

Information Frictions and Heterogeneity in Valuations of Personal Data

Avinash Collis* Alex Moehring[†] Ananya Sen[‡] Alessandro Acquisti^{§¶}

November 2021

Abstract

We investigate how consumer valuations of personal data are affected by real world information interventions. Proposals to compensate users for the information they disclose to online services have been advanced in both research and policy circles. These proposals are hampered by information frictions that limit consumers' ability to assess the value of their own data. We use an incentive compatible mechanism to capture consumers' willingness to share their social media data for monetary compensation, and estimate distributions of valuations of social media data before and after an information treatment. We find evidence of significant dispersion and heterogeneity in valuations before the information intervention, with women and Black and low income individuals reporting systematically lower valuations than other groups. After an information intervention, we detect significant revisions in valuations, concentrated among individuals with low initial valuations. Dispersion and heterogeneity in valuations across these demographic groups decrease, but persist after the information intervention. The findings suggest that strategies aimed at reducing information asymmetries in markets for personal data may increase consumer welfare. At the same time, the findings highlight how consumer valuations of personal data are only in part influenced by market information.

*McCombs School of Business, University of Texas at Austin, avinash.collis@mcombs.utexas.edu

[†]Sloan School of Management, Massachusetts Institute of Technology, moehring@mit.edu

[‡]Heinz College of Information Systems and Public Policy, Carnegie Mellon University, ananyase@andrew.cmu.edu

[§]Heinz College of Information Systems and Public Policy, Carnegie Mellon University, acquisti@andrew.cmu.edu

[¶]A.C, A.M, and A.S are listed in alphabetical order. This research was funded by research grants from the McCombs Research Excellence Fund; the Center for Technology, Innovation and Competition at the University of Pennsylvania, and the NET Institute (www.NETinst.org). The study received IRB approval from Carnegie Mellon University, Massachusetts Institute of Technology, and the University of Texas at Austin. The study was preregistered in the American Economic Association Registry for randomized control trials (RCT ID: AEARCTR-0006727). We thank Dana Budzyn from YouGov and Enoch Liang and Mark Mao from the Data Dividend Project for helping us recruit participants for our study. We thank Idris Adjerid, Laura Brandimarte, Erik Brynjolfsson, Dave Holtz, Xiang Hui, Tesary Lin, Catherine Tucker, Glen Weyl, Heather Yang, and participants at the University of Minnesota, University of Nottingham, CODE@MIT, Carnegie Mellon University (Heinz College), IIM (Ahmedabad), PEEEX Lab, and the MIT Digital Economics Working Group for helpful comments.

1 Introduction

Personal data generate significant value for digital platforms, as source of revenue or as input facilitating algorithmic targeting (Elsaify and Hasan, 2020). To address concerns over the potentially unequal allocation of value between data holders (the platforms) and data subjects (the users), proposals to share data-based revenues with consumers have emerged in both policy and academic circles.¹ Gavin Newsom, the Governor of California, has proposed data “dividends” to compensate consumers who create online footprints (Ulloa, 2019; Au-Yeung, 2019). Academics have argued that users’ online data should be viewed as “labor” and compensated accordingly (Arrieta-Ibarra et al., 2018). And a growing number of companies have pledged not to sell personal data to ad brokers.²

Markets for personal data have a storied history in the literature. Kenneth Laudon first proposed, in 1996, “national information markets” through which consumers could trade rights over the usage of their data (Laudon, 1996). The design of any framework for data markets or data dividends, however, may face challenges arising from information frictions. In data markets, consumers provide companies access to their data through, potentially, a variety of selling mechanisms including negotiated prices (Spiekermann et al., 2015; Yang, 2020b). Whereas platforms could quantify the value of user data to the platforms, users face a more significant hurdle in pinpointing the value of data to *themselves*. Empirical research on consumer valuation of data and data privacy across disciplines has highlighted that individuals’ valuations of personal information are highly uncertain and—crucially—marred by endemic problems of asymmetric information (Tomaino et al., 2021; Acquisti et al., 2016). Not only do consumers rarely know how their data are used, but they also often lack information on the value that other entities extract from their data (Shiller, 2021), or the costs they may ultimately bear when their data are misused. Thus, when interacting with platforms or data intermediaries, consumers may be uninformed and fail to extract optimal levels of surplus

¹We use “users” and “consumers” interchangeably to designate individuals who utilize digital platforms.

²As an example, Neeva is an ad-free search engine that does not sell user data to third parties launched by the ex-SVP of Google ads.

from those markets. Unlike established markets for goods (such as cars, or shares of a company), where consumers have access to a plethora of information, markets for personal data could consist of transactions that are unique to an individual, exacerbating consumers' costs to learning their own preferences (Cao and Zhang, 2021). In such contexts, data markets may fail to achieve the very purpose that, according to proponents, they ought to fulfill: enabling a more equitable allocation of the benefits extracted from consumer data. Such failure would be particularly problematic if the inability to determine fair valuations was disparately distributed across different socioeconomic groups.

Whether those information frictions can be offset with information treatments (such as making consumers better aware of the value that others extract from their data, or the costs that may accrue when their data are compromised) is an open question. In this manuscript, we study how consumers' valuations of personal data change under the influence of information interventions, based on actual data points involving personal information, and how valuations ex ante and ex post information treatments vary across demographic groups. In our data market setting, we use an incentive compatible mechanism and a mixed between- and within-subjects experimental design to estimate participants' willingness to share personal data for monetary compensation (their Willingness to Accept, or WTA) before and then again after an information treatment. We focus on social media data, and capture the distribution of compensations participants require to share the entirety of their Facebook profile data with the researchers (those data include public profile information, pictures, and private messages).

We randomly assign participants to one of two information treatments, each providing participants with information from actual scenarios involving users' social media data. We rely on estimates of data valuation from two scenarios that have become central to the debate over data dividends and data privacy: the value companies extract from utilizing user data, and the compensation users can receive from data holders when their data is abused. Thus, in one condition, participants are provided accurate information summarizing Facebook's projections of revenues per North-American profile. In the other condition, participants are

provided accurate information summarizing the monetary compensation that some Facebook users received following the improper harvesting of user data. In each case, we theorize that such exposure to ‘market’ information may reduce value uncertainty regarding personal data, and thus affect participants’ own valuations.

We recruit experimental participants from two groups—one nationally representative, and one expected to be, on average, more data and privacy-conscious than the representative sample. The first group is a representative sample of U.S. internet users recruited in collaboration with YouGov (YouGov sample). The second group, based on our collaboration, comprises members of the Data Dividend Project (DDP), a data advocacy group started by former Democratic presidential candidate Andrew Yang (DDP sample). Members of the DDP are interested in ensuring that technology companies share a part of their revenue when they monetize data and are more likely to believe in digital privacy as a fundamental right (Yang, 2020a). The two samples allow us to compare whether and how information treatments differentially affect individuals who are likely heterogeneous in their fundamental views on personal data.

The objective of our analysis is not to pinpoint or estimate a “true” value of data. First, prior work has firmly established that personal data valuations (as well as associated privacy concerns over data) are context dependent (Xu and Zhang, 2020) and vary based on the reference frame (Acquisti et al., 2016). Second, the experimental treatments—which inform participants about Facebook’s projected revenues per user profile and users’ monetary compensation for breach of their profiles—are contiguous to, but not congruent with, the valuation captured in the experiment (participants’ willingness to accept compensation to share their profiles with researchers). We are instead interested in the distribution of data valuations before and after information treatments that update participants’ beliefs about possible market valuations of their social media profiles.

Before the information treatment, we find evidence of significant dispersion and heterogeneity in valuations in the nationally representative (YouGov) sample. Valuations are bimodal,

with a large mass of respondents clustered at less than \$250 and another mass reporting valuations of at least \$10,000. Not only are ex ante valuations highly dispersed — there is substantial heterogeneity in valuations by demographic traits, with historically marginalized groups reporting significantly lower valuations.³ For instance, the distribution of valuations for White users first-order stochastically dominates the distribution for Black users; the distribution for male users first-order stochastically dominates the valuation distribution for female users. We find that the racial and gender divide persists also after controlling for education, income, privacy beliefs and Facebook usage suggesting that certain demographic groups appear to systematically undervalue their data relative to others. Moreover, the direction of the gender divide in user valuations in our study are at odds with the direction of the divide in the market for data, where female data have shown to command a higher price in online advertising markets (Lambrecht and Tucker, 2019). Additionally, in our data, we find that valuations of data decrease in unison with decrease in income, with lower income individuals asking as little as half the amount of money to share their data with researchers, compared to higher income individuals. We find that the distribution of valuations in the DDP sample is similarly bimodal, but first-order stochastically dominates the distribution of the nationally representative sample. Moreover, there is qualitatively similar heterogeneity across demographics in the DDP sample. The broad consistency in results across the two samples provides external validity and credibility to our estimates.

Following the information treatments, close to one third of participants in both samples revise their data valuations. In the YouGov sample, 28.6% of individuals revise their valuations following the treatments. The probability of revision is highly asymmetric, with individuals with a $WTA < \$400$ (below the dollar amount mentioned in the treatment) driving the effect with a 53% probability of revision. Furthermore, 98.2% of the individuals who update their valuations do so by revising up to a higher valuation. Results for the DDP sample are very similar: 29.4% of participants revise their valuations, again predominantly driven by those

³Historically marginalized groups are those that might be discriminated against in social, cultural, or economic life. Examples of such marginalized populations include groups excluded due to race, gender identity, or sexual orientation—among others (Sevelius et al., 2020).

with baseline WTA < \$400 revising up. The asymmetric revisions suggest that valuations post treatment are not merely a function of anchoring within a Bayesian framework, and that other, more subtle, factors are in fact at play (Section 4). In both samples, the provision of information leads to a reduction in, but not elimination of, dispersion in data valuations. The reduction takes the form of increasing valuations by low valuation individuals—in which women, low income, and Black participants are over-represented. Taken together, both the ex ante distributions of valuations and the ex post variations suggest that information frictions related to privacy could partially explain low personal data valuations in the literature (Athey et al., 2017). Furthermore, the asymmetric revisions suggest that valuations post treatment are not *only* a function of anchoring or Bayesian updating, and that other, more subtle, factors are in fact at play (Section 4). In fact, as dispersion in valuations persists following information treatments, it is reasonable to conclude that consumer valuations of personal data are only in part influenced by market information. Consistent with recent work on privacy decision making (Lin, 2020), our analysis (Section 4) suggests that consumers’ valuations of their data are the composite of *objective* or instrumental factors—such as knowledge of the fair market value of one’s data, which information interventions can affect—and inherently and deeply subjective or intrinsic ones—such as individuals’ personal stances on data privacy, or the psychological harm different individuals associate with violations of their data (Calo, 2011).

Our findings have direct policy implications. Although some policy makers and some scholars have viewed data markets as means to compensate consumers fairly for their data as well as reduce economic inequality across demographic groups (Feygin et al., 2021), our results suggest that information frictions may impair the ability of those markets to build a more equitable data economy. Our findings suggest, however, that strategies aimed at reducing information asymmetries may be helpful to consumers, especially those from historically marginalized groups and would aid the functioning of data markets. They may allow consumers to negotiate higher paybacks and extract more surplus in personal data markets, moving closer to

existing market valuations and increasing consumer welfare from data markets. Various recent regulatory efforts in the privacy field have, in fact, aimed at addressing and reducing informational asymmetries (Shiller, 2020). As noted, provisions in the draft regulations of the CCPA stipulate conditions under which businesses should share with consumers information on their data valuation methods; similarly, the EU General Data Protection Regulation (GDPR) requires firms to disclose how collected consumer data is used. At the same time, the finding that a demographic divide in valuations persists after the information treatments suggests that information campaigns might be useful, but are not a silver bullet in aiding marginalized socio-economic groups to extract surplus from personal data markets. Our findings also have managerial implications for strategy around pricing and marketing. A number of companies are aiming to establish data markets and becoming data brokers (e.g., YouGov, Permission, Brave, etc.) using a variety of pricing strategies (Yang, 2020b). Additionally, on the other end of the spectrum, some firms are pledging not to collect or sell data generated by users' online activity (Holtrop et al., 2017). Our results help firms understand how to price their product, in line with the recent literature (e.g., Huang et al. (2020); Cao and Zhang (2021)), taking this heterogeneity and frictions into account, as well use marketing strategies to inform individuals about the value of their data. The potential disparity in data and privacy valuations across demographics has started to be recognized by the industry, leading to ventures such as Streamlytics — the largest first-party provider of African-American data (Streamlytics, 2021).

Related Work This study contributes to a few strands of the academic literature. First, we contribute to the literature on economic valuations of personal data and online privacy. Over the years, several studies have investigated both individuals' willingness to pay to protect personal information (WTP; for instance, Beresford et al. (2012)), and individuals' willingness to accept payments to share it (WTA; for instance, Danezis et al. (2005) and Hui et al. (2007)). Related to our study, Benndorf and Normann (2018) studied the WTA of college students to divulge their Facebook contacts and timeline details, finding a median WTA of 25 Euros. Athey et al. (2017) find privacy-concerned individuals willing to divulge personal information

in exchange for small amounts of money or rewards. The findings of this body of work support the notion that individuals’ valuations of personal data reflect a combination of factors—from rational privacy calculus to heuristics, cognitive biases, and information asymmetries (Acquisti et al., 2015). Accordingly, recent work has started disentangling the extent to which preferences for privacy are influenced by “intrinsic” (subjective) versus “instrumental” (more quantifiable and objective) factors (Lin, 2020). Our paper differentiates from these studies by focusing on information frictions as a key factor impacting data valuations, and by analyzing whether information provision can reduce the dispersion in data valuations. Moreover, we focus on and provide clear regularities in systematically lower valuations by marginalized groups.

Second, our paper is related to studies that explicitly focus on how privacy concerns (or, in a few cases, valuations) vary across demographic characteristics and different contexts. Prince and Wallsten (2020) measure the (hypothetical) WTA for different types of data and for populations across the world. Christofides et al. (2012) analyze privacy attitudes of adults and adolescents towards Facebook activity. Hoy and Milne (2010) analyze gender differences in privacy beliefs associated with Facebook use. Our work differentiates by focusing on *incentive compatible* valuations for the entire stock of Facebook data, including private messages and photos, with a representative sample of US internet users that allows us to investigate regularities across demographic groups and highlight the potential role of information provision in this context. Our results show that data conscious individuals in the DDP sample, while having higher valuations, revise their valuations at the same rate relative to those in the YouGov sample suggesting that information frictions are an important issue for data valuations, even for engaged users.

Our results are related to the literature that attempts to provide information treatments to increase awareness and make individuals take more informed decisions in a variety of contexts. Guriev et al. (2020) analyze the impact of providing fact-checked political information on an individual’s beliefs and propensity to share misinformation. Alesina et al. (2021) analyze the

impact of providing economic statistics to Black and White respondents on their perceptions of why racial inequities persist. In finance, Beshears et al. (2009) provide financial information to consumers so that they can make better daily financial decisions. These results are also broadly related to studies that analyze the impact of salience in privacy-related information on consumer behavior (Beke et al., 2018; Adjerid et al., 2013; Tsai et al., 2011). This literature, together with our results, suggests that there could be a large payoff to information campaigns by policy makers and data advocates. Finally, the manuscript is also related to a growing body of work on data markets and data propertization in the economics and law literature (Laudon, 1996; Schwartz, 2003; Arrieta-Ibarra et al., 2018). In the context of the psychology of data ownership (Spiekermann et al., 2012), Tomaino et al. (2021) highlight that users may underestimate their privacy valuations since transactions with companies happen through barter rather than money.

2 Experimental Design

We conduct a pre-registered incentive compatible online experiment in which we solicit consumers' willingness to accept (WTA) monetary compensation to share their Facebook data. We recruit participants from two populations. First, we partner with YouGov and recruit from its population of US based adult respondents (YouGov sample). YouGov screens out respondents for our study so that they are representative of the US internet population in terms of age, gender, region, race, and education based on the US Census Current Population Survey (2018). Second, we recruit from members of the Data Dividend Project (DDP) through email solicitations sent by the DDP to its entire member base, inviting them to take part in our study (DDP sample). Participants from both populations were required to have a Facebook account in order to participate in the study.⁴

Participants are provided a link to access an online survey. First, they are asked to

⁴Facebook membership and usage is verified and was provided by YouGov to us. For the DDP sample, respondents self-report using Facebook.

provide demographic information.⁵ Next, respondents' WTA valuations are captured using a Becker–DeGroot–Marschak (BDM) mechanism (Becker et al., 1964) (BDM mechanisms are common in recent literature estimating welfare effects of social media and the value of online services (Allcott et al., 2020; Brynjolfsson et al., 2019)). Respondents are asked for the minimum amount of money they would require to share the entirety of their Facebook data with the researchers.⁶ The data include posts, photos, private messages, likes, and comments (See Figure A.5 in the Online Appendix). Facebook offers a simple way to download a copy of a user's data, and the process is explained to the participants (Facebook, 2021b). To make responses incentive compatible, we explain the BDM mechanism to respondents. They are informed that at the end of the study a random amount of payment will be generated, and if the randomly generated payment is greater than the minimum WTA entered by the respondent, the participant will be asked to upload their data in exchange for the payment.

After baseline valuations are elicited, respondents are randomly assigned to one of two information treatments: revenue treatment and settlement treatment. Across treatments, the monetary amount associated with the value of a Facebook's user data is held constant. In the revenue treatment, participants are informed that Facebook is expected to earn around \$400 per North American user in the next three years. This information is based on Facebook's 10-K filings in January 2021 (Facebook, 2021a). In the settlement treatment, participants are informed that each affected Facebook user in a Facebook data settlement in 2020 was paid around \$400 (Sun-Times Staff, 2020). (see Figures A.6, and A.7 in the Online Appendix). This information is based on a lawsuit that was settled by Facebook in Illinois. We provide a link to the source of the information in case an individual wants to access more details. We focus on a revenue frame, and a data misuse frame as those two dimensions have become salient in public and academic discourses around data and privacy (Elsaify and Hasan, 2020; Feygin et al., 2021). Similar to the methodology and approach in Hjort et al. (2021), we use

⁵For the YouGov sample, demographic data of participants are directly collected by YouGov and shared with us.

⁶Our approach to measure WTA in dollars is consistent with a recent study by Tomaino et al. (2021) that shows that individuals understand data values better in dollars than in return for a product like in a barter system.

these treatments as one way to assess whether real-world information impacts user valuations, and do not regard them as the ultimate “ground truth” on data valuations. Although our information treatments are accurate summaries of true information, we are careful to not tell users that this is what their data are *exactly* worth to Facebook or *exactly* reflective of the costs if their data are used improperly. In our incentivized, anonymous setting that deals with sensitive information about personal data, we expect experimenter demand effects to be minimal (Haaland et al., 2020) as has been evidenced in the experimental literature (De Quidt et al., 2018). Additionally, a potential taste for consistency would work against our information treatments impacting valuations (Falk and Zimmermann, 2013). We carry out a placebo check to demonstrate clearly that the effects we detect are not an artifact of experimenter demand.

Following the information treatment, we allow individuals to revise their valuations. These revised valuations are considered for the BDM lottery, and hence are incentive compatible. We also ask respondents to explain why they revised their valuations or why they did not revise their valuations, in an open ended text box. Then, we ask endline questions related to Facebook usage, and participants’ views about data privacy.

Finally, all respondents are entered in the BDM lottery. If the randomly generated payment is greater than the minimum WTA entered by the respondent, the participant is asked to upload their data in exchange for the randomly generated payment. The data upload page shows up for those respondents who qualify. Participants receive their payment if and only if they upload their Facebook data. We verify that uploaded files are authentic by checking their metadata (directory names, sizes, and formats). If the randomly generated payment is less than the minimum WTA, the selected respondents do not receive any payment and they do not have to upload their data. We ask comprehension questions to verify respondents’ understanding of the BDM mechanism. Respondents are unaware of the payment distribution, in line with best practices for BDM research (Allcott et al., 2020).

3 Baseline Results

We conducted our study in June-July 2021. The study was completed by 4,149 participants from the YouGov sample (this number includes respondents who passed the comprehension tests and provided both baseline and revised valuations).⁷ Baseline covariates for the YouGov sample (as well as the DDP sample) are shown in the Online Appendix (Table A.1). Roughly half of respondents use Facebook less than 30 minutes per day; the average account was created in 2009.

The distribution of baseline valuations is plotted in Figure 1. This figure provides a few takeaways. First, valuations are highly dispersed — in fact, their distribution is bi-modal, with bunching of WTA at low dollar values (less than \$250) or at very high values (at least \$10,000). For ease of representation, we truncate valuations at \$10,000 in the histogram (at the 75th percentile). The CDF of these valuations is plotted in Figure 1 (median WTA is \$750). The spikes in the distribution take place around whole numbers (e.g., \$1,000, \$5,000 etc), suggesting that users utilize heuristics while attempting to determine an otherwise highly uncertain value. The figure demonstrates how focusing on only single summary statistic, as much of the prior work has done, hides substantial heterogeneity in valuations. We interpret extremely high valuations as an expression of the respondents’ unwillingness to part with their data. All the respondents displayed an understanding of the BDM mechanism by correctly answering the comprehension questions; hence, they knew that an extremely large WTA would make it unlikely that they would have to upload their data.⁸

Valuations are also highly heterogeneous across fundamental dimensions of race, gender, and income. Figure 2 shows the CDF of valuations across different racial groups (see Online Appendix Figure A.2 for additional race and ethnicities). The racial divide in valuations across

⁷We discuss results from the DDP sample and how they compare to the YouGov sample in Section 5.

⁸This point is reinforced by the observation that 14 of the 18 participants (approximately 80%) selected to upload their Facebook data based on their BDM bids did upload their data and subsequently received payment in line with the BDM draw. In this sample of data uploads, we do observe individuals making their private messages and pictures available. Given the small size of this sample, we cannot carry out further heterogeneity analysis based on the size of the files across different data dimensions.

White and Black individuals is stark.⁹ Black Facebook users value their data significantly less than White users, as the distribution of valuations for White respondents first-order stochastically dominates that of Black respondents (median WTA for a White user is \$1,000, whereas it is \$500 for a Black user). The WTA for Hispanic Facebook users is very similar to the distribution for Black users. Asian users value their the data the most, although their sample size is small (at about 3%, it mirrors the proportion in the US population).

Women value their data less than men (Figure 2), with the WTA of male respondents first order stochastically dominates the distribution of WTA of female respondents (median valuation for female respondents of \$558 relative to \$1,000 for men). To test whether race and gender differences are driven by income, education, or patterns in Facebook usage, we analyze the logarithm of the valuations within a regression framework in Table 1. As can be seen, the descriptive results hold in univariate regressions for gender (column (1)), race (column (2)), gender and race together in column (3). We find that racial and gender differences also persist after adjusting for income, education, age — as well as Facebook usage and privacy beliefs — in column (4).

We also look at income differences in valuations. The differences are stark and consistent: individuals in higher income groups value their data significantly more. Interestingly, individuals who decide not to disclose their income are the ones who value their data the most across all income groups, consistent with the interpretation that these individuals are very high income individuals and/or value their income data and do not want to disclose it.

While we do not focus on a single dollar valuation in our analysis, we examine the plausibility of our baseline valuations in the context of the literature. We focus on the median valuation (\$750) to account for outliers. This valuation corresponds to the entire "stock" of an individual's Facebook data.¹⁰ The median respondent in our sample has been using Face-

⁹At any dollar value on the horizontal axis, the distribution plots the share of people with valuations less than that dollar value. This implies that a CDF shifted towards the right will have higher valuations associated with it.

¹⁰In a pilot, we analyzed the WTA for the entire stock of Facebook, Instagram, and Twitter data. The median valuations for Facebook and Instagram were higher than Twitter, potentially due to more personal information being available on Facebook and Instagram.

book for 12 years, since 2009. This leads to a median value of \$73 per year and a monthly valuation of \$6.1 per month for their data. Putting these values in the context of prior literature, we see that our estimates exceed those in existing studies (Prince and Wallsten, 2020; Benndorf and Normann, 2018; Lin, 2020). Relative to Lin (2020) and Benndorf and Normann (2018), the data we asked for is larger in quantity as well as sensitivity (because it consists of participants’ entire Facebook data, including personal messages, since they first created the account). Prince and Wallsten (2020) find, within a hypothetical setting, that the average US respondent has a WTA of \$5 per month to let Facebook share information from texts sent using Facebook Messenger. In contrast, our incentive compatible setting presents participants with an actual possibility of having to share all their Facebook data, likely leading to higher (and potentially more realistic) valuations.¹¹

Overall, the bi-modality and heterogeneity in valuations highlight divides across demographic groups. The lower data values for under-represented groups suggest that such data dividends based on individual valuations could further exacerbate existing inequalities. These findings are relevant to the scholarly and policy discussion around using data dividends to create an equitable digital economy by reducing inequalities across groups.

4 Information Treatments and Updating Behavior

4.1 Baseline Treatment Effects

Our baseline valuations offer insights over individuals’ reasoning around the value of their digital data, and highlight significant dispersion and heterogeneity. Our information treatments test the extent to which the provision of objective information on data market valuations can influence individuals’ WTA and reduce dispersion and heterogeneity in valuations across demographic groups.

¹¹Previous research on valuing digital goods finds that incentive compatible valuations are significantly higher than hypothetical valuations (Brynjolfsson et al., 2019; Athey et al., 2017).

We first present aggregate results across the two experimental conditions.¹² Then, in Section 4.2, we discuss differences across the treatments.

Figure 4 shows that individuals revise their valuations in response to information interventions—but do so asymmetrically. As a reminder, the dollar amount mentioned in both the settlement and the revenue treatment is \$400.¹³ The majority of participants who revise have baseline WTAs lower than \$400. Towards the right end of the distribution, there is no difference between the WTA and revised WTA. This holds even when we extend the distribution to \$10,000 (as seen in the Online Appendix in Figure A.1). Focusing on the individuals with baseline WTA less than \$400, we note that 98.2% revise their valuations upwards. This makes the overall distribution of the revised valuations less dispersed than the distribution of the initial baseline valuations. In particular, after the treatments, the proportion of individuals with a WTA of \$400 or more increases from 60.9% to 70.1%.¹⁴

We find that women are more likely (by 5%) to update their valuations in response to the information treatment than men (column (3), Table 2). Similarly, Black people are more likely to update their valuations than White. The effect is similar for low vs. high income, measured at the \$50K annual income threshold, with low income individuals more likely to revise. These results suggest that information treatments can lead individuals to reassess their valuations, and that providing actual market information can reduce dispersion and heterogeneity in valuations. Given that disadvantaged groups respond significantly more to these interventions, we conclude that information frictions do play a role in the ex ante gaps in WTA for digital data across these demographic groups of interest. In fact, while the information interventions help marginalized groups largely because there are more of them below the \$400 threshold, column (4) shows that the probability to revise valuations persists even after controlling for initial WTA. This suggests that participants from these groups

¹²We provide evidence that the randomization worked as expected through a randomization check (Table A.3) and series of balance tests (Table A.4).

¹³In a pilot study, we found that different dollar values associated with the information treatments lead to qualitatively similar results. Results available from the authors upon request.

¹⁴The number of individuals whose valuations were exactly \$400 increases from 0.9% of the sample to 7.3% post treatment.

appear to value the information treatment more, relative to others. It is important to note that the asymmetric updating outcome of the informational treatments cannot be explained as a simple reference point effect within a Bayesian setting. A reference point effect would have implied a symmetric revision of valuations, consistent with a Bayesian framework (we would have witnessed users with valuations above \$400 revise their WTA downward, in addition to those whose with valuations lower than \$400 revising upward). Indeed, these results suggest that an individual’s data valuations are driven in part by objective information, and in part by subjective beliefs about data and data privacy.

In fact, though the information treatments increase the valuations of marginalized groups, dispersion and heterogeneity in valuations persist *ex post*. Objective information alone does not eliminate differences in valuations. This suggests that data valuations might comprise both objective factors (including, for instance, market information about transactions involving personal data) and subjective factors (e.g., views on privacy). To examine this further, we utilize participants’ self-reported beliefs about data and privacy. We analyze how an individual’s baseline valuation and propensity to revise varies with their beliefs about privacy being a fundamental human right and the ability of the free market to value data correctly.¹⁵ We code each variable as one if the individual either agrees or strongly agrees with a certain statement. In column (1) of Table 2, we find that individuals who think that privacy is a fundamental human right value their data significantly more. Similarly, those who think that the free market provides an appropriate degree of privacy protection value their data less. Consistent with the intuition of valuations being driven by a composite of objective and subjective factors, individuals who think that privacy is a fundamental human right are less likely to revise their beliefs, whereas market-oriented individuals are more likely to revise (column (3)). Finally, in column (4), we show that the higher propensity to revise valuations for women and Black participants persists even after controlling for baseline valuations.

To ensure that this updating behavior is not a simple artifact of experimenter demand,

¹⁵The exact statements used in the endline survey are: (1) Privacy is a fundamental human right and (2) I trust that the free market leads to appropriate privacy protection.

in that individuals update simply because we ask if they would like to revise their valuations (following the information treatment), we carry out a robustness check. We run an additional online study (N=251) where we do not give individuals any information about valuations but simply ask them if they would like to revise their valuations (without any information treatment) to see if they are confident in their stated valuations. In this placebo check, we find that only 2.3% of the individuals (6 participants) revise their valuations. The revision probability in our main analysis is about 12 times higher with clear regularities in line with the information provided and our placebo estimates are significantly lower than estimates of similar checks in the literature (Allcott and Taubinsky, 2015).

4.2 Differences across Information Treatments

Next, we analyze whether it is in fact the information presented in the treatments that led individuals to update their valuations. For individuals with a valuation below \$400, both treatments are very effective: the settlement treatment led 55% of individuals to update, while the revenue treatment induced 49% of individuals to update their baseline valuations. Since both treatments refer to the same monetary value (\$400), the mechanism behind the updating of valuations might emanate from how individuals view the information provided to them.

To understand this, we analyze the answers to the free-form question asked immediately after respondents were given the chance to revise their valuations. We train an algorithm to cluster responses into groups and the full details of this algorithm can be found in the Online Appendix. The algorithm produced four clusters, and we look at the relative distributions of responses in those clusters across treatment groups (Figure 5).¹⁶ The four clusters (in Tables A.5-A.8) are a "revenue cluster" that contains responses discussing the value of one's data; a "data use cluster" that contains responses discussing how data is used by firms; a "careful sharing cluster" that contains responses from individuals who state they are very careful about

¹⁶Brief and uninformative responses (shorter than 20 words) were not used in this analysis.

the information they share online; and an "extreme privacy cluster" that contains responses emphasizing the sensitivity of their data and an unwillingness to part with it except at extreme prices. The results in Figure 5 show that the distribution of responses across clusters are different between the two information treatments. First, we find that those in the revenue treatment are substantially more likely to be in the revenue cluster ($p < 0.01$). This suggests that the revenue treatment induces individuals to revise by updating their beliefs about the value of their data to the firm, which also impacts their own user valuations. The settlement treatment, however, induces more individuals to be members of the extreme privacy clusters ($p < 0.01$). This suggests that the settlement treatment induces respondents to revise their valuations by updating beliefs about the potential harms of sharing their data. Taken together, we find that although both information treatments induce a substantial share of respondents to revise their valuations, the two treatments work by updating beliefs on different aspects of a user's value of data.¹⁷

The baseline updating results also show that the settlement treatment led to a slightly higher increase in updating relative to the revenue treatment, even though the absolute increase across the two treatments was high (55% vs. 49% with $p < 0.01$). To understand the quantitative difference in revision across the two treatments, we conducted a post-hoc survey on Amazon Mechanical Turk ($n=250$). We test whether this difference might arise due to the fact that individuals are more uncertain about settlement amounts relative to how much revenue Facebook earns. We asked subjects to predict the monetary values associated with Facebook revenues as well as the data settlement lawsuit. We found that there was significantly more uncertainty about the settlement treatment, consistent with the hypothesis that the settlement treatment induced larger updates in individuals' beliefs about the value of their data. In particular, subjects were 10.5 percentage points more likely ($p=0.002$) to respond saying "I don't know" when trying to predict the settlement amount ($n=79$) relative to the

¹⁷This analysis also shows that individuals do not go into the details of the data settlement lawsuit or the specifics of Facebook's revenue streams. The responses of users are generic and centered around the broad themes found in the text analysis. This is also consistent with the limited time individuals spend on the information treatment page.

revenue amount (n=53).¹⁸ That said, there was still substantial uncertainty and underestimation about the amount of money Facebook would earn per user, consistent with the smaller, but still substantial, share of revisions we observed in the revenue treatment. This survey lends credence to the fact that individuals are especially uncertain about settlement values, leading to a higher effect of the settlement treatment relative to the revenue treatment.

Overall, this section provides evidence on how information frictions might contribute to low data valuations and how this might explain some differences in valuations across demographic groups of interest. Moreover, we show that information interventions can help individuals revise their beliefs about data valuations, especially low WTA individuals. Such interventions can become an important part of the policy toolkit to spread awareness about the value of social media data. It appears that such informational interventions, while important, cannot eliminate the heterogeneity and dispersion in data valuations across groups¹⁹. This is because an individual's valuations are based on objective (market-based) information as well as subjective beliefs about data and privacy.

5 Results for the DDP Sample

In this section we report the results of the survey experiment fielded with the Data Dividend Project. Members of the DDP are those who are interested in being part of a movement that ensures that technology companies share a part of their revenue when they monetize data. Access to this sample allows us a rare insight into information frictions associated with personal data valuations for individuals whom we would expect to be more conscious and better informed about data and privacy.

Six hundred and fifty-two respondents completed the experiment. Table A.2 shows this DDP sample to be more concerned about privacy than the YouGov sample, as DDP members provide significantly different responses to most of the endline privacy attitude questions. For

¹⁸These elicitations were not incentivized, so as to avoid subjects searching for the exact answer.

¹⁹See Figure A.4 in the Online Appendix, where the distributions of revised valuations are closer than the distributions of baseline distributions.

example, DDP members are more likely to believe that privacy is a fundamental right and that tech companies earn too much, and less likely to believe that the free market will lead to the appropriate amount of privacy. In terms of demographics, DDP members are more likely to be Asian and less likely to be Black, Table A.1. In addition, men are overrepresented in the DDP sample.

Figure 6 presents the distribution of valuations across the DDP and the YouGov sample. The WTA distribution in the DDP first order stochastically dominates that of the representative sample respondents. Figure 6 shows an additional nuance: individuals in the DDP sample are more likely to have very high WTA and less likely to have very low valuations relative to the YouGov sample (in fact, the median baseline valuation for the DDP sample is \$1000, which is 33% higher than the median valuation of \$750 in the YouGov sample). These results are qualitatively in line with intuition, and provide confidence in our general approach.

Analyzing the WTA heterogeneity by race,²⁰ we find that the median Black respondent (n=29) has a WTA of \$500, whereas that of a White respondent is \$1000. In line with the YouGov sample, WTA increases with income, with a median WTA of \$600 for the lowest income bracket (less than \$30K) and a WTA of \$1000 for those with income greater than \$100K. Also in line with the YouGov sample, those individuals who prefer not to report their income (n=65) have the highest median valuations: \$10,000. Finally, the median valuations by gender are the same at \$1000, though female WTA is lower after the 75th percentile.

We find remarkably similar responses to the information treatments across the two sample of respondents. The probability of revision in the DDP sample is 29.4% (28.6% in the YouGov sample). As in the YouGov sample, the probability of revision in the DDP sample is asymmetric, with 58.7% of individuals whose WTA is less than \$400 revising their valuations (Figure 7). All participants who revise their WTA revise it higher. Additionally, in line with the YouGov sample, we see that individuals are 12% more likely to respond to the settlement treatment relative to the revenue treatment. This effect is noisy (and not statistically signif-

²⁰This analysis is based on a smaller sample than the YouGov analysis. Therefore, some of its results should be taken as suggestive.

icant), due to the smaller sample size and hence should be interpreted as suggestive. These results are notable and may shed light on the potential information problem associated with individuals being asked to evaluate their own data. They show that also individuals who care about the core issue of data dividends and privacy also have a hard time valuing their personal data for which there is no active market.

The DDP analysis replicates important results from the YouGov sample, providing confidence in our general measurement approach. Moreover, it provides estimates from a sample of users who are highly engaged in the data and privacy policy conversation. These individuals are hard to access (access to them was possible through our collaboration with the DDP). In line with intuition, the DDP sample has a higher WTA than the YouGov sample, but the information treatment results highlight that this sample suffers, too, from information frictions.

6 Conclusion

As policymakers explore introducing data dividends and companies experiment with new business models around data markets, it is essential to understand the economic valuations of consumers' personal data. In this paper, we provide evidence documenting substantial dispersion and heterogeneity by gender, race, and income in users' data valuations for their social media data through incentive compatible studies on a representative sample of US internet population as well as a data conscious sample. Marginalized individuals (women, Black, and lower income) have significantly lower data valuations in both samples even after controlling for income and education. Through a randomized intervention, we find evidence that participants respond to information giving them a signal about the value of their data from legal settlements and revenue projections. Specifically, we find that low WTA users in both samples revise their valuations upwards towards the settlement amount while high WTA users do not revise downwards. These revisions significantly reduce the observed heterogeneity in baseline valuations for marginalized individuals. Dispersion and heterogeneity in valuations,

however, persist following the information treatments, consistent with theories of data and privacy valuations that construed them as amalgams of objective and subjective factors.

Our research is not without limitations. First, we explore data from only one platform, that is, Facebook, and analyze the value for the entire stock of data, which includes both sensitive data (private messages, photos) and less sensitive data (public likes). Although studying Facebook data is important because it is the largest social media platform in the world, it would be informative to replicate these results for other types of data that might range in their degree of sensitivity. The aim of the paper is not to focus on an exact dollar valuation for data; rather, we focus on comparing valuations across different demographic groups. Future research could study how firms value user data and how these valuations vary based on user characteristics.

References

- Acquisti, A., Brandimarte, L., and Loewenstein, G. (2015). “Privacy and human behavior in the age of information.” *Science*, 347(6221), 509–514.
- Acquisti, A., Taylor, C., and Wagman, L. (2016). “The economics of privacy.” *Journal of economic Literature*, 54(2), 442–92.
- Adjerid, I., Acquisti, A., Brandimarte, L., and Loewenstein, G. (2013). “Sleights of privacy: Framing, disclosures, and the limits of transparency.” In *Proceedings of the ninth symposium on usable privacy and security*, 1–11.
- Alesina, A., Ferroni, M. F., and Stantcheva, S. (2021). “Perceptions of racial gaps, their causes, and ways to reduce them.” Tech. rep., National Bureau of Economic Research.
- Allcott, H., Braghieri, L., Eichmeyer, S., and Gentzkow, M. (2020). “The welfare effects of social media.” *American Economic Review*, 110(3), 629–76.
- Allcott, H., and Taubinsky, D. (2015). “Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market.” *American Economic Review*, 105(8), 2501–38.
- Angelov, D. (2020). “Top2vec: Distributed representations of topics.” *arXiv preprint arXiv:2008.09470*.
- Arrieta-Ibarra, I., Goff, L., Jiménez-Hernández, D., Lanier, J., and Weyl, E. G. (2018). “Should we treat data as labor? moving beyond” free”.” In *aea Papers and Proceedings*, vol. 108, 38–42.
- Athey, S., Catalini, C., and Tucker, C. (2017). “The digital privacy paradox: Small money, small costs, small talk.” Tech. rep., National Bureau of Economic Research.
- Au-Yeung, A. (2019). “California wants to copy alaska and pay people a ‘data dividend.’ is it realistic?” <https://www.forbes.com/sites/angelaueung/2019/02/14/california-wants-to-copy-alaska-and-pay-people-a-data-dividend-is-it-realistic/?sh=76a5d1d4222c>, accessed: 2021-11-30.
- Becker, G. M., DeGroot, M. H., and Marschak, J. (1964). “Measuring utility by a single-response sequential method.” *Behavioral science*, 9(3), 226–232.
- Beke, F. T., Eggers, F., Verhoef, P. C., et al. (2018). “Consumer informational privacy: Current knowledge and research directions.”
- Benndorf, V., and Normann, H.-T. (2018). “The willingness to sell personal data.” *The Scandinavian Journal of Economics*, 120(4), 1260–1278.
- Beresford, A. R., Kübler, D., and Preibusch, S. (2012). “Unwillingness to pay for privacy: A field experiment.” *Economics letters*, 117(1), 25–27.
- Beshears, J., Choi, J. J., Laibson, D., and Madrian, B. C. (2009). “The importance of default options for retirement saving outcomes: Evidence from the united states.” In *Social security policy in a changing environment*, 167–195, University of Chicago Press.

- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). “Latent dirichlet allocation.” *the Journal of machine Learning research*, 3, 993–1022.
- Brynjolfsson, E., Collis, A., and Eggers, F. (2019). “Using massive online choice experiments to measure changes in well-being.” *Proceedings of the National Academy of Sciences*, 116(15), 7250–7255.
- Calo, R. (2011). “The boundaries of privacy harm.” *Ind. LJ*, 86, 1131.
- Cao, X., and Zhang, J. (2021). “Preference learning and demand forecast.” *Marketing Science*, 40(1), 62–79.
- Christofides, E., Muise, A., and Desmarais, S. (2012). “Hey mom, what’s on your facebook? comparing facebook disclosure and privacy in adolescents and adults.” *Social Psychological and Personality Science*, 3(1), 48–54.
- Danezis, G., Lewis, S., and Anderson, R. J. (2005). “How much is location privacy worth?” In *WEIS*, vol. 5, Citeseer.
- De Quidt, J., Haushofer, J., and Roth, C. (2018). “Measuring and bounding experimenter demand.” *American Economic Review*, 108(11), 3266–3302.
- Elsaify, M., and Hasan, S. (2020). “Some data on the market for data.” *Available at SSRN 3568817*.
- Facebook (2021a). “Facebook inc, 10-k filed on 1/28/2021.” <https://d18rn0p25nwr6d.cloudfront.net/CIK-0001326801/63ed9556-914c-4fbe-9603-09900dbaf7d9.html>, accessed: 2021-11-30.
- Facebook (2021b). “How do i download a copy of my information on facebook?” <https://www.facebook.com/help/212802592074644>, accessed: 2021-11-30.
- Falk, A., and Zimmermann, F. (2013). “A taste for consistency and survey response behavior.” *CEifo Economic Studies*, 59(1), 181–193.
- Feygin, Y., Li, H., Lala, C., Hecht, B., Vincent, N., Scarcella, L., and Prewitt, M. (2021). “A data dividend that works: steps toward building an equitable data economy.” *Berggruen Institute White Paper*.
- Guriev, S., Henry, E., and Zhuravskaya, E. (2020). “Checking and sharing alt-facts.” *CEPR Discussion Paper No. DP14738*.
- Haaland, I., Roth, C., and Wohlfart, J. (2020). “Designing information provision experiments.”
- Hjort, J., Moreira, D., Rao, G., and Santini, J. F. (2021). “How research affects policy: Experimental evidence from 2,150 brazilian municipalities.” *American Economic Review*, 111(5), 1442–80.
- Holtrop, N., Wieringa, J. E., Gijzenberg, M. J., and Verhoef, P. C. (2017). “No future without the past? predicting churn in the face of customer privacy.” *International Journal of Research in Marketing*, 34(1), 154–172.

- Hoy, M. G., and Milne, G. (2010). “Gender differences in privacy-related measures for young adult facebook users.” *Journal of interactive advertising*, 10(2), 28–45.
- Huang, Y., Ellickson, P. B., and Lovett, M. J. (2020). “Learning to set prices.” *Available at SSRN 3267701*.
- Hui, K.-L., Teo, H. H., and Lee, S.-Y. T. (2007). “The value of privacy assurance: An exploratory field experiment.” *Mis Quarterly*, 19–33.
- Lambrech, A., and Tucker, C. (2019). “Algorithmic bias? an empirical study of apparent gender-based discrimination in the display of stem career ads.” *Management science*, 65(7), 2966–2981.
- Laudon, K. C. (1996). “Markets and privacy.” *Communications of the ACM*, 39(9), 92–104.
- Lin, T. (2020). “Valuing intrinsic and instrumental preferences for privacy.” *Available at SSRN 3406412*.
- Prince, J., and Wallsten, S. (2020). “How much is privacy worth around the world and across platforms?” *Available at SSRN*.
- Schwartz, P. M. (2003). “Property, privacy, and personal data.” *Harv. L. Rev.*, 117, 2056.
- Sevelius, J. M., Gutierrez-Mock, L., Zamudio-Haas, S., McCree, B., Ngo, A., Jackson, A., Clynes, C., Venegas, L., Salinas, A., Herrera, C., et al. (2020). “Research with marginalized communities: challenges to continuity during the covid-19 pandemic.” *AIDS and Behavior*, 24(7), 2009–2012.
- Shiller, B. (2021). “Optimized sticky targeted pricing.” *Available at SSRN 3845138*.
- Shiller, B. R. (2020). “Approximating purchase propensities and reservation prices from broad consumer tracking.” *International Economic Review*, 61(2), 847–870.
- Spiekermann, S., Acquisti, A., Böhme, R., and Hui, K.-L. (2015). “The challenges of personal data markets and privacy.” *Electronic markets*, 25(2), 161–167.
- Spiekermann, S., Korunovska, J., and Bauer, C. (2012). “Psychology of ownership and asset defense: Why people value their personal information beyond privacy.” *Available at SSRN 2148886*.
- Streamlytics (2021). “Streamlytics crushes target, 1 billion data points processed.” <https://www.globenewswire.com/news-release/2021/10/19/2316656/0/en/Streamlytics-Crushes-Target-1-Billion-Data-Points-Processed.html>, accessed: 2021-11-30.
- Sun-Times Staff (2020). “Illinois facebook users can apply for up to \$400 settlement as part of privacy lawsuit.” <https://chicago.suntimes.com/2020/9/22/21451907/illinois-facebook-users-can-apply-for-up-to-400-settlement-as-part-of-privacy-lawsuit>, accessed: 2021-11-30.
- Tomaino, G., Wertenbroch, K., and Walters, D. J. (2021). “Intransitivity of consumer preferences for privacy.” *INSEAD Working Paper*.

- Tsai, J. Y., Egelman, S., Cranor, L., and Acquisti, A. (2011). “The effect of online privacy information on purchasing behavior: An experimental study.” *Information systems research*, 22(2), 254–268.
- Ulloa, J. (2019). “Newsom wants companies collecting personal data to share the wealth with californians.” <https://www.latimes.com/politics/la-pol-ca-gavin-newsom-california-data-dividend-20190505-story.html>, accessed: 2021-11-30.
- Xu, H., and Zhang, N. (2020). “From contextualizing to context-theorizing: Assessing context effects in privacy research.” *Available at SSRN 3624056*.
- Yang, A. (2020a). “Op-ed: Andrew yang: Make tech companies pay for your data.” <https://www-latimes-com.libproxy.mit.edu/opinion/story/2020-06-23/andrew-yang-data-dividend-tech-privacy>, accessed: 2021-11-30.
- Yang, K. H. (2020b). “Selling consumer data for profit: Optimal market-segmentation design and its consequences.”

Tables and Figures

Table 1: Race and Gender Regression results

	(1)	(2)	(3)	(4)
	Log(WTA)	Log(WTA)	Log(WTA)	Log(WTA)
Female	-0.378** (0.157)		-0.376** (0.157)	-0.398** (0.161)
Black		-0.557** (0.217)	-0.553** (0.217)	-0.438** (0.217)
Controls	N	N	N	Y
Obs	4141	4141	4141	4140
Adj. R ²	0.001	0.001	0.002	0.031

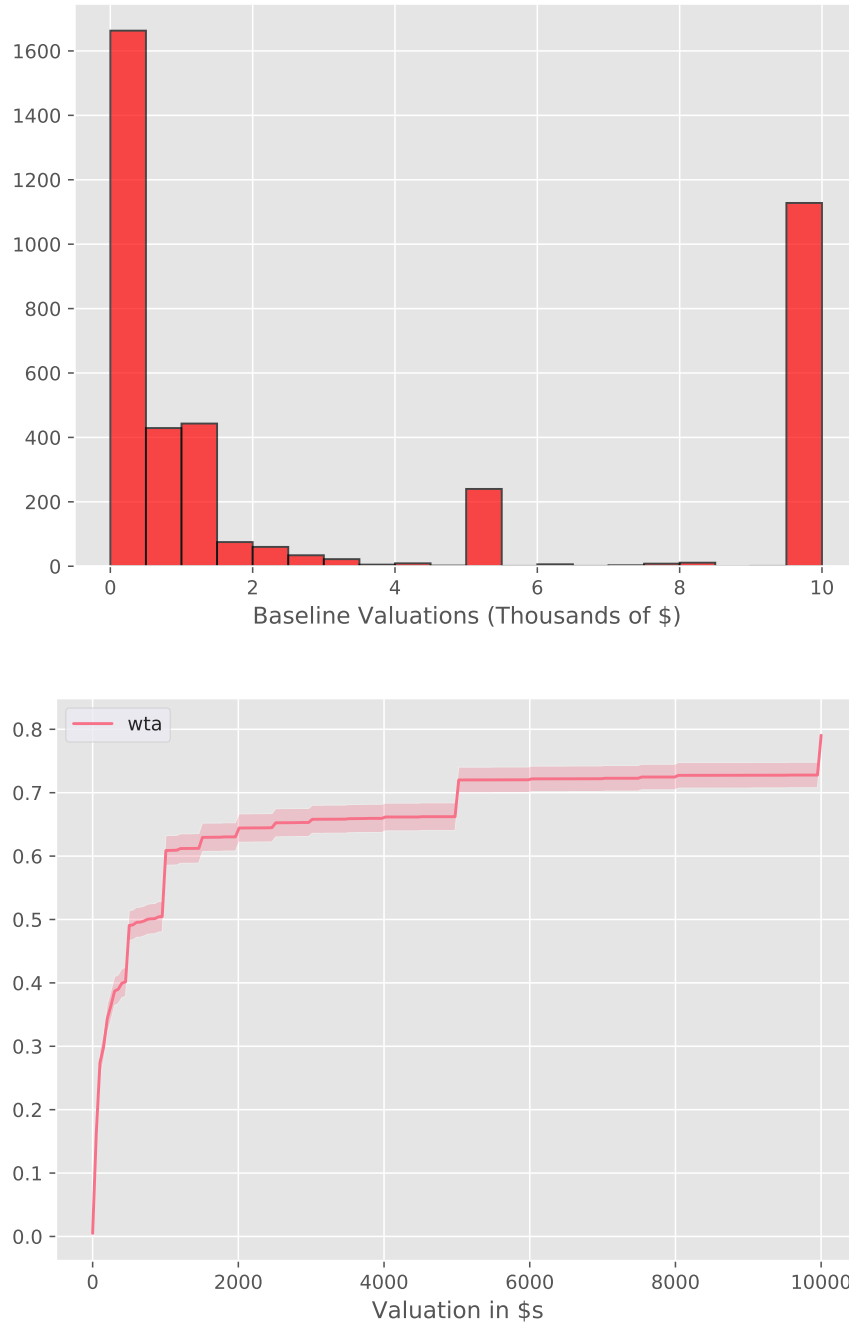
The unit of observation is the individual respondent. Robust standard errors in parentheses. The estimates are based on an OLS regression. Column (4) uses controls for income, education, age, Facebook usage and privacy beliefs. The mean of the dependent variable is 7.55. * $p < 0.1$, ** $p < 0.05$ *** $p < 0.01$.

Table 2: Mechanism Regression results

	(1)	(2)	(3)	(4)
	Log(WTA)	Log(revised WTA)	P(Revise)	P(Revise)
Female	-0.429*** (0.162)	-0.361** (0.159)	0.050*** (0.014)	0.039*** (0.014)
Black	-0.431** (0.217)	-0.355* (0.210)	0.049** (0.024)	0.038* (0.023)
Privacy is a Right	0.949*** (0.153)	0.932*** (0.151)	-0.002 (0.017)	0.022 (0.017)
Market is Correct	-0.826*** (0.188)	-0.884*** (0.182)	0.099*** (0.022)	0.078*** (0.022)
Age	0.009** (0.004)	0.011*** (0.004)	-0.001 (0.000)	-0.000 (0.000)
High Income	0.645*** (0.143)	0.549*** (0.140)	-0.060*** (0.015)	-0.044*** (0.015)
Prefer Not to Report Income	2.171*** (0.371)	2.153*** (0.368)	-0.063*** (0.023)	-0.010 (0.023)
Baseline WTA Control	N	N	N	Y
Obs	4140	4140	4140	4140
Adj. R ²	0.030	0.031	0.014	0.084



The unit of observation is the individual respondent. Robust standard errors in parentheses. All regressions also control for intensity of Facebook usage. The estimates are based on an OLS regression. * $p < 0.1$, ** $p < 0.05$ *** $p < 0.01$.

Figure 1: Distribution of Baseline Valuations: Full Sample




The figure in the top panel shows a histogram with the distribution of valuations for the YouGov sample at \$250 intervals. In this figure, all values above \$10,000 are displayed in the bar at \$10,000. The bottom panel shows the cumulative distribution function (CDF) of baseline valuations for the YouGov sample. The shaded area represents the 95% uniform confidence interval for the distribution.

Figure 2: Distribution of Baseline Valuations: Race and Gender Heterogeneity




The figure in the top panel shows CDF of the baseline WTA split up by race. The Online Appendix contains distributions of valuations for additional race and ethnicities. The bottom panel shows the CDF of WTA by gender. The shaded area represents the 95% uniform confidence interval for the distribution.

Figure 3: Distribution of Baseline Valuations: Income Heterogeneity



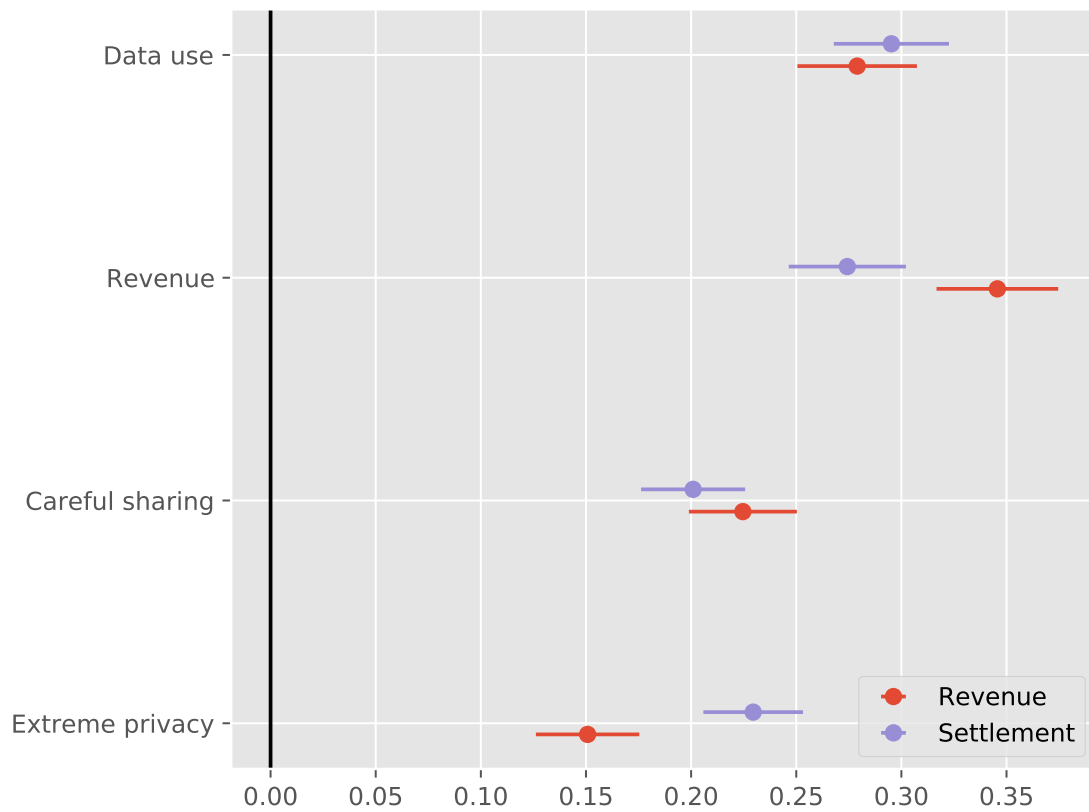
The figure shows the CDF of the baseline WTA by Income. The shaded area represents the 95% uniform confidence interval for the distribution

Figure 4: Revision of Baseline Valuations



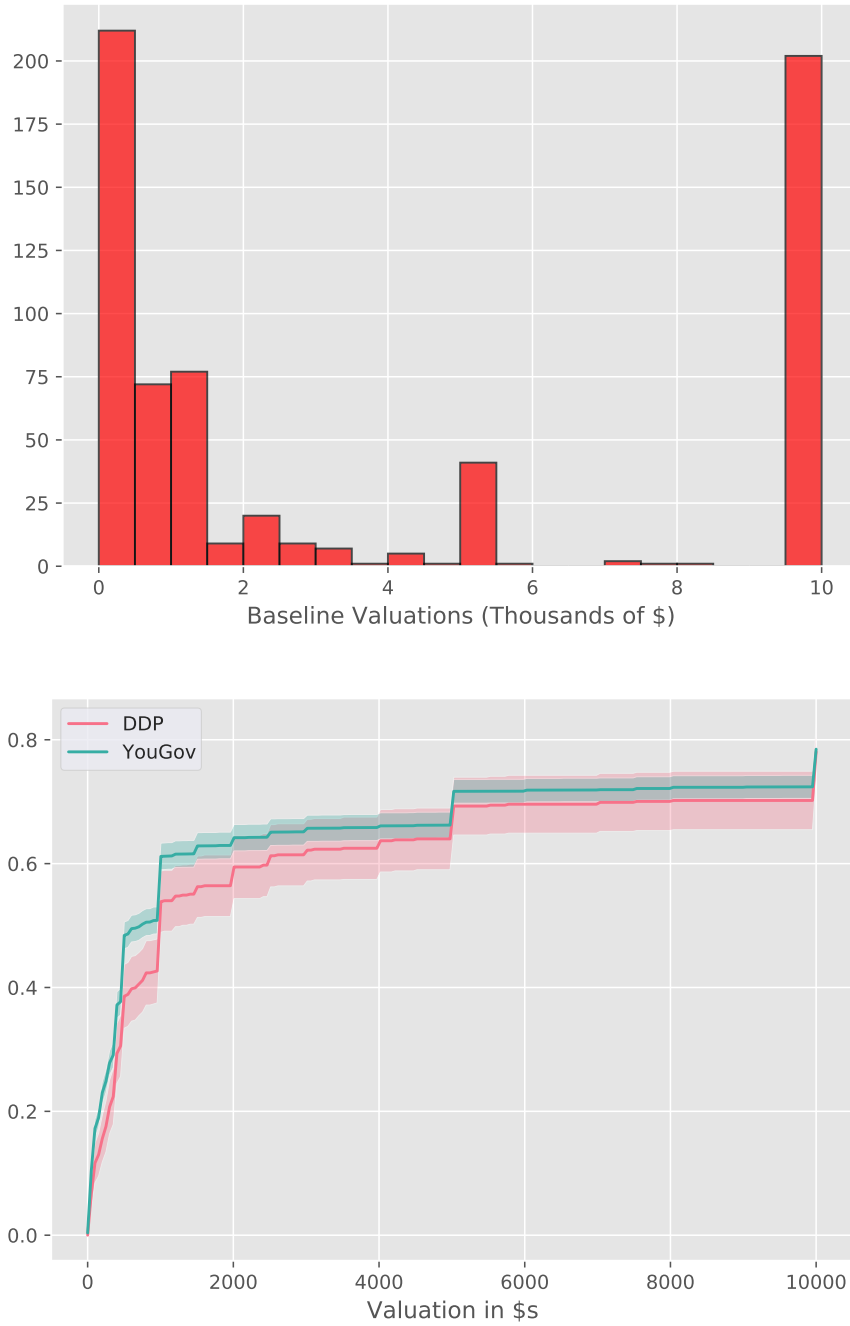
The figure shows the CDF of baseline and revised WTA, where the revised WTA is the valuation measured after the information interventions. The shaded area represents the 95% uniform confidence interval for the distribution

Figure 5: Text Explanation Clustering by Revision Group




Here we plot coefficients from regressions of an indicator for cluster membership on treatment. The horizontal bars represent 95% confidence intervals.

Figure 6: Distribution of Baseline Valuations: DDP



The figure in the top panel shows a histogram of the distribution of valuations for the DDP sample. All values greater than \$10,000 are included in the bar at \$10,000. The bottom panel shows the cumulative distribution function (CDF) of baseline valuations for the DDP and YouGov samples. The shaded area represents the 95% uniform confidence interval for the distribution.

Figure 7: CDF of Valuations with Revisions: DDP vs. YouGov

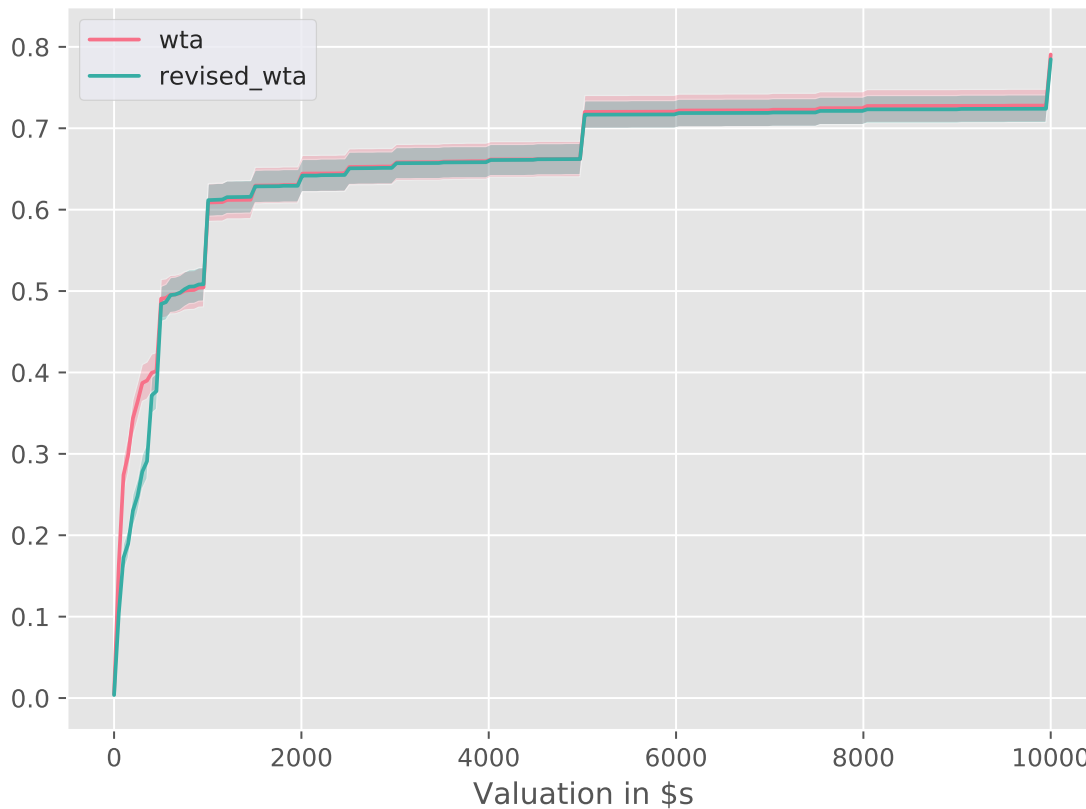


The figure shows the CDF of baseline and revised WTA for the YouGov and DDP samples, where the revised WTA is the valuation measured after the information interventions. The shaded area represents the 95% uniform confidence interval for the distribution

Online Appendix


First, we plot the distribution of baseline and revised valuations in more detail. Figure A.1 plots the distribution of baseline and revised valuations for the YouGov sample extended out to \$10,000. Figure A.2 plots the distribution of baseline valuations by race / ethnicity for all groups in the survey.

Figure A.1: Revision of Baseline Valuations




The figure shows the CDF of baseline and revised WTA for the YouGov sample, where the revised WTA is the valuation measured after the information interventions. The shaded area represents the 95% uniform confidence interval for the distribution

Figure A.2: Distribution of Baseline Valuations: Full Sample and Heterogeneity by Race



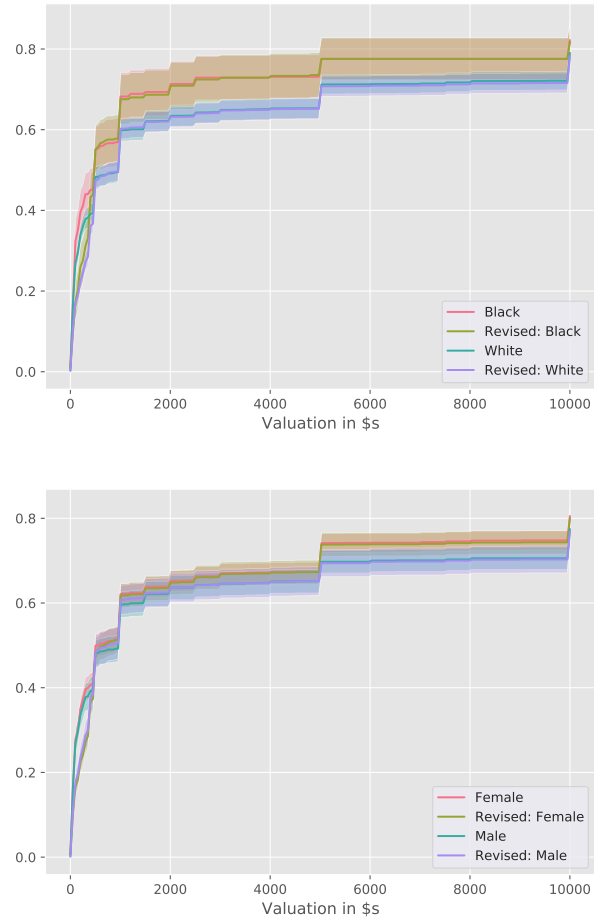
The figure in the top panel shows CDF of the baseline WTA for the whole sample while in the bottom panel it is the CDF split up by race. The shaded area represents the 95% uniform confidence interval for the distribution

Figure A.3: Valuation Distribution Regressions



Distribution regressions of an indicator if an individual's valuation is greater than a threshold (y-axis) on demographic variables. The bars represent 95% confidence intervals.

Figure A.4: Baseline and Revised Valuations by Demographic



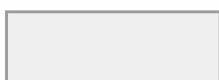
The figure shows the CDF of baseline and revised WTA for the YouGov sample by Race (top figure) and Gender (bottom figure), where the revised WTA is the valuation measured after the information interventions. The shaded area represents the 95% uniform confidence interval for the distribution

Survey Materials

Here we share screenshots of key parts of the survey instrument. Figure A.5 shows the text of the question where subjects entered their valuations, Figure A.6 shows the Revenue treatment and the opportunity for subjects to revise their valuations, and Figure A.7 shows the Settlement treatment and the opportunity for subjects to revise their valuations.

Figure A.5: Baseline Question Screenshot

What is the minimum amount of money (in US Dollars) you would require to share all your Facebook data? This includes your posts, photos, messages, likes and comments.



The screenshot shows the question asked to the respondents to elicit their baseline valuations.

Figure A.6: Revenue Information Treatment Screenshot

To provide some additional context, Facebook is expected to earn over \$400 per North American user over the next three years as reported in their Securities and Exchange Commission (SEC) filing ([source](#)).

You answered that you will share your Facebook data for \$.

Do you want to change your answer?

Yes

No

The screenshot shows the revenue information treatment after eliciting baseline valuations.

Figure A.7: Settlement Information Treatment Screenshot

To provide some additional context, Facebook recently lost a class action lawsuit for harvesting user data and violating privacy laws and agreed to pay around \$400 per user for eligible users ([source](#)).

You answered that you will share your Facebook data for \$.

Do you want to change your answer?

Yes

No

The screenshot shows the settlement information treatment after eliciting baseline valuations.

Sample Characteristics

In Table A.1 we plot the distribution of demographic characteristics by sample. Notice that the DDP sample has substantially more Asian respondents, and fewer Black respondents, relative to the YouGov sample. In addition, the DDP sample has more Male respondents, uses Facebook less on average, and is more Liberal than the YouGov sample. Table A.2 shows the difference in privacy attitudes between the two samples.

Table A.1: Demographics Across Samples

	DDP		YouGov		P-Val
Race / Ethnicity: Asian	0.157	(0.364)	0.034	(0.181)	0.000
Race / Ethnicity: Black	0.044	(0.205)	0.109	(0.311)	0.000
Race / Ethnicity: Hispanic	0.071	(0.257)	0.089	(0.285)	0.105
Race / Ethnicity: Middle Eastern	0.000	(0.000)	0.002	(0.047)	0.003
Race / Ethnicity: Native American	0.000	(0.000)	0.012	(0.108)	0.000
Race / Ethnicity: Other	0.047	(0.212)	0.043	(0.204)	0.697
Race / Ethnicity: White	0.681	(0.467)	0.711	(0.453)	0.119
Gender: Female	0.322	(0.468)	0.539	(0.499)	0.000
Gender: Male	0.657	(0.475)	0.440	(0.497)	0.000
Gender: Non-Binary / Third Gender	0.015	(0.122)	0.021	(0.142)	0.303
Gender: Prefer Not To Say	0.006	(0.078)	0.000	(0.000)	0.045
Fb Usage: 10-30	0.221	(0.415)	0.284	(0.451)	0.000
Fb Usage: 31-60	0.148	(0.356)	0.203	(0.402)	0.000
Fb Usage: Less Than 10	0.458	(0.499)	0.247	(0.431)	0.000
Fb Usage: More Than 60	0.172	(0.378)	0.267	(0.442)	0.000
Political Views: Conservative	0.017	(0.128)	0.145	(0.352)	0.000
Political Views: Extremely Conservative	0.003	(0.055)	0.068	(0.251)	0.000
Political Views: Extremely Liberal	0.132	(0.338)	0.138	(0.345)	0.659
Political Views: Liberal	0.375	(0.485)	0.218	(0.413)	0.000
Political Views: Moderate	0.147	(0.354)	0.232	(0.422)	0.000
Political Views: Other	0.083	(0.276)	0.052	(0.223)	0.007
Political Views: Slightly Conservative	0.044	(0.205)	0.064	(0.245)	0.021
Political Views: Slightly Liberal	0.200	(0.400)	0.083	(0.276)	0.000

Pre-treatment covariate means and standard deviations for all respondents who completed the survey for both samples. The p-values come from a test of equality of means across the two treatments.

Table A.2: Endline Privacy Attitudes Across Samples

	DDP		YouGov		P-Val
Endline Fundamental Human Right	5.295	(1.110)	5.065	(1.307)	0.000
Endline Careful	4.637	(1.240)	4.705	(1.290)	0.192
Endline Free Market	1.360	(1.594)	2.253	(1.809)	0.000
Endline Misuse	5.077	(1.176)	4.514	(1.403)	0.000
Endline Earn Too Much	5.135	(1.371)	4.683	(1.417)	0.000

Means and standard deviations for all respondents who completed the survey for both samples. The p-values come from a test of equality of means across the two treatments.

Randomization Checks

Below we provide evidence that the randomization between Revenue and Settlement treatments occurred as expected. Table A.3 shows that roughly half of the final sample received each treatment and we cannot reject the null that each are shown with equal probability. Second, Table A.4 shows the distribution of covariates across the two treatments and we cannot reject the null that these are equal in the large majority of variables.

Table A.3: Randomization Check

	Revenue	Settlement	P-Val
Treatment	0.490 (2031)	0.510 (2110)	0.225

Here we plot the share of respondents who received the Revenue and Settlement information treatments and the results from a test of the null that these shares are equal.

Table A.4: Balance Tests

	Revenue		Settlement		P-Val
Race / Ethnicity: Asian	0.029	(0.167)	0.039	(0.193)	0.066
Race / Ethnicity: Black	0.115	(0.319)	0.103	(0.304)	0.220
Race / Ethnicity: Hispanic	0.083	(0.276)	0.095	(0.293)	0.172
Race / Ethnicity: Middle Eastern	0.000	(0.022)	0.004	(0.061)	0.021
Race / Ethnicity: Native American	0.012	(0.108)	0.012	(0.108)	0.993
Race / Ethnicity: Other	0.047	(0.212)	0.040	(0.196)	0.240
Race / Ethnicity: White	0.714	(0.452)	0.708	(0.455)	0.651
Gender: Female	0.538	(0.499)	0.540	(0.498)	0.866
Gender: Male	0.440	(0.497)	0.441	(0.497)	0.970
Gender: Non-Binary / Third Gender	0.022	(0.147)	0.019	(0.136)	0.469
Age: 10000 - 10019	0.000	(0.022)	0.000	(0.000)	0.317
Age: 18 - 39	0.372	(0.483)	0.386	(0.487)	0.336
Age: 40 - 59	0.327	(0.469)	0.321	(0.467)	0.652
Age: 60 - 79	0.290	(0.454)	0.284	(0.451)	0.689
Age: 80 - 99	0.011	(0.104)	0.009	(0.094)	0.553
Income: \$30K-\$50K	0.168	(0.374)	0.184	(0.387)	0.191
Income: \$50K-\$100K	0.284	(0.451)	0.289	(0.453)	0.721
Income: <\$30K	0.223	(0.416)	0.209	(0.406)	0.257
Income: >\$100K	0.211	(0.408)	0.199	(0.399)	0.332
Income: Prefer Not To Say	0.114	(0.318)	0.120	(0.325)	0.537
Fb Usage: 10-30	0.301	(0.459)	0.267	(0.442)	0.015
Fb Usage: 31-60	0.197	(0.398)	0.208	(0.406)	0.417
Fb Usage: Less Than 10	0.236	(0.425)	0.258	(0.438)	0.109
Fb Usage: More Than 60	0.265	(0.442)	0.268	(0.443)	0.862
Fb Age	2009.757	(3.925)	2009.740	(3.828)	0.891
Political Views: Conservative	0.150	(0.357)	0.140	(0.347)	0.344
Political Views: Extremely Conservative	0.067	(0.251)	0.068	(0.251)	0.968
Political Views: Extremely Liberal	0.135	(0.342)	0.141	(0.348)	0.585
Political Views: Liberal	0.210	(0.408)	0.225	(0.418)	0.246
Political Views: Moderate	0.232	(0.422)	0.232	(0.422)	0.991
Political Views: Other	0.050	(0.218)	0.055	(0.227)	0.536
Political Views: Slightly Conservative	0.069	(0.253)	0.060	(0.237)	0.227
Political Views: Slightly Liberal	0.086	(0.281)	0.081	(0.272)	0.515

Pre-treatment covariate means and standard deviations for all respondents who completed the survey for both information treatments. The p-values come from a test of equality of means across the two treatments.

Analysis of text responses

After subjects were presented with the informational intervention and given the opportunity to revise their valuations, we asked subjects why they did or did not change their answer and about their general attitudes towards privacy and how their data is used. To better understand what was inducing subjects to revise their valuations, we conducted a post-hoc text analysis.

To do so, we used the top2vec algorithm proposed in Angelov (2020) that automatically handles stop-word removal, lemmatization, and selection of the number of clusters. Moreover, this algorithm, relative to bag-of-words approaches such as Latent Dirichlet Allocation (Blei et al., 2003), takes advantages of advances in word embeddings that account for the context of a word in a document. We find four relatively coherent topics emerge from the data.

We labeled these four clusters through manual inspection of the top 10 most closely aligned responses within each cluster. As shown in Table A.5, we see that this cluster contains many statements concerning the amount of money Facebook earns off of their data. As such, we refer to this cluster as the Revenue cluster. Turning to Table A.6, we see that this cluster contains statements concerning the use of an individual’s data. The algorithm failed to separate individuals who are both concerned and those who are not concerned about how their data is used, and this cluster contains both types. Therefore, we refer to this as the Data Use cluster. The third cluster, shown in Table A.7, contains individuals who are privacy conscious and, for that reason, do not share much information on Facebook. We refer to this as the Careful Sharing cluster. Finally, Table A.8 contains many responses that indicate an individual is extremely privacy conscious and is unwilling to share their data for anything but a very high price. We label this the Extreme Privacy cluster.

Table A.5: Sample of responses from revenue cluster

Sample	Text Response
1	i changed my facebook data valuation because i figured if facebook is willing to pay that amount for certain users then that must mean my data alon...
2	if i cared about my privacy, i would not be using facebook. but my data is worth a lot to facebook, so i should be paid just an insane amount of mo...
3	yes, it made me think my data is worth more. facebook must be using my data for a lot of purposes if it is worth so much to them, even up to \$400. ...
4	because i realized how much revenue facebook is making from people sharing their data on facebook, so i decided my data is worth more than i previo...
5	i believe people underestimate how much money facebook makes off our data. i would not want to share my data for a modest amount. less than the num...
6	i changed the amount when i saw how much money facebook will make on each north american user over three years. i am not happy with the way facebo...
7	i care about privacy and if my data is misused but in this case i don't use facebook much and there isn't much data i care about. however i do nee...
8	i'm not happy with the way my data is used, but because of that i don't put much data on facebook in the first place. i know it will likely be misu...
9	i care about privacy, but facebook is definitely misusing our data. i don't know what my data is worth. i don't like the idea of sharing it but it ...
10	in my opinion my data has a value which i feel is a fair amount related to what facebook does with my data and information. they are making a treme...

Table A.6: Sample of responses from data use cluster

Sample	Text Response
1	i'm happy with the way my data is being used to the best of my knowledge. i'm not really worried that my data has been or will be misused. i do care...
2	i changed my mind because the information i would be sharing contains some very personal data. i am ok with the way my data is being used. i am not...
3	i am not concerned about my data as there is not enough data that would affect me on facebook. i am not worry about my data being misused. i do care...
4	i'm not exactly sure how my data is being used, but it doesn't bother me that much. i'm not overly concerned with my data being misused. kinda care...
5	i am not happy with the way my data is being used, nor how it is gathered. i definitely worry about my data being misused, and i care about my priv...
6	i am worried about my data being misused. i care about my privacy. i am not happy with my data being used the way it is. i am worth at least that a...
7	i wouldn't because i feel that my information is worth at least \$500. i'm not happy with the way my data is being used because i feel like i'm bein...
8	no, i am not happy with how my data is being used and i am always worried that my data will be misused. my privacy is very important to me and i do...
9	my data is information about myself and shouldn't be shared with people i don't know. i am worried about someone misusing my data. yes i do care ab...
10	i changed my evaluation because you said facebook settled a lawsuit for around \$400 per user. i am not terribly concerned about my privacy because ...

Table A.7: Sample of responses from careful sharing cluster

Sample	Text Response
1	i think \$10,000 is a fair amount to be sharing all my information. i think the \$400 offered by fb is way too little. privacy is a big concern for m...
2	i worry about data misuse and i value my privacy highly. for that reason, i do not share sensitive information on facebook and only very few pictu...
3	i worry about ever facebook does and do all i can to maintain my privacy and not allow facebook to use my information without my consent.
4	i don't really care about privacy on facebook because i am not ashamed of anything tied to my activity. but if facebook wants access to my data, i ...
5	i really care about privacy . i'm always very careful when sharing stuff on facebook especially my location . i don't fully trust facebook company ...
6	i am very uncomfortable sharing my facebook info indiscriminately. i'm not happy with how facebook treats my information. i do not believe they ha...
7	honestly i don't really know how facebook uses my data . i do care about privacy and i am curious as to what they use the data on. i would like it ...
8	i really enjoy facebook but it would take alot of money for me to share my data it is personal i dont want anything i share misused and it does wor...
9	i just don't think i have enough down on facebook to worry about it. i barely share anything online because i'm a very private person to begin with...
10	facebook has made billions by sharing my data. i find this disconcerting and i prefer to keep it relatively private and have the ability to share o...

Table A.8: Sample of responses from extreme privacy cluster

Sample	Text Response
1	i care deeply about my privacy which is why i put such a high number. at that payout price i could afford to handle any repercussions from my data
2	... after seeing that each person was only getting \$400 as compensation for data that was used illegally then it really can't be worth that much. i'd s...
3	i care about my privacy and don't want to share my data. i chose a fairly high dollar amount in order to price my data high enough that it would no...
4	if you want my data you will pay for it. my data is not worth that much but not sharing it for less.
5	i just don't feel comfortable with sharing my data unless it's for a certain sum of money because it does have value
6	i honestly put 0 before because i wasn't going to share my data, but if i'm getting paid for it i might as well share the data.
7	i believe sharing my data is worth more money. i am worried about my data being misused.
8	if i'm going to voluntarily give my data, i'd want to be well paid for it. \$400 is not what that many years of data is worth.
9	most of the data is being used without my permission anyway. the amount i chose seems reasonable for the one-year profit from my information.
10	i don't want to share my data for any monetary amount, so that's why i picked such a high borderline insane number like \$1,000,000. you did not giv...