

The Digital Welfare of Nations: New Measures of Welfare Gains and Inequality

Erik Brynjolfsson 
Stanford University
and NBER

Avinash Collis 
Carnegie Mellon
University

Asad Liaqat
Meta

Daley Kutzman
Meta

Haritz Garro
Meta

Daniel Deisenroth
Meta

Nils Wernerfelt*
Northwestern
University

Abstract

Digital goods can generate large benefits for consumers, but these benefits are largely unmeasured in the national accounts, including GDP and productivity. In this paper, we measure welfare gains from 10 popular digital goods across 13 countries by conducting large-scale incentivized online choice experiments on representative samples of nearly 40,000 people. We estimate that these goods—many of which are free to users—generate over \$2.5 trillion in aggregate consumer welfare across these countries per year, which is roughly equivalent to 6% of their combined GDP. We find that lower-income individuals and lower-income countries obtain relatively larger welfare gains from these goods compared to higher-income individuals and countries. This suggests that digital goods may reduce inequality in welfare within and across countries by disproportionately benefiting lower-income groups.

1. Introduction

Digital technologies create challenges for economic measurement. On one hand, with the spread of the Internet, time spent on digital goods has increased dramatically, and these goods are affecting more and more aspects of daily life. For instance, the average person in both the US and the UK now spends almost 24 hours a week online (Coyle and Nakamura, 2022), while the number of photos shared increased from 350 billion to 2.5 trillion between 2010 and 2017

*Authors listed in order of contribution with E.B. and A.C. contributing equally. A.L. and D.K. led the project internally at Meta. A.L., D.K., and H.G. designed and implemented the survey, and analyzed and validated the data. A.L. and A.C. wrote the paper, with contributions from E.B., D.K., H.G., D.D., N.W.; E. B. and A. C. were the joint external principal investigators; D. D. and N. W. were the joint internal principal investigators. We thank Shane Greenstein and Sonny Tambe for helpful comments.

Competing interests: A.L., D.K., H.G., and D.D. are employees of Meta Platforms and hold a financial interest in Meta. N.W. was an employee of Meta while this research was conducted though no longer holds a financial interest in the company. All other authors declare no current competing interests, although E.B. and A.C. have in the past been awarded unrestricted gifts from Meta.

 Corresponding authors. Email: erikb@stanford.edu; avinashcollis@cmu.edu.

(Brynjolfsson et al., 2019b). On the other hand, the officially measured size of the information sector as a share of total GDP has remained almost unchanged at around 4-5% for the past four decades. The discrepancy is a reflection of the fact that, regardless of the value they create for consumers, most digital goods are available at zero price (Brynjolfsson and Collis, 2019).¹

To better understand the welfare effects of digital goods, new metrics are needed to complement existing production-based metrics such as GDP and productivity (Masood, 2022).² Recent research has proposed new ways of directly measuring consumer surplus from digital goods using massive online choice experiments (Brynjolfsson et al., 2019a; Brynjolfsson et al., 2019b). However, existing research in this area has only looked at a select few regions and goods using convenience samples that were not representative of the population of digital users. Moreover, how these welfare gains from digital goods are distributed across different income levels is unexplored.

In this paper, we conduct large-scale incentivized choice experiments involving nearly 40,000 representative users of the Facebook digital service in 13 countries to estimate the welfare gains generated by 10 popular digital goods. While previous research looks at specific digital goods in a single country using smaller non-representative samples (Allcott et al., 2020; Brynjolfsson et al., 2019a), the scale and scope of our sample enable us to estimate valuations for ten leading digital goods with sufficient precision to compare welfare gains from digitization across countries. Similar contingent valuation methods have been used in the past to value non-market goods, including environmental goods (Bishop et al., 2017) and accepted as evidence in legal cases (Lowensohn, 2012).

We find that digital goods generate substantial welfare for consumers across these countries. Specifically, our analysis implies that digital goods create \$2.52 trillion of value across all 13 countries, which corresponds to 5.95% of their aggregate GDP. We also find that lower-income individuals within countries and lower-GDP countries obtain disproportionately more welfare from these digital goods compared to higher-income individuals and higher-GDP countries—not only relative to income but in some instances even in absolute terms. Because the free digital goods we examine (e.g. search engines or instant messaging platforms) are available for free to both higher- and lower-income individuals, they serve to reduce welfare inequality both within and across countries in our sample.

¹ Most are instead supported by advertising revenues, while others are supported by volunteers.

² Production based metrics may also not properly include welfare gains from non-digital goods. For e.g., Nordhaus (1996) calculates that lighting generated \$275 billion in consumer surplus between 1800 and 1992, and this is not measured in existing metrics. Jones and Klenow (2016) propose a more comprehensive measure of welfare accounting for consumption, leisure, mortality, and inequality. In this work we focus on digital goods because the changes in welfare over the past few decades may be the highest for digital goods compared to other types of goods.

Since none of the digital goods we examine existed a few decades ago, our findings suggest that economic growth may have been underestimated by conventionally measured GDP.³ Furthermore, since labor productivity is typically defined as GDP per hour worked, it also fails to reflect the full contribution of digital goods. As the digital economy becomes relatively more important, it will be increasingly important to explicitly measure the trillions of dollars of value created by these sorts of goods. The reduction in inequality is consistent with the idea that greater access to free public goods can be effective in reducing inequality (Fischer, 2017) and that relative price changes have an important effect on inequality (Slottje, 1987). In particular, as argued by Goolsbee and Klenow (2006), low-wage workers should be expected to consume relatively higher quantities of free digital goods, not only because they have less money to spend on other goods but also because their opportunity cost of time may be lower.

The rest of the paper is organized as follows. Section 2 provides background on other measurement statistics and past literature, Section 3 describes our study and methods in depth, Section 4 discusses the results, and Section 5 concludes.

2. Background

It has long been recognized that traditional measures of productivity only measure a subset of outcomes that social planners would care about.⁴ In response, various initiatives have sought to provide a more comprehensive assessment of well-being. For example, the United Nations Sustainable Development Solutions Network, since 2012, publishes an annual World Happiness Report, evaluating countries based on happiness metrics. Jones and Klenow (2016) propose a measure considering consumption, leisure, mortality, and inequality for gauging economic well-being. Despite progress in subjective well-being and life satisfaction measurement (Krueger and Stone 2014), there is a significant gap in achieving a consensus among macroeconomists for reliable policymaking indicators (den Haan et al. 2017).

This paper diverges from these efforts, opting for a more traditional microeconomic approach. Specifically, it explores measuring consumer welfare gains resulting from digital goods consumption by leveraging methods to directly elicit consumers' valuations of products (Wertenbroch and Skiera, 2002). We apply these methods to a variety of digital goods at scale, allowing us to estimate their valuations across a large sample of users. Brynjolfsson et al. (2019a) highlight that the influence of the information age is conspicuously absent from GDP statistics, a sentiment echoed by the survey of leading macroeconomists (den Haan et al.

³ As discussed in Brynjolfsson et al (2019b), the fact that some of these goods are paid for by advertising or other means does not fundamentally change this fact. That said, it is also possible, even likely, that economic growth was misestimated in earlier eras as well, as other unmeasured or poorly measured goods and services entered (and exited) the economy. Our paper is only one in a long line of efforts to improve on past measures.

⁴ See, for example, Robert Kennedy's speech at the University of Kansas in 1968, where he said "Gross National Product counts air pollution and cigarette advertising...[yet] does not allow for the health of our children, the quality of their education or the joy of their play" (<http://www.jfklibrary.org/Research/Research-Aids/Ready-Reference/RFK-Speeches/Remarks-of-Robert-F-Kennedy-at-the-University-of-Kansas-March-18-1968.aspx>).

2017). The paper confronts the fundamental issue that GDP and productivity are not direct measures of welfare and well-being.

The gap between GDP and well-being has persisted since the former's inception, becoming a more substantial concern in the current digital era. For example, taking the music industry as a case study, the shift from physical to digital goods challenges the adequacy of GDP in reflecting well-being (Brynjolfsson and Saunders (2009)). Despite the decline in industry revenue, consumers enjoy more and arguably better music (Waldfogel (2018)), challenging the traditional correlation between GDP and well-being. Taking the encyclopedia industry as another case study, as we shifted from purchasing physical copies of Britannica to accessing Wikipedia for free online, consumers are clearly better off since they get access to better quality and more quantity of information at a lower price of zero (Brynjolfsson et al. (2018)).

Recent work by Brynjolfsson et al. (2019a) estimates changes in consumer surplus from digital goods. They propose a methodology involving choice experiments to directly measure consumer surplus from digital goods. By eliciting consumers' valuations through choices between keeping a digital good or accepting a monetary equivalent, this approach addresses the shortcomings of market data where the price of many digital goods is zero. This paper builds on their approach to conduct the first-ever large-scale cross-country measurement of consumer welfare for the ten most popular goods across 13 countries. By partnering with a major digital platform, we estimate valuations using an incentive-compatible approach on representative samples of users.

3. Study and Methods

We surveyed 39,717 Facebook users in 13 countries on the Facebook internal survey platform from March 25 to April 07 2022. The two main components of the survey were (i) a best-worst scaling task (N=23,752) where users select their most and least preferred options from a list (Louviere et al., 2015), and (ii) an incentivized willingness-to-accept measurement (N=39,717) using single binary discrete choice experiments (Brynjolfsson et al., 2019a). The former gives us relative valuations across digital goods, while the latter gives us valuations for our benchmark good using a revealed preference approach.⁵ In this paper, we use Facebook's valuation as the benchmark to calibrate valuations of the other goods.⁶

⁵ These are well-established methods for valuing non-market goods. Note that there are other valuation methods available, including, for example, Becker-DeGroot-Marschak's willingness-to-pay approach ("BDM" - Becker et al., 1964). Prior to running our main survey, we piloted several such methods and found that our preferred approach worked best.

⁶ These valuations capture the consumer surplus generated by these goods since they measure consumers' valuation of digital goods net of costs associated with consuming the goods. In most instances, the costs associated with consuming digital goods are very small or nonexistent as most users have a phone with an internet plan. However, some users might purchase a phone with an internet plan with the main purpose of being able to use Facebook or a similar app. In these cases, the purchase of the phone and internet plan will be reflected in the GDP numbers, and our measure of consumer surplus will be an underestimate.

The survey sample for each component was weighted to be representative of the population of “monthly active” Facebook users in each country (i.e., users who have been active on Facebook within the last 30 days). The survey invitation was shown to respondents at the top of their Facebook News Feed (Appendix Figure 1). Facebook has nearly 3 billion monthly active users, and using a sample that is representative of Facebook users in each country is a stark improvement over existing estimates that were based on laboratory experiments or off-platform surveys of the general public.⁷

In the following section, we first discuss our sampling frame and weighting procedures; we then provide details on the design of the best-worst scaling and the willingness-to-accept methodologies.

Sampling & Weighting

We recruited our sample by sending in-app survey invitations to Facebook users. All 18+ Facebook users in these 13 countries who had been active on the Facebook platform in the month before the start of the survey and whose accounts had been created at least 30 days ago were eligible to be included in the sampling frame. The 13 countries included in our study are the United States, Canada, Mexico, Germany, the United Kingdom, Ireland, France, Belgium, Norway, Spain, Romania, Japan, and Korea. We selected these 13 countries based on a combination of research interest and the availability of survey pool resources in each country.

At a high level, Meta routinely runs surveys for advertisers and for internal purposes. We follow similar sampling and weighting protocols as to what is used in these instances (e.g., see Athey et al., 2023): we employ non-response and design weights to ensure our sample is representative of the relevant populations. First, we use inverse probability weighting models to account for selection in our sample, such that all our estimates are representative of the population of Facebook users in each country. Our ability to weight results to match populations of Facebook users is an improvement over existing work. For all analyses, we account for both unit and item non-response bias. Second, the probability a monthly active user is in our sampling frame varies across countries; to correct for this, we use design weights for any analysis that pools data across countries. Further details on weighting are available in the online appendix.

⁷ The monthly active user (MAU) population in Facebook in these 13 countries constitutes a large share of the total population. The average Facebook penetration rate (defined as MAU divided by population) among these countries is 59.4%. These users are likely to be fairly representative of the users of digital technologies in each of these countries, though not necessarily of people who do not use any digital goods.

Measuring Relative Welfare from Digital Goods using Best-Worst Scaling

We used a best-worst scaling (BWS) methodology to measure the relative willingness to accept for stopping the use of 10 digital goods or not meeting friends in person for one month. The digital goods include social media / messaging tools (Facebook, Twitter, Instagram, WhatsApp, Snapchat, and TikTok), and other digital tools (Google Search, Google Maps, YouTube, and Amazon Shopping). We chose the most popular digital goods in our sample countries and restricted our attention to goods that exist in all of our sample countries so that we can do cross-country comparisons. Since many non-digital activities also generate welfare, we chose meeting friends in person as an example of a non-digital activity to benchmark our valuations of digital goods and compare the relative valuation of in-person communication versus online social networks.

Given this list of digital goods (and the meeting friends in person option), the next task is to derive sets of questions to ask individual users. Using a balanced-incomplete block design to ensure that all pairs of goods are evaluated together sufficiently (Hanani, 1961), we generated 70 questions with the 11 items mentioned above plus 10 monetary amounts.⁸ We dropped 7 questions that included only monetary amounts, resulting in 63 questions from which we presented 3 random questions to each respondent. Within each question, respondents were asked to indicate their ‘best’ and ‘worst’ options from among three options presented to them. Screenshots of the survey invitations and specific questions are in the Appendix (e.g., Appendix Figure 2 contains an example BWS question).

We fit the following weighted conditional logit model to estimate the relative value that respondents place on these digital goods:

$$y_{ijk} = B_1 GoogleSearch + B_2 Facebook + \dots + B_{20} \$500 \quad (1)$$

where y is a binary variable indicating whether or not respondent i prefers option j to option k , where the two options are from among the 21 items (10 digital goods, not meeting friends, and 10 monetary amounts). Each row represents a pairwise comparison between two options, with one option designated as best (taking on the value 1 under the maxdiff model), another designated as worst (taking on the value -1 under the maxdiff model), the remaining items taking the value 0, and y indicating whether the best/worst combination in the row is true for the respondent. Since a best-worst scaling question contains three options, each question a respondent receives is represented by six rows in our dataset; each of the three pairwise comparisons is repeated twice with alternate options designated as best or worst. The dichotomous response variable y takes the value ‘TRUE’ if the row corresponds to the respondent’s actual choice for that question and ‘FALSE’ otherwise. We calculate separate

⁸ The amounts were \$5, \$10, \$20, \$30, \$40, \$50, \$65, \$80, \$100 and \$500 for the United States. For the other countries we converted the amounts to local currencies and rounded the amounts up.

weights for best-worst scaling respondents to ensure that the best-worst scaling respondents are representative of the Facebook population in each country.

We employ an attention check embedded in the research design to exclude respondents who did not rationally evaluate the relative utility of monetary amounts. This attention check appears in questions that have two monetary amounts and one digital good as the two options. Out of 23,752 respondents⁹, 16061 (68%) answered such a question. Out of these, 10,752 (67%) passed the attention check and 5309 (33%) failed. We exclude these 5309 respondents from our analysis, which leaves 18,443 respondents.

Calibrating Welfare Gains from digital goods using Valuations for Facebook Obtained via Single Binary Discrete Choice Experiments

An important aspect of our experimental setup is that we are able to monitor deactivation behavior for Facebook. We use this to run incentivized, willingness-to-accept questions about how much users would have to be paid in order to deactivate Facebook for one month.¹⁰ As aforementioned, these estimates can then be combined with the BWS results to calibrate the valuations for the other digital goods.

Specifically, in the same survey in which the BWS questions are asked, we used an incentivized single binary discrete choice questionnaire to elicit respondents' willingness to accept to stop using Facebook for one month (see Appendix Figures 3 and 4 for survey screenshots). Respondents were asked: "Would you be willing to stop using Facebook for one month in exchange for X?", where X was chosen randomly from a set of 9 monetary values from \$5 to \$100. We clarified to the respondents that they could be randomly selected for their choices to be fulfilled, and if so, they would actually be eligible to receive the offer amount if they deactivated their Facebook account for a month.

Respondents were given offers in their own currency. For instance, if a respondent in France was chosen to receive an offer equivalent to \$50, they were given an offer of 45 Euros, which was equivalent to US\$50 at the time. These offers were incentivized, and participants had to agree to a set of terms and conditions drafted by legal experts to be in compliance with local laws in each country. Each choice made by respondents of whether to accept or reject an offer had a 2% probability of being selected. This probability was not known to respondents.

⁹ The sample size for the BWS questions is slightly lower than the sample size for the Facebook SDBC question due to survey attrition. The BWS questions appeared later in the survey. Respondents could stop answering the survey at any time.

¹⁰ We chose Facebook for this study because we can ensure compliance of choices for selected users (we are able to observe their true Facebook usage).

To arrive at the median willingness-to-accept value for Facebook, we first estimate a weighted logit model to ascertain how the probability of rejecting the offer depends on the amount of the offer:

$$P(\text{Reject Offer}_i) = \Lambda(\beta_0 + \beta_1 \text{Offer}_i) \quad (2)$$

where *Offer* is equal to the randomly assigned offer amount from \$5 to \$100, and *Reject Offer* is an indicator equal to 1 when a respondent answers that they are not willing to stop using Facebook in exchange for that offer amount. To find the median value, we solve for the offer amount that sets this probability equal to 0.5, which yields:

$$\text{Offer}^* = -\beta_0/\beta_1 \quad (3)$$

We use bootstrapped standard errors to estimate the confidence intervals.

4. Results

We start with some descriptives about sample sizes across countries (Table 1). Next, we provide results for relative valuations using the best-worst scaling methodology. We then analyze absolute valuations of Facebook using incentive compatible single binary discrete choice experiments. Using these absolute valuations of Facebook, we estimate valuations (in \$) for the other digital goods in our sample and calculate welfare gains across and within countries across different users. Finally, we benchmark our valuations with other ways of measuring valuations of digital goods (time spent and ad revenues).

Country	Sample size for Facebook valuation	Sample size for best-worst scaling analysis
Ireland	3698	1929
United Kingdom	3674	1899
Canada	3460	2067
Norway	3263	2060
South Korea	3252	1802
Mexico	3061	2018
Germany	3048	1648
Spain	2937	1762
Belgium	2776	1787

France	2700	1534
Romania	2658	1508
United States	2635	1641
Japan	2555	2097

Table 1: Sample Sizes

Table 1 Notes: The table displays sample sizes for our survey sample for both valuation methodologies outlined above.

Table 1 shows that we have between 2555 and 3698 respondents for the Facebook valuation and between 1508 and 2097 respondents for the best-worst scaling analysis. The sample size is a bit lower for the latter since these questions appeared towards the end of the survey and experienced higher attrition once respondents finished the incentivized single binary discrete choice experiments to measure valuation of Facebook. Our weighting methodology accounts for this attrition to ensure that samples for both analyses in each country are representative of the respective Facebook populations. Our sample has a higher proportion of male respondents which is reflective of the user base of digital platforms broadly.

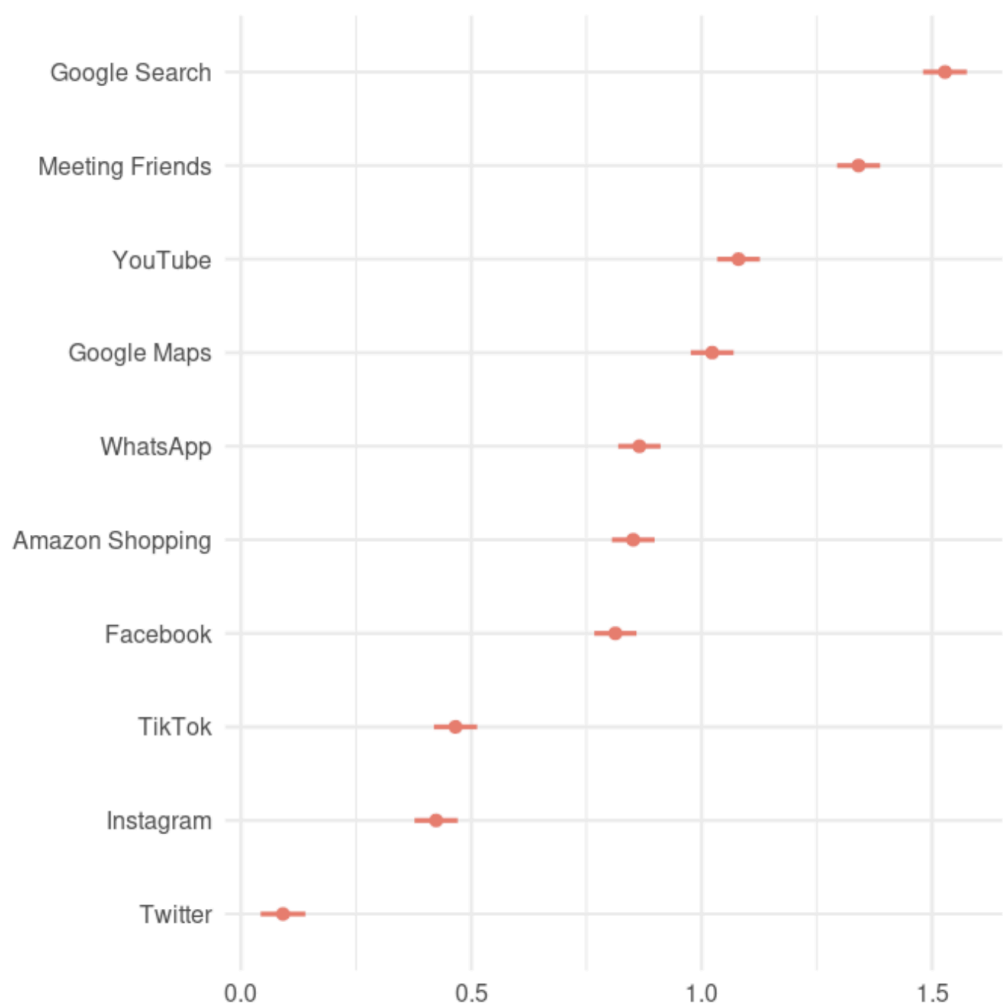
Once the survey was completed, a total of 381 respondents were selected to deactivate their account. Out of these, 170 (45%) attempted to deactivate, based on log data. 113 (30%) successfully deactivated and were paid. The amounts actually paid out ranged from US\$5 to US\$510. A total of approximately US\$14,400 were paid out to 113 participants. There may be multiple reasons why compliance with deactivation was not higher. Deactivation was a difficult five-step process, involving at least 8 clicks. Furthermore, we were only able to notify respondents that they had been selected to deactivate via email. Respondents may not have checked their emails in time, they may have missed them, or they may have been filtered as spam.

Measuring relative valuations of digital goods using best-worst scaling

Figure 1 shows the ranking of items from most preferred to least preferred based on the disutility of giving up access to that item, as estimated using equation (1). The omitted category is Snapchat.¹¹ The figure depicts the disutility from stopping use of the digital service (i.e. the relevant coefficient from eq. 1) relative to stopping use of Snapchat, which is the least preferred

¹¹This approach also allows us to analyze heterogeneity in valuations based on various user characteristics.

digital service.¹² For instance, stopping use of Google Search for a month causes the largest disutility—relative to stopping use of Snapchat—among users of the 13 countries.¹³ That means that Google Search is the most preferred good and is even preferred to meeting friends in person. YouTube is also highly ranked, along with Google Maps, WhatsApp, Amazon Shopping, and Facebook. Other apps such as TikTok, Instagram, Twitter, and Snapchat have more limited broad-based appeal among the suite of digital goods.¹⁴



¹² The plotted coefficients can be interpreted as the effect of one unit change in the predictor on the log-odds of the outcome. For example, the Google Search coefficient of 1.52 indicates that if Google Search was framed as the best possible choice, the log-odds of the respondent preferring that option would be 1.52 higher than if Snapchat was framed as the best possible choice. Probabilities of preferring that option are a monotonic function of the log-odds, with $p = \exp(\ln(\text{odds})) / 1 + \exp(\ln(\text{odds}))$. The estimated coefficients can therefore be interpreted as the relative disutility of giving up the good in question, relative to other goods (and with Snapchat as the omitted category benchmarked at 0).

¹³ Note that these figures reflect the valuation of the average Facebook user and not necessarily the valuation of the average user of these other platforms or the valuation of individuals who don't use Facebook.

¹⁴ Some of the digital goods in the list, such as Snapchat, are not broadly used in many countries, which contributes to them being least preferred.

Figure 1: Relative disutility from stopping use, estimated using a conditional logit model. Snapchat is the omitted category.

Figure 1 Notes: The figure depicts the disutility from stopping use of the digital service relative to stopping use of Snapchat, which is the least preferred digital service. For instance, stopping use of Google Search for a month causes the largest disutility—relative to stopping use of Snapchat—among users of the 13 countries.

Calibrating welfare gains using Facebook

Figure 2a below shows, per each offer value, the proportion of respondents who rejected the offer to stop using Facebook for one month. Offer values were randomly assigned, with 81% of users rejecting the lowest offer (equivalent to US\$5) and 24% of respondents rejecting the highest offer (equivalent to US\$100). Reassuringly, the proportion of respondents who reject the offer decreases monotonically as the offer value increases.

The median monthly Facebook valuation in a population is the amount of money that half of respondents would accept and half would reject. As aforementioned, we first estimate a weighted binary logit model to measure how the probability of rejecting the offer depends on the offer amount and then solve for the offer amount that sets the probability of rejecting to 50%. To ascertain how the value that users derive from Facebook is distributed across countries, we first examine how the median value varies across surveyed countries (Figure 2b). This figure shows estimates from separate binary logit models for the median willingness-to-accept (WTA) value for each country. The estimates and confidence intervals in blue are unweighted, while those in orange are weighted in line with the weighting strategy described above. Countries are arranged in order of decreasing weighted median WTA value of Facebook. The overall weighted median value is \$31 per month, and ranges from \$11 per month in Romania to \$57 per month in Norway.¹⁵

¹⁵ Note that we only display the weighted median WTA for the pooled sample on the leftmost side of the plot. The reason is that the unweighted estimate does not have a clear interpretation: the sample size was similar across countries, regardless of the number of monthly active Facebook users in the country.

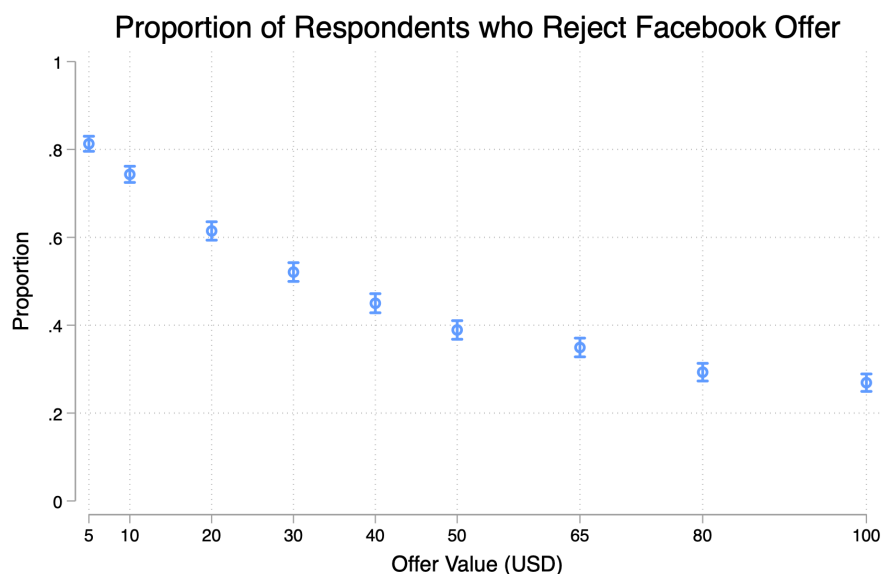


Figure 2a: Facebook offer rejection rates by offer value across all countries in the sample.

Figure 2a Notes: This figure shows, per each offer value, the proportion of respondents who reject the offer to stop using Facebook for one month. Offer values were randomly assigned. The proportion of respondents who reject the offer decreases as the offer value increases.




Figure 2b: Facebook median willingness to accept (WTA) by Country

Figure 2b Notes: This figure shows estimates from separate binary logit models for the median WTA monthly value for each country. The estimates and confidence intervals in blue are unweighted, while those in orange are weighted in line with the weighting strategy described in the Appendix. Countries are arranged in order of decreasing weighted median WTA value of Facebook.

Welfare gains from digital goods across countries

We use data on Facebook valuations to first calculate total welfare generated by Facebook across the 13 countries in our sample. We do this by multiplying the weighted median willingness-to-accept value (shown in Figure 2b) by the number of monthly active Facebook users in that country, and then multiply this number by 12 to annualize it (Appendix Figure 6).¹⁶ Using this approach, Facebook generates a total of \$246 billion in welfare across these countries (ranging from \$137 billion in the US to \$1.2 billion in Ireland). Appendix Figure 8 shows the valuation of Facebook as a percent of GDP per capita in each country.

In turn, using the incentivized single binary discrete choice (SBDC) Facebook valuations as our benchmark and applying the relative utilities across goods estimated with the best-worst scaling (BWS) method, we calibrate the valuations of other goods and estimate total welfare generated by all these 10 goods in our sample across all countries (Appendix Figure 11 shows the calibrated valuations of each digital good in each country¹⁷). This is done by multiplying the aggregate annual value of Facebook in each country by the utility of each digital goods relative to Facebook—where the utility for Facebook is normalized to 1—obtained through the best-worst scaling (BWS) estimation. Our analysis implies that, among Facebook users, the 10 digital goods selected in this study generate a combined total of \$2.52 trillion in welfare across these countries—ranging from \$1.29 trillion in the US to \$13 billion in Romania (Figure 3a). Appendix Figure 9 shows the valuation of the 10 digital goods as a percent of GDP per capita in each country.

¹⁶ We assume that the willingness-to-accept (WTA) value to stop using Facebook for a year is approximately 12 times as large as the analogous willingness-to-accept for a month. Arguments can be made to support the hypothesis that willingness-to-accept increases more than proportionally or less than proportionally with respect to disconnection time. Previous research (Brynjolfsson et al., 2019a) shows that WTA to stop using Facebook for 1 month (4 weeks) is slightly more than four times WTA to stop using Facebook for 1 week, and WTA to stop using Facebook increases more than proportionally over time (from 1 month to 1 year). Therefore, our assumption might slightly underestimate the welfare generated by Facebook over a period of 12 months.

¹⁷ There is substantial heterogeneity in valuations of digital goods across countries. For example, Google Search is the most valued good in the US, WhatsApp is the most valued good in Mexico and YouTube is the most valued good in South Korea.




Figure 3a: Aggregate annual value of 10 digital goods by country

Figure 3a Notes: This figure presents country-level estimates of the aggregate annual value of all ten digital goods studied using relative valuations from the best-worst scaling study calibrated with incentivized Facebook valuations using the single-binary discrete choice study. Countries are arranged in order of decreasing weighted median WTA value of Facebook. The y-axis includes a discontinuity in order to visually accommodate the US estimate which is far higher than the estimate for other countries.

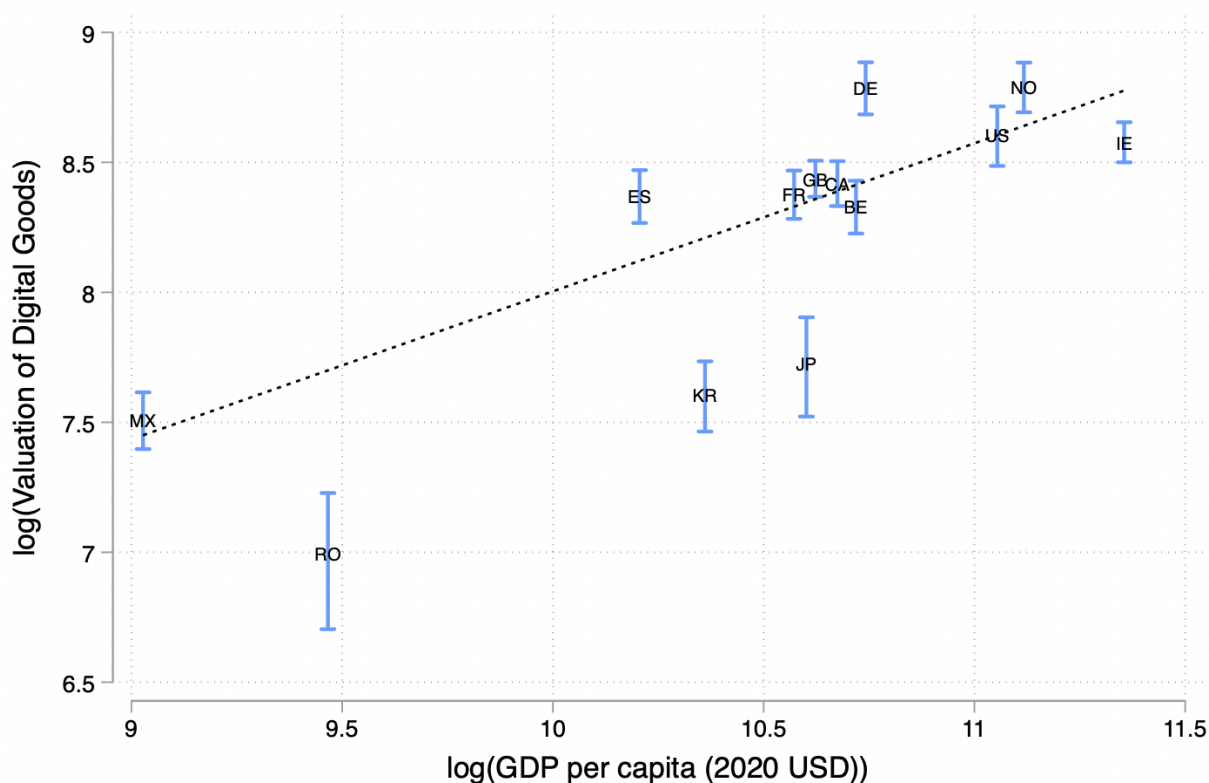
To explore whether the welfare gains from digital goods vary across higher-income and lower-income countries we calculate, for each country, Facebook users' logged median valuations of all the digital goods. We then regress this variable on a country's logged GDP per capita:

$$\log(\text{Valuation of Digital Goods})_c = \alpha_c + \beta \log(\text{GDPpc}_c) + \varepsilon_c \quad (4)$$

where c denotes country and β denotes the parameter to be estimated. We run a weighted least squares regression in which the weights are given by the monthly active Facebook users in each country. The rationale for including the weights is to equally consider all Facebook users in the regression.

We find a highly significant positive relationship between users' logged digital goods valuation in a country and logged GDP per capita in that country (Figure 3b, see Appendix Figure 8 for an analogous plot with the valuation of Facebook alone in the Y axis). The figure depicts logged 2020 GDP per capita (in US dollars) in the X axis and Facebook users' logged valuation of the

10 digital goods in the Y axis.¹⁸ A weighted regression of Y on X with weights given by logged Facebook monthly active users (MAU) in each country yields a point estimate of 0.68 with a p-value equal to 0.002. An unweighted regression yields a point estimate of 0.69 with a p-value equal to 0.002.¹⁹ The fitted lines in the plot correspond to the weighted regression. A 1% increase in a country's GDP per capita is associated with a 0.68% increase in users' valuation of the 10 digital goods relative to GDP per capita.²⁰ In other words, among Facebook users, increases in income are associated with less than 1-to-1 increases in digital goods valuation, which implies that the welfare gains from digital goods represent a higher share of income in lower-income countries compared to higher-income countries.



¹⁸ The dependent variable is calculated for active Facebook users. This means that the variable is invariant to penetration levels of Facebook in a country. If we use an alternative variable that is not invariant to FB penetration, such as the valuation of digital goods as a percent of total GDP (calculated by multiplying the valuation of digital goods of the median FB user by the number of monthly active FB users in the country and dividing by total GDP), we find a negative relationship between users' digital goods valuation and GDP per capita, which is consistent with the findings in Figure 3b.

¹⁹ Appendix Figure 7 is analogous to Figure 3b but excluding Mexico. The coefficient associated with GDP per capita in the weighted regression is 0.90 and the p-value is equal to 0.001.

²⁰ In the Appendix we discuss a simulation procedure that accounts for the statistical uncertainty around our estimates of the value of the 10 digital goods as a percent of GDP per capita. The simulation exercise shows that the results of the regression are robust to variations in the country estimates of the value of digital goods as a percent of GDP per capita.

Figure 3b: Association between GDP per capita and Facebook users' valuation of 10 digital goods

Figure 3b Notes: The figure depicts logged 2020 GDP per capita (in US dollars) in the X axis and Facebook users' logged valuation of the 10 digital goods in the Y axis. A weighted regression of Y on X with weights given by Facebook monthly active users (MAU) in each country yields a point estimate of 0.68 with a p-value equal to 0.002. An unweighted regression yields a point estimate of 0.69 with a p-value equal to 0.002. The fitted lines in the plot correspond to the weighted regression.

Welfare gains from digital goods within countries

To analyze heterogeneity in welfare gains across income levels *within* each country, we calculate valuations of Facebook and other digital goods by income levels and education. For income, we use the relative wealth index of the location of a user's residence as a proxy (Chi et al. (2022)).²¹ For education, we rely on the responses that users provided in our survey.

We find that the monetary value that users derive from Facebook does not tend to vary by these indicators of material welfare. Figure 4a shows weighted estimates from separate binary logit models for the median WTA value for each relative wealth index tercile within each country. For instance, users who live in the least wealthy localities classified by the Relative Wealth Index have a value of Facebook that is similar to that of users in the highest wealth localities (Figure 4a).²² Similarly, users with less than secondary education or with just secondary education (and who therefore tend to earn less on average) have a median value of Facebook that is not statistically distinguishable from the value that users with a college degree derive from Facebook (see Appendix Figure 5).²³

These findings imply that the value of Facebook represents a higher *share* of their income and wealth for users who currently have lower income and wealth. Extending to other digital goods, we estimate three separate conditional logit models (on users who are in the bottom, middle, and top tercile of relative wealth in their respective countries) of the utility of each digital goods

²¹ The relative wealth index calculated in Chi et al. (2021) estimates the relative wealth and poverty of an area at 2.4 km resolution. From the paper: "The estimates are built by applying machine-learning algorithms to vast and heterogeneous data from satellites, mobile phone networks, and topographic maps, as well as aggregated and deidentified connectivity data from Facebook. We train and calibrate the estimates using nationally representative household survey data from 56 LMICs and then validate their accuracy using four independent sources of household survey data from 18 countries."

²² We do not have perfect information about users' locations. The more granular the geographic area, the less accurate the location predictions are. However, the accuracy of Facebook's zip-code level predictions in the US is as high as 68%—although note that the 2.4-km microregions defined in Chi et al. (2022) do not neatly correspond with zip-codes. To the extent that relative wealth index levels are geographically clustered, small imprecisions in the location predictions should not substantially impact the accuracy of the RWI tercile categorizations. That said, any noise in the relative weight indices may bias the relationship between Facebook valuations and the RWI index toward zero.

²³ We also explore heterogeneity in Facebook valuations based on home ownership and gender (Appendix Figure 5), which are both correlated with wealth. Facebook valuations do not significantly vary based on home ownership (owned vs. rented home). Women value Facebook significantly higher than men (and women's wealth is lower than men in all of our sample countries).

relative to Facebook (which is the omitted good, with utility set at 0) in Figure 4b (Appendix Figure 10 plots this for the US alone).²⁴

Interestingly, we find that the poorest and wealthiest users often have greater value for most digital goods than those in the middle of the relative wealth distribution, though the pattern is inconsistent across digital goods. For some digital goods (e.g. Google Search), users in the top tercile derive the highest absolute welfare in dollar terms while for other digital goods (e.g. TikTok) those in the lowest tercile benefit the most. The differences in valuations are rarely significantly different from each other across terciles which suggests that, on balance, these digital goods tend to lower welfare inequality within countries.

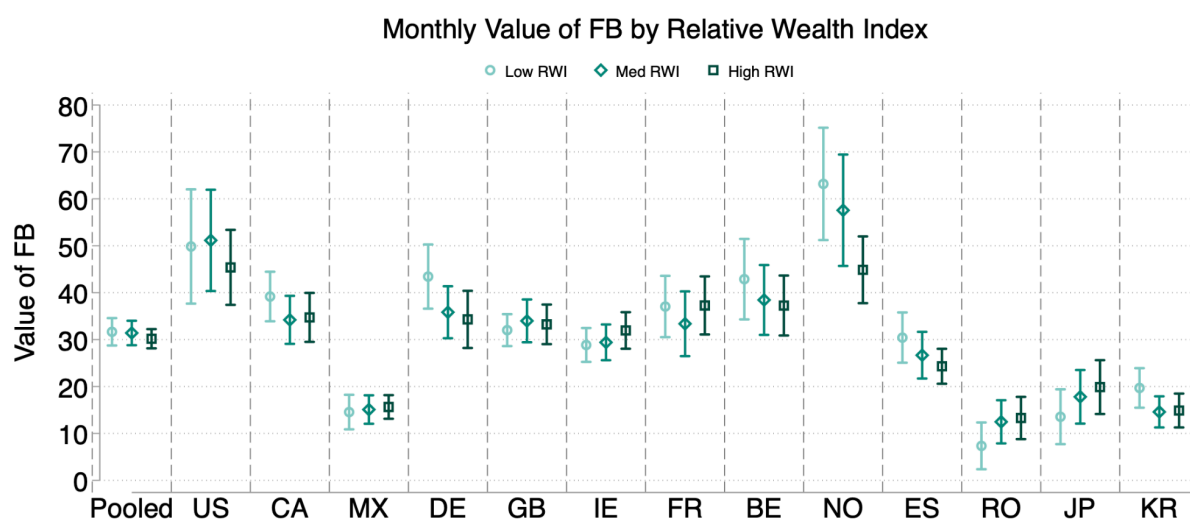


Figure 4a: Value of Facebook by relative wealth index

Figure 4a Notes: This figure shows weighted estimates from separate binary logit models for the median WTA value for each relative wealth index tercile within each country.

²⁴ These results should be interpreted in conjunction with those in Figure 4a.




Figure 4b: Relative valuation of 10 digital goods by relative wealth index

Figure 4b Notes: This figure shows estimates from three separate conditional logit models (each containing users belonging to the low, medium and high RWI tertiles within their respective countries) of the utility of each digital good relative to Facebook (which is the omitted good, with utility set at 0). These results should be interpreted in conjunction with those in Figure 4a.

Comparing consumer welfare with time spent and firm revenue

Previous research has estimated consumer welfare generated by free digital goods using measures of time spent (Goolsbee and Klenow, 2006; Brynjolfsson et al., 2023) and advertising revenues (Nakamura et al., 2017). Do our measures of consumer welfare capture additional information beyond measures of time spent and advertising revenues? To explore this, we

compare our estimates of valuation of Facebook with time spent on Facebook. Figure 5 plots these comparisons for Facebook valuations across all the 13 countries pooled together. We split our study respondents into three terciles—low, medium, and high—based on time spent on Facebook. For each of these terciles, we calculate the median valuation of Facebook.

We find that users in the first tercile—i.e. users who spend the least time on the platform—have a median valuation of \$19.72 per month. Users in the second tercile have a median valuation of \$32.97 per month. Finally, users in the third tercile have a median valuation of \$40.61 per month. Moving from the first tercile to the third tercile, the valuation of Facebook increases by 2.06 times while time spent increases by 12.84 times. Thus, valuation increases much more slowly than time spent, implying that value is distributed across a broad user base rather than concentrated on a few very active users.

We can also compare consumer welfare gains to revenues for the producers. Nordhaus (2005) estimated that only a small portion of the total welfare generated by technological advances in the 1948-2001 time period was ultimately captured by producers. Instead, consumers enjoyed the vast majority of the welfare gains. Brynjolfsson (1996) also finds that IT investments generate three times their cost in welfare gains to consumers. Our study finds results that are consistent with these analyses: when we compare our welfare estimates to advertising revenues, we find that user value for just the Facebook app in the 13 countries studied (\$246 billion) is more than double the global advertising revenue of Meta Platforms' (\$115 billion, including Facebook, Instagram, and WhatsApp). Tadelis et al. (2023) find that each dollar spent on Meta ads leads to over three dollars in revenues for advertisers. Our findings imply that the vast majority of the welfare gains from using Facebook go to consumers and not to Facebook.

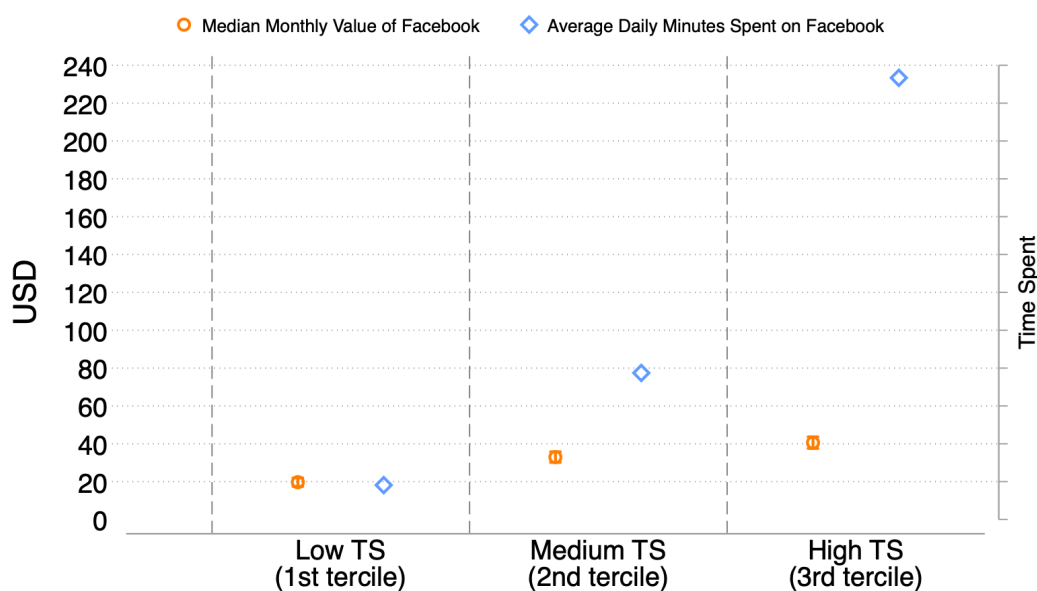


Figure 5: Comparing monthly valuation of Facebook with daily time spent

5. Discussion

Digital goods generate large benefits for consumers, but because most of these goods are free to use, these benefits are largely invisible in standard government statistics such as GDP and productivity. In this paper, we provide estimates of the value digital goods create for users in 13 countries around the world by conducting large-scale incentivized online choice experiments on representative samples of nearly 40,000 people.

We find that the 10 selected digital goods across the 13 countries generate more than \$2.5 trillion in aggregate consumer welfare per year, roughly equivalent to 6% of GDP in these countries. We also find that lower-income individuals and countries disproportionately benefit from these digital goods. These findings suggest that digital goods reduce inequality in welfare within and across countries.

Our approach is subject to a number of limitations. Compared to GDP, which can be measured with high precision, our estimates are relatively noisy. We are confident in our qualitative findings that these digital goods create trillions of dollars of value and reduce welfare inequality, but the exact values are not precisely estimated. The large sample size in our study partly mitigates this problem via the law of large numbers, but there may remain systematic biases in our estimates for a variety of reasons.²⁵ Relatedly, we study a particular sample of countries and of digital goods. While a pattern is evident within this large sample, different effects may occur for other sets of countries and goods. While our sample accounts for a substantial fraction of the likely value of global GDP and of value from digital goods, the out-of-sample implications can best be addressed by simply expanding the sample.

Furthermore, even when the valuations we obtain are accurate, they may reflect irrational choices that are not in the consumers' genuine self-interest for some of the digital goods (Allcott et al., 2022) or other errors in judgment (Kahneman et al., 1982). In addition, these goods may create positive or negative externalities on other people—from shared memories and connections to misinformation and polarization—which means that the total welfare gains are not necessarily equal to the sum of individual valuations. While these concerns are important, they may not apply to all of the digital goods in our sample (especially some of the most valued digital goods such as search engines and maps). It should also be noted that the same concerns apply to standard measures of GDP, which also reflect consumer values which may be irrational or omit important externalities. Future work should seek to address these concerns, and online choice experiments can also be used to quantify these externalities (e.g., Bishop et al., 2017; Collis and Eggers, 2022).

For the case of Facebook and for the subset of participants in our study who were chosen for deactivation, we are able to observe activity on Facebook once the study was completed. When we compare the percentage of users who are monthly active users of Facebook 2 years after the study among the users who deactivated their accounts versus all survey respondents, we

²⁵ See Brynjolfsson et al. (2019a) for a more detailed discussion of potential biases and shortcomings of massive online choice experiments.

find a small, 0.2 percentage point difference, with the rate being higher for the user group who deactivated their accounts. These figures suggest that users who were incentivized to deactivate Facebook were just as likely to be active on Facebook once the study was completed compared to other users who were not asked to deactivate. This suggests that our consumer welfare estimates captured through our willingness-to-accept valuations reflect consumers' rational choices for the most part. If valuations severely suffered from irrational choices or errors of judgment, we would have expected significantly more users who were offered monetary rewards to deactivate Facebook to keep their Facebook accounts deactivated once the study was completed.

Having demonstrated the feasibility of running massive online choice experiments to estimate valuations of multiple goods across multiple countries, future work can expand this line of research in at least three dimensions: i) More goods, including non-digital goods like breakfast cereal, improved healthcare, or cleaner water, ii) More countries or regions, and iii) More respondents per item (which will increase the precision of our estimates). Furthermore, by conducting online choice experiments such as this one at a regular cadence, e.g. annually or even more frequently, and with consistent methods, we can better understand not only levels but also changes in welfare as the basket of goods and other variables change over time. Analyzing these changes may also help to overcome some of the limitations mentioned earlier.

This paper provides a first step toward systematically estimating welfare using massive online choice experiments. Since the contribution of digital goods to welfare is likely to continue to grow in the twenty-first century, establishing a reliable baseline will provide a foundation for understanding the magnitude and nature of future changes in the economy.

References

- Allcott, H., Braghieri, L., Eichmeyer, S., Gentzkow, M. (2020). The Welfare Effects of Social Media. *American Economic Review*, 110(3), 629-676.
- Allcott, H., Gentzkow, M., & Song, L. (2022). Digital addiction. *American Economic Review*, 112(7), 2424-63.
- Athey, S., Grabarz, K., Luca, M., & Wernerfelt, N. (2023). Digital public health interventions at scale: The impact of social media advertising on beliefs and outcomes related to COVID vaccines. *Proceedings of the National Academy of Sciences*, 120(5), e2208110120.
- Becker, G. M., M. H. DeGroot and J. Marschak (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, 9(3), 226-232.
- Bishop, R. C., Boyle, K. J., Carson, R. T., Chapman, D., Hanemann, W. M., Kanninen, B., ... & Scherer, N. (2017). Putting a value on injuries to natural assets: The BP oil spill. *Science*, 356(6335), 253-254.

Brynjolfsson, E. (1996). The contribution of information technology to consumer welfare. *Information Systems Research*, 7(3), 281-300.

Brynjolfsson, E., & Collis, A. (2019). How should we measure the digital economy? *Harvard Business Review*, 97(6), 140-148.

Brynjolfsson, E., Collis, A., & Eggers, F. (2019a). Using massive online choice experiments to measure changes in well-being. *Proceedings of the National Academy of Sciences*, 116(15), 7250-7255.

Brynjolfsson, E., Collis, A., Diewert, W. E., Eggers, F., & Fox, K. J. (2019b). *GDP-B: Accounting for the value of new and free goods in the digital economy* (No. w25695). National Bureau of Economic Research.

Brynjolfsson, E., Eggers, F., & Gannamaneni, A. (2018, May). Measuring welfare with massive online choice experiments: A brief introduction. *AEA Papers and Proceedings* (Vol. 108, pp. 473-476).

Brynjolfsson, E., Kim, S. T., & Oh, J. H. (2023). The Attention Economy: Measuring the Value of Free Goods on the Internet. *Information Systems Research*.

Brynjolfsson, E., & Saunders, A. (2009). What the GDP gets wrong (why managers should care). *MIT Sloan Management Review*, 51(1), 95.

Chi, G., Fang, H., Chatterjee, S., & Blumenstock, J. E. (2022). Microestimates of wealth for all low-and middle-income countries. *Proceedings of the National Academy of Sciences*, 119(3), e2113658119.

Collis, A., & Eggers, F. (2022). Effects of restricting social media usage on wellbeing and performance: A randomized control trial among students. *Plos one*, 17(8), e0272416.

Coyle, D., & Nakamura, L. (2022). Time Use, Productivity, and Household-centric Measurement of Welfare in the Digital Economy. *International Productivity Monitor*, 42, 165-186.

den Haan, W., Ellison, M., Ilzetki, E., McMahon, M., & Reis, R. (2017). Happiness and wellbeing as objectives of macroeconomic policy: Views of economists. *VOX EU*.

Fischer, C. (2017). Inequality Is About Access to Public Goods, Not Income. *Boston Review*. April.

Goolsbee, A., & Klenow, P. J. (2006). Valuing consumer products by the time spent using them: An application to the Internet. *American Economic Review*, 96(2), 108-113.

Hanani, H. (1961). The existence and construction of balanced incomplete block designs. *The Annals of Mathematical Statistics*, 32(2), 361-386.

Jones, C. I., & Klenow, P. J. (2016). Beyond GDP? Welfare across countries and time. *American Economic Review*, 106(9), 2426-2457.

Krueger, A. B., & Stone, A. A. (2014). Progress in measuring subjective well-being. *Science*, 346(6205), 42-43.

Louviere, J. J., Flynn, T. N., & Marley, A. A. J. (2015). Best-worst scaling: Theory, methods and applications. *Cambridge University Press*.

Lowensohn, J. (2012) Apple Features worth up to \$100 to Samsung Buyers, Witness Says, *CNET*, August 10. <https://www.cnet.com/tech/tech-industry/apple-features-worth-up-to-100-to-samsung-buyers-witness-says/>

Masood, E. (2022). GDP is getting a makeover—what it means for economies, health and the planet. *Nature (News Feature)*, 611, 224-226.

Nakamura, L. I., Samuels, J., & Soloveichik, R. H. (2017). Measuring the Free Digital Economy within the GDP and productivity accounts. *FRB of Philadelphia Working Paper*.

Nordhaus, W. D. (1996). Do real-output and real-wage measures capture reality? The history of lighting suggests not. In *The economics of new goods* (pp. 27-70). University of Chicago Press.

Nordhaus, W. D. (2005). Schumpeterian profits and the alchemist fallacy. *Yale Economic Applications and Policy Discussion Paper*, (6).

Sarig, T., Galili, T., and R. Eilat (2023). balance - - a Python package for balancing biased data samples, (<https://arxiv.org/abs/2307.06024>).

Slottje, D. J. (1987). Relative price changes and inequality in the size distribution of various components of income: A multidimensional approach. *Journal of Business & Economic Statistics*, 5(1), 19-26.

Tadelis, S., Hooton, C., Manjeer, U., Deisenroth, D., Wernerfelt, N., Greenbaum, L., and N. Dadson (2023) Learning, Sophistication, and the Returns to Advertising: Implications for Differences in Firm Performance, NBER Working Paper.

Waldfoegel, J. (2018). Digital Renaissance: What Data and Economics Tell Us about the Future of Popular Culture. Princeton University Press.

Wertenbroch, K., & Skiera, B. (2002). Measuring consumers' willingness to pay at the point of purchase. *Journal of marketing research*, 39(2), 228-241.

Online Appendix

Weighting Details

Our weighting strategy consists of three key building blocks: (1) Design weights to account for differential probability of inclusion into the sample by country, (2) Unit non-response weights to account for the probability of a user responding to the survey, (3) Item non-response weights to account for the probability of a user who started the survey responding to the survey item in question (i.e. either the Facebook valuation or the best-worst scaling). For estimates at the country-level, we rely on building blocks (2) and (3) only, whereas for estimates where we pool data across countries, we rely on all three pieces to construct the final weights.

1. Design weights

The probability of inclusion in our sample is different across countries. Our weighting strategy adjust for the different probability of selection into the sample by country due to the uneven sample allocation across countries:

$$\text{Weight for responses from country } i = 1 / (\text{Number of users included in the sampling frame from country } i / \text{monthly active user population from country } i)$$

The weights obtained using this approach can be interpreted as the number of monthly active users (i.e., population) represented by a user included in the sampling frame (i.e., sample).

2. Unit non-response weights within country

We adjust the weights to account for unit non-response by modeling the probability that a user in the country's sampling frame responds to the survey as a function of observable user characteristics. For each country, the target population is the sampling frame. To account for the possibility that the response model differs across countries, we model this probability separately by country.

$$\text{Weight} = 1 / \text{Est. Pr}(\text{user starts survey})$$

We use logistic regression to estimate the probability of a user starting the survey. We model this probability as being conditional on : gender, primary phone operating system, whether the user has a profile picture, age (quartile bins), friend count (quartile bins), the number of days within the last 28 days that the user was active, an indicator for whether the user was active for all days within the last 28 days, and time since the user created their account (quartile bins). Note that we use an internal version of the publicly available 'balance' package (Sarig et al. 2023) to implement the inverse probability weighting. This package uses a regularized logistic model using LASSO.

3. Item non-response weights for the valuation questions

We account for item non-response by modeling the probability that a user in the country who started the survey responds to the relevant survey question (i.e., answered at least one of the best-worst scaling questions, or answered the Facebook incentivized valuation question). For each country, the target population is the set of users from that cluster who started the survey. To account for the possibility that the response model differs across countries, we model this probability separately by country.




$$\text{Weight} = 1 / \text{Est. Pr}(\text{user answers question} \mid \text{user starts survey})$$

Similar to the unit non-response weights, we use logistic regression to estimate the probability of responding to the item. We model this probability as being conditional on the same user characteristics as unit non-response weights above, in addition to an indicator for whether they indicated they would be willing to stop using Facebook for one month if they were offered money in return.

Simulation Procedure to Account for the Statistical Uncertainty Around the Valuation Estimates

In Figure 3b in the main text, we regress Facebook users' logged median valuation of the 10 digital goods (in the Y axis) on logged 2020 GDP per capita (on the X axis). We weight country observations according to the number of monthly active users on Facebook. However, our estimates of the users' valuation of digital goods are themselves uncertain. To account for this uncertainty, we ran a simulation analysis. The simulation exercise shows that the results of the regression are robust to variations in the country estimates of the value of digital goods.

In the simulations, we draw different values of the Y variable (Facebook users' logged valuation of the 10 digital goods) for each country from a Normal distribution with mean and variance matching each country's point estimate and confidence intervals around Y. For each realization, we run a weighted regression that is analogous to the regression in the main text. We then compute the percent of such regressions that yielded a negative and significant β coefficient. For both the weighted and unweighted regressions, we find that 100 percent of regressions yield a negative and significant β . These results are not surprising insofar as the p-value associated with the weighted and unweighted regressions in the main text are 0.002 and 0.001, respectively.



Asad, you have an opportunity to receive a payment* for sharing your opinion with us! Could you take a few minutes to answer a short survey?

*Ends 3/21/22 at 11:59:59pm PST. Open to individuals who: (1) are legal residents of US, UK, FR, DE, ES, BE, RO, NO, IE, CA (excl. Quebec), KR, JP, ID, MX, TH; (2) 18+ and age of majority; (3) receive authorized invitation; & (4) are a registered Facebook user with valid email address & Internet access. Subject to full Terms and Conditions (research.fb.com/dss-epr-survey-terms). Void where prohibited.

[Start Survey](#)

Appendix Figure 1: Survey invitation

Which of these three situations are you MOST WILLING to experience and which are you LEAST WILLING to experience?

MOST WILLING		LEAST WILLING
<input type="radio"/>	You don't use YouTube for 1 month	<input type="radio"/>
<input type="radio"/>	You don't use WhatsApp for 1 month	<input type="radio"/>
<input type="radio"/>	You don't use Snapchat for 1 month	<input type="radio"/>

Submit

Appendix Figure 2: Example BWS task



If you agree to participate, we may offer to pay you to stop using Facebook and to temporarily deactivate this Facebook account for one month.

If you temporarily deactivate, you could continue using Messenger, and nothing on your Facebook account would be deleted.*

In order to participate, please confirm that you agree to the terms and conditions.

- I agree to the terms and conditions
- I do not agree to the terms and conditions
-

***You could reactivate your account at any time, but we would check that your account stays deactivated for the entire month before paying you. View full terms at research.fb.com/dss-epr-survey-terms in a separate browser window.**

Ends 4/7/22 at 11:59:59pm PST. Open to individuals who: (1) are legal residents of US, UK, FR, DE, ES, BE, RO, NO, IE, CA (excl. Quebec), KR, JP, ID, MX, TH; (2) 18+ and age of majority; (3) receive authorized invitation; & (4) are a registered Facebook user with valid email address & Internet access. Subject to full Terms and Conditions. Void where prohibited.

Continue

Appendix Figure 3: Terms and conditions screen

Here is an example of the offer made to respondents after they had agreed to the terms and conditions:



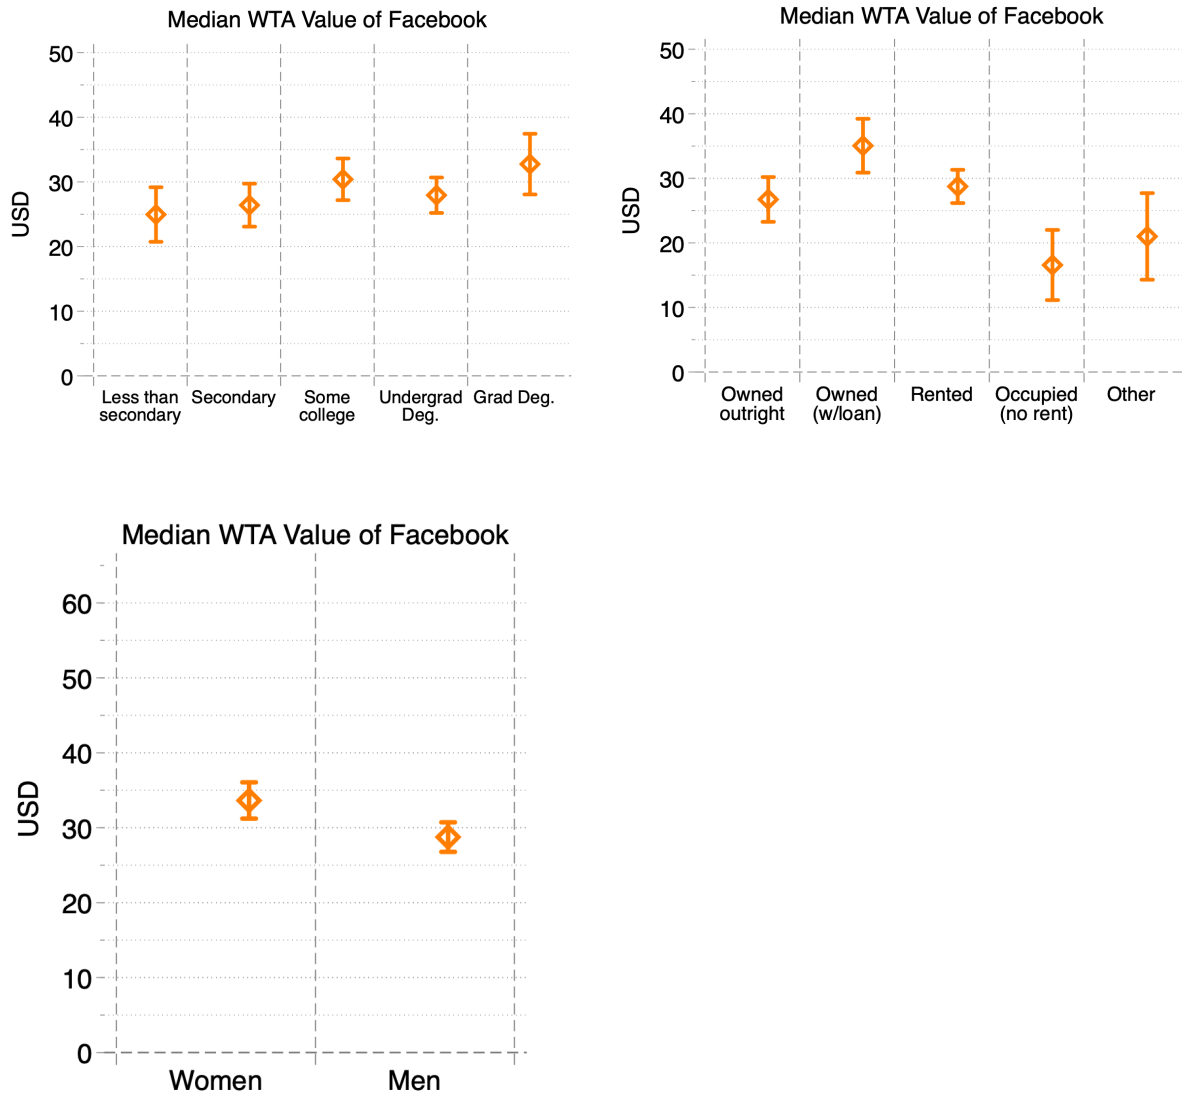
Would you be willing to stop using Facebook for one month in exchange for \$40? You may be randomly chosen and offered \$40 to stop using Facebook based on your answer to this question.

- No, I am not willing to stop using Facebook for one month in exchange for \$40
- Yes, I am willing to stop using Facebook for one month in exchange for \$40


Continue

Appendix Figure 4: Example Facebook offer screen


Additional Results and Robustness Checks



Appendix Figure 5: Median Value of Facebook by Education, Home Ownership & Gender

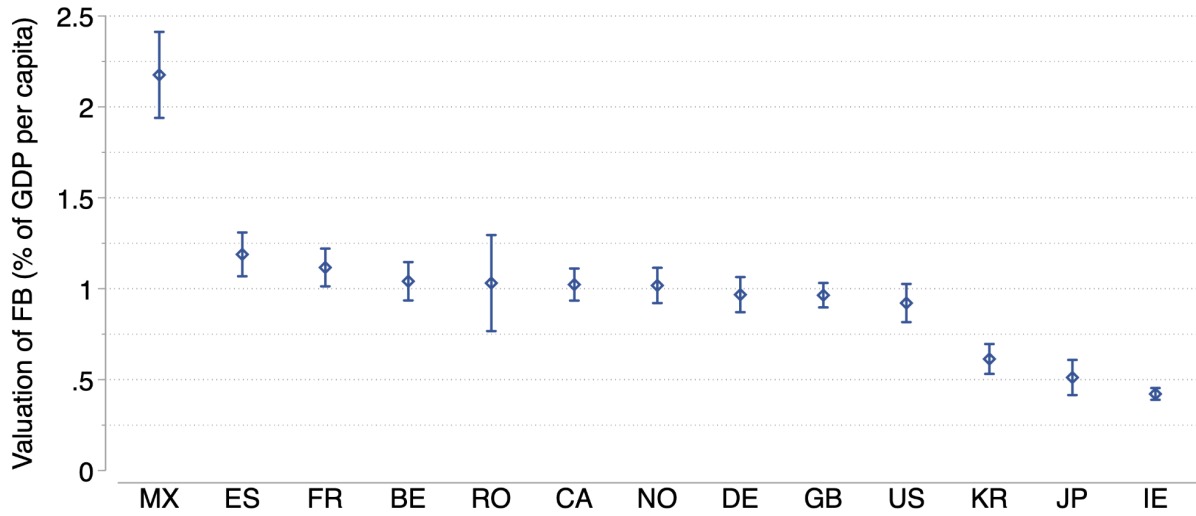


Appendix Figure 6: Facebook aggregate annual value by country from incentivized single binary discrete choice experiments




Appendix Figure 7: Relationship between valuation of digital goods and GDP per capita (excluding Mexico).


Appendix Figure 7 Notes: The figure is analogous to Figure 3b but excludes Mexico. A weighted regression of Y on X with weights given by Facebook monthly active users (MAU) in each country yields a point estimate of 0.90 with a p-value equal to 0.001.



Appendix Figure 8: Valuation of Facebook as a share of GDP per capita




Appendix Figure 9: Aggregate annual value of 10 digital goods as share of GDP per capita




Appendix Figure 10: Relative utility by relative wealth index, with Facebook as the omitted category (US Only)

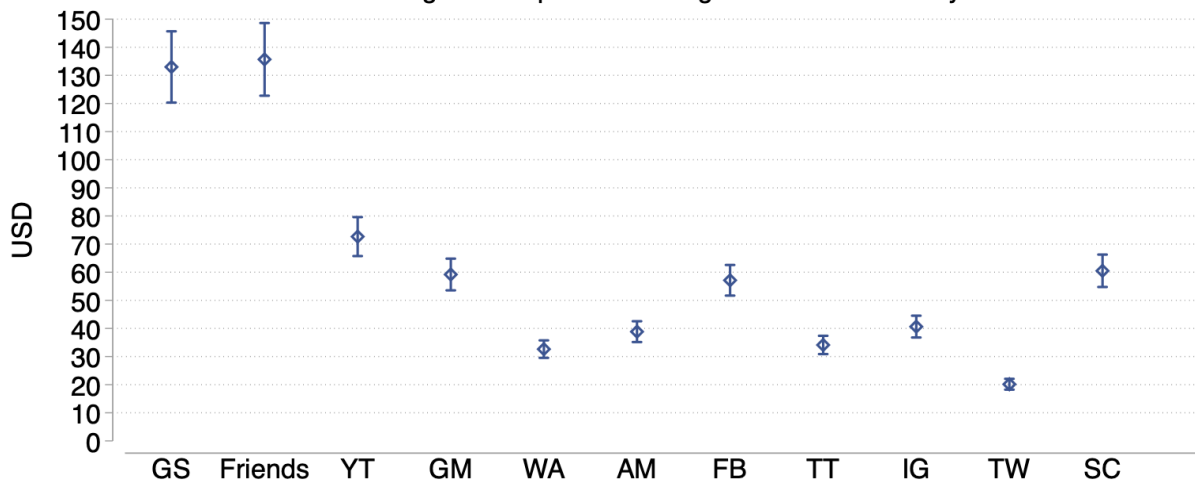
Willing-to-Accept Value of Digital Goods in United States



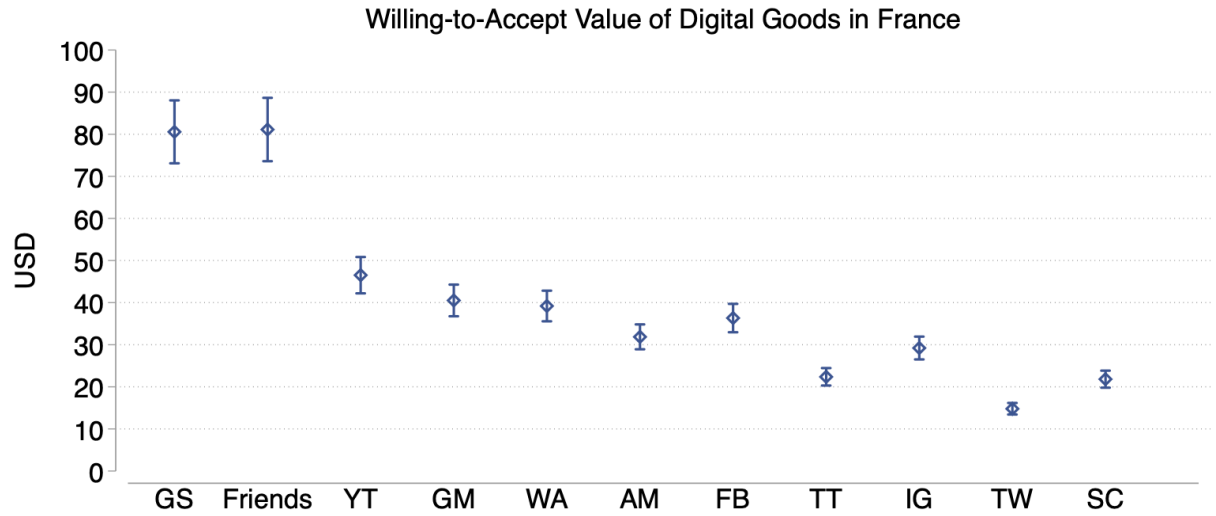





Willing-to-Accept Value of Digital Goods in Germany

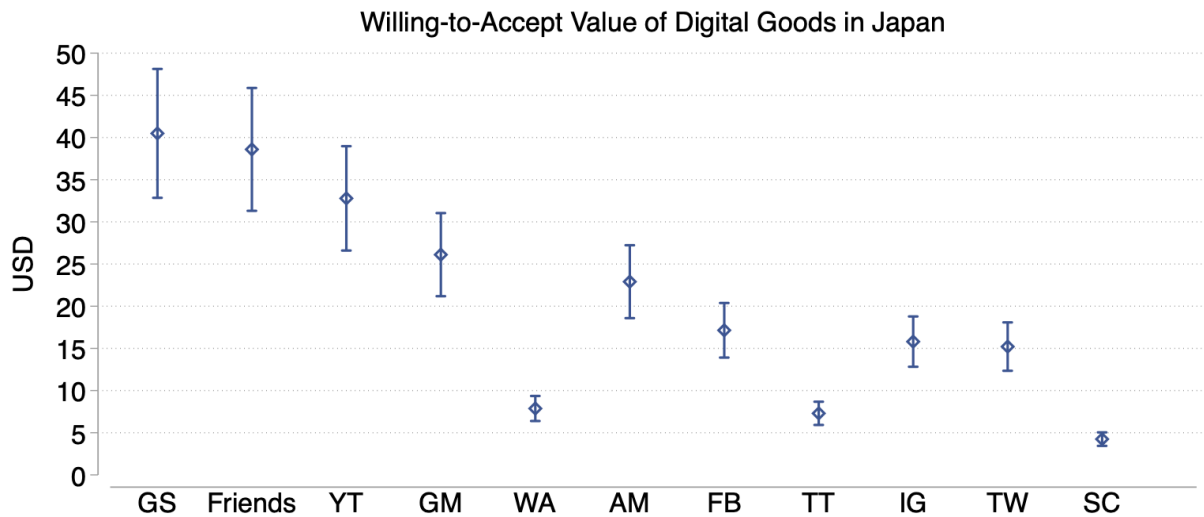





Willing-to-Accept Value of Digital Goods in Norway











Appendix Figure 11: Implied Median WTA Value of Each Digital Good by Country