit experiments', *Marketing Science* **36**(4), 500–522.
- 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.
- Strulov-Shlain, A. (2023), 'More than a penny's worth: Left-digit bias and firm pricing', *The Review of Economic Studies* .
- 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'.
- Thaler, R. H. & Sunstein, C. R. (2008), *Nudge: Improving decisions about health, wealth, and happiness*, Penguin.
- Tomaino, G., Wertenbroch, K. & Walters, D. J. (2022), 'Express: Intransitivity of consumer preferences for privacy', *Journal of Marketing Research* .
- 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.
- Veiga, A. & Weyl, E. G. (2016), 'Product design in selection markets', *The Quarterly Journal of Economics* **131**(2), 1007–1056.
- Wernerfelt, N., Tuchman, A., Shapiro, B. & Moakler, R. (2022), 'Estimating the value of offsite data to advertisers on meta', *Available at SSRN 4176208* .
- Westin, A. F. (2003), 'Social and political dimensions of privacy', *Journal of social issues* **59**(2), 431–453.



Woodruff, A., Pihur, V., Consolvo, S., Brandimarte, L. & Acquisti, A. (2014), Would a privacy fundamentalist sell their dna for \$ 1000... if nothing bad happened as a result? the westin categories, behavioral intentions, and consequences, *in* '10th Symposium On Usable Privacy and Security (SOUPS 2014)', pp. 1–18.

## A Experiment Appendix

Figure A.1: Example Recruiting Ad



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

(a) Accepting an offer price too low

! If you choose less than \$14.50, you might be selling the pen less than what's worth for you.

Will you sell your pen for \$14?

☒ Yes

☐ No

(b) Rejecting an offer price too high

! If you choose not to sell for a price more than \$14.50, you might not sell the pen even though you want to.

Will you sell your pen for \$16?

☐ Yes

☒ No

Figure A.3: The Multiple Choice List Procedure: Step 1

In this question, we are asking for your price to share **Posts**: your Facebook posts and feed history from the last month.

**If the computer chose any of the following prices, will you share your data?**

	Your choice	
	Yes	No
Will you share your <b>Posts</b> for \$50?	<input type="radio"/>	<input type="radio"/>
Will you share your <b>Posts</b> for \$60?	<input type="radio"/>	<input type="radio"/>
Will you share your <b>Posts</b> for \$70?	<input type="radio"/>	<input type="radio"/>
Will you share your <b>Posts</b> for \$80?	<input type="radio"/>	<input type="radio"/>
Will you share your <b>Posts</b> for \$90?	<input type="radio"/>	<input type="radio"/>
Will you share your <b>Posts</b> for \$100?	<input type="radio"/>	<input type="radio"/>

*Note:* This screenshot shows what a participant sees in the first step of our MPL procedure in the *high price range + forced choice* condition. In the opt-in condition, all the “yes” options are preselected; in the opt-out condition, the “no” options are pre-selected. If a participant gets randomized into the *low price range* condition, the prices listed will range from \$0 to \$50, instead of \$50 to \$100.

## Figure A.4: The Multiple Choice List Procedure: Step 2

(a) Scenario 1: participant says “yes” to some prices but not all

You said you are willing to share your Posts for \$90 but not for \$80, is there a more exact price that you are willing to share the data for?



(b) Scenario 2: participant says “no” to all prices listed

You have indicated that you will not share Posts for prices listed on the previous page. Can you tell us why?

☐ The price I am willing to share Posts was not listed. I am willing to share it for (please type the desired price without \$)

☐ I do not want to share Posts for any price

*Note:* The screenshots above show two examples of what a participant sees in the second step of our MPL procedure. In the first scenario, a participant agrees to share their posts data for \$90 or above but not for \$80 or below. In the second scenario, a participant chooses not to share their *posts* data for all the prices listed in the first step. In this case, we give them the option to indicate that their valuation for the *posts* data is infinity.

## 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: Baseline Variables

	Treatment						T-test stat	
	low: opt-out	low: active	low: opt-in	high: opt-out	high: active	high: opt-in	opt-in	p-value
<b>Number of participants</b>								
n	786	824	817	832	815	831		0.321
<b>Race (percentage)</b>								
White	0.838	0.8	0.819	0.802	0.801	0.812		0.457
Black	0.055	0.073	0.059	0.073	0.071	0.059		0.773
Asian	0.106	0.112	0.105	0.108	0.118	0.125		0.42
Other	0.034	0.045	0.049	0.056	0.052	0.047		0.012
<b>Gender (percentage)</b>								
Female	0.592	0.586	0.595	0.594	0.574	0.592		0.191
<b>Age</b>								
Age	39.398	37.885	38.359	38.864	39.124	38.526		0.146
<b>Income (\$)</b>								
Income	66724	67433	72080	70354	73156	66347		0.962
<b>Education (percentage)</b>								
College education and above	0.644	0.623	0.647	0.637	0.659	0.643		0.87
Some college	0.19	0.184	0.177	0.198	0.179	0.182		
Less than college	0.167	0.193	0.175	0.165	0.162	0.176		
<b>Facebook and Internet usage</b>								
Avg time spent on FB (h)	-0.051	0.061	0.01	-0.046	0.039	-0.027		0.755
FB membership duration (y)	0.037	0.028	0.095	0.019	0.023	0.014		0.712
Avg time spent on internet (h)	0.009	0.054	-0.013	-0.021	0.002	-0.016		0.872
Active user (percentage)	0.298	0.301	0.296	0.321	0.326	0.289		0.236
Purchase from FB or Instagram (times / mo)	-0.023	-0.032	0.003	-0.012	-0.02	-0.065		0.365
FB or Instagram ad click (times / mo)	-0.007	-0.024	0.01	-0.024	-0.002	0.042		0.019

Table B.3: Covariate Balance Across Treatments: Endline Variables

	Treatment						T-test stat	
	low: opt-out	low: active	low: opt-in	high: opt-out	high: active	high: opt-in	high: opt-in	p-value
<b>Number of participants</b>								
n	786	824	817	832	815	831		0.321
<b>Information available to the public</b>								
Info from the about page	0.692	0.704	0.689	0.702	0.714	0.745		0.457
Posts	0.375	0.396	0.381	0.375	0.404	0.397		0.773
Photos	0.427	0.428	0.414	0.401	0.415	0.438		0.42
Lists of likes	0.309	0.317	0.321	0.312	0.324	0.337		0.012
Friends and followers	0.469	0.527	0.504	0.493	0.482	0.51		0.962
Don't know	0.162	0.126	0.151	0.16	0.144	0.144		0.191
Other information	0.085	0.104	0.075	0.079	0.106	0.066		0.146
<b>Information available to Facebook advertisers</b>								
Name and email	0.725	0.718	0.704	0.69	0.733	0.729		0.755
Info from the about page	0.822	0.816	0.815	0.831	0.847	0.836		0.712
Posts	0.531	0.535	0.526	0.507	0.531	0.521		0.872
Photos	0.468	0.45	0.454	0.425	0.456	0.452		0.236
Lists of likes	0.711	0.708	0.71	0.69	0.739	0.71		0.365
Friends and followers	0.585	0.631	0.624	0.588	0.632	0.611		0.019
Don't know	0.109	0.103	0.115	0.105	0.109	0.105		0.353
Other information	0.038	0.038	0.037	0.035	0.022	0.022		0.484
<b>Participants looking up information</b>								
Looked up information when answering questions	0.019	0.023	0.021	0.023	0.028	0.031		0.891
<b>The type of information participants look up</b>								
How advertisers use data for targeting	0.008	0.007	0.004	0.01	0.011	0.011		0.659
What Facebook shares with advertisers	0.006	0.01	0.011	0.013	0.017	0.011		0.429
How much each data is worth	0.003	0.004	0.006	0.006	0.01	0.01		0.182
Other information	0.01	0.012	0.006	0.008	0.007	0.017		0.974
Frequency of encountering cookie banners/day	1.523	1.657	1.575	1.561	1.609	1.552		0.15

Figure B.1: Consistency between final price and MPL choice





Figure B.2: Data Supply Curves by Data Type (First-Round Answer Only)



## C Empirical Results Appendix

Table C.1: Treatment effects with OLS Specification

	WTA	log(WTA)	WTA	log(WTA)
Intercept	60.438 *** (0.723)	3.884 *** (0.022)	48.375 *** (0.760)	3.556 *** (0.024)
Price Anchor = Low	-16.522 *** (0.762)	-0.536 *** (0.023)	-16.522 *** (0.762)	-0.536 *** (0.023)
Default = Active	1.763 + (0.927)	0.042 (0.029)	1.763 + (0.927)	0.042 (0.029)
Default = Opt-in	4.185 *** (0.929)	0.105 *** (0.029)	4.185 *** (0.929)	0.105 *** (0.029)
Likes			6.964 *** (0.381)	0.226 *** (0.012)
About Me Page			14.833 *** (0.421)	0.409 *** (0.013)
Posts			16.283 *** (0.429)	0.446 *** (0.013)
Friends and Followers			22.236 *** (0.467)	0.562 *** (0.014)
Num. Obs.	25140	25140	25140	25140
R2	0.060	0.071	0.111	0.108

+ p < 0.1, \* p < 0.05, \*\* p < 0.01, \*\*\* p < 0.001. The outcomes are top-coded at \$100; standard errors are clustered at the participant level.

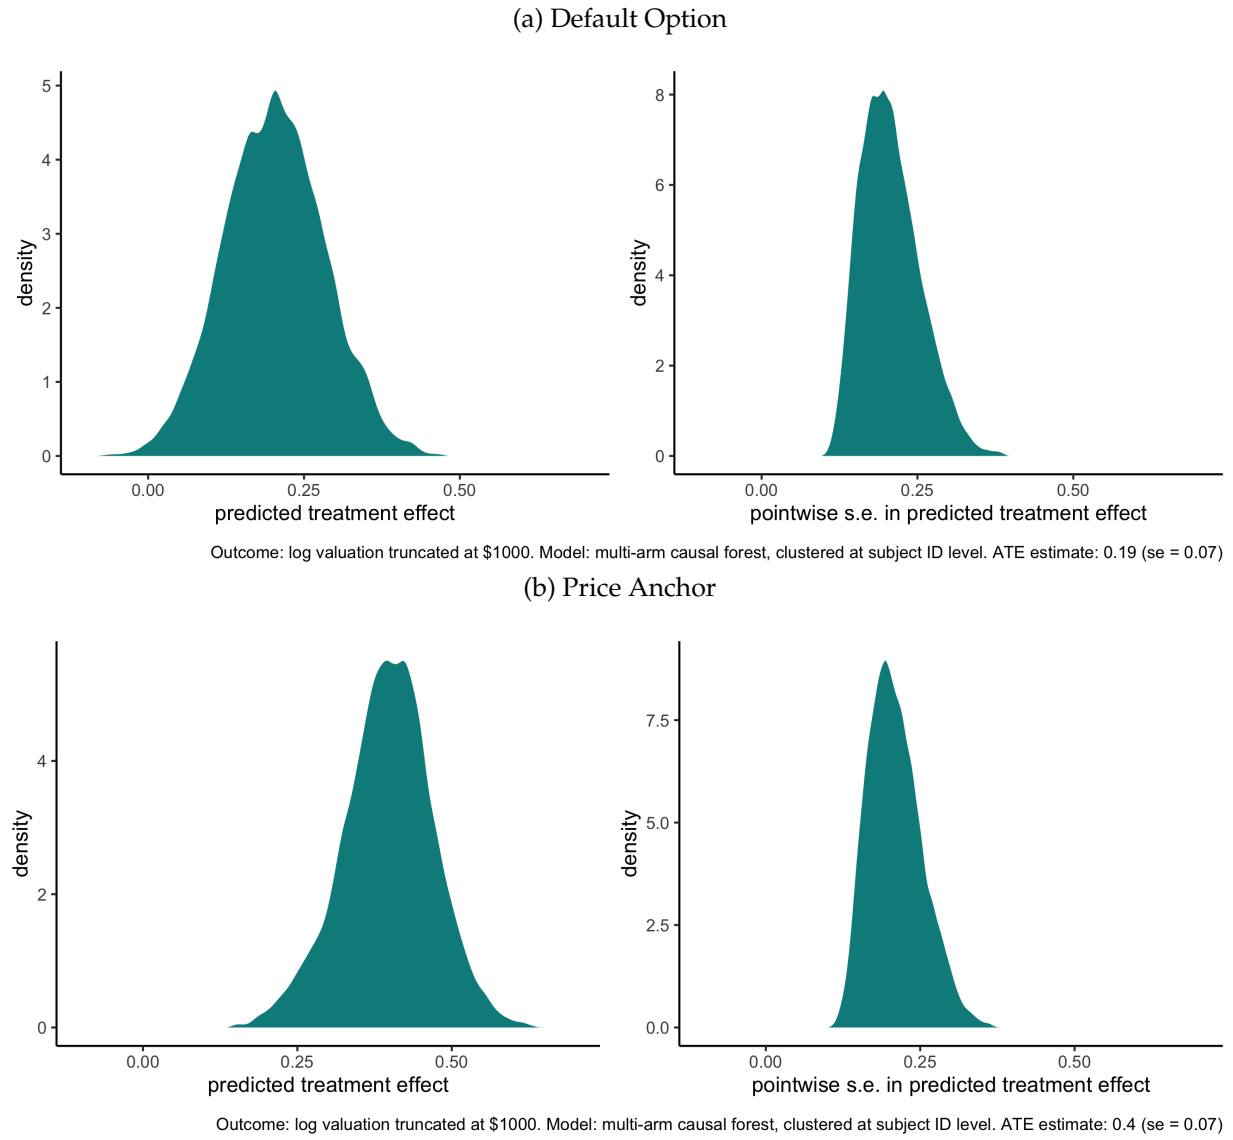


Table C.2: Treatment Effects with Additional Controls

	WTA	WTA	WTA	log(WTA)	log(WTA)	log(WTA)
Intercept	63.884 *** (0.876)	58.839 *** (0.802)	71.972 *** (1.048)	3.985 *** (0.025)	3.848 *** (0.024)	4.194 *** (0.029)
Price Anchor = Low	-16.112 *** (0.947)	-4.917 *** (0.991)	-16.155 *** (0.938)	-0.525 *** (0.028)	-0.216 *** (0.029)	-0.526 *** (0.027)
Default = Active	2.377 * (1.139)	1.930 + (1.029)	2.607 * (1.131)	0.058 + (0.034)	0.046 (0.031)	0.064 + (0.033)
Default = Opt-in	5.178 *** (1.147)	2.978 ** (1.026)	5.315 *** (1.135)	0.135 *** (0.034)	0.076 * (0.031)	0.139 *** (0.034)
Practice round deviation		0.695 *** (0.021)			0.019 *** (0.001)	
Believes data is available			-12.304 *** (0.861)			-0.317 *** (0.024)
Num. Obs.	25140	25140	25140	25140	25140	25140

+ p < 0.1, \* p < 0.05, \*\* p < 0.01, \*\*\* p < 0.001. The outcomes are top-coded at \$100; standard errors are clustered at the participant level.

Figure C.1: Heterogeneous Treatment Effect Estimates and Standard Errors: Multi-Arm Causal Forests



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.1 shows that the default estimates from the multi-arm causal forest are slightly biased upwards while the price anchor effect is biased downwards.