

Facebook isn't free: zero-price companies overcharge consumers with data

MATT SUMMERS *

Harvard Law School Antitrust Association, Cambridge, MA, USA

Abstract: The emergence of dominant companies that don't charge money to their primary consumers poses serious challenges to current antitrust law around the world. This paper suggests an approach to regulating these 'zero-price' companies that considers the data consumers give up to use them as the 'price' they pay. The 'data as price' model acts as a starting point to assess whether consumers are being 'overcharged' by Facebook in the status quo compared to how much data they would give up in a more competitive social media landscape. By surveying thousands of participants and assessing a litany of relevant behavioural considerations, this paper finds that customers are overpaying for Facebook, and that this may come at a serious welfare cost to millions of consumers. While further analysis is warranted, there is substantial cause for concern, and for critical re-evaluation of the standards generally used by antitrust regulators around the world to regulate companies such as Facebook.

Submitted 17 June 2019; revised 15 August 2020; accepted 23 August 2020

Introduction

This paper assesses whether consumers give up too much data to 'free' technology platforms by asking the following questions:

Phase I: What are the welfare effects of current data extraction?

Phase II: What data extraction would persist in a competitive, de-shrouded social media market?

Phase III: How much better off would consumers be amidst a de-shrouded, competitive market?

This paper surveys Americans about Facebook, but can apply globally to any 'zero-price-to-consumers' company. The data show substantial consumer harm from data extraction and suggest that de-shrouding and increased

* Correspondence to: Founder, Harvard Law School Antitrust Association, Cambridge, MA, USA.
Email: msummers@jd21.law.harvard.edu

competition (as would be provided by sound privacy and antitrust regulatory measures) can help.

Srinivasan (2019) chronicles Facebook's increasing data extraction as competitors were beaten and acquired. Because of Facebook's 'zero-price' business model, 'it does not make much sense to focus the competitive assessment on prices' (Valletti, 2019). Situating 'zero-price' companies into a traditional antitrust framework requires two things: first, a careful assessment of what consumers give up, which behaviourally informed surveys are best positioned to contribute; and second, a strong finding that, under the right circumstances, competition could remedy the harm done. This paper offers both.

Definitions and concepts

Data extraction: Collecting data for third-party monetization. Data extraction is distinct from data collection, which occurs when companies collect individualized data for purposes besides third-party monetization (e.g., Netflix personalizing recommendations) (Rodriguez, 2020).

Supra-competitive extraction: More data extraction than would exist in a competitive market.

De-shrouded extraction: Transparent data extraction where consumers fully understand what data are collected and what happens post-collection.

Data as price: Measuring the amount consumers pay for goods using the data extracted from them. While prior work has discussed using data as price (Eben, 2018), none has quantified it with behaviourally informed consumer-preference surveys, nor operationalized it within antitrust law.

While data, unlike money, can be duplicated, it remains a useful proxy in an antitrust context. Other possible substitutes for money (e.g., a barter market where legal analysis is a pseudo-currency) are equally duplicable. An hour of legal analysis can be copied or widely distributed simultaneously via webinar. Data could also differ in that they are necessary for interacting with social networks (making it an input). While this is true for data collection (collecting your name/university), data extraction (third-party monetization) is severable from social media. Another possible difference is lack of scarcity. However, marginally increasing data extraction is still costly to individuals.

Others use data extraction as a quality metric (OECD, 2018). However, global antitrust law rarely employs quality theories (Sage & Hammer, 1999). Additionally, quality is reduced by features of goods themselves,

rather than what consumers give up to access them. Lastly, people generally tolerate free low-quality products more than overpriced ones.

Other models use attention (not data) as a form of price (Wu, 2017). These models, however, suffer because consumers cannot use the product without providing attention.

Survey methodology

Qualtrics/Mechanical Turk

This paper follows Campbell and Kay (2014) and Sunstein (2018) in using Qualtrics/Mechanical Turk. To avoid ordering effects, individuals were each only asked one substantive question.

Participants

A total of 2326 US participants (slightly more males than females; median age = 30–49 years) were paid \$0.17. Participants of previous studies were excluded. Questions received 100+ responses, enabling cross-question comparison.

Representativeness

Participants were recruited from the pool of US Mechanical Turk users. The sample appears fairly nationally representative of US Facebook habits. In total, 79% of respondents use Facebook daily (compared to 74% nationally), 14% weekly (compared to 17% nationally) and 7% monthly or less (compared to 10% nationally) (Pew Research Center, 2018; US Census Data).

Randomization

Randomized question order was unnecessary because each respondent received one substantive question. Answer choices were rotated.

Testing

Surveys were tested a dozen times, establishing baselines of attentive respondents.

Duration

Average duration was 70 seconds, consistent with pre-launch tests (median = 69 seconds) and with respondents only answering one substantive question and demographic questions.

Limitations

In-person experiments were infeasible for the single-substantive-question surveys, though they would have ensured participant focus. Similarly, introducing willingness-to-pay (WTP) and test questions would improve consistency and understanding. However, median values minimize individual outlier disruption. In-person replication can test robustness.

Phase I: current welfare effects of data extraction

Overview

This section measures welfare loss in the status quo compared to a hypothetical non-extractive ideal using three similar questions (to assess framing effects) (Kahneman *et al.*, 1991). While WTP and willingness-to-accept (WTA) questions can be challenging, respondents were largely social media users, avoiding challenges of familiarizing respondents with the underlying material.

Willingness to pay

The WTP question summarizes Facebook's current privacy policies (using Mahmoodi *et al.*, 2018) (Appendix A1). Then, individuals answered how much they would pay for perfectly protective data policies.

Worth to you

The next design avoided an unwillingness to pay for previously free products. Individuals imagined receiving more protective policies and stated what that change was worth to them (Appendix A2).

Willingness to accept

Respondents were told that their data were currently perfectly protected and asked their WTA more extractive policies (Appendix A3). This triggers an endowment effect by endowing strong protections (Knesier *et al.*, 2014).

Results

Demographics and results are appended (Appendices A4 & A5). The medians are fairly low (\$10–25) considering the lifelong benefits. This may be anchored by 'software as a service' offerings (Spotify, Netflix, etc.) charging ~\$10 monthly. Replication can frame payments as monthly, soliciting higher lifetime values.

Endowment effects

These findings reveal a 1.75 endowment effect for privacy (WTA/WTP). This is small compared to other studies: '[A] meta-analysis ... found an average [endowment effect] ... of 7.2 [over a range of goods]' (Knieser *et al.*, 2014). 'Superendowment effects' have also been found for Facebook use generally (Sunstein, 2018). Replication can use discrete choice experiments in order to 'avoid some of the distortions of [WTA and WTP]' (Sunstein, 2018).

An endowment effect of 1.75 is also smaller than in other privacy experiments (Acquisti *et al.*, 2014, finding a 5.7 endowment effect). However, privacy endowment 'depends critically on the context ... and how the problem is framed' (Acquisti *et al.*, 2014).

Endowment effects may be depressed because it is ineffective to endow privacy protections that people know they don't have. Distrusting the WTA question's premise (that Facebook is currently very privacy protective) reduces anger at losing privacy protections. While the analysis could be limited to people who believe Facebook is currently privacy protective, they may be preference outliers. An approach identifying people who accept the question's premise by analysing the 75th percentiles yields a (very large) 17.5 endowment effect.

Endowment effects may also be depressed because of difficulties in clawing back privacy that we have lost. Technically, disclosed data are (largely) permanent and protections are less meaningful if other providers already have similar information. Additionally, individuals might be unconcerned about Internet privacy if others are in the same boat. Historically, as well, privacy incursions are rarely reversed and 'most people ... view [privacy erosion] as inevitable' (LaFrance, 2016).

Possible harm

Notably, the WTA's 75th percentile is \$875. This staggeringly suggests that 25% of Facebook users may think that they are being injured nearly \$1000 by Facebook's *current* policies. Importantly, these individuals are not minimal Facebook users. Four-fifths of individuals in the 75th percentile use Facebook daily. When casual 'I Agree' clicks generate over \$1.5 billion in aggregate harm (Appendix A6), regulators should take notice.

This shows that individuals may be flagrantly ignorant of Facebook's data policies – a failure of law and policy. Alternatively, these individuals may feel that they gain more than \$875 in value (Sunstein, 2018).

International effects likely mirror these US results. Facebook is widely used globally. This analysis should be replicated broadly to assess the harm done to consumers around the world across diverse regulatory and privacy expectation

environments. Similarly, while respondents were over 18, children suffer equal harm from privacy infringements. Using data as price provides a unique mechanism to recover for data-related antitrust injuries against children and adults.

Considerations

Survey results can differ from real-world behaviour (Sunstein, 2018). Replications should observe market behaviour by, for example, examining discounts necessary to attract customers to extractive rewards programmes. Existing work shows consumers will pay 'about \$2.28 to conceal their browser history, \$4.05 to conceal their contacts, \$1.19 to conceal their location, \$1.75 to conceal their phone's identification number, and \$3.58 to conceal their texts' (Savage & Waldman, 2013).

Worth to you and the 'untethered effect'

Worth-to-you values, strikingly, exceeded WTA values (Appendix A5). Why?

This paper proposes an 'untethered effect', whereby valuations increase when they are untethered from real-life financial situations. WTP relates to actual ability to pay (Sunstein, 2007) and income (Brefle *et al.*, 2015), and both WTP and WTA implicitly encourage people to consider personally paying/receiving money. By contrast, worth to you solicits abstract worth without regard to ability to pay or financial need. Notably, the untethered effect here is even larger than the endowment effect.

The 'untethered effect' is important to understand and quantify. For one, disentangling valuations from wealth may reduce anti-poor biases in regulation. However, untethered estimates may produce inaccurate estimates. For example, US juries may calculate inefficient awards because they do not personally imagine paying or receiving money. Experts and attorneys should be mindful of untethered estimates (though facing real litigants may adequately tether). Additionally, further analysis can explore the impact of combining the untethered and endowment effects ('how much is it worth to you to maintain current privacy?').

Phase II: data extraction in a competitive, de-shrouded social media market

Overview

Demand and supply simulations estimated consumer drop-off and increased revenue from increased extraction. These results approximated the likely level of data extraction in a competitive, de-shrouded market.

De-shrouding

Even with more competitors, markets with misleading or mystifying pricing do not produce efficient outcomes. Current data extraction is highly shrouded: '[It takes] 244 hours per year to read [most consumers'] privacy policies' (OECD, 2018). Therefore, 97% of young adults 'consent to legal terms and services conditions without reading them' (Cakebread, 2017). Extraction 'is ubiquitous and often invisible ... [precluding] informed choices' (FTC, 2012).

Here, data extraction was de-shrouded with brief, clear explanations of data policies based on Mahmoodi *et al.* (2018).

Demand simulation

A conjoint analysis simulated a competitive social media market (Appendices A7 & A8). Respondents chose between social media platforms with different data policies and a ~5% quality difference. Five percent fewer friends on the platform was a proxy for a small but significant quality difference that might persist in a competitive market, loosely analogized from US antitrust market definition practice (Salop *et al.*, 2017).

Response rates were compared to a baseline question with a 5% quality difference but no data policy differences (in both, advertisers could target based on public posts). Analysing differences from the baseline isolates the effect of each data policy on consumer behaviour amidst a 5% quality difference.

Demand simulation results

Demographics and results (Appendices A9 & A10) show tremendous drop-off from extraction across the board. This suggests that people would flee extractive platforms if comparable options existed. However, several concerns exist. It is possible that the 5% quality difference was too small to represent actual expected quality differences under competition. Additionally, more analysis can explore the lack of (and occasionally peculiar) variation between different types of extraction.

What explains large drop-offs compared to real-life dominance of extractive platforms? For one, consumers may face quality reductions exceeding 5% to guarantee data privacy. Additionally, a focusing illusion may have inflated drop-off by adding artificial salience to parts of the online experience that are generally unimportant. People may have also rejected the premise of the questions (that the companies were otherwise identical).

Additionally, individuals may profess a commitment to privacy that they wouldn't act on: '91% of consumers think companies shouldn't be accessing their data without their consent' (Fast Company, 2019). But '[many] consumers

will not consciously spend time to select higher privacy settings for their services' (OECD, 2018). De-shrouding, however, may obviate this 'privacy paradox' (Kokolakis, 2017). Kummer and Schulte (2016) studied an app market with simplified expressions of privacy preferences, and they 'observe[d] fewer downloads for apps which ask for more privacy-sensitive permissions'.

At base, the results here should alarm regulators. The status quo dominance of extractive services may illustrate massive market failure considering that ~60% of consumers rejected even minimally extractive options amidst de-shrouding and viable alternatives.

Analysing demographic effects

Demographic questions are shown in Appendix A11. To analyse demographic effects on the demand simulation, fixed effects on the 'question #' variable helped control for the expected variation of answers between questions.

The regression was framed as follows:

- *Dependent variable*: Answer to the substantive question (coded as binary)
- *Independent variables*: Gender, race, politics, school and age
- *Fixed effects*

The results are shown in Appendix A12.

Preliminary results indicate that only age was statistically significant at a 5% level. The positive coefficient shows older respondents choosing less extractive policies over more friends, an important but unsurprising finding. It is also noteworthy that none of the other coefficients were statistically significant (even after correcting for collinearity). These findings suggest that extraction is equally unfavourable across gender, racial, educational and political lines.

Supply simulation

To estimate extraction revenue, respondents answered questions as hypothetical business owners advertising on social media (Appendices A13 & A14). Because of the difficulty of answering these questions, the supply simulation is a first approximation that should be verified with (currently unavailable) data from actual ad buyers. Categories addressing more complex extraction necessitated modified framing (Appendix A15).

Demographics and results are appended, but they are primarily useful for integration with the demand simulation (Appendices A16 & A17).

Integrating results from the demand and supply simulations

To conceptualize integrating the simulations, imagine 100 Facebook users, each of whom generates \$27.61 of advertising revenue (Flynn, 2018). Total

quarterly revenue equals \$2761. If increasing data extraction reduces demand by 20%, 80 users remain (in a simplified model). If that same policy increases revenue 50%, Facebook users each generate \$41.42 in ad revenue. After considering user drop-off and revenue increases, \$3313 of total revenue results. The ~\$500 revenue increase (from \$2,761) suggests that Facebook would rationally adopt the policy.

For this analysis, the baseline result was that 83% of consumers chose the platform with more friends when data policies were identical. Additional work is needed to understand this outcome, but for now, assume 17% of people preferred a platform with fewer friends (perhaps for anonymity or exclusivity).

The results shown in Appendix A18 indicate that nearly every increase in extraction reduces total revenue. The only revenue-positive extraction is 'age, gender and city' data (contingent on using the mean value). If the median is more accurate (which it likely is, in order to minimize outlier influence), even rudimentary data extraction may be unprofitable in a competitive, de-shrouded market.

This suggests that widely used data-extraction business models may be based on deception, oligopoly or both. It would also imply that the values obtained from Phase I of this analysis regarding consumer harm (to the tune of billions of dollars) may accurately represent consumer harm.

Replication and limitations

Replications should experiment with quality gaps exceeding 5% and use real ad-revenue data (if available). Replications should also explore the additive effects of different extractions. A company using certain extractive techniques (that caused some users to leave) may not lose additional users by implementing new extractive techniques (if sensitive consumers have already left). Further analysis should also ask whether extraction would cause consumers to switch from platforms they are *already using* (rather than assessing new platform choices). Myriad other challenges exist when modelling highly complex industries replete with network effects, stickiness, minimal data portability, etc.

One key objection is that companies may become more extractive when facing steeper competition (to preserve revenue). This makes sense in the heavily shrouded status quo. However, in a de-shrouded data extraction environment, as modelled in the demand simulation, consumers who receive understandable data policies and have comparable-quality alternatives steer clear of data extraction. The integration of the supply and demand simulations suggests that a de-shrouded, competitive environment could yield a dramatically less extractive equilibrium than the status quo.

Phase III: assessing consumer harm under de-shrouded competition

The Phase I survey was modified to study the likely harm under the extraction policies that Phase II suggests came closest to being viable in a competitive, de-shrouded market. Individuals were faced with one WTP (Appendix A19) or WTA (Appendix A20) question. Demographics and results are shown in Appendices A21 and A22.

To compare consumer harm under the status quo to consumer harm under de-shrouded competition, Phase III values were subtracted from Phase I values (Appendix A23). The results are striking and discordant with expectations. It appears that for half of the relevant measures, individuals found extracting 'the age, city and gender you listed on your profile' to be more harmful than extracting a wider range of personalized information (Appendix A24).

These peculiar outcomes likely occurred because age, gender and city feel especially concrete to consumers. It is easier for consumers to imagine their age, gender and city being extracted than picturing every post and message they ever wrote. Future research should explore the impact of the focusing effect and the ways in which policy wordiness and concreteness affect outcomes.

Conclusion

This paper finds that there would be dramatically less extraction (and likely less harm) in a world with de-shrouded data policies and more robust social media competition. Regulators should think concretely about analysing data as price in order to leverage antitrust law and behavioural insights towards regulating zero-price markets. Shrouding and insufficient competition in the market for data are concerning, and behaviourally informed antitrust interventions have the potential to improve consumer welfare dramatically.

These findings stand apart from the existing literature. First, this paper is the first to calculate the value that consumers put on Facebook privacy. Second, no prior work has modelled consumer choice amidst a competitive, de-shrouded social media market.

This paper constitutes a jumping-off point for regulators around the world regulating zero-price companies that do not fit neatly into existing antitrust frameworks. While data do not represent a perfect proxy for price, they provide an analytically useful tool that empowers regulators to leverage traditional antitrust tools in new competition environments.

References

- Acquisti, A., L. John and G. Loewenstein (2014), 'What Is Privacy Worth?', *Journal of Legal Studies*, 42(2): 249–274.

- Breffle, W.S., M.E. Eiswerth, D. Muralidharan and J. Thornton (2015), 'Understanding how income influences willingness to pay for joint programs: A more equitable value measure for the less wealthy', *Ecological Economics*, 109: 17–25.
- Cakebread, C. (2017), 'You're not alone, no one reads terms of service agreements', *Business Insider*. Retrieved from <https://www.businessinsider.com/deloitte-study-91-percent-agree-terms-of-service-without-reading-2017-11?r=US&IR=T>
- Campbell, T.H. and A.C. Kay (2014), 'Solution Aversion: On the Relation between Ideology and Motivated Disbelief', *Journal of Personality and Social Psychology*, 107(5): 809–824.
- Eben, M. (2018), 'Market Definition and Free Online Services: The Prospect of Personal Data as Price', *IS: A Journal of Law and Policy*, 14(2): 227.
- Fast Company (2019), 'Instead of surveillance, what if we told advertisers what they wanted to know?'. Retrieved from <https://www.fastcompany.com/90324911/could-zero-party-data-end-the-scourge-of-third-party-data>
- Flynn, K. (2018), 'Facebook's making more money per user in North America than ever before', *Digiday*. Retrieved from <https://digiday.com/marketing/facebooks-making-money-per-user-north-america-ever/>
- FTC (2012), 'FTC Report, Protecting Consumer Privacy in an Era of Rapid Change', 1, 2. Retrieved from <https://www.ftc.gov/reports/protecting-consumer-privacy-era-rapid-change-recommendations-businesses-policy-makers>
- Kahneman, D., J.L. Knetsch and R.H. Thaler (1991), 'Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias', *Journal of Economic Perspectives*, 5(1): 193–206.
- Knieser, T.J., W.K. Viscusi and J.P. Ziliak (2014), 'Willingness to Accept Equals Willingness to Pay for Labor Market Estimates of the Value of Statistical Life', *Journal of Risk and Uncertainty*, 48: 187–205.
- Kokolakis, S. (2017), 'Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon', *Computers & Security*, 64: 122–134.
- Kummer, M. and P. Schulte (2016), 'When Private Information Settles the Bill: Money and Privacy in Google's Market for Smartphone Applications', *Management Science*, 65(8): 3470–3469.
- Lafrance, A. (2016), 'The Convenience–Surveillance Tradeoff', *The Atlantic*. Retrieved from <https://www.theatlantic.com/technology/archive/2016/01/the-convenience-surveillance-tradeoff/423891/>
- Mahmoodi, J., J. Čurdová, C. Henking, M. Kunz, K. Matic, P. Mohr and M. Vovko (2018), 'Internet Users' Valuation of Enhanced Data Protection on Social Media: Which Aspects of Privacy are Worth the Most?', *Frontiers in Psychology*, 9: 1516.
- OECD (2018), 'Quality considerations in digital zero-price markets', Directorate for Financial and Enterprise Affairs Competition Committee. Retrieved from [https://one.oecd.org/document/DAF/COMP\(2018\)14/en/pdf](https://one.oecd.org/document/DAF/COMP(2018)14/en/pdf)
- Pew Research Center (2018), 'Social Media Fact Sheet', 5 February. Retrieved from <https://www.pewresearch.org/internet/fact-sheet/social-media/>
- Rodriguez, A. (2020), 'Netflix gave a detailed explanation of why it won't sell advertising, but it's still working with brands in other ways', *Business Insider*. Retrieved from <https://www.businessinsider.in/tech/news/netflix-gave-a-detailed-explanation-of-why-it-wont-sell-advertising-but-its-still-working-with-brands-in-other-ways/articleshow/73526843.cms>
- Sage, W.M. and P.J. Hammer (1999), 'Competing on Quality of Care: The Need to Develop a Competition Policy for Health Care Markets', *University of Michigan Journal of Law Reform*, 32: 1069–1118.
- Salop, S.C., S. Moresi and J.R. Woodbury (2017), *Market Definition*. Washington, DC: Georgetown Law Faculty Publications.
- Savage, S.J. and D.M. Waldman (2013), 'The Value of Online Privacy', Working paper. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2341311

- Srinivasan, D. (2019), 'The Antitrust Case Against Facebook', *Berkeley Business Law Journal*, 16: 39–101.
- Sunstein, C. (2007), 'Willingness to Pay versus Welfare', *Harvard Law Policy Review*, 1: 303–330.
- Sunstein, C. (2018), 'How Much Would You Pay to Use Facebook? A Behavioral Perspective', Harvard Public Law Working Paper No. 18-30.
- Valletti, T. (2019), 'Online Platforms and Market Power Part 3: The Role of Data and Privacy in Competition', House Judiciary Committee Subcommittee on Antitrust, Commercial, and Administrative Law. Retrieved from <https://www.congress.gov/116/meeting/house/110098/witnesses/HHRG-116-JU05-Wstate-VallettiT-20191018.pdf>
- Wu, T. (2017), 'Blind Spot: The Attention Economy and the Law', *Antitrust Law Journal*, 82: 3.

Appendices

A1. Phase I: willingness to pay

- Facebook stores data on everything that you:
 - Create and share while using Facebook.
 - Privately send on Messenger.
- Duration:
 - Facebook stores data as long as it sees fit.
 - Information others share about you will not be deleted even if you delete your account.
- Third-party sharing
 - Facebook exchanges information about you with third-party apps and websites (e.g., free Wi-Fi and apps where you use 'Login with Facebook').
 - Facebook shares your information with companies it owns (e.g., Instagram and WhatsApp).

How much would you be willing to pay (as a one-time payment in US dollars) for Facebook to change their data policies regarding your account such that:

- You can delete any data stored about you at any time?
- Your data isn't shared with third-party vendors without your permission?
- You can control what data is collected about you on and off Facebook?

This analysis excludes back-end data security. Data security is complex and absolutes (e.g., 'Your data will be perfectly secure') are largely unachievable in the real world.

The ‘regarding your account’ language makes clear that the policy changes would affect only the user’s account (avoiding free-riding or altruistic effects).

A2. Phase I: worth to you

How much would it be worth to you (as a one-time amount in US Dollars) if Facebook changed their data policies regarding your account such that:

A3. Phase I: willingness to accept

You recently learned the following about your Facebook account:

- You can delete any data stored about you at any time.
- Your data isn’t shared with third party vendors without your permission.
- You can control what data is collected about you on and off Facebook.

What is the lowest amount (as a one-time payment in US Dollars) that you would accept to allow Facebook to replace the above data policies with the following:

- Facebook will store data on everything that you:
 - Create and share while using Facebook.
 - Privately send on Messenger.
- Duration:
 - Facebook will store data as long as it sees fit.
 - Information others share about you will not be deleted even if you delete your account.
- Third-party sharing
 - Facebook will exchange information about you with third-party apps and websites (e.g., free Wi-Fi and apps where you use ‘Login with Facebook’).
 - Facebook will share your information with companies it owns (e.g., Instagram and WhatsApp).

The information provided in this section is intentionally not an accurate representation of Facebook’s data policies (in order to endow a right to privacy). Alternative constructions of the survey can validate that this mismatch with reality did not influence results. Note that a disclaimer was

provided at the end of the survey to avoid spreading misinformation about data policies.

A4. Phase I: demographics and results for other platforms

Total respondents (<i>n</i>)	
With demographic data	501
Respondents without demographic data (excluded from analysis)	0
Age	
18–29	176
30–49	244
50–64	67
65+	14
Gender	
Female	238
Male	258
Non-binary	5
Race	
Black	37
White	374
Latinx	16
Asian	38
Native Hawaiian or Pacific Islander	0
American Indian or Alaska Native	9
Other	6
Multi-racial	21
Facebook account	
Yes, I have one	435
No, I never have	9
I did have one, but I deleted/deactivated it	57

Results for other platforms

Users who had deleted/deactivated their Facebook accounts were also asked to ‘Please briefly explain why you stopped using Facebook’. Interestingly, many users who stopped using Facebook but now primarily use Instagram seemed to have left Facebook due to concerns surrounding privacy and data security. One such user specifically stated: ‘I don’t trust them with my data’. This is fascinating, considering that Instagram is owned by Facebook and the two share data. Unfortunately, the 11 respondents who don’t have a Facebook account and now primarily use Instagram did not offer a statistically reliable sample to make inferences from regarding their willingness to pay (WTP) or willingness to accept (WTA) different data policies.

Results for other users who don't have Facebook accounts also mentioned privacy concerns on Facebook, but may have been primed by the earlier question regarding their WTP/WTA for new data policies. Sixteen respondents who don't have Facebook accounts primarily use Twitter. Many of these users referenced privacy concerns for why they stopped using Facebook. One colourfully wrote: 'I don't trust Zuckerberg and his minions, and I believe Facebook is not what it ostensibly stands for, but rather a gigantic data mining operation'.

Of the 57 people who used to have a Facebook account but deleted/deactivated it, only 4 don't use one of Reddit, Twitter, Snapchat or Instagram. This strongly suggests that social media platforms are substitute goods, or at least that when users stop using one, they shift consumption to another platform.

In addition, Facebook's huge presence in the social media market is represented by the fact that only 9 of the 500+ respondents had never had a Facebook account. Of those 9, the only 3 who use another social media platform use Reddit. Only 2 of those 9 were between 18 and 29 years old. One side effect of this is that it might be challenging to study the behaviour of non-Facebook users on Mechanical Turk without soliciting them specifically, since only 9 of a random sample of over 500 respondents have never had an account.

A5. Phase I: baseline results

	Duration (seconds)	Standard WTP	Worth to you	WTA
Median	63	\$10	\$25	\$18
25th percentile	43	\$1	\$5	\$1
75th percentile	93	\$50	\$100	\$875
Observations (<i>n</i>)	435	111	99	110

A6. Phase I: results calculating possible harm

	Standard WTP	WTA
Median	\$1,721,947,640	\$3,013,408,370

Approximately 172 million US adults use Facebook (Pew Research Center 2018; US Census Data).

A7. Phase II: conjoint analysis

You are signing up for a social media platform. Which of these do you choose (assuming other attributes are identical)?

- Social media platform with the following attributes:

Quality:

- 90% of your friends use this platform.

Data policy:

- Advertisers CAN target you based on public posts you make.

- Social media platform with the following attributes:

Quality:

- 85% of your friends use this platform

Data policy:

- Advertisers CANNOT target you based on public posts you make

Framing the question as whether advertisers can target the user made it possible to 'bake in' the effect of consumers who prefer to get more targeted ads while still studying the primary effect: discomfort with extraction.

Further research should explore the effect of different approaches to de-shrouding data policies in order to determine which leads to the most efficient market outcome in terms of consumer welfare. More radical methods of de-shrouding (which might be more manipulative than informative) include listing salient stories of negative consequences of targeted advertising or showing examples of things people are surprised to find out Facebook knows about them when they download the data Facebook has collected on them. These examples likely push the barrier between creating salience and manipulation too far.

A8. Phase II: conjoint analysis – alternative framings

- Advertisers CANNOT target you based on public posts you make.
- Advertisers CANNOT target you based on private messages you send.
- Advertisers CANNOT target you based on your activity on other platforms owned by the same company.
- Your data IS NOT exchanged with third-party vendors (e.g., apps, websites and Wi-Fi you log into using your account).
- Information others share about you WILL be deleted when you delete your account.
- You CAN select posts to delete permanently.
- Advertisers CANNOT target you based on the age, gender and city you list on your profile.

A9. Phase II: demand simulation demographics

Total respondents (<i>n</i>)	
With demographic data	855
Respondents without demographic data (excluded from analysis)	0
Age	
18–29	337
30–49	389
50–64	107
65+	22
Gender	
Female	440
Male	411
Non-binary	4
Race	
White	591
Black	86
Latinx	29
Asian	87
Native Hawaiian or Pacific Islander	1
American Indian or Alaska Native	7
Other	5
Multi-racial	49
Political views	
0 (extremely liberal)	67
1	103
2	119
3	167
4	150
5	108
6	97
7 (extremely conservative)	44
Education	
High school	246
Associate's degree	120
Bachelor's degree	381
Graduate school	108
Social media use	
Facebook	101
Twitter	7
Instagram	17
Snapchat	1
Reddit	32
Facebook and at least one other platform above	577
Multiple of the above (not including Facebook)	92
None	28

A10. Phase II: demand simulation results

Policy	Percentage who chose more friends	Decrease from baseline
Baseline	83%	0%
Public posts	21%	62%
Private messages	19%	64%
Other owned companies	34%	49%
Third-party vendors	19%	64%
Delete posts about you	19%	64%
Delete your posts	19%	64%
Age, gender and city	26%	57%

A11. Phase II: demographic questions

What is your age?

- 18–29
- 30–49
- 50–64
- 65+

What is your gender?


- Male
- Female
- Non-binary/non-conforming

What is your race? (check all that apply)

- White
- Black
- Latinx
- Asian
- Native Hawaiian or Pacific Islander
- American Indian or Alaska Native
- Other

For statistical analysis, individuals who checked more than one box were grouped into a back-end mixed-race category.

Here is a 7-point scale on which the political views that people might hold are arranged from extremely liberal (0) to extremely conservative (7). Where would you place yourself on this scale?



Please select the highest level of school that you have completed:

- High school
- Associate's degree
- Bachelor's degree
- Graduate school

A12. Phase II: demographic effects

Observations: 749

Groups: 7

Observations per group: 103–110

	Coefficient	Standard error	<i>t</i>	<i>P</i> > <i>t</i>	95% confidence interval	Answer
Gender	0.01	0.03	0.19	0.850	−0.05	0.06
Race	0.00	0.00	0.40	0.691	−0.00	0.00
Politics	−0.02	0.01	−2.37	0.018	−0.03	−0.00
School	−0.01	0.01	−0.95	0.344	−0.03	0.01
Age	0.08	0.02	3.94	0.000***	0.04	0.12
Constant	1.72	0.07	24.37	0.000	1.58	1.85

****p* < 0.001.

A13. Phase II: supply simulation prompt

You manage the social media marketing of your medium-sized business. The social media platform you advertise on now allows you to target customers **based on their public posts**. How much, if at all, would you increase your advertising spending (as a percentage of current spending)?

Again, each participant was given a blank text box to put their answer into (which was validated to make sure all answers were numeric) and respondents were only allowed to answer one substantive question.

To simplify the scenario, advertisers were not told what substitute advertising options were available to them, so a fully competitive market was not simulated. Similarly, it might be useful in the future to ask individuals to estimate downward changes to advertising spending based on losing targeting mechanisms. This would have the added benefit of bounding the maximum decrease to 100% (but might be affected by an endowment effect). Ultimately, respondents were not given an operating budget to work from or a current value of advertising spending to avoid anchoring around particular values. However, future analysis could do this while varying the values that might be anchored on.

A14. Phase II: supply simulation prompt alternatives

- ... based on their public posts.
- ... based on their private messages.
- ... based on their activity on other platforms owned by the social media company.
- ... based on data gained from third-party vendors (e.g., apps, websites and Wi-Fi the individual signs into using their account).
- ... based on the age, gender and city they list on their profiles.

A15. Phase II: supply complex extraction prompt

You manage the social media marketing of your medium-sized business. The social media platform you advertise on no longer allows customers to select posts to delete permanently. This allows you to continue targeting individuals **based on posts they would have deleted**. How much, if at all, would you increase your advertising spending (as a percentage of current spending)?

You manage the social media marketing of your medium-sized business. The social media platform you advertise on has had a recent change in policy. When individuals delete their accounts, **posts that friends made about them are no longer deleted**. This allows you to target customers based on these posts, which would otherwise be deleted. How much, if at all, would you increase your advertising spending (as a percentage of current spending)?

A16. Supply simulation demographics

Total respondents (<i>n</i>)	
With demographic data	663
Respondents without demographic data (excluded from analysis)	0
Age	
18–29	243
30–49	322
50–64	78
65+	20
Gender	
Female	347
Male	311
Non-binary	5
Race	
White	469
Black	49
Latinx	34
Asian	71
Native Hawaiian or Pacific Islander	1
American Indian or Alaska Native	5
Other	8
Multi-racial	26
Political views	
0 (extremely liberal)	49
1	98
2	83
3	124
4	123
5	97
6	66
7 (extremely conservative)	23
Education	
High school	180
Associate's degree	103
Bachelor's degree	274
Graduate school	106
Social media use	
Facebook	97
Twitter	10
Instagram	8
Snapchat	2
Reddit	14
Facebook and at least one other platform above	451
Multiple of the above (not including Facebook)	62
None	19

A17. Phase II: supply simulation results

Policy	Median	5th percentile	25th percentile	75th percentile	95th percentile
Public posts	20	0	10	40	94
Private messages	20	0	10	43	90
Other owned companies	20	0	10	30	79
Third-party vendors	20	4	10	31	142
Delete posts about you	10	0	5	21	92
Delete your posts	10	0	3	25	50
Age, gender and city	20	0	10	40	1,150

A18. Phase II: integrated results

	Users	Using mean			Using median		
		Revenue per user	Total revenue	Change from starting point	Revenue per user	Total revenue	Change from starting point
Starting point	100	\$27.61	\$2761	0%	\$27.61	\$2761	0%
Public posts	38	\$35.78	\$1370	-50%	\$33	\$1268	-54%
Private messages	36	\$40.23	\$1464	-47%	\$33	\$1206	-56%
Other owned companies	51	\$37.91	\$1943	-30%	\$33	\$1698	-39%
Third-party vendors	36	\$60.90	\$2194	-21%	\$33	\$1194	-57%
Delete posts about you	36	\$43.11	\$1553	-44%	\$30	\$1094	-60%
Delete your posts	36	\$71.34	\$2586	-6%	\$30	\$1101	-60%
Age, gender and city	43	\$348.01	\$15,084	446%	\$33	\$1436	-48%

Total revenue was calculated using 100 users to make the numbers manageable, but the number of users mathematically has no effect on the results.

The mean values were somewhat distorted by right tail outliers, so all calculations were done (as indicated) with both mean and median values.

A19. Phase III: willingness to pay

- Facebook ONLY lets advertisers target you based on the age, gender and city you list on your profile.
- Facebook does NOT let advertisers target you based on:
 - Public posts
 - Private messages
 - Third-party data
 - Information from Instagram and WhatsApp
- Duration:
 - Facebook PERMITS you to delete any posts by or about you at any time.

How much would you be willing to pay (as a one-time payment in US dollars) for Facebook to change their data policies regarding your account such that:

- Advertisers will NOT be able to target you based on the age, city and gender you listed on your profile?

A20. Phase III: willingness to accept

You recently learned the following about your Facebook account:

- Facebook does NOT let advertisers target you based on:
 - Public posts
 - Private messages
 - Third-party data
 - Information from Instagram and WhatsApp
- Duration:
 - Facebook PERMITS you to delete any posts by or about you at any time.

What is the lowest amount (as a one-time payment in US Dollars) that you would accept to allow Facebook to add the following data policy to the above:

- Advertisers can target you based on the age, gender and city you list on your profile?

A21. Phase III: demographics

Total respondents (<i>n</i>)	
With demographic data	260
Demographic data not solicited (non-social media users)	47
Age	
18–29	79
30–49	150
50–64	26
65+	5
Gender	
Female	148
Male	112
Non-binary	0
Race	
White	207
Black	22
Latinx	9
Asian	14
Native Hawaiian or Pacific Islander	0
American Indian or Alaska Native	0
Other	0
Multi-racial	8
Political views	
0 (extremely liberal)	29
1	43
2	28
3	35
4	42
5	37
6	31
7 (extremely conservative)	15
Education	
High school	66
Associate's degree	34
Bachelor's degree	128
Graduate school	32

A22. Phase III: willingness to pay and willingness to accept results

	Mean	Median	5th percentile	25th percentile	75th percentile	95th percentile
WTP	\$25	\$2	\$0	\$0	\$10	\$50
WTA	\$2172	\$25	\$0	\$5	\$100	\$5000

A23. Phase III: results – harm compared to the status quo

	Mean	Median
Standard WTP	-\$9	\$8
WTA	\$10,223	-\$16

A24. Phase III: more extensive extraction

- Facebook will store data on everything that you:
 - Create and share while using Facebook.
 - Privately send on Messenger.
- Duration:
 - Facebook will store data as long as it sees fit.
 - Information others share about you will not be deleted even if you delete your account.
- Third-party sharing
 - Facebook will exchange information about you with third-party apps and websites (e.g., free Wi-Fi and apps where you use ‘Login with Facebook’).
 - Facebook will share your information with companies it owns (e.g., Instagram and WhatsApp)