

NBER WORKING PAPER SERIES

WHEN PRODUCT MARKETS BECOME COLLECTIVE TRAPS:
THE CASE OF SOCIAL MEDIA

Leonardo Bursztyn
Benjamin R. Handel
Rafael Jimenez
Christopher Roth

Working Paper 31771
<http://www.nber.org/papers/w31771>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2023, Revised July 2024

We thank the Editors (Pietro Ortoleva and Erzo F.P. Luttmer) and three anonymous referees for very insightful comments which helped us improve the paper. We are also grateful to Daron Acemoglu, Eliana La Ferrara, Matthew Gentzkow, Alex Imas, Ro'ee Levy, Matthew Lowe, Pietro Ortoleva, Ricardo Perez-Truglia, Aakaash Rao, Gautam Rao, Jean Tirole, as well as many seminar participants for useful comments. Caroline Boone, Pietro Ducco, Maximilian Fell, Paul Grass, Anastasiya Nebolsina, Ewan Rawcliffe, Christoph Ziegler, and especially Beyza Gulmezoglu provided excellent research assistance. Roth acknowledges funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC 2126/1-390838866. The research described in this article was approved by the University of Chicago Social and Behavioral Sciences Institutional Review Board and pre-registered in AsPredicted (#137878, #142247, and #144630). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Leonardo Bursztyn, Benjamin R. Handel, Rafael Jimenez, and Christopher Roth. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

When Product Markets Become Collective Traps: The Case of Social Media
Leonardo Bursztyn, Benjamin R. Handel, Rafael Jimenez, and Christopher Roth
NBER Working Paper No. 31771
October 2023, Revised July 2024
JEL No. D62,D91

ABSTRACT

Individuals might experience negative utility from not consuming a popular product. With such externalities to non-users, standard consumer surplus measures, which take aggregate consumption as given, fail to appropriately capture consumer welfare. We propose an approach to account for these externalities and apply it to estimate consumer welfare from two social media platforms: TikTok and Instagram. Incentivized experiments with college students indicate positive welfare based on the standard measure, but negative welfare when accounting for these non-user externalities. Our findings highlight the existence of product market traps, where large shares of active users prefer each platform not to exist.

Leonardo Bursztyn
Department of Economics
University of Chicago
1126 E. 59th Street
Chicago, IL 60637
and NBER
bursztyn@g.uchicago.edu

Benjamin R. Handel
Department of Economics
University of California, Berkeley
508-1 Evans Hall #3880
Berkeley, CA 94720
and NBER
handel@berkeley.edu

Rafael Jimenez
Bocconi University
Via Sarfatti 25
Milano
Italy
rafael.jimenez@unibocconi.it

Christopher Roth
Department of Economics
University of Cologne
Albertus-Magnus-Platz
Cologne
Germany
cproth89@gmail.com

A randomized controlled trials registry entry is available at AsPredicted #137878, #142247, and #144630.

1 Introduction

Much of consumption is highly social in nature. In many contexts, the utility that an individual derives from consuming a product or service increases as more people consume it. Going to a concert or dinner with friends is more enjoyable than going alone. Yet, consumption can also negatively affect others (Frank, 2005). Indeed, the literature on conspicuous consumption and positional externalities (Frank, 1985; Bursztyn et al., 2018; Imas and Madarász, 2022; Pesendorfer, 1995) has highlighted that one’s utility can be negatively impacted by others’ incomes or consumption, for instance, as a result of social comparisons (Bottan and Perez-Truglia, 2022; Clark and Oswald, 1996; Cullen and Perez-Truglia, 2022; Luttmer, 2005; Perez-Truglia, 2020).

These social forces play a vital role in the context of social media. For a given platform, a larger number of users may increase the benefits of joining, by expanding the network of individuals available for interaction. Beyond that, the size of the network may also affect the utility of potential non-users. Such externalities to non-users can be driven by mechanisms such as social exclusion or a fear of missing out (Gupta and Sharma, 2021). As the total number of platform users increases, marginal users may participate because they want to avoid the negative externalities imposed on non-users but may still have negative overall utility from the platform’s existence.

In the presence of such externalities to non-users, standard measures of consumer surplus that take aggregate consumption of a product as given do not appropriately capture the welfare of its *users*. In particular, when non-user utility is negative, these measures overstate the total welfare associated with the product because they use an incorrect outside option; namely, not consuming the product while holding fixed others’ consumption. Instead, the relevant outside option for calculating welfare in the presence of negative non-user utility is the non-existence of the product market. Negative non-user utility can also give rise to *product market traps*: a situation similar to a Prisoner’s Dilemma where some users would prefer the product not to exist, yet they find it optimal to consume it. In such traps, some users’ utility is negative but would have been even more negative had they not used the product, which is why they continue using it. Such traps can arise from social forces even with fully rational expectations and without behavioral frictions, such as a lack of self-control and naivete.

In this paper, we propose an approach to measure consumer welfare in the presence of such externalities to non-users and network effects, and apply it to the welfare analysis

of social media platforms. We implement our methodology in pre-registered online experiments with more than 1,000 students from various colleges in the US. We focus on two prominent social media platforms: TikTok and Instagram. These platforms have been the subject of concern, among other reasons, due to their potential adverse effects on mental health (Faelens et al., 2021).¹

In the experiment, we employ standard tools to measure consumer welfare, leveraging an incentivized Becker-DeGroot-Marschak (BDM) mechanism (Becker et al., 1964), which we implement using an iterative multiple price list. The experiment proceeds in three main steps. In Step 1, we measure individual-level willingness to accept (WTA) to deactivate one’s social media for four weeks while keeping constant others’ social media consumption. This step provides us with the standard measure of individual consumer surplus (*Valuation Keeping Network*). In Steps 2 and 3, we plausibly reduce the size of our respondents’ networks by presenting the possibility of a large-scale deactivation where all participating students at their university deactivate their accounts. Participants are told that this large-scale deactivation will be conducted if we recruit two-thirds of students at their university. To measure network effects, in Step 2 we measure individual WTA to deactivate conditional on all other participating students having been asked to deactivate their account in exchange for monetary compensation (*Valuation Removing Network*). Finally, in Step 3, we measure welfare differently, taking as the outside option the non-existence of the market. To do so, we elicit individuals’ preferences over the deactivation of the social media accounts of all participating students, including themselves. In particular, we measure whether individuals are willing to forego payment or instead require a payment to deactivate all participating students’ accounts (*Product Market Valuation*).

Our main results highlight the importance of accounting for externalities to non-users. Our individual-level elicitation reveals an average individual consumer surplus of \$55 and \$47 for TikTok and Instagram, respectively, with 92% and 86% of users deriving positive welfare from the products. These findings are in the ballpark of estimates in the literature (Mosquera et al., 2020; Allcott et al., 2020) and indicate that users require substantial compensation to stop using social media when others in their network keep using it.

We next turn to product market surplus, our preferred measure of welfare that accounts for externalities to non-users. Our main finding is that 60% and 46% percent of active

¹We measured TikTok’s welfare in a survey conducted in July 2023, and Instagram’s welfare in a survey conducted in August and September 2023. Both surveys are virtually identical, except that in the second survey we added more questions and clarified some of the instructions. We describe the differences in Section 3.2 and present the full set of instructions in Appendix G.

TikTok and Instagram users, respectively, experience negative welfare from the product’s existence. Average product market surplus is significantly lower compared to individual consumer surplus for both TikTok ($p < 0.01$) and Instagram ($p < 0.01$). Users have an average willingness to pay (WTP), rather than a willingness to accept, of \$24 and \$6 to have others, including themselves, deactivate TikTok and Instagram, respectively. Overall, our evidence shows the existence of a *social media trap* for a large share of consumers, who find it individually optimal to use the product even if they derive negative welfare from it.

Finally, we present our estimates of network effects by comparing the valuation removing the network against the valuation keeping the network. The fraction of users with positive welfare drops to approximately 72% and 69% of users for TikTok and Instagram, respectively. Compared to the valuation keeping network, the average willingness to accept significantly drops by approximately 29% and 21%, to \$39 and \$37 for TikTok and Instagram, respectively. This drop provides evidence that network effects are positive and quantitatively significant, consistent with canonical theoretical frameworks (Rohlf, 1974; Katz and Shapiro, 1985). Moreover, the fact that the valuation removing the network is larger than the product market valuation is consistent with recent evidence of preferences for exclusion (Imas and Madarász, 2022).

To ensure high levels of understanding, we restrict our main analysis to respondents who pass several attention checks and do not regret their choices, though our results are robust to considering different samples (including regretters and inattentive respondents) and measurement error corrections as in Luttmer and Samwick (2018). Moreover, we confirm our findings using a hypothetical qualitative question in which we ask respondents whether they would prefer to live in a world without the social media platform. Indeed, most respondents and a large share of users in our samples would prefer to live in a world without TikTok and Instagram, respectively. Similarly, another hypothetical question reveals that respondents favored the option of everyone deactivating their accounts over only deactivating their own account or no one deactivating their accounts, for both platforms.

One possible concern with our empirical design is that respondents may think that it is unlikely that we will actually conduct the large-scale deactivation study. However, the perceived likelihood that the large-scale Instagram deactivation study will be implemented is high, at approximately 45%. Moreover, for respondents deeming the large-scale deactivation study more likely, the estimated product market surplus is even more negative, suggesting that our elicitation provides a conservative estimate of how negative the product market surplus is. More broadly, given that even in the case of the large-scale deactivation

study not all users would deactivate their account, our study plausibly identifies lower bounds for the size of the negative product market surplus.

One conjecture is that the drop in welfare from the individual to the product market surplus could be fully driven by factors such as a “repugnance” (Roth, 2007) towards digital products, animus against big tech companies, or a distaste for others spending time on their phone. To rule out this possibility, we conduct an experiment with an identical design but with a product that creates plausibly less pronounced negative externalities for non-users: navigation and maps smartphone apps (hereafter referred to simply as “Maps”). For these apps, our estimates of product market surplus remain positive, large, and highly significant and highly significant ($p < 0.01$). Besides elucidating the underlying mechanisms, the positive product market surplus for Maps also suggests that the negative product market surplus we document for TikTok and Instagram is not driven by mechanical factors such as the way we frame our elicitation.

The wedge between individual consumer surplus and product market surplus highlights an important role of externalities to non-users. To shed light on the motives behind active users’ preferences for living in a world without their social media platform, we ask them an open-ended question on why they still use the platform. This data indicates that the fear of missing out is the most prevalent motive for both TikTok and Instagram. Paired with our main estimates, the evidence of these underlying mechanisms supports the notion that accounting for externalities to non-users is crucial to assessing the welfare effects of social media platforms. These externalities to non-users may arise from anticipated social exclusion that would actually occur in case of deactivation or could arise from misperceptions or other psychological biases.

One implication of our framework is that producers have incentives to use technologies or marketing campaigns that decrease non-user utility—increasing the cost of not consuming the product. Indeed, large tech companies commonly use tools that might decrease non-consumer surplus, such as increasing the salience of being a non-consumer or tying messaging apps and social media platforms. An example of such technology is the case of the “green bubble” messages on iPhones, which make it salient for iPhone users when they exchange text messages with non-iPhone users. The social stigma arising from this green-bubble culture has received widespread attention in the mass media.² In this case, the green bubbles feature might increase demand for the iPhone, not because it improves

²For popular press coverage on the “green bubble culture:” See “Why Apple’s iMessage Is Winning: Teens Dread the Green Text Bubble” Higgins, Tim. *The Wall Street Journal*, January 8, 2022.

its “intrinsic” value, but because it avoids the social stigma of not using an iPhone. More generally, our findings challenge the standard revealed-preference argument that the mere existence of a product implies positive welfare for its consumers, even if they are fully rational. Indeed, we provide evidence of a product that is consumed by a large share of individuals, even when it creates negative welfare for many of them. This finding suggests a heightened need for regulators to assess whether different products create traps for consumers and, potentially, diminish competition between platforms. More broadly, these patterns could apply to other markets. We provide suggestive hypothetical survey evidence of a large fraction of consumers preferring to live in a world without luxury goods and to slow down the release frequency of products with different vintages.

Our paper also speaks to work assessing the welfare generated by social media (Brynjolfsson et al., 2019; Mosquera et al., 2020; Allcott et al., 2020, 2022; Brynjolfsson and Oh, 2023; Brynjolfsson et al., 2023). The papers in this space measure consumer surplus by either taking the aggregate level of consumption as given or assuming that externalities to non-users are zero. Existing work finds large user valuations for social media, consistent with the large amount of time spent on these platforms (2.5 hours per day on average (Kemp, 2022)), while at the same time documenting that the expansion and use of these platforms can harm individual well-being and mental health (Allcott et al., 2020; Braghieri et al., 2022). Our results on the switch in signs of consumer welfare after accounting for non-user utility help reconcile these seemingly contradictory findings and paint an integrated and more pessimistic picture of the welfare effects of social media. Additionally, we provide the first incentivized evidence of network effects in the context of social media, which has proven difficult aside from hypothetical estimates (Benzell and Collis, 2022).

We also contribute to a long-standing literature in industrial organization that models consumer choice in the presence of network effects (Rohlfs, 1974; Katz and Shapiro, 1985; Farrell and Klemperer, 2007; Rochet and Tirole, 2003). Our work differs from this literature in a few key ways. First, a standard procedure in this literature is to normalize the utility from not using a product to zero, effectively ruling out externalities to non-users.³ We develop an experimental framework to elicit the magnitude of network effects and externalities to non-users. We simultaneously identify and quantify both positive network effects for users and negative externalities to non-users. Second, the literature has pointed

³One exception is Bhattacharya et al. (2023) who also relax this assumption and show that welfare effects are not point identified in models with externalities. They apply their model to the evaluation of the welfare effects of bed nets in a discrete-choice econometric framework with a focus on health externalities arising from contagious diseases.

to coordination failures that arise in the presence of network effects in cases where one firm becomes dominant despite not being the most efficient supplier (Farrell and Saloner, 1985; Farrell and Klemperer, 2007). While this coordination failure occurs in the presence of multiple competing platforms and externalities among product users, our work highlights the possibility of a Prisoner’s Dilemma that can arise even with a single platform due to the presence of externalities to non-users that lock consumers into using the product.⁴

This paper proceeds as follows: Section 2 provides a simple conceptual framework for measuring welfare in the presence of externalities to non-users. Section 3 provides the empirical design. In Section 4, we present results for individual and product market surplus and network effects. Finally, Section 5 discusses the policy implications of our findings.

2 Conceptual Framework: Product Market Traps

Setup. There is a set of individuals and an indivisible product. Individual i derives utility from their own consumption of the product, $x_i \in \{0, 1\}$, and from the fraction of other individuals who consume it, X . For ease of exposition, we assume that utility is quasilinear in income, given by $u_i(x_i, X) - p$, where p is the price of the product.

We leverage the presence of X in the utility to model two distinct phenomena: consumption externalities and network effects. First, we allow for consumption externalities; the extent to which utility changes in response to others’ consumption. Concretely, i exhibits positive (negative) consumption externalities from the product if their utility increases (decreases) when the fraction of others consuming it increases. Without loss of generality, we normalize to zero the utility that i receives when no one else consumes the product, $u_i(0, 0) = 0$. Besides this normalization, most prior work on network effects, and prior empirical work on social media, assumes a constant *non-user utility*; that is, $u_i(0, X) = u_i(0, X')$ for all X, X' , which implies that $u_i(0, X) = 0$ for all X . We relax this assumption and allow for consumption externalities for non-users so that, generically:

$$u_i(0, X) \neq u_i(0, X'), \quad X \neq X'.$$

⁴While we focus on the latter coordination failure, in the presence of multiple platforms both kinds of coordination failures could be present simultaneously. However, the empirical patterns of social media use, paired with our finding of negative welfare among single- and multi-homers (those who use one platform or multiple platforms, respectively), suggest product market traps above and beyond the coordination failure in Farrell and Saloner (1985). This relates to recent literature documenting forms of Prisoner’s Dilemma in the industry generated by different mechanisms (Cheyre and Acquisti, 2024; Sullivan, 2022).

For instance, if the utility of non-users decreases when more people use the product, we say that the product exhibits negative consumption externalities for non-users.

Second, given this relaxation, we need to distinguish between consumption externalities and network effects. We use a definition of direct network effects based on strategic complementarities in consumption. Concretely, i 's utility exhibits positive network effects when their marginal utility of consumption increases with others' consumption:

$$u_i(1, X') - u_i(0, X') > u_i(1, X) - u_i(0, X),$$

for $X' > X$. In the absence of consumption externalities for non-users, this definition is equivalent to the standard definition of network effects in the literature of users having positive consumption externalities, $u_i(1, X') > u_i(1, X)$. However, in our setting, it is possible that positive network effects (as defined above) coexist with negative consumption externalities; that is, $u_i(1, X) < u_i(1, 0)$, for $X > 0$. This flexibility allows, for example, the presence of both preferences for exclusivity (which manifest as negative consumption externalities) and positive network effects.

Welfare Measures. The standard measure of individual consumer surplus compares the utility that i gets relative to their utility when they do not consume the product, given a fraction of others consuming it. We refer to this measure as the *individual* consumer surplus, ICS , in the sense that it only accounts for i 's individual choice:

$$ICS_i(p, X) := \begin{cases} u_i(1, X) - p - u_i(0, X) & \text{if } i \text{ consumes, } u_i(1, X) - p \geq u_i(0, X) \\ u_i(0, X) - u_i(0, X) = 0 & \text{if } i \text{ does not consume, } u_i(1, X) - p < u_i(0, X). \end{cases}$$

In practice, researchers estimate this measure by eliciting individuals' willingness to accept to give up a product in exchange for a monetary payment, or their willingness to pay to get it, holding constant the others' consumption.

We define *product market surplus* as i 's utility from consuming relative to their utility when no one consumes:

$$PMS_i(p, X) := \begin{cases} u_i(1, X) - p - u_i(0, 0) = u_i(1, X) - p & \text{if } i \text{ consumes} \\ u_i(0, X) - u_i(0, 0) = u_i(0, X) & \text{if } i \text{ does not consume.} \end{cases}$$

The key difference between these measures is the outside option each uses. Product market

surplus is better suited to measure welfare from the product’s existence because it correctly compares the utility that consumers and non-consumers get from a product to their utility in the absence of it. Without consumption externalities, both measures are identical. More generally, however, the individual consumer surplus will be biased upwards or downwards depending on whether $u_i(0, X)$ —the non-user utility—is negative or positive, respectively:

$$ICS_i(p, X) = PMS_i(p, X) - u_i(0, X).$$

For example, when i has a fear of missing out, their individual surplus will be biased upward, as it reflects not only their valuation of the product but also their distaste for being left out when they do not consume it.

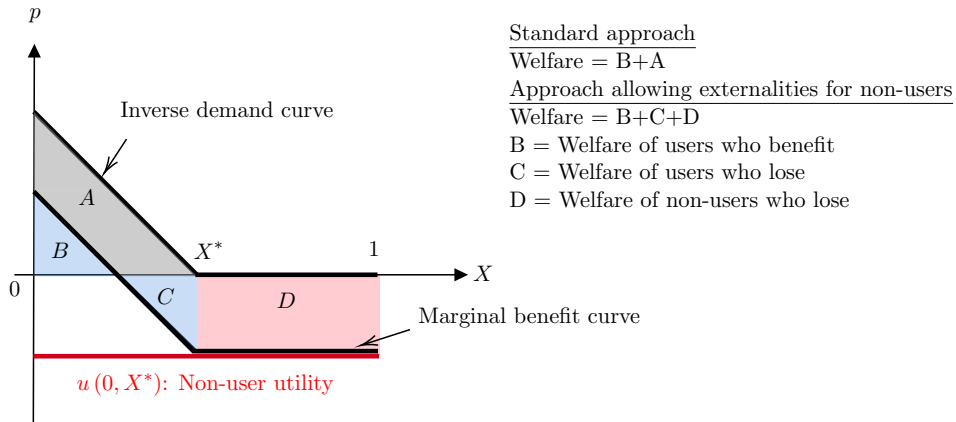
Product Market Traps. Our framework allows for the possibility of a *Product Market Trap* for individuals; where, for a given price and aggregate consumption of the product:

- (i) i chooses to consume the product: $ICS_i(p, X) > 0$.
- (ii) i would be better off if no one consumed it (i ’s welfare is negative): $PMS_i(p, X) < 0$.

Note that these conditions imply that consumer i experiences negative non-consumer surplus ($u_i(0, X) < 0$); negative externalities to non-users are necessary to generate product market traps. Individuals in a product market trap would like to coordinate with others to not consume, but they cannot commit. They are “trapped” into consuming because others do so. Hence, the revealed-preference argument that the existence of a product implies that users benefit from it fails to apply.

The experiments we describe below seek to estimate individual consumer surplus and product market surplus in the context of social media. Figure 1 illustrates how our welfare calculations differ from the standard setting that ignores externalities for non-users, given the observed market share, X^* , and a price fixed at 0. For simplicity, we assume a mass unit of individuals with homogeneous non-user utility $u(0, X^*)$ and uniformly distributed user utility $u(1, X^*)$. The standard welfare analysis calculates the average individual consumer surplus, equal to areas A + B. This approach concludes that those who use the product benefit from it and that everyone who does not use it gets zero welfare. In our framework, the welfare impact is: positive and equal to B for users who benefit from the product, negative and equal to C for users who lose out from the product, and negative and equal to D for non-users who lose out from the product’s existence, giving a total welfare B+C+D.

Figure 1: Comparison between Welfare Measures Given Market Share X^*



Notes: This figure illustrates how welfare calculations differ between our setting and the standard setting that ignores consumption externalities for non users, conditional on an observed market share equal to X^* and a price equal to zero. For ease of exposition, this figure assumes a mass unit of individuals with homogeneous non-user utility, $u(0, X^*)$, and uniformly distributed user utility $u(1, X^*)$. The inverse demand curve describes $ICS_i(0, X^*)$ at every level of demand. The marginal benefit curve describes $PMS_i(0, X^*)$ at every level of demand.

3 Measuring Individual and Product Market Surplus

3.1 Sample

College Student Sample. We recruited college students to participate in our experiments through a partnership with College Pulse, a company specialized in recruiting college students for online experiments with a panel of 650,000 college students at the time of the study. We focus on college students for various reasons. First, they are of high policy relevance as they are among the most active on social media.⁵ Second, social media usage has been linked to the increasing prevalence of depression among college students (Braghieri et al., 2022). Third, even if other fellow college students might not represent the entire network of friends of our participants (corresponding to X in our theoretical framework), they constitute a significant subset of students' social networks.⁶

⁵In an April 2021 survey, 84% of adults between the ages of 18 and 29 reported using social media (Pew Research Center, 2021). Among our respondents, 92% stated to have used Instagram at least once in the past month, with 74% using it daily. Meanwhile, 68% mentioned using TikTok at least once in the past month, and 52% did so daily.

⁶Our respondents estimate an average 60% of their mutual friends on Instagram are fellow college students, indicating that the college social network constitutes a majority of our respondents' social networks.

Pre-registration. The pre-registrations include the experimental design, hypotheses, analysis, sample sizes, and exclusion criteria. The pre-registrations of the two data collections can be found on AsPredicted #137878 and #142247.⁷ While we pre-registered pooling 28 pilot responses with the pre-registered data in the Instagram and Maps experiment, we deviate from this plan and instead only report the pre-registered data. Appendix F shows that our main results remain unchanged including this pre-registered data.

TikTok. In July 2023, we recruited 1,713 respondents who began our experiment, out of which 66% had used TikTok in the past month, our measure of activity on the platform.⁸ All active users are then asked whether they are willing to participate in the deactivation study. Fifty-seven percent of TikTok users in our sample were willing to provide their handle to participate in the study. Much of this selection does not simply arise from an unwillingness to deactivate their accounts, with 40% of participants mentioning privacy concerns and 32% mentioning the fear of missing out as motives for not being willing to participate in the study (see Appendix Figure A2). Nonetheless, this selection into our study severely limits our ability to measure average welfare among the population of college students.

We restrict our sample to respondents aged between 18 and 30 and we exclude respondents who failed any attention checks or regretted their valuations for a second time.⁹ Our final sample consists of 707 college students, 371 TikTok users, 336 non-users.

Instagram and Maps. To provide evidence for a second social media platform and for another smartphone application that is not a social media platform, we recruited college students who had not taken our TikTok experiment to participate in a second wave in August and September 2023.¹⁰ Respondents were randomly assigned to complete a version of the experiment about (i) Instagram or (ii) Maps (the following navigation and maps

⁷We also pre-registered running an in-person experiment at the University of Chicago at the end of May 2023 (https://aspredicted.org/WDF_DFH). However, we failed to recruit the minimum pre-registered number of participants and only managed to collect 12 pre-registered responses from active users who passed the sample inclusion criteria. The very small sample thus makes it difficult to draw meaningful conclusions.

⁸This sample size excludes incomplete survey responses and duplicate respondents. Including such respondents, our initial sample is 2,688.

⁹To further ensure high data quality, we also exclude six respondents who leave the compulsory open-ended questions in our survey blank.

¹⁰Our experiments were conducted during a time when universities were on summer break. This design choice has implications for respondents' valuation of social media as it may depend on the ease of in-person interaction. It seems conceivable that social media has a higher value in situations when respondents cannot readily substitute online social interactions with in-person interactions.

smartphone apps: Google Maps, Apple Maps, and Waze). All active users are asked for their willingness to participate in the deactivation study, while non-users proceed directly to the practice questions and the *Product Market Valuation*.

We randomize a total of 935 and 854 respondents into the Instagram and Maps experiments, respectively.¹¹ Out of those, 94% reported actively using Instagram and 99.8% reported using Maps.¹² All active users are then asked whether they are willing to participate in the deactivation study. Forty-six percent of Instagram users in our sample were willing to provide their handle to participate in the study, while 58% of Maps users were willing to participate. As with TikTok, much of this selection is not simply a result of their unwillingness to deactivate their accounts, with 32% and 27% of participants mentioning privacy concerns for Instagram and Maps, respectively, and 41% and 17% mentioning “not wanting to be without their account while their friends are still on the platform” as motives for not participating in the study, for Instagram and Maps, respectively.¹³ While this selection into the study makes it difficult to make statements about average welfare among the population of students, the hypothetical survey questions that we elicited also among those unwilling to participate allow us to examine the nature of selection. Even among respondents unwilling to participate in our Instagram deactivation study, a large share (39%) prefer living in a world without Instagram, compared to 57% among those users willing to participate. While this data provides evidence that those unwilling to participate likely derive higher welfare from the product, the fraction deriving negative utility is still quite high.¹⁴

Our main sample consists of 235 Instagram users, 25 respondents not active on Instagram, and 272 Maps users. As with TikTok, we restrict our sample to respondents aged between 18 and 30 and exclude—as pre-specified—respondents failing any attention checks

¹¹This number excludes incomplete survey responses and duplicate respondents. Including such respondents, our initial sample is 1,444 for Instagram and 1,446 for Maps.

¹²To assess how representative our sample is in terms of social media usage, we compare it to data obtained from the American Trends Panel of the Pew Research Center (2021). To increase comparability with our sample, we filter the data by age and education, to approximate a sample of college students. Specifically, we narrow the data to those in the age category of “18-29” and the education category of “Some college, no degree.” Among respondents in this filtered sample, 54% and 75% reported to use TikTok and Instagram, respectively.

¹³Participation in the deactivation study required respondents to provide their TikTok/Instagram handles and submitting screenshots of their phone’s usage statistics. For Maps, only the usage statistics were required to verify compliance with the deactivation study. Given these hassle costs, the elicitation measures the joint effect of not using the service and the hassle costs.

¹⁴Three percent of the respondents unwilling to participate in the Maps deactivation study prefer living in a world without Maps.

or regretting any of their final valuations.¹⁵

Sample Characteristics. Respondents in our samples are all undergraduate students from 365 universities attending 4-year colleges. On average, 7% of the undergraduate student body is part of College Pulse in the colleges in our sample. This is above the entire College Pulse average of approximately 4%. Sixty-four percent of our sample are from public universities, while 36% are from private universities.¹⁶ Moreover, our sample is well-spread out across the United states, making it geographically fairly representative. The majority of our students (59%) attend universities in the top 150 in the U.S. News ranking of universities. Only a relatively small fraction of students (9%) from our sample attend top 20 universities. Our sample is mostly female: among respondents, 68% are female, while among active users 72% are female. Their average age is 21 years.¹⁷

3.2 Design

The purpose of the experiment is to measure welfare while accounting for externalities to non-users. Below we describe the core experimental instructions. The full set of instructions can be found in Appendix G.

TikTok and Instagram. Our main evidence focuses on consumers’ valuation of two popular social media platforms, TikTok and Instagram, that have been the subject of concern regarding their impact on individual well-being. While both TikTok and Instagram are social media platforms that focus on visual content, they differ in several key ways. TikTok specializes in short-form video content, often featuring music, dance, and challenges, and utilizes a unique algorithm that prioritizes content discovery, allowing even unknown creators to go viral. Instagram, on the other hand, started as a photo-sharing platform and its discovery mechanisms are more reliant on existing social networks and hashtags, making it generally harder for new creators to gain visibility.

¹⁵All of our attentive respondents use Maps. As opposed to the case of Instagram and Maps, our TikTok pre-registration did not specify dropping inattentive users or those who regret their choices, but we add these filters to increase comparability across samples. However, results are similar without these filters (see Appendix Figure A8).

¹⁶Among the respondents enrolled in private universities in our sample, 98% are not-for-profit and 2% are for-profit.

¹⁷We assess the representativeness of our sample based on these observables against the American Trends Panel of the Pew Research Center (2021). In the ATP data, 68% of TikTok users and 55% of Instagram users that are college students and aged 18-29 identify as female. In our final sample, 75% of TikTok users and 68% of Instagram users identify as female.

Overview. We now turn to the structure of our experiments, which is also summarized in Figure 2 for the case of the Instagram and Maps experiment.¹⁸ For active social media users, the experiment proceeds in four steps. In Step 0, we measure individual-level WTA to deactivate an example product: a ride-sharing app. This elicitation considers the individual-level decision conditional on aggregate consumption and is meant to accustom respondents to the instructions. In Step 1 (*Valuation Keeping Network*), we measure individual level WTA to deactivate one’s social media account for a period of four weeks taking others’ social media consumption as given. In Steps 2 and 3, we present respondents with the possibility of a large-scale deactivation study where all participating students at their university deactivate their accounts. In Step 2 (*Valuation Removing Network*), we measure individual WTA conditional on all participating students being asked to deactivate their account in exchange for monetary compensation. In Step 3 (*Product Market Valuation*), we measure individuals’ preferences over the deactivation of social media accounts of all participating students, including themselves. In particular, we elicit students’ WTP or their WTA to deactivate everyone’s account.

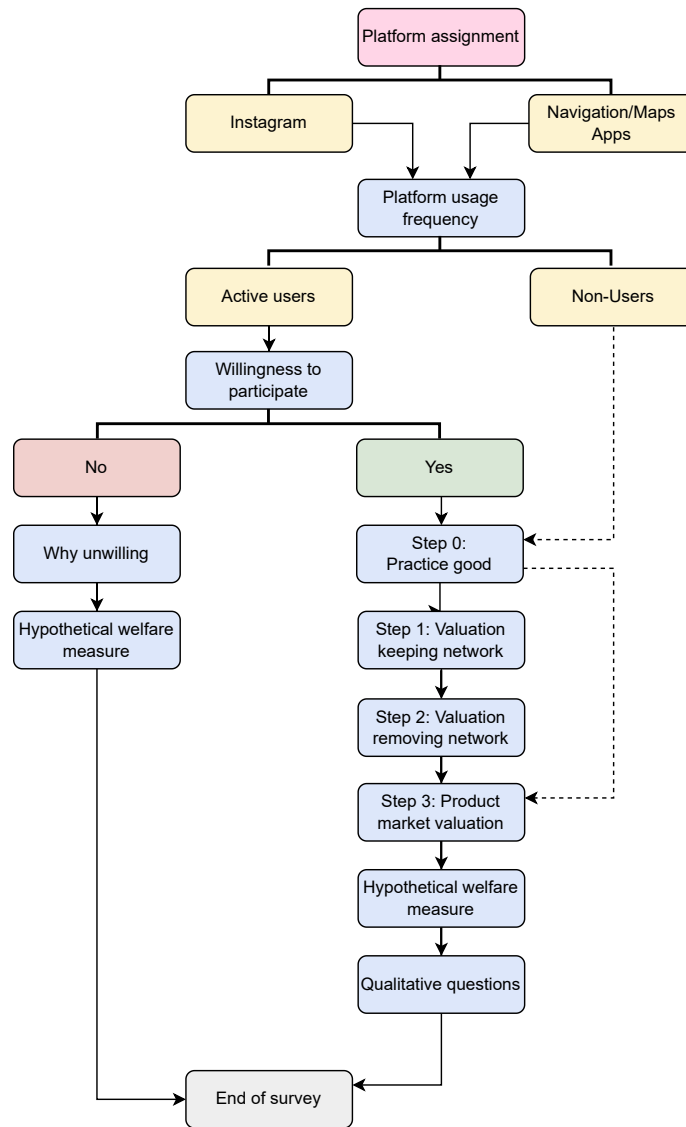
Respondents who are not active social media users take a modified version of the experiment. After completing the practice, they proceed to a customized *Product Market Valuation*, where we measure their preferences over the deactivation of social media accounts of all participating students who are active social media users.

Introduction. We inform all respondents that we will conduct a deactivation study in which we will ask students at their university to deactivate their social media accounts for four weeks in exchange for monetary compensation. To enhance the credibility of our deactivation study we inform them that “deactivation studies like this have been conducted in the past (e.g., by Allcott et al. (2020) and Mosquera et al. (2020)).” We explain that they can go back to using their account whenever they want, with their content and network unchanged, but they would then forgo any monetary payment. We also tell respondents that, to verify that they deactivate their accounts, we will visit their profiles and require them to upload screenshots of their app usage.¹⁹ To ensure high levels of attention, we inform respondents they will receive an additional bonus payment if they correctly respond

¹⁸The TikTok experiment has a similar structure but with only one platform. This structure applies to both social media (TikTok and Instagram) and Maps users, but for simplicity, we refer to all these platforms as “social media.” We reintroduce the distinction when we talk separately about Maps in section 4.4.1.

¹⁹The use of screenshots of participants’ app usage prevents them from substituting between different accounts of the same platform.

Figure 2: Structure of Experiment: Instagram and Maps



Notes: Figure 2 presents the structure of the experiment. At the beginning of the experiment, the platform is cross-randomized between Instagram and Maps. Active users and non-active users are directed to a distinct path. Active users are asked whether they are willing to participate in a deactivation study. The experiment ends for those unwilling to participate after two subsequent questions. The active users willing to participate are directed to Steps 0 to 3, followed by the hypothetical welfare measure and a series of qualitative questions. Non-users proceed to Steps 0 and 3, as indicated by the dashed arrows. The yellow boxes indicate embedded data, the blue boxes indicate question blocks, and the pink box indicates randomization. The flow of the TikTok experiment from July 2023 is identical except that there was no initial random platform assignment and that we did not elicit hypothetical welfare measures among respondents unwilling to participate in the study. Interested readers can access an identical example of our Instagram/Maps survey at the following link: https://ssd.az1.qualtrics.com/jfe/form/SV_br1qh5JMfdw0Bfg.

to all comprehension questions included in the experiment.

Willingness to Accept Elicitation. The core object of interest in our experiment concerns people’s willingness to accept the deactivation of their social media accounts for four weeks. We combine an incentivized BDM elicitation (Becker et al., 1964) with an iterative multiple price list.

Our MPL places participants’ valuation in one of 12 ranges, with lower and upper limits at \$0 and \$200 and internal increments of \$20: $(-\infty, \$0]$, $[\$0, \$20]$, \dots , $[\$180, \$200]$, $[\$200, \infty)$. In Step 3, we expand the limits to $-\$200$ and $\$200$, to account for the possibility of having a WTP as well as a WTA, resulting in 22 ranges. The algorithm proceeds sequentially, starting from an initial monetary offer and upper and lower bounds for the valuation. In each step, we present respondents with two options: either deactivating their social media account and receiving the monetary offer, or keeping their social media account active. If the respondent accepts the offer (i.e., chooses to deactivate), her upper bound is set to that amount. Similarly, if she rejects the offer (i.e., keeps her account active), her lower bound is set to that amount. The algorithm then selects the next offer as the midpoint between her new bounds, resulting in progressively narrower valuation ranges with each response. The elicitation ends once we can narrow down the respondent’s WTA to a \$20 range or once we surpass one of the upper or lower limits, which can take between 1 and 6 choices depending on the initial random offer and the respondent’s answers.

To ensure that choices are incentive-compatible, we inform respondents that a computer will generate an amount of money to offer them to participate in the deactivation study. We further tell them that we will ask them a series of questions offering them different payment scenarios in case they are selected for the deactivation study. If they accept any price scenario lower than the computer’s offer, we will invite them to the deactivation study and give them the computer’s offer. If, on the other hand, they do not accept any price scenario lower than the computer’s offer, we will not invite them to the deactivation study even if they are the selected participant. To examine comprehension, we ask respondents whether demanding a higher amount affects their likelihood of receiving any payment. Reassuringly, 87% of respondents pass this comprehension check.

Step 0: Practice Good. To enhance comprehension, we start with a hypothetical example good (Dizon-Ross and Jayachandran, 2022). We measure individual-level willingness to accept the deactivation of respondents’ ride-sharing (Uber) accounts, taking aggregate

consumption as given.

Step 1: Valuation Keeping Network. In Step 1, we measure individuals' WTA to deactivate their social media accounts, taking aggregate consumption as given. We tell respondents that, to establish appropriate payment amounts for the deactivation study, we will ask them to decide whether to deactivate their social media account in exchange for different monetary amounts. We also reiterate that one student from their university will be randomly selected to participate in the study. We start the MPL with a randomly drawn offer between \$0 and \$200 in \$20 increments.²⁰ Respondents then proceed to the MPL procedure where they decide between either (i) deactivating their social media account (with none of the students at their university deactivating) and sequentially varying amounts of money or (ii) not deactivating their account.²¹

Step 2: Valuation Removing Network. To assess the role of network effects in shaping individual consumer surplus, we measure individuals' valuation of their social media accounts when all participating students at their university are asked to deactivate their social media accounts. We start by presenting our participants with the possibility of a large-scale deactivation study at their university, where all participating students are asked to deactivate their accounts. In particular, we tell our respondents:

College Pulse has a panel exceeding 650,000 university students. We are targeting universities with a high penetration of College Pulse.

We will now ask you to consider two additional options for a large-scale deactivation of TikTok [Instagram] at your university. One of them will be randomly implemented if we manage to recruit more than two-thirds of the students at your university.

We expect 90% of students to comply with deactivation based on previous studies (e.g., by Mosquera et al., 2018 and Allcott et al., 2020).

²⁰For Instagram, we draw the random offer between \$20 and \$200.

²¹To enhance comprehension and make the choices more intuitive, we added an explanation to the decisions of all 3 steps in terms of "taking a break from social media" or "not taking a break from social media" in the August and September 2023 collection. During the price elicitation, this collection also emphasizes more saliently on the decision screen that the respondent would not receive any monetary payment in cases where the offer equals \$0.

Thereafter, we inform respondents that we will randomly choose one of two options for conducting this larger-scale deactivation study. We then proceed with describing the first option: We tell respondents that we will ask all participating students at their university sequentially whether they would like to deactivate their accounts. We then measure respondents' WTA to deactivate their social media accounts, conditional on us having asked all participating students at their university to deactivate their accounts in exchange for monetary payment. Respondents choose between (i) deactivating their account (when all other participating students have also been asked to deactivate) and receiving varying amounts of money sequentially and (ii) keeping their account active. To economize time, we randomize the initial offer between the lower and upper bounds of the respondent's valuation from Step 1 (unless respondents are at the lower or upper ends of the WTA interval, in which case we offer them again this bound).²²

Step 3: Product Market Valuation. In Step 3, we measure the product market valuation by eliciting individuals' preferences over the deactivation of the social media accounts of all participating students, including themselves.

Respondents are told that we know how much we need to pay every participating student at their university to deactivate their accounts for four weeks.²³ We inform respondents that we will randomly select one of the students to anonymously choose between the following two options: (i) keep things as they are or (ii) deactivate the accounts of all participating students.²⁴ We clarify that if they decide for all participating students to deactivate their accounts, the researchers will pay the other students the amount they require. Moreover, they are told that we will establish their payment, if any, below.

To clarify the incentive compatibility of the mechanism, respondents learn that the deactivation study will be stopped for everyone only if the chosen respondent goes back to using the platform before the end of the four weeks. In particular, the chosen respondent will not receive payment and the other students will be paid based on the actual time

²²We randomize locally, as the Step 1 valuation plausibly constitutes a more precise starting point for the Step 2 elicitation compared to a fully randomized offer.

²³Note that the information we collect through Step 2 provides us with the necessary information to compensate respondents for their individual deactivation in the scenario of the large-scale deactivation that respondents decide upon in Step 3. Since respondents in Step 2 did not anticipate Step 3, both elicitations are incentive-compatible.

²⁴Being pivotal is a low probability event, meaning that there is relatively little expected incentivization. Reassuringly, prior literature finds evidence that stake size many times does not significantly alter behavior (Enke et al., 2023).

they spend in the study. For example, if the study was stopped after two weeks because the chosen respondent did not comply with the deactivation, then another participating student who required a payment of \$40 for the deactivation in Step 2, who complied with the deactivation for the two weeks, would receive a payment of \$20. Finally, if a participating student who was not chosen to decide on the large-scale deactivation goes back to using the platform before the end of the study, they will not receive any payment.²⁵ Table A1 in Appendix A illustrates the incentive compatibility of our elicitation in more detail.

Subsequently, we remind people that the choice they make is incentivized and that their final payoffs, in case Step 3 is implemented, will depend on their valuation as well as the randomly drawn offer. Respondents then proceed to the first main decision screen where they decide between (i) all participating students at their university deactivating their accounts (Option A) and (ii) all participating students at their university keeping their accounts active (Option B) when the deciding participant does not receive payment. This incentivized choice effectively splits the participants' valuation into the positive or negative range. Consider the scenario where a respondent prefers Option A, of all participating students deactivating their accounts. In the following screen, they make a decision between all participating students deactivating their accounts vs. all participating students keeping their accounts active plus a random dollar amount, \$X, drawn between \$20 and \$200 in steps of \$20. If the respondent chooses Option A once again, then she is willing to forgo a payment worth \$X. As in the previous steps, the subsequent offers are made iteratively to narrow down the respondent's WTP. Symmetrically, if she chooses Option B in the first screen, we then employ the iterative MPL algorithm to elicit her WTA to have all participating students deactivate their accounts.

Computing Welfare Based on the MPL. Responses to the MPL questions establish the lower and upper bounds of each respondent's WTA/WTP, effectively assigning them to one of the MPL ranges. For simplicity, we assign the mean of the endpoints for each range in order to have a unique WTA/WTP value; for instance, a range of [\$60, \$80] is assigned a value of \$70. In Section 4.3, we consider an alternative way of assigning valuations.

²⁵This design choice allows our elicitation to account for the option value that individuals might have of reactivating their account even when they initially reported that they wished to deactivate for everyone. By interrupting the study (returning to the status quo) in case the respondent exercises this option, the elicitation compares the utility of the status quo with the utility of joint deactivation. Moreover, the option value is present in all three elicitations. We do not expect individuals to have a differential consideration of these option values across elicitations.

3.3 Discussion of the Design

Elicitation Scales. The scale of our elicitation in Step 3 differs from those used in the other steps. In Step 3, we elicit respondents' positive WTA or their negative WTA (WTP) to deactivate their accounts depending on their response to the first question. In Steps 1 and 2, on the other hand, we only elicit respondents' WTA to deactivate their social media accounts. This difference reflects that it is likely unnatural and inconceivable for individuals to pay to deactivate their accounts individually, when they could deactivate their own accounts for free. Indeed, in contrast to measuring WTA, eliciting WTP in Steps 1 and 2 would not be incentive compatible as there is no way of penalizing respondents with positive WTP if they deactivate their account. In Step 3, on the other hand, WTP is incentive compatible as we can stop the large-scale deactivation in case the respondent chosen to decide deviates from their deactivation. Section 4.3.2 shows that differences in scales are very unlikely to explain differences in valuations across the different steps.

Concerns about Borderline Deception. A key challenge for our design concerns the large-scale deactivation study. The instructions in our experiment rely on language suggesting to participants that the implementation of the large-scale deactivation study is likely. Given that we did not manage to recruit two-thirds of students at any participant's university, no large-scale deactivation study was implemented. While we do not lie to participants, our approach may come close to the boundary of deception.²⁶ We decided to adopt this approach because it appeared to us as the only practically feasible way to elicit valuations for the occurrence of the large-scale deactivation study while maintaining incentive compatibility.

4 The Social Media Trap

4.1 Main Results

We next proceed with presenting our main results for both TikTok and Instagram. First, we present the traditional welfare measure which does not account for externalities to non-users. We then report the results of our preferred measure of welfare accounting for these spillovers. Finally, we present estimates of network effects.

²⁶We received ethical approval for this study from the University of Chicago Social and Behavioral Sciences Institutional Review Board.

Individual Consumer Surplus. Panels (a) and (b) of Figure 3 display the distribution of valuations of the individual consumer surplus for TikTok and Instagram, respectively. These panels illustrate that there is substantial variation in valuations for the individual consumer surplus. As Panels (a) and (b) of Figure 4 show, roughly 90% of users derive positive welfare from both platforms, while the remaining respondents indicate requiring no payment for deactivating their account.²⁷

We next turn to average welfare effects. The dark blue bars in Panels (c) and (d) of Figure 4 show that the individual consumer surplus is large and positive, with a WTA to deactivate of approximately \$50 on average (with a median of \$30) for TikTok and Instagram.²⁸

Product Market Surplus. We next turn to product market surplus, our preferred measure of welfare that accounts for externalities to non-users. Panels (a) and (b) of Figure 3 illustrate that there is strong heterogeneity, with a significant portion of consumers deriving positive and negative welfare from the platforms, respectively.

We then present results on the fraction of active users deriving negative welfare from the platforms. Compared to our estimates of average welfare, this statistic is less susceptible to the framing of response options and to the scales used in this elicitation. Figure 4 shows that 60% and 46% of active TikTok and Instagram users, respectively, have a negative product market surplus. A similar pattern emerges for non-users: 83% and 56% of non-users of TikTok and Instagram, respectively, are willing to pay to have others deactivate their accounts.

Panels (c) and (d) of Figure 4 report the average PMS. Average PMS is significantly lower compared to individual consumer surplus for both TikTok ($p < 0.01$) and Instagram ($p < 0.01$). Hence, the standard individual consumer surplus measure overstates welfare in this context, which is confirmed by Figure A3, where the inverse demand curve lies almost uniformly above the product market surplus curve. On average, users are *willing to pay* \$24 and \$6 to have others, including themselves, deactivate TikTok and Instagram, respectively. The median valuation is -\$10 and \$10 for TikTok and Instagram, respectively. Non-users have an average WTP of \$65 and \$39 to have others deactivate TikTok and Instagram,

²⁷This positive fraction of users requiring no payment could signal that some of them are partly aware of self-control problems and demand commitment devices, as documented by Allcott et al. (2022).

²⁸These estimates are lower than those in Allcott et al. (2020), who find a \$100 median valuation for Facebook. Aside from measuring welfare for different platforms, a possible explanation for the higher valuation uncovered in Allcott et al. (2020) is that their sample consists of more active participants given their recruitment with Facebook ads, while we recruit college students.

respectively. These estimates are statistically significantly below zero for both TikTok ($p < 0.01$) and Instagram ($p < 0.05$).

Our results highlight an important role of externalities to non-users. Many users report a large individual consumer surplus not because they derive positive welfare from the platform, but because they would experience negative utility if they were the only ones to be excluded from it. In that sense, a large fraction of active users are in a *social media trap*. Overall, these findings are evidence that the revealed-preference argument that users of a product derive positive welfare from it fails to apply in the presence of negative externalities to non-users. Our findings also suggest that college students are sophisticated about how others’ social media consumption affects their own valuation.

Network Effects. Finally, we present our estimates of network effects, as defined in Section 2. We first present results on the fraction of respondents whose valuation removing network (Step 2) is lower than their valuation keeping network (Step 1). Fifty-eight percent of TikTok users have lower valuations in Step 2, 36% have the same valuation, and 6% have a higher valuation. Similarly, 42% of Instagram users have lower valuations in Step 2, 52% have the same valuation, and 6% have a higher valuation. Moreover, and strikingly, only approximately 70% of users on both platforms derive positive welfare from the product in Step 2, compared to approximately 90% in Step 1. Turning to average effects, Figure 4 uncovers a significant drop in average valuations between Step 1 and Step 2 of 29% ($p < 0.01$) and 21% ($p < 0.01$) for TikTok and Instagram users, respectively. Taken together, these results indicate that network effects are positive and quantitatively important, consistent with canonical theoretical frameworks (Rohlfs, 1974; Katz and Shapiro, 1985).

Our estimates also reveal that participants’ average utility from using these platforms is positive when they are among the few ones in their college using it, but negative in the status quo case where the rest of their school uses it as well. This pattern suggests that there are negative consumption externalities conditional on use ($u(1, X)$ is decreasing in X) but that non-user consumption externalities are even larger ($u(0, X)$ is decreasing in X faster than $u(1, X)$).

These results could be partly driven by, e.g., our participants having a preference for the status of being one of the few ones in their school with access to the social media world, consistent with important work on preferences for exclusion (Imas and Madarász, 2022). In this setting, aggregate use is not really zero (since the rest of the world keeps using the platform) and people might enjoy being the “gatekeepers” with access to the

new trends when part of their network is excluded from the platform. These results could also be driven by a preference to use social media for broader informative purposes paired with a relative distaste for direct social interactions on the platform (with even stronger FOMO or negative externalities when not using social media). These results suggest that the product market trap could be difficult to resolve by only introducing a coordination device that moves to a lower use equilibrium. Instead, it could result from a dominant-strategy in a Prisoner’s Dilemma where people are initially interested in using social media but become worse off (both as users and non-users) as overall use increases. It is important to note that a full characterization of a dynamic network formation equilibrium is beyond the scope of this paper and would depend on (i) heterogeneity in how user and non-user utility functions depend on X and (ii) the precise notion of equilibrium.

4.2 Correlates of Consumer Surplus

Table A2 examines heterogeneity in our different surplus measures along several demographics and displays regression coefficients from multivariate regressions. There are no significant correlations between gender and any of the surplus measures for both TikTok and Instagram. Individual valuations of TikTok, both with and without network, slightly increase with age, although the correlation is only marginally significant; whereas for Instagram there are no significant correlations with any of the surplus measures. As one would expect, the frequency in platform usage is positively and significantly correlated with individual welfare measures for TikTok. Indeed, using the platform daily, as opposed to less frequently, is associated with a \$23 and \$14 increase in respondent’s valuation for TikTok and Instagram, respectively. Coefficients of daily platform usage are lower and more noisily measured for the product market valuation, compared to the individual measures.


4.3 Robustness

4.3.1 Measurement Differences Across Steps


Our *Valuation Keeping Network* does not allow for negative welfare, as it would require people being willing to pay for the deactivation of their account individually. As a result, there is naturally an asymmetry in measurement between *Valuation Keeping Network*, where those requiring no payment for deactivation are coded as having a valuation of -\$10,

Figure 3: Distribution of Consumer Surplus Across Welfare Measures

(a) TikTok

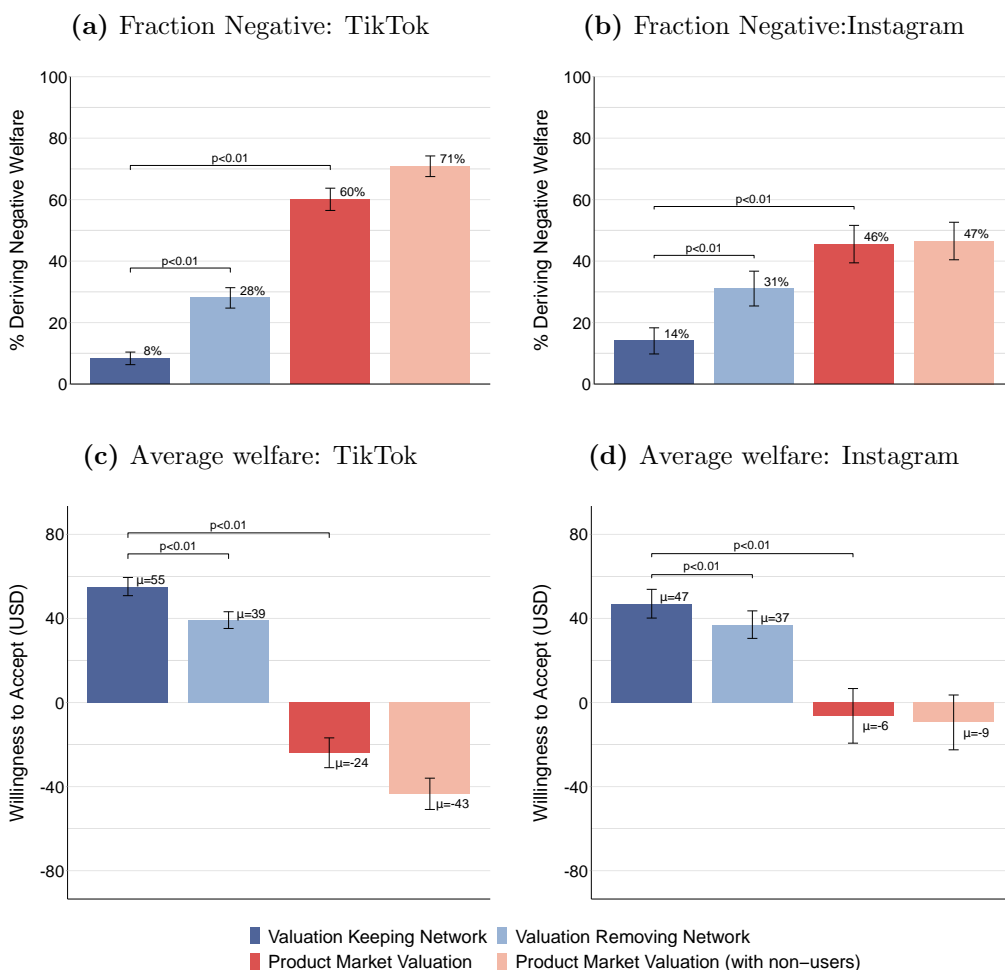


(b) Instagram



Notes: Figure 3 presents the probability density function of valuations for the different welfare measures. Panel (a) presents the results for TikTok and Panel (b) present the results for Instagram. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included.

Figure 4: Consumer Surplus across Welfare Measures



Notes: Panels (a) and (b) of Figure 4 present the percentage of respondents with negative product valuations across our different welfare measures. Panels (c) and (d) present averages. Panels (a) and (c) present the results for TikTok and Panels (b) and (d) present the results for Instagram. The first three bars in each panel represent valuations exclusively for active users. The dark blue bar denotes *Valuation Keeping Network*; the light blue bar denotes *Valuation Removing Network*; the red bar denotes *Product Market Valuation* for users. The pink bar represents the average *Product Market Valuation* of active users and non-users. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Error bars represent 95% confidence intervals.

and the *Product Market Valuation*, where negative values of up to -\$210 are possible.²⁹

To provide a very conservative way to examine whether this asymmetry in measurement can explain the sharp differences in valuation across *Valuation Keeping Network* and *Product Market Valuation*, we conduct a simple bounding exercise. In this exercise, we assume that participants that require no payment for the deactivation of their account in *Valuation Keeping Network* have a valuation of -\$210. Reassuringly, even under this very conservative bounding exercise, *Valuation Keeping Network* remains positive and large at \$38 for TikTok ($p < 0.01$) and \$19 for Instagram ($p < 0.01$), respectively.

4.3.2 Measurement Error

We designed our survey to reduce the importance of measurement error. First, we measure respondents' agreement with their valuations and give them the opportunity to revise their valuations. Second, we provide our respondents with a series of binary questions, which are commonly perceived to be easier to understand than questions directly eliciting respondents' reservation price on a continuous scale.

Despite these design choices, one concern with our elicitation is that measurement error could affect valuations in Step 1 and Step 3 differentially as these valuations are measured on different scales. Valuations in Step 1 (and Step 2) are elicited on a scale from -\$10 to \$210, while valuations in Step 3 are measured on a scale from -\$210 to \$210. Noise might therefore upwardly bias estimates for Step 1 if the true distribution is close to zero, but no such upwards bias will occur in Step 3.

To gauge the importance of noise, we examine to what extent our respondents' final valuations in Steps 1 and 2 are affected by the initial offers they receive.³⁰ Following the approach in Luttmer and Samwick (2018), which we explain in more detail in Appendix C.1, we regress final valuations from Step 1 on individuals' initial offers, which were randomized. As Table A5 shows, the initial offers are not significantly related to the final valuations for both TikTok ($p = 0.33$) and Instagram ($p = 0.55$), respectively. Moreover, initial offers only explain 0.3% and 0.1% of the variation in final valuations for TikTok and Instagram, respectively. In Step 2, we did not fully randomize the set of offers, but we can leverage that we randomized whether respondents receive the upper or lower bound of the Step 1

²⁹Reassuringly, only a small fraction of respondents (8% and 14% for TikTok and Instagram, respectively) require no payment for the deactivation of their account in *Valuation Keeping Network*.

³⁰We do not conduct these exercises for Step 3, since all participants started with an initial offer of zero. Additionally, we did not record data on the subsequent randomized offer.

valuation as the initial offer, for some individuals. Our analysis suggests that the initial offer significantly increases valuations by \$9 for TikTok ($p = 0.02$), and by an insignificant \$3 for Instagram ($p = 0.48$). In other words, we find evidence for measurement error in the case of TikTok.

We then adjust our estimates for measurement error following the approach in Luttmer and Samwick (2018). This approach assumes that the measured valuation is a weighted average of the respondent’s true underlying valuation and the initial offer they receive. Table A6 reports these adjustments. In the case of Step 1, both the average WTA and the proportion of individuals with a negative valuation remain stable after measurement error correction, across both platforms. For Step 2, the WTA in the case of Instagram also remains virtually unchanged after the correction, but the WTA for TikTok is revised downwards. This evidence makes it unlikely that measurement error generated by our elicitation drives the differences in valuations that we observe between Step 1 and Step 3. If anything, the presence of measurement error might have resulted in an underestimation of network effects in the case of TikTok, since valuation removing network might be overestimated.

4.3.3 Regret

To ensure data quality, we ask respondents whether they agree with a statement about what their choices mean in terms of their preferences over social media accounts for each of the four steps. For example, in the case of the practice good, a respondent with an implied valuation of between \$X1 and \$X2 is asked whether they agree with the statement that “According to your answers to the previous questions, you would require a payment worth between \$X1 and \$X2 to deactivate your Uber account for four weeks.” If respondents do not agree with this statement, they are asked to complete the multiple price list questions one more time. Prior to the redirection, we inform participants that this will be their last chance to modify their answers. Our main sample is restricted to respondents who do not regret their final answers in any of the steps, but we discuss below that results still hold when we also include those that regret their final choice.

Overall, we find that 33% of respondents regret their choices once and a smaller fraction of 5% regret their choices twice. Figure A4 illustrates that this pattern holds for each step: after being redirected to the MPL questions, fewer participants disagree with their elicited WTA. The extent of regret fluctuates across steps. The percent of respondents regretting their choices is relatively high at the practice section with 22% disagreeing with their

elicited WTA initially and 5% after completing the MPL questions a second time. In the subsequent steps, the fraction of respondents regretting their final answers fluctuates around 3%. These patterns suggest that comprehension and data quality is high and that the practice questions helped improve comprehension.

4.3.4 Perceived Stakes and Credibility

The key design challenge for our paper concerns the large-scale deactivation study. One concern with our empirical design is that respondents may not find it likely that we will manage to recruit two-thirds of university students at their university, the condition for the large-scale deactivation study. To examine whether people perceived as credible that the large-scale deactivation study would take place, we asked respondents in our Instagram and Maps experiments about the percent chance that the researchers will recruit more than two-thirds of the students at their university. On average, participants perceive this likelihood to be quite substantial, at 45% for Instagram.³¹ This in turn implies that respondents perceived the likelihood that the large-scale deactivation study would take place as substantial.³² Moreover, we examine how this perceived probability is correlated with our estimated welfare effects. Panel (e) of Figure A5 illustrates that individuals who deem the large-scale deactivation study as more probable do not have a less negative product market valuation. This heterogeneity suggests that our design is conservative and likely underestimates the extent of negative welfare. More broadly, given that even in the case of the large-scale deactivation study, not all users would deactivate their account, our study plausibly identifies lower bounds for the size of negative product market surplus and for the size of network effects. Future work should measure people’s preferences over deactivation studies in settings with potentially higher credibility. For example, it seems conceivable to achieve higher credibility when small groups of individuals coordinate on deactivating social media together.

4.3.5 Hypothetical Welfare Measures

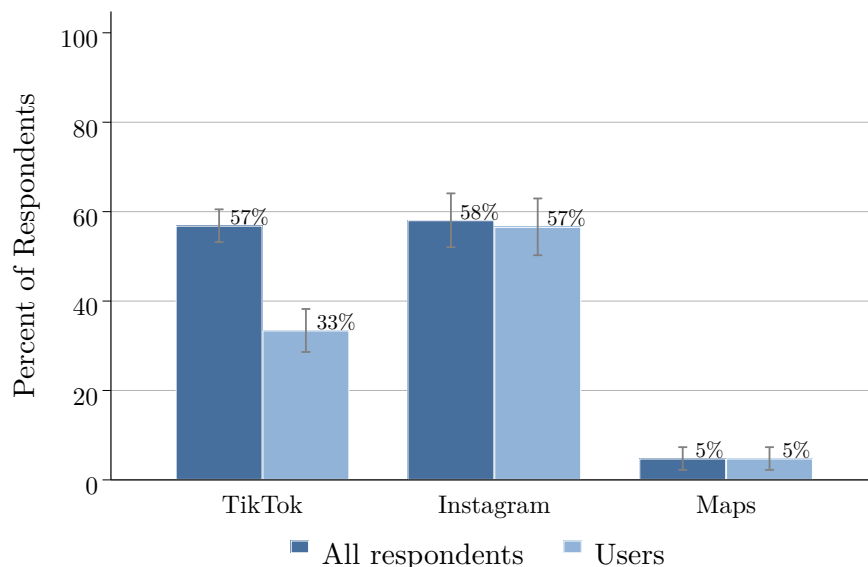
Live in a World without. After the price elicitation, we present our respondents with a series of hypothetical qualitative questions. To assess the boundary conditions of our results (i.e., extrapolating to a hypothetical case where every user in the world stops using

³¹The perceived likelihood for Maps is reassuringly similar at 46%.

³²It is conceivable that social desirability bias or experimenter demand effects may somewhat distort respondents’ responses to this question.

their social media), we ask respondents whether they would prefer to live in a world with or without the social media platform. As Figure 5 shows, 57% and 58% of respondents (including users and non-users) prefer to live in a world without TikTok and Instagram, respectively. Even among users, 33% and 57% prefer to live in a world without TikTok and Instagram, respectively.³³ Panels (a) and (b) of Appendix Figure A5 validate these hypothetical survey questions with the incentivized measure of product market surplus. The figures illustrate that the hypothetical question is strongly correlated with the incentivized measure for both TikTok ($p < 0.01$) and Instagram ($p < 0.01$).

Figure 5: Percentage of Respondents that Prefer to Live in a World without the Platform



Notes: Figure 5 displays the percent of the respondents that stated they would prefer to live in a world without the platform for TikTok, Instagram and Maps separately. The dark blue bar represents the fraction among all respondents and the light blue bar represents the fraction among active users of the respective platform. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Error bars represent 95% confidence intervals.

Preference Rankings. To understand the preferences of university students regarding social media platform usage among their peers, we ask them to rank three hypothetical scenarios: (i) they deactivate the platform and every other student at their university

³³As suggestive evidence against experimenter demand effects, Figure 5 also shows that only less than 5% of respondents prefer to live in a world without Maps.

keeps using it, (ii) every student at their university, including themselves, deactivates the platform, and (iii) no one deactivates the platform.

The results based on these rankings support our main findings. The most preferred scenario among our respondents is the scenario where every student at their universities, including themselves, deactivates their social media account, respectively (Figure A6). Among TikTok users, 40% prefer this option, while 49% of Instagram users prefer this option. In contrast, the least preferred scenario is where no one deactivates their account, with 50% and 52% of respondents citing this option as their least preferred one for TikTok and Instagram, respectively.³⁴

Panels (c) and (d) of Appendix Figure A5 validate these hypothetical measures with the incentivized measure of product market surplus. This figure illustrates that the hypothetical question is strongly correlated with the incentivized measure. Indeed, while respondents who preferred deactivation for everyone have highly negative product market valuation for both Instagram and TikTok, respondents for whom deactivation for everyone was the least preferred option have positive product market valuation. The differences in product market surplus across these survey measures are highly significant for both Instagram and TikTok ($p < 0.01$).

4.3.6 Substitution Across Social Media Platforms

One concern is that people’s product market valuation of social media platforms is so low given their opportunity to substitute their social media consumption to another platform, based on an argument similar to the one in Farrell and Saloner (1985). Specifically, even in the absence of non-consumer surplus, if there are two technologies, where an “alternative” technology is superior to the predominant technology, individuals’ welfare could be improved if they all stopped using the inferior technology. However, several pieces of evidence can help rule out this story and shed light on how substitution affects our estimates.

First, structural estimates of diversion ratios suggest that the outside option (offline or other online activities) is the most important substitution channel for social media, including Instagram and TikTok (Aridor, 2022). This pattern suggests that individuals do

³⁴We also find that “Only I deactivate” is the most preferred option for 33% and 29% of TikTok and Instagram users, respectively. This is somewhat higher than the fraction of users who would accept to deactivate their accounts without compensation. This difference could arise from differences in the elicitations, such as the lack of incentives and the time horizon of the deactivation.

not all substitute towards a “better” platform. Second, as has been documented in the literature (Allcott et al., 2022; Aridor, 2022), the vast majority of users in our data are multi-homers, with very few users having only a TikTok (n=6 in the first survey) or an Instagram account (n=17 in the second survey). This adoption pattern makes it unlikely that users are trapped in Instagram or TikTok because they cannot switch to a better alternative. Third, our estimates show that both respondents who multi-home and those that single-home have a negative product market valuation, which alleviates concerns that the negative product market surplus is driven by cross-platform substitution.³⁵

To further understand the mechanisms underlying participants’ valuations, it would be helpful to measure their beliefs about substitution patterns, i.e., what they thought would happen following the temporary shutdown of these platforms. For example, future work could examine whether students think that the large-scale deactivation would entail more offline interactions on campus or better study outcomes in those four weeks and subsequently.

4.3.7 Other Robustness Checks

Distributional Assumptions. As a robustness check against potential censoring in valuations, we assume a triangular distribution for those values that lie in these ranges, following the methodology of Allcott and Kessler (2019). Estimates constructed this way give more weight to the upper and lower bounds in the elicitation and thus allow us to gauge the sensitivity of our results to extreme valuations at the tails. Given that we see more mass at the lower end of the distribution (see Figure 3) this means that the welfare estimates based on the triangular distribution are more negative than our main estimates (see Appendix Figure A11). This, in turn, suggests that censoring, if anything, makes us overestimate the welfare effects of social media. See Appendix C.3 for details on the triangular distribution.

Robustness to Sample Restrictions. In our main analysis, we reported results for respondents who passed all attention checks and did not regret any of their final choices. Appendix Figure A8 confirms our results for the full sample without those exclusions;

³⁵Among individuals that only have a TikTok account (n=6) and only an Instagram account (n=17), the product market surplus is even more negative at -\$43 and -\$44, respectively. Naturally, these estimates are noisy given that most respondents have both a TikTok and an Instagram account. Among users that multi-home, the estimates are -\$24 for TikTok and -\$3 for Instagram.

Appendix Figure A9 confirms our findings with a sample that includes inattentive respondents and excludes respondents regretting their final choice. Finally, Appendix Figure A10 demonstrates the robustness of our findings to including respondents regretting their final choice and only excluding inattentive respondents.

4.4 Mechanisms

4.4.1 Repugnance Towards Digital Products

One possible mechanism that might explain the drop in welfare from the individual to the product market surplus could be repugnance towards digital products. To test this conjecture, we run a deactivation study experiment with a digital good that plausibly does not cause strong externalities to non-users; Maps. These applications likely have more muted externalities on non-users as they do not create social costs of exclusion and are less likely to impact relative social standing. The instructions are virtually identical to our main experiment, except for the product name and the way the deactivation is monitored. Deactivation of Maps is only monitored through screenshots.

Figure A7 shows that both the individual consumer surplus and the product market surplus are positive and significantly different from zero in the case of Maps.³⁶ The product market surplus is significantly lower than the individual consumer surplus ($p < 0.01$), which might result from various motives. First, respondents may dislike “Big Tech” companies and their associated market power and therefore prefer a ban of products that underlie the market power of big tech companies. Second, respondents may feel repugnance towards digital products, such as mobile phones and modern technologies. Third, respondents may have a distaste for others using their phone.

To formally test whether the drop in welfare between individual consumer surplus and product market surplus is larger for Instagram compared to Maps, we conduct a simple difference-in-differences exercise, where we compute the change between valuation keeping network and product market valuation between Instagram and Maps. Table A4 shows that the coefficient on the interaction term (of an indicator of Instagram and the product market valuation) is of substantial magnitude and significant ($p < 0.01$). This corroborates that negative externalities to non-users are larger on Instagram than they are for Maps.

³⁶These findings are also qualitatively in line with the hypothetical ranking question, where we find that the most preferred scenario is for no one to quit maps (46%), while only 25% of respondents have a preference for everyone to quit maps (see Appendix Figure A7d).

The evidence on the positive product market valuation of Maps and the significant difference-in-differences estimate also alleviates concerns that the question wording we used in the product market valuation mechanically induces negative welfare estimates, i.e., respondents providing positive willingness to pay to ban others using the product. Naturally, given the many differences between Maps apps and social media platforms, this evidence remains limited in disentangling mechanisms underlying our main findings.

4.4.2 Motives Behind Consumption

Social Media Platforms. To provide additional evidence on mechanisms, we asked active users of the platform who said that they prefer to live in a world without the platform an open-ended question to better understand the motives behind their usage.³⁷ We asked them the following question:

You mentioned you would prefer to live in a world without [platform]. Why do you still use it?

To quantitatively analyze the data, we devised a simple hand-coding scheme, which comprises five categories.³⁸ “FOMO” responses usually mention feeling left out (“I feel like if I stop using it, I will be completely out of the loop”).³⁹ “Entertainment” responses talk about the high entertainment value of the platform (“It’s a very good source of entertainment and it’s always something to do when bored”). “Addiction” responses mention self-control problems and addiction (“It’s very addicting and I cannot stop”). “Information” responses indicate receiving useful information (“I follow pages that keep me up to date with the largest news”). Finally, “Productivity/Convenience” responses mention using the platform for productive use or convenience (“I still use Instagram for business purposes”). Appendix Table A7 provides an overview of the hand-coding scheme and provides further example responses.⁴⁰

³⁷Open-ended questions are increasingly used to better understand the hidden motives behind people’s choices, see, e.g., Bursztyn et al. (2022, 2023b). These questions avoid priming respondents on particular motivations and better capture what naturally comes to mind compared to more structured questions (Haaland et al., 2023).

³⁸A given response can fall into multiple categories.

³⁹Social media platforms like TikTok and Instagram may cultivate FOMO among users through specific platform features and content dynamics. TikTok’s video format and evolving trends may create social pressure to remain continually informed. Conversely, Instagram’s feature of shared content that is only available for a limited time on the platform may foster a fear of being left out.

⁴⁰We validate our hand-coded open-ended data with data coded by a large language model (LLM). We

Figure 6 illustrates the quantitative distribution of the hand-coded data for TikTok, Instagram and Maps, respectively. It reveals that the fear of missing out is the most prevalent motive both for Instagram (76%) and TikTok (40%). Moreover, entertainment motives also play an important role in driving people’s social media consumption (31% of TikTok and 21% of Instagram users), consistent with evidence on people’s news preferences (Bursztyn et al., 2023a). Consistent with prior evidence (Allcott et al., 2022), addiction is an important reason for TikTok (33%), though somewhat less important for Instagram (11%). Finally, only a very small fraction of users (6% and 8% on TikTok and Instagram, respectively) cite productivity/convenience as a reason for using the platform. While this evidence suggests that FOMO is an important mechanism underlying our findings, it does not conclusively show that FOMO is the main mechanism underlying the wedge between the traditional welfare measure and our measure of product market surplus. One important limitation of the open-ended data is that its analysis necessarily involves subjective judgment calls in the coding manual and from the coders.

Maps. We also conducted a similar coding procedure for the Maps experiment. Figure 6 reveals a very different distribution of motives: 69% of respondents mention productivity reasons, 23% mention information and only 8% mention the fear of missing out (“[...] still use navigation maps because it is what everyone uses”).


4.4.3 Direct Evidence on Mechanisms Behind Externalities to Non-Users

Social Media Platforms. To provide direct evidence on the mechanisms behind externalities to non-users, we asked all of our respondents an open-ended question to describe the nature and motives behind their non-user utility. In particular, we asked them “How would you feel if you were the only one who deactivated [platform] and everyone else kept using it?”

Based on the open-ended responses, we devised a coding scheme to capture the most common topics. “FOMO” responses talk about the fear of missing out (“I would definitely feel a bit left out”). “Negative” responses express negative emotions without explicitly mentioning the fear of missing out (“[...] it would be a little unfair”). “Indifferent” responses indicate that they do not expect the deactivation to have strong effects on them

show that the LLM-based measure yields similar frequencies of the categories (See Panel (a) of Figure A12). Moreover, the LLM-based categories are highly correlated with the hand-coded measure (see Panel (a) of Table A9 for details).

Figure 6: Motives for Social Media Consumption Despite a Preference to Live in a World without It



Notes: Figure 6 presents the fraction of respondents mentioning different motives in their open-ended responses. Active users who said that they prefer to live in a world without the platform were asked the following open-ended question: *You mentioned you would prefer to live in a world without [platform]. Why do you still use it?* “FOMO” denotes responses mentioning the fear of missing out or related social concerns. “Entertainment” denotes responses mentioning the entertainment value of the platform. “Addiction” denotes responses indicating the addictive nature of the platform and self-control problems. “Information” denotes responses mentioning informational purposes such as following the news or keeping abreast of college events. “Productivity” denotes responses mentioning productivity benefits, such as using the platform for business purposes. The categorization of the open-ended answers is not mutually exclusive. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Nonsensical responses were dropped from the analysis. The underlying sample sizes are 121 for TikTok, 131 for Instagram, and 13 for Maps. Error bars represent 95% confidence intervals.

(“That wouldn’t be a big deal”). “Beneficial” responses mention the benefits of not using the respective platforms (“I would be able to focus on more important things”).⁴¹

Figure 7 illustrates the results. Panel (a) shows results for respondents who prefer to live in a world without and Panel (b) shows results for respondents who prefer to live in a world with the platforms.⁴² Among the TikTok users who would prefer to live in a world without TikTok, 35% express FOMO, 33% are indifferent, 7% have generic negative feelings and 17% see it as beneficial. Among the Instagram users who would prefer to live in a world without Instagram, 38% express FOMO, 35% are indifferent, 8% have generic negative feelings, and a 18% see it as beneficial. Among Instagram users who would prefer to live in a world with Instagram, 41% express FOMO, 25% are indifferent, 13% have generic negative feelings, and 15% see it as beneficial. These data clearly highlight that the fear of missing out is a prevalent motive behind non-user utility.

Maps. We hypothesized that the nature of externalities for navigation and maps smartphone applications would be different from social media platforms. Figure 7 provides direct evidence for this conjecture. Among respondents preferring to live in a world without Maps, FOMO is mentioned fairly infrequently (8%), while “Indifferent” responses and “Negative” responses are more prevalent at 31% each, and 15% of responses fall into the “Beneficial” category.⁴³ Among respondents preferring to live in a world with Maps, the patterns are similar: 14% mention FOMO, 27% are “Indifferent” responses, 39% are “Negative” responses, and 8% are “Beneficial” responses.

4.4.4 Other-regarding Preferences

Our design in Step 3 tries to hold constant other-regarding preferences by telling respondents that other participating students would receive *just enough* money to deactivate their accounts.⁴⁴ As such, we tried to make it clear to respondents that there would be no surplus left for other respondents as we would compensate them at their indifference

⁴¹We again show that the hand-coded measure is highly correlated with analogous data annotated by a large language model (see Panel (b) of Table A9 and Panel (b) of Figure A12).


⁴²We did not collect the open-ended data for respondents who preferred to live in a world with TikTok.

⁴³An example FOMO response is: “would feel a bit isolated, maybe excluded from certain conversations involving travel plans, etc.”

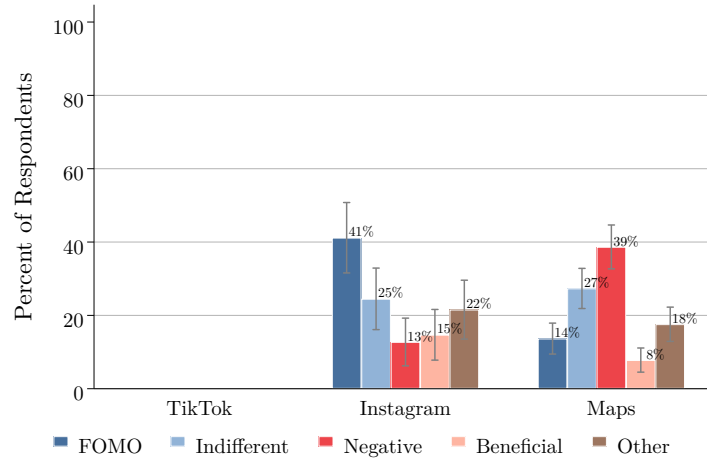
⁴⁴The compensation that the other participating students would require to deactivate their accounts if the respondent chooses (and is selected) to deactivate everyone in Step 3 is their valuation removing network, as measured in Step 2. This is one of the main reasons why we elicited this object, besides its usefulness to measure network effects.

Figure 7: Evidence on Mechanisms Behind Externalities to Non-Users

(a) Active Users that Prefer to Live in a World without Platform



(b) Active Users that Prefer to Live in a World with Platform



Notes: Figure 7 presents the fraction of respondents expressing different emotions in their open-ended responses. Panel (a) shows results for respondents who prefer to live in a world without the respective platform, while Panel (b) shows results for those who prefer to live in a world with the respective platform. Active users were asked the following open-ended question: *How would you feel if you were the only one who deactivated [platform] and everyone else kept using it?* Data for TikTok is missing for Panel (b) as this question was only directed to TikTok users who stated they would rather live in a world without TikTok. “FOMO” denotes responses mentioning the fear of missing out or related social concerns. “Indifferent” denotes responses expressing they would not be particularly affected. “Negative” denotes responses expressing negative emotions, whereas “Beneficial” denotes responses where respondents mention a potential benefit of deactivation. “Other” denotes a diverse set of responses that mention different motives. The categorization of the open-ended answers is not mutually exclusive. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Nonsensical responses were dropped from the analysis. Responses indicating indifference conditional on payment/contribution to research were placed in the “Other” category. The underlying sample sizes are 121 for TikTok, 233 for Instagram, and 269 for Maps. Error bars represent 95% confidence intervals.

point. This should be quite salient given that we asked respondents who regretted their initial valuation to retake the questions. Of course, other-regarding preferences could still be a mechanism underlying the empirical patterns we document. For example, people may have paternalistic preferences and could believe that other respondents have self-control problems and therefore prefer the large-scale deactivation as they think this will be beneficial to other participants' mental health. The evidence presented in Section 4.4.3 suggests that paternalistic concerns are at least not top-of-mind when respondents are asked how they would feel if they were the only one who deactivated [platform] and everyone else kept using it. Hand-coding the open-ended responses reveals that less than 1% of the answers expressed paternalistic sentiments ("Self-deactivation can benefit other people") for TikTok and Instagram.

4.5 Other Applications: Luxury Goods and Product Vintages

The previous evidence is specific to social media. To probe the external validity of our findings, we provide suggestive survey evidence from contexts in which positional concerns are plausible drivers of externalities to non-users: luxury goods and products with different vintages. We use pre-registered surveys with a sample of 500 respondents from the US conducted on Prolific. Appendix Section E provides additional details on the sample and design.

Among respondents that owned any luxury brands, 44% preferred to live in a world without those brands. Among respondents not owning any of these brands, the fraction preferring to live in a world without them is higher, at 69%. We next examine preferences regarding the frequency of product variations. Among iPhone owners, a striking 91% of respondents would prefer Apple to release the iPhone every other year rather than every year. Among respondents not owning the iPhone, 94% prefer Apple to release the iPhone every other year rather than every year.

Overall, the findings from this survey, though just suggestive, point to the possibility that negative non-user utility may not be specific to the case of social media, and might also extend to luxury consumption and certain high-end technology products. We believe that examining the importance of product market traps in other contexts is a fruitful avenue for future research.

5 Conclusion

In the conventional assessment of consumer welfare, the emphasis is predominantly on individual-level evaluations, holding aggregate consumption fixed. However, such measures do not accurately reflect welfare in settings with externalities to non-users. We introduce a new method to gauge welfare in these contexts, which we apply to widely used social media platforms through incentivized trials involving college students. While traditional measures of individual consumer surplus suggest positive welfare, the *Product Market Valuation* that accounts for externalities to non-users tells a different story: it reveals negative welfare, with a notable portion of users experiencing a disutility from the platform.

Our conceptual framework defines *product market traps*, a phenomenon whereby consumers prefer the product not to exist, but cannot avoid using it. Intriguingly, such product market traps can arise even with fully rational expectations and without any behavioral frictions. In the context of social media, our empirical evidence highlights the existence of a *social media trap* for a large fraction of consumers, who derive large individual consumer surplus but, simultaneously, experience negative welfare from the product. These results could help reconcile the seemingly contradictory findings in the social media literature of a large consumer surplus coexisting with negative effects on well-being. More generally, these findings challenge the standard revealed-preference argument that the mere existence of a product implies that its consumers derive positive welfare. The presence of product market traps underscores the need for more research on whether companies introduce features that exacerbate non-user utility and diminish consumer welfare, rather than enhance it, increasing people’s need for a product without increasing the utility it delivers to them.

Our framework also highlights a few important levers for policymakers: first, policymakers should regulate markets to counteract producers’ incentives to use technologies that decrease non-user utility. Second, given that larger networks decrease non-user utility, optimal anti-trust policy may involve reducing the size of networks.

While our evidence shows that a large fraction of active users of social media platforms derive negative utility from the platform’s existence, it is an open question whether people would be willing to jointly deactivate their social media accounts above and beyond the temporary four week deactivation. Future work should examine whether coordination devices could enable users to actually deactivate their social media accounts persistently or whether they would go back to using social media platforms after an initial period of deactivation.

References

- Allcott, Hunt and Judd B Kessler**, “The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons,” *American Economic Journal: Applied Economics*, 2019, 11 (1), 236–276.
- , **Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow**, “The Welfare Effects of Social Media,” *American Economic Review*, 2020, 110 (3), 629–676.
- , **Matthew Gentzkow, and Lena Song**, “Digital Addiction,” *American Economic Review*, 2022, 112 (7), 2424–2463.
- American Trends Panel of the Pew Research Center**, “Wave 93,” 2021. Accessed: 2023-10-02.
- Aridor, Guy**, “Drivers of Digital Attention: Evidence from a Social Media Experiment,” 2022.
- Becker, Gordon M, Morris H DeGroot, and Jacob Marschak**, “Measuring Utility by a Single-Response Sequential Method,” *Behavioral Science*, 1964, 9 (3), 226–232.
- Benzell, Seth and Avinash Collis**, “Regulating Digital Platform Monopolies: The Case of Facebook,” *Available at SSRN 3619535*, 2022.
- Bhattacharya, Debopam, Pascaline Dupas, and Kanaya Shin**, “Demand and Welfare Analysis in Discrete Choice Models with Social Interactions,” *The Review of Economic Studies*, 2023.
- Bottan, Nicolas L and Ricardo Perez-Truglia**, “Choosing Your Pond: Location Choices and Relative Income,” *Review of Economics and Statistics*, 2022, 104 (5), 1010–1027.
- Braghieri, Luca, Ro’ee Levy, and Alexey Makarin**, “Social Media and Mental Health,” *American Economic Review*, 2022, 112 (11), 3660–3693.
- Brynjolfsson, Erik and JooHee Oh**, “The Attention Economy: Measuring the Value of Free Digital Services on the internet,” *Information Systems Research*, 2023.
- , **Avinash Collis, and Felix Eggers**, “Using Massive Online Choice Experiments to Measure Changes in Well-Being,” *Proceedings of the National Academy of Sciences*, 2019, 116 (15), 7250–7255.
- , – , **Asad Liaqat, Daley Kutzman, Haritz Garro, Daniel Deisenroth, Nils Wernerfelt, and Jae-Joon Lee**, “The Digital Welfare of Nations: New Measures of Welfare Gains and Inequality,” Working Paper 31670, National Bureau of Economic Research 2023.

- Bursztyn, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott**, “Opinions as Facts,” *The Review of Economic Studies*, 2023, 90 (4), 1832–1864.
- , **Bruno Ferman, Stefano Fiorin, Martin Kanz, and Gautam Rao**, “Status Goods: Experimental Evidence from Platinum Credit Cards,” *The Quarterly Journal of Economics*, 2018, 133 (3), 1561–1595.
- , **Georgy Egorov, Ingar Haaland, Aakaash Rao, and Christopher Roth**, “Scapegoating during Crises,” *AEA Papers and Proceedings*, 2022, 112, 151–55.
- , – , – , – , and – , “Justifying Dissent,” *The Quarterly Journal of Economics*, 2023, 138 (3), 1403–1451.
- Cheyre, Cristobal and Alessandro Acquisti**, “Online Intermediation in Legacy Industries: Evidence from the Adoption of Restaurant Reservation Platforms,” *Available at SSRN 4721874*, 2024.
- Clark, Andrew E and Andrew J Oswald**, “Satisfaction and Comparison Income,” *Journal of Public Economics*, 1996, 61 (3), 359–381.
- Cullen, Z. and R. Perez-Truglia**, “How Much Does Your Boss Make? The Effects of Salary Comparisons,” *Journal of Political Economy*, 2022, 130 (3), 766–822.
- Dizon-Ross, Rebecca and Seema Jayachandran**, “Improving Willingness-to-Pay Elicitation by Including a Benchmark Good,” *AEA Papers and Proceedings*, 2022, 112, 551–555.
- Enke, Benjamin, Uri Gneezy, Brian Hall, David Martin, Vadim Nelidov, Theo Offerman, and Jeroen Van De Ven**, “Cognitive Biases: Mistakes or Missing Stakes?,” *Review of Economics and Statistics*, 2023, 105 (4), 818–832.
- Eyal, Peer, Rothschild David, Gordon Andrew, Evernden Zak, and Damer Ekaterina**, “Data Quality of Platforms and Panels for Online Behavioral Research,” *Behavior Research Methods*, 2021, pp. 1–20.
- Faelens, Lien, Kristof Hoorelbeke, Ruben Cambier, Jill van Put, Eowyn Van de Putte, Rudi De Raedt, and Ernst HW Koster**, “The Relationship Between Instagram Use and Indicators of Mental health: A Systematic review,” *Computers in Human Behavior Reports*, 2021, 4, 100121.
- Farrell, Joseph and Garth Saloner**, “Standardization, Compatibility, and Innovation,” *the RAND Journal of Economics*, 1985, pp. 70–83.
- and **Paul Klempner**, “Coordination and Lock-In: Competition with Switching Costs and Network Effects,” *Handbook of Industrial Organization*, 2007, 3, 1967–2072.

- Frank, Robert H**, *Choosing the Right Pond: Human Behavior and the Quest for Status.*, Oxford University Press, 1985.
- , *Luxury Fever: Money and Happiness in an Era of Excess*, Princeton University Press, 2000.
- , “Positional Externalities Cause Large and Preventable Welfare Losses,” *American Economic Review*, 2005, *95* (2), 137–141.
- , “The Darwin Economy,” in “The Darwin Economy,” Princeton University Press, 2012.
- Gupta, Mayank and Aditya Sharma**, “Fear of Missing Out: A Brief Overview of Origin, Theoretical Underpinnings and Relationship with Mental Health,” *World Journal of Clinical Cases*, 2021, *9* (19), 4881.
- Haaland, Ingar, Christopher Roth, and Johannes Wohlfart**, “Designing Information Provision Experiments,” *Journal of Economic Literature*, 2023, *61* (1), 3–40.
- Imas, Alex and Kristóf Madarász**, “Superiority-Seeking and the Preference for Exclusion,” *Review of Economic Studies*, 2022.
- Katz, Michael L and Carl Shapiro**, “Network Externalities, Competition, and Compatibility,” *The American Economic Review*, 1985, *75* (3), 424–440.
- Kemp, Simon**, “Digital 2022: Global Overview Report,” Available at <https://datareportal.com/reports/digital-2022-global-overview-report> 2022.
- Luttmer, Erzo FP**, “Neighbors as Negatives: Relative earnings and Well-Being,” *The Quarterly Journal of Economics*, 2005, *120* (3), 963–1002.
- and **Andrew A Samwick**, “The Welfare Cost of Perceived Policy Uncertainty: Evidence from Social Security,” *American Economic Review*, 2018, *108* (2), 275–307.
- Mosquera, Roberto, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie**, “The Economic Effects of Facebook,” *Experimental Economics*, 2020, *23*, 575–602.
- Perez-Truglia, Ricardo**, “The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment,” *American Economic Review*, 2020, *110* (4), 1019–1054.
- Pesendorfer, Wolfgang**, “Design Innovation and Fashion Cycles,” *The American Economic Review*, 1995, pp. 771–792.
- Pew Research Center**, “Social Media Use in 2021,” 2021. Accessed: 2023-08-22.
- Rochet, Jean-Charles and Jean Tirole**, “Platform Competition in Two-Sided Markets,” *Journal of the European Economic Association*, 2003, *1* (4), 990–1029.

Rohlf, Jeffrey, “A Theory of Interdependent Demand for a Communications Service,” *The Bell Journal of Economics and Management Science*, 1974, pp. 16–37.

Roth, Alvin E., “Repugnance as a Constraint on Markets,” *Journal of Economic Perspectives*, September 2007, 21 (3), 37–58.

Sullivan, Michael, “Price Controls in a Multi-Sided Market,” Technical Report, Working paper 2022.


Online Appendix: Not for publication

Our supplementary material is structured as follows. Section A provides additional details related to the conceptual framework. Section B includes additional tables and figures. Section C provides additional evidence on robustness. In Section D, we provide additional details on the open-ended data, including the coding scheme. Section F provides evidence including the pilot data from Instagram. Section E Finally, Appendix G presents the instructions for all experiments described in the paper.

A Additional framework details

Changes in Market Share. Our model specified above allows for comparative statics on how welfare changes for distinct groups (e.g., users who lose out from the product) as the observed market share X^* changes. Figure A1 provides an illustrative example of how an increase in the number of users can, simultaneously, 1) increase the marginal utility from using the product relative to not using it (positive network effects) and 2) decrease the utility of both using and not using the product. Our empirical application to social media elicits portions of these two curves using experimental methods, allowing us to assess key comparative statics related to how welfare changes with total equilibrium use.

Appendix Figure A1: Utility and Negative Consumption Spillovers from Product Use as a Function of Aggregate Use X



Notes: Figure A1 presents an example of individual utilities as a function of aggregate use X , as well as individual consumer surplus, product market surplus, and network effects. In this example there are negative consumption spillovers (to both users and non-users), negative product market surplus, positive individual consumer surplus, and positive network effects.

Incentive Compatibility. Table A1 helps illustrate the incentive compatibility of our empirical procedure. It presents, for each of the three steps in our survey, the payoffs that individual i gets when given an offer y to deactivate their social media account and the conditions under which i chooses to select “deactivate” or “not deactivate” in the decision screen. For simplicity and without loss of generality, we assume that participants face a single decision screen, as opposed to the multiple decision screens that they face in our procedure. Moreover, we assume that participants face some small hassle cost such that they choose to not deactivate if they are indifferent between deactivating and not

deactivating. Note that the table incorporates the option-value that participants have when choosing to deactivate or to not deactivate in the decision screen; they can always change their mind afterward, although doing so implies that they can lose their monetary compensation or cancel the deactivation experiment for everyone else (in the case of step 3 in our experiment). For example, a participant that receives positive offer $\$y$ to individually deactivate their account in step 1 will have a payoff of $\max\{u_i(0, X) + y, u_i(1, X)\}$ if they accept to deactivate. The reason is that they receive $u_i(0, X) + y$ if they follow through with deactivation but a payoff $u_i(1, X)$ if they change their mind and reactivate. Note that the table shows that there is no incentive compatibility for negative offers in the case of steps 1 and 2 (which estimate the valuation keeping and removing network, respectively). The reason is that the individual can always deactivate for free, so, when faced with the option to deactivate vs. not deactivate and receive $\$y$ (how one would implement in practice a negative offer to deactivate), it is a dominant strategy to choose to not deactivate (even if i were willing to pay to deactivate).

Appendix Table A1: Payoffs and Decisions for Different Offers to Deactivate

	Offer $y > 0$	Offer $y \leq 0$
<i>Valuation keeping network</i>		
Payoff of “deactivate”	$\max\{u_i(0, X) + y, u_i(1, X)\}$	$\max\{u_i(0, X), u_i(1, X)\}$
Payoff of “do not deactivate”	$\max\{u_i(0, X), u_i(1, X)\}$	$\max\{u_i(0, X), u_i(1, X)\} + y $
i clicks “deactivate” if	$y > u_i(1, X) - u_i(0, X)$	N/A
i clicks “do not deactivate” if	$y \leq u_i(1, X) - u_i(0, X)$	For all $y \leq 0$
<i>Valuation removing network</i>		
Payoff of “deactivate”	$\max\{u_i(0, 0) + y, u_i(1, 0)\}$	$\max\{u_i(0, 0), u_i(1, 0)\}$
Payoff of “do not deactivate”	$\max\{u_i(0, 0), u_i(1, 0)\}$	$\max\{u_i(0, 0), u_i(1, 0)\} + y $
i clicks “deactivate” if	$y > u_i(1, 0) - u_i(0, 0)$	N/A
i clicks “do not deactivate” if	$y \leq u_i(1, 0) - u_i(0, 0)$	For all $y \leq 0$
<i>Product market valuation</i>		
Payoff of “deactivate”	$\max\{u_i(0, 0) + y, u_i(1, X)\}$	$\max\{u_i(0, 0), u_i(1, X)\}$
Payoff of “do not deactivate”	$\max\{u_i(0, 0), u_i(1, X)\}$	$\max\{u_i(0, X), u_i(1, X)\} + y $
i clicks “deactivate” if	$y > \max\{u_i(1, X), u_i(0, X)\} - u_i(0, 0)$	
i clicks “do not deactivate” if	$y \leq \max\{u_i(1, X), u_i(0, X)\} - u_i(0, 0)$	

Notes: The table presents the payoffs that individual i gets when given an offer y to deactivate their social media, and the conditions under which i selects “deactivate” or “not deactivate” in the decision screen. “Payoff of ‘deactivate’” and “Payoff of ‘do not deactivate’” corresponds to the case when i selects the option to deactivate and not deactivate in the survey, respectively. Without loss of generality, we that i faces a single decision. For simplicity, we assume quasilinear preferences, a static framework, and perfect monitoring on our end (that is, that we can perfectly follow through when participants do not comply with deactivation). Moreover, we assume that i faces some small hassle costs such that they choose to not deactivate if they are indifferent between deactivating and not deactivating.

B Additional tables and figures

Appendix Table A2: Correlates of Consumer Surplus

	Valuation Keeping Network (1)	Valuation Removing Network (2)	Product Market Valuation (3)
Panel A: TikTok			
Age	3.13* (1.76)	2.94* (1.62)	1.10 (3.20)
Female	-6.09 (6.94)	-3.63 (6.56)	-5.94 (12.26)
Daily usage	23.21*** (5.99)	23.28*** (5.59)	13.58 (10.62)
Dep. var. mean	55.18	39.22	-23.91
Dep. var sd	58.68	53.89	96.02
Observations	371	371	371
Panel B: Instagram			
Age	-0.06 (2.27)	0.47 (2.24)	3.01 (4.49)
Female	-1.18 (8.39)	-9.60 (8.86)	-8.88 (17.70)
Daily usage	14.16* (7.41)	14.91** (7.55)	-1.82 (14.96)
Network size	0.09 (0.13)	0.12 (0.12)	0.19 (0.30)
Dep. var. mean	46.02	36.94	-5.61
Dep. var sd	55.20	53.56	105.63
Observations	196	196	196

Notes: The table presents coefficient estimates from OLS regressions. Panel A displays the results for TikTok and Panel B displays the results for Instagram. Columns 1-3 correspond to the elicitation steps 1-3 in our survey, respectively. The independent variables are age, dummy variables for identifying as female and self-reported daily platform usage, and self-reported fraction of college students who are mutual friends on Instagram, labeled *Network size* in the table. *Network size* is only available in the Instagram survey and contains missing observations. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Standard errors are given in parentheses. Standard errors are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Appendix Table A3: Summary Statistics

	Obs.	Mean	Std. dev.	Median	Min	Max
Panel A: Willingness to accept (WTA) elicitations						
<i>Panel A.1: TikTok</i>						
Valuation Keeping Network	371	55.18	58.68	30	-10	210
Valuation Removing Network	371	39.22	53.89	30	-10	210
Product Market Valuation	371	-23.91	96.02	-10	-210	210
Product Market Valuation (with non-users)	707	-43.44	100.83	-10	-210	210
<i>Panel A.2: Instagram</i>						
Valuation Keeping Network	235	47.02	55.99	30	-10	210
Valuation Removing Network	235	37.06	53.58	10	-10	210
Product Market Valuation	235	-6.34	106.39	10	-210	210
Product Market Valuation (with non-users)	260	-9.46	106.79	10	-210	210
<i>Panel A.3: Navigation/maps apps</i>						
Valuation Keeping Network	272	48.82	50.98	30	-10	210
Valuation Removing Network	272	41.18	52.58	30	-10	210
Product Market Valuation	272	16.62	98.60	30	-210	210
Panel B: Comprehension checks						
<i>Panel B.1: TikTok</i>						
% Regretted elicited preferences	1,174	5.37	22.54	0	0	100
% Passed attention checks	1,174	63.97	48.03	100	0	100
<i>Panel B.2: Instagram</i>						
% Regretted elicited preferences	436	5.50	22.83	0	0	100
% Passed attention checks	436	62.84	48.38	100	0	100
<i>Panel B.3: Navigation/maps apps</i>						
% Regretted elicited preferences	468	4.27	20.25	0	0	100
% Passed attention checks	468	60.68	48.90	100	0	100
Panel C: Sample demographics						
<i>Panel C.1: TikTok</i>						
% Active user	707	52.48	49.97	100	0	100
% Female	707	66.05	47.39	100	0	100
Age	707	20.89	2.03	21	18	30
<i>Panel C.2: Instagram</i>						
% Active user	260	90.38	29.54	100	0	100
% Female	260	68.08	46.71	100	0	100
Age	260	20.84	2.15	20.5	18	30
<i>Panel C.3: Navigation/maps apps</i>						
% Active user	272	100.00	0.00	100	100	100
% Female	272	70.96	45.48	100	0	100
Age	272	20.87	1.98	21	18	30


Notes: The table presents summary statistics across all platforms, TikTok, Instagram, and navigation/maps applications. The data collection for TikTok took place in July and the data collection for Instagram and navigation/maps apps took place in August in a cross-randomized survey. The statistics depicting % of respondents are derived from dummy variables multiplied by 100. The % active user represents the fraction of respondents in the final sample who have used the platform at least once in the past month, after filtering those who do not wish to participate in the study and applying regret and attention checks.

Appendix Table A4: Effect of Consumption Spillovers on Welfare Estimates

	Consumer surplus			Negative surplus		
	(1)	(2)	(3)	(4)	(5)	(6)
Instagram	-1.98 (4.87)	-1.87 (3.97)		0.06* (0.03)	0.06* (0.03)	
Product Market Valuation	-32.21*** (5.71)	-32.21*** (5.72)	-32.21*** (5.70)	0.19*** (0.03)	0.19*** (0.03)	0.19*** (0.03)
Instagram × Product Market Valuation	-21.16** (8.20)	-21.16** (8.21)	-21.16*** (8.18)	0.13*** (0.04)	0.13*** (0.04)	0.13*** (0.04)
Uber Valuation		0.70*** (0.07)			-0.00*** (0.00)	
Dep. var. mean	26.98	26.98	26.98	0.23	0.23	0.23
Dep. var sd	84.51	84.51	84.51	0.42	0.42	0.42
Observations	1,014	1,014	1,014	1,014	1,014	1,014
Individual controls	Yes	Yes	No	Yes	Yes	No
Individual FEs	No	No	Yes	No	No	Yes

Notes: The table presents Difference-in-Differences (DiD) coefficient estimates, comparing the elicited individual and product market surplus across two platforms: Instagram and navigation/maps applications. The two dependent variables are (i) the quantitative measure of consumer surplus, denoted as *Consumer surplus*, and (ii) the direction of the consumer surplus, denoted as *Negative surplus*, represented by a binary variable coded as 1 if the surplus is negative and 0 otherwise. Columns 1 and 4 include the following individual control variables: age, gender, and the frequency of platform use, which is determined through a set of qualitative questions. Columns 2 and 5 additionally control for the valuation of the practice good, Uber. Columns 3 and 6 include individual fixed effects. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Standard errors are given in parentheses. Standard errors are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

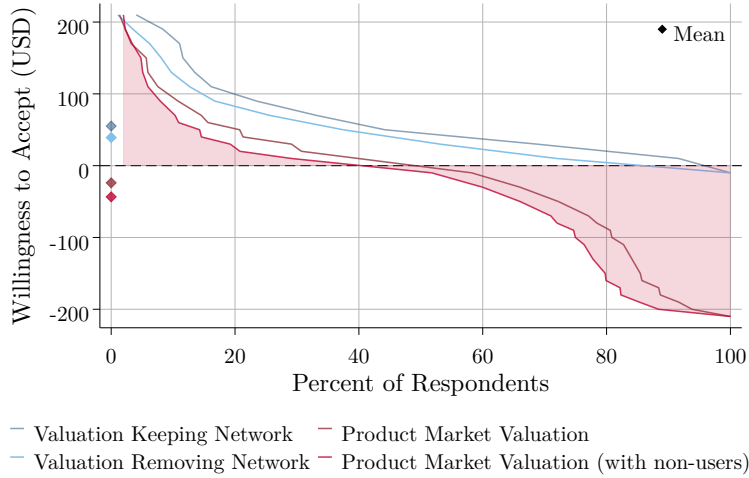
Appendix Figure A2: Reasons for Unwillingness to Participate in Deactivation Study



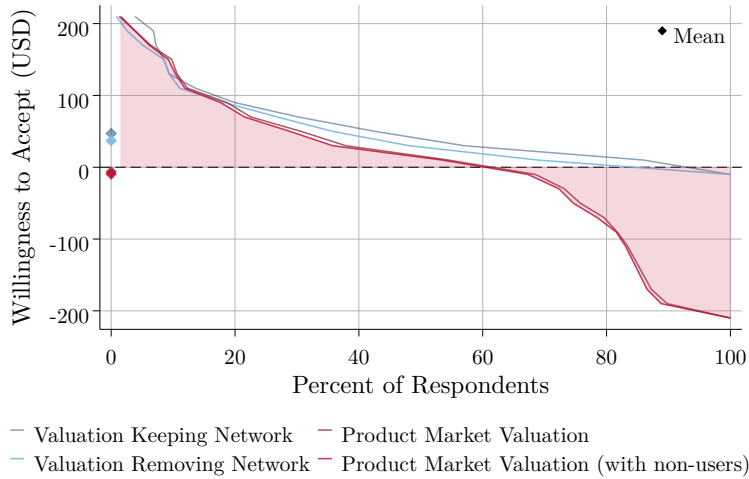
Notes: The figure presents data on the motives behind people’s unwillingness to participate in the deactivation study. The respondents who declined participating in the study were asked the following question: “Why were you unwilling to participate in the study? Please select all that apply.” The figure displays the fraction of respondents that were unwilling to participate because they (i) had privacy concerns, (ii) were unwilling to deactivate their account, and (iii) had a fear of missing out (FOMO). An additional “Other reason not listed above” option was available to ensure genuine feedback. As multiple selections were allowed, the categories presented above are not mutually exclusive. Error bars represent 95% confidence intervals.

Appendix Figure A3: Inverse Demand Function Across Welfare Measures

(a) TikTok




(b) Instagram



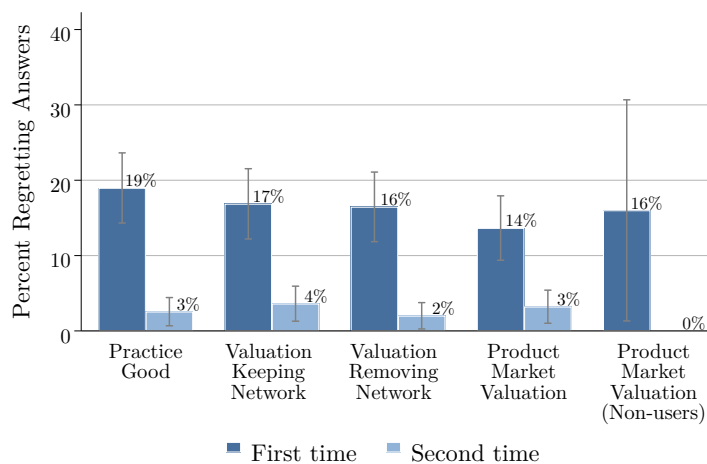
Notes: Figure A3 displays the inverse demand function of respondents' valuation for our different welfare measures. Panel (a) presents the results for TikTok and Panel (b) presents the results for Instagram. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included.

Appendix Figure A4: Regretters Across Steps

(a) TikTok



(b) Instagram

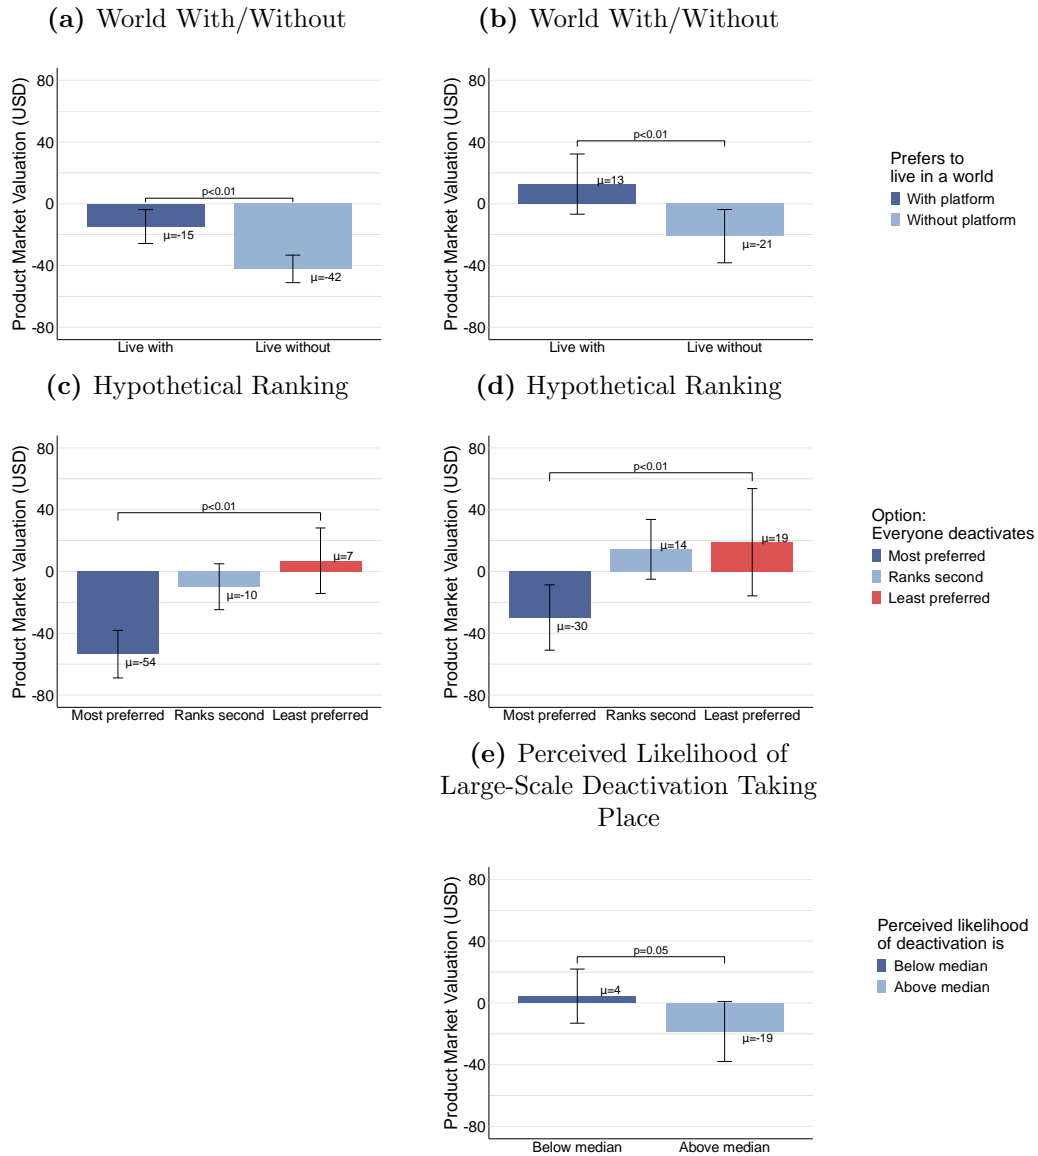


Notes: Figure A4 presents the fraction of respondents that regret their choices across the different measures. Panel (a) presents the results for TikTok and Panel (b) present the results for Instagram. Dark blue bars indicate the fraction of respondents regretting their choices the first time they completed a given valuation. Light blue bars indicate the fraction of respondents regretting their choices the second time they completed a given valuation. Error bars represent 95% confidence intervals.

Appendix Figure A5: Validation of Hypothetical Survey Questions

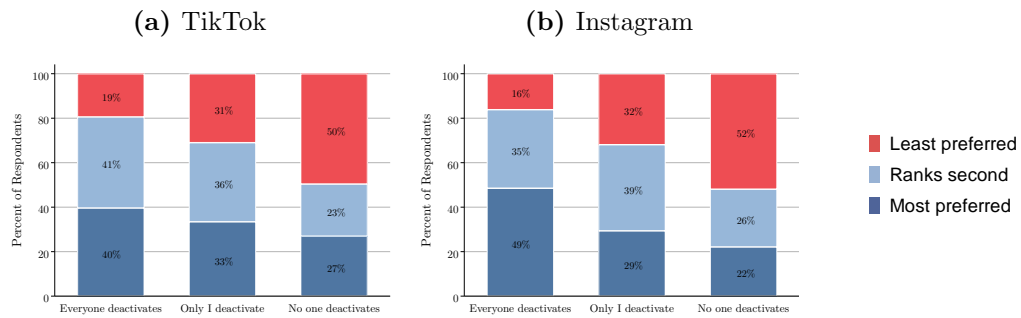
TikTok

Instagram



Notes: Figure A5 presents a validation of the hypothetical survey questions. The outcome variable in the figures is the *Product Market Valuation*. Panels (a) and (c) show results for TikTok. Panels (b), (d) and (e) present results for Instagram. Panels (a) and (b) present the *Product Market Valuation* by people’s preference to live in a world with or without the platform. Panels (c) and (d) presents the *Product Market Valuation* by people’s hypothetical ranking of the deactivation for everyone. Panel (e) presents the *Product Market Valuation* by respondents’ perceived likelihood of the large-scale deactivation study taking place. Error bars represent 95% confidence intervals.

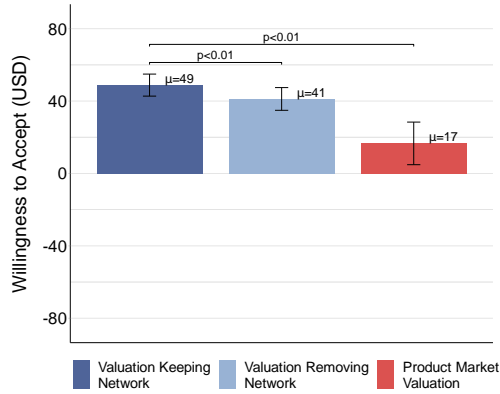
Appendix Figure A6: Hypothetical Ranking of Alternatives



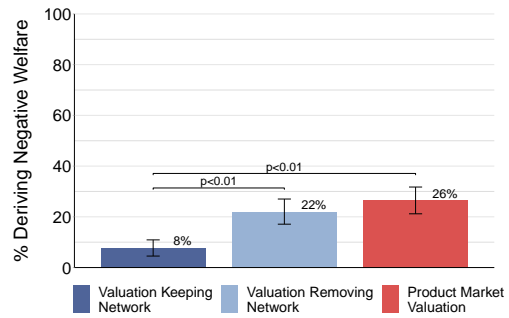
Notes: Figure A6 presents participants' ranking of three hypothetical scenarios about the deactivation of the social media platform: (i) Everyone deactivates (ii) Only I deactivate (iii) No one deactivates. Panel (a) displays results for TikTok and Panel (b) shows results for Instagram. The area in dark blue indicates people's most preferred option; the area in light blue indicates the option ranked second; the area in red shows the least preferred option.

Appendix Figure A7: Consumer Welfare: Navigation and Maps Smartphone Apps


(a) Consumer Surplus Across Welfare Measures



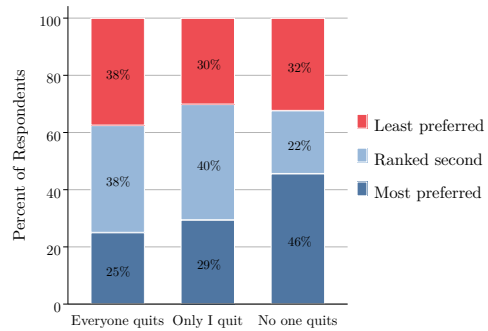
(b) Fraction of Respondents Deriving Negative Welfare



(c) Distribution of Consumer Surplus Across Welfare Measures



(d) Ranking



Notes: Figure A7 presents the survey results for navigation and maps smartphone apps. Figure A7a average valuations for the different welfare measures. Figure A7b presents the fraction of users with negative welfare across the different welfare measures for navigation and maps smartphone apps. Figure A7c presents the probability density function of valuations for the different welfare measures for navigation and maps smartphone apps. Figure A7d presents participants' responses ranking of three hypothetical scenarios: (i) All participating students quit using navigation apps (ii) Only I quit using navigation apps (iii) No one quits using navigation apps. The area in dark blue indicates people's most preferred option; the area in light blue indicates the option ranked second; the area in red shows the least preferred option. In all figures, respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Error bars represent 95% confidence intervals.

C Additional Robustness Checks

C.1 Measurement Error Correction

In this section, we implement the same measurement error detection and correction as in Luttmer and Samwick (2018). We assume that the measured WTA in Step 1 (Valuation Keeping Network, $WTA_i^{1,m}$) relates to the true WTA ($WTA_i^{1,t}$) in the following way:

$$WTA_i^{1,m} = (1 - \beta^1)WTA_i^{1,t} + \beta^1 Offer_i^1,$$

where $Offer_i^1$ is the random initial offer that participants receive in Step 1. Because we randomize the initial offer, we estimate the β^1 coefficient from an OLS regression of Valuation Keeping Network on the initial offer (plus controls).

We follow a similar approach to correct the Valuation Removing Network estimates. We assume that the measured WTA ($WTA_i^{2,m}$) relates to the true WTA ($WTA_i^{2,t}$) in the following way:

$$WTA_i^{2,m} = (1 - \beta^2)WTA_i^{2,t} + \beta^2 Offer_i^2,$$

where $Offer_i^2$ is the initial offer that participants receive in Step 2. Note that these initial offers are not fully random; what we randomize is whether participants receive the lower or upper limit of their Step 1 valuation (unless respondents are at the lower or upper ends of the WTA interval, in which case we offer them again this bound). For these individuals, the offer is fully deterministic and follows the equation $Offer_i^2 = -10 + WTA_i^{1,m} + 20Upper_i$, where $Upper_i$ is the indicator for receiving the upper or lower limit of their Step 1 valuation. Therefore, we estimate the β^2 coefficient by regressing $WTA_i^{2,m}$ on $WTA_i^{1,m}$ and $20Upper_i$. By our randomization of the upper vs. lower limit, the coefficient of $20Upper_i$ gives a consistent estimate of β^2 .

Table A5 presents our estimates for β^1 and β^2 ; with and without controls, for both platforms. Regarding Step 1, we find that initial offers are unrelated to the valuation keeping network for both Instagram ($p=0.55$) and TikTok ($p=0.33$). The results are not only statistically insignificant; the magnitude of the coefficient is close to zero for both platforms, which eases concerns regarding the presence of measurement error in Step 1.

Regarding Step 2, we find that initial offers are positively and significantly related to final valuations for TikTok ($p=0.02$), suggesting the presence of measurement error. These estimates suggest that an extra \$20 in the initial offer induced by our randomization results in an additional \$9 in our participants' valuation removing network. As opposed to TikTok, the estimates for Instagram are not statistically significant ($p=0.48$). While the magnitude of the coefficient is relatively large (0.17), below we show that our estimates do not change substantially after correcting for noise in the case of Instagram.

We report the results from the measurement error correction exercise in Table A6. As the table shows, both the average WTA and the proportion of individuals with a negative valuation in Step 1 remain stable after correcting for the initial offer that individuals

receive, across both platforms. Importantly, the fraction of people with a negative valuation remains unchanged after the correction. These results reduce concerns that our main results are driven by noise in the elicitation.

The valuation removing network (and the fraction of individuals with a negative valuation) in the case of Instagram also remains virtually unchanged after accounting for noise. In the case of TikTok, the measured values are substantially larger, and the fraction of individuals with negative valuation is substantially smaller, than the adjusted values. In particular, accounting for measurement error reduces the average valuation removing network (in the population in which we randomize the bounds) from \$50 to \$32 and increases the percent of individuals with negative valuation from 9% to 24%. We think that this change is not a concern for our main results; it merely indicates that our measure of network effects is conservative in the case of TikTok.

To summarize, these results suggest that measurement error is an unlikely driver of the difference in WTAs that we observe between Steps 1 and 3. If anything, the presence of measurement error might have resulted in an underestimation of network effects in the case of TikTok, since Valuation Removing Network might be overestimated.

Appendix Table A5: Detecting Measurement Error

	WTA	
	(1)	(2)
Panel A: TikTok		
<i>Valuation Keeping Network</i>		
Initial Offer	0.05 (0.05) [0.33]	0.05 (0.05) [0.33]
Observations	371	371
R^2	0.003	0.044
<i>Valuation Removing Network</i>		
Initial Offer	0.45* (0.20) [0.02]	0.47* (0.20) [0.02]
Observations	240	240
R^2	0.570	0.574
Panel B: Instagram		
<i>Valuation Keeping Network</i>		
Initial Offer	-0.03 (0.07) [0.65]	-0.04 (0.07) [0.55]
Observations	235	235
R^2	0.001	0.037
<i>Valuation Removing Network</i>		
Initial Offer	0.20 (0.23) [0.37]	0.17 (0.24) [0.48]
Observations	125	125
R^2	0.735	0.738
Demographic controls	No	Yes

Notes: The table presents a series of regressions assessing the measurement error introduced by the initial offers on respondents' final valuations. Panel A focuses on TikTok and Panel B focuses on Instagram. *Valuation Keeping Network* denotes OLS regressions of Valuation Keeping Network on the random initial offer (plus demographic controls). *Valuation Removing Network* represents OLS regressions of Valuation Removing Network on the "upper bound" dummy (indicating whether participants were offered their upper or lower bound of their valuation in step 1) multiplied times 20, controlling for Valuation Keeping Network (plus demographic controls). For Valuation Keeping Network, we use the full sample of users. For Valuation Removing Network, we focus on those users for whom we randomly offered the upper or lower bound of their Valuation Keeping Network. Demographic controls are age, gender, and platform usage frequency. Robust standard errors are presented in parentheses and p-values are in presented brackets. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

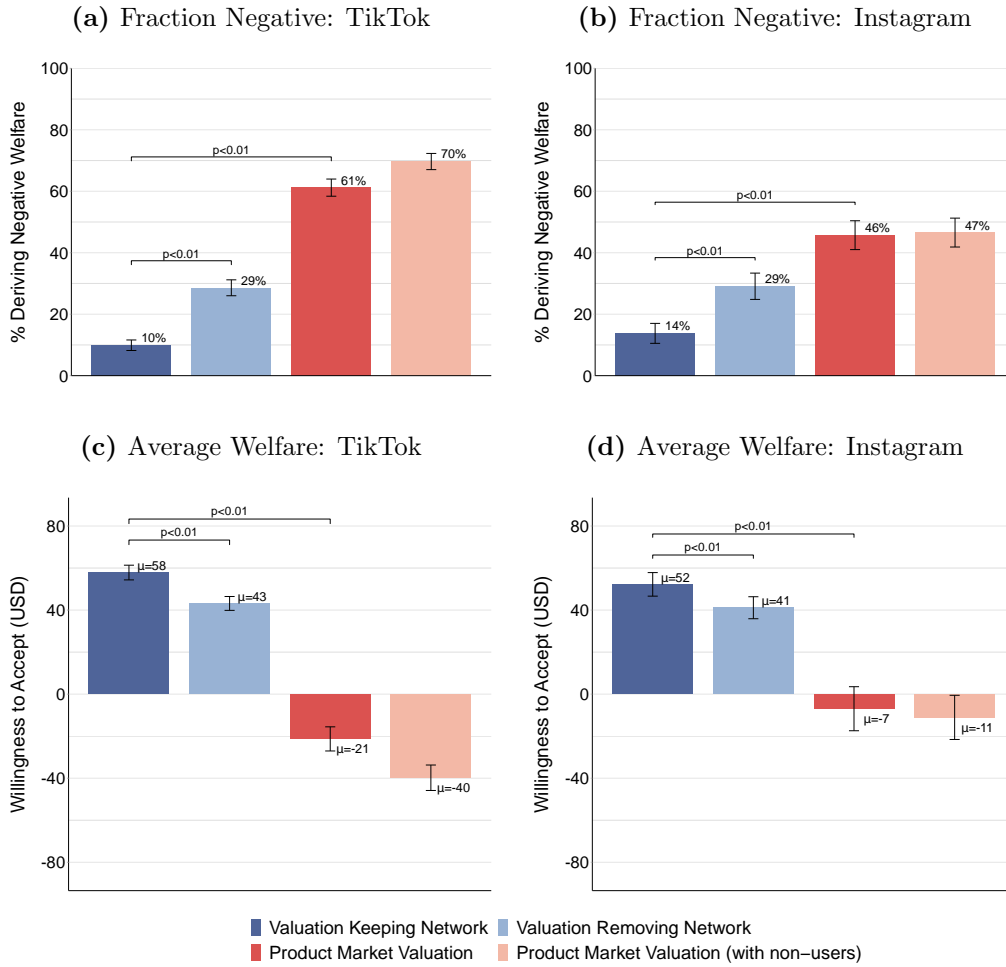
Appendix Table A6: Adjustments for Starting Value

	Valuation Keeping Network		Valuation Removing Network	
	Mean	Proportion < 0	Mean	Proportion < 0
Panel A: TikTok				
Unadjusted	55.18 (3.05)	0.08 (0.01)	50.17 (2.95)	0.09 (2.95)
Initial offer adjustment	-2.20 (2.39)	0.00 (0.05)	-16.19 (24.05)	0.14 (0.14)
Initial offer adjustment with controls	-2.18 (2.36)	0.00 (0.05)	-17.89 (34.23)	0.15 (0.15)
Observations	371	371	240	240
Panel B: Instagram				
Unadjusted	47.02 (3.65)	0.14 (0.02)	60.88 (4.38)	0.06 (4.38)
Initial offer adjustment	1.70 (3.64)	0.00 (0.06)	-2.40 (5.34)	0.02 (0.06)
Initial offer adjustment with controls	2.30 (3.67)	0.00 (0.06)	-1.87 (4.71)	0.00 (0.06)
Observations	235	235	125	125

Notes: The table presents corrected estimations for Valuation Keeping Network and Valuation Removing Network using the formula in Luttmer and Samwick (2018) and the coefficients reported in Table A5. Panel A focuses on TikTok and Panel B focuses on Instagram. Columns 1 and 3 display means and columns 2 and 4 display the % of the sample with negative valuations. For Valuation Keeping Network, we use the full sample of users. For Valuation Removing Network, we focus on those users for whom we randomly offered the upper or lower bound of their Valuation Keeping Network. The Unadjusted row reports statistics before the measurement error correction. The Initial offer adjustment rows report the difference between the adjusted value and the measured value. Initial offer adjustment uses the coefficient reported in Column 1 of Table A5 and Initial offer adjustment with controls uses the coefficient reported in Column 2 of Table A5. Bootstrapped standard errors are presented in parentheses (based on 10,000 replications). When computing the bootstrapped standard errors, we remove those replications that result in β greater than 0.95: 100 and 9 replications (out of 10,000) are removed for TikTok and Instagram, respectively. This only affects Valuation Removing Network. The reason for this is that the adjusted WTA is undefined when β approaches 1 (and hence the standard errors become very large).

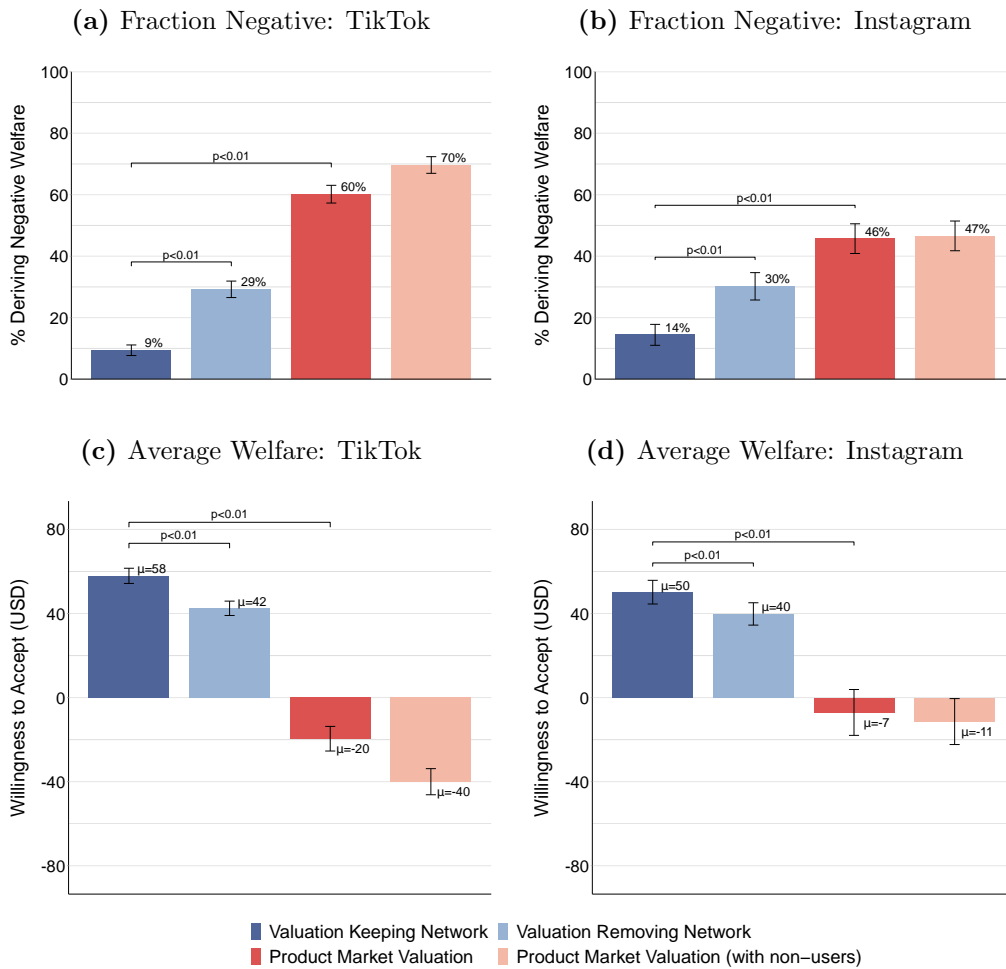
C.2 Sample Selection

Appendix Figure A8: Consumer Surplus across Welfare Measures: Full Sample



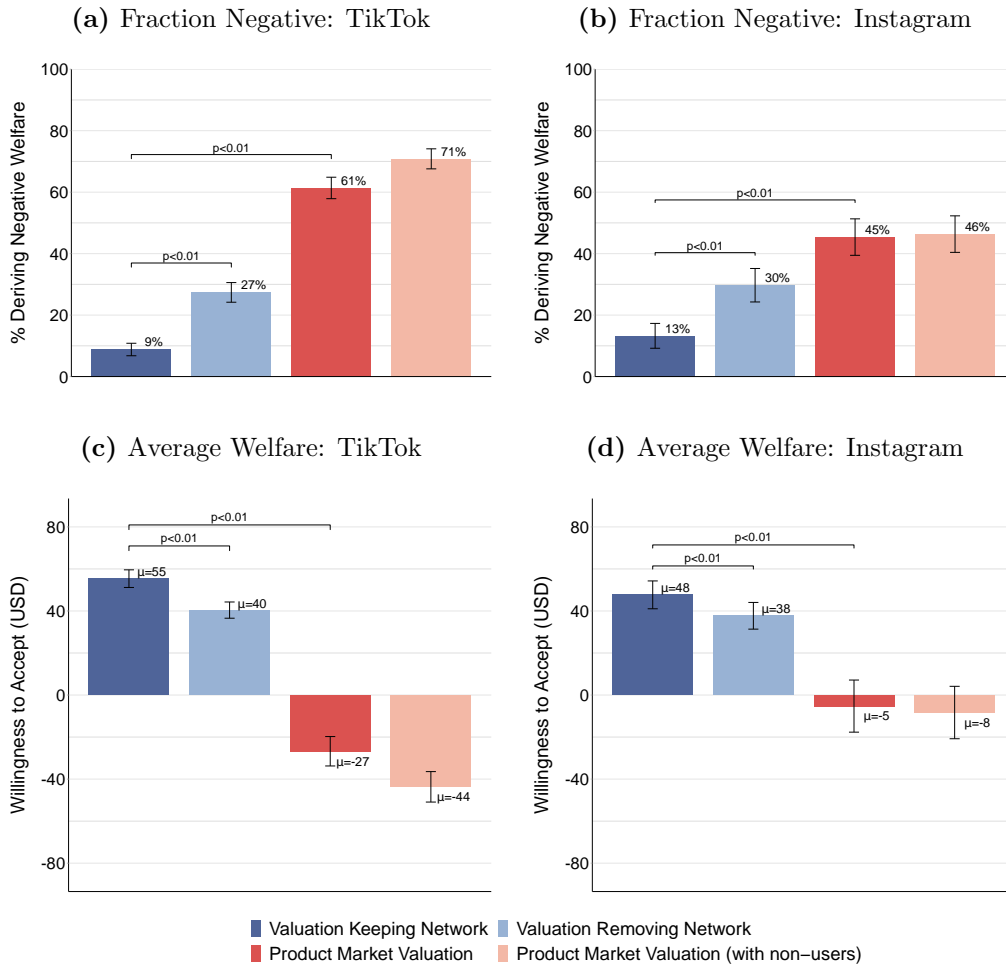
Notes: Panels (a) and (b) of Figure 4 present the percentage of respondents with negative product valuations across our different welfare measures. Panels (c) and (d) present averages. Panels (a) and (c) present the results for TikTok and Panels (b) and (d) present the results for Instagram. The first three bars in each panel represent valuations exclusively for active users. The dark blue bar denotes *Valuation Keeping Network*; the light blue bar denotes *Valuation Removing Network*; the red bar denotes *Product Market Valuation* for users. The pink bar represents the average *Product Market Valuation* of active users and non-users. Error bars represent 95% confidence intervals.

Appendix Figure A9: Consumer Surplus across Welfare Measures: Excluding Regretters and Including Inattentive Respondents



Notes: Panels (a) and (b) of Figure 4 present the percentage of respondents with negative product valuations across our different welfare measures. Panels (c) and (d) present averages. Panels (a) and (c) present the results for TikTok and Panels (b) and (d) present the results for Instagram. The first three bars in each panel represent valuations exclusively for active users. The dark blue bar denotes *Valuation Keeping Network*; the light blue bar denotes *Valuation Removing Network*; the red bar denotes *Product Market Valuation* for users. The pink bar represents the average *Product Market Valuation* of active users and non-users. Only respondents who regretted any of their choices are excluded, while inattentive respondents are included. Error bars represent 95% confidence intervals.

Appendix Figure A10: Consumer Surplus across Welfare Measures: Excluding Inattentive Respondents and Including Regretters



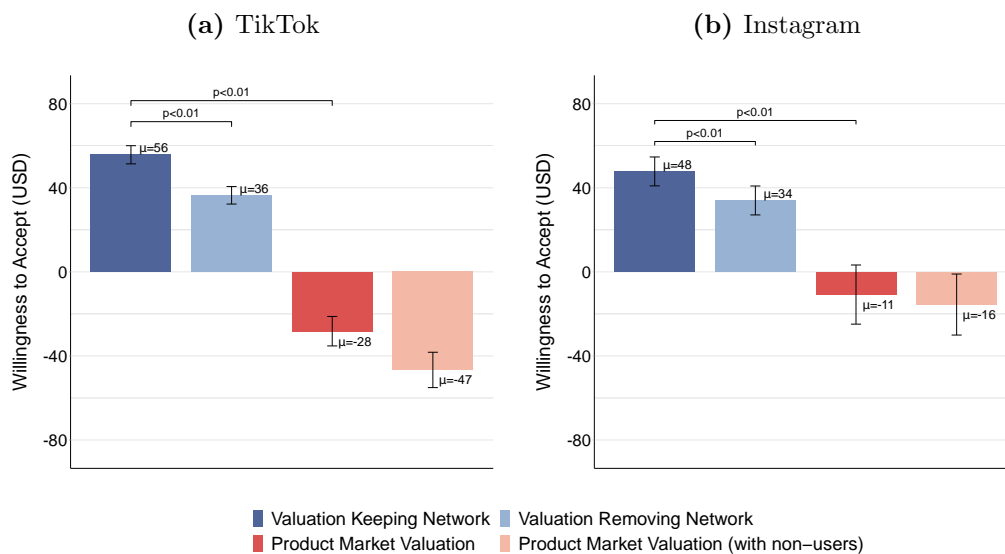
Notes: Panels (a) and (b) of Figure 4 present the percentage of respondents with negative product valuations across our different welfare measures. Panels (c) and (d) present averages. Panels (a) and (c) present the results for TikTok and Panels (b) and (d) present the results for Instagram. The first three bars in each panel represent valuations exclusively for active users. The dark blue bar denotes *Valuation Keeping Network*; the light blue bar denotes *Valuation Removing Network*; the red bar denotes *Product Market Valuation* for users. The pink bar represents the average *Product Market Valuation* of active users and non-users. Only respondents who failed to pass all attention checks are excluded, while those who regretted their choices are included. Error bars represent 95% confidence intervals.

C.3 Triangular Distribution Estimates

In our primary analysis, we employ a multiple price list to narrow down the WTA range of our respondents to \$20 and assign the mean of each respondent's lower and upper bounds to obtain a unique WTA. As a robustness check, we employ an alternative distributional assumption: a triangular distribution to account for potential biases stemming from the unbounded nature of our lowest and highest intervals.

Following Allcott and Kessler (2019) we assume a triangular distribution at the unbounded ranges. To determine a new upper bound, we compute the mass at the upper unbounded interval, [\$200, \$220], and the density at the preceding interval, [\$180, \$200]. Then, using the formula for the probability density function (PDF) for a triangular distribution, we determine the alternative upper bound of the distribution. Subsequently, we compute the mean for the upper unbounded range. Analogously, for the lower unbounded interval, using the same principles we determine a new lower bound and substitute it with \$-10. Figure A11 shows the willingness to accept means for each category, estimated assuming a triangular distribution.

Appendix Figure A11: Consumer Surplus Across Welfare Measures: Triangular Distribution Estimates



Notes: Figure A11 presents average valuations for the different welfare measures assuming triangular distributions for unbounded intervals. Panel (a) presents the results for TikTok and Panel (b) present the results for Instagram. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. The first three bars in each panel represent valuations exclusively for active users. The fourth bar represents the average valuation of active users and non-users. Reported p-values correspond to one-sided t-tests testing the null hypothesis that individual welfare estimates are lower than the aggregate welfare estimate. Error bars represent 95% confidence intervals.

D Open ended responses

Our surveys included two open-ended questions to provide direct evidence on the mechanisms and motives driving consumption. The purpose of this section is two-fold; first, we present an overview of the hand-coding schemes we employed for the categorization of open-ended responses. Second, we summarize the validation of our manual hand-coding with artificial intelligence methods, presenting results from both techniques.

Hand-coding schemes. Table A7 presents the hand-coding scheme applied to open-ended responses for the question: “You mentioned you would prefer to live in a world without [platform]. Why do you still use it?”. As implied by the phrasing, this question targeted only those respondents who previously expressed a desire to live without the platform. Table A7 presents the hand-coding scheme used for the question: “How would you feel if you were the only one who quit using [platform] and everyone else kept using it?”. Note that the categories “Negative”, “Beneficial”, and “Other”, encompass several subcategories. Specifically, “Negative” includes responses mentioning unfairness, impracticality, feeling inferior, dependent, bad, stressed, or lost; whereas “Beneficial” includes responses mentioning self-improvement and feeling positively challenged, unpressured, or good.

Certain respondents expressed a conditional indifference based on compensation or the deactivation’s duration. Given that this does not truly signify ‘indifference’, such responses were categorized under “Other”.⁴⁵ The category also includes relatively infrequent subcategories such as a stated preference to deactivate themselves in order to prevent inconvenience to others, and deriving satisfaction from going against the norm. For both open-ended questions, responses that were non-sensical were excluded from the analysis (N=8, 1.27%).

Validation with Artificial Intelligence. To corroborate our manual categorization, we employed recent artificial intelligence methods, in particular a powerful large language model (GPT-4). We structured a validation exercise with the prompt: “You will be supplied with a list of responses. The responses refer to the usage of different platforms, the platform will be indicated in parentheses at the end of the response. Please classify responses based on the coding scheme below. Please note that each open-ended response can fall into multiple categories or even none.” Subsequent to this, we supplied GPT-4 with the hand-coding scheme, complete with category names, definitions, and illustrative examples. To maintain methodological consistency between our manual coding and GPT-4’s process, we provided GPT-4 with definitions and examples for each subcategory in the subsequent question. These subcategories were subsequently grouped under the primary categories.

Figure A12 displays the category distributions by platform and coding methods. Panel A presents the results for the open-ended question aimed at eliciting the motives for social media consumption despite a preference to live in a world without it; while Panel B

⁴⁵These open-ended questions followed the willingness-to-accept elicitation questions, potentially leading some respondents (9%, N=57) to mistakenly believe that the “deactivation” pertained to the study’s deactivation, which was in exchange for monetary payment and had a four-week duration.

presents the results for second question aimed at unraveling the mechanisms behind non-user consumption spillovers. The juxtaposition of the results of the two coding methods demonstrates that both methods yield remarkably similar results.

To further validate our hand-coding, we conduct a correlational exercise for each category. Once again, the results are presented per question. Each column represents the categories employed for the coding schemes. As displayed in Table A9, all categories have large and statistically significant correlation coefficients across the two methods.

Appendix Table A7: Overview of hand-coding scheme for reasons to use TikTok/Instagram/Maps despite preferring a world without it

Category	Definition	Example(s)
FOMO	Respondent mentions fear of missing out, feeling out of the loop, their wish to stay connected, or justifies usage through others' usage	"I feel compelled to keep 'in touch' with what I perceive as being the culturally relevant 'thing' at the moment. It breeds a sense of FOMO when you don't use it." (<i>TikTok</i>); "Everyone else uses it so I feel that I will be missing out if I don't." (<i>Instagram</i>); "I still use navigation maps because it is what everyone uses [...]" (<i>Maps</i>)
Entertainment	Respondent mentions they use it to be entertained	"It's a very good source of entertainment and it's always something to do when bored." (<i>TikTok</i>); "It's a default way to pass time when I'm bored." (<i>Instagram</i>);
Addiction	Respondent mentions inability to let go or directly mentions addiction	"I use TikTok as a habit. I hate TikTok and know that I have other things I need to do, but I subconsciously click on it, then scroll for hours. It's very hard to control it." (<i>TikTok</i>); "Because I am addicted to the scrolling and tired of wasting valuable time on the app." (<i>Instagram</i>)
Information	Respondent mentions informational purposes such as following the news, keeping abreast of college events, or getting directions.	"for information on current events because i do not watch the news" (<i>TikTok</i>); "I use it to keep inform about my university events and news" (<i>Instagram</i>); "I don't know where to go" (<i>Maps</i>)
Productivity/ Convenience	Respondent mentions convenience of use or states to use platform for productive/business purposes.	"It's easy to see stuff I like (art, new art news, movie reviews, etc)." (<i>TikTok</i>); "I still use insta-gram for business purposes." (<i>Instagram</i>); "It's more convenient than pulling out a map and I have a terrible sense of direction" (<i>Maps</i>)

Notes: The table displays an overview of the hand-coding scheme used for categorizing the open-ended answers given to the question: "You mentioned you would prefer to live in a world without TikTok/Instagram/navigation apps. Why do you still use it/them?". The question was only asked to participants that are active users of the respective platforms and stated they would prefer to live in a world without said platform. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included.

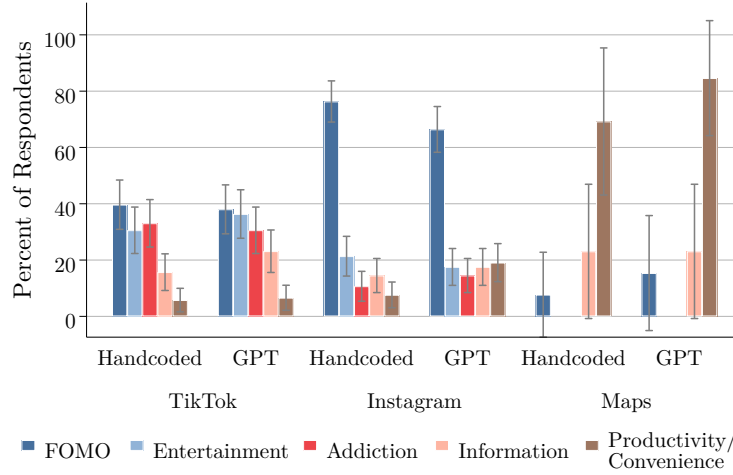
Appendix Table A8: Overview of hand-coding scheme for how respondent would feel if only they were to quit using platform

Category	Definition	Example
FOMO	Respondent mentions fear of feeling missing out, left out, or being out of the loop	“I would probably feel somewhat out of the loop when it comes to trends, with a consistent feeling of FOMO.” (<i>TikTok</i>); “I would feel really left out since a lot of people use it to communicate about events and parties and with one another” (<i>Instagram</i>); “I would feel a bit isolated, maybe excluded from certain conversations involving travel plans, etc” (<i>Maps</i>)
Negative	Respondent expresses negative emotions; that it would be unfair, impractical, etc.	“it would be a little unfair” (<i>TikTok</i>); “[...] feel discouraged and jealous of everyone else.” (<i>Instagram</i>); “I would feel lost and not confident in my ability to navigate” (<i>Maps</i>)
Indifferent	Respondent states that they would not be particularly affected	“No different, because I don’t use tiktok often anyway”; (<i>TikTok</i>); “I would be fine. I don’t really post on Instagram. It wouldn’t be much of a change.” (<i>Instagram</i>); “Wouldn’t mind as long as knew my way around” (<i>Maps</i>)
Beneficial	Respondent mentions deriving a benefit	“relieved, probably.” (<i>TikTok</i>); “I would feel free”; (<i>Instagram</i>); “It will be an awesome experiment and experience, asking everyone for directions” (<i>Maps</i>)
Other	Diverse set of motives; including substituting platform, indifference conditional on getting paid, or fondness to spare others the struggle	“I would just use other social media” (<i>TikTok</i>); “It’s okay as long as I have some monetary benefit in it.”; (<i>Instagram</i>); “... I don’t want everyone else to struggle especially since people have different like circumstances” (<i>Maps</i>)


Notes: The table displays an overview of the hand-coding scheme used for categorizing the open-ended answers given to the question: “How would you feel if you were the only one who deactivated (quit using) TikTok/Instagram (navigation/maps apps) and everyone else kept using it (them)?”. The question was only asked to participants that are active users of the respective platform. In the TikTok survey, the question was further restricted to respondents who stated they would prefer to live in a world without TikTok. We did not apply this restriction for the Instagram/maps survey. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included.

Appendix Figure A12: Validation based on Large Language Model

(a) Motives for social media consumption despite a preference to live in a world without it



(b) Mechanisms behind non-user consumption spillovers



Notes: Figure A12 presents the distribution of categories based on open-ended responses separately for the hand-coded data and the data coded by a large language model (GPT4). Panel (a) details the results for the question, “You mentioned you would prefer to live in a world without TikTok/Instagram/navigation apps. Why do you still use them?” Meanwhile, Panel (b) showcases the results for the question, “How would you feel if you were the only one who stopped using TikTok/Instagram/navigation apps, while everyone else continued their use?”

Appendix Table A9: Validation of hand-coded data from Large Language Model

Panel A: Motives for social media consumption despite a preference to live in a world without it					
	FOMO	Entertainment	Addiction	Information	Productivity/ Convenience
Correlation coefficient	0.762 (0.040)	0.698 (0.044)	0.863 (0.031)	0.664 (0.046)	0.535 (0.052)
<i>Hand-coded responses:</i>					
Mean	0.562	0.245	0.204	0.155	0.098
Std. dev.	0.497	0.431	0.404	0.362	0.298
<i>GPT-4 coded responses:</i>					
Mean	0.509	0.253	0.211	0.204	0.166
Std. dev.	0.501	0.435	0.409	0.404	0.373
Observations	265	265	265	265	265
Panel B: Evidence on mechanisms behind non-user consumption spillovers					
	FOMO	Indifferent	Negative	Beneficial	Other
Correlation coefficient	0.885 (0.019)	0.860 (0.020)	0.693 (0.029)	0.718 (0.028)	0.613 (0.032)
<i>Hand-coded responses:</i>					
Mean	0.273	0.297	0.241	0.146	0.140
Std. dev.	0.446	0.457	0.428	0.353	0.347
<i>GPT-4 coded responses:</i>					
Mean	0.254	0.281	0.241	0.162	0.220
Std. dev.	0.435	0.450	0.428	0.369	0.415
Observations	623	623	623	623	623

Notes: The table presents the correlation coefficients between our manual categorization and the GPT-4 categorization of open-ended responses. Each column corresponds to a specific category used in the classification process. Correlation coefficients were calculated using dummy variables: for every coding technique, a dummy variable is set to 1 if the open-ended response fits within a particular category. These coefficients then show the correlation between these dummy variables. Panel A details the results for the question, “You mentioned you would prefer to live in a world without TikTok/Instagram/navigation apps. Why do you still use them?” Meanwhile, Panel B showcases the results for the question, “How would you feel if you were the only one who stopped using TikTok/Instagram/navigation apps, while everyone else continued their use?” Standard errors are given in parentheses and are computed based on the Pearson correlation coefficient formula.

E Other Applications

E.1 Luxury Goods

We start with evidence on luxury goods, where positional externalities are a plausible driver of negative non-user utility.

Sample. We fielded pre-registered surveys with 500 US participants from Prolific, a widely used online labor market used for social science experiments (Eyal et al., 2021), in September 2023.⁴⁶

Survey. Our survey consists of two randomly ordered blocks: one block on luxury goods discussed in this section and another block on vintage goods presented in Section E.2. In the luxury block, we ask respondents to indicate whether they owned products from luxury brands they personally purchased.⁴⁷ We then ask respondents whether they prefer to live in a world with or without any of these luxury brands. The full set of instructions can be found in Appendix G.3.

Results. In our survey, 32% of respondents own luxury brands. Conditional on owning any luxury brand, they owned 2.04 luxury brands on average. Figure A13 shows that among respondents who owned any goods of luxury brands, 44% preferred to live in a world without those brands. Among respondents not owning any of these brands, the fraction preferring to live in a world without them is higher, at 69%. While the literature on luxury goods has emphasized the negative externalities these goods impose on non-consumers (Frank, 1985, 2000, 2012), our evidence highlights that large shares of consumers of these products would prefer them not to exist.⁴⁸ Given that these results are in line with our social media estimates, it is plausible that status concerns might be an important mechanism driving negative non-consumer surplus.

E.2 Frequency of Product Variations

Product market traps lead to a situation where the existence of a product is harmful to consumers. This can manifest as excessive consumption by users or the production of an excessive number of product variations or vintages (Pesendorfer, 1995).

⁴⁶The pre-registration can be found on AsPredicted #144630.

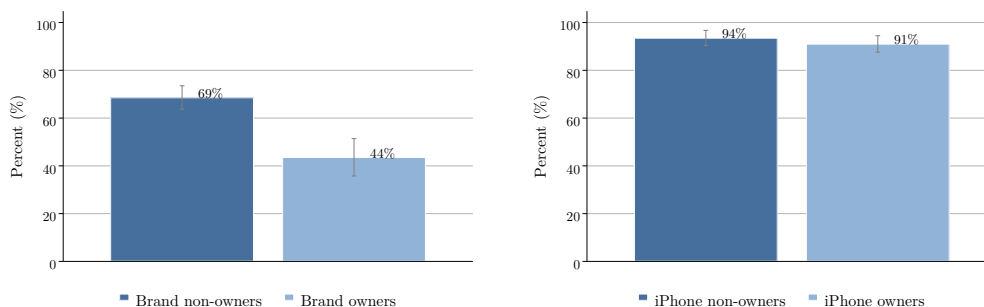
⁴⁷The brands we used are: Louis Vuitton, Gucci, Chanel, Yves Saint Laurent (YSL), Balenciaga, Versace, Rolex, Tiffany & Co., Burberry, Givenchy, and Swarovski. Respondents could also fill in any other luxury brand not part of this list.

⁴⁸An alternative interpretation of our findings is that a participant has a positive product-market valuation for brand X, but valuations for the remaining brands are so negative that the overall preference would still be to live in a world without any of the brands.

Appendix Figure A13: Luxury and Vintage Goods

(a) Percentage of Respondents that Prefer to Live in a World Without Any Luxury Brands

(b) Percentage of Respondents that Prefer to Live in a World Where iPhone is Released Every Other Year



Notes: Panel (a) of Figure A13 displays the fraction of respondents preferring to live in a world without any luxury brands separately for brand owners and brand non-owners. Panel (b) displays the fraction of respondents that prefer to live in a world where Apple releases the new iPhone every other year rather than every year, separately for iPhone owners and iPhone non-owners. Error bars represent 95% confidence intervals.

To examine people’s preferences regarding the frequency of product variations, we asked respondents whether they would prefer to live in a world where Apple releases the iPhone every year or every other year in the survey presented in the previous section. We document that, among iPhone owners, a striking 91% of respondents would prefer Apple to release the iPhone every other year rather than every year.⁴⁹ Among respondents not owning the iPhone, 94% prefer Apple to release the iPhone every other year rather than every year. This finding provides suggestive evidence that consumers consider the number of product variations or vintages of the iPhone as excessive and thus harmful to consumer welfare. Overall, the findings from this survey suggest that negative non-consumer surplus is not specific to the case of social media, but also extends to luxury consumption and particular high-end technology products.⁵⁰

⁴⁹It is worth highlighting that among iPhone owners only 8% prefer to live in a world without iPhones, while among respondents not owning the iPhone this fraction is 49%.


⁵⁰An alternative explanation for these consumer preferences could be environmental motives (along with the belief that others will tend to purchase the newest version).

F Main Exhibits Including Pilot Data

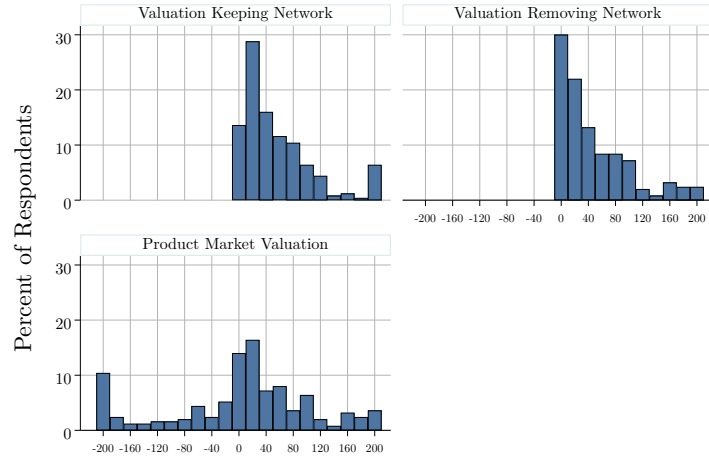
For the Instagram and Maps experiment, we pre-registered that we would pool the final data with 28 pilot responses. While we deviated from this plan and instead only report the pre-registered data in our main exhibits, we present the main results pooling the final data with the pilot responses in this section. As there are no changes in the TikTok sample, this section solely focuses on Instagram and Maps.

Appendix Figure A14: Distribution of Consumer Surplus Across Welfare Measures

(a) Instagram

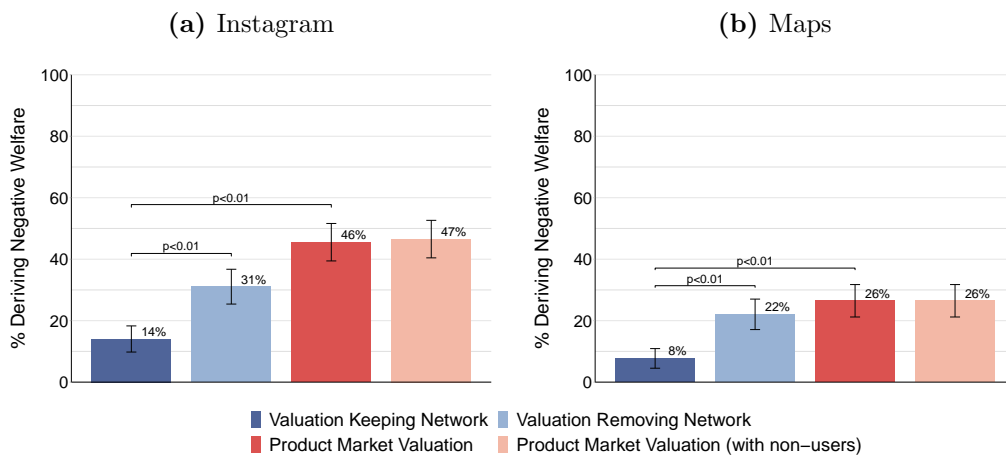


(b) Maps



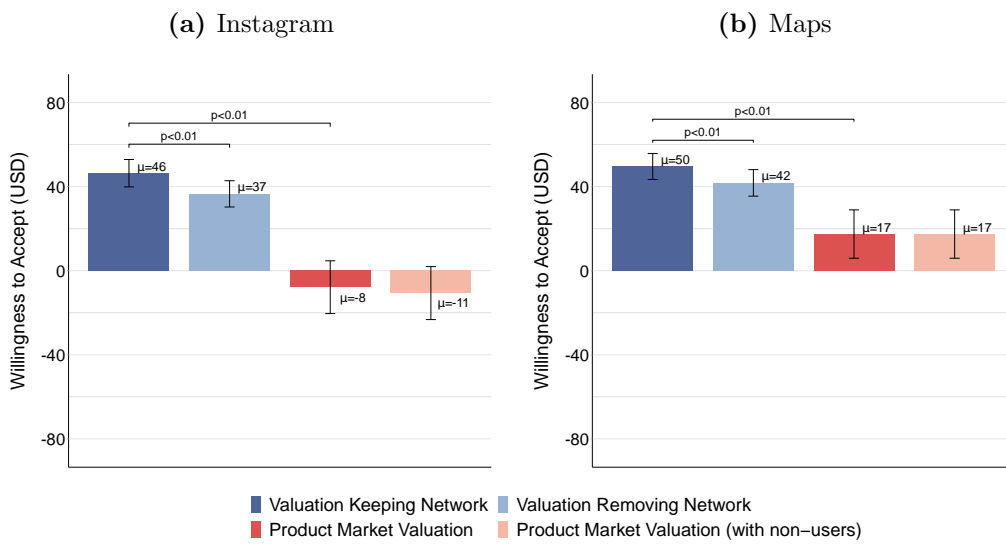
Notes: Figure A14 presents the probability density function of valuations for the different welfare measures. Panel (a) presents the results for Instagram and Panel (b) present the results for Maps. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included.

Appendix Figure A15: Fraction with Negative Welfare across Welfare Measures




Notes: Figure A15 presents the percent of respondents with negative product valuations across our different welfare measures. Panel (a) presents the results for Instagram and Panel (b) presents the results for Maps. The first three bars in each panel represent valuations exclusively for active users. The dark blue bar denotes *Valuation Keeping Network*; the light blue bar denotes *Valuation Removing Network*; the red bar denotes *Product Market Valuation* for users. The pink bar represents the average *Product Market Valuation* of active users and non-users. For Maps, the red and pink bars are identical as there are no non-users in our sample. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Error bars represent 95% confidence intervals.

Appendix Figure A16: Consumer Surplus across Welfare Measures



Notes: Figure A16 presents average valuations for the different welfare measures. Panel (a) presents the results for Instagram and Panel (b) presents the results for Maps. The first three bars in each panel represent valuations exclusively for active users. The dark blue bar denotes *Valuation Keeping Network*; the light blue bar denotes *Valuation Removing Network*; the red bar denotes *Product Market Valuation* for users. The pink bar represents the average *Product Market Valuation* of active users and non-users. For Maps, the red and pink bars are identical as there are no non-users in our sample. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Reported p-values correspond to one-sided t-tests testing the null hypothesis that individual welfare estimates are lower than the aggregate welfare estimate. Error bars represent 95% confidence intervals.

Appendix Figure A17: Percentage of Respondents that Prefer to Live in a World without the Platform



Notes: Figure 5 displays the percent of the respondents that stated they would prefer to live in a world without the platform for Instagram and Maps separately. The dark blue bar represents the fraction among all respondents and the light blue bar represents the fraction among active users of the respective platform. For Maps, the dark and light blue bars are identical as there are no non-users in our sample. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Error bars represent 95% confidence intervals.

G Experimental Instructions

G.1 TikTok: July 2023

G.1.1 Introduction to Survey and Deactivation Study Instructions

University of Chicago

Online Consent Form for Research Participation

Study Number: IRB23 - 0797

Researcher: Leonardo Bursztyn

Description: You will be asked to fill out a short survey. Participation is voluntary and takes about 8 minutes.

Incentives: Upon completion of this survey, you will be compensated by your survey provider.

Risks and Benefits: There are no foreseeable risks associated with this study beyond those involved in answering a survey. The research team cannot and do not guarantee or promise that you will receive any benefits from this study.

Confidentiality: Confidentiality of your research records will be strictly maintained by storing any personally identifiable data in secure accounts that can only be accessed by researchers in this study. After the experiment is over, we will delete any personally identifiable information from our dataset and replace it with an arbitrary participant number. This will allow us to maintain your privacy in all published and written data resulting from the study. Information not containing identifiers may be used in future research or shared with other researchers without your additional consent. If you decide to withdraw, data collected up until the point of withdrawal may still be included in the analysis.

Contacts & Questions: If you have questions or concerns about the study, you can contact the researcher at: bursztyn@uchicago.edu.

If you have any questions about your rights as a participant in this research, feel you have been harmed, or wish to discuss other study-related concerns with someone who is not part of the research team, you can contact the University of Chicago Social & Behavioral Sciences Institutional Review Board (IRB) Office by phone at (773) 702-2915, or by email at sbs-irb@uchicago.edu

Consent: Participation is voluntary. Refusal to participate or withdrawing from the research will involve no penalty or loss of benefits to which you might otherwise be entitled. By clicking "Agree" below, you confirm that you have read the consent form, are at least 18 years old, and agree to participate in the research. Please print or save a copy of this page for your records.

I agree to participate in the research. I confirm that I am above 18 years of age or older.

I do NOT agree to participate in the research. You will be directed to an exit screen.

This survey is directed at university students' social media preferences. It is designed by our research team at the University of Chicago and is administered via partnership with College Pulse.*

Thank you for participating!

*College Pulse is a research and analytics company that specifically aims to understand the attitudes, preferences, and behaviors of today's college students.

How frequently did you use each of the following **social media platforms** in the past month?

	Not at all	Once	Once a week	Twice a week	Every day
Instagram	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Facebook	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
TikTok	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Twitter	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Please read the instructions of this survey carefully. We will ask you some questions to ensure your understanding of its content and give you a **bonus payment** if you answer correctly!



As mentioned, this survey is directed at university students' social media preferences.

After this survey, we will conduct a study in which we will ask **students at your university to deactivate their TikTok accounts for four weeks** in exchange for a **monetary payment**.

"Deactivation" studies like this have been conducted in the past (e.g., by Mosquera et al., 2018 and Allcott et al., 2020) with close to 90% compliance.

Should students deactivate their TikTok, they can go back to using it whenever they want, with their content and network unchanged, but they will then forgo any monetary payment.

To verify that students deactivate their TikTok accounts, we will **periodically visit their profiles** and **require them to upload screenshots of their app usage**.

The next part of the study involves asking you if you want to deactivate your TikTok account for four weeks in exchange for monetary payment.

For that we will need to collect your TikTok handle and we would ask you to submit periodic screenshots of your phone's time use statistics if you are selected.

Would you be willing to participate?

Yes

No



How will we verify that selected users deactivate their TikTok accounts?

By periodically visiting their profiles

By requiring them to upload screenshots

By asking them to install an app

By periodically visiting their profiles and requiring them to upload screenshots



For how long will we ask selected users to deactivate their TikTok accounts?

Eight weeks

Four weeks

One week

Ten weeks

G.1.2 Step 0: Practice Good

Before we begin, we will ask you a series of hypothetical practice questions for you to get accustomed to our survey.



To understand how much people value their Uber account, **we ask university students, including you**, to decide whether to **deactivate their Uber account** for four weeks, in exchange for different monetary payments.

A computer will **randomly select one student from your university** to be eligible for the deactivation study.



The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions offering you different payment scenarios in case you are selected for the deactivation study.

- If you accept any price scenario lower than the computer's offer, we will invite you to the deactivation study and give you the computer's offer.
- If you do not accept any price scenario lower than the computer's offer, we will not invite you to the deactivation study even if you are selected.

This rule means that the higher the amount you require the lower the chance that you will receive the computer's offer.

Therefore, while answering the following questions, **please choose carefully as each answer could be "the one that counts."**

We will now ask you a comprehension question based on the text above:

Which of the following statements is true?

- The amount I require does not affect the chance that I receive the computer's offer
- The higher the amount I require the lower the chance that I receive the computer's offer
- The higher the amount I require the higher the chance that I receive the computer's offer

Uber

Which of the following options would you prefer?

I deactivate my Uber account AND I receive \$120

I keep my Uber account active

Uber

Which of the following options would you prefer?

I deactivate my Uber account AND I receive \$60

I keep my Uber account active

Uber

Which of the following options would you prefer?

I deactivate my Uber account AND I receive \$20

I keep my Uber account active

Uber

Which of the following options would you prefer?

I deactivate my Uber account AND I receive \$40

I keep my Uber account active

Uber

According to your answers to the previous questions, you would require a payment worth between \$20 and \$40 to deactivate your Uber account for four weeks.

Do you agree with the above statement about your **valuation**?

Yes

No

G.1.3 Step 1: Valuation Keeping Network

You have now completed the practice section. We will now ask you a series of questions about your **social media preferences**.



To establish appropriate payment amounts for the deactivation study, **we ask university students, including you**, to decide whether to deactivate their TikTok accounts for different monetary amounts.

A computer will **randomly select one student from your university** to be eligible for the deactivation study.



The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions offering you different payment scenarios in case you are selected for the deactivation study.

Therefore, while answering the following questions, **please choose carefully as each answer could be "the one that counts."**



Which of the following options would you prefer?

I deactivate my TikTok account AND I receive \$60

I keep my TikTok account active



Which of the following options would you prefer?

I deactivate my TikTok account AND I receive \$140


I keep my TikTok account active



Which of the following options would you prefer?

I deactivate my TikTok account AND I receive \$180


I keep my TikTok account active



Which of the following options would you prefer?

I deactivate my TikTok account AND I receive \$160

I keep my TikTok account active




According to your answers to the previous questions, you would require a payment worth between \$160 and \$180 to deactivate your TikTok account for four weeks.

Do you agree with the above statement about your **valuation**?

Yes

No

G.1.4 Step 2: Valuation Removing Network




College Pulse has a panel exceeding 650,000 university students.* We are targeting universities with a high penetration of College Pulse.

We will now ask you to consider two additional options for a large-scale deactivation of TikTok at your university. **One of them will be randomly implemented** if we manage to recruit **more than two-thirds of the students at your university**.

We expect 90% of students to comply with deactivation based on previous studies (e.g., by Mosquera et al., 2018 and Allcott et al., 2020).

*See <https://collegepulse.com>



Option 1 for the large-scale deactivation study:

In collaboration with College Pulse, we will ask students at your university sequentially whether they would like to deactivate their TikTok accounts.

Consider the scenario where we have asked **all participating students at your university** to deactivate their TikTok accounts for four weeks in exchange for a payment.



The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions offering you different payment scenarios in case you are selected for the deactivation study.

Therefore, while answering the following questions, **please choose carefully as each answer could be "the one that counts."**



Consider the scenario where we have asked **all participating students at your university** to **deactivate their TikTok accounts** for four weeks in exchange for monetary payment.

If you had the choice to **also** deactivate your TikTok account for the next four weeks, which of the following options would you prefer?

I deactivate my TikTok account if the other students deactivate their TikTok accounts AND I receive \$180

I keep my TikTok account active if the other students deactivate their TikTok accounts



Consider the scenario where we have asked **all participating students at your university** to **deactivate their TikTok accounts** for four weeks in exchange for monetary payment.

If you had the choice to **also** deactivate your TikTok account for the next four weeks, which of the following options would you prefer?

I deactivate my TikTok account if the other students deactivate their TikTok accounts AND I receive \$80

I keep my TikTok account active if the other students deactivate their TikTok accounts



Consider the scenario where we have asked **all participating students at your university** to **deactivate their TikTok accounts** for four weeks in exchange for monetary payment.

If you had the choice to **also** deactivate your TikTok account for the next four weeks, which of the following options would you prefer?

I deactivate my TikTok account if the other students deactivate their TikTok accounts AND I receive \$40

I keep my TikTok account active if the other students deactivate their TikTok accounts



Consider the scenario where we have asked **all participating students at your university** to **deactivate their TikTok accounts** for four weeks in exchange for monetary payment.

If you had the choice to **also** deactivate your TikTok account for the next four weeks, which of the following options would you prefer?

I deactivate my TikTok account if the other students deactivate their TikTok accounts AND I receive \$60

I keep my TikTok account active if the other students deactivate their TikTok accounts



According to your answers to the previous questions, you would require a payment worth between \$40 and \$60 to deactivate your TikTok account for four weeks, if we ask all participating students at your university to deactivate their TikTok accounts.

Do you agree with the above statement about your **valuation**?

Yes

No

G.1.5 Step 3: Product Market Valuation



Option 2 for the large-scale deactivation study:

We know how much we would need to pay **every participating student at your university** to deactivate their TikTok accounts for four weeks.

To give everyone an equal chance to decide, we will randomly choose one of the students.

This student's identity will remain anonymous, and they can choose one from the following options:

1. **We ask all participating students with a TikTok account to deactivate it, or**
2. **We keep things as they are.**

If you decide for all participating students to deactivate their TikTok accounts:

- **We will pay the other students the amount they required** and we will **establish your payment, if any, below.**
- The deactivation study will be stopped for everyone only if you go back to using TikTok before the end of the four weeks.
- If the study is stopped early, you will not receive payment and we will pay the other students based on the actual time they spend in the study.
- If someone from the other participating students goes back to using TikTok before the end of the study, they themselves will not receive any payment.



Suppose that you decide for us to ask all participating students to deactivate their TikTok accounts:

Which of the following statements is correct?

- We will force the other students to deactivate their TikTok accounts
- We will pay the other students what they required to deactivate their TikTok accounts
- We will pay the other students more than what they required to deactivate their TikTok accounts



The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions with different payment scenarios in case you are selected.

Therefore, while answering the following questions, **please choose carefully as each answer could be "the one that counts."**



Which of the following options would you prefer?

- All participating students at my university deactivate their TikTok accounts
- All participating students at my university keep their TikTok accounts active



Which of the following options would you prefer?

- All participating students at my university deactivate their TikTok accounts
- All participating students at my university keep their TikTok accounts active AND I receive \$160



Which of the following options would you prefer?

- All participating students at my university deactivate their TikTok accounts
- All participating students at my university keep their TikTok accounts active AND I receive \$180



Which of the following options would you prefer?

- All participating students at my university deactivate their TikTok accounts
- All participating students at my university keep their TikTok accounts active AND I receive \$200



According to your answers to the previous questions, you would forgo a payment above \$200 to have all participating students at your university, including you, deactivate their TikTok accounts for four weeks.

Do you agree with the above statement about your **valuation**?

Yes

No

As you have disagreed with the payment amount, you will be redirected to the beginning of the previous section.

This is the last chance to modify your answers.



Which of the following options would you prefer?

All participating students at my university deactivate their TikTok accounts

All participating students at my university keep their TikTok accounts active



Which of the following options would you prefer?

All participating students at my university deactivate their TikTok accounts AND I receive \$40

All participating students at my university keep their TikTok accounts active



Which of the following options would you prefer?

All participating students at my university deactivate their TikTok accounts AND I receive \$20

All participating students at my university keep their TikTok accounts active



According to your answers to the previous questions, you would require a payment worth between \$20 and \$40 to have all participating students at your university, including you, deactivate their TikTok accounts for four weeks.

Do you agree with the above statement about your **valuation**?

Yes

No

G.1.6 Qualitative Questions

Please rank the following options from your most preferred (1) to your least preferred (3).

1 I deactivate TikTok and every other student at my university keeps using it

2 Every student at my university, including me, deactivates TikTok

3 No one deactivates TikTok

Would you prefer to live in a world with or without TikTok?

I would prefer to live in a world with TikTok

I would prefer to live in a world without TikTok

You mentioned you would prefer to live in a world without TikTok.

Why do you still use it?

How would you feel if you were the only one deactivating your TikTok and everyone else kept using it?

Now some **demographic** questions.

Which of the following describes you more accurately?

Male

Female

Other / Prefer not to say

What is your age?

G.2 Instagram and Maps: August and September 2023

G.2.1 Introduction to Survey and Deactivation Study Instructions

University of Chicago

Online Consent Form for Research Participation

Study Number: IRB23 - 0797

Researcher: Leonardo Bursztyn

Description: You will be asked to fill out a short survey. Participation is voluntary and takes around 10 minutes.

Incentives: Upon completion of this survey, you will be compensated by your survey provider.

Risks and Benefits: There are no foreseeable risks associated with this study beyond those involved in answering a survey. The research team cannot and does not guarantee or promise that you will receive any benefits from this study.

Confidentiality: Confidentiality of your research records will be strictly maintained by storing any personally identifiable data in secure accounts that can only be accessed by researchers in this study. After the experiment is over, we will delete any personally identifiable information from our dataset and replace it with an arbitrary participant number. This will allow us to maintain your privacy in all published and written data resulting from the study. Information not containing identifiers may be used in future research or shared with other researchers without your additional consent. If you decide to withdraw, data collected up until the point of withdrawal may still be included in the analysis.

Contacts & Questions: If you have questions or concerns about the study, you can contact the researcher at: bursztyn@uchicago.edu.

If you have any questions about your rights as a participant in this research, feel you have been harmed, or wish to discuss other study-related concerns with someone who is not part of the research team, you can contact the University of Chicago Social & Behavioral Sciences Institutional Review Board (IRB) Office by phone at (773) 702-2915, or by email at sbs-irb@uchicago.edu

Consent: Participation is voluntary. Refusal to participate or withdrawing from the research will involve no penalty or loss of benefits to which you might otherwise be entitled.

By clicking "Agree" below, you confirm that you have read the consent form, are at least 18 years old, and agree to participate in the research. Please print or save a copy of this page for your records.

I **agree** to participate in the research. I confirm that I am above 18 years of age or older.

I **do NOT agree** to participate in the research. You will be directed to an exit screen.

This survey is directed at university students' social media preferences. It is designed by our research team at the University of Chicago and is administered via partnership with College Pulse.*

Thank you for participating!

*College Pulse is a research and analytics company that specifically aims to understand the attitudes, preferences, and behaviors of today's college students.

How frequently did you use each of the following social media platforms in the past month?

	Not at all	Once	Once a week	Twice a week	Every day
Instagram	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Facebook	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
TikTok	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Twitter	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Please read the instructions of this survey carefully. We will ask you some questions to ensure your understanding of its content and give you a **bonus payment** if you answer correctly!



As mentioned, this survey is directed at university students' social media preferences. **We are studying how much students like you value Instagram.**

After this survey, we will conduct a study in which we will ask **students at your university to deactivate their Instagram accounts for four weeks** in exchange for a **monetary payment**.

"Deactivation" studies like this have been conducted in the past (e.g., by Mosquera et al., 2018 and Allcott et al., 2020) with close to 90% compliance.

Should students deactivate their Instagram accounts, they can go back to using it whenever they want, with their content and network unchanged, but they will then forgo any monetary payment.

To verify that students deactivate their Instagram accounts, we will **periodically visit their profiles and require them to upload screenshots of their app usage.**



The next part of the study involves asking you if you want to deactivate your Instagram account for four weeks in exchange for monetary payment.

For that we will need to collect your Instagram handle so that we can check whether you are active. We would also ask you to upload periodic screenshots of your phone's time use statistics if you are selected.

Would you be willing to

1. **deactivate your Instagram account for four weeks in exchange for monetary payment,**
2. provide your Instagram handle, and
3. provide screenshots of your phone's time use statistics?

Yes

No



How will we verify that selected users deactivate their Instagram accounts?

By periodically visiting their profiles and requiring them to upload screenshots

By requiring them to upload screenshots

By asking them to install an app and requiring them to upload screenshots

By periodically visiting their profiles



For how long will we ask selected users to deactivate their Instagram accounts?

One week

Four weeks

Eight weeks

Ten weeks

G.2.2 Step 0: Practice Good

Before we begin, we will ask you a series of hypothetical practice questions for you to get accustomed to our survey.

Uber

To understand how much people value their Uber account, **we ask university students, including you**, to decide whether to **deactivate their Uber account** for four weeks, in exchange for different monetary payments.

A computer will **randomly select one student from your university** to be eligible for the deactivation study.

Uber

Suppose the next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions offering you different payment scenarios in case you are selected for the deactivation study.

- If you accept any price scenario lower than the computer's offer, we will invite you to the deactivation study and give you the computer's offer.
- If you do not accept any price scenario lower than the computer's offer, we will not invite you to the deactivation study even if you are selected.

This rule means that the higher the amount you require the lower the chance that you will receive the computer's offer.

Therefore, while answering the following questions, **please choose carefully as each answer counts**.

We will now ask you a comprehension question based on the text above:

Which of the following statements is true?

- The higher the amount I require the higher the chance that I receive the computer's offer
- The higher the amount I require the lower the chance that I receive the computer's offer
- The amount I require does not affect the chance that I receive the computer's offer

Uber

Over the next 4 weeks, would you like to take a break from ride-sharing apps?

Which of the following options would you prefer?

I take a break: I deactivate my Uber account AND I receive \$60

I do not take a break: I keep my Uber account active

Uber

Over the next 4 weeks, would you like to take a break from ride-sharing apps?

Which of the following options would you prefer?

I take a break: I deactivate my Uber account AND I receive \$20

I do not take a break: I keep my Uber account active

Uber

Over the next 4 weeks, would you like to take a break from ride-sharing apps?

Which of the following options would you prefer?

I take a break: I deactivate my Uber account AND I receive \$40

I do not take a break: I keep my Uber account active

Uber

According to your answers to the previous questions, you would require a payment worth between \$20 and \$40 to deactivate your Uber account for four weeks.

Do you agree with the above statement about your **valuation**?

Yes

No

G.2.3 Step 1: Valuation Keeping Network

You have now completed the practice section. We will now ask you a series of questions about your **social media preferences**.



To establish appropriate payment amounts for the deactivation study, **we ask university students, including you**, to decide whether to deactivate their Instagram accounts for different monetary amounts.

A computer will **randomly select one student from your university** to be eligible for the deactivation study.



The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions offering you different payment scenarios in case you are selected for the deactivation study.

Therefore, while answering the following questions, **please choose carefully as each answer counts**.



Over the next 4 weeks, would you like to take a break from social media?

Which of the following options would you prefer?

- I take a break:** I deactivate my Instagram account AND I receive \$20
- I do not take a break:** I keep my Instagram account active



Over the next 4 weeks, would you like to take a break from social media?

If you were not receiving payment, which of the following options would you prefer?

- I take a break:** I deactivate my Instagram account
- I do not take a break:** I keep my Instagram account active



According to your answers to the previous questions, you would deactivate your Instagram account for four weeks without monetary payment.

Do you agree with the above statement about your **valuation**?

Yes

No

G.2.4 Step 1: Valuation Removing Network



College Pulse has a panel exceeding 650,000 university students.* We are targeting universities with a high penetration of College Pulse.

We will now ask you to consider two additional options for a large-scale deactivation of Instagram at your university. **One of them will be randomly implemented** if we manage to recruit **more than two-thirds of the students at your university**.

We expect 90% of students to comply with deactivation based on previous studies (e.g., by Mosquera et al., 2018 and Allcott et al., 2020).

*See <https://collegepulse.com>



Option 1 for the large-scale deactivation study:

In collaboration with College Pulse, we will ask students at your university sequentially whether they would like to deactivate their Instagram accounts.

Consider the scenario where we have asked **all participating students at your university** to deactivate their Instagram accounts for four weeks in exchange for a payment.



The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions offering you different payment scenarios in case you are selected for the deactivation study.

Therefore, while answering the following questions, **please choose carefully as each answer counts.**



Over the next 4 weeks, would you like to take a break from social media if we ask all participating students at your university to take a break from social media in exchange for monetary payment?

If you had the choice to **also** deactivate your Instagram account without payment, which of the following options would you prefer?

- I take a break:** I deactivate my Instagram account if the other students deactivate their Instagram accounts
- I do not take a break:** I keep my Instagram account active if the other students deactivate their Instagram accounts



Over the next 4 weeks, would you like to take a break from social media if we ask all participating students at your university to take a break from social media in exchange for monetary payment?

If you had the choice to **also** deactivate your Instagram account, which of the following options would you prefer?

- I take a break:** I deactivate my Instagram account if the other students deactivate their Instagram accounts AND I receive \$120
- I do not take a break:** I keep my Instagram account active if the other students deactivate their Instagram accounts



Over the next 4 weeks, would you like to take a break from social media if we ask all participating students at your university to take a break from social media in exchange for monetary payment?

If you had the choice to **also** deactivate your Instagram account, which of the following options would you prefer?

- I take a break:** I deactivate my Instagram account if the other students deactivate their Instagram accounts AND I receive \$60
- I do not take a break:** I keep my Instagram account active if the other students deactivate their Instagram accounts



Over the next 4 weeks, would you like to take a break from social media if we ask all participating students at your university to take a break from social media in exchange for monetary payment?

If you had the choice to **also** deactivate your Instagram account, which of the following options would you prefer?

- I take a break:** I deactivate my Instagram account if the other students deactivate their Instagram accounts AND I receive \$40
- I do not take a break:** I keep my Instagram account active if the other students deactivate their Instagram accounts



According to your answers to the previous questions, you would require a payment worth between \$40 and \$60 to deactivate your Instagram account for four weeks, if we ask all participating students at your university to deactivate their Instagram accounts.

Do you agree with the above statement about your **valuation**?

- Yes
- No

G.2.5 Step 3: Product Market Valuation



Option 2 for the large-scale deactivation study:

We know how much we would need to pay **every participating student at your university** to deactivate their Instagram accounts for four weeks.

To give everyone an equal chance to decide, we will randomly choose one of the students.

This student's identity will remain anonymous, and they can choose one from the following options:

1. **We ask all participating students with an Instagram account to deactivate it, or**
2. **We keep things as they are.**

If you decide for all participating students to deactivate their Instagram accounts:

- **We will pay the other students the amount they required** and we will **establish your payment, if any, below.**
- The deactivation study will be stopped for everyone only if you go back to using Instagram before the end of the four weeks.
- If the study is stopped early, you will not receive payment and we will pay the other students based on the actual time they spent in the study.
- If someone from the other participating students goes back to using Instagram before the end of the study, they themselves will not receive any payment.



Suppose that you decide for us to ask all participating students to deactivate their Instagram accounts:

Which of the following statements is correct?

- We will force the other students to deactivate their Instagram accounts
- We will pay the other students what they required to deactivate their Instagram accounts
- We will pay the other students more than what they required to deactivate their Instagram accounts



The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions with different payment scenarios in case you are selected.

Therefore, while answering the following questions, **please choose carefully as each answer counts.**



Over the next 4 weeks, would you like students at your university to take a break from social media?

If you were not receiving payment, which of the following options would you prefer?

- Everyone takes a break:** All participating students at my university deactivate their Instagram accounts
- No one takes a break:** All participating students at my university keep their Instagram accounts active



Over the next 4 weeks, would you like students at your university to take a break from social media?

Which of the following options would you prefer?

- Everyone takes a break:** All participating students at my university deactivate their Instagram accounts
- No one takes a break:** All participating students at my university keep their Instagram accounts active AND I receive \$160



Over the next 4 weeks, would you like students at your university to take a break from social media?

Which of the following options would you prefer?

- Everyone takes a break:** All participating students at my university deactivate their Instagram accounts
- No one takes a break:** All participating students at my university keep their Instagram accounts active AND I receive \$180



Over the next 4 weeks, would you like students at your university to take a break from social media?

Which of the following options would you prefer?

- Everyone takes a break:** All participating students at my university deactivate their Instagram accounts
- No one takes a break:** All participating students at my university keep their Instagram accounts active AND I receive \$200



According to your answers to the previous questions, you would forgo a payment above \$200 to have all participating students at your university, including you, deactivate their Instagram accounts for four weeks.

Do you agree with the above statement about your **valuation**?

- Yes
- No

As you have disagreed with the payment amount, you will be redirected to the beginning of the previous section.

This is the last chance to modify your answers.

The **computer** will randomly generate an amount of money to offer you.

We will now ask you a series of questions offering you different payment scenarios in case you are selected.

- If you accept any price scenario lower than the computer's offer, we will give you the computer's offer
- If you do not accept any price scenario lower than the computer's offer, we will not give you the computer's offer even if you are selected.

Therefore, while answering the following questions, **please choose carefully as each answer counts.**



The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions with different payment scenarios in case you are selected.

Therefore, while answering the following questions, **please choose carefully as each answer counts.**



Over the next 4 weeks, would you like students at your university to take a break from social media?

If you were not receiving payment, which of the following options would you prefer?

- Everyone takes a break:** All participating students at my university deactivate their Instagram accounts
- No one takes a break:** All participating students at my university keep their Instagram accounts active



Over the next 4 weeks, would you like students at your university to take a break from social media?

Which of the following options would you prefer?

- Everyone takes a break:** All participating students at my university deactivate their Instagram accounts AND I receive \$100
- No one takes a break:** All participating students at my university keep their Instagram accounts active



Over the next 4 weeks, would you like students at your university to take a break from social media?

Which of the following options would you prefer?

- Everyone takes a break:** All participating students at my university deactivate their Instagram accounts AND I receive \$60
- No one takes a break:** All participating students at my university keep their Instagram accounts active



Over the next 4 weeks, would you like students at your university to take a break from social media?

Which of the following options would you prefer?

- Everyone takes a break:** All participating students at my university deactivate their Instagram accounts AND I receive \$40
- No one takes a break:** All participating students at my university keep their Instagram accounts active



According to your answers to the previous questions, you would require a payment worth between \$40 and \$60 to have all participating students at your university, including you, deactivate their Instagram accounts for four weeks.

Do you agree with the above statement about your **valuation**?

- Yes
- No

G.2.6 Qualitative Questions

What is the percent chance that we will recruit more than two-thirds of the students at your university?

Enter a number between 0 and 100.

Would you prefer to live in a world with or without Instagram?

- I would prefer to live in a world with Instagram
- I would prefer to live in a world without Instagram

You mentioned you would prefer to live in a world without Instagram.

Why do you still use it?

Please rank the following options from your most preferred (1) to your least preferred (3).

- 1 Every student at my university, including me, deactivates Instagram
- 2 I deactivate Instagram and every other student at my university keeps using it
- 3 No one deactivates Instagram

How would you feel if you were the only one who deactivated Instagram and everyone else kept using it?

What fraction of your mutual friends on Instagram are fellow college students? Please enter your response in percent.

To what extent do you agree with the following statements?

	Strongly disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Strongly agree
I use social media immediately upon waking up in the morning.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I would worry that I would be socially isolated if I quit social media.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I would worry about missing out on new trends or memes on social media if I were to quit it.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I feel left out when I see posts or videos about gatherings or events on social media that I wasn't part of.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I feel anxious if I haven't checked social media for a while.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I feel pressure to share or post interesting content on social media.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

How frequently did you share content from social media with your friends in the past month?

	Not at all	Once	Once a week	Twice a week	Every day
Instagram	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Facebook	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
TikTok	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Twitter	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Now some **demographic** questions.

Which of the following describes you more accurately?

Male

Female

Other / Prefer not to say

What is your age?

G.3 Luxury and Vintage Good Survey

University of Chicago

Online Consent Form for Research Participation

Study Number: IRB23 - 1227

Researcher: Leonardo Bursztyn

Description: You will be asked to fill out a short survey. Participation is voluntary and takes around 5 minutes.

Incentives: Upon completion of this survey, you will be compensated by your survey provider.

Risks and Benefits: There are no foreseeable risks associated with this study beyond those involved in answering a survey. The research team cannot and does not guarantee or promise that you will receive any benefits from this study.

Confidentiality: Confidentiality of your research records will be strictly maintained by storing any personally identifiable data in secure accounts that can only be accessed by researchers in this study. After the experiment is over, we will delete any personally identifiable information from our dataset and replace it with an arbitrary participant number. This will allow us to maintain your privacy in all published and written data resulting from the study. Information not containing identifiers may be used in future research or shared with other researchers without your additional consent. If you decide to withdraw, data collected up until the point of withdrawal may still be included in the analysis.

Contacts & Questions: If you have questions or concerns about the study, you can contact the researcher at: bursztyn@uchicago.edu. If you have any questions about your rights as a participant in this research, feel you have been harmed, or wish to discuss other study-related concerns with someone who is not part of the research team, you can contact the University of Chicago Social & Behavioral Sciences Institutional Review Board (IRB) Office by phone at (773) 702-2915, or by email at sbs-irb@uchicago.edu

Consent: Participation is voluntary. Refusal to participate or withdrawing from the research will involve no penalty or loss of benefits to which you might otherwise be entitled. By clicking "Agree" below, you confirm that you have read the consent form and agree to participate in the research. Please print or save a copy of this page for your records.

I agree to participate in the research.

I do NOT agree to participate in the research. You will be directed to an exit screen.

Do you currently own products from luxury brands that you purchased **yourself**? Please tick all that apply.

<input type="checkbox"/> Tiffany & Co.	<input type="checkbox"/> Versace
<input type="checkbox"/> Gucci	<input type="checkbox"/> Louis Vuitton
<input type="checkbox"/> Rolex	<input type="checkbox"/> Chanel
<input type="checkbox"/> Givenchy	<input type="checkbox"/> Balenciaga
<input type="checkbox"/> Swarovski	<input type="checkbox"/> I do not own products from any luxury brands
<input type="checkbox"/> Burberry	<input type="checkbox"/> Other:
	<input type="text"/>
<input type="checkbox"/> Yves Saint Laurent (YSL)	

Do you prefer to live in a world with or without any luxury fashion brands (such as Louis Vuitton, Gucci, Chanel, Yves Saint Laurent (YSL), Balenciaga, Versace, Rolex, Tiffany & Co., Burberry, Givenchy, and Swarovski)?

I prefer to live in a world with luxury fashion brands

I prefer to live in a world without any luxury fashion brands

Do you currently own an iPhone that you purchased yourself?

Yes

No

Which iPhone model do you own?

Do you prefer to live in a world where Apple releases the iPhone every year or every other year?

I prefer to live in a world where Apple releases the iPhone **every year**

I prefer to live in a world where Apple releases the iPhone **every other year**

Are you planning to buy or have you already bought the new iPhone 15 this year?

Yes

No

Do you prefer to live in a world with or without any iPhones?

- I prefer to live in a world with iPhones
- I prefer to live in a world without any iPhones

Which of the following describes you more accurately?

- Male
- Female
- Other / Prefer not to say

What is your age?

What was your TOTAL household income, before taxes, last year?

- \$0 - \$20 000
- \$20 000 - \$50 000
- \$50 000 - \$90 000
- \$90 000 - \$150 000
- \$150 000 - \$200 000
- \$200 000+

Which category best describes your highest level of education?

Some High School

High School Degree

Some College

College Degree

Master's Degree

Doctoral Degree

What racial or ethnic group best describes you?

White

Black or African-American

Hispanic or Latino

Asian or Asian-American

Middle Eastern

Other