Drivers of Digital Attention: Evidence from a Social Media Experiment*

Guy Aridor[†]

December 17, 2022

Abstract

I study demand for social media services by conducting an experiment where I monitor how participants spend time on digital services and shut off access to Instagram or YouTube on their mobile phones. I characterize how participants substitute their time during and after the restrictions, which motivates estimating a model of time usage with inertia. I apply the substitution patterns observed during the restriction period and implied by the model with and without inertia to a relevant market definition exercise to conclude that relevant markets may be larger than those considered by regulatory authorities for social media applications.

Keywords: Social Media, Attention Markets, Field Experiment, Relevant Markets
JEL Codes: L00; L40; L86.

*I am indebted to my advisors Yeon-Koo Che, Tobias Salz, Andrey Simonov, and Michael Woodford for their guidance. I thank Maayan Malter for partnering on the data collection. I thank Matt Backus, Mark Dean, Laura Doval, Andrey Fradkin, Matthew Gentzkow, Garrett Johnson, Ro'ee Levy, Andrea Prat, Silvio Ravaioli, Pietro Tebaldi, and Xiao Xu for their valuable comments. I am grateful to Ro'ee Levy who provided advice on recruitment as well as to NYU, HKUST, and Chicago Booth for providing access to their lab pools. I thank seminar participants at Columbia University, Columbia GSB, Department of Justice, Federal Trade Commission, Hebrew University, HEC Paris, Imperial College London, Microsoft, Northwestern Kellogg, Sciences Po, Stanford GSB, Stanford Digital Economy Seminar, Tel Aviv University, Tel Aviv Coller, Toulouse School of Economics, University of Michigan Ross, and University of Toronto Rotman as well as conference participants at NBER Summer Institute (IT/Digitization), IC2S2, QME conference, MIT Conference on Digital Experimentation, and CEPR Workshop on Media, Technology, Politics & Society for helpful comments. I thank Tyler Redlitz and Rachel Samouha for excellent research assistance. The experimental design was pre-registered in the AEA RCT Registry as AEARCTR-0007256. Finally, I thank the Columbia University Program for Economic Research for providing funding for this research. All errors are my own.

[†]Northwestern University, Kellogg School of Management. Email: guy.aridor@kellogg.northwestern.edu.

1 Introduction

In the past two decades social media has evolved from a niche online tool for connecting with friends to an essential aspect of people's lives. Indeed, the most prominent social media applications are now used by a majority of individuals around the world and these same applications are some of the most valuable companies in the modern day. Due to the sheer amount of time spent on these applications and concentration of this usage on only a few large applications, there has been a global push towards understanding whether and how to regulate these markets (Scott Morton et al., 2019; CMA, 2020).[1] At the heart of the issue is that consumers pay no monetary price to use these applications, which renders the standard antitrust toolkit difficult to apply as the lack of prices complicates the measurement of demand and identification of plausible substitutes.[2] The demand measurement problem is further compounded by the fact that some fraction of usage may be driven by addiction to the applications or, more broadly, various aspects of consumer inertia (Hou et al., 2019; Morton and Dinielli, 2020). This aspect of demand renders it difficult to assess the substitutability between applications as it inflates substitution toward more prominent applications. These two complications have led to substantial difficulties in understanding the core aspects of consumer demand that are crucial for market evaluation and merger analysis.

In this paper I empirically study demand for these applications and characterize substitution patterns between prominent social media applications. I conduct a field experiment where, using parental control software installed on their phone and a Chrome Extension installed on their computer, I continuously track how participants spend time on digital services for a period of 5 weeks. I use the parental control software to shut off access to YouTube or Instagram on their phones for a period ranging from one to two weeks. I explicitly design the experiment so that there is variation in the length of the restriction period and continue to track how participants allocate their time for two to three weeks following the restrictions. The time usage substitution patterns observed during the restriction period allow me to determine plausible substitutes, despite the lack of prices. The extent to which there are persistent effects in the post-restriction period allows me to uncover the role that inertia plays in driving demand for these applications. I use this variation to estimate a demand model with consumer inertia in order to quantify the role that inertia plays in driving substitution patterns towards more prominent applications. I relate both the reduced form restriction

[1] As pointed out by Prat and Valletti (2021), the increased concentration of consumer attention can have ramifications far beyond this market alone since increased concentration in this market influences the ability for firms to enter into product markets that rely on advertising for product discovery.

[2] This issue was at the heart of the Facebook-Instagram and Facebook-WhatsApp mergers. Without prices, regulatory authorities resorted to market definitions that only focused on product characteristics, as opposed to substitution patterns of usage. For instance, Instagram's relevant market was only photo-sharing applications and WhatsApp's relevant market was only messaging applications. This issue continues to play a role in the ongoing FTC lawsuit against Facebook where a similar debate is ongoing.

period substitution results and model estimates to different measures of relevant market definition.

Broader antitrust concerns motivate the following two questions about substitution patterns: what types of activities do participants substitute to and is this substitution concentrated on prominent applications or dispersed among the long tail? I characterize substitution towards other mobile phone applications as well as more broadly towards non-digital activities, but focus primarily on substitution towards other mobile phone applications since these are most plausibly substitutes for advertisers and most relevant from an antitrust perspective.[3] This has been at the heart of regulatory debates with both CMA (2020) and FTC (2021) arguing for narrow relevant market definitions for the largest social media applications.

Within this type of substitution, the most relevant question is whether there is evidence that consumers substitute across application categories and the extent of substitution within category. This has featured prominently in debates between these applications and regulators since the degree to which applications such as YouTube and Instagram are substitutable is important for monopolization claims about Facebook and mergers between different types of applications. Furthermore, the extent of substitution within a category is also relevant as, for instance, FTC (2021) argues that the only relevant substitute for Facebook is Snapchat and not even the broader set of social applications. Thus, it is also important to understand to what extent substitution is concentrated towards more popular applications such as Snapchat, within the vast Facebook ecosystem which spans application categories, or dispersed towards smaller applications in order to assess the magnitude of substitution even within category.

I consider two approaches to assessing the set of plausible substitutes and the resulting relevant market definition. First, I argue that the set of applications that consumers substitute to during the restriction period serves as a conservative market definition since it measures consumer substitution at the "choke" price – the price which is sufficiently high so that no one would use the application at all. Thus, even with zero consumer prices, the product unavailability variation alone allows me to assess the plausibility of claims that applications such as YouTube and Instagram directly compete against each other for consumer attention. Second, I produce model-based estimates of diversion between applications and consider a more precise relevant market definition using these estimates. Both approaches lead to the conclusion that the relevant market for the set of applications is relatively broad compared to those posited by regulatory authories.

I measure average treatment effects of substitution across product categories during the application restriction period in order to implement the first proposed market definition assessment. I manually pair each observed application in the data with the category it is assigned to on the

[3] Lacking detailed data on advertiser spending across applications I cannot estimate the substitution patterns of advertisers. However, substitutability is more plausible across mobile phone applications relative to off phone activities and so I focus primarily on characterizing consumer substitution patterns across mobile phone applications.

Google Play Store. For the Instagram restriction group, I find a 22.7% increase in time spent on other social applications, but also a marginally significant 10.4% increase in time spent on communication applications and a positive, but imprecise, substitution towards entertainment applications. For the YouTube restriction group, I find that there is a null effect of substitution towards other entertainment applications, but also find a 15.1% increase in time spent on social applications. Furthermore, for both restrictions I find a large amount of substitution towards non-digital activities, which suggests that, despite observed substitution towards other mobile applications, they don't fully substitute for the restricted application. While this provides evidence for cross-category substitution, there is a notable asymmetry where blocking Instagram, a social media application, does not lead to statistically significant substitution towards entertainment applications such as YouTube, whereas blocking YouTube, an entertainment application, leads to substitution towards social applications such as Instagram and Facebook.

Pairing these results with the conservative relevant market definition test implies that market definitions ought to span across the application categories between which I observe substitution. I show that, under this market definition, concentration is meaningfully lower relative to only using application categories as the relevant market definition. However, I elicit a subjective measure of how each participant uses the set of prominent social media, entertainment, and communication applications and find that, especially for social media applications, participants use the applications for different reasons ranging from social connection to pure entertainment. This points to the application categories not necessarily capturing the different uses of these applications and partially explaining some of the observed cross-category substitution.

The experimental design further allows me to test for the presence of consumer inertia. This is important to characterize in order to assess the degree of substitutability between applications since the baseline substitution patterns are likely inflated towards more prominent applications and over-estimate the degree of substitution towards these applications, whereas from an antitrust perspective longer run substitution patterns are more relevant and parsing out inertia enables a more plausible estimate of this. I find the following experimental experience that is consistent with inertia playing a role in demand for these applications. There is a persistent reduction in time spent on the restricted applications and this is primarily driven by the participants that had the two week restriction. For the Instagram restriction, the two week restriction group reduced average daily usage relative to the control group by 5 minutes and had a similar reduction relative to the one week restriction group. A survey sent after the study indicates that this reduction in time spent persists even a month following the conclusion of the study. For the YouTube restriction, there is suggestive evidence of a similar difference between the one and two week restriction group, but the resulting difference in average daily usage is not statistically significant. However, I find that participants in the YouTube restriction spent more time on applications installed during the restriction period relative to the control group and persisted to use these applications even in the post-restriction period.

I use the insights from the experimental results to construct a discrete choice model of time usage with inertia between prominent social media and entertainment applications. I model inertia by including past usage into consumer utility similar to state-dependent demand estimation models (Dubé, Hitsch and Rossi, 2010; Bronnenberg, Dubé and Gentzkow, 2012). I directly observe the currently installed set of applications for each participant which dictates their consideration set. Furthermore, I incorporate the heterogeneity in subjective usage of the applications and number of connections on a given application into the utility function in order to capture the preference heterogeneity indicated by the experimental results. The model also exploits the granular time usage data to have a flexible outside option that varies across time.

I use the resulting demand model to compute second-choice diversion ratios – diversion with respect to a change in availability. The diversion ratio from application i to application j is defined as the fraction of sales / consumption that gets diverted from application i to application j as a result of a change in price / quality / availability of application i. Diversion ratios play a prominent role in the current US horizontal merger guidelines for measuring possible unilateral effects and can further be used for relevant market definition. I compute the resulting set of diversion ratios with and without inertia, which shows that shutting off the inertia channel softens diversion towards more prominent applications and increases diversion towards less popular applications. This highlights that without accounting for inertia, estimates of substitutability are biased towards more prominent applications. Furthermore, the magnitude of the inertia channel itself is large as shutting down the inertia channel decreases the share of time spent on the prominent social media and entertainment applications by nearly 30%.

In order to further assess the quantitative importance of the inertia channel, I apply the resulting diversion estimates to a simple relevant market definition exercise: for a given application x, sort the other applications in terms of their diversion from x and add them into the relevant market until the sum crosses a threshold τ. The results are consistent with the reduced-form results by highlighting that, even in the baseline, the relevant market is relatively broad and not narrowly defined. Furthermore, it suggests that the inertia channel is quantitatively large enough that by shutting it off the relevant market for each application meaningfully increases. This also suggests that regulatory policy targeting this inertia channel, e.g. digital addiction policy, can be a useful policy tool to increase competition in this market. Finally, I compare these results to a counterfactual where I simulate choices if participants always had the prominent social media and entertainment applications installed and show that the inertia channel plays a larger role than the lack of adoption in driving usage.

More broadly, this paper highlights the usefulness of product unavailability experiments paired with granular consumer-specific tracking for demand and merger studies between digital goods. I exploit the insight that digital goods enable individual level, randomized controlled experiments of product unavailability that are difficult to conduct with other types of goods and in other markets. Furthermore, the same digital tracking that these firms use to collect fine-grained data on individual consumers can similarly be collected by researchers and regulators. These experiments enable causal estimates of substitution patterns and identify plausible substitutes even when consumers pay no prices and can be used to estimate the relevant portions of consumer demand that are difficult to estimate using only observational data. As a result, they serve as a practical and powerful tool for antitrust regulators in conducting relevant market definition and merger assessments in digital markets.

The results provided in the paper have to be interpreted with caution for several reasons. First, the sample of participants that I conduct the experiment on draws primarily from university lab pools, which makes them plausibly less selected on usage of the applications of interest relative to other recruitment options, but is not demographically representative. Second, the intervention targets a particular individual's usage of applications, while plausibly holding fixed the rest of the social network on the applications. Thus, the estimated diversion ratios capture the partial equilibrium substitution patterns and provide an assessment for what types of applications participants view as being substitutable. However, by not capturing general equilibrium effects that would occur from an extended shutdown of an application for all consumers, my estimates are a lower bound on the magnitude of substitution. Despite these limitations, the qualitative implications of the results – such as the large presence of consumer inertia and cross-category / offline substitution – are likely not impacted by them, though the quantitative estimates of diversion would likely be different once these effects are taken into account.

2 Related Work

This paper contributes to four separate strands of literature, which I detail below.

Economics of Social Media: The first is the literature that studies the economic impact of social media. Methodologically my paper is closest to Brynjolfsson, Collis and Eggers (2019); Allcott et al. (2020); Mosquera et al. (2020) who measure the psychological and economic welfare effects of social media usage through restricting access to services. Allcott et al. (2020); Mosquera et al. (2020) restrict access to Facebook and measure the causal impact of this restriction on a battery of psychological and political economy measures. Allcott et al. (2020) similarly studies substitution and post-restriction reduction in usage through self-reported time estimates. Brynjolfsson, Collis and Eggers (2019) measures the consumer surplus gains from free digital services by asking participants how much they would have to be paid in order to give up such services for a period of time. This paper utilizes a similar product unavailability experiment, but uses this variation in order to precisely measure substitution patterns and relate them to relevant issues in antitrust as opposed to quantifying welfare effects.

A concurrent paper that is also methodologically related is Allcott, Gentzkow and Song (2021). They utilize similar tools to do automated and continuous data collection of phone usage.[4] They focus on identifying and quantifying the extent of digital addiction by having separate treatments to test for self-control and habit formation. I argue that my experimental design also enables me to understand the persistent effects of the restriction, which I use to identify a demand model of time usage with inertia. While my experiment does not allow me to identify the precise mechanism behind this inertia effect, I rely on Allcott, Gentzkow and Song (2021) to argue that the most likely possible mechanism is tied to digital addiction. Thus, I view Allcott, Gentzkow and Song (2021) as being complementary to my work as I focus on the competition aspect between these applications, but also find patterns consistent with their results.[5]

Finally, there is a burgeoning literature on the broader economic and social ramifications of the rise of social media applications. Collis and Eggers (2022) study the impact of limiting social media usage to ten minutes a day on academic performance, well-being, and activities and observes similar substitution between social media and communication applications. The broader literature has focused on political economy issues associated with social media (Bakshy, Messing and Adamic, 2015; Corrigan et al., 2018; Enikolopov, Makarin and Petrova, 2020; Levy, 2021) as well as its psychological impact (Levy, 2016; Burke and Kraut, 2016; Kuss and Griffiths, 2017; Bailey et al., 2020; Baym, Wagman and Persaud, 2020; Braghieri, Levy and Makarin, 2022). Relative to this broad literature, my paper focuses more prominently on antitrust related issues in these markets.

Product Unavailability and State-Dependent Demand Estimation: The second is the literature in marketing that studies brand loyalty and, more broadly, state-dependent demand estimation. The discrete choice model of time usage closely follows the formulation in this literature where past consumption directly enters into the consumer utility function and the empirical challenge is to disentangle the inertia portion of utility from preference heterogeneity (Shum, 2004; Dubé, Hitsch

[4] An important antecedent of this type of automated data collection is the "reality mining" concept of Eagle and Pentland (2006) who first used mobile phones to comprehensively digitize activities done by experimental participants. Allcott, Gentzkow and Song (2021) relies on a custom-made application, whereas the primary data collection done in my paper relies on a (relatively) cheap, publicly available, parental control application and an open source Chrome extension which is more accessible to other researchers. Furthermore, unlike Allcott, Gentzkow and Song (2021), I can comprehensively track substitution towards other devices without having to rely on self-reported data.

[5] In the theory literature, Ichihashi and Kim (2021) study competition between addictive platforms where platforms trade off application quality for increased addictiveness. Hoong (2021) also studies the role of self-control issues in driving usage and commitment devices to reduce usage through a randomized experiment.

and Rossi, 2010; Bronnenberg, Dubé and Gentzkow, 2012; Simonov et al., 2020). I consider that consumers have a usage stock that enters directly into the utility function, which I interpret as inertia that drives usage of the applications.

Relative to this literature, I exploit the fact that I conduct an experiment and induce product unavailability variation as a shock to the usage stock in order to identify this portion of consumer utility. Conlon and Mortimer (2013, 2021); Conlon, Mortimer and Sarkis (2022); Raval, Rosenbaum and Wilson (2022) explore the value of product unavailability in identifying components of consumer demand. In this paper I use this variation to understand the role of inertia as well as provide extensive reduced-form and qualitative analysis of consumer demand resulting from individual-level availability changes. Finally, Goldfarb (2006) studies a natural experiment of product unavailability due to website outages in order to understand the medium term effects of inertia on overall usage.

Attention Markets: The third is the literature that studies "attention markets" (see Calvano and Polo (2020), Section 4 for an overview). An important modeling approach taken in the theoretical literature, starting from Anderson and Coate (2005) and continuing in Ambrus, Calvano and Reisinger (2016); Anderson, Foros and Kind (2018); Athey, Calvano and Gans (2018) is modeling the "price" faced by consumers in these markets as the advertising load that the application sets for consumers. In the legal literature a similar notion has emerged in Newman (2016); Wu (2017) who propose replacing consumer prices in the antitrust diagnostic tests with "attention costs." Relative to the theoretical literature in economics, Newman (2016); Wu (2017) interpret these "attention costs" as being broader than just advertising quantity and including reductions in application quality. I use this notion to interpret product unavailability as being informative about the relevant market definition exercise through observing substitution at the choke value of attention costs.

Mobile Phone Applications: The fourth is the literature that studies the demand for mobile applications, which typically focuses on aggregate data and a broad set of applications. This paper, on the other hand, utilizes granular individual level data to conduct a micro-level study of the most popular applications. Ghose and Han (2014) study competition between mobile phone applications utilizing aggregate market data and focus on download counts and the prices charged in the application stores, as opposed to focusing on time usage. Han, Park and Oh (2016); Yuan (2020) study the demand for time usage of applications in Korea and China respectively building off the multiple discrete-continuous model of Bhat (2008). Relative to these papers there are two important differences. First, I exploit the granularity of the data to model time allocation as a panel of discrete choices instead of a continuous time allocation problem. Second, I exploit my experimental variation to study the role of inertia as opposed to complementarity / substitutability.

3 Experiment Description and Data

3.1 Recruitment

I recruit participants from a number of university lab pools in March 2021, including the University of Chicago Booth Center for Decision Research, Columbia Experimental Laboratory for Social Sciences, New York University Center for Experimental Social Science, and Hong Kong University of Science and Technology Behavioral Research Laboratory. A handful of participants came from emails sent to courses at the University of Turin in Italy and the University of St. Gallen in Switzerland. Furthermore, only four participants were recruited from a Facebook advertising campaign.[6] The experimental recruitment materials and the Facebook advertisements can be found in Online Appendix A.1. Participants earned $50 for completing the study, including both keeping the software installed for the duration of the study as well as completing the surveys. Participants had an opportunity to earn additional money according to their survey responses if they were randomly selected for the additional restriction.

Preliminary data indicated that there was a clear partition in whether participants utilized social media applications such as Facebook, Instagram, Snapchat, and WhatsApp as opposed to applications of less interest to me such as WeChat, Weibo, QQ, and KakaoTalk.[7] As a result, the initial recruitment survey ensured that participants had Android phones as well as used applications such as Facebook/Instagram/WhatsApp more than applications such as WeChat/Weibo/QQ/KakaoTalk. I had 553 eligible participants that filled out the interest survey. The resulting 553 eligible participants were then emailed to set up a calendar appointment to go over the study details and install the necessary software. This occurred over the period of a week from March 19th until March 26th. At the end, 410 participants had agreed to be in the study, completed the survey, and installed the necessary software.

There are two points of concern that are worth addressing regarding recruitment. The first is whether there is any selection into the experiment due to participants seeking limits on their use of social media applications. In the initial recruitment it was emphasized that the purpose of the study was to understand how people spend their time with a particular focus on the time spent in their digital lives, in order to dissuade such selection into the experiment. Once the participants had already registered, they were informed about the full extent of the study. However, they were still broadly instructed that the primary purpose of the study was to understand how people spend

[6] While these participants only ended up making up a small fraction of overall participants, in order to ensure that the nature of selection was consistent across the different recruiting venues the Facebook advertisements were geographically targeted towards 18-26 year olds that lived in prominent college towns (e.g. Ann Arbor in Michigan, Ames in Iowa, Norman in Oklahoma). This was to ensure that there was similar demographic selection as those implicitly induced by recruitment via university lab pools.

[7] This was from another experiment that collected mobile phone data from the same participant pool.

their time and that they may face a restriction of a non-essential phone application. The precise application that would be restricted was not specified in order to further ensure there were no anticipatory effects that would bias baseline usage. The second is that I do not exclusively recruit from Facebook or Instagram advertisements as is done in several other studies (e.g. Allcott et al. (2020); Levy (2021); Allcott, Gentzkow and Song (2021)), but instead rely on university lab pools. This leads to an implicit selection in the type of participants I get relative to a representative sample of the United States (e.g. younger, more educated), however it does not induce as much selection in the intensity of usage of such applications that naturally comes from recruiting directly from these applications. For a study such as this some degree of selection is inevitable, but in this case I opted for selection in terms of demographics instead of selection on intensity of application usage as for a study on competition this was more preferable.

3.2 Automated Data Collection

The study involved an Android mobile phone application and a Chrome Extension. Participants were required to have the Android mobile phone application installed for the duration of the study and were recommended to install the Chrome Extension. Despite being optional, 349 of the participants installed the Chrome Extension. It is important that I collect objective measures of time allocations for the study as subjective measurements of time on social media are known to be noisy and inaccurate (Ernala et al., 2020).

The Android mobile phone application is the ScreenTime parental control application from ScreenTime Labs.[8] This application allows me to track the amount of time that participants spend on all applications on their phone, the exact times they're on the applications, and the set of installed applications on the phone. Furthermore, it allows me to restrict both applications and websites so that I can completely restrict usage of a service.[9] This application is only able to collect time usage data on Android, which is why I only recruit Android users.[10] I pair the data with manually collected data on the category of each application pulled from the Google Play Store.

The Chrome Extension collects information on time usage on the Chrome web browser of the desktop/laptop of participants.[11] All the restrictions for the study are only implemented on the mobile phone so that participants have no incentive to deviate to different web browsers on their computers at any point during the study. The software is setup with the participants over Zoom

[8]For complete information on the application see `https://screentimelabs.com`.

[9]For instance, if I want to restrict access to Instagram then it's necessary to restrict the Instagram application as well as `www.instagram.com`. It does this by blocking any HTTP requests to the Instagram domain, so that the restriction works across different possible browsers the participant could be using.

[10]Technological limitations of iOS at the time of the experiment prevent similar types of data collection on iPhones.

[11]The source code for the Chrome Extension is available here: `https://github.com/rawls238/time_use_study_chrome_extension`. The extension is modified and extended based off David Jacobowitz's original code. Some participants had multiple computers (e.g. lab and personal computers) and installed the extension on multiple devices.

where they were instructed that the restriction was only on the phone and they should feel free to use the same service on the computer if they wished to do so. Thus, it was important that participants did not feel as though they should substitute between web browsers on the computer as this would lead me to not observe their true computer usage. The final data that I make use of from the extension are time data aggregated at the daily level as well as time period data (e.g. 9:50 - 9:55, 10:30-10:35 on Facebook). The full details on the data collection software can be found in Appendix A.

3.3 Survey Data

In order to supplement the automated time usage data, I elicit additional information via surveys. The surveys allow me to validate the software recorded data, to get information about how participants spend time on non-digital devices, and to elicit qualitative information about how participants use the set of prominent social media and entertainment applications.

Baseline Survey: The first is the baseline survey that participants complete at the beginning of the study. This survey is intended to elicit participants' perceived value and use of social media applications as well as basic demographic information. The full set of questions is provided in Online Appendix A.2. The main question which requires additional explanation and is crucial for the participants' incentives is that I elicit the monetary value that participants assign to each application using a switching multiple price list between \$0 and \$500 (Andersen et al., 2006). This elicitation is incentive-compatible since the participants are made aware that, at the end of the study period, two participants will have one application and one offer randomly selected to be fulfilled and thus have an additional restriction beyond the main portion of the study.[12] For the additional survey questions, I will refer them as relevant throughout the text.

Weekly Surveys: Every week throughout the study there are two weekly surveys that participants complete. The first is sent on Thursdays, which contains a battery of psychology questions and was part of the partnership for this data collection and not reported on in this paper.[13] The second is sent on Saturday mornings and asks participants to provide their best guess at how much time they are spending on activities off their phones. It is broken down into three parts: time spent on applications of interest on other devices, time spent on necessities off the phone, and time spent on leisure activities off the phone.

Endline Survey: The endline survey contains the following questions geared towards understanding participants' response to the restrictions. The goal is to try to disentangle the mechanisms at play in potential dynamic effects of the restrictions. The questions are all multiple choice questions

[12] I do not directly use the answer to this question in the analysis, but mainly use it to provide additional incentives for participation in the study.

[13] The questions that participants answered are presented with the survey instruments in Online Appendix A.2.

that ask how participants think they reallocated their time during the week of the restrictions and how they think their time spent after the restrictions changed relative to before the restrictions. The full details of the questions and possible responses can be found in Online Appendix A.3.

One Month Post-Experiment Survey: I send the participants a survey one month following the conclusion of the main study period. They are told that if they fill out the survey they will have an opportunity to receive a $100 Amazon Gift Card, but it is separate from the experimental payment. The survey asks if they think they are spending a lot less, somewhat less, similar, somewhat more, or a lot more time compared to the pre-experiment levels of usage on their phone, social media in general, and each of the applications of interest. There are also a number of psychology questions asked in the survey, which I do not report here.

3.4 Experiment Timeline

The experiment timeline is as follows. There is an initial week where the software is set up on the devices and I remove participants where the software does not work at all with their phone. During this week we meet all the participants on Zoom to ensure the software is working properly and they understand the extent of data collection done in the study.

After all of the participants have the software set up on their devices, there is a week where I collect baseline, pre-restriction, time usage data. Following this, there is a two week restriction period, but some participants have no restrictions at all or restrictions that last only a week. Participants do not know whether they will have a restriction at all or which applications I target for the restrictions beyond the fact that it will be a non-essential social media or entertainment application. They are only informed of the restriction and its duration via SMS two hours before the restriction went into effect at 11:59 PM on Friday night so that they have limited time to anticipate the restriction.

After the restrictions, there are two weeks where I collect time allocations when there are no restrictions, so that I can measure any persistent effects on behavior for the participants. Finally, the participants complete the endline survey and then, to ensure a degree of incentive compatibility for the WTA elicitations, two participants are randomly selected and potentially have an additional week of restriction depending on their survey responses and the randomly selected offer. The following summarizes the timeline:

- March 19th - March 26th: Complete baseline survey and install software
- March 27th- April 2nd: Baseline usage period
- April 3rd - April 17th: Restriction period

- April 18th - May 2nd: Post-Restriction period
- May 3rd - May 10th: Additional restriction for two participants

3.5 Experimental Restrictions

Figure 1: Experiment Timeline

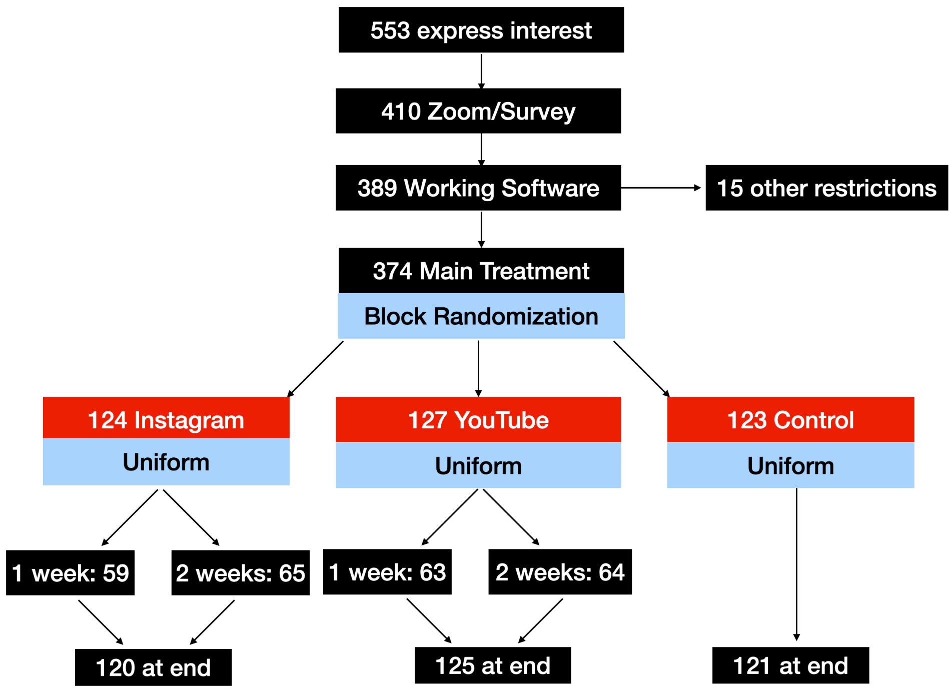

For the main experimental intervention, I restrict to participants that make use of either YouTube or Instagram. From the original 410 participants, 21 had phones that were incompatible with the parental control software and so were dropped from the study. There were 15 participants that did not use either YouTube or Instagram and so were given idiosyncratic applications restrictions.[14] The remaining 374 of the participants are the primary focus – 127 of which have YouTube restricted, 124 of which have Instagram restricted, and 123 which serve as a control group. For participants in the Instagram treatment, 59 and 65 participants have it restricted for one and two weeks respectively. For participants in the YouTube treatment, 63 and 64 have it restricted for one

[14]For most participants in this group this restriction comprised of Facebook or WhatsApp, but for some subset of participants this restriction was Twitch, Twitter, or Facebook Messenger.

and two weeks respectively. There was minimal attrition from the experiment with only 2 participants from the control group, 2 participants from the YouTube restriction group, and 4 participants from the Instagram restriction group dropping from the experiment – in most cases due to reasons orthogonal to treatment (e.g. getting a new phone, tired of surveys). The experimental timeline, treatment assignments, and participant attrition are summarized in Figure 1.

Table 1: Summary Statistics on Usage

Application	Phone (Mean)	Phone (Median)	Computer (Mean)	Computer (Median)	Phone Users	Computer Users
YouTube	48.71	17.79	32.74	11.50	334	266
Instagram	29.82	19.00	4.05	0.43	295	86
WhatsApp	26.57	15.54	8.52	6.07	268	6
Facebook	21.90	7.36	5.21	1.57	234	176
Messenger	10.47	1.96	13.96	6.21	208	32
Snapchat	9.30	3.86	0.00	0.00	151	0
Reddit	21.62	5.36	7.73	1.00	138	127
Twitter	13.41	3.79	6.79	0.86	134	93
TikTok	50.71	28.86	0.95	0.36	68	12

Notes: Each row reports the statistics for the specified application. I report average daily minutes spent during the baseline period for participants that I observe using the application. Columns 1 and 2 report the mean and median daily time on the phone conditional on observed usage on the phone during the baseline. Columns 3 and 4 report the mean and median daily time on the computer conditional on observed usage on the computer during the baseline. Columns 5 and 6 report the total users of the application of the phone and computer.

In order to ensure that the experimental groups are balanced on usage of the applications of interest, I employ block randomization utilizing the baseline usage data from March 27th until April 1st. I categorize the quartile of usage for Instagram and YouTube for each participant and assign each participant into a block defined as the following tuple: (Instagram quartile, YouTube quartile). Within each block, I determine the treatment group uniformly at random (Instagram, YouTube, Control) and then again to determine whether the restriction is one or two weeks. The resulting distribution of usage across the treatment groups for the applications of interest can be found in Figure A2. It shows that the resulting randomization leads to balanced baseline usage between the groups both on the restricted applications as well as other social media applications. Furthermore, the average time spent on the applications of interest is displayed in Table 1 with the full set of descriptive statistics on time allocations in Appendix B. Finally, in order to get additional power for my experimental estimates, I will sometimes pool data with the pilot experiment that was conducted between 9/29/2020 and 12/4/2020 with details of this study provided in Online Appendix A.5.

4 Experimental Results

In this section I analyze the substitution patterns of time allocations throughout the study period. I characterize what applications are considered substitutes for the restricted applications by measuring substitution during the restriction period and relate these substitution patterns to a conservative measure of relevant market definition. I then explore the extent to which there were persistent effects of the restriction by investigating how time allocations differ after the treatment period relative to before it. The insights from this section will be used to guide the demand model estimated in Section 5.

4.1 Empirical Specification

The primary empirical specification that I utilize to estimate the average treatment effect of the experimental interventions is as follows, with i representing a participant and j representing an application / category:

$$Y_{ijk} = \beta T_i + \kappa X_i + \gamma Y_{ij,-1} + \alpha_t + \epsilon_{ijk} \tag{1}$$

where Y_{ijk} represents the outcome variable of interest k weeks after their restriction, $Y_{ij,-1}$ represents the outcome variable of interest during the baseline period (i.e. the first week), T_i represents a treatment dummy, X_i represents a dummy variable for the block participant i was assigned to, and α_t denotes week fixed effects. The main parameter of interest is β; $Y_{ij,-1}$ controls for baseline differences in the primary outcome variable and X_i controls for the block assigned to the participant in the block randomization, which is standard for measuring average treatment effects of block randomized experiments (Gerber and Green, 2012).

For analyzing substitution patterns during the restriction period, I consider Y_{ijk} as the average daily time spent on applications / categories during the days when the participant's software was active and logging data. I focus on these outcome variables during the first week of the restriction.[15] Due to this, I omit the week fixed effects and report heteroskedasticity-robust standard errors. For this analysis I will pool data, when available, with the pilot study that included two separate restriction periods for different subsets of participants and was specifically designed to measure restriction period substitution.[16] In this case, I additionally control for the experimental period as well as cluster standard errors at the participant level. When I consider multiple weeks of usage, as in subsection 4.3, I include week fixed effects and cluster standard errors at the participant level. I also consider Y_{ijk} as the number of newly installed applications, but for this outcome variable, I do not have any baseline data and so estimate the specification omitting the baseline usage term.

I am interested in not just the average treatment effects, but also effects across the distribution since one might imagine that power users of an application or category would respond differently than infrequent users at the baseline. As a result, I also estimate quantile treatment effects using the same specification with a quantile regression since the fact that treatment status is exogenous

[15] This enables consideration of the same substitution interval across all participants.
[16] For the details on the pilot experiment see Online Appendix A.5.

allows for identification of the conditional QTE with a quantile regression (Abadie, Angrist and Imbens, 2002). Finally, Figure A3 indicates that the distribution of phone usage is skewed, which motivates me to consider the specifications in both logs and levels. In order to accommodate occasional zeros in my data, I use the inverse hyperbolic sine transform in lieu of logs, which leads to a similar interpretation of coefficient estimates (Bellemare and Wichman, 2019).

4.2 Time Substitution During the Restriction Period

4.2.1 Conceptual Framework

There are a wide range of possible activities that participants could substitute towards and it is challenging to define the precise substitution patterns that are most relevant to the question of consumer demand and merger analysis. There are two broad questions of interest that guide the analysis. The first is what *types of activities* do participants substitute to and the second is *how dispersed* across different applications are the substitution patterns. These questions are at the heart of the debate about monopolization arguments surrounding Facebook and, more generally, in merger evaluation between applications in this market.

Substitutable Activities: A directly relevant question to the ongoing debate between Facebook and regulators is which types of applications are most substitutable for the restricted applications. For instance, in CMA (2020) Facebook contends that it competes with a broad range of applications that compete for consumer time such as YouTube, which is not traditionally considered a social media application, whereas regulators contend that the most relevant competitors are other social media applications such as Snapchat (e.g. see FTC (2021)). One of the challenges underlying this debate has been the lack of prices in these markets as standard market definition tests rely on understanding substitution with respect to price. Despite the lack of prices, the theoretical literature on two-sided media markets (starting from Anderson and Coate (2005)) and the legal literature (Newman, 2016; Wu, 2017) have noted that in these markets consumers face implicit costs on their time and attention that are direct choice variables for the application. This indicates that one alternative harm in lieu of higher prices is an increased cost on consumer attention, which can take the form of increased advertising load or decreased quality.[17]

Under this interpretation, the substitution observed during the restriction period is a limit case of taking "attention costs" to their choke values where no one would consume the application. Thus, it can serve as a conservative test of substitutability and, in particular, can function as a conservative market definition – only including the applications and activities that are at all substitutable. This has appeal as a tool for practitioners as well since, in practice, variation in "attention

[17] Newman (2016); Wu (2017) propose modifications of the standard Small but significant and non-transitory increase in price (SSNIP) test explicitly considering this harm in lieu of the standard price test.

costs" is substantially more ambiguous and difficult to come by relative to price variation in other markets. Furthermore, experiments such as the one analyzed in this paper are feasible due to the nature of digital goods.[18]

I primarily focus on substitution within the set of mobile applications since these are most relevant for advertiser substitution, but also characterize substitution towards non-digital activities. Since the default approach taken by regulators has been to consider only applications within the same application category as relevant substitutes, I use the categories assigned to the applications in the Google Play Store and characterize substitution across these different application categories during the restriction period. Summary statistics of the time spent on different categories, displayed in Table A3, indicate that the vast majority of time is spent on social, entertainment, and communication categories and, as such, I measure substitution across these categories. Given this, the empirical framing of the YouTube vs. Facebook question is whether there is substitution towards social applications during the YouTube restriction and substitution towards entertainment applications during the Instagram restriction. If I observe no cross-category substitution during the restriction period, then the implication is that smaller increases in "attention costs" would similarly not lead to considerable substitution between these categories. If I do observe cross-category substitution, then it only says that such a market definition is not entirely unreasonable.

Substitution Dispersion: Another important question is the extent to which substitution is concentrated towards a small number of prominent applications or dispersed among the long tail of applications. This captures a different dimension of competition relative to category substitution. This is because it focuses on understanding whether the set of substitutable applications are prominent applications that are likely more attractive to advertisers relative to smaller applications in the long tail. Furthermore, with the data collected during the study, I am able to observe whether participants actively seek out new applications in the long tail, indicating that the presence of these applications prevents this search process and that participants are unsure about appropriate substitutes. For instance, a participant that uses YouTube to keep up with the news or to get trading advice may not have a readily available substitute on their phone and go search in the Google Play Store for a new application if they are restricted from YouTube.

4.2.2 Category Market Definition and Cross-Category Substitution

Cross-Category Substitution: I test the extent of cross-category substitution by measuring the average treatment effect of time substitution towards other categories as a result of the restriction.

[18] Even without directly implemented experiments, natural experiments caused by product outages would induce similar variation and enable similar estimates. For example, extended outages such as the Facebook, WhatsApp, Messenger, and Instagram outage on 10/4/2021 could be utilized to a similar extent, `https://www.nytimes.com/2021/10/04/technology/facebook-down.html`. However, these outages would need to be sufficiently long in order to capture meaningful substitution.

I consider the effects of each restriction on category usage separately and report the results of the analysis pooled with data from the pilot experiment.[19] For these results I focus my interpretation on the log specification as, due to the skewed distribution of usage, this is more representative of the average participant's behavior and is not driven by the most intensive users of the applications.

Table 2: Instagram Category Substitution

	Dependent variable:					
	Social	Social (No IG)	Communication	Entertainment	Other	Overall Phone Time
	(1)	(2)	(3)	(4)	(5)	(6)
Category Time - Pooled	-18.557^{***}	4.202^{*}	3.223	-0.607	-3.318	-16.336^{*}
	(3.100)	(2.424)	(2.769)	(3.872)	(4.093)	(9.081)
asinh(Category Time) - Pooled	-0.594^{***}	0.227^{***}	0.104^{*}	0.071	-0.037	-0.047
	(0.100)	(0.076)	(0.057)	(0.098)	(0.064)	(0.048)

Table 3: YouTube Category Substitution

	Dependent variable:					
	Social	Communication	Entertainment	Entertainment (No YT)	Other	Overall Phone Time
	(1)	(2)	(3)	(4)	(5)	(6)
Category Time - Pooled	3.989	-2.566	-46.685^{***}	-3.608	-4.277	-51.381^{***}
	(2.909)	(3.346)	(5.686)	(2.917)	(4.621)	(11.282)
asinh(Category Time) - Pooled	0.151^{**}	-0.041	-1.484^{***}	0.049	-0.054	-0.154^{***}
	(0.067)	(0.051)	(0.123)	(0.112)	(0.063)	(0.045)

*p<0.1; **p<0.05; ***p<0.01

Notes: These tables report the average treatment effect of average daily minutes spent on applications in different categories during the Instagram and YouTube restrictions respectively. I only consider participants with software active at least 3 days in the baseline and treatment periods. In Table 2, the columns show time spent on social, social (without Instagram), communication, and entertainment. In Table 3, the columns show time spent on social, communication, entertainment, entertainment (without YouTube), other categories, and overall phone time respectively. The entertainment category includes applications marked as entertainment or video players/editors. The column with entertainment (without YouTube) aggregates entertainment time excluding time spent on YouTube, both in the baseline and treatment periods and similarly for social (without Instagram). The estimates display the primary specification estimated on data pooled from the main experiment and the pilot experiment. The reported standard errors for these regressions are clustered standard errors at the participant level.

Table 2 displays the results for the Instagram restriction. The overall amount of time spent on all social applications drops across all specifications (column 1), but the time spent on non-Instagram social applications increases by 22.7% (column 2). This means that there was considerable substitution towards other social applications, but not enough to entirely counteract the loss of Instagram. Column (3) indicates that there is some cross-category substitution to communication applications with the logs specification pointing to a marginally significant 10-12% increase in time spent on such applications. This is consistent with the qualitative evidence from the participants reported in Online Appendix C. For instance, one participant stated *"Instagram was restricted for me and because I mainly use it as a communication app, I was not significantly affected. I just used regular text, video call, and Snapchat to keep up socially."* I observe a fairly precise null

[19] Table A6 and Table A7 present the full set of results, including using only the data from main experiment, and specifications studying changes in market shares. The reported results are consistent with those from these specifications.

result for substitution from Instagram to other applications, but find a positive, though statistically insignificant, increase in substitution to entertainment applications.

Table 3 displays the results for the YouTube restriction. Similar to the results for the Instagram restriction, there is a sharp decrease in own-category time during the restriction period (see column 1). However, unlike the results of the Instagram restriction, there is a precise null of substitution towards other applications within the same category (see column 4). Column (1) points to an increase in time spent on social applications with a 15.1% increase in time spent on these applications, while columns (3) and (5) suggest little increase in time spent on communication and other applications. Finally, Figure A7 displays the effects of the restriction along the entire distribution and shows that the own-category substitution for both applications is upward sloping across deciles, indicating that more intensive overall users of social media and entertainment applications respectively were more likely to look for close substitutes.

Table 4: Stated Activities

Application	Entertainment	Keep up with Friends	Communication	Get Information	Shopping
Facebook	0.26	0.36	0.14	0.20	0.04
Messenger	0.01	0.08	0.88	0.02	0.02
Instagram	0.37	0.47	0.08	0.07	0.01
YouTube	0.78	0.002	0.002	0.22	0.002
TikTok	0.92	0.02	0.05	0.02	0.0
WhatsApp	0.01	0.06	0.92	0.02	0.0
Twitter	0.22	0.03	0.06	0.67	0.01
Snapchat	0.09	0.31	0.58	0.02	0.0
Reddit	0.38	0.0	0.02	0.60	0.01
Netflix	0.97	0.004	0.01	0.02	0.004

Notes: Each row reports the stated activities for the specified application. The cells report the proportion of participants who use the application and report using the application for the column purpose.

Survey Evidence of Cross-Category Substitution: The presence of cross-category substitution and the asymmetry of substitution patterns across the restriction groups requires further inquiry. One possible explanation is that even for applications in the same category, participants use them for different purposes. Table 4 displays the self-reported purpose for using the most prominent social media and entertainment applications, which displays a clear pattern indicating that applications in the social category are used for different purposes. For instance, TikTok is primarily used for entertainment purposes, Twitter for getting information, Snapchat for communication, and Facebook/Instagram's usage is spread across entertainment, keeping up with friends, getting information, and communication. These patterns are broadly consistent with the responses for the hypothetical switching question asked in the baseline survey (see Table A5). The fact that the uses of the applications are heterogeneous and intersect with applications that are not in the same application category therefore helps to understand the observed asymmetry. This is since, if participants view applications such as Instagram or TikTok as primarily for entertainment, then it's not surprising that we observe substitution from an entertainment application such as YouTube towards these social applications. It further suggests a broader issue with using the functional application categories as a crude measure of substitutability as content is personalized to consumer tastes, enabling the same application to serve different purposes for different consumers.

Table 5: Herfindahl–Hirschman Index Across Market Definitions

	Social	Entert.	Comm.	Social + Entert.	Social + Comm.	Social + Entert. + Comm.
Current Ownership	0.344	0.591	0.232	0.225	0.271	0.186
Independent Ownership	0.203	0.591	0.163	0.184	0.094	0.103

Notes: This table displays the Herfindahl–Hirschman Index (HHI) based on different application category market definitions using the baseline period data. I take the category(s) in each column as the definition of the market and compute the HHI of this market. The first row displays the HHI under the current ownership structure (i.e. Facebook owns Facebook, Instagram, Messenger, and WhatsApp). The second row displays the HHI if each of these applications was independently owned.

Implications for Market Concentration: A natural question is whether the observed cross-category substitution would be lead to any differences in assessments of the degree of market concentration. I consider a conservative relevant market definition that includes categories between which substitution was observed and compute the most common market concentration index, the Herfindahl–Hirschman Index (HHI).[20] Table 5 displays the results, separating out the measures by applications individually and by incorporating Facebook ownership. An HHI above 0.25 generally indicates excessively high concentration. There are two main observations. First, multi-category market definitions leads to substantially lower estimated concentration than the application category market definitions alone. For instance, the HHI reduces from 0.591 to 0.225 (from entertainment to social and entertainment) and from 0.344 to 0.271 (from social to social and communication) for YouTube and Instagram respectively – indicating the importance of taking into account the actual substitution patterns in assessing market concentration. Second, despite this, market concentration would be substantially lower if each of the Facebook-owned applications was independently owned, regardless of whether the market definition was single or multiple categories.

4.2.3 Newly Installed Applications and Long-Tail Substitution

In this section, I analyze whether the restrictions induce the participants to substitute towards prominent applications or explore new applications and substitute towards the long tail of applications available in the Google Play Store. I use the fact that I observe the set of installed applications on the phone every day to construct a measure of the number of newly installed applications and the corresponding time spent on them. Furthermore, I characterize whether participants substitute

[20] HHI is defined as follows: $HHI = \sum_j s_j^2$.

towards applications in the Facebook ecosystem – Facebook, Messenger, WhatsApp, Instagram –, “major” applications, or “long tail” applications as a proxy to understand whether substitution is directed towards larger applications or scattered across the long tail of applications. I define “major” applications as those that are not in the Facebook ecosystem or core phone applications, but are in the top 25 applications in terms of average time usage in the baseline period.[21]

Table 6: Newly Installed Applications During the Restriction Period – YouTube

	Dependent variable:				
	Number of Applications Installed	asinh(Number of Applications Installed)	% change in Applications Installed	Time on New Applications	asinh(Time on New Applications)
	(1)	(2)	(3)	(4)	(5)
YouTube Treatment	0.908	0.176*	0.005	3.532**	0.394**
	(0.732)	(0.105)	(0.004)	(1.471)	(0.163)
Block Controls	Yes	Yes	Yes	Yes	Yes
Observations	243	243	243	243	243

Table 7: Newly Installed Applications During the Restriction Period – Instagram

	Dependent variable:				
	Number of Applications Installed	asinh(Number of Applications Installed)	% change in Applications Installed	Time on New Applications	asinh(Time on New Applications)
	(1)	(2)	(3)	(4)	(5)
Instagram Treatment	0.223	0.009	0.003	1.432	0.078
	(0.344)	(0.101)	(0.004)	(1.145)	(0.150)
Block Controls	Yes	Yes	Yes	Yes	Yes
Observations	242	242	242	242	242

*p<0.1; **p<0.05; ***p<0.01

Notes: Columns (1) and (2) report the regression with the dependent variable as the total number of newly installed applications in levels and logs respectively. Column (3) reports the regression with the dependent variable as the % increase in new applications. Columns (4) and (5) report the regression with the dependent variable as the average daily minutes spent on these new applications in levels and logs respectively. Reported standard errors are heteroskedasticity-robust standard errors.

Newly Installed Applications: I construct a measure of the number of newly installed applications as follows. For each week, I collect the set of applications that had been detected to be installed on the phone at any point during the week.[22] Then, for each week following the baseline week, I compute the number of applications that were present on the participant’s phones this week that

[21]The set of major applications comprises of the applications: Reddit, YouTube, TikTok, Netflix, Twitter, Discord, Snapchat, Twitch, LinkedIn, Spotify, Zoom, Telegram, Hulu, Prime Video, Signal, Google, Amazon Shopping. I exclude time spent on Messages, Phone, Gmail, Clock, Gallery, Google Play Store, Camera, Browser, Chrome Beta, Drive.

[22]Recall that the set of installed applications is pulled at the same time that the data is pulled from the parental control application and so occurs late at night. There was an issue pulling the installed applications for seven of the participants during the baseline period and so their data is dropped for only this part of the analysis. An issue with the script prevented collection of this data for all participants for the first couple of days of the baseline period.

were not present in the previous week, the time spent on these new applications during the week, and the percentage increase in total applications between the weeks.

I estimate specification (1) with the dependent variables as the number of newly installed applications and the amount of time spent on them. Similar to before, I focus on the first week of the restriction period with the results are reported in Table 6 and Table 7, respectively.[23] I find that there is an imprecise increase in the number of newly installed applications for YouTube, but that there is a statistically and economically significant increase of 3.5 minutes per day in time spent on these applications. For Instagram, there is neither an increase in the number of installed applications nor a difference in the time spent on them. One interpretation of this result is that for Instagram the substitutes are more apparent to participants (e.g. Facebook), which leads to less need to install new applications. For YouTube, the substitutes are less apparent so participants are less likely to have readily available substitutes and thus spend more time off the phone as well as be more likely to explore new alternatives.

Table A14 further shows that a substantial proportion of participants not only believe they substituted towards other applications during the restriction, but also actively "invested" in them so that they could source better content from them. For instance, one participant wrote *"I had to figure out what I want from other applications I didn't know offered similar content before time, after the restriction elapsed, I had adjusted to sourcing for such content on both apps."* This suggests that there was active adjustment in the extensive margin of installing new applications as well as that participants more fully explored the capabilities of other applications.

Substitution to the Long-Tail: I now study whether participants are substituting to a few prominent applications or dispersed amongst the long tail of applications. To investigate this question, I use the same empirical specification as the cross-category substitution regressions, but consider the categories as overall time on the Facebook ecosystem, major applications, and long tail applications. Table A8 displays the results for Instagram. Indeed, while there is little observed substitution to "long tail" applications or other major applications, there is weak evidence for substitution towards Facebook-owned applications, with an imprecise 16.3% increase in the main experiment and a more precisely estimated, statistically significant, 17.4% increase when pooled with the pilot experiment. Table A9 displays the results for YouTube. The effects in this case are more muted with a clear drop in "major applications" due to the drop in YouTube time, but only a small amount of substitution towards the other categories. Once I condition on phone usage, I find that the largest share gain is to the Facebook ecosystem and the long tail applications. Thus, substitution for Instagram is more concentrated, in particular concentrated within the Facebook ecosystem, compared to the more dispersed substitution patterns observed for YouTube.

[23] As I do not observe the week before the baseline, this regression does not control for baseline usage.

4.2.4 Off Phone Substitution

One possible concern is that since the restriction is only on the phone, participants may substitute to the restricted application on other devices, which would bias the previous estimates. This would understate substitution towards other applications since such substitution would replace time spent on other phone applications, implying that the previous results are a lower bound. In order to assess the extent of cross-device substitution, I rely on the weekly subjective time use surveys and the data from the Chrome Extension. In the weekly surveys, the participants self-report how much time they spent on several applications off their phone. Table A10 displays the results on non-phone Instagram and YouTube time, which show negative point estimates on the time spent on both applications. Indeed, the estimates point to a statistically significant *reduction* in time spent on YouTube off the phone.[24]

Table 8: Substitution Towards the Computer During Treatment Week

	Dependent variable:					
	Overall Computer Time	asinh(Overall Computer Time)	YouTube Computer Time	asinh(YouTube Computer Time)	Instagram Computer Time	asinh(Instagram Computer Time)
	(1)	(2)	(3)	(4)	(5)	(6)
Instagram Treatment	8.294	−0.072			1.583**	0.387***
	(13.794)	(0.115)			(0.796)	(0.093)
YouTube Treatment	18.011	−0.094	9.226*	0.108		
	(13.464)	(0.112)	(5.182)	(0.166)		
Baseline Time Controls	Yes	Yes	Yes	Yes	Yes	Yes
Block Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	331	331	225	225	216	216
R^2	0.699	0.676	0.484	0.624	0.156	0.365

*p<0.1; **p<0.05; ***p<0.01

Notes: Heteroskedasticity-robust standard errors are reported in parentheses. The table presents the estimated ATE on average daily computer minutes during the first week of the restriction period using the recorded data from the Chrome Extension. The first and second columns present the estimated ATE of overall computer usage for levels and logs respectively. The third and fourth columns present the estimated ATE of computer YouTube usage for levels and logs respectively. The fifth and sixth columns present the estimated ATE of computer Instagram usage for levels and logs respectively.

The result that time on the restricted applications potentially *decreases* on non-phone devices is possibly driven by biases in self-reported time usage data. The biases in such data has been pointed out by Ernala et al. (2020) in the context of social media usage. I use the data from the Chrome Extension in order to get an objective measure of how participants substituted, which allows me to

[24]One possible worry is that participants are misinterpreting the survey and reporting aggregate time spent on the application across all devices. However, the survey was explicitly designed to include a grayed out column for phone time saying that it was automatically collected and then next to it including a time for other device time in order to minimize the likelihood of this occurring. Furthermore, I obtained the same result in the pilot experiment and this was the main reason I added the Chrome Extension in order to have a non-self reported measure of this quantity.

validate whether the self-reported data is indeed biased or if it was the case that participants did not substitute at all across devices. Table 8 considers the same specification for the subset of participants that have installed the Chrome Extension. I estimate whether there was a change in overall computer time, Instagram time on the computer, and YouTube time on the computer. Table 8 finds little evidence that overall computer time changed as a result of the treatment. However, there is a marginally significant increase of 9.3 minutes of computer time on YouTube during the YouTube treatment and a statistically significant increase of 1.58 minutes of computer time on Instagram during the Instagram treatment. These point estimates indicate that there was a small amount of cross-device substitution and suggest that I am slightly underestimating the degree of substitution towards other applications on the phone.[25] In order to interpret the magnitude of the cross-device substitution, it is important to note that the baseline usage of Instagram computer usage is only 1 minute a day on average and, conditional on usage in the baseline, Table 1 reports that the average and median usage are 4.05 and 0.43 average daily minutes respectively. Furthermore, Figure A10 shows the time series of usage of the restricted applications across both devices and indicates that the aggregate usage of the applications drops dramatically during the treatment week.

Beyond the extent of cross-device substitution towards the restricted application, there is a broader question of whether there are non-digital substitutes to the restricted applications. Column (6) of Table 2 and Table 3 displays the estimated average treatment effects for overall phone usage during the Instagram and YouTube treatments respectively. It shows that there is a reduction of 27 minutes and 44 minutes per day of phone time as a result of the Instagram and YouTube treatments respectively. The logs specification shows a lesser effect with a statistically significant and meaningful drop in phone time for YouTube, but an imprecise, negative point estimate for Instagram. Consistent with this, I find that this is primarily driven by reductions in phone usage of participants in the upper deciles of phone usage.[26] Figure A8 shows that while the YouTube restriction leads to fairly depressed phone usage throughout the entirety of the day, the reduction in phone usage for the Instagram treatment is largely in the afternoon and evening hours. Thus, it is plausible that, especially for Instagram, participants are substituting to non-digital substitutes during these hours. It is unclear what activities off the phone participants are substituting to as Table A11 displays the estimated average treatment effect on the most natural off-phone substitutes, such as cable television, video games and streaming services, and finds no effect on time spent on these services.

[25]Furthermore, the discrepancy in the sign of the effect between the survey-based and objective measures highlights the importance of collecting objective data for time allocations.

[26]Figure A9 plots the quantile treatment effects for each decile for logs of overall phone time. It shows that the QTE of the Instagram treatment is quite similar across deciles, whereas for YouTube it is more likely to be driven by reductions in the lower deciles.

4.3 Time Substitution After Restriction Period

In this section I explore the extent to which there are persistent effects as a result of the restrictions. This is important for understanding whether there are potentially dynamic elements of demand for such services and will be used to guide the demand model in Section 5.

Figure 2: Time on Restricted Applications

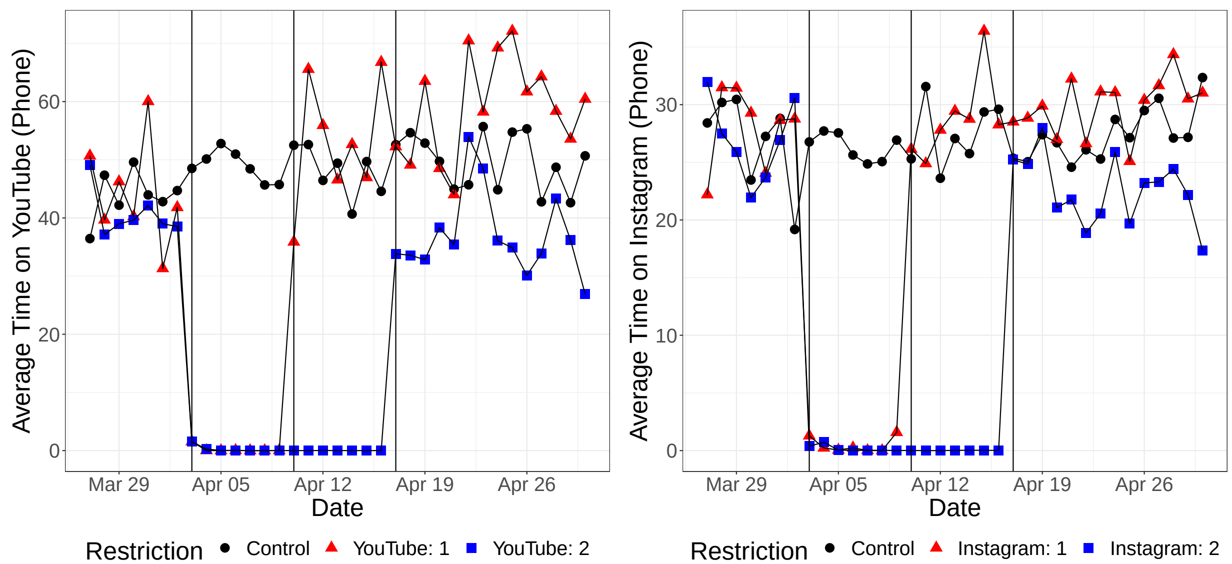

Notes: This figure plots the average daily minutes on Instagram (left) and YouTube (right) on the phone. Each figure plots the usage of the control group, one week and two week restriction group for the application.

Persistent Effects on Restricted Applications: Figure 2 plots the time series of usage of the restricted applications through the study period. There are two striking patterns. First, in both treatments, the one week restriction group usage returns to pre-experiment levels immediately after the restriction is lifted. Second, in both treatments, the two week restriction group usage immediately jumps to a lower level than the pre-restriction period and persists at this level until the end of the study period. Figure A10 shows that the same trend holds when time spent on the computer is included. Motivated by these observations, I then estimate specification (1) with heterogeneous effects across restriction lengths and for the 2 week group alone with the results reported in Table A12 and Table A13. Columns (1) and (2) of Table A12 and Table A13 show the change in restricted application time in levels and logs for Instagram and YouTube respectively. For Instagram, there is a statistically significant difference in post-restriction time between these two for the levels specification and the 2 week restriction group. Furthermore, columns (3) and (4) drops the 1 week group entirely and estimates the treatment effect for only the 2 week group. This further confirms that there is a drop of approximately 5 minutes of time on Instagram on average for the Instagram restriction group. For YouTube, there is a negative, but imprecise point, estimate for both specifications. Given the skewed usage distribution and the discrepancy between logs and levels, one might expect that the changes in post-restriction usage are driven by those at the high end of the usage distribution. Figure A11 estimates the QTE of post-restriction effects and confirms this intuition.

A natural question is, if such post-restriction effects exist, how persistent are they? It is plausible that these effects dissipate very quickly, but I only observe participants for 2-3 weeks following the restriction. In order to understand how much longer the effects last, I rely on an optional survey that was sent one month following the conclusion of the study asking how they had been spending their time relative to before the experiment.[27] Participants could mark whether they were spending a lot less time (1), somewhat less time (2), the same amount of time (3), somewhat more time (4), or a lot more time (5). They could also mark if they did not use the application or had started to use it during the study period. I estimate the impact of the restrictions on overall phone, overall social media, Instagram, and YouTube usage. Table A15 displays the estimated average treatment effect, which shows that there is still a large drop in the Instagram treatment group's overall social media and Instagram usage.

This result must be caveated for the following two reasons. First, there is potential for selection bias since participants with stronger responses to the treatment may be more willing to respond. However, roughly an equal number of participants from both the treatment and control group responded, indicating this may not be a large concern.[28] Second, these are unverifiable survey responses, so it is possible that some of the results are driven by experimenter demand. Subject to these caveats, these results show that a one or two week restriction led to a reduction in usage *nearly two months* later. Combined with the other results, this provides evidence that there were persistent effects of the restrictions on the restricted application.

Persistent Effects on Non-Restricted Applications: Table A16 and Table A17 provide estimates for persistent changes on usage of non-restricted applications as a result of the Instagram and YouTube treatments. I focus on applications / categories where I observed substitution towards during the restriction period and the applications installed during the restriction period.[29] I find little evidence of persistent changes in usage along these dimensions. The only notable persistent increase is in the amount of time spent on applications that were installed during the restriction period for YouTube. This is important since it highlights that participants in the YouTube restriction

[27]Participants were incentivized by being able to enter the chance to win a $100 Amazon Gift Card by completing the survey. However, they had already been paid their experimental payment after the conclusion of the study period.

[28]Since there is only partial response I do not include the controls for randomization block.

[29]In order to economize on space I do not include the interaction term in the reported estimates. Instead, I estimate the ATE of the persistent for both restriction lengths (without an interaction) and then report point estimates for the 14-day treatment group alone.

discovered new applications during the restriction period that they had not sought out before and will be important motivation for the modeling approach we take in Section 5. Furthermore, there is a marginally significant increase in time spent on Instagram for participants in the YouTube treatment. However, beyond this, there are minimal persistent changes on other applications.

It must be noted that these are average treatment effects and I observed heterogeneous substitution during the restriction period itself, so it does not rule out that there were persistent changes in time usage but that these are so heterogeneous that they would not be picked up by this specification. Indeed, Table A14 indicates that participants self-report having persistent effects on other applications, but the effect sizes may be too small for them to be detectable given the power of the experiment.

5 Model of Time Usage with Inertia

In this section I estimate a discrete choice model of time usage for prominent social media and entertainment applications. Motivated by the experimental results, I consider the role of inertia in driving the time allocation choices of the participants and directly include this in the model. The main output of interest from the model is an essential output of a demand model needed for market definition and merger analyses - diversion ratios. I use the demand model to quantify the role of inertia in driving usage and produce diversion ratios with the inertia channel shut off.[30] Finally, I apply the estimated diversion ratios with and without inertia to a less conservative relevant market definition test than in Section 4.2.2 and assess the quantitative significance of the inertia channel.

5.1 Model and Identification

I model participant's choices as a panel of discrete choices. Informally, I consider this as participants in each time period deciding whether to use a social media application in order to "kill time" or directly seeking information from the applications.[31] This modeling assumption is consistent with the survey responses of the participants as they said they were habituated to open up the application to take a break and sometimes attempted to do so though even though they knew the application was restricted:

[30] In Online Appendix B, I report diversion ratio estimates using only the estimated experimental treatment effects of substitution between the applications that I consider in this section. This enables a less parametric estimate of the diversion ratios, but does not allow me to quantify the role of inertia in usage.

[31] Other models of time demand for applications such as Han, Park and Oh (2016) and Yuan (2020) consider a multiple discrete-continuous framework. Byzalov (2008) takes a similar approach as mine when considering time allocation demand for channels on cable television. One benefit to the discrete choice approach is that it enables me to flexibly control for variation in usage throughout the day and week, which is apparent in Figure A4, as well as directly incorporate past usage into the utility function.

- "*At first restricting instagram was frustrating as I had the application on my homescreen and built muscle memory for the past 4 years to press that part of the screen where the instagram shortcut is...*"

- "*It's strange, because I didn't feel like I needed YouTube, I just knew I had spent a lot of time on it. ... It felt weird knowing that my instinct was to immediately press the YouTube button when I got bored...*"

I use the experimental results to guide the key assumptions of the consumer utility function. The first assumption is that the choice of application is driven by *consumer inertia*. There are two experimental results that point to the importance of inertia. First, participants spent time on newly installed applications and persisted to use these applications, even once the restriction period was over. This indicates that search/inattention plays a role in driving usage. Second, there is a reduction in usage of the restricted application in the post-restriction period, especially for the power users of the applications. This indicates that habit formation plays a role in driving usage. The second assumption is that participants are myopic and thus do not consider how current period usage will impact their future usage.[32]

There is a set of participants $\mathcal{I} = \{1, ..., I\}$, indexed by i, and a set of applications $\mathcal{J} = \{0, 1, ...J\}$, indexed by j, where 0 denotes the outside option. I consider each distinct choice set observed across participants as a separate "market", denoted by k. This includes the set of currently installed applications on their phone minus any applications that are experimentally restricted. Participant i receives the following utility from application j in market k and time period t:

$$u_{ijkt} = \beta^{q(i)} \cdot h_{ijt} + \zeta^{q(i)} \cdot r_{ijt} + \omega^{q(i)} \cdot r_{ijt}^2 + \gamma_j^{q(i)} + \kappa^{q(i)} \cdot ac_{ij} + \alpha^{q(i)} \cdot asinh(cn_{ij}) + \epsilon_{ijkt} \quad (2)$$

where $\gamma_j^{q(i)}$ denotes application fixed effects, ac_{ij} incorporates the subjective usage of application j, which comes from Table 4, for participant i, cn_{ij} denotes the self-reported number of connections for participant i on application j, and ϵ_{ijkt} is the Type-1 Extreme Value error.[33] $q(i)$ denotes the type of participant i that is determined by running k-means on the aggregated baseline data in order to group participants into different types. Thus, the specification accommodates preference heterogeneity across participants by having type-specific estimates of the coefficients, incorporating the subjective uses of the applications directly into the utility function, and accounting for the number of accounts each participant follows.

[32] This is justified since Figure 2 indicates a lack of a spike of usage on the day that the restricted application is unrestricted, which we would expect if the degree of intertemporal substitution was strong. Similarly, Allcott, Gentzkow and Song (2021) find that consumers are aware that current period usage impacts future usage, but behave as though they are inattentive to it and thus are effectively myopic.

[33] cn_{ij} is the self-reported range of the number of accounts that the participants follow on each application. See Online Appendix A for the details.

The main parameters of interest are those that relate to consumer inertia. There are broadly two types of inertia effects that are present – short-term and long-term inertia. I model long-term inertia as a continuous stock of past usage directly entering into the utility function in a similar manner to the state-dependent demand estimation literature (e.g. see Dubé, Hitsch and Rossi (2010)).[34] Motivated by the apparent difference in long-run behavior between the one and two week restriction groups, I define the usage stock, h_{ijt}, as the total amount of time participant i has spent on application j in the past two weeks.[35] It is important to note that this formulation broadly captures multiple mechanisms that can induce state-dependence, several of which there is experimental evidence for, (e.g. see MacKay and Remer (2022) for discussion), which limits the welfare claims that I can make. However, it allows me to understand how inertia influences substitution patterns and characterize its overall importance in driving usage.

Due to the discrete choice formulation, it is important to further account for short-term inertia, which is that a participant is more likely to choose application j in period t if they used the application in period $t-1$. I include a term, r_{ijt}, which is defined as the number of consecutive periods which participant i has used application j. Since this short-term component potentially has satiation effects, it enters the utility function both linearly and quadratically. It is important to emphasize that the short-term inertia term is largely a nuisance term to get a more precise estimate of longer term inertia.

The granularity of the data allows me to vary the outside option flexibly across time.[36] For any time index t, I allow the outside option to vary across the week of the experiment $w(t)$, day of the week $d(t)$, and hour of the day $o(t)$. I collapse the hours of the day into morning (7 AM - 12 PM), afternoon (12 PM - 6 PM), evening (6 PM - 1 AM), and late night (1 AM - 7 AM). I normalize the outside option to zero at late night, Sundays, and the final week of the experiment. Thus, the utility for the outside option is denoted as follows where $\alpha_{o(t)}$ denotes hour of day fixed effects, $\iota_{d(t)}$ denotes day of week fixed effects, and $\mu_{w(t)}$ denotes week fixed effects:

$$u_{i0tk} = \alpha_{o(t)} + \iota_{d(t)} + \mu_{w(t)} + \epsilon_{i0tk}$$

[34] Directly considering a continuous stock of past usage in the utility specification is similar to the formulation used in Bronnenberg, Dubé and Gentzkow (2012) as well as papers focused on characterizing demand for addictive goods (Becker and Murphy, 1988; Gordon and Sun, 2015).

[35] There is an initial conditions problem at the beginning of the experiment since there is no previous data to use to define this. Because of this I drop the first two days of data entirely from the estimation and, for any date in the first two weeks, I multiply the accumulated "stock" by the inverse of the fraction of the current time period by the time period exactly 2 weeks from the start of the experiment. I chose two days since the descriptive statistics point to usage not varying drastically across the days of the week and preliminary experiments showed that after two days the usage stock variable is fairly constant in the baseline period.

[36] Figure A4 shows how phone usage varies across the hours of the day and days of the week. The modeling assumption captures that this variation is likely not driven by changes in the value of e.g. Facebook throughout the day, but variation in the value of non-phone activities throughout the day and the week.

The assumption that ϵ_{ijkt} is independent and identically distributed according to a Type-1 extreme value distribution induces the following probability that application j will be chosen by participant i if it is available to them in market k:

$$\frac{\exp(\beta^{q(i)} \cdot h_{ijt} + \zeta^{q(i)} \cdot r_{ijt} + \omega^{q(i)} \cdot r_{ijt}^2 + \gamma_j^{q(i)} + \kappa^{q(i)} \cdot ac_{ij}) + \alpha^{q(i)} \cdot asinh(cn_{ij}))}{\exp(\alpha_{o(t)} + \iota_{d(t)} + \mu_{w(t)}) + \sum_{j'} \exp(\beta^{q(i)} \cdot h_{ij't} + \zeta^{q(i)} \cdot r_{ij't} + \omega^{q(i)} \cdot r_{ij't}^2 + \gamma_{j'}^{q(i)} + \kappa^{q(i)} \cdot ac_{ij'} + \alpha^{q(i)} \cdot asinh(cn_{ij'}))}$$

Identification: The primary parameter of interest is $\beta^{q(i)}$. The typical identification challenge is to disentangle the effect of past usage on current usage from preference heterogeneity. The model flexibly captures preference heterogeneity by having type-specific parameter estimates and capturing the self-reported type of usage and number of connections for each application.[37] The subjective usage of the applications is important for interpreting the substitution patterns in the restriction period and thus captures an important dimension of preference heterogeneity directly. Furthermore, by directly exploiting the set of currently installed applications, I have variation in the choice sets across different participants and this separates out the case when a participant has no usage stock because they don't have the application installed. The experimental restrictions provide exogenous variation in the usage stock of the restricted applications as well as the other applications (via substitution during the restriction period). Thus, the core assumption for identification is that the restriction induces a shock to the usage stock and does not impact the intrinsic preferences for the applications.

Estimation: I use the session data aggregated to the time interval of 5 minutes (see Appendix A for additional details on the session data). In order to map the session data to a discrete choice, I compute the time allocations allotted to each application in each interval, including off the phone time, and assign the chosen application in this time period as the maximum of these quantities. I restrict myself to the most prominent social media and entertainment applications – Facebook, TikTok, Twitter, Reddit, YouTube, Instagram, and Snapchat – and denote every other application or off phone activity as the outside option. For these applications, I collect the average daily usage in the baseline period for each participant and cluster the participants according to k-means. I then estimate the model separately for each type. Since my model is likelihood-based, I estimate the parameters using maximum likelihood estimation and construct standard errors using bootstrap.

[37]The biggest worry about unobserved heterogeneity in usage comes from the power users of specific applications or bundles of applications. The clustering formulation is able to capture the differences in preference intensity for these participants and considers separate estimates for them. The approach of discretizing a potentially continuous distribution of unobserved heterogeneity through k-means has precedent in Bonhomme, Lamadon and Manresa (2022).

5.2 Model Estimates and Validation

The first step of estimation requires classifying the participants into different types using k-means. There is a large literature in data mining and statistics about choosing the "optimal" k that trades off the parsimony of having fewer clusters against the reduction in within-cluster variance that arises from additional clusters. In this case an additional consideration is that it is important to ensure that the clusters have sufficiently many participants to allow for estimation of the parameters of interest for this group, but also having sufficiently many clusters to capture the unobserved preference heterogeneity. I consider an index of these measures for choosing the "optimal" k which reports $k \in \{3, 6\}$. In order to accommodate additional heterogeneity in consumer preferences, I utilize $k = 6$.[38] Figure A12 displays the clusters and time allocations within each of them. The clustering of participants identifies sets of power users. Cluster 1 captures power users of Facebook. Cluster 2 identifies participants who are power users of TikTok, but also use the other social media applications extensively. Clusters 3 and 6 capture the YouTube intensive participants. Cluster 4 captures Instagram power users. Cluster 5 captures the typical users of these applications who have moderate usage of each of the applications and consists of the majority of participants.

The estimates from the model are presented in Table A18. I report the estimates of each type separately. The first observation is that the coefficient on h_{ijt} is fairly consistent across the different types as well as the estimate for the influence of short-term inertia, r_{ijt} and r_{ijt}^2. Both of these terms are statistically different from 0, indicating that both the short-term and long-term inertia channels play a role. The coefficient on r_{ijt}^2 is negative, indicating satiation effects. The differences in the natural usage of each of the applications across the different types, which is reflected in Figure A12, naturally translates to differences in the estimated application fixed effects. The coefficients on the different subjective uses of the applications varies across the types in accordance with the most used applications by participants classified as that type. I validate the in-sample fit of the model by comparing how well the model is able to match the actual market shares throughout the study period. Table A19 validates that the model fits the data reasonably well as it matches the non-restriction period market shares and predicts the extent of substitution towards other applications and the outside option as a result of the experimental restrictions.

The primary output of the estimated model is the second-choice diversion ratio. The second-choice diversion ratio between application j and k provides an estimate of what fraction of consumption of application k would shift from application k to application j if application k was removed from the choice set. Typically, regulatory authorities use second-choice diversion ratios coming from switching surveys as a critical input to relevant market definition and merger evalua-

[38] I additionally consider density-based spatial clustering of applications with noise (DBSCAN) (Ester et al., 1996) and spectral clustering (Von Luxburg, 2007) which are clustering algorithms that do not restrict themselves to convex regions. Following best practices for the methods, I find that they do not result in a substantially different clustering.

tion (Reynolds and Walters, 2008; Conlon and Mortimer, 2021). In order for the model to provide reasonable estimates for this quantity it is important that it is able to predict how participants would substitute towards the other applications if the application was not available. The model validation exercises showed that the model is able to do this for the Instagram and YouTube restrictions and thus ought to provide a reasonable estimate of this quantity.

Table 9 displays the estimated diversion ratios, which are given by $D_{jk} = \frac{s_j(\mathcal{J}\setminus\{k\})-s_j(\mathcal{J})}{s_k(\mathcal{J})}$. Each of the diversion ratios is computed as before, by a weighted average over the different types according to the fraction of participants assigned to a type. The diversion ratios across each of the different applications predict a large amount of diversion to the outside option, with Instagram and YouTube having the highest diversion towards the outside option and reflect intuitive patterns of substitution.

Table 9: Second-Choice Diversion Ratios

	Facebook	Instagram	Reddit	Snapchat	TikTok	Twitter	YouTube	Outside Option
Facebook	-	0.0242	0.00292	0.00504	0.00404	0.00486	0.0187	0.94
Instagram	0.0126	-	0.00311	0.00728	0.0053	0.00494	0.0198	0.947
Reddit	0.00592	0.012	-	0.0058	0.00238	0.00292	0.0287	0.942
Snapchat	0.00989	0.0269	0.00559	-	0.0148	0.00427	0.0214	0.917
TikTok	0.0103	0.0235	0.00281	0.0156	-	0.0091	0.0303	0.908
Twitter	0.0131	0.0257	0.00417	0.00579	0.00997	-	0.0216	0.92
YouTube	0.00944	0.0199	0.007	0.0052	0.00625	0.00393	-	0.948

Notes: This table displays the estimated second-choice diversion ratios that come from the estimated model. The cell in each row k and column j is computed by $D_{kj} = \frac{s_j(\mathcal{J}\setminus\{k\})-s_j(\mathcal{J})}{s_k(\mathcal{J})}$.

5.3 Counterfactuals: No Inertia and Relevant Market Definition

In this section I utilize the model estimates to conduct several counterfactuals related to better understanding the relevant market definition and the role of consumer inertia in driving usage. First, I characterize how consumer inertia influences the resulting market shares and diversion between the applications. I then apply these estimates to a stylized relevant market definition exercise to understand the quantitative relevance of the inertia channel. Finally, I explore the quantitative significant of the inertia channel versus full adoption of the set of considered applications.

5.3.1 Usage and Diversion Ratios without Inertia

I characterize the role of long-term inertia in usage by imposing $\beta^{q(i)} = 0$ and compute the change in the resulting market shares and diversion ratios. It is important to understand the interpretation of this counterfactual since it is not a direct policy counterfactual. The inertia channel comprises a number of different aspects of usage – ranging from addictive impulses to more natural mechanisms such as switching costs — and I provide an interpretation that is directly motivated by regulatory and antitrust concerns.

The primary interpretation of the counterfactual is that it provides a more direct measure of substitutability between the applications of interest. The observed substitution patterns are a mix of direct substitutability and inertia, where the latter channel naturally favors more prominent incumbent applications. Thus, the diversion ratios without inertia provide a more nuanced view of the set of substitutes for each application that does not depend as strongly on the prominence of the given application. This enables a more accurate measurement of substitution between prominent and niche applications. A secondary interpretation is that, while some aspects of inertia are natural components of application choice, there are features of these applications that can inflate the importance of this inertia channel and a number of policy instruments have been proposed for alleviating this issue. For instance, it has been argued that the objective function of content curation algorithms and design patterns such as infinite scroll news feeds result in excessive usage of these applications (Narayanan et al., 2020). Thus, this counterfactual also provides an upper bound for how policies targeting these aspects would impact the substitution between applications and what impact they would have on the competitiveness of the market.

Table 10: Market Shares (No Inertia)

Application	No Inertia: Weeks 1,4,5	Baseline Weeks 1,4,5	No Inertia: Weeks 4,5	Baseline: Weeks 4,5	No Inertia: Week 1	Baseline: Week 1
Facebook	0.00735	0.0105	0.00706	0.00989	0.008	0.0117
Instagram	0.0139	0.0209	0.0132	0.0196	0.0154	0.0237
Reddit	0.00453	0.00638	0.00447	0.00644	0.00465	0.00625
Snapchat	0.00412	0.0049	0.00395	0.00471	0.00451	0.00534
TikTok	0.00428	0.00702	0.00422	0.00697	0.00441	0.00712
Twitter	0.00311	0.00379	0.00306	0.00371	0.0032	0.00396
YouTube	0.0219	0.0318	0.0217	0.0314	0.0224	0.0327
Outside Option	0.941	0.915	0.942	0.917	0.937	0.909

Notes: Columns 1 and 2 display the predictions of the model over week 1, 4, and 5 including the long-term inertia term and without. Columns 3 and 4 display the prediction of the model only over weeks 4 and 5. Columns 5 and 6 display the prediction of the model only over week 1. Each cell displays the market share of the row application under the specification designated by the column.

Table 10 compares the average market shares with and without the inertia term across different weeks of the experiment when participants had the full set of applications available to them. Since the results across the different subsets of weeks are quantitatively very similar, I restrict focus to the first two columns which compare the differences across all weeks in the experiment. The first observation is that the overall market share of considered applications drops by nearly 30% when this channel is shut down. Table A22 displays the reduction of usage in percentages, showing that TikTok has the largest percentage reduction, nearly 39%, in average usage when this channel is shut down. Recall that TikTok in particular has a smaller number of users in my sample relative to the other applications, but, conditional on using the application, has one of the highest average time allocations. As a result, it is not too unsurprising that the model predicts that inertia is a large driver of usage for this application. Instagram and YouTube also have a sizeable reduction of 33% and 31%, respectively.

I further compute the estimated second-choice diversion ratios when the inertia channel is shut down. The estimates are displayed in Table A20 with the percentage differences between the baseline and no inertia case presented in Table A21. There is a drop in the diversion ratios from other applications towards the most prominent applications such as Instagram or YouTube, but there is not a reduction across the entire matrix of diversion ratios. For instance, there is an increase in diversion from Facebook to Reddit as well as from Reddit to Twitter, which indicates that the smaller applications in my sample can actually benefit from the lack of inertia for the larger applications such as Instagram or YouTube.

5.3.2 An Application to Relevant Market Definition

Table 11: Summary of Market Definition Analysis

Threshold	$\tau = 0.025$	$\tau = 0.05$
Instagram	YouTube, Facebook, *Snapchat*	YouTube, Facebook, Snapchat, Twitter, TikTok, Reddit
Twitter	Instagram, *YouTube, Facebook*	Instagram, YouTube, Facebook, *Snapchat, Reddit, TikTok*
YouTube	Instagram, Facebook, *Reddit*	Instagram, YouTube, Facebook, Snapchat, Reddit, TikTok
TikTok	YouTube, *Instagram*	YouTube, Instagram, *Snapchat, Facebook, Twitter*
Reddit	YouTube, *Instagram*	YouTube, Instagram, Facebook, Snapchat, *Twitter, TikTok*
Snapchat	Instagram, *YouTube*	Instagram, YouTube, TikTok, *Facebook, Reddit, Twitter*
Facebook	Instagram, YouTube	Instagram, YouTube, Snapchat, Twitter, *Reddit, TikTok*

Notes: This table presents the relevant market definition for the reported threshold τ (column) and application (row). The green highlighted applications are in the relevant market using the diversion ratios with and without inertia. The red highlighted (and italicized) applications are only in the relevant market using the diversion ratios without inertia.

In order to assess whether the change in diversion ratios is quantitatively meaningful, I apply both sets of diversion ratios to a simple relevant market definition exercise: for each application x, sort the other applications in descending order of diversion from x and add them one by one into the relevant market until the sum crosses a pre-specified threshold τ.

Table 11 summarizes the results of this exercise using the estimates of diversion with and without inertia (see Table 9 and Table A20, respectively). Table 11 reports the results for $\tau \in \{0.025, 0.05\}$ and indicates that, in most cases, the relevant market is larger using the set of diversion ratios with inertia compared to without. For example, the relevant market definition for TikTok in the baseline is only YouTube and Instagram, whereas using the no inertia diversion ratios leads to Snapchat, Facebook, and Twitter additionally being in the market for $\tau = 0.05$. This highlights that the inertia channel is quantitatively meaningful and, based on the interpretation of the counterfactual, results in two key takeaways. The first is that these more fundamental measurements of diversion indicate a broader set of competitors. Indeed, one striking result is that in contrast to the analysis in FTC (2021), which defined the relevant market for Facebook as being Facebook, Instagram and Snapchat, the relevant market under both the baseline and no inertia counterfactual is broader than this and consistent with the results from the reduced form analysis. The second is that it validates that policies regulating application features that increase the role of consumer inertia can increase competitiveness in this market.

5.3.3 Counterfactual: Full Adoption of Prominent Applications

Table 12: Market Shares (Full Adoption)

Application	Baseline	Full Adoption	No Inertia	Full Adoption & No Inertia
Facebook	0.0117	0.0149	0.008	0.0112
Instagram	0.0237	0.0275	0.0154	0.0186
Reddit	0.00625	0.00891	0.00465	0.00743
Snapchat	0.00534	0.00796	0.00451	0.00724
TikTok	0.00712	0.0105	0.00441	0.00783
Twitter	0.00396	0.00654	0.0032	0.00583
YouTube	0.0327	0.033	0.0224	0.0229
Outside Option	0.909	0.891	0.937	0.919

Notes: This table displays the predicted market shares in the baseline week under four different scenarios. Column 1 (Baseline) is the predicted market share without any counterfactual restrictions. Column 2 (Full Adoption) is the predicted market share when participants are forced to install all the considered applications. Column 3 (No Inertia) is the predicted market share when $\beta^{q(i)} = 0$. Column 4 (Full Adoption & No Inertia) is the predicted market share when participants are forced to install all the considered applications and $\beta^{q(i)} = 0$.

I consider an additional counterfactual that exploits the fact that the model utilizes the set of installed applications on the phone. This allows me to consider how market shares would shift under a counterfactual where there is full adoption of the set of prominent social media and entertainment applications. This is particularly interesting in the case of an application such as TikTok, which has a limited number of users but, conditional on installation, has heavy usage (e.g. see Table 1). There are claims that the expansion of these applications will serve as a strong competitive pressure to applications such as Instagram and YouTube, which this counterfactual enables me to assess. The currently installed applications have an existing usage stock and so I also characterize how shutting down the inertia channel paired with full adoption would influence market shares.

Table 12 displays the results. Comparing the results in columns (1) and (2) leads to the observation that the applications with high baseline levels of adoption in my sample, YouTube and Instagram, do not benefit greatly from full adoption, whereas more niche applications such as Twitter and TikTok gain a large market share. In line with this, the HHI of the considered applications drops from 0.231 to 0.198 moving from the baseline to full entry. As comparison, column (3) displays the market shares with the baseline set of installed applications, but shutting down the long-term inertia channel, as was reported previously. The resulting HHI within this set of applications also decreases from the baseline to 0.223. However, it is important to point out that in this case both the concentration within the set of applications decreases and the market share of the outside option, which includes even more niche applications, increases. Indeed, the increase in the outside option is fairly large compared to the change in the full adoption scenario. Finally, column (4) displays the cleanest comparison between the different applications as it shuts down the inertia channel and forces adoption of each application with a resulting HHI of 0.182. Overall, these results suggest that, while full adoption would decrease concentration, it likely wouldn't alter the competitive balance of the market and further validate that inertia plays a quantitatively large role in driving market shares and diversion in this market.

6 Conclusion

In this paper I report the results of an experiment where I continuously monitor how participants spend time on digital services and shut off their access to Instagram or YouTube on their phones for one or two weeks. I use the resulting data on how participants substitute their time during and after the restrictions in order to uncover a rich picture of the demand for social media and entertainment applications. I illustrate how the estimated substitution patterns can be used to guide questions of market definition that have troubled regulators.

I find heterogeneous substitution patterns across categories and towards non-digital activities during the restrictions. Furthermore, participants with the YouTube restriction spend time on applications installed during the restriction period and that participants with the two week Instagram restriction reduce their time spent on Instagram even after the restrictions are lifted. Motivated by this, I estimate a discrete choice model of time usage with inertia and find that inertia accounts for a substantial fraction of usage. I apply the estimates of second-choice diversion ratios coming from the model with and without inertia to show that the inertia channel is quantitatively large enough to meaningfully expand the relevant market definition for the social media applications. Overall, my results highlight the usefulness of product unavailability experiments for demand and merger analysis in attention markets. These experiments provide a clean way of measuring substitution patterns as well as identifying addiction/inertia effects, which allow for a comprehensive picture of demand for these applications that are relevant to antitrust issues. These experiments are feasible to conduct for regulatory authorities since the nature of digital goods enables individual level, randomized controlled experiments of product unavailability.

My results point to a broad competition for time between social media applications, but also emphasize that inertia drives a substantial fraction of their usage and diversion towards larger applications such as Facebook and YouTube. There are two broader policy takeaways from these results. The first is that, due to the personalized nature and importance of user-generated content on these applications, determining plausible substitutes according to similarities in product characteristics alone – as has been done in several prominent merger cases – is likely to be insufficient. This leads to a broad competition for consumer attention across these applications. The second is that due to the role of consumer inertia in driving usage, policies aimed at curbing aspects of these applications that are conducive to addictive usage are an important policy tool at the disposal of regulators aiming to promote competition in these markets. I believe that the insights from this paper can help push forward the regulatory debate and lead to a better understanding of these zero price attention markets.

References

Abadie, Alberto, Joshua Angrist, and Guido Imbens. 2002. "Instrumental variables estimates of the effect of subsidized training on the quantiles of trainee earnings." *Econometrica*, 70(1): 91–117.

Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. 2020. "The welfare effects of social media." *American Economic Review*, 110(3): 629–76.

Allcott, Hunt, Matthew Gentzkow, and Lena Song. 2021. "Digital Addiction." Working Paper.

Ambrus, Attila, Emilio Calvano, and Markus Reisinger. 2016. "Either or both competition: A" two-sided" theory of advertising with overlapping viewerships." *American Economic Journal: Microeconomics*, 8(3): 189–222.

Andersen, Steffen, Glenn W Harrison, Morten Igel Lau, and E Elisabet Rutström. 2006. “Elicitation using multiple price list formats.” *Experimental Economics*, 9(4): 383–405.

Anderson, Simon P, and Stephen Coate. 2005. “Market provision of broadcasting: A welfare analysis.” *The Review of Economic Studies*, 72(4): 947–972.

Anderson, Simon P, Øystein Foros, and Hans Jarle Kind. 2018. “Competition for advertisers and for viewers in media markets.” *The Economic Journal*, 128(608): 34–54.

Andreassen, Cecilie Schou, Torbjørn Torsheim, Geir Scott Brunborg, and Ståle Pallesen. 2012. “Development of a Facebook addiction scale.” *Psychological Reports*, 110(2): 501–517.

Athey, Susan, Emilio Calvano, and Joshua S Gans. 2018. “The impact of consumer multihoming on advertising markets and media competition.” *Management Science*, 64(4): 1574–1590.

Bailey, Erica R, Sandra C Matz, Wu Youyou, and Sheena S Iyengar. 2020. “Authentic self-expression on social media is associated with greater subjective well-being.” *Nature Communications*, 11(1): 1–9.

Bakshy, Eytan, Solomon Messing, and Lada A Adamic. 2015. “Exposure to ideologically diverse news and opinion on Facebook.” *Science*, 348(6239): 1130–1132.

Baym, Nancy K, Kelly B Wagman, and Christopher J Persaud. 2020. “Mindfully Scrolling: Rethinking Facebook after Time Deactivated.” *Social Media + Society*, 6(2): 2056305120919105.

Becker, Gary S, and Kevin M Murphy. 1988. “A theory of rational addiction.” *Journal of Political Economy*, 96(4): 675–700.

Bellemare, Marc F, and Casey J Wichman. 2019. “Elasticities and the inverse hyperbolic sine transformation.” *Oxford Bulletin of Economics and Statistics*.

Berry, Steven, and Philip Haile. 2016. “Identification in differentiated products markets.” *Annual Review of Economics*, 8: 27–52.

Bhat, Chandra R. 2008. “The multiple discrete-continuous extreme value (MDCEV) model: role of utility function parameters, identification considerations, and model extensions.” *Transportation Research Part B: Methodological*, 42(3): 274–303.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa. 2022. “Discretizing unobserved heterogeneity.” *Econometrica*, 90(2): 625–643.

Braghieri, Luca, Ro’ee Levy, and Alexey Makarin. 2022. “Social media and mental health.” *American Economic Review*, 112(11): 3660–93.

Bronnenberg, Bart J, Jean-Pierre H Dubé, and Matthew Gentzkow. 2012. “The evolution of brand preferences: Evidence from consumer migration.” *American Economic Review*, 102(6): 2472–2508.

Brynjolfsson, Erik, Avinash Collis, and Felix Eggers. 2019. “Using massive online choice experiments to measure changes in well-being." *Proceedings of the National Academy of Sciences*, 116(15): 7250–7255.

Burke, Moira, and Robert E Kraut. 2016. "The relationship between Facebook use and well-being depends on communication type and tie strength." *Journal of Computer-Mediated Communication*, 21(4): 265–281.

Byzalov, Dmitri. 2008. "Unbundling Cable Television: An Empirical Investigation." working paper.

Calvano, Emilio, and Michele Polo. 2020. "Market Power, Competition and Innovation in digital markets: A survey." *Information Economics and Policy*, 100853.

CMA. 2020. "Online platforms and digital advertising: Market Study Final Report."

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

Conlon, Chris, Julie Mortimer, and Paul Sarkis. 2022. "Estimating Preferences and Substitution Patterns from Second Choice Data Alone."

Conlon, Christopher, and Julie Holland Mortimer. 2021. "Empirical properties of diversion ratios." *The RAND Journal of Economics*, 52(4): 693–726.

Conlon, Christopher T, and Julie Holland Mortimer. 2013. "Demand estimation under incomplete product availability." *American Economic Journal: Microeconomics*, 5(4): 1–30.

Corrigan, Jay R, Saleem Alhabash, Matthew Rousu, and Sean B Cash. 2018. "How much is social media worth? Estimating the value of Facebook by paying users to stop using it." *Plos One*, 13(12): e0207101.

Dubé, Jean-Pierre, Günter J Hitsch, and Peter E Rossi. 2010. "State dependence and alternative explanations for consumer inertia." *The RAND Journal of Economics*, 41(3): 417–445.

Eagle, Nathan, and Alex Sandy Pentland. 2006. "Reality mining: sensing complex social systems." *Personal and Ubiquitous Computing*, 10(4): 255–268.

Enikolopov, Ruben, Alexey Makarin, and Maria Petrova. 2020. "Social media and protest participation: Evidence from Russia." *Econometrica*, 88(4): 1479–1514.

Ernala, Sindhu Kiranmai, Moira Burke, Alex Leavitt, and Nicole B Ellison. 2020. "How well do people report time spent on Facebook? An evaluation of established survey questions with recommendations." 1–14.

Ester, Martin, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. 1996. "A density-based algorithm for discovering clusters in large spatial databases with noise." Vol. 96, 226–231.

FTC. 2021. "FTC vs. Facebook 1:20-cv-03590-JEB." `https://www.ftc.gov/system/files/documents/cases/ecf_75-1_ftc_v_facebook_public_redacted_fac.pdf`.

Gerber, Alan S, and Donald P Green. 2012. *Field experiments: Design, analysis, and interpre-*

tation. WW Norton.

Ghose, Anindya, and Sang Pil Han. 2014. "Estimating demand for mobile applications in the new economy." *Management Science*, 60(6): 1470–1488.

Goldfarb, Avi. 2006. "The medium-term effects of unavailability." *Quantitative Marketing and Economics*, 4(2): 143–171.

Gordon, Brett R, and Baohong Sun. 2015. "A dynamic model of rational addiction: Evaluating cigarette taxes." *Marketing Science*, 34(3): 452–470.

Han, Sang Pil, Sungho Park, and Wonseok Oh. 2016. "Mobile app analytics." *MIS Quarterly*, 40(4): 983–1008.

Hoong, Ruru. 2021. "Self control and smartphone use: An experimental study of soft commitment devices." *European Economic Review*, 140: 103924.

Hou, Yubo, Dan Xiong, Tonglin Jiang, Lily Song, and Qi Wang. 2019. "Social media addiction: Its impact, mediation, and intervention." *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*, 13(1).

Ichihashi, Shota, and Byung-Cheol Kim. 2021. "Addictive Platforms."

Kuss, Daria J, and Mark D Griffiths. 2017. "Social networking sites and addiction: Ten lessons learned." *International Journal of Environmental Research and Public Health*, 14(3): 311.

Levy, David M. 2016. *Mindful tech: How to bring balance to our digital lives.* Yale University Press.

Levy, Ro'ee. 2021. "Social media, news consumption, and polarization: Evidence from a field experiment." *American Economic Review*, 111(3): 831–70.

MacKay, Alexander, and Marc Remer. 2022. "Consumer inertia and market power." *Available at SSRN 3380390*.

Morton, Fiona M Scott, and David C Dinielli. 2020. "Roadmap for an antitrust case against facebook." Technical report, Omidyar Network.

Mosquera, Roberto, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie. 2020. "The economic effects of Facebook." *Experimental Economics*, 23(2): 575–602.

Narayanan, Arvind, Arunesh Mathur, Marshini Chetty, and Mihir Kshirsagar. 2020. "Dark Patterns: Past, Present, and Future: The evolution of tricky user interfaces." *Queue*, 18(2): 67–92.

Newman, John M. 2016. "Antitrust in zero-price markets: applications." *Wash. UL Rev.*, 94: 49.

Prat, Andrea, and Tommaso M Valletti. 2021. "Attention oligopoly." *American Economic Journal: Microeconomics, Forthcoming*.

Raval, Devesh, Ted Rosenbaum, and Nathan E Wilson. 2022. "Using disaster-induced closures to evaluate discrete choice models of hospital demand." *The RAND Journal of Economics*, 53(3): 561–589.

Reynolds, Graeme, and Chris Walters. 2008. "The use of customer surveys for market definition and the competitive assessment of horizontal mergers." *Journal of Competition Law and Economics*, 4(2): 411–431.

Scott Morton, Fiona, Pascal Bouvier, Ariel Ezrachi, Bruno Jullien, Roberta Katz, Gene Kimmelman, A Douglas Melamed, and Jamie Morgenstern. 2019. "Committee for the study of digital platforms: Market structure and antitrust subcommittee report." *Chicago: Stigler Center for the Study of the Economy and the State, University of Chicago Booth School of Business*.

Shum, Matthew. 2004. "Does advertising overcome brand loyalty? Evidence from the breakfast-cereals market." *Journal of Economics & Management Strategy*, 13(2): 241–272.

Simonov, Andrey, Jean-Pierre Dubé, Günter Hitsch, and Peter Rossi. 2020. "State-dependent demand estimation with initial conditions correction." *Journal of Marketing Research*, 57(5): 789–809.

Von Luxburg, Ulrike. 2007. "A tutorial on spectral clustering." *Statistics and Computing*, 17(4): 395–416.

Wu, Tim. 2017. "Blind spot: The attention economy and the law." *Antitrust Law Journal*, 82(3).

Yuan, Han. 2020. "Competing for Time: A Study of Mobile Applications." *Mimeo*.

Appendix

A Data Collection Appendix

In this section I provide additional details on the data collection procedures for extracting the required data from the ScreenTime parental control software and the functioning of the Chrome Extension.

Figure A1: Software Reliability

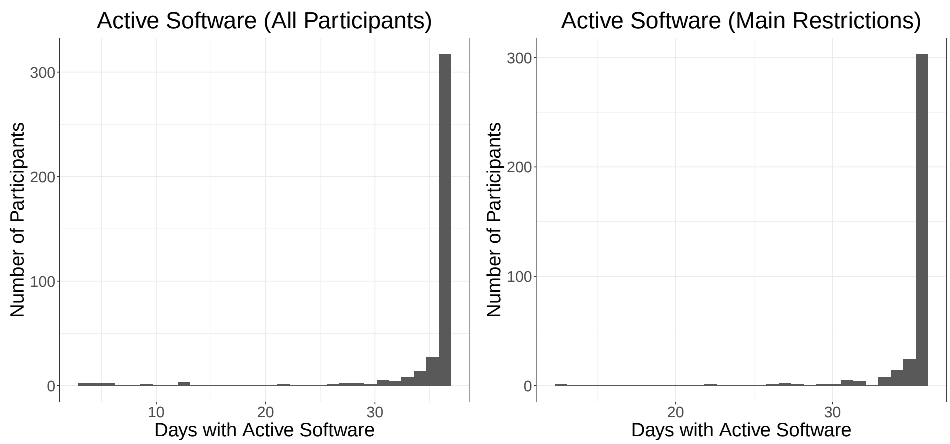

Notes: The figure on the left shows the number of days with active software for all participants, including those who dropped out but whose data I do not drop entirely. The figure on the right shows the number of days with active software for participants in the main experimental group and who stayed through the entirety of the study.

Phone Data: The data from the parental control application was extracted by a script that would run daily at 2 AM EST. There is a maximum of 5 "children" per parental control account and there are a total of 83 separate accounts. The script logs into each account separately and for each "child" it pulls the aggregated and time period data for the previous day. For the subset of devices where it is available, it pulls the web history information which is then used to convert browser time into time on the application that it maps to. The parental control application provides two different aggregations of time allocations for each "child". The first is the aggregated daily usage per application that I utilize in the reduced-form analysis. The second is a breakdown of each application used throughout the day and the precise timing of the sessions. This latter data is used for the model estimation, but is rounded to the nearest minute of the beginning and end of the session. I normalize the session data using the aggregated daily data to ensure consistency. The interface also enables the "parent" to restrict any application on the child's phone. The script ensures that the restrictions for the current child is in place as well as pulls the set of currently installed applications when parsing this list.

At the conclusion of the script, it logs any accounts that logged no data or had abnormally low usage. Typically around 8 AM EST, I manually check these accounts and then reach out to participants who are flagged and ask them to either restart their phone or reinstall the application if it's confirmed to be an issue with the application. When I reach out to a participant, I drop their data from the days where it is determined that the application was not logging properly. The primary reason for the instability is usually based on the device type. Huawei devices have specific settings that need to be turned off in order for the software to run properly. The vast majority of issues with Huawei devices were resolved in the setup period of the study. OnePlus and Redmi devices, however, have a tendency to kill the usage tracking background process unless the application is re-opened every once in a while. As a result, participants with these phones were instructed to do so when possible. Figure A1 plots a histogram of the number of active days with the software working across participants and shows that this issue only impacts a small fraction of participants. Beyond this, I drop two participants entirely from the analysis – one since the scripts failed to detect that they evaded the YouTube restriction and another since a bug with their particular type of phone resulted in no valid baseline data.

Chrome Extension: By default, the Chrome Extension only collects time spent on entertainment and social media domains with the rest of the websites logged under other. In particular, it only logs time spent on the following domains: instagram.com, messenger.com, google.com, facebook.com, youtube.com, tiktok.com, reddit.com, pinterest.com, tumblr.com, amazon.com, twitter.com, pandora.com, spotify.com, netflix.com, hulu.com, disneyplus.com, twitch.tv, hbomax.com. This is made clear to participants during the setup period. Participants can optionally allow time tracking on all websites and can view how much time they've spent on an application in the Chrome Extension itself (see Figure OA7). The time tracking done by the Chrome Extension is crude due to limitations on how Chrome Extensions can interact with the browser. The Chrome Extension script continually runs in the background and wakes up every minute, the lowest possible time interval, observes what page it is on, and then ascribes a minute spent to this page. This process induces some measurement error in recorded time, but gives me a rough approximation of time spent on each domain. The recorded data is continually persisted to a server, which allows me to see what the recorded website was for every minute as well as aggregated by day.

B Descriptive Statistics

Participant Demographics: I report the gender and age of the participants in the study in Table A1 and Table A2 respectively. Given that the participants were recruited primarily through university lab pools, they are younger relative to the national average with an average age of 26 years old and a median age of 23 years old.[39] The participants, especially due to the fact that this study was conducted during the COVID-19 pandemic, were geographically distributed not just around the United States, but also the world.

Time Allocations: Figure A3 plots the distribution of daily phone and computer usage across participants during the baseline period. For both devices, the distribution is right-skewed and usage is quite substantial with participants averaging 3-4 hours of usage on each device per day. When considering the aggregate time spent across the devices, participants spend around 6 hours on average per day across their phone and computer. Figure A4 displays phone usage across the week, indicating that there isn't substantial variation in usage patterns across days. However, there is variation in usage patterns within the day with peak usage around lunch and in the later evening hours. Finally, Figure A5 displays self-reported time allocations throughout the experiment on other forms of media and life activities and shows that they are fairly constant over the course of the experiment.

Table A3 displays the summary statistics of the different phone categories and shows that most of the time on the phone is spent on communication, entertainment, or social media applications. Furthermore, within the set of prominent social media, communication, and entertainment applications there is extensive multi-homing across these applications as observed in Figure A6, which shows that most participants use between 4 and 7 of the applications of interest. Table A4 shows that most participants are mainly consumers of content on applications such as YouTube, Reddit, and TikTok, while they most often post content on Instagram and Snapchat. However, even on these applications, there are not many participants who post at a relatively high frequency.

Table A1: Gender Distribution

Female	Male	Non-Binary
180	216	11

Table A2: Age Distribution

Minimum	25th Percentile	50th Percentile	Mean	75th Percentile	Maximum
18	21	23	25.92	27.0	73

[39]There were some exceptions to this, primarily from participants drawn from the Chicago Booth lab pool which attracts a more representative sample of the population relative to other lab pools. Thus, from this lab pool several older participants were recruited.

Figure A2: Distributions of Application Usage Across Treatment Groups

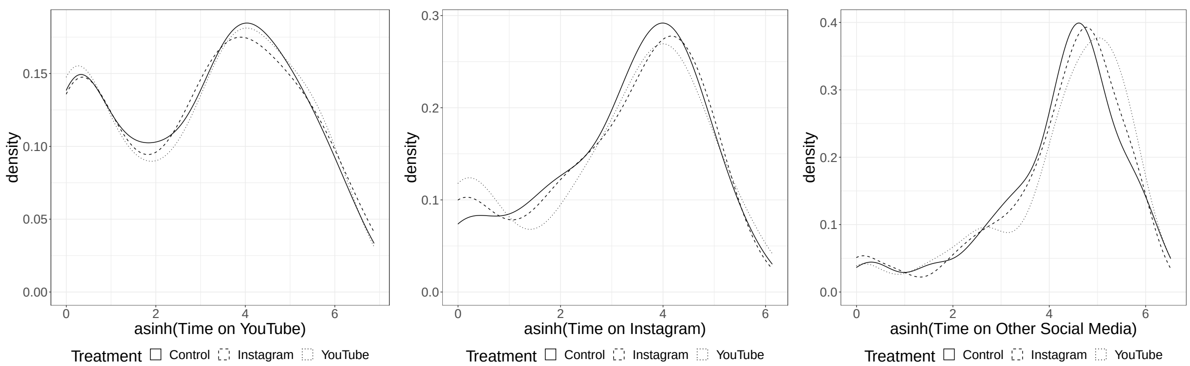

Notes: The figures show the distribution of usage on YouTube (left), Instagram (middle), and other social media (right) during the baseline period across the different experimental treatment groups.

Figure A3: Distribution of Daily Phone Usage

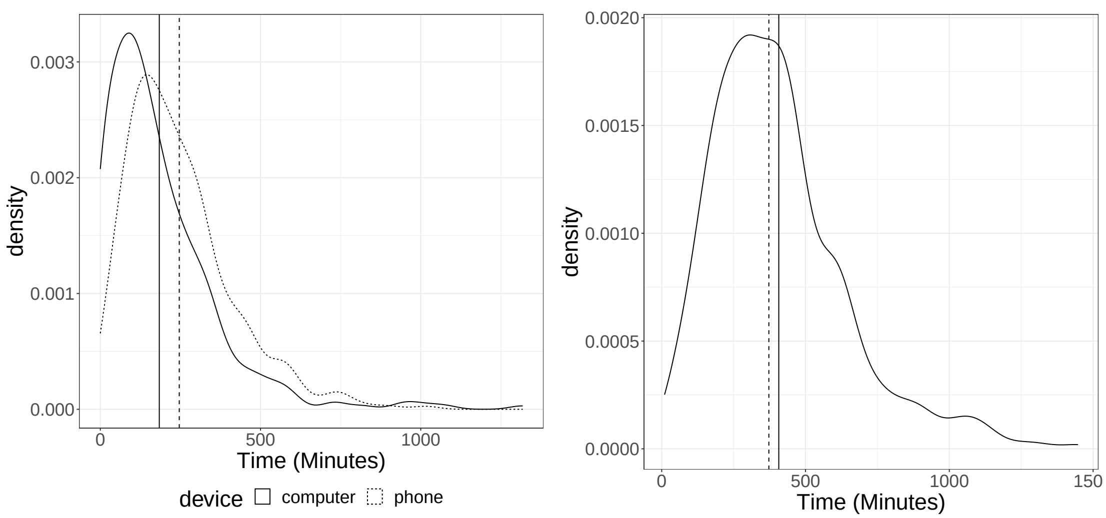

Notes: Both figures plot a kernel density fit of the observed average daily phone usage over the baseline week of the experiment. The figure on the left plots the distribution of phone and computer data separately with the dashed vertical line representing the mean phone time and the solid vertical line representing the mean computer time. The figure on the right displays the distribution of time spent across both computer and phone. The solid line represents the mean time and the dashed line represents the median time.

Figure A4: Time on Phone Across the Week

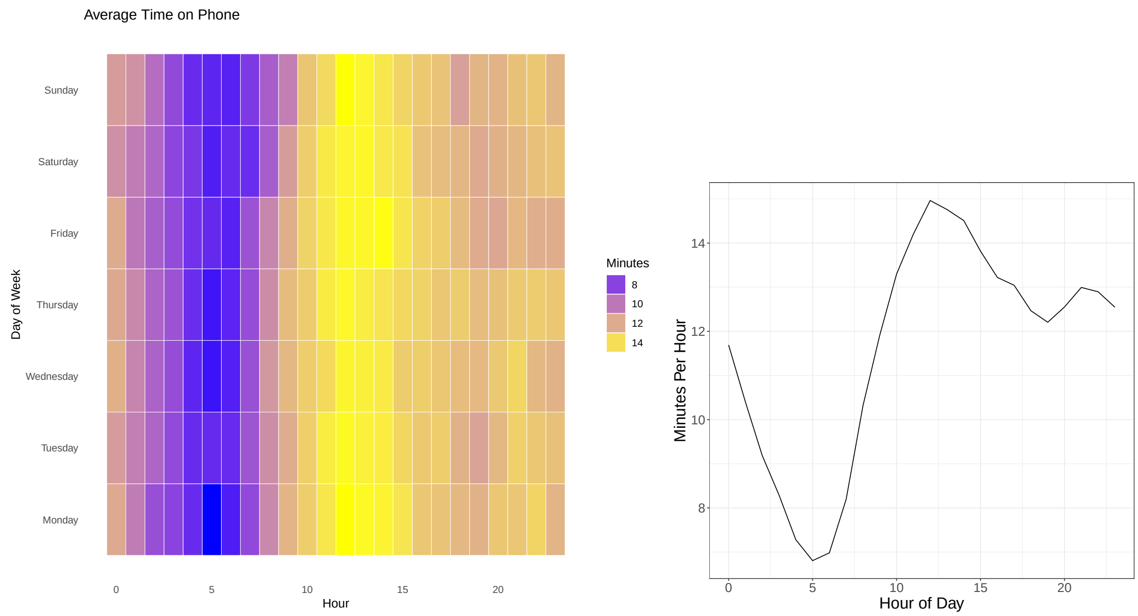

Notes: The figure on the left plots the heatmap of average minutes of usage throughout the entire study period across days of the week and hours of the day. The figure on the right plots the average minutes of usage across hours of the day.

Figure A5: Time Off Digital Devices

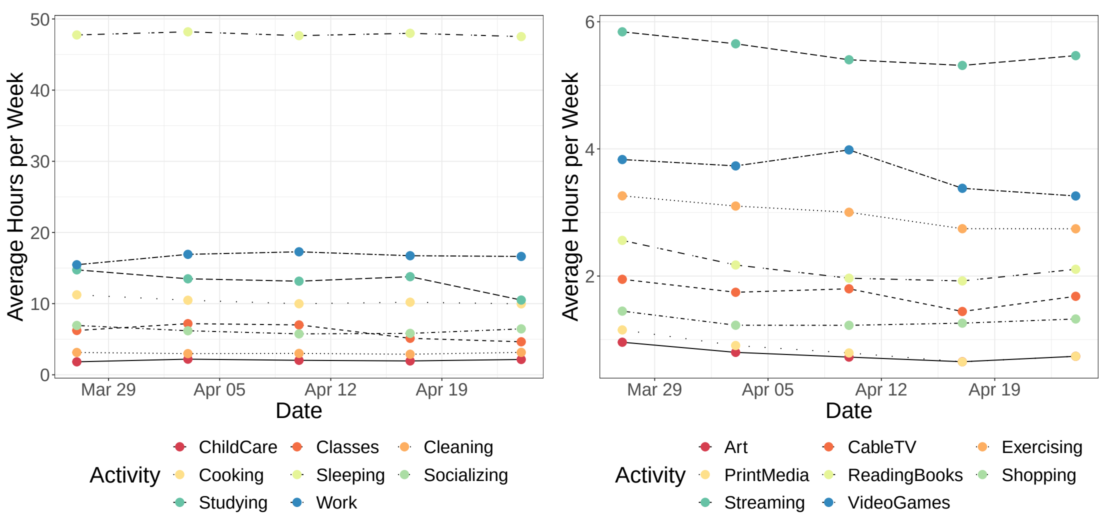

Notes: A single point on the graph represents the average reported hours spent on a category and week. Each reported data point comes from the weekly time use survey filled out by participants. The figure on the left displays the amount of time spent on necessities in life such as sleeping and working. The figure on the right displays the amount of time spent on leisure activities such as streaming movies, reading books, or playing video games.

Table A3: Time Spent on Application Categories on Phone

Category	Average Time	Median Time	Average Time | Usage	Average Time | Usage	Numbers of Users
social	66.38	52.36	68.88	53.75	373
entertainment	55.34	20.54	59.13	24.21	365
communication	54.95	40.86	55.38	41.00	387
game	23.85	0.57	42.38	16.93	175
tools	11.59	6.54	11.74	6.64	385
education	5.28	0.14	8.69	1.00	215
maps	4.52	0.83	6.40	2.07	275
business	4.48	0.50	6.55	2.36	253
productivity	4.33	1.43	4.73	1.64	357
art	3.92	1.43	4.44	1.83	345
news	3.80	0.00	8.51	1.50	130
shopping	3.33	0.29	5.28	1.50	229
sports	3.11	0.07	5.71	1.25	54
lifestyle	2.70	0.14	4.62	0.64	211
finance	2.19	0.71	2.64	1.29	315
dating	2.03	0.07	3.41	0.57	218
food	1.76	0.29	2.80	1.29	189
health	1.60	0.07	3.03	0.43	176
music	1.56	0.00	4.15	0.61	144

Notes: This table displays the time allocations for the product categories on the phone. The product categories are those assigned to the applications in the Google Play Store. I report average daily minutes spent on each category during the baseline week for the days when there were no known issues with application usage logging. The first column displays the name of the category. The second and third columns display the average and median minutes per day, respectively, across all participants. The fourth and fifth columns display the same quantities respectively, but conditional only on the participants that make use of those applications. The sixth column displays the number of participants that use the category.

Table A4: Post Frequency on Applications of Interest

Application	Never	Less Than Once a Month	At least once a month	At least once a week	2 or 3 times per week	Every day
Facebook	0.36	0.41	0.10	0.04	0.04	0.05
Instagram	0.16	0.44	0.20	0.08	0.07	0.05
YouTube	0.81	0.11	0.03	0.02	0.02	0.02
TikTok	0.76	0.13	0.08	0.01	0.01	0.02
Twitter	0.32	0.31	0.10	0.11	0.11	0.05
Snapchat	0.24	0.28	0.09	0.12	0.11	0.16
Reddit	0.51	0.27	0.07	0.07	0.07	0.01

Notes: Each cell represents the fraction of users of the row application that reported the column post frequency. A post means that the participant actively contributes content to the selected application (including ephemeral content such as stories). For each row, I only report the proportion of participants who stated in the survey that they use this application or if there is observed time usage of the application in the baseline period of the study.

Figure A6: Multihoming

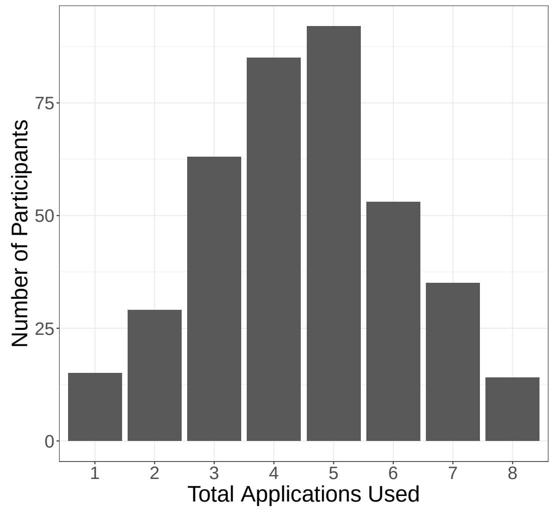

Notes: This figure computes the set of participants that make use of Facebook, Messenger, Instagram, YouTube, Reddit, WhatsApp, TikTok, and Snapchat. It plots how many participants used 1, 2, 3, etc. of these applications over the course of the experiment.

C Additional Experimental Results

Table A5: Stated Substitution Patterns

Application	Social Media	Entertainment	News	Messaging	In-person	Other Hobbies
Facebook	0.33	0.25	0.12	0.11	0.04	0.16
Messenger	0.09	0.10	0.02	0.57	0.11	0.12
Instagram	0.23	0.32	0.05	0.12	0.06	0.22
YouTube	0.10	0.46	0.10	0.02	0.03	0.31
TikTok	0.18	0.43	0.01	0.04	0.04	0.28
WhatsApp	0.10	0.08	0.003	0.55	0.16	0.10
Twitter	0.27	0.11	0.41	0.08	0.04	0.10
Snapchat	0.29	0.08	0.02	0.38	0.10	0.13
Reddit	0.17	0.19	0.36	0.03	0.01	0.22
Netflix	0.07	0.57	0.02	0.03	0.05	0.25

Notes: Each row corresponds to the response for each application about what the participant believes they would substitute their time with if the application was no longer available. Each cell in the row corresponds to the fraction of participants who selected the column option. For each row, I only report the proportion of participants who stated in the survey that they use this application or if there is observed time usage of the application in the baseline period of the study as well as if they did not mark no change in response to the question.

Table A6: Instagram Category Substitution

	Dependent variable:					
	Social	Social (No IG)	Communication	Entertainment	Other	Overall Phone Time
	(1)	(2)	(3)	(4)	(5)	(6)
Category Time	-18.781^{***}	4.273	3.691	-7.454	-6.646	-27.653^{**}
	(4.343)	(3.467)	(3.720)	(5.195)	(5.648)	(12.351)
Category Time - Pooled	-18.557^{***}	4.202^{*}	3.223	-0.607	-3.318	-16.336^{*}
	(3.100)	(2.424)	(2.769)	(3.872)	(4.093)	(9.081)
asinh(Category Time)	-0.461^{***}	0.225^{**}	0.129^{*}	-0.040	-0.105	-0.053
	(0.099)	(0.092)	(0.073)	(0.135)	(0.082)	(0.051)
asinh(Category Time) - Pooled	-0.594^{***}	0.227^{***}	0.104^{*}	0.071	-0.037	-0.047
	(0.100)	(0.076)	(0.057)	(0.098)	(0.064)	(0.048)
Category Share	-0.059^{***}	0.048^{***}	0.052^{***}	0.006	0.0005	-
	(0.014)	(0.013)	(0.013)	(0.016)	(0.015)	
Category Share - Pooled	-0.065^{***}	0.042^{***}	0.054^{***}	0.006	0.010	-
	(0.013)	(0.011)	(0.011)	(0.012)	(0.012)	

*p<0.1; **p<0.05; ***p<0.01

Notes: This regression reports the average treatment effect of average daily minutes spent on applications in different categories during the Instagram restriction. I only consider participants with software active at least 3 days in the baseline and treatment periods. The columns show time spent on social, social (without Instagram), communication, entertainment, other categories, and overall phone time respectively. The entertainment category includes applications marked as entertainment or video players/editors. The column with social (without Instagram) aggregates social time excluding time spent on Instagram, both in the baseline and treatment periods. The first, third, and fifth rows display the primary specification estimated on data from the main experiment. The reported standard errors for these regressions are heteroskedasticity-robust standard errors. The second, fourth, and sixth rows display the primary specification estimated on data pooled from the main experiment and the pilot experiment. The reported standard errors for these regressions are clustered standard errors at the individual level, to accommodate the multiple treatments during the pilot study. The category share row measures the on phone share of time spent on the category.

Table A7: YouTube Category Substitution

	Dependent variable:					
	Social	Communication	Entertainment	Entertainment (No YT)	Other	Overall Phone Time
	(1)	(2)	(3)	(4)	(5)	(6)
Category Time	2.901	-0.143	-43.676^{***}	1.634	-4.050	-44.204^{***}
	(4.471)	(3.695)	(6.788)	(3.984)	(6.746)	(14.409)
Category Time - Pooled	3.989	-2.566	-46.685^{***}	-3.608	-4.277	-51.381^{***}
	(2.909)	(3.346)	(5.686)	(2.917)	(4.621)	(11.282)
asinh(Category Time)	0.164^{*}	0.017	-1.609^{***}	0.176	-0.052	-0.151^{***}
	(0.084)	(0.069)	(0.160)	(0.142)	(0.074)	(0.051)
asinh(Category Time) - Pooled	0.151^{**}	-0.041	-1.484^{***}	0.049	-0.054	-0.154^{***}
	(0.067)	(0.051)	(0.123)	(0.112)	(0.063)	(0.045)
Category Share	0.056^{***}	0.044^{***}	-0.138^{***}	0.011	0.035^{**}	-
	(0.014)	(0.012)	(0.016)	(0.009)	(0.015)	
Category Share - Pooled	0.056^{***}	0.027^{***}	-0.129^{***}	0.007	0.043^{***}	-
	(0.012)	(0.009)	(0.013)	(0.008)	(0.012)	

*p<0.1; **p<0.05; ***p<0.01

Notes: This regression reports the average treatment effect of average daily minutes spent on applications in different categories during the YouTube restriction. I only consider participants with software active at least 3 days in the baseline and treatment periods. The columns show time spent on social, communication, entertainment, entertainment (without YouTube), other categories, and overall phone time respectively. The entertainment category includes applications marked as entertainment or video players/editors. The column with entertainment (without YouTube) aggregates entertainment time excluding time spent on YouTube, both in the baseline and treatment periods. The first, third, and fifth rows display the primary specification estimated on data from the main experiment. The reported standard errors for these regressions are heteroskedasticity-robust standard errors. The second, fourth, and sixth rows display the primary specification estimated on data pooled from the main experiment and the pilot experiment. The reported standard errors for these regressions are clustered standard errors at the individual level, to accommodate the multiple treatments during the pilot study. The category share row measures the on phone share of time spent on the category.

Figure A7: Quantile Treatment Effects of Category Substitution

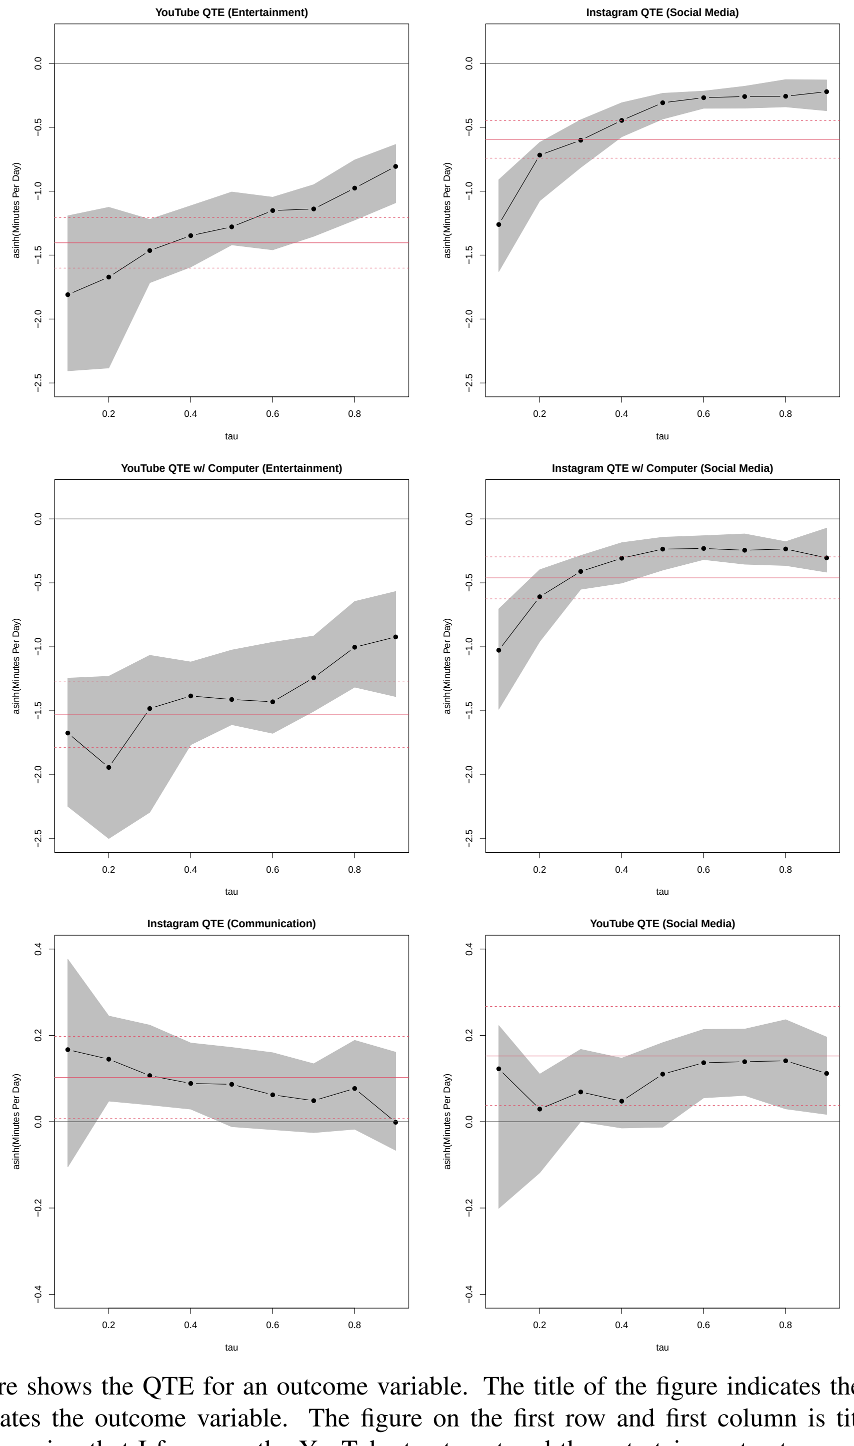

Notes: Each figure shows the QTE for an outcome variable. The title of the figure indicates the treatment and the parentheses indicates the outcome variable. The figure on the first row and first column is titled YouTube QTE (Entertainment) meaning that I focus on the YouTube treatment and the entertainment category. The figures in the middle row include time from the Chrome extension, whereas the rest only include time from the phone.

Table A8: Instagram Type of App Substitution

	Dependent variable:			
	Facebook Ecosystem	Facebook Ecosystem (No IG)	Major	Minor
	(1)	(2)	(3)	(4)
Category Time	-21.776^{***}	1.905	-4.486	-3.270
	(4.107)	(3.180)	(6.062)	(5.230)
Category Time - Pooled	-21.494^{***}	0.842	2.352	2.378
	(3.185)	(2.604)	(4.527)	(3.915)
asinh(Category Time)	-0.573^{***}	0.163	0.062	0.015
	(0.121)	(0.108)	(0.097)	(0.089)
asinh(Category Time) - Pooled	-0.644^{***}	0.174^{**}	0.076	0.080
	(0.099)	(0.075)	(0.075)	(0.070)
Category Share	-0.059^{***}	0.051^{***}	0.027^{*}	0.024
	(0.015)	(0.014)	(0.016)	(0.016)
Category Share - Pooled	-0.067^{***}	0.044^{***}	0.024^{*}	0.029^{**}
	(0.012)	(0.011)	(0.013)	(0.012)

*p<0.1; **p<0.05; ***p<0.01

Notes: I consider the degree of substitution to Facebook-owned applications (WhatsApp, Facebook, Instagram, Messenger), "major" applications (Reddit, YouTube, TikTok, Netflix, Twitter, Discord, Snapchat, Twitch, LinkedIn, Spotify, Zoom, Telegram, Hulu, Prime Video, Signal, Google, Amazon Shopping), and the rest of the applications (excluding core phone applications). Each cell reports the estimated average treatment effect on average daily minutes for the participants who have the software active for at least 3 days in the baseline and restriction periods. The first, third, and fifth rows display the primary specification estimated on data from the main experiment with heteroskedacity-robust standard errors reported in parentheses. The second, fourth, and sixth rows display the primary specification estimated on data pooled from the main experiment and the pilot experiment with standard errors clustered at the individual level reported in parentheses. The category share row measures the on phone share of time spent on the category.

Table A9: YouTube Type of Application Substitution

	Dependent variable:			
	Facebook Ecosystem	Major	Major (No YT)	Minor
	(1)	(2)	(3)	(4)
Category Time	4.132	-46.954^{***}	-0.862	1.108
	(4.279)	(7.017)	(4.991)	(6.597)
Category Time - Pooled	1.981	-47.806^{***}	-4.005	-0.285
	(4.009)	(6.176)	(3.876)	(4.616)
asinh(Category Time)	0.026	-0.711^{***}	0.033	0.034
	(0.082)	(0.113)	(0.100)	(0.080)
asinh(Category Time) - Pooled	0.029	-0.689^{***}	0.054	0.026
	(0.069)	(0.088)	(0.077)	(0.065)
Category Share	0.060^{***}	-0.123^{***}	0.021	0.055^{***}
	(0.013)	(0.016)	(0.013)	(0.016)
Category Share - Pooled	0.044^{***}	-0.104^{***}	0.028^{**}	0.057^{***}
	(0.011)	(0.014)	(0.012)	(0.012)

*p<0.1; **p<0.05; ***p<0.01

Notes: I consider the degree of substitution to Facebook-owned applications (WhatsApp, Facebook, Instagram, Messenger), "major" applications (Reddit, YouTube, TikTok, Netflix, Twitter, Discord, Snapchat, Twitch, LinkedIn, Spotify, Zoom, Telegram, Hulu, Prime Video, Signal, Google, Amazon Shopping), and the rest of the applications (excluding core phone applications). Each cell reports the estimated average treatment effect on average daily minutes for the participants who have the software active for at least 3 days in the baseline and restriction periods. The first, third, and fifth rows display the primary specification estimated on data from the main experiment with heteroskedacity-robust standard errors reported in parentheses. The second, fourth, and sixth rows display the primary specification estimated on data pooled from the main experiment and the pilot experiment with standard errors clustered at the individual level reported in parentheses. The category share row measures the on phone share of time spent on the category.

Table A10: Survey of Time on Restricted App During Treatment Week Off Phone

	Dependent variable:			
	Other Device Instagram Time	Other Device YouTube Time	asinh(Other Device Instagram Time)	asinh(Other Device YouTube Time)
	(1)	(2)	(3)	(4)
YouTube Treatment		−8.151		−0.409**
		(6.850)		(0.207)
Instagram Treatment	−1.941		−0.042	
	(1.964)		(0.181)	
Baseline Time Controls	Yes	Yes	Yes	Yes
Block Controls	Yes	Yes	Yes	Yes
Observations	231	238	231	238
R^2	0.103	0.182	0.316	0.311

*p<0.1; **p<0.05; ***p<0.01

Notes: Heteroskedasticity-robust standard errors are reported in parentheses. The first and third columns present the results of a regression of self-reported average daily minutes on Instagram on other devices between the Instagram restriction group and the control group. The second and fourth columns present the results of a regression of self-reported average daily minutes on YouTube on other devices between the YouTube restriction group and the control group.

Figure A8: Time Spent on Phone Throughout the Week (During Treatment Period)

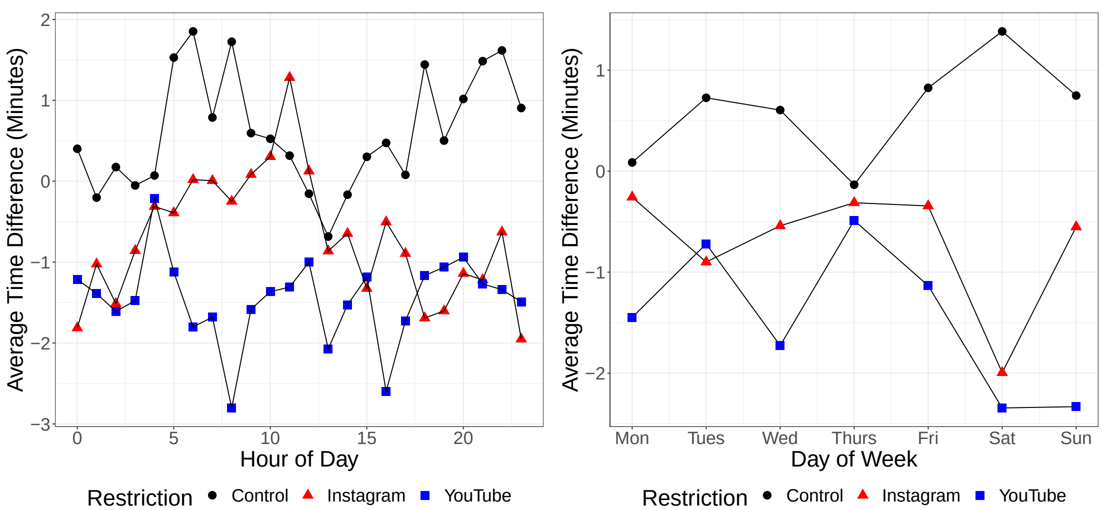

Notes: The figures plot the difference between the first week and the treatment week for each treatment group. The figure on the left plots the difference across different hours of the day and the figure on the right plots the difference across different days of the week.

Figure A9: Quantile Treatment Effects of Overall Phone Time

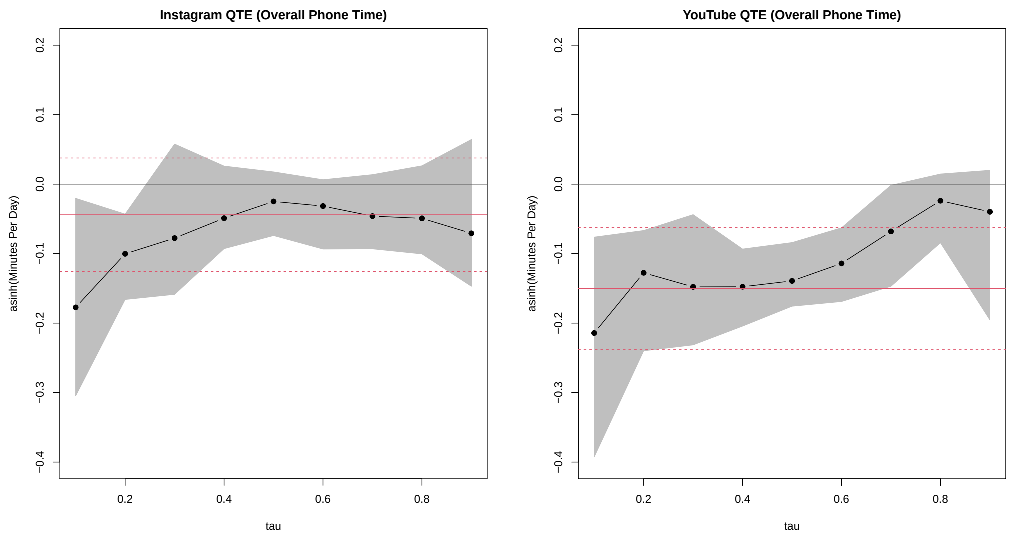

Notes: The figures present the estimated QTE of the log of overall phone usage across both treatment groups during the restriction period.

Table A11: Survey of Time Spent on Other Media During Restriction Period

	Dependent variable:			
	asinh(Time on Cable TV)	asinh(Time on Video Games)	asinh(Time on Streaming Services)	asinh(Time on Other Media Composite)
	(1)	(2)	(3)	(4)
YouTube Treatment	0.015	0.258	−0.381	−0.076
	(0.185)	(0.205)	(0.248)	(0.208)
Instagram Treatment	−0.290	0.217	−0.292	−0.079
	(0.187)	(0.207)	(0.251)	(0.210)
Baseline Time Controls	Yes	Yes	Yes	Yes
Block Controls	Yes	Yes	Yes	Yes
Observations	357	357	357	357
R^2	0.471	0.565	0.344	0.386

*p<0.1; **p<0.05; ***p<0.01

Notes: Heteroskedasticity-robust standard errors are reported in parentheses. This table reports the estimated ATE on time spent on non-phone media during the restriction period. The data for this come from the weekly time use survey. The first column reports the impact of the treatment on average daily minutes on cable TV. The second column reports the impact of the treatment on average daily time on video games. The third column reports the impact of the treatment on average daily time on non-phone video streaming services. The fourth column reports the impact of the treatment on the sum of the average daily time on cable TV, video games, and non-phone video streaming services.

Figure A10: Time on Restricted Applications

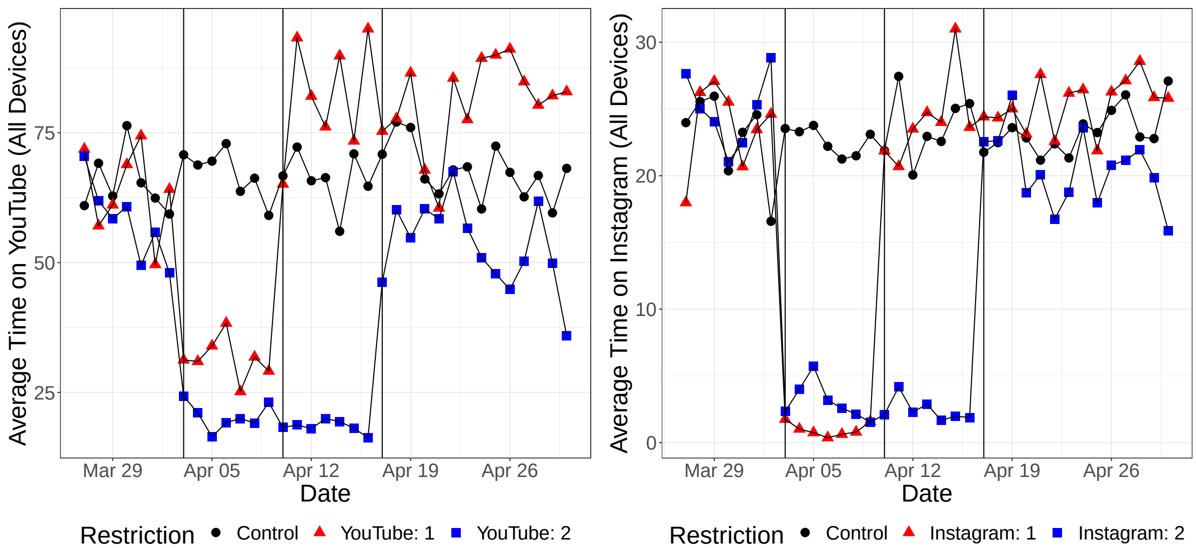

Notes: This figure plots the average daily minutes on the restricted applications on the phone and computer across the different treatment groups for the YouTube (left) and Instagram (right) restriction group.

Table A12: Instagram Post-Restriction Usage

	Dependent variable:			
	Instagram Time	asinh(Instagram Time)	Instagram Time	asinh(Instagram Time)
	(1)	(2)	(3)	(4)
Instagram Treatment	4.845	0.177	-5.164^{**}	-0.061
	(3.438)	(0.166)	(2.483)	(0.134)
2 week restriction	3.180	0.038		
	(3.048)	(0.179)		
Instagram Treatment $\times$ 2 week restriction	-10.452^{**}	-0.231		
	(4.746)	(0.232)		
Baseline Usage	Yes	Yes	Yes	Yes
Week Fixed Effects	Yes	Yes	Yes	Yes
Block Controls	Yes	Yes	Yes	Yes
Observations	410	410	312	312
R^2	0.696	0.731	0.707	0.732

*p<0.1; **p<0.05; ***p<0.01

Notes: The standard errors for the regression are clustered at the participant level. The regression is estimated on the data of average daily minutes of Instagram in the two weeks following the restriction period. The dependent variables reported are both the levels and logs of Instagram usage. The first two columns report the regression across all restriction groups with heterogeneous effects across restriction lengths. The last two columns report the regression on the entire control group and restricting focus to the 2 week Instagram restriction group.

Table A13: YouTube Post-Restriction Usage

	Dependent variable:			
	YouTube Time	asinh(YouTube Time)	YouTube Time	asinh(YouTube Time)
	(1)	(2)	(3)	(4)
YouTube Treatment	1.387	−0.047	−8.930	−0.167
	(10.374)	(0.162)	(6.747)	(0.190)
2 week restriction	−6.974	−0.030		
	(9.021)	(0.213)		
YouTube Treatment × 2 week restriction	−6.640	0.004		
	(10.639)	(0.273)		
Baseline Usage	Yes	Yes	Yes	Yes
Week Fixed Effects	Yes	Yes	Yes	Yes
Block Controls	Yes	Yes	Yes	Yes
Observations	482	482	360	360
R^2	0.558	0.674	0.531	0.619

*p<0.1; **p<0.05; ***p<0.01

Notes: The standard errors for the regression are clustered at the participant level. The regression is estimated on the data of average daily minutes of YouTube in the two weeks following the restriction period. The dependent variables reported are both the levels and logs of YouTube usage. The first two columns report the regression across all restriction groups with heterogeneous effects across restriction lengths. The last two columns report the regression on the entire control group and restricting focus to the 2 week YouTube restriction group.

Figure A11: Quantile Treatment Effects of Post-Restriction Usage

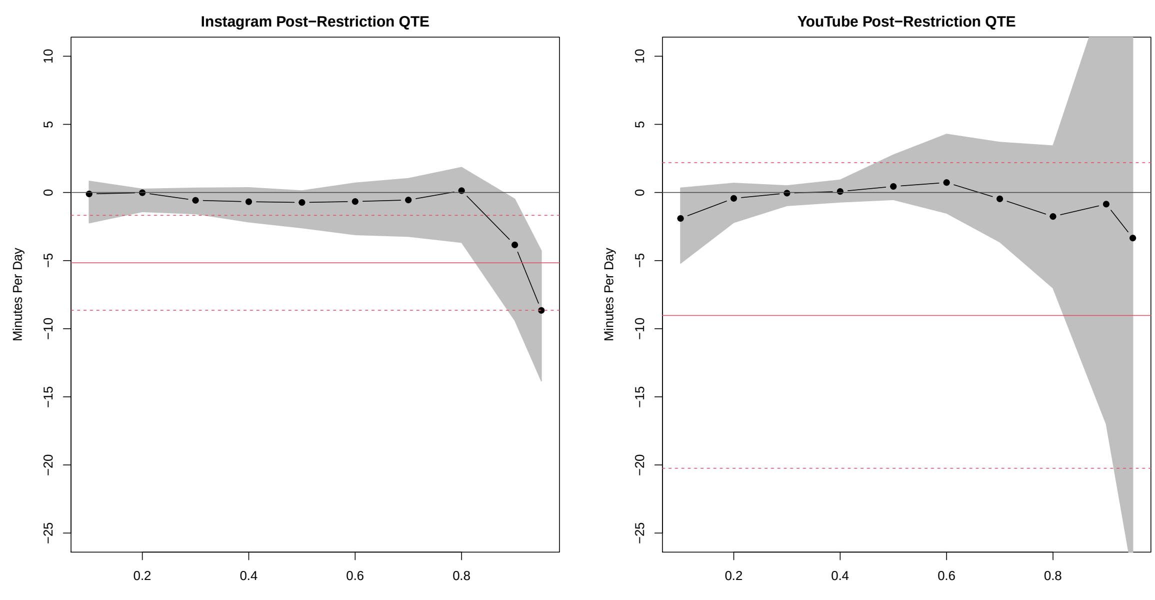

Notes: The figures present the estimated QTE of post-restriction usage on the restricted applications across both treatment groups.

Table A14: Perceived Endline Substitution Patterns

Restricted Application	New Apps	Invested in Other Apps	Time on Other Apps	Computer Time	Offline	No Change
During Restriction - Instagram	0.05	0.19	0.26	0.20	0.18	0.11
After Restriction - Instagram	0.04	0.08	0.16	0.17	0.15	0.41
During Restriction - YouTube	0.10	0.15	0.30	0.22	0.15	0.08
After Restriction - YouTube	0.05	0.11	0.13	0.17	0.13	0.41

Notes: This table shows the proportion of participants in each treatment group that report their perceived substitution during the experiment. The first and third rows show the perceived changes in behavior during the restriction period. The second and fourth rows show the perceived changes in behavior following the restriction period. Column 2 represents primary substitution towards newly installed applications. Column 3 represents primary substitution towards installed applications that participants "invested" in sourcing better content from. Column 4 represents primary substitution towards other installed applications but without significant additional "investment" in them. Column 5 represents primary substitution towards the computer. Column 6 represents primary substitution towards non-digital activities. Column 7 represents no change in behavior.

Table A15: One Month Post-Experiment Survey Results

	Dependent variable:			
	Phone Time	Social Media Time	Instagram Time	YouTube Time
	(1)	(2)	(3)	(4)
Instagram Restriction	−0.115	-0.305^{**}	-0.316^{*}	−0.113
	(0.147)	(0.150)	(0.189)	(0.178)
YouTube Restriction	0.087	−0.003	0.189	−0.268
	(0.145)	(0.148)	(0.186)	(0.176)
Block Controls	No	No	No	No
Constant	2.811^{***}	2.698^{***}	2.756^{***}	3.113^{***}
	(0.106)	(0.107)	(0.137)	(0.127)
Observations	168	168	149	167
R^2	0.012	0.033	0.051	0.014

*p<0.1; **p<0.05; ***p<0.01

Notes: Heteroskedasticity-robust standard errors are reported in parentheses. The data comes from the survey sent one month after the study concluded where participants indicated whether they were spending significantly less time (1), somewhat less time (2), the same time (3), somewhat more time (4), or significantly more time (5) on each outcome variable. The dependent variable in column 1 is the overall phone time, in column 2 is overall social media time, in column 3 is Instagram time, and column 4 is YouTube time. For the YouTube and Instagram time dependent variables, I drop participants who marked that they do not use the respective application or started to use it during the study.

Table A16: Instagram Post-Restriction Usage of Non-Restricted Applications

	Dependent variable:			
	Time	asinh(Time)	Time: 2 Weeks	asinh(Time): 2 Weeks
	(1)	(2)	(3)	(4)
Social Category	-1.960	-0.067	-1.037	-0.133
	(4.305)	(0.079)	(5.855)	(0.109)
Communication Category	-3.395	0.049	0.435	0.071
	(3.833)	(0.073)	(4.644)	(0.090)
TikTok	0.367	-0.199	5.876	-0.222
	(7.047)	(0.196)	(9.376)	(0.247)
Facebook	1.063	0.032	1.198	0.117
	(2.516)	(0.137)	(3.213)	(0.161)
Snapchat	-0.053	-0.063	-0.532	-0.156
	(1.049)	(0.101)	(1.520)	(0.145)
WhatsApp	-2.585	0.101	-1.315	0.259^{*}
	(2.838)	(0.118)	(3.401)	(0.146)
Messenger	-0.161	-0.010	-0.008	0.017
	(0.943)	(0.116)	(1.345)	(0.154)
YouTube	-6.288	-0.177	-5.244	-0.176
	(6.843)	(0.141)	(7.383)	(0.183)
Apps Installed During Restriction	4.293	0.262	1.554	0.123
	(2.605)	(0.194)	(1.993)	(0.255)

*p<0.1; **p<0.05; ***p<0.01

Notes: Reported standard errors in parentheses are clustered at the participant level. This table presents the ATE estimates of usage on the row application / category after the restriction period. A single data point is the average daily minutes on the row application one or two weeks following the restriction. The first two rows consider the average time on social and communication categories respectively. The following rows consider the average time on the specified application. The final row considers the average time on the applications that were installed during the restriction period. The first two columns report the ATE on time usage both in levels and logs. The final two columns report the ATE on time usage restricting to the 2 week restriction group.

Table A17: YouTube Post-Restriction Usage of Non-Restricted Applications

	Dependent variable:			
	Time	asinh(Time)	Time: 2 Weeks	asinh(Time): 2 Weeks
	(1)	(2)	(3)	(4)
Social Category	0.541	0.014	1.951	-0.028
	(4.182)	(0.081)	(6.208)	(0.115)
Entertainment Category	-4.405	-0.133	-12.623	-0.282^{*}
	(7.120)	(0.121)	(8.024)	(0.170)
TikTok	-3.313	-0.041	-6.529	0.012
	(8.271)	(0.288)	(12.526)	(0.371)
Facebook	-0.735	-0.031	-0.954	0.068
	(2.010)	(0.127)	(2.522)	(0.146)
Instagram	3.999^{*}	0.190^{*}	2.876	0.161
	(2.262)	(0.112)	(3.276)	(0.133)
Snapchat	0.473	0.003	0.747	0.040
	(1.115)	(0.114)	(1.628)	(0.158)
WhatsApp	-0.253	0.091	0.725	0.156
	(2.483)	(0.101)	(3.233)	(0.122)
Apps Installed During Restriction	3.366^{**}	0.406^{**}	1.990	0.201
	(1.450)	(0.186)	(1.837)	(0.243)

*p<0.1; **p<0.05; ***p<0.01

Notes: Reported standard errors in parentheses are clustered at the participant level. This table presents the ATE of usage on the row application / category after the restriction period. A single data point is the average daily minutes on the row application one or two weeks following the restriction. The first two rows consider the average time on social and entertainment categories respectively. The following rows consider the average time on the specified application. The final row considers the average time on the applications that were installed during the restriction period. The first two columns report the ATE on time usage both in levels and logs. The final two columns report the ATE on time usage restricting to the 2 week restriction group.

D Additional Figures / Tables for Time Usage Model

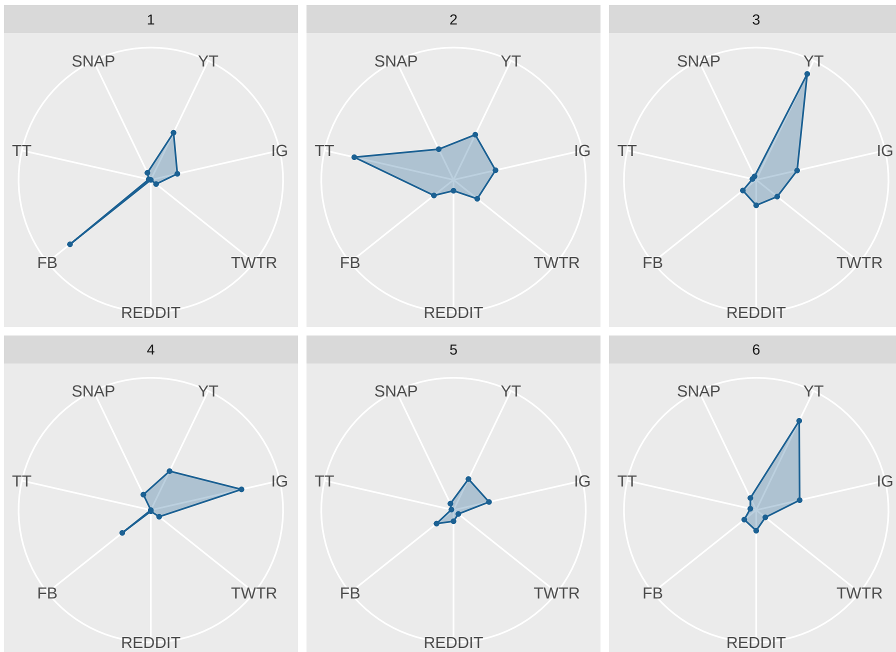

Notes: The figures display the results of k-means clustering for $k = 6$. Each pane shows the average (log) time allocations across the different applications for the participants in the cluster. For instance, if a point is closer to the outer edge for an application A than application B then that indicates that that application A has more usage on average than application B. The application names are abbreviated so that the figure is readable. TT is TikTok, FB is Facebook, YT is YouTube, IG is Instagram, SNAP is SnapChat, REDDIT is Reddit, and TWTR is Twitter.

Table A18: Demand Model Parameter Estimates

Type	(1)	(2)	(3)	(4)	(5)	(6)
h_{ij}	0.00066	0.00074	0.00032	0.00058	0.00079	0.00048
	(1.2e-5)	(1.1e-5)	(1.0e-5)	(9.6e-6)	(7.0e-6)	(3.7e-6)
r_{ij}	1.0	1.0	0.8	1.2	1.4	0.97
	(0.018)	(0.019)	(0.01)	(0.013)	(0.01)	(0.0088)
r_{ij}^2	-0.016	-0.0088	-0.0092	-0.021	-0.012	-0.0055
	(0.00059)	(0.00027)	(0.00015)	(0.00059)	(0.00032)	(0.00015)
asinh(c_{ij})	0.11	0.022	0.13	0.17	0.15	0.072
	(0.0056)	(0.007)	(0.0061)	(0.0044)	(0.0023)	(0.0024)
App - Facebook	-6.0	-6.5	-7.9	-8.1	-7.8	-6.8
	(0.072)	(0.088)	(0.1)	(0.078)	(0.036)	(0.047)
App - Instagram	-5.9	-6.4	-8.3	-7.6	-7.4	-6.0
	(0.072)	(0.088)	(0.11)	(0.079)	(0.035)	(0.048)
App - Reddit	-8.3	-8.0	-8.4	-8.9	-7.4	-6.7
	(0.17)	(0.11)	(0.12)	(0.095)	(0.036)	(0.048)
App - Snapchat	-6.5	-6.6	-9.6	-8.0	-7.4	-6.0
	(0.077)	(0.079)	(0.14)	(0.075)	(0.033)	(0.044)
App - TikTok	-7.8	-6.6	-8.7	-8.1	-7.6	-6.2
	(0.17)	(0.095)	(0.11)	(0.08)	(0.032)	(0.045)
App - Twitter	-6.4	-6.5	-9.7	-8.4	-7.9	-6.8
	(0.07)	(0.086)	(0.13)	(0.076)	(0.035)	(0.051)
App - YouTube	-6.0	-6.9	-8.1	-8.0	-7.7	-6.2
	(0.072)	(0.088)	(0.12)	(0.078)	(0.035)	(0.047)
a_{ij} - Communicate with my friends	0.0	0.0	0.0	0.0	0.0	0.0
	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)
a_{ij} - Entertainment content	0.14	0.03	1.8	-0.08	0.36	0.63
	(0.026)	(0.039)	(0.082)	(0.028)	(0.014)	(0.019)
a_{ij} - Get Information	0.34	-0.22	1.6	-0.00085	0.029	0.67
	(0.03)	(0.05)	(0.088)	(0.029)	(0.017)	(0.022)
a_{ij} - Keep up with my friends' lives	-0.23	-0.33	1.0	-0.11	0.2	0.18
	(0.039)	(0.042)	(0.09)	(0.028)	(0.015)	(0.021)
a_{ij} - Online Shopping	0.0	0.0	-11.0	-0.02	-0.071	0.68
	(0.0)	(0.0)	(1.5)	(0.082)	(0.048)	(0.036)
h_t - Afternoon	-0.98	-1.2	-1.9	-1.5	-1.3	-0.77
	(0.037)	(0.034)	(0.048)	(0.037)	(0.019)	(0.023)
h_t - Late Night	0.0	0.0	0.0	0.0	0.0	0.0
	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)
h_t - Morning	-0.61	-0.6	-1.6	-1.3	-1.1	-0.52
	(0.038)	(0.034)	(0.051)	(0.041)	(0.019)	(0.021)
h_t - Evening	-0.79	-1.3	-1.9	-1.4	-1.2	-0.7
	(0.035)	(0.035)	(0.048)	(0.04)	(0.019)	(0.018)
d_t - Monday	-0.19	-0.74	-0.48	-1.1	-0.54	-0.51
	(0.036)	(0.046)	(0.036)	(0.047)	(0.021)	(0.027)
d_t - Tuesday	-0.27	-0.81	-0.43	-0.98	-0.46	-0.45
	(0.043)	(0.049)	(0.043)	(0.048)	(0.019)	(0.03)
d_t - Wednesday	-0.2	-0.79	-0.38	-0.97	-0.46	-0.49
	(0.042)	(0.05)	(0.043)	(0.047)	(0.021)	(0.03)
d_t - Thursday	-0.2	-0.74	-0.35	-1.0	-0.44	-0.45
	(0.047)	(0.043)	(0.044)	(0.045)	(0.023)	(0.026)
d_t - Friday	-0.29	-0.71	-0.43	-1.1	-0.44	-0.41
	(0.037)	(0.049)	(0.043)	(0.045)	(0.022)	(0.028)
d_t - Saturday	-0.27	-0.71	-0.42	-0.99	-0.52	-0.51
	(0.042)	(0.051)	(0.046)	(0.044)	(0.021)	(0.031)
d_t - Sunday	0.0	0.0	0.0	0.0	0.0	0.0
	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)
w_t - Week 1	-0.18	-0.69	-0.25	-0.59	-0.43	-0.45
	(0.04)	(0.035)	(0.029)	(0.031)	(0.014)	(0.02)
w_t - Week 2	-0.18	-0.55	-0.3	-0.54	-0.43	-0.45
	(0.032)	(0.032)	(0.029)	(0.027)	(0.011)	(0.02)
w_t - Week 3	-0.26	-0.56	-0.4	-0.69	-0.41	-0.42
	(0.031)	(0.037)	(0.03)	(0.029)	(0.012)	(0.02)
w_t - Week 4	-0.27	-0.47	-0.37	-0.69	-0.35	-0.45
	(0.034)	(0.032)	(0.028)	(0.031)	(0.013)	(0.021)
w_t - Week 5	0.0	0.0	0.0	0.0	0.0	0.0
	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)

Notes: This table presents the estimated parameters of the demand model. The estimates for each type are presented in a separate column. Standard errors in parentheses are computed by 50 bootstrap samples.

Table A19: Model Validation

Application	Baseline (Predicted)	Baseline (Actual)	Instagram Restriction (Predicted)	Instagram Restriction (Actual)	YouTube Restriction (Predicted)	YouTube Restriction (Actual)
Instagram	0.0237	0.0234	-	-	0.0266	0.0273
YouTube	0.0327	0.0311	0.0376	0.0353	-	-
Facebook	0.0117	0.0116	0.0106	0.0122	0.0125	0.0121
Reddit	0.00625	0.00565	0.00912	0.00876	0.00634	0.00669
Snapchat	0.00534	0.00514	0.00612	0.00599	0.00606	0.00508
Twitter	0.00396	0.00428	0.00509	0.00504	0.00364	0.00432
TikTok	0.00712	0.00717	0.00845	0.00824	0.00683	0.00672
Outside Option	0.909	0.912	0.923	0.924	0.938	0.937

Notes: Columns 1 and 2 compare the true market shares in week 1, 4, 5 to the predicted market shares from this model during this time period. Columns 3 and 4 compare the true to predicted market shares in the week 2 restriction period for the Instagram restriction group. Columns 5 and 6 compare the true to predicted market shares in the week 2 restriction period for the YouTube restriction group.

Table A20: Second-Choice Diversion Ratios (No Inertia)

	Facebook	Instagram	Reddit	Snapchat	TikTok	Twitter	YouTube	Outside Option
Facebook	-	0.0162	0.00302	0.00479	0.00283	0.00352	0.0123	0.957
Instagram	0.0089	-	0.00305	0.00579	0.00325	0.00388	0.0127	0.962
Reddit	0.00541	0.0103	-	0.00461	0.00201	0.00296	0.0155	0.959
Snapchat	0.00784	0.0167	0.00401	-	0.00688	0.00388	0.0137	0.947
TikTok	0.00833	0.016	0.00288	0.0113	-	0.00637	0.0133	0.942
Twitter	0.00821	0.0164	0.00397	0.00543	0.00498	-	0.0136	0.947
YouTube	0.00685	0.0136	0.005	0.00452	0.00294	0.00319	-	0.964

Notes: This table displays the estimated second-choice diversion ratios that come from the estimated model with $\beta^{q(i)} = 0$. The cell in each row k and column j is computed by $D_{kj} = \frac{s_j(\mathcal{J} \setminus \{k\}) - s_j(\mathcal{J})}{s_k(\mathcal{J})}$.

Table A21: Percentage Change in Diversion Ratio (No Inertia)

	Facebook	Instagram	Reddit	Snapchat	TikTok	Twitter	YouTube	Outside Option
Facebook	-	-33%	3.3%	-4.9%	-30.1%	-27.5%	-34.1%	1.8%
Instagram	-29.2%	-	-1.9%	-0.20.5%	-38.7%	-21.5%	-35.5%	1.6%
Reddit	-8.6%	-14.0%	-	-20.5%	-15.5%	1.4%	-46%	1.8%
Snapchat	-20.8%	-37.8%	-28.4%	-	-53.5%	-9.1%	-36.0%	3.3%
TikTok	-19.0%	-32.1%	2.4%	-27.1%	-	-30.1%	-56.3%	3.7%
Twitter	-37.1%	-36.0%	-4.7%	-6.4%	-50.1%	-	-37.1%	3.0%
YouTube	-27.4%	-31.7%	-28.5%	-13.0%	-52.9%	-18.9%	-	1.7%

Notes: This table presents the percentage change in the second-choice diversion ratios when $\beta^{q(i)} = 0$ relative to the baseline diversion ratios.

Table A22: Percentage Change in Market Share (No Inertia)

Facebook	Instagram	Reddit	Snapchat	TikTok	Twitter	YouTube
-29.7%	-33.3%	-29.0%	-15.9%	-39.0%	-18.0%	-31.0%

Notes: This table presents the percentage reduction in predicted average market share for the column application when $\beta^{q(i)} = 0$. The predicted average market share is computed over weeks $1, 4, 5$ of the experiment when all the participants faced no restrictions.

Online Appendix: For Online Publication Only

A Experiment Materials

A.1 Recruitment Materials

The following are the recruitment materials that were used for the study. Participants were either recruited from university lab pools or Facebook advertisements. For the participants who came from university lab pools they received the invitation in Section A.1.1 via email. The Facebook advertisement that was used for recruitment is shown in Figure OA1.

A.1.1 Recruitment Letter

Hello [NAME OF PARTICIPANT]!

We are inviting you to participate in an Internet experiment via Zoom. You will be able to earn money while contributing to science and hopefully having fun!

We are running an experiment to better understand how people spend their time online. We will ask you to install an application that will allow us to track how much time you spend on your phone and computer and periodically restrict access to certain applications on your phone [we only observe the time spent, not what happens on the app itself]. We will meet with you on zoom for five minutes to make sure the app is set up on your phone properly and then you will take a fifteen minute intro survey. You will not have to actively do anything during the rest of the experiment, beyond answering a short 4-minute survey once a week for five weeks.

Participants will earn $50 for successfully completing the experiment (i.e. keeping the application installed and completing all the survey questions each week). Note that only individuals with Android phones can participate in this experiment.

To sign up for the study, please click the link below to express your interest and we will follow up via email to schedule an initial meeting to set up the software and start the study: [link]

Thanks for your interest in participating in this study.

Figure OA1: Facebook Advertisement

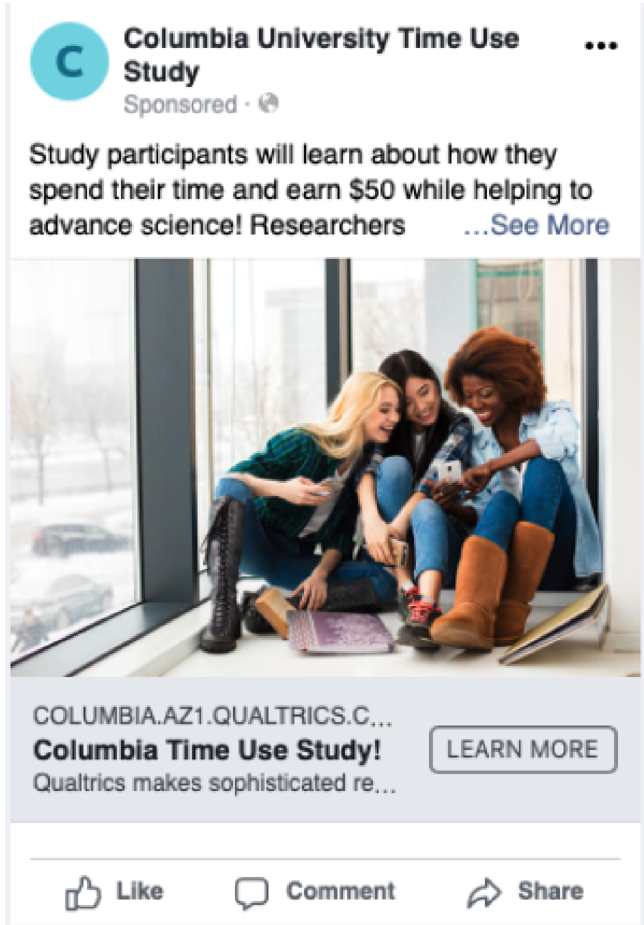

A.1.2 Recruitment Survey

Once the participants clicked on the link in the email sent from the lab pool or the Facebook advertisement, they were sent to an interest survey to complete. The recruitment survey had two pages. The first described the study in more detail, as shown in Figure OA2, and still emphasized that the main purpose of the study was to understand how participants spent their time online. The second page elicited information on social media habits and preferences with participants who stated that they used Facebook/Instagram/WhatsApp more than WeChat/Weibo/QQ/KakaoTalk were invited to the study.

Figure OA2: Recruitment Survey

<u>We are recruiting Android users for a five-week experiment!</u>

We are inviting you to participate in an Internet experiment via Zoom. You will be able to earn money while contributing to science and hopefully having fun!

We are running an experiment to better understand how people spend their time online. We will ask you to install an application that will allow us to **<u>track how much time you spend on your phone and computer</u>** [we only observe the time spent, not what happens on the app itself]. Additionally, **<u>there may be a period of 1-2 weeks in the middle of the study where we restrict your usage of a single social media application on your phone</u>**. This means that you will not be able to use that social media platform on your phone for that period of time, but **<u>will be able to do so on other devices</u>**. We will meet with you on zoom for five minutes to make sure the app is set up on your phone properly and then you will take a fifteen minute intro survey. You will not have to actively do anything during the rest of the experiment, beyond **answering a short 2-minute survey once a week for <u>five weeks</u>.**

Participants will earn $50 for successfully completing the experiment (i.e. keeping the application installed and completing all the survey questions each week). If you only complete a portion of the study you will receive $5 payment as compensation for your time and effort. **<u>Note that only individuals with Android phones can participate in this experiment.</u>**

If you are interested in participating, please fill out your contact information (phone number and email) and we will send a separate email about scheduling a time to get you enrolled into the experiment. This should happen sometime in early to mid March.

If you have more questions, you can email the researchers directly at msm2254@columbia.edu

What kind of phone do you have?

Android | iPhone | Other

1. Question # 1: Which set of social media platforms and apps do you use more often?
   - Facebook/Instagram/WhatsApp
   - WeChat/Weibo/QQ/KakaoTalk
2. Question # 2: Which of these apps do you use frequently (at least once a week)? Select all that are applicable.
   - Facebook, Instagram, Messenger, YouTube, WhatsApp, TikTok, Reddit, Snapchat, Twitter, WeChat, QQ, Weibo, KakaoTalk, Line, Telegram
3. Question # 3: Which web browser do you use most often?
   - Google Chrome, Safari, Internet Explorer, Firefox, Other

4. Question # 4: Contact Information - name, phone number, email

A.2 Baseline Survey

The baseline survey that participants fill out when they set up the software starts with the standard experimental consent form and study details. It then proceeds to ask a number of questions about their usage of social media applications.

Figure OA3: Consent Form and Study Details

Welcome to the study!

The study you are about to participate in is an economics and marketing study. The purpose of the study is to understand how people utilize applications on their phones and spend their time more generally. In order to do so, we will ask that you install software on your phone and computer. We will restrict **a single** social media or entertainment application on your phone for a time period ranging from one to two weeks during the course of the study.

Procedure
(Must read in order to know what is going on)

Overview
(1) You will set up the software on your phone and complete the initial long survey. (This is today)
(2) We will restrict a single application from your phone, for either one week or two weeks, starting on April 3rd. We will text you on April 2nd informing you which application will be restricted.
(3) The applications will remain installed and you will complete weekly surveys until May 2nd. You will receive two short surveys every week, one on Thursday and one on Saturday. Both will take 1-3 minutes to complete.
(4) Depending on your answer to a question later in this survey, you may have the opportunity to earn \$0-\$500 on top of the \$50. We will randomly select two participants to have an additional restriction and receive additional payment.

Details
The study will start with a Zoom meeting to set up the ScreenTime application, the desktop chrome extension, and a survey (which you should currently be in). The survey will ask you about how you use several popular social media and entertainment applications as well as some personality questions. The survey should take approximately ten to twenty minutes.

The majority of the study will make use of the installed ScreenTime application on your phone. This application will allow us to collect data on how much time you spend on applications on your phone. This application will not enable us to see what you do on the phone (i.e. the actual content within the applications), but only record how much time you spend on individual applications. This portion of the study will run until May 2nd (approximately 5 weeks).

If we do not text you about an application being blocked, then all the applications on your phone should be available. We will **only block entertainment and social media applications, not any essential components of your phone (i.e. maps, SMS, calling)**. At the end of the five weeks, you will be texted a password that will enable you to delete the application from your phone and receive your payment for completing the study.

It is important to note that all personal identifiers will be removed and researchers on the project will be the only ones who will have access to the data. If you complete **ALL** parts of the study, you will receive **\$50** in compensation. Based on your survey responses, you can earn additional compensation if you are selected at the end of the study to have an additional restriction. This will become clear when you complete the current survey. If you do not complete all parts, you will be compensated \$5 for completion of this initial survey. If you wish to opt-out of the study at any point, you can contact Guy Aridor at g.aridor@columbia.edu or Maayan Malter at mmalter22@gsb.columbia.edu but, if you do so, you will be forgoing the additional \$45 payment.

The questions were then as follows:

1. Question #1: Subjective Time Use. For each application write in your best guess for the number of hours you spend on it each week (in 30 minute increments, e.g. 1.50 hours for 1 hour and 30 minutes per week). The first column asks how much time you think you spend on the application on your phone and the second column asks how much time you think you spend on the application on your other devices.

- Facebook, Twitter, WhatsApp, TikTok, Instagram, Snapchat, Facebook Messenger, Attention Check. Write 99., YouTube, Reddit, Netflix

2. Question #2: Content Production. How frequently do you post content (including stories, resharing posts) on each of the following applications? For each of the following applications, the participants were asked to select one of the following options.

   - Applications: Facebook, Instagram, YouTube, Reddit, TikTok, Netflix, Snapchat, Twitter
   - Options: Never, Less than once a month, At least once a month, At least once a week, 2 or 3 times per week, Every day

3. Question #3: Subjective Activity on Application. The main activity I do on each application on my phone is as follows. For each of the following applications the participants were asked to select one of the following options.

   - Applications: Facebook, Instagram, YouTube, Reddit, TikTok, Netflix, Snapchat, Twitter, Messenger, WhatsApp
   - Options: Get Information (e.g. news about politics, sports, business, etc.), Online Shopping, Keep up with my friends' lives, Communicate with my friends, Entertainment content (e.g. memes, influencers, videos, etc.), I don't use this application

4. Question #4: Connections. For each application, write in the number of people you are connected to on the application. Please put your best guess for the range, there is no need to check for the exact values. For applications with followers / following, please let us know approximately how many individuals you follow on the application. For applications without direct connections, please let us know approximately how many individuals you interact with each week on the application.

   - Facebook (Friends): 0, 1-49, 50-149, 150-299, 300-499, 500-749, 750-999, 1000-2499, 2500-4999, 5000+
   - Twitter (Following): 0, 1-49, 50-149, 150-299, 300-499, 500-749, 750-999, 1000-2499, 2500-4999, 5000+
   - WhatsApp (Contacts): 0, 1-4, 5-9, 10-19, 20-29, 30-39, 40-49, 50-99, 100-249, 250+
   - TikTok (Following): 0, 1-9, 10-24, 25-49, 50-99, 100-199, 200-299, 300-399, 400-499, 500+
   - Instagram (Accounts Followed): 0, 1-49, 50-149, 150-299, 300-499, 500-749, 750-999, 1000-2499, 2500-4999, 5000+

- Snapchat (Friends): 0, 1-9, 10-24, 25-49, 50-99, 100-249, 250-499, 500-999, 1000-2499, 2500+
- YouTube (Channels Subscribed): 0, 1-9, 10-24, 25-49, 50-99, 100-249, 250-499, 500-999, 1000-2499, 2500+
- Reddit (Sub-reddits Subscribed): 0, 1-9, 10-24, 25-49, 50-99, 100-249, 250-499, 500-999, 1000-2499, 2500+

5. Question #5: WTA. See Figure OA4 for the interface and description presented to participants.

6. Question #6: Hypothetical Consumer Switching. For this question suppose the application in each row was no longer available on your phone. How do you think you would use the time you can no longer spend on that application? For each row application, let us know the category where you would spend most of your freed up time instead. For instance, if your Facebook is restricted and you think you would spend most of the gained time on other social media such as Twitter or TikTok then you would select “Social Media.” If you think you would spend your most of your time painting instead, then you would select “Other Hobbies.” If you don’t use the blocked app on a regular basis, then select “No Change.” The interface presented to participants can be seen in Figure OA5.

7. Remaining Questions: A battery of psychology questions and demographic questions. The only one reported in this paper is a social media addiction question, see Figure OA6, adapted from Andreassen et al. (2012).

Figure OA4: WTA Elicitation Interface

In this part, we ask you to state your monetary value for keeping access to each of your applications. Responding allows you to **earn additional money** on top of the $50 payment.

We present a series of offers from $0-$500 and ask you to select a **cutoff point** which indicates your true valuation for each application. All offers above this amount of money will be automatically filled in with "lose access" and all offers below this amount will be filled in with "keep access". Thus, the cutoff point you select indicates the minimum amount of money you'd be willing to get in exchange for having the application restricted.

For example, see the interface below and focus on the row $30 for the column Snapchat. If your chosen cutoff point was lower than $30 then you lose access to the application Snapchat and receive an additional $30 on top of the $50 experimental payment. If your cutoff point was equal to or higher than $30 then you retain access to Snapchat and receive no additional money.

We utilize the following procedure to determine whether you are selected to receive payment and which offer we consider. We will randomly select two participants. For these participants, we will **randomly select one of the applications (columns) and one of the offers (rows)**. If, for the selected row, you had chosen **keep access** then nothing will happen and you will receive no additional payment. If, for the selected row, you had chosen **lose access** then you will have the application restricted for a week and receive the additional payment.

Because we select any of the given rows randomly, the higher the cutoff point you state the less likely it is that you receive money. Conversely, the lower the cutoff point you set the more likely you are to receive it. **The procedure is constructed so that it in balance it is best for you to report your true valuation for keeping access.**

It is important to note that this is **in addition to the restrictions in the study** and will take place on May 2nd to May 9th extending the total duration of the study by one week. You will receive a text message if you are one of the selected participants.

	Facebook		Twitter		WhatsApp		Snapchat		Reddit	
	Keep Access + $0	Lose Access + Offer	Keep Access + $0	Lose Access + Offer	Keep Access + $0	Lose Access + Offer	Keep Access + $0	Lose Access + Offer	Keep Access + $0	Lose Access + Offer
$0										
$5										
$10										
$15										
$20										
$25										
$30										
$35										
$40										
$45										
$50										
$60										
$70										
$80										
$90										
$100										
$125										
$150										
$175										
$200										
$250										
$300										
$350										
$400										
$450										
$500										

Figure OA5: Hypothetical Consumer Switching Interface

	Social Media	Messaging Applications (Messenger, WhatsApp, etc.)	Entertainment Applications (e.g. Netflix, YouTube, Twitch, etc.)	News Sources (e.g. WSJ, NYT, WashPo, etc.)	Other Hobbies	In-person socializing	No Change
If **Facebook** were blocked on your phone, which activity (to the right) would you do instead?							
If **Instagram** were blocked on your phone, which activity (to the right) would you do instead?							
If **Messenger** were blocked on your phone, which activity (to the right) would you do instead?							
If **YouTube** were blocked on your phone, which activity (to the right) would you do instead?							

Figure OA6: Social Media Addiction Scale

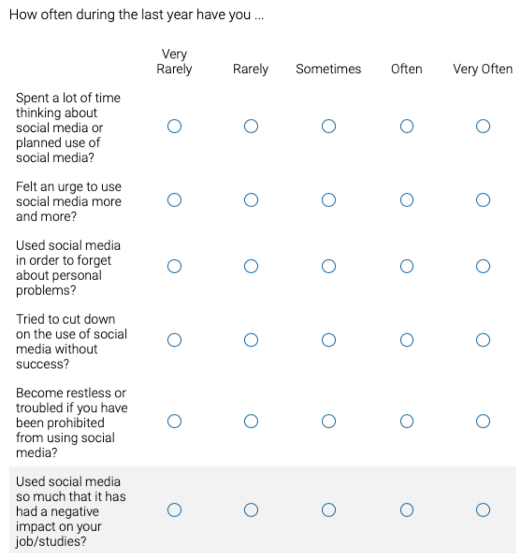

How often during the last year have you ...

	Very Rarely	Rarely	Sometimes	Often	Very Often
Spent a lot of time thinking about social media or planned use of social media?	○	○	○	○	○
Felt an urge to use social media more and more?	○	○	○	○	○
Used social media in order to forget about personal problems?	○	○	○	○	○
Tried to cut down on the use of social media without success?	○	○	○	○	○
Become restless or troubled if you have been prohibited from using social media?	○	○	○	○	○
Used social media so much that it has had a negative impact on your job/studies?	○	○	○	○	○

A.3 Additional Surveys

There are two weekly surveys throughout the study. The first is during the week and sent on Thursdays as part of the data collection partnership for this study. It is meant to capture instantaneous psychology measures, which is why it is sent during the week while the application restrictions are ongoing. The second is sent on Saturday mornings and is meant to record subjective perceptions of time usage throughout the week.

The Thursday survey asks the participants how fast they felt the week had passed, questions about their social connectedness and well-being, a question about whether they made any big purchases in the past week, and finally whether there were any major life events in the past week.

The Saturday survey is broken into three separate components. The first component asks participants how much time they felt they spent off their phones on Facebook, Instagram, YouTube, Facebook Messenger, WhatsApp, Netflix, TikTok, Twitter, and Reddit. The second component asks participants how much time they spent on life necessities, including sleeping, studying, attending classes, paid work, cooking/eating, cleaning, socializing in person, and child care. The final component asks participants how much time they spent on leisure activities off the phone, including playing video games, reading books, watching cable TV, streaming on TV / tablet, exercising, shopping (in person), artistic hobbies, and reading print media.

Finally there is an endline survey that is attached to the final weekly time use survey, which asks the following questions:

1. Question #1: Ability to revise WTA. The participants are given the same WTA question as the initial survey, but the results are pre-filled based on their initial survey responses. They are instructed to revise the values if they wish.

2. Question #2: Reason for revision. The participants are asked why they revised the WTA value.

   - Applications: Facebook, Instagram, YouTube, Reddit, TikTok, Netflix, Snapchat, Twitter, Messenger, WhatsApp
   - Options: Have a better idea of how much time I spend on the application, Reduced my usage of the application during the study period, Started using the application during the study period, Increased my usage of the application during the study period, Realized the application is more/less important to me than I thought, I realized I misunderstood this question when I first answered it, No Change

3. Question #3: What did you think the purpose of the study and the restrictions was? Open-Response.

4. Question #4: During the restriction period, select the following statement which you think most accurately describes your behavior. Multiple choice.

   - I downloaded new applications and spent most of the gained time using them.
   - I spent more time on applications I already had installed and spent time curating better content on these applications (e.g. following more accounts/channels on YouTub/TikTok/Instagram, figuring out how different features worked).
   - I spent more time on applications I already had installed, but did not significantly invest time in improving my experience on them.
   - I spent more time on my computer.
   - I spent more time off my devices.
   - I had no restrictions.
   - No change.

5. Question #5: After the restriction period, I started to use the restricted application on my phone. Multiple choice with the following possible responses: More time than before the restrictions, the same time as before the restrictions, Less time than before the restrictions, I had no application restriction.

6. Question #6: Select the following statement which you think most accurately how your behavior after the restrictions compares to before the restrictions. Multiple choice.
   - I spent my time more or less the same.
   - I spent more time on applications I downloaded during the restriction period.
   - I spent more time on applications I already had installed but did not significantly invest time in improving my experience on them during the restriction period.
   - I spent more time on applications I already had installed, but had invested time in making my experience on them better.
   - I spent more time on my computer.
   - I spent more time off my devices.
   - I had no application restrictions
7. Question #7: (Optional) If you want to describe in words how you responded to the restrictions, feel free to elaborate below.
8. Question #8: (Optional) How do you think you will change your behavior with respect to social media applications going forward?

A.4 Software

Figure OA7: Chrome Extension Interface

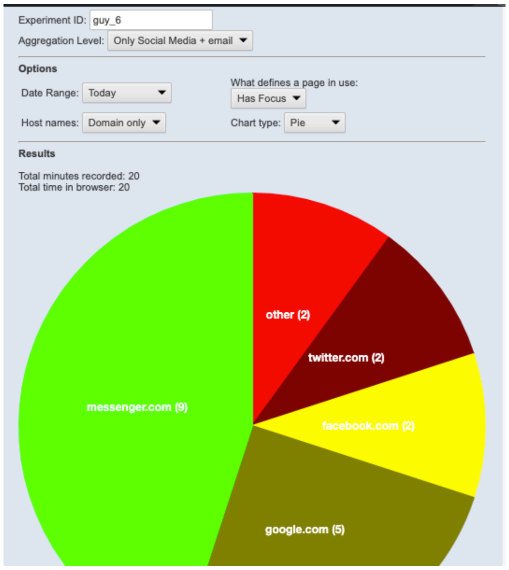

Figure OA8: Parental Control Interface

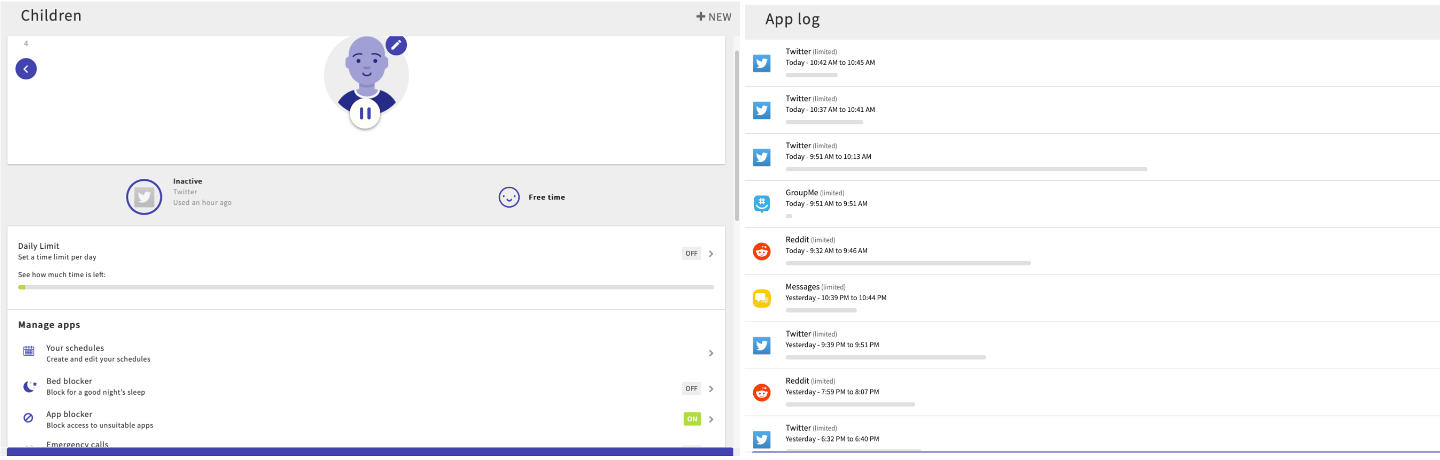

A.5 Pilot Experiment

This section contains information on the details of the pilot study. The phone data collection software is the same as the main experiment, but there was no Chrome Extension for this version of the study. The primary differences between the two experiments are that the pilot experiment included several restrictions for each participant and the sample size was substantially smaller. The study consisted of 123 participants recruited from the Columbia Business School Behavioral Research Lab. Participants were similarly paid $50 for completing the study.[1]

The timeline for the study was as follows. Participants had a virtual meeting to set up the software from 9/29 - 10/10. The vast majority of participants were set up before 10/3, but a handful were set up between 10/3-10/10. There are two experimental blocks. The first block runs from 10/3 until 11/7. The period between 10/3 and 10/10 serves as the baseline usage for this block. Participants were randomized into group A and B on 10/10. Group A had a restriction on Facebook and Messenger together from 10/10-10/17, followed by a week of no restrictions, a week of YouTube restriction, and finally a week of no restrictions. Group B had no restrictions for 10/10-10/17, followed by week of Instagram restriction, a week of no restrictions, and finally a week of Snapchat and TikTok restricted together. In the second experimental block that runs from 11/7 - 12/4, participants were randomly assigned each week to either have a restriction or be in the control group. The period from 11/7-11/14 serves as a second week of baseline usage and the order of the restrictions across the weeks is as follows: Facebook/Messenger, YouTube, Instagram.

[1]In order to ensure that there was little cross-contamination of participants from the pilot study in the larger study, different lab pools were utilized for the pilot vs. main study. However, to my knowledge, there were only 3 participants who overlapped between the two different experiments.

B Alternative Estimation of Diversion Ratios

In this section I provide an alternative method of estimating the diversion ratios. I follow the methods proposed in Conlon and Mortimer (2021); Conlon, Mortimer and Sarkis (2022) that directly exploit the experimental product unavailability variation to estimate the diversion ratios. The method proceeds by first using the estimated average treatment effects between the restricted applications and the other applications of interest to estimate the diversion ratios from the restricted applications to other applications. This provides a nonparametric estimate of the diversion ratio between these applications. Then, I impose a semiparametric logit assumption and, using the aggregate market shares and the estimated diversion ratios, an MPEC procedure enables the estimation of the rest of the matrix of diversion ratios.

B.1 Estimation Procedure

I restrict to the same set of prominent social media and entertainment applications as in the main text: Snapchat, Facebook, Reddit, TikTok, Instagram, YouTube, and Twitter. The outside option is defined as time not on the phone and other applications on the phone. Thus, I have a choice set of $J = 7$ applications plus an outside option and the goal is to estimate the $J \times (J + 1)$ matrix of diversion ratios. I aggregate time spent during these time periods as follows. I consider that each time of day has T time periods of half minute units. I aggregate time spent during these time periods in order to compute market shares for each individual. I drop the “late night” hours so that I only consider 17 hours in the day.

There are I individuals, $J + 1$ applications (including outside option), and T time periods. I denote the choice decision of each individual i for application j at time period t as a discrete choice:

$$d_{ij,t} = \begin{cases} 1, & \text{if } u_{ij,t} > u_{ij',t} \quad \forall j' \in \mathcal{J} \setminus j \\ 0, & \text{otherwise} \end{cases}$$

Thus, the individual choice shares for individual i as well as the aggregate choice shares for application j are given as follows:

$$s_{ij}(\mathcal{J}) = \frac{1}{T}\sum_{t=1}^{T} d_{ij,t} \quad s_j(\mathcal{J}) = \frac{1}{IT}\sum_{i=1}^{I}\sum_{t=1}^{T} d_{ij,t}$$

B.2 Estimating Diversion Ratios of the Restricted Applications

I estimate the diversion ratios for the restricted applications. I denote $\mathcal{S}$ as the vector of aggregate market shares. Following Conlon and Mortimer (2021), I can directly compute the diversion ratios from the restricted application to other applications of interest using the estimated treatment effect of the application restrictions:

$$\tilde{D}_{kj} = \frac{\mathcal{S}_j(\mathcal{J} \setminus k) - \mathcal{S}_j(\mathcal{J})}{\mathcal{S}_k(\mathcal{J})}$$

In order to compute the numerator, I estimate the baseline specification (1) for each application of interest and, for the denominator, I use the average share of application k in the baseline period. However, this formulation does not guarantee that the resulting diversion ratios sum to 1 or are non-negative. I impose the assumption that the resulting diversion ratios must be non-negative (i.e. the applications are substitutes) and that they sum to 1. Thus, given the resulting estimates of the diversion ratio, I first impose that they are non-negative and then normalize them so that the resulting estimated diversion ratios all sum to 1.

For additional precision in the estimates of the diversion ratios, I make use of the empirical Bayesian shrinkage estimator used by Conlon, Mortimer and Sarkis (2022). The estimator is given as follows where q_j denotes the share of daily time on application j:

$$\hat{D}_{kj} = \lambda \cdot \mu_{kj} + (1 - \lambda) \cdot \tilde{D}_{kj}, \quad \lambda = \frac{m_{kj}}{m_{kj} + q_j}$$

The idea is that one can view the diversion ratio as a binomial with $\mathcal{S}_k(\mathcal{J})$ "trials" and $\mathcal{S}_j(\mathcal{J} \setminus k) - \mathcal{S}_j(\mathcal{J})$ successes in terms of consumers who chose application k but now switch to application j. Viewed in this manner, I specify a prior belief on D_{kj} and parameterize this prior as Dirichlet$(\mu_{j0}, \mu_{j1}, .., \mu_{jK}, m_{kj})$.[2] The reason to make use of this estimator as opposed to the estimate itself is that some of the estimates may be large, but noisy, especially for applications that have smaller number of users such as TikTok and we want the estimator to account for this. Note that this procedure makes no parametric assumptions about the functional form of consumer utility beyond the substitutes assumption.

B.2.1 Estimating the Other Entries of the Diversion Ratio Matrix

The challenge now is to estimate the remaining cells of the matrix of diversion ratios. However, of course, I do not have direct experimental variation for all of the applications of interest. Following Conlon, Mortimer and Sarkis (2022), I assume that consumer utility follows a semi-parametric

[2] Since $0 \leq D_{kj} \leq 1$ and $\sum_k D_{kj} = 1$.

logit, $u_{ij} = V_{ij} + \epsilon_{ij}$ where ϵ_{ij} is the standard type-1 extreme value error. Given this assumption, then Conlon and Mortimer (2021) show that the average second-choice diversion ratio is given by:

$$D_{kj} \equiv \mathbb{E}[D_{kj,i} \mid i \text{ chooses } k] = \sum_{i=1}^{N} \frac{\pi_i \cdot s_{ik}}{s_k} \cdot \frac{s_{ij}}{1 - s_{ik}} \tag{3}$$

Under this parameterization, Conlon, Mortimer and Sarkis (2022) propose the following MPEC matrix completion procedure in order to estimate the rest of the diversion ratios by using the aggregate shares and the estimated diversion ratios from the experimental data. One intuition as to why this procedure works is that the logit assumption induces full support so that everything weakly substitutes with everything else (i.e. the "connected substitutes" notion discussed in Berry and Haile (2016)) so that it's possible to get information on the substitution between Facebook and Snapchat even if I observe no experiments with these items removed. I simplify their procedure since in my case the time spent on the outside option is pinned down due to the fact that there are a limited number of minutes in the day. The notation is as follows: $\hat{D}_{kj}$ denotes the estimated diversion ratios from second choice data, $\mathcal{S}_j$ denotes the aggregate shares, and π_i denotes the probability that a consumer is of type i, OBS denotes the pairs of applications for which I have second-choice measures of diversion.

$$\min_{s_{ij},\pi_i} \sum_{(k,j)\in OBS} (\hat{D}_{kj} - D_{kj})^2 + \lambda \sum_j (\mathcal{S}_j - s_j)^2 \tag{4}$$

$$\text{subject to: } s_j = \sum_i \pi_i \cdot s_{ij} \tag{5}$$

$$D_{kj} = \sum_i \pi_i \cdot \frac{s_{ij}}{1 - s_{ik}} \cdot \frac{s_{ik}}{s_k} \tag{6}$$

$$0 \leq s_{ij}, \pi_i, s_j, D_{kj} \leq 1, \sum_i \pi_i = 1, \sum_j s_{ij} = 1 \tag{7}$$

This procedure involves an exogenous selection of I latent types of individuals each with different preferences as well as the penalization parameter $\lambda > 0$. The idea is that, as in standard random coefficients logit demand models, the resulting aggregate market shares come from a mixture of different types of consumers whose preferences each follow a different logit. Thus, (3) pins down the average second-choice diversion ratio and the MPEC procedure optimizes over the space of possible mixtures of different possible types of individuals in order to best fit the observed diversion ratios and aggregate market shares.

I implement this procedure and choose the exogenous parameter λ by the model with the best in-sample fit according to the mean-squared error or mean absolute error.[3] I consider the set of

[3] I alternatively considered a cross-validation procedure where the model is estimated holding out one set of di-

$I \in \{1, 2, ..., 8, 9\}$ and for each I choose $\lambda \in \{0.2, 0.6, 1.0, ..., 9.2, 9.6, 10\}$. Given the fixed λ for each I, I then choose across I by comparing whether the resulting estimate correctly fits the market shares and whether the resulting estimated diversion ratios could reasonably be implied by the noisier experiments from the pilot experiment which included two applications (Facebook-Messenger and Snapchat-Tiktok) and a smaller sample size. The nonparametric diversion ratios from the joint restrictions in the pilot experiment are reported in Table OA2.

B.3 Diversion Ratio Estimates

I report the nonparametric diversion ratio estimates for Instagram and YouTube that I compute directly using the experimental variation. I pool together the data from the pilot and larger-scale experiment in order to get more precise estimates. For the estimates I use an informative prior so that the prior follows the predictions of logit and the diversion is proportional to market shares, $\mu_{kj} = \frac{s_j}{1-s_k}$ and $m_{kj} = 10$. I compute standard errors using simple block bootstrap with the blocks being participants and utilizing the bootstrap percentile 95% confidence interval with 20000 replications. Table OA2 reports the estimated diversion ratios as I vary the value of m_{kj}. Recall that increasing m_{kj} places additional weight on the prior, which is the predicted diversion from logit, at the expense of the experimental estimates.[4] Furthermore, I also report the estimated diversion ratios from the joint removal of Snapchat and TikTok as well as Facebook and Messenger. I do not directly incorporate these into $\mathcal{D}$ since they contain multiple application restrictions and are less precisely estimated due to smaller sample sizes and multiple restrictions, but rather use them to choose between resulting estimates.

Table OA1 reports the estimated diversion ratios for the rest of the applications using the MPEC procedure for $I = 3$ and $\lambda = 6.6$. For lower values of I, the selected λ do a poor job at fitting the market shares, whereas for the higher values of I the selected λ predict very little diversion to the outside option for the other applications. The resulting estimates to the outside option for the reported specification in Table OA1 are in line with what one would expect given the nonparametric diversion estimates in Table OA2 for the joint Snapchat and TikTok as well as the joint Facebook and Messenger restrictions.

version ratios (i.e. holding out one of the two experiments) but found that this led to unreasonable estimates, likely since I have a small number of experiments and in this case the procedure is relying only on the estimates from one experiment.

[4]This also varies across applications since, for instance, Snapchat and TikTok have lower aggregate usage the estimator naturally places more weight on the prior for diversion to these applications relative to diversion to diversion for more used applications like YouTube and Instagram.

Table OA1: Diversion Ratio Estimates

	Instagram	YouTube	Facebook	TikTok	Snapchat	Reddit	Twitter	Outside Option
Instagram	-	$2.8e{-5}$	0.05	0.046	$5.4e{-5}$	$2.8e{-5}$	0.014	0.89
YouTube	0.052	-	0.033	0.019	0.0035	0.0039	$6.5e{-5}$	0.89
Facebook	0.024	0.077	-	0.012	0.0062	0.0092	0.0072	0.86
TikTok	0.022	0.061	0.017	-	0.0065	0.0098	0.0079	0.88
Snapchat	0.017	0.019	0.014	0.011	-	0.012	0.0099	0.92
Reddit	0.016	0.014	0.014	0.011	0.0076	-	0.01	0.93
Twitter	0.015	0.00033	0.013	0.011	0.0079	0.012	-	0.94

Notes: The presented table is of the matrix of diversion ratios, D_{kj}, where a cell in the table is the diversion from application k (row) to application j (column). The diversion ratios are estimated using the MPEC procedure.

Table OA2: Nonparametric Diversion Ratio Estimates

	Instagram	YouTube	Facebook	TikTok	Snapchat	Reddit	Twitter	Outside Option	m_{kj}
Instagram	–	0.0	0.07	0.08	0.0	0.0	0.027	0.82	0
		(0.0, 0.22)	(0.0, 0.18)	(0.0, 0.17)	(0.0, 0.03)	(0.0, 0.04)	(0.0, 0.11)	(0.51, 0.95)	
YouTube	0.06	–	0.05	0.03	0.002	0.0	0.0	0.86	0
	(0.0, 0.13)		(0.0, 0.10)	(0.0, 0.09)	(0.0, 0.03)	(0.0, 0.04)	(0.0, 0.02)	(0.72, 0.94)	
Instagram	–	0.0	0.05	0.05	0.0	0.0	0.01	0.89	10
		(0.0, 0.22)	(0.0, 0.13)	(0.003, 0.10)	(0.0, 0.01)	(0.0, 0.02)	(0.0, 0.05)	(0.61, 0.97)	
YouTube	0.05	–	0.03	0.02	0.003	0.004	0.0	0.89	10
	(0.005, 0.10)		(0.002, 0.07)	(0.0, 0.05)	(0.0, 0.01)	(0.0, 0.02)	(0.0, 0.01)	(0.80, 0.95)	
Snapchat and TikTok	0.03	0.04	0.03	–	–	0.003	0.002	0.90	0
	(0.01, 0.07)	(0.004, 0.09)	(0.006, 0.07)			(0.0, 0.01)	(0.0, 0.01)	(0.78, 0.95)	
Facebook and Messenger	0.08	0.0	–	0.0	0.01	0.01	0.0	0.90	0
	(0.0, 0.31)	(0.0, 0.31)		(0.0, 0.08)	(0.0, 0.10)	(0.0, 0.10)	(0.0, 0.04)	(0.42, 1.0)	

Notes: The presented table is of the matrix of diversion ratios, D_{kj}, where a cell in the table is the diversion from application k (row) to application j (column). This displays different estimates of diversion from Instagram to other applications and YouTube to other applications, depending on the value m_{kj}. I additionally compute the diversion during the Snapchat-TikTok and Facebook-Messenger restrictions which were run in the pilot study. 95% confidence intervals are constructed by simple block bootstrap and using the percentile confidence interval calculation with 20000 replications and are reported in parentheses.

C Collection of Survey Responses

In this section are the responses to the optional question in the endline survey which asked the participants to describe in words how they responded to the restrictions.

- Addiction
  - I hated it while it happened, but it really broke the app's addictive nature.
  - I never realized that I am tsuch addicting to instagram until I found myself opened it absentmindedly several times during my restrictions period. my usage time of ig has decreased from averagely 6.5 hrs before the restrictions to 3 hr in the first week, but bounce back to 7 hrs this week, even exceeding the number before.
  - It's strange, because it didn't feel like I needed YouTube, I just knew I had spent a lot of time on it. However, when it became restricted, I noticed how much time I had spent simply laying about and watching YouTube. It felt weird knowing that my instinct was to immediately press the YouTube button when I got bored, and I realized I perhaps need/use it more than I think.
  - It was crazy how addicted I am to these apps. During the restrictions, I kept accidentally trying to open the app -all the time. I didn't realize how much time I spent on them.
  - I kept opening instagram time after time forgetting that is was blocked
  - I had one restriction on Instagram and it was weird breaking the habit of accessing and took some getting used to avoiding the app
  - When the restriction started I got a feeling I was gonna be a little anxious. I was wrong.
  - It was frustrating - did not know I was so addicted to YouTube
  - I felt out of the loop so I often tried to access Instagram using my laptop.
  - At first restricting instagram was frustrating as i had the application on my home screen and built muscle memory for the past 4 years to press that part of the screen where the instagram shortcut is. I removed instagram from my home screen and after 5 days of the restriction i completely realized instagram was nor important at all for me and only time i open it is when i receive a direct message.
- Shifted Towards Other Apps
  - It wasn't easy at first as I tried to access the restricted application about two different times but I received the restriction message from screen time app with a grin on my face....lol. I had to figure out what I want from other applications I didn't know offered

similar content before time, after the restriction elapsed, I had adjusted to sourcing for such content on both apps.

- Well at first after my YouTube was restricted, I thought I could access it using my browser but then i realised that was also impossible. I was like, how will I cope without streaming videos on YouTube? But after some time I adjusted and got used to it.
- At the beginning i felt like damn this is an important application (Youtube) and what if i need it for anything Turns out i dont need it as much and there are other options available
- Pre-COVID, I would listen to a lot of podcasts when driving, walking to class, etc. So when Youtube was restricted, I mostly just listened to more podcasts like I used to. I think I also probably watched more Youtube on my PC and smart TV during this period.
- At the beginning it felt like something was missing but eventually I started using other apps and filled that vacancy
- I spent time on twitch watching streamers vs. Youtube where I had watched them before.
- I think the restriction gave me the opportunity to spend more time on other applications i had already installed but hardly use.
- I often use youtube for music on my phone when I don't want to pay for Spotify premium, but during the restriction period I ended up resubscribing to Spotify Premium for $5 so I could listen to music on my phone easily

• Realized Value of Application

- It was a bit hard to adapt at first but I eventually got used to it.Eventually I realized I am better off without it so I ended up deleting it and till now am okay with my decision.
- After the restriction I definitely started spending more time on the app that was restricted. I started to use the app more because I wanted to track local businesses which can be hard to discover by googling. I'm not sure if it was a coincidence that I developed an interest in small businesses and increased my app usage or if it was the restriction that caused me to appreciate what I could do on the app more.
- I felt that I missed using it I realized I was spending to much time on the app
- Struggling to access Instagram, but when there's no restrictions, i found that the content i wanna access previously is very trivial

  - I felt minorly inconvenienced since I could still access on my computer if it were an emergency like an insta dm I needed to respond to. Having time away from insta definitely helped me mentally.
  - Sometimes I misses to use but nothing as bad as I thought. Most of time I have not importante to do, Its just a way to spend time
  - I felt after restrictions that I need this application more and I can't take this restrictions for a long.
  - YouTube was restricted, so it was a little difficult when my baby was having a meltdown in public, but it also wasn't as often as usual, thankfully. It was difficult also if I needed to learn something off of YouTube pertaining to my career like a how-to or new technique.

- Shifted Towards Non-Digital Activities
  - Honestly I spent more time outdoor and with friends.
  - I initially felt bored, since a common reflex I had was to open up Youtube whenever I had nothing to do. However, within a few days, I started doing other things instead, such as reading. It was actually a good experience.
  - At first it was difficult because YouTube is the most used app by me.Whatever it is YouTube is a go to for me in my daily life.After that I made up my mind to concentrate in different things and spent more time off the devices.I tried to concentrate more on my studies and spent time with my family.
  - I was surprised my youtube was restricted. For me its a big part of the content i consume and it is was hard to not have it on my phone. Initially I tried watching it on my computer but it was something i couldn't keep up all the time. Over time my useage dropped from watching a lot to, mainly watching when i am on my computer taking a small break (even then only watching the videos i really like and not wasting time on YT)

- Impact on Socializing
  - I realized I spent a lot of time on an app establishing really ineffective communication. I changed the way in which I communicate online.
  - I didn't think I used Instagram very much but the restriction turned out to be very annoying as friends would message me there and wonder why I wasn't responding

– I used Instagram to communicate with friends less frequently when it was restricted, but used WhatsApp more instead. These were reverted after restrictions were lifted

– I felt frustrated because I feel like I was missing out. I wasn't able to keep up with the people I followed on Instagram as much because the app was restricted

– I felt it was a very interesting experience. I don't feel like I have an addiction to certain applications and could probably live my life without it. The only limit I faced was that I could not contact certain people, who I only talk with on that application. But to be honest, I could live even without those conversations or certain people and would probably find other apps to contact them on. But I did not do that.

– Instagram was restricted for me and because I mainly use it as a communication app, I was not significantly affected. I just used regular text, video call, and Snapchat to keep up socially.

– It was a little annoying especially whenever my friend shared something that can only open on that platform. But after a couple of days I was able to make my peace with it

– I did a bit of communication on Instagram, so told the person I was chatting to to switch and that didn't really happen so it ended up reducing how much we messaged