Regulating Digital Platform Monopolies: The Case of Facebook

Seth G. Benzell*, Avinash Collis†‡

September 2022

Abstract

Although they often offer services to consumers for low prices, or even freely, many digital platforms are thought to have significant market power due to their network effects and supply side economies of scale. However, the existence and magnitude of this harm to social welfare, and the distributional impact of possible antitrust remedies to ameliorate it, are disputed. We construct and analyze a general model of digital platforms, determining conditions under which government interventions raise welfare. We calibrate our model for the case of Meta's Facebook using a survey of over 57,000 US Internet users. Facebook creates $14 billion in surplus per month, concentrated among female and older users. We simulate six proposed policy interventions. We find a 3% tax on Facebook's ad-revenue raises welfare by 1.1%, by shifting Meta's incentives towards maintaining a larger platform. Achieving perfect competition, while preserving network effects, would raise surplus from Facebook by 4.8%. On the other hand, a horizontal or vertical breakup of Facebook that failed to maintain interoperability or promote competition could decrease social surplus by up to 84.7%. Finally, a "data-dividend" rebate of profits to users would increase surplus by 30.3%.

*Argyros School of Business and Economics, Chapman University, Stanford HAI Digital Economy Lab & MIT Initiative on the Digital Economy, benzell@chapman.edu

†McCombs School of Business, The University of Texas at Austin, avinash.collis@mccombs.utexas.edu

‡We thank Sinan Aral, Daniel Ershov, Lindsey Raymond, Fiona Scott Morton, Erik Brynjolfsson, Dean Eckles, Geoffrey Parker, Xiang Hui, Zhou Zho, Daniel Rock and Marshall Van Alstyne for their invaluable comments. We thank Andrey Fradkin and Chris Forman for fantastic discussions. We additionally thank the MIT Initiative on the Digital Economy for research funding. Our study was determined to be exempt by MIT's IRB. After this research was substantively completed, A.C. received funding for other research from Meta.

1 Introduction

Although many successful digital platforms are inexpensive or free for consumers, their strong network effects and supply side economies of scale can give them significant market power. This has led to several FTC antitrust lawsuits and European Commission proposals seeking remedies to resulting harms. However, the magnitude of the social welfare loss due to the market power of digital platforms is still poorly understood. The same holds for the distributional consequences of proposed digital platform taxes, regulations, and antitrust remedies. One reason for this is the unique role that network effects play in digital platforms. The existence of network effects significantly modifies the social welfare consequences of changes to the size of a platform's user base. To better understand these issues, we construct and analyze a general model of digital platforms, solving for the monopolist's optimal strategy and determining conditions under which government interventions raise welfare. We then calibrate the model for the case of Meta's Facebook social media platform, and simulate the consequences of several proposed policies.

Our paper begins by introducing a model of platform participation that allows for several dimensions of heterogeneity. Users vary in their opportunity cost for using the platform, the disutility they receive from advertising, and the value they get from interacting with other types of users on the platform over different lags of time. In the context of a social media platform, this last dimension of heterogeneity can represent how the value of reading a newsfeed is a function of the different types and ages of posts. It is a model of a multi-sided platform in the sense that user types, or market segments, benefit (or are harmed) by others' participation on the network heterogeneously.[1] The platform monopolist cares both about its profits as well as the size of its user base. To achieve these objectives the monopolist sets a level of monetization, which corresponds to a fee or advertising level for each side of the platform.

We employ a comparative static approach to analyze the model, where we consider the steady state impact of nudges to an initial equilibrium. In terms of platform strategy, we find that profit-maximizing firms should have increased monetization on market segments that demand the platform inelastically and receive stronger positive network effects than they create (and vice-versa). At the social level, the nature of platforms' incentives to maintain a large user base is critical to optimal taxation and regulation. When network effects are positive, and distortionary taxation is nonexistent, small de-

[1] When conceived in this way, any platform, including a one-sided platform, can be thought of as an multi-sided platform once we account for the heterogeneity in users within a side. For example, a telephone network, which is the classic example of a one-sided network, can be thought of as consisting of multiple sides that can be distinguished based on various characteristics including business versus personal use, demographics, regional location, heterogeneity in activity (frequent users or not) and type of activity (always callers, callers and receivers, or always receivers).

creases in monetization from the monopoly level are social welfare improving so long as the firm's incentive to have a large user base does not exist or is socially positive (e.g., it represents an intangible asset the firm is accumulating from data), but not necessarily so otherwise. If the monopolist's incentive to grow their user base is sufficiently socially detrimental (e.g., it represents dumping to deter entry or does not account for an unpriced externality), then the monopoly or competitive price may be too low from a social welfare perspective. The competitive price may also still be too high, from a social welfare perspective, when network participation creates positive externalities.

After introducing and analyzing our model, we proceed to an empirical illustration. We collected information from 57,000 Internet users in the U.S. on their demand for Facebook using choice experiments conducted through Google Surveys. We categorize the participants into twelve demographic groups based on their age and gender. To collect information on demand for and network effects on the social network, we use an experimental choice approach in the spirit of Brynjolfsson et al. (2019) and Allcott et al. (2019). These authors measure the consumer surplus generated by digital goods by conducting discrete choice experiments where they offer consumers the choice to give up access to the good in exchange for monetary compensation. We build on these papers by asking about a new type of free good (the value of digital network connections) as well as by using information from the full distribution of responses to estimate elasticities of demand rather than solely focusing on median and average valuations. We adapt this approach to our case by giving participants the choice to give up access to a subset of their network in exchange for monetary compensation.

Using this information about the demand for Facebook, we estimate the parameters of a demand curve for each of the twelve demographic groups, as well as the twelve by twelve matrix of their network externalities. We complement this with additional survey questions about friend frequency, the disutility of advertising, and publicly available data on Facebook's advertising revenue by demographic group.

With this model of individual participation, we can then calculate the effects of counterfactual pricing policies, government policies, and demand shocks. Because our model explicitly simulates the impact of policy changes on each demographic group, we estimate both the total and distributional welfare impact of changes. We begin by simulating Facebook's revenue maximizing strategy under the assumption that they do not derive any non-monetary value from accumulating a large user base. We find that Facebook could raise revenue by \$2.38 billion a month (from a baseline of \$1.79 billion) if it did not care about the size of its user base. This strategy entails squeezing value from inelastic users, reducing Facebook usage by 49.1%, and lowering total consumer surplus by 42.1%. Overall, women are more harmed by this than men. Women aged 45-54 are the group harmed most by this in percentage terms, with a 56% decrease in consumer surplus. 35 percentage points of their reduction in surplus is due to the direct effect of advertisements, and the remainder is because of the value from connections they derive from more elastically participating men, who leave the platform. We infer that in addition to maximizing current revenues, Facebook values maintaining a large user base. We impute the shadow value Facebook places on maintaining a large user base as the one that justifies their current level of advertising as optimal. In subsequent simulations, we take into account this shadow value when simulating Facebook's response to policy changes.

We then evaluate three taxation and redistributive policies. The first is a tax on ad revenue. We show theoretically that a tax on ad revenues would not change Facebook's optimal advertising level, so long as Facebook has no other considerations. However, if Facebook values a large user base, then a tax on advertising redirects it from raising high levels of advertising revenue to cultivating a large user base. A tax on the number of users has the opposite effect, leading Facebook to squeeze a smaller group of users with a higher level of fees. Quantitatively, we find that a 3% tax on ad revenues would raise consumer surplus by 1.3%, and raise 2.4% of current Facebook ad revenue in taxes. A tax on the number of users of Facebook, which raised the same amount of revenue, would lower consumer surplus by -0.1%. The third proposed policy for redistributing the wealth from Facebook is the "Data as Labor" framework, where Internet users would be compensated for their "labor" in viewing targeted advertisements (Posner and Weyl, 2018). We conceive of this policy as a rebate of Facebook's current advertising revenues to users. We estimate one possible "Data as Labor" regime — i.e. rebating all current ad revenues to users — as raising social welfare by 30.3%. This policy is better than the first best, because it allows Facebook to continue showing 'productive' advertisements[2] while still providing a net subsidy to use the platform. "Data as Labor" therefore represents the best of both worlds regarding welfare maximization, if somehow obvious obstacles, especially the creation of fake accounts to steal the subsidy, could be overcome.

We also simulate three proposed regulatory interventions. The first is a successful regulatory intervention that achieves perfect competition, while preserving cross-firm network effects, in the market for social media. In practice, this might entail lowering barriers to entry and enforcing interoperability (i.e., allowing users on a Facebook competitor to view posts by and communicate with users of Facebook and other Facebook competitors). Finally, we simulate two regulatory interventions under the assumption that they fail to increase competition. The first is a horizontal breakup of Facebook which leaves the U.S. with two monopolies over half of the population each. The

[2]Productive, that is, in the sense that they create less than an a dollar of direct disutility per dollar of revenue.

second envisages a vertical breakup which erodes Facebook quality with no offsetting increase in competitiveness. We predict that perfect competition would raise social welfare by 4.8%. Breaking Facebook into two non-competitive successors would lower social welfare by 84.7%. A vertical breakup which made Facebook unattractive to 5% of the population, without increasing competition, would lower social welfare by about 10.1%.

2 Related Literature

Following the seminal work of Rochet and Tirole (2003) and Parker and Van Alstyne (2005), platform researchers have extensively studied the impact of direct and indirect network effects on various strategic issues including pricing (Hagiu, 2009), launch (Evans and Schmalensee, 2010) and openness (Boudreau, 2010). The core insight of this research is that it can be optimal for a two-sided platform to subsidize one side and increase fees one the other (Eisenmann et al., 2006).

The above papers all focus on what are known as one or two-sided platforms. Examples of two-sided platforms are Uber (riders and drivers) and Ebay (sellers and buyers). In a two-sided platform, it can make sense to price discriminate based on side, because sides differ in both their elasticity of demand and in the network effects they provide. For example, an additional Uber driver in a region provides a positive externality to riders (they will get a ride faster) but a negative externality to other drivers (they will have to wait longer between fares). However, a large literature suggests that even within a "side" of a one- or two-sided platform, users are heterogenous in the effect their actions have on the network. The empirical literature on network effects uses several techniques for their estimation, including studying exogenous shocks to the network (e.g., (Tucker, 2008)), using an instrumental variable approach (e.g., (Aral and Nicolaides, 2017)) and conducting field experiments (e.g., (Aral and Walker, 2012)).

Several more recent papers model pricing in the presence of multidimensional network effects. The paper with a model most similar to ours is Weyl (2010). That paper, like ours, introduces and analyzes a model of an indivisible platform good with network effects. That model, also like the one in this paper, allows for groups to vary in both their network effect on other groups and in their opportunity cost for using the platform. Unlike our model, it does not explicitly account for preferences over lagged participation. It finds that a wedge exists between the profit maximizing and social welfare maximizing pricing strategy.[3] While our setting is similar, we go beyond Weyl

[3] The exact nature of this wedge – as a marginal, not an average distortion – was clarified in a published comment (Tan and Wright, 2018).

(2010) to derive theorems on the welfare implications of the effect of nudges from an initial equilibrium. This is in contrast to Weyl's solution which entails using an "insulating tariff" to allow the monopolist to coordinate agents on the preferred global equilibrium, an approach under-which it is harder to derive similar results.[4] While any equilibrium in our model is also an equilibrium in Weyl, our approach, which incorporates transition dynamics, also allows us to give necessary (and in special cases, sufficient) conditions for equilibrium stability.

Another closely related paper is Bernstein and Winter (2012). In that paper, Bernstein and Winter determine a monopolist or city-planner's optimal strategy for attracting complementary businesses or agents. Their focal example is renting out storefronts in a shopping mall. Unlike our model, in Bernstein and Winter (2012) the relevant characteristics of all agents are directly observable. Because of this, if the monopolist were allowed to coordinate agents on the "all participate" equilibrium (such as via an insulating tariff like Weyl (2010)) it could first-degree price discriminate and claim all the surplus from the most efficient outcome.[5] If forced to settle for a Nash Equilibrium, Bernstein and Winter (2012) shows the monopolist's solution is to employ a "divide and conquer" strategy. The business with the strongest positive network effects is given the most attractive offer, and less important agents given decreasingly attractive offers. The least positively-influential agent to participate is given an offer of exactly their surplus, conditional on all the more influential agents participating. Our setting differs from Bernstein and Winter (2012) in two main ways. First, like our difference from Weyl (2010) we focus on local comparative statics rather than on determining global maxima. Second, our analysis treats sides of the platform as non-observably heterogeneous in their opportunity costs. Because our monopolist cannot make individually tailored offers (only those tailored to a side, such as a demographic group), participation is probabilistic at the individual level and the monopolist loses a tool for extracting surplus. In the setting of a social media platform, which has millions of users, our approach may be more appropriate.

Two other related papers are Candogan et al. (2012) and Fainmesser and Galeotti (2015). These consider monopolistic pricing of a divisible network good, where utility from the good is quadratic in the amount consumed and linear in the impact of neighbors' consumption. In Candogan et al. (2012), the platform firm has perfect knowledge about all individuals' utility functions, but allows for individuals to vary in their utility from the platform good, although this utility must be quadratic. They show that the problem of determining profit maximizing prices is NP-hard, but derive

[4] An insulating tariff in this context is a commitment by the the monopolist platform to pay out a large fine to participants if the platform fails to achieve the equilibrium it promised. This payment is never made in equilibrium.

[5] See Segal (2003) for a discussion of this for the case of homogeneous network effects.

an algorithm guaranteeing 88% of the maximum. Fainmesser and Galeotti (2015) considers a similar model but assume that all individuals have the same demand for the network good, while allowing for a random distribution of network connections. They find that allowing for the network to lower prices on "influencers" must increase social welfare, but allowing firms to fully price discriminate might be harmful.

Our paper builds on these prior papers along several additional dimensions. First, our model features more realistic monetization, allowing for different types of users to face different levels of disutility from the firm increasing their level of advertising. This is in contrast to Candogan et al. (2012) and Fainmesser and Galeotti (2015), who do not allow for such variation, and Weyl (2010), who features an unrealistic pricing scheme, where users are charged based on the level of participation of other users (i.e., the insulating tariff). Weyl's use of insulating tariffs in pricing forces users to immediately jump to a desired equilibrium in response to a price change, which prevents a dynamic analysis of a pricing change. Second, unlike the models presented in Candogan et al. (2012) and Fainmesser and Galeotti (2015), our model has a realistic amount of uncertainty within a side of a model, meaning that first-degree price discrimination that drives consumer surplus to zero is impossible.[6] We are also the first to explicitly allow for heterogeneity in the lagged network effects over time. The most important contribution of our model is that it is the first to allow for straightforward calibration. To the best of our knowledge, no previous paper has made quantitative model-based recommendations about multi-sided platform pricing, or quantitatively evaluated the welfare consequences of a platform regulation market structure change.

The illustration in our paper is of Meta's Facebook app, a digital platform primarily monetized through advertising.[7] Most platforms keep the quantity of ads ("ad load" to those in the industry) shown per user fixed while showing different ads to different users based on their characteristics and bid outcomes of ad auctions, e.g. Google (Hohnhold et al., 2015); Pandora (Huang et al., 2018). Platforms with a newsfeed, such as Facebook, WeChat and Linkedin, understand the trade-off between ad load and user engagement. Some of them show the same number of ads per person (see Huang et al. (2020) for advertising on WeChat), while others fix the number of ads a user sees based on the expected revenue generated by the user in the long term (Yan et al. (2019) describe Linkedin's ad load strategy). While this optimization takes user engagement into account, network externalities generated by a user are not explicitly

[6]The fact that platforms cannot fully first-degree price discriminate is testified to by papers which show that users benefit considerably on average from joining a platform. For example, Ceccagnoli et al. (2012) find that independent software publishers experience an increase in sales and a greater likelihood of issuing an IPO after joining a major platform ecosystem, and Brynjolfsson et al. (2019) find large consumer surplus from the use of digital platforms.

[7]In 2019, Meta, then Facebook, made 5.8 billion dollars in advertising revenues monthly from its "family of apps," 99.2% of total revenue from that source (Meta, 2022).

modeled and users generating different levels of network externalities end up seeing the same number of ads.[8] In estimating structurally the impact of market structure on social welfare in the presence of network effects, our paper is in the tradition of Rysman (2004). Rysman's model is of an analog two-sided platform: the yellow pages. He uses instruments to find the spillover effects of additional advertisements on phone-book quality. Rysman finds that small *decreases* in competition might increase welfare, as there would be fewer better phone books with more utilitous advertisements.

Our paper also addresses the growing literature on the optimal regulation of digital platform monopolies. An important paper summary and discussion of this literature can be found in Scott Morton et al. (2019). The authors ultimately calls for a special 'Digital Authority' that would have access to platforms' internal information for the purpose of making regulatory decisions and boosting platform competition. We hope that this structural model and its descendants could be tools for such an agency.

3 Model

In this section, we lay out our model of multi-sided platform participation. In the model, everyone's platform participation is a function of all other individuals' decisions to participate on the platform at different times in the past. For a social network, this could represent how the value of reading a newsfeed is a function of the different types and ages of posts.

We begin by characterizing the response of participation to an infinitesimal change in platform quality or monetization. When the model's initial equilibrium is stable, this leads to a finite change in long-term equilibrium platform participation. We provide intuition for this and discuss conditions for equilibrium stability.

Next, we micro-found the model of participation as the discrete choice of heterogeneous agents who choose between platform participation and their outside option. The value of platform participation is a function of the history of participation of all other users, and the outside option is a random variable generating an elasticity of demand.

We then introduce a monopolist platform owner to the model and solve for its objective maximizing strategy. The monopolist platform cares about post-tax net-revenues and may also be motivated to maintain a large user base. To achieve these goals the monopolist chooses how intensely to monetize each type of user on the platform. When network effects are positive and elasticities of demand are negative, higher monetization means more revenue per user, but may reduces platform participation.

[8]Based on informal conversation with researchers who have worked with Facebook, our understanding is that in constructing its newsfeed, Facebook gives every potential entry a score, based on the amount of engagement the entry is expected to create in the user who sees the ad, the amount of revenue that might be generated (if it is an advertisement) and a penalty for being similar to a recently displayed entry.

The monopolist's optimal strategy has a deep connection to Katz-Boniach centrality, which was found also to figure in the monopolist's optimal strategy in Ballester et al. (2006). We provide an intuitive discussion of the monopolist's first order condition.

Finally, with both the monopolist strategy and a model of consumer welfare in hand, we provide comparative static analyses of the monopoly equilibrium. We find that, under a set of general conditions which likely apply to Facebook, a decrease in monetization increases social welfare. Under the same conditions, we find that taxes on ad-revenues can increase platform participation and social welfare, while a tax on the number of users has the opposite effect. On the other hand, raising or lowering prices from the competitive equilibrium has an ambiguous effect on social welfare. This comparative static analysis then leads into the applied portion of the paper, where we estimate the magnitude of these effects for the case of Facebook, as well as the direction of the effect of policy interventions with ambiguous effects under the comparative static analysis.

3.1 Model Summary and Key Terms

General Model

1. $\vec{P}_t$
2. $\vec{\Phi}_t$
3. $\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}}$
4. $\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}$
5. $\vec{\lambda}$
6. $\frac{\partial \vec{P}_{ss}}{\partial \vec{\Phi}} =$

$(I - \sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1} \frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}$

1. Participation on the platform at time t of each group i. It is an I dimensional vector.
2. Net revenue at time t from each user of group i. An I dimensional vector. A choice variable for the platform. Also referred to as monetization.
3. Matrices of local network effects at Y time delays. It is a Jacobian IxI matrix giving the change in participation for group i due to a change in participation of group j, y periods in the past.
4. The Jacobian IxI matrix giving change in participation for group i due to a contemporaneous change in monetization vector $\vec{\Phi}$.
5. An I dimensional vector indicating the non-pecuniary (i.e. in addition to $\vec{\Phi}$) value of marginal users of the platform, to the platform. The shadow value of users.
6. The IxI matrix giving the steady state ($\lim_{T\to\infty}$) response of participation of each group to a permanent change in each type of monetization ($\vec{\Phi}$)

Microfoundation

1. $\vec{\mu}_t$
2. $U_{t,i,\mathrm{i}}(P_{t,i,\mathrm{i}} = 1) =$
   $\mu_{t,i} = \mu_i(\vec{P}_{t-1}, ..., \vec{P}_{t-Y}, \vec{\Phi}_t)$
3. $U_{t,i,\mathrm{i}}(P_{t,i,\mathrm{i}} = 0) = \epsilon_{t,i,\mathrm{i}}$
   $\epsilon_{t,i,\mathrm{i}} \sim \epsilon_i$
4. $\frac{\partial \vec{\mu}_t}{\partial \vec{\phi}_t}$
5. $\frac{\partial \vec{P}}{\partial \vec{\mu}}$
6. $\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t} = \frac{\partial \vec{\mu}_t}{\partial \vec{\Phi}_t} \frac{\partial \vec{P}}{\partial \vec{\mu}}$
7. $\boldsymbol{B} := \sum_{y=1}^{Y} \frac{\partial \vec{\mu}_t}{\partial \vec{P}_{t-y}}$
8. $\frac{\partial \vec{P}_{ss}}{\partial \vec{\Phi}} = (I - \boldsymbol{B} \frac{\partial \vec{P}_t}{\partial \vec{\mu}_t})^{-1} \frac{\partial \vec{\mu}_t}{\partial \vec{\phi}_t} \frac{\partial \vec{P}_t}{\partial \vec{\mu}_t}$

1. The value, in dollars, of platform participation for each group i. An I dimensional vector
2. The value of platform participation for an individual i in group i is a function of participation of all groups up to Y periods in the past, and current monetization level $\vec{\Phi}_t$
3. The opportunity cost of participation for an individual i in group i is a random, with distribution determined by group.
4. The matrix of marginal disutilities from additional monetization on each group (e.g. each group's disutility from marginal advertising). It is typically diagonal.
5. The diagonal matrix of elasticities of demand (i.e. for each group i, additional users due to a marginal increase in platform quality, measured in dollars). An IxI matrix.
6. An IxI matrix of the immediate change in participation of each group i due to a change in the monetization vector $\vec{\Phi}_t$. It is the product of the marginal disutility from monetization and elasticity of demand.
7. $\boldsymbol{B}$ is the matrix of steady state, local, network effects in dollars. It is the sum of Y IxI matrices which give the marginal value of each group's participation on the platform $\vec{P}_t$, y periods in the past, to each other group's platform value $\vec{\mu}_t$
8. An IxI matrix giving the steady state response of participation of each group to a change in monetizaiton vector Φ.

3.2 Platform Participation

Let $\vec{P}_t$ be a vector of platform participation by users of different types, with t denoting the time period. $\vec{P}_t$ is a vector with I elements, where I is the number of sub-populations, or sides, of the multi-sided platform.

Let $\vec{\Phi}$ be an I dimensional variable, representing net revenue per user (which we also refer to as "monetization level"). In the typical case, where $\vec{\Phi}$ is a choice variable of the platform monopolist, $\vec{\Phi}$ represents revenue from fees, subscriptions, advertising, personal data sales etc. net of marginal per-user costs. The vector $\vec{\Phi}$ can take negative values, corresponding to the platform deciding to subsidize a particular side of the platform.

Participation on the platform can be very flexibly modeled as a function of all other agents' participation over Y previous periods,[9] as well as the level of monetization. This can be written as

$$\vec{P}_{t+1} = f(\vec{P}_t, \vec{P}_{t-1}, ..., \vec{P}_{t-Y}, \vec{\Phi}_{t+1}). \tag{1}$$

We assume that the network starts in equilibrium (i.e. $\vec{P}_0 = \vec{P}_{-y} \ \forall \ Y > y > 0$).

Assuming that the function in equation (1) is differentiable in all dimensions near the initial equilibrium, it can be linearly approximated as:

$$\vec{P}_{t+1} = \sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}} P_{t-y+1} + \frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t} \Phi_{t+1} + \vec{P}_0. \tag{2}$$

where the matrices $\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}}$ are I dimensional square matrices indicating the elasticity of every type of user's participation to the participation of all others y periods in the past.[10] $\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}$ is an I dimensional square matrix representing the local elasticity of platform participation by users of different types to monetization of different types. In many practical settings, for example, where monetization represents fees on a particular group, $\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}$ will be a diagonal matrix of negative elements. Finally, $\vec{P}_0$ is a constant intercept vector.

Using equation (2) we can evaluate the elasticity of steady state participation on a platform to different parameters. Consider a permanent, infinitesimal change in

[9] Allowing for a user's participation to be a function of many lags of others' participation is a flexible way of allowing for dynamic changes in the value of network effects over time. One particularly important reason to allow for this, in the context of a social network, is that different types of posts may have usefulness that decays at different rates (e.g. a thoughtful "effort post" may create value on a newsfeed for a longer period than a faddish meme-post). Additionally, as an alternative microfoundation to the one we posit later, if user participation is a function of the expected simultaneous user participation of others, then this is a flexible way of allowing for expectations to be a function of historical patterns in participation.

[10] Therefore the matrices $\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}}$ are Jacobian matrices: Matrices of first derivatives of participation with respect to participation in previous periods

$\vec{\Phi}$, the vector of monetization. This is a particularly important partial derivative to understand, because $\vec{\Phi}$ is the platform monopolist's choice parameter. Taking the derivative of equation (2) with respect to a one-time, permanent change in monetization $\vec{\Phi}$ in period 1 yields the following recursive formulae for the change in participation in any period in the future due to this change:

$$\begin{aligned}\frac{\partial \vec{P}_1}{\partial \vec{\Phi}} &= \frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t} \\ \frac{\partial \vec{P}_2}{\partial \vec{\Phi}} &= \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-1}} \frac{\partial \vec{P}_1}{\partial \vec{\Phi}} + \frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t} \\ &\vdots \\ \frac{\partial \vec{P}_{t+1}}{\partial \vec{\Phi}} &= \sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}} \frac{\partial \vec{P}_{t-y+1}}{\partial \vec{\Phi}} + \frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}.\end{aligned} \tag{3}$$

And, if the initial equilibrium is stable, a permanent increase in $\vec{\Phi}$ will have the following impact on participation in the long run steady state:

$$\frac{\partial \vec{P}_{ss}}{\partial \vec{\Phi}} = lim_{T \to \infty} \frac{\partial \vec{P}_T}{\partial \vec{\Phi}} = \big(\boldsymbol{I} - \sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}}\big)^{-1} \frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}. \tag{4}$$

where $\boldsymbol{I}$ is the identity matrix.[11] Note that this equation has deep connections to Katz-Bonacich centrality, which was shown by Ballester et al. (2006) to arise in a similar network goods setting.[12]

[11]It also yields the following effect in T periods:

$$\frac{\partial \vec{P}_T}{\partial \vec{\Phi}} = \Big(\sum_{t=1}^{T} \big(\sum_{y=1}^{t-1} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}}\big)^{T-t}\Big) \frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}, \tag{5}$$

although we are primarily interested in the steady state impacts of policies.

[12]The vector of Katz-Bonacich centralities, for an adjacency matrix $\boldsymbol{G}$ is, discounted by $\boldsymbol{a}$ is:

$$\vec{C} = (\boldsymbol{I} - \boldsymbol{aG})^{-1}\vec{1} \tag{6}$$

where $\vec{1}$ is the vector of 1's. This measure of centrality assigns a number to each node according to how much value it derives, in the game's equilibrium, from all other nodes, both directly and indirectly. Our related equation, (4), gives the matrix of long-term changes in participation in response to a change in any element of the monetization vector $\vec{\Phi}$. Later we discuss the connection of this measure to our model's monopoly solution.

3.2.1 Equilibrium Stability

But why (or when) should the partial derivative of long-term platform participation with respect to a shock be finite? To get an intuition, consider the one-dimensional, one-lag special case (i.e., $I = 1$, $Y = 1$) of equation (5). Then participation in T, after an increase in monetization in period 1, can be written as,

$$\frac{\partial P_T}{\partial \Phi} = \left(\sum_{t=1}^{T} \left(\frac{\partial P_t}{\partial P_{t-1}} \right)^{t-1} \right) \frac{\partial P_t}{\partial \Phi_t}. \tag{7}$$

This is just the geometric series, and it converges so long as $|\frac{\partial P_t}{\partial P_{t-1}}| < 1$. When $|\frac{\partial P_t}{\partial P_{t-1}}| < 1$, the limit of the series is:

$$lim_{T \to \infty} \frac{\partial P_T}{\partial \Phi} = (1 - \frac{\partial P_t}{\partial P_{t-1}})^{-1} \frac{\partial P_t}{\partial \Phi_t}. \tag{8}$$

A small change in monetization will have a finite impact on participation so long as $|\frac{\partial P_t}{\partial P_{t-1}}|$ is less than one. Thus, when the value of a social media platform is not very dependent on others' participation (i.e., $\frac{\partial P_t}{\partial P_{t-1}} = 0$), the total effect of a change in advertising will be approximately equal to the direct effect of advertising alone. When the value of social media is more tightly tied to others' participation ($|\frac{\partial P_t}{\partial P_{t-1}}|$ close to 1 or larger), a small change in advertising can have massive effects.[13] This effect can also be visualized graphically, where $\frac{\partial P_t}{\partial P_{t-1}}$ is displayed as the slope relating P_{t+1} to P_t, as in Figures 1 and 2.

In the I dimensional, one-lag case (Y=1), participation on the platform will be stable as long as all eigenvalues of the matrix $\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-1}}$ are less than one in absolute value.[14] To understand why, recall that the eigenvalue associated with an eigenvector of a given linear transformation determines how much a vector on the span of that eigenvector will be stretched. If all the eigenvalues of a matrix are less than one, then every time the linear transformation is applied to a starting vector, that vector will shrink closer to the origin. As the equations above show, calculating the long-term effect of a change in fees Φ requires taking the sum of $\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-1}}$ projected on itself repeatedly, with each additional term corresponding to another cascade of the network effect. Having all eigenvalues less than one guarantees that these terms get increasingly small in magnitude, and that the infinite series has a finite sum. In the I dimensional

[13] An analogy can be made to epidemiology. "R_0" is the term for the rate at which individuals infect others with a disease. When R_0 is less than 1, exogenously infecting one additional individual causes a finite amount of others to get infected. Eventually the cascades of the infected infecting others peter out. However, when R_0 is greater than one, a single infection can in principle go on to infect the entire world. Hence the emphasis on NPIs to try to reduce R_0; ideally to a level below 1.

[14] This is a well established result for discrete time linear dynamic systems. See, for example, Knill (2019), theorem 22.3 or Åström and Wittenmark (1997) theorem 3.1.

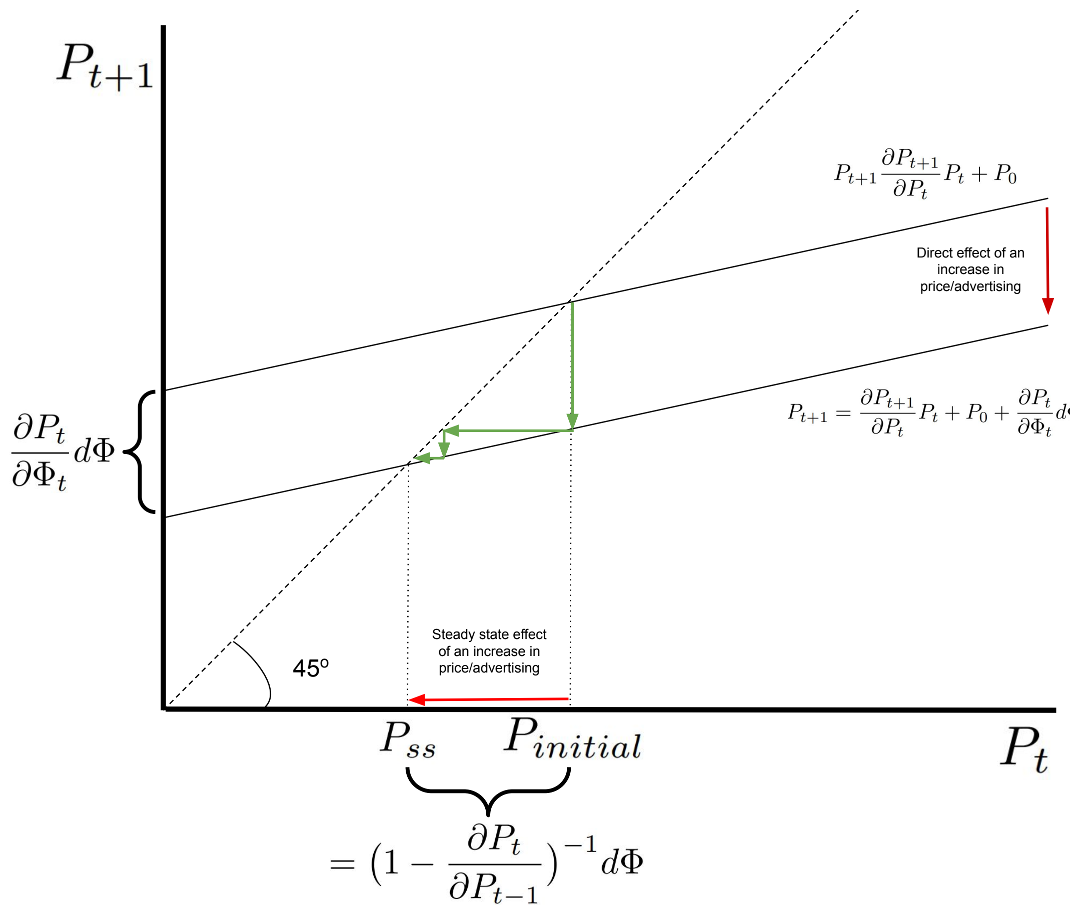

Figure 1: An example of a platform in a stable equilibrium because $0 < \frac{\partial P_t}{\partial P_{t-1}} < 1$. An increase in monetization Φ shifts down the participation curve relating participation in time t to participation in $t+1$. As the slope of the curve is less than 1, the system is stable, and an increase in advertising leads to a finite decrease in steady-state participation. The steady state decrease in participation is larger than the direct effect, because the initial decrease in participation induces others to stop using the platform. However, the size of these effects dampen rapidly over the course of several cascades.

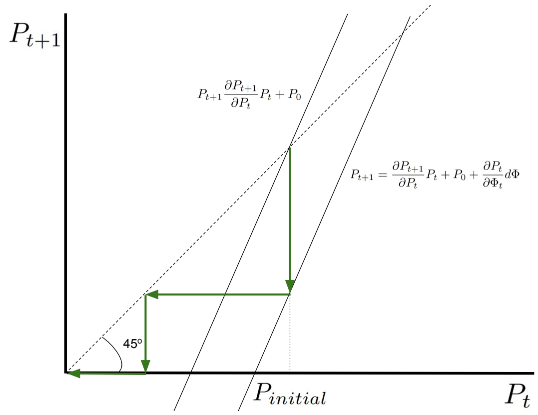

Figure 2: An example of an unstable participation system because $\frac{\partial P_t}{\partial P_{t-1}} > 1$. An increase in monetization Φ shifts down the participation curve relating participation in time t to participation in $t+1$. As the slope of the curve is greater than 1, the system is unstable, and an increase in advertising leads participation on the platform to completely collapse. The steady-state decrease in participation is larger than the direct effect, because the initial decrease in participation induces others to stop using the platform. The size of these cascades increases over the course of iterations, leading all participants to stop using the platform in the steady state if $\frac{\partial P_t}{\partial P_{t-1}}$ is constant.

case of many lags (i.e. $Y > 1$), a necessary condition for stability is that the eigenvalues of the matrix $\sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}}$ all be less than one.[15]

3.3 Microfounding Platform Participation in Discrete Choice

In this section, we ground the above model of platform participation in the discrete choice of heterogenous individuals between platform participation and an outside option. This microfoundation is necessary for developing a model of consumer welfare, which in turn is necessary for evaluating how different proposed policies will impact welfare. We decompose the elasticity of platform participation into a vector of elasticities of demand for the platform (generated by a distribution of opportunity costs), a matrix of disutility from different sorts of monetization, and matrixes of network effects indicating how platform quality (measured in dollars) depends on the participation of others of different types at different lags. We measure these terms directly with experiments that we use to calibrate our model and analyze the welfare effect of policies.

Formally, a potential platform participant i on side i (a potential user, for short) chooses whether to participate in the platform at time t ($P_{t,i,\text{i}} = 1$) or not ($P_{t,i,\text{i}} = 0$).

[15] For a further discussion of necessary and sufficient conditions in discrete time delay difference equations, an ongoing area of research, see Sun (2006) and Stanković et al. (2013).

If the potential user chooses to use the platform, the user receives:

$$U_{t,i,\text{i}}(P_{t,i,\text{i}} = 1) = \mu_{t,i} = \mu_i(\vec{P}_{t-1}, ..., \vec{P}_{t-Y}, \vec{\Phi}_t). \tag{9}$$

A potential participant i on side i also faces a distribution of opportunity costs ϵ_i, the resolution of which is unknown by the platform monopolist, who could otherwise first-degree price discriminate.[16]

The distribution of opportunity costs generates a demand function for each side $P_{t,i}(\mu_{t,i})$ (i.e. a relationship between the quality of the platform and the number of individuals who participate). Locally, this relationship is summarized by the elasticity of platform participation to platform quality. Elasticities of demand can be written as the following diagonal matrix:

$$\frac{\partial \vec{P}}{\partial \vec{\mu}} = \begin{bmatrix} \frac{\partial P_1}{\partial \mu_1} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \frac{\partial P_I}{\partial \mu_I} \end{bmatrix}, \tag{10}$$

where all terms are positive, indicating increasing usage as platform quality increases.

With this model of choice in hand, we now decompose the general model's elasticities into the product of two terms: The impact of changes on platform quality and the elasticity of demand. In other words, the direct elasticity of participation to a change in monetization is:

$$\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t} = \frac{\partial \vec{P}}{\partial \vec{\mu}} \frac{\partial \vec{\mu}_t}{\partial \vec{\Phi}_t}, \tag{11}$$

where $\frac{\partial \vec{\mu}_t}{\partial \vec{\Phi}_t}$ is the matrix of marginal disutilities due to monetization. In our comparative static analysis, we will assume that this latter term is a diagonal matrix of only negative terms.[17] For an ad-supported platform, this means that ads degrade platform quality at the margin for those who view it, but has no direct impact on those who do not.[18]

[16]The ex-post user demand function is therefore

$$\begin{cases} P_{t,i,\text{i}} = 1 & \text{if } \mu_{t,i} \geq \epsilon_{t,i}(\text{i}) \\ P_{t,i,\text{i}} = 0 & \text{otherwise.} \end{cases}$$

[17]This is still a very flexible assumption. Generally, a firm's net-revenue per-user is determined by multiple inputs (for example, a level of investment in quality and a level of advertising). From any given starting point, the firm will face a production frontier of different levels of net-revenue per-user and perceived platform quality. Collapsing this to a one dimensional decision nests an underlying decision by the platform to locate at the most efficient part of this tradeoff. In other words, $\frac{\partial \vec{\mu}_t}{\partial \vec{\Phi}_t}$ is the marginal change in perceived platform quality when net-revenue per user changes, when the firm does so with the most efficient balance of instruments. In the case of Facebook, we assume that the level of advertising is the marginal tool Meta uses to balance platform quality and net revenue per user.

[18]That being said, a non-diagonal matrix for $\frac{\partial \vec{\mu}_t}{\partial \vec{\Phi}_t}$ still has an economic interpretation. If $\vec{\Phi}$ represents

Similarly, the general model's matrices describing the elasticity of participation to previous periods' participation can be decomposed into the product of the elasticity of demand matrix, and a matrix of network effects, describing the value a user of a type at a certain lag provides to platform quality, in dollars. So,

$$\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}} = \frac{\partial \vec{P}}{\partial \vec{\mu}} \frac{\partial \vec{\mu}_t}{\partial \vec{P}_{t-y}}. \tag{12}$$

The matrix giving the marginal contribution to platform quality of a user of a given type at a given lag can be written as:

$$\frac{\partial \vec{\mu}_t}{\partial \vec{P}_{t-y}} = \begin{bmatrix} \frac{\partial \mu_{1,t}}{\partial P_{1,t-y}} & \cdots & \frac{\partial \mu_{1,t}}{\partial P_{I,t-y}} \\ \vdots & \frac{\partial \mu_{i,t}}{\partial P_{j,t-y}} & \vdots \\ \frac{\partial \mu_{I,t}}{\partial P_{1,t-y}} & \cdots & \frac{\partial \mu_{I,t}}{\partial P_{I,t-y}} \end{bmatrix}. \tag{13}$$

We can now re-write equation (4) as:[19]

$$\vec{P}_{ss} = lim_{T\to\infty} \frac{\partial \vec{P}_T}{\partial \vec{\Phi}} = \big(\boldsymbol{I} - \frac{\partial \vec{P}}{\partial \vec{\mu}} \boldsymbol{B}\big)^{-1} \frac{\partial \vec{P}}{\partial \vec{\mu}} \frac{\partial \vec{\mu}_t}{\partial \vec{\Phi}_t}, \tag{14}$$

where $\boldsymbol{B}$ is the sum of matrices of network effects in dollars over all lags

$$\boldsymbol{B} = \sum_{y=1}^{Y} \frac{\partial \vec{\mu}_t}{\partial \vec{P}_{t-y}} \tag{15}$$

3.4 The Platform Monopolist

Having microfounded the model, the last piece we need before a welfare analysis is a model of how to determine the level of monetization. Determining the competitive price is straightforward, the level of net-revenue for each type of user is set to zero ($\vec{\Phi} = \vec{0}$). Here we characterize how a monopolist, taking every user's demand for platform services as exogenous, would set monetization levels to maximize their utility.[20]

The platform monopolist wishes to maximize their steady state utility Π. In general, this can be a complex function of the number of users of every type, as well as the

a vector of different types of quality and monetization investments that a platform might make, then the matrix $\frac{\partial \vec{\mu}_t}{\partial \vec{\Phi}_t}$ would correspond to how each side of the platform feels about these multi-dimensional quality investments. In this way our model can also nest the one presented in Veiga et al. (2017).

[19] We could further decompose $\vec{P}$ into a participation rate for each side of the platform and a total eligible population for each side of the platform. Similarly, we could decompose $\boldsymbol{B}$ into a value per connection and a rate at which a pair of users by side form connections (on Facebook these connections are called 'friends' and are entered into by mutual agreement). We do this in the calibration of the model, but it does not add any clarity or usefulness in the theoretical analysis.

[20] Additional derivation details are provided in Appendix A.

platform's total net revenues. Such a utility function can be locally approximated as the sum of post-tax net revenues from each user, and the marginal non-pecuniary value of each user type. This non-pecuniary value from users can arrise for many reasons, both socially positive (it represents some intangible asset, e.g., training data for AI, that the firm is accruing) or socially ambiguous (it is a reduced form way of representing the firm's incentive to under-price to deter entry). Setting aside the nature of the non-pecuniary motivation for a moment, and letting $\vec{\lambda}$ denote this shadow value of a marginal user, we have, in the steady state:

$$\Pi = (1-\tau_1)\Phi'\vec{P}_{ss} + (\vec{\lambda} - \vec{\tau}_2)'\vec{P}_{ss} + F, \tag{16}$$

where τ_1 is a scalar corresponding to a tax on marginal profits (i.e., a corporate income tax) and $\vec{\tau}_2$ is a vector of per-user taxes. $\vec{\tau}_2$ might have different elements because governments might want to only apply a special excise tax on child users, for example. F is the fixed cost of operation and $'$ indicates the transpose operator.

Taking the linear approximation of the general model of platform participation, equation (2) and solving for the steady state yields:

$$P_{ss} = (I - \sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1}(\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}.\vec{\Phi} + P_0) \tag{17}$$

Plugging this into equation (16), taking the derivative of the above with respect to $\vec{\Phi}$, and rearranging, yields a vector of first-order conditions for profit maximization:

$$\begin{aligned} \frac{\partial \vec{\Pi}}{\partial \vec{\Phi}} = \vec{0} = (1-\tau_1)\Big(\frac{\partial \vec{P}_{ss}}{\partial \vec{\Phi}}\vec{\Phi} + (I - \sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1}\vec{P}_0 + \frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}'(I - \sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1\prime}\vec{\Phi}\Big) \\ + \frac{\partial \vec{P}_{ss}}{\partial \vec{\Phi}}'(\vec{\lambda} - \vec{\tau}_2). \end{aligned} \tag{18}$$

This is a vector of equations, because at the optimum none of the I monetization vectors being tweaked should increase profits. Let $\bar{\boldsymbol{B}} = \sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}}$. We can now characterize the monopolist's optimal strategy as follows.

Theorem 3.1. *The monopolist's local steady-state profit maximizing strategy* $\vec{\boldsymbol{\Phi}}$, *for the special case where* $\vec{\lambda} - \vec{\tau}_2 = \vec{0}$ *(i.e., firms have no marginal incentive beyond net revenues to acquire users) and* $\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t} = -I$ *(i.e., at the margin, a dollar of additional monetization creates a dollar of disutility for platform users) and using the linearization in equation (17) is:*

$$\vec{\mathbf{\Phi}} = \left(\frac{1}{2}I - \frac{1}{4}(\bar{\boldsymbol{B}}' - \bar{\boldsymbol{B}})\left(I - \frac{\bar{\boldsymbol{B}}' + \bar{\boldsymbol{B}}}{2}\right)^{-1}\right)\vec{P}_0. \tag{19}$$

Equation (19) has several interesting features. First, note that τ_1 does not appear in the equation at all. So long as $\vec{\lambda} - \vec{\tau}_2 = \vec{0}$, this term drops out, meaning the monopolists' profit maximizing strategy does not change with a tax on monetization.

Second, note the deep connection between Katz-Bonacich centrality, previously found occurring in the solution to network optimization problems, and equation (19). The term $\left(I - \frac{\bar{\boldsymbol{B}}' + \bar{\boldsymbol{B}}}{2}\right)^{-1}$ corresponds to the Katz-Bonacich centrality of the average of the *influencing* and *influencable* matrixes.[21] This quantity, which summarizes the importance to cumulative network effects of a node to the matrix is multiplied by the difference $-(\bar{\boldsymbol{B}}' - \bar{\boldsymbol{B}})$. Intuitively, a monopolist should set the price lower for (or even subsidize) market segments that have large positive externalities and set the price higher on those who receive lots of benefits from others' externalities, and this effect should be more extreme for more central market segments.

To make the managerial implications even clearer, it is useful to substitute back in the microfoundation terms. Equation (20) displays the first-order condition for profit maximization for group i, but only for the first cascade of the network effect. In other words, this equation is exactly correct for the level of participation one period after a change in monetization on group i:

$$\begin{aligned} \frac{\partial \Pi_{ss}}{\partial \phi_i} \approx \frac{\partial \Pi_2}{\partial \phi_i} &= \underbrace{(1-\tau_1)P_i - ((1-\tau_1)\phi_i + \lambda_i - \tau_{2,i})\frac{\partial \mu_i}{\partial \phi_i}\frac{\partial P_i}{\partial \mu_i}}_{\text{Direct Effect}} \\ &\quad \underbrace{- \frac{\partial \mu_i}{\partial \phi_i}\frac{\partial P_i}{\partial \mu_i}\sum_{j\neq i}^{J}((1-\tau_1)\phi_j + \lambda_j - \tau_{2,j})\frac{\partial P_j}{\partial \mu_j}\frac{\partial \mu_j}{\partial P_{t-1,i}}}_{\text{First Cascade Network Effect}}. \end{aligned} \tag{20}$$

The simplified first-order condition presented in equation (20) says that at the firm's utility maximizing price two effects are equal. These effects are denoted "Direct Effect" and "First Cascade Network Effect". The direct effect consists of two terms. Together, they show the immediate consequence of raising the amount of advertising on side i by one dollar. This will raise revenue, based on that side's current likelihood of participation ($(1-\tau_1)P_i$), and lose utility based on how important those users currently are to the firm ($((1-\tau_1)\phi_i + \lambda_i - \tau_{2,i})$) and how elastic that side's participation is to

[21] $\bar{\boldsymbol{B}}$ tells us how much P_i will change in the steady state if P_j changes. In other words, it is a measure of the sensitivity of each group to network effects. From our microfoundation, we know the elasticity of each group to network effects is the product of its elasticity of demand and how much they value having friends on the platform. Conversely, $\bar{\boldsymbol{B}}'$ speaks to the influence of P_i on any given P_j.

advertising ($\frac{\partial \mu_i}{\partial \phi_i}\frac{\partial P_i}{\partial \mu_i}$). The two direct effect terms are what normal firms have to consider when pricing their products. Note that when $\frac{\partial \mu_j}{\partial P_{t-1,i}} = 0 \; \forall \; i, j$, i.e. when no network effects are present at the current margin, equation (20) reduces to this pair of terms.

The last term in equation (20) is the network effect of an advertising increase after a single cascade. The increase in advertising makes i less likely to participate (by an amount $\frac{\partial \mu_i}{\partial \phi_i}\frac{\partial P_i}{\partial \mu_i}$) which leads others to stop participating (by an amount $\frac{\partial P_j}{\partial \mu_j}\frac{\partial \mu_j}{\partial P_{t-1,i}}$). When these third parties stop participating, the platform loses on the current revenues that they were paying ϕ_j, net of taxes and plus any non-pecuniary value.

In other words, to maximize utility one period into the future, the fee or level of disutilitous advertising should be increased on user i if the increased revenue (P_i) is greater than the decreased revenue from the person directly impacted possibly dropping out plus the decreased revenue from all the charged person's friends potentially dropping out.

3.5 Social Welfare

Social welfare is the sum of producer surplus and consumer surplus:[22]

$$U = \big((\Phi' + \omega\lambda')P + \vec{CU}\big)'\vec{1} + F, \tag{21}$$

where ω is a scalar indicating what share of the platform's shadow value from a large user base "counts" in social welfare, and $\vec{CU}$ is a vector giving consumer surplus for each group.

Consumer surplus for a group i is μ_i, platform quality for that group, times the number of individuals on side i who participate, less opportunity costs for every individual who decides to participate:

$$CU_i = \int_0^{P_i} \mu_i - \epsilon_i(\mathrm{i}) d_{\mathrm{i}}. \tag{22}$$

Assuming, as we do throughout, that demand is locally linear, we can derive a new linear approximation for consumer surplus:[23]

$$CU_i = \frac{1}{2}\Big(\frac{\partial P_i}{\partial \mu_i}\Big)^{-1} P_i^2, \tag{23}$$

which, when plugged into 21 yields the following long term equilibrium change in

[22]Note that advertiser surplus does not explicitly appear in this equation. One option is for advertisers (and other suppliers of inputs to the platform) to be incorporated in the model as a side, creating positive or negative network effects. In our theorems and calibration, we treat the monetization decision of the firm as only changing platform and consumer welfare. This can be justified by assuming markets which consume advertising are competitive, so Meta captures any value from the additional advertising slots they create.

[23]See Appendix B for details of derivation.

welfare due to a change in monetization for some side i, written out in summation notation:

$$\frac{\partial U}{\partial \Phi_i} = P_i + \sum^J \left((\phi_j + \omega\lambda_j)\frac{\partial P_{ss,j}}{\partial \phi_i} + \sum^J \left(\frac{\partial P_j}{\partial \mu_j}^{-1}\frac{\partial P_{ss,j}}{\partial \phi_i} P_j\right), \tag{24}$$

which when evaluated at the monopolist's goal maximizing price yields:

$$\frac{\partial U}{\partial \Phi_i}\Big|_{\vec{\Phi}^M} = \tau_1 P_i + \sum^J \left((\tau_1\phi_j + (\omega - 1)\lambda_j)\frac{\partial P_{ss,j}}{\partial \phi_i} + \sum^J \left(\frac{\partial P_j}{\partial \mu_j}^{-1}\frac{\partial P_{ss,j}}{\partial \phi_i} P_j\right), \tag{25}$$

and when evaluated at the competitive price, but allowing for cross-platform externalities (by assumption $\vec{\Phi} = 0$) yields:

$$\frac{\partial U}{\partial \Phi_i}\Big|_{\vec{\Phi}=0} = P_i + \sum^J \omega\lambda_j\frac{\partial P_{ss,j}}{\partial \phi_i} + \sum^J \left(\frac{\partial P_j}{\partial \mu_j}^{-1}\frac{\partial P_{ss,j}}{\partial \phi_i} P_j\right). \tag{26}$$

3.6 Motivations for Maintaining a Large User Base

Before closing the model's theoretical section with a comparative static welfare analysis, it makes sense to take a second to consider the terms ω and $\vec{\lambda}$. Recall that λ represents the firm's motivation to maintain a large user base, over and above the impact on net-revenue. ω, represents what share of this value should figure in social welfare. If a platform creates large negative externalities on society, $\omega\vec{\lambda}$ may even be negative. We show that the welfare implications of different social regulations will depend heavily on the nature of these terms.

A platform's operator may value a large user base intrinsically for several reasons. These may be distinguished into two categories: pro-competitive and anti-competitive. A large user base enables data collection for analysis that will lead to better products. Alternatively, or additionally, a large user base may create opportunities for profiting off future products. In the case of Facebook, this may represent the desire to maximize the sales of Metaverse and Oculus VR services, Libra digital currencies, or other future products. This motivation is a kind of intangible investment and could be viewed as a social good that is being internalized by the firm.

On the other hand, a platform operator's desire for a large user base may be anti-competitive. A platform may keep prices artificially low (or quality artificially high) to deter the entry of competitors or to fend off regulators. In either view, Facebook's courtship of a large user base at the expense of profits is a sign that the threat of entry (or regulation) induces its actions. It does not make the low level of monetization itself bad, but it suggests that an even more socially beneficial outcome (e.g., a more competitive or differently regulated social media industry) might manifest in the absence of these actions.

Also ambiguously, a large user base might be cultivated to prevent the network from "unraveling." Unraveling is a process by which a shock to a platform or market leads to increasingly larger cascades of users leaving the platform. This can occur due to negative platform quality shocks when platform participation is in an unstable equilibrium, as discussed above. It is, in essence, the opposite of the salutary cycle of positively reinforcing network effects that drives the early growth of large platforms. Instead of additional users positively reinforcing the network and one another, a steady withdrawal of users leads to a run on the platform. It is theoretically and empirically possible for a moderate departure of users to trigger such a cascade. If Facebook has estimated its collapse once its user base fell below some critical mass, the platform might build up many more users than that minimum amount as a buffer. Whether to think of this last motivation as anti-competitive is unclear. While some research has found that taking steps against unraveling is Pareto-efficient (Hałaburda, 2010), the chaos of a collapsing dominant platform could allow a more competitive or otherwise more socially beneficial industry to emerge.

Finally, participation on a social network might cause all sorts of externalities that are not internalized in producer or consumer surplus. Use of a social network might help individuals better coordinate against corruption, or might promote the spread of fake news. Negative externalities would be a reason for the social planner to choose an $\omega \leq 1$ and vice versa. Our preferred interpretation is $\omega = 1$, but in the calibrated model, we present results for both $\omega = 1$ and $\omega = 0$. These are labeled the change in social welfare with and without (SV) or "shadow values" respectively.

3.7 Comparative Static Analysis

We now evaluate how different sorts of government interventions may impact welfare.

Theorem 3.2. *If all elements in the matrix $\frac{\partial \vec{P}_{ss}}{\partial \vec{\Phi}}$ are negative, $\tau_1 = 0$ and $(\omega - 1)\lambda_j \geq 0$ $\forall j$; Or if $\vec{\lambda} = \vec{0}$ and $\vec{\tau}_2 = \vec{0}$ then:* ***a marginal decrease in monetization from its monopoly level weakly increases welfare.***

Proof. For the first condition, direct from equation (25). Given this condition all terms in equation (25) are weakly negative as the demand curve $\frac{\partial P_j}{\partial \mu_j}$ must be weakly positive, as is participation P_j. For the second condition, given the assumption, the first two terms in equation (25) cancel (because these two terms are just the monopolist's first-order condition, which must sum to zero at the monopoly price, multiplied by the constant $\frac{\tau_1}{1-\tau_1}$), leaving only the final necessarily negative term. □

Intuitively, the monopoly price is usually too high because the monopolist ignores the infra-marginal impact of lowering prices. The monopolist only cares about how prices impact marginal users. The final term in equation (25) corresponds to this non-internalized consumer surplus. This term shows that the social benefit of decreasing prices from their monopoly level is particularly high when P is large and when $\frac{\partial P}{\partial \mu}$ is small. When P is large, there are more infra-marginal individuals to benefit from lower prices, and greater participation and therefore network effects. When $\frac{\partial P}{\partial \mu}$ is large, the marginal user is more elastic, and the firm needs to keep prices lower to satisfy them, so there's less of a benefit from government intervention.

It is important to note as well that the monopoly price can be too low. This can be the case when the firm places too high a weight on a non-socially positive, non-pecuniary motivation to retain customers (i.e., $(\omega - 1)\lambda_j \leq 0$). The reasons for this may include the platform keeping prices low to deter entry (e.g., dumping), or that the platform has harmful externalities.

Corollary 3.2.1. *Given our microfoundation, a sufficient condition for for all elements of* $\frac{\partial \vec{P}_{ss}}{\partial \vec{\Phi}}$ *to be less than or equal to 0 is for the following to hold:*

1. *All elements of* $\boldsymbol{B}$ *are greater than 0 and,*
2. *All elements of* $\frac{\partial \vec{\mu}_t}{\partial \vec{\phi}_t}$ *are weakly negative, with the diagonal terms strictly negative.*

Proof. All elements of the elasticity of demand matrix $\frac{\partial \vec{P}}{\partial \vec{\mu}}$ are weakly greater than 0 (i.e., positive elasticities of demand to increases in quality). Therefore, if the above conditions hold, an increase in any monetization Φ_i must weakly decrease participation, as the first cascade of the network effect will weakly reduce participation for all groups i and all future recursions will see further reductions in participation (from equation (3), inserting equations (11) and (12)). □

We argue that both of these conditions hold for Facebook at the margin. All elements of $\boldsymbol{B}$ being greater than 0 means that all steady state network effects are positive. This is plausible given that all relationships on Facebook are voluntary. Second, the fact that advertising is disutilitous at the margin is necessary due to the assumption the firm is profit maximizing.

Theorem 3.3. *If* $\frac{\partial \mu_i}{\partial \Phi_{i,t}}$ *is sufficiently close to 0, or* $\omega\lambda_j$ *sufficiently large and negative, then, allowing for network effects to spillover across firms,* ***an increase in monetization above the competitive level will increase social welfare***

Proof. From equation (26), inserting equations (11) and (12).

□

Alternatively, the effect of a change in monetization level on welfare starting from the competitive price is generally ambiguous, even in the $\tau_1 = \tau_{2,i} = \lambda_i = 0 \ \forall \ i$ case; assuming that network effects are preserved across the competitive firms. A flat, low level of monetization may be good for participation, but even more social welfare may be achieved by subsidizing one side and increasing fees on another. This is especially true if $\frac{\partial \vec{\mu}_t}{\partial \vec{\Phi}_t}$ is close to zero, indicating that users only slightly care about marginal monetization. More social welfare may be generated in that case by raising the level of advertising, which would only slightly lower consumer surplus and participation, but have a larger positive effect on producer surplus. In other words, a competitive market for social media could, theoretically, lower the level of advertising below the Kaldor-Hicks efficient level.

On the other hand, one simple reason for the monetization of digital platform to be too low is if it has large negative social externalities ($\omega\lambda_j$ very negative). In this case it is clear that a Pigouvian tax would raise social welfare.

Theorem 3.4. *In the absence of a per-capita incentive to acquire users (i.e. $\vec{\lambda} - \vec{\tau}_2 = \vec{0}$), marginal increases in corporate income taxes τ_1*

1. *do not change monetization levels $\vec{\Phi}$,*
2. *do not change platform participation, and*
3. *do not change social welfare*

i.e., the incidence of a corporate income tax is entirely on monopoly platform profits

Proof. Direct from the monopolist's solution in equation (19), which does not include a τ_1 term directly or indirectly. If $\vec{\Phi}$ does not change, then platform participation and social welfare does not change. □

Theorem 3.5. *If the platform has a positive non-monetization incentive to acquire users (i.e. all elements of the vector $\vec{\lambda} - \vec{\tau}_2$ are greater than 0), and if all elements in the matrix $\frac{\partial \vec{P}_{ss}}{\partial \vec{\Phi}}$ are negative, then marginal increases in corporate taxes τ_1:*

1. *weakly decrease monetization levels $\vec{\Phi}$,*
2. *weakly increase platform participation rates $\vec{P}$, and*
3. *weakly increase consumer surplus.*

The opposite holds when all elements of the vector $\vec{\lambda} - \vec{\tau}_2$ are less than than 0

Proof. Consider equation (18). An increase in corporate tax rates decreases the first set of terms. All terms in the vector $\frac{\partial \vec{P}_{ss}}{\partial \vec{\Phi}}'(\vec{\lambda} - \vec{\tau}_2)$ must be negative, as the vector $\vec{\lambda} - \vec{\tau}_2$ only has positive elements and the matrix $\frac{\partial P_{ss}}{\partial \Phi}$ only has negative elements. Therefore, for the firm to have been utility maximizing, the first term must be weakly positive.

Therefore, an increase in τ_1 must make $\frac{\partial \vec{\Pi}}{\partial \vec{\Phi}}$ weakly negative, until the platform reacts. This will lead the platform to lower prices, which will increase platform participation and consumer surplus.

□

Theorem 3.6. *If all elements in the matrix* $\frac{\partial \vec{P}_{ss}}{\partial \vec{\Phi}}$ *are negative, all elements of* $\omega\vec{\lambda} \geq 0$*, and all elements in* B *are positive, keeping* $\vec{\Phi}$ *fixed and rebating the revenues to platform participants:*

1. *weakly increases platform participation,*
2. *weakly increases consumer welfare, and*
3. *weakly increases social welfare.*

Proof. Rebating $\vec{\Phi}$ revenues to users has a neutral direct effect on social welfare (profits decrease by the same amount as the transfer) but increases platform participation and consumer welfare. $\omega\vec{\lambda} \geq \vec{0}$ guarantees that non-pecuniary effects are positive as well. □

4 Calibrating the Model for Facebook

While the direction of many policies can be signed through theory alone, we can also simulate their effects in a calibrated model to quantify their relative magnitude. The setting for our empirical illustration is Facebook. Facebook is an ad-supported social network. It was selected because it is used by a very large percentage of the US population, and previous research has determined that many value it highly (Brynjolfsson et al., 2019).

We calibrate our model on estimated demand for Facebook, the value of it's characteristics, as well as complementary data from the Facebook Ad API and other sources. To produce these estimates, we conducted 57,195 surveys on a representative sample of the U.S. population that uses the internet in the summer of 2019. We estimate the potential population, current user-base, elasticity of demand, network effects, and disutility of advertising separately for twelve market segments: two genders and six age brackets. Full details on the calibration and simulation of the structural model are provided in Appendix C and details on the survey instrument are provided in Appendix D.

Figure 3 presents Facebook usage and network externalities by market segment. The size of each node represents the relative current size of the Facebook user base by demographic. The thickness of the arrows corresponds to the value received by a Facebook user of the demographic the arrow is pointing towards from an additional

Facebook friend of the source demographic (i.e., $\frac{\partial \mu_i}{\partial P_j}$). The top figure reports all 12x12 network externalities, the middle figure reports the eight strongest externalities, and the bottom figure reports the three strongest externalities. The three friendship externalities in the bottom figure are all worth more than 50 cents a month on average, with the typical friendship worth much less.

As can be seen, there are more female users of Facebook overall and within each age group. The thickest lines in Figure 3 flow from right to left, and from the bottom to the top. In other words, on Facebook, value tends to flow from younger and male users to older and female users. Appendix figure E15 restricts attention to the Facebook friendship network effects experienced and caused two nodes of interest: Females 65+ and Males 18-24. Females age 65+ most value connections to men 18-24, perhaps corresponding to connections to grandchildren and nephews. They provide the most value to middle-aged women, with young people hardly valuing the connections at all. Males 18-24 provide the most value to the oldest women, with middle-aged women next. They value individual connections of all types only slightly, but amongst these most valuable connections are to Males age 55-64.

Figure 3: Facebook usage and network externalities by market segment. The size of each node represents the relative current size of the Facebook user base by demographic. The thickness of the arrows corresponds to the relative value received by a Facebook user of the demographic the arrow is pointing towards from a friendship with a user of the source demographic (i.e., $U_i(j)$ with i being where the arrow is pointing towards, and j being the source of the arrow). The top figure displays all bilateral connection values, the middle figure only eight most valuable connections, and the bottom only the three most valuable connections.

5 Simulation Results

With the calibrated model in hand, we can now proceed to simulating counter-factual pricing strategies and potential government policies. We begin by estimating Facebook's profit maximizing strategy, under the assumption that they have no non-monetary size motivation. We then calculate λ, the non-monetary value Facebook places on users which justifies their current monetization strategy as optimal. This shadow value is included in all further policy simulations.

5.1 Facebook Profit Maximization

We begin by calculating Facebook's profit maximizing level of monetization. To calculate this, we iterate through guesses of different ϕ_i's for each demographic group until we identify a global maximum. We find that Facebook's profit maximizing strategy entails a large increase in the level of monetization. Therefore, for this analysis only, we assume that the marginal disutility from increased monetization $\frac{\partial \mu_i}{\partial \phi_i} a_i$ is equal to 1 for each group.[24] While our model and simulation program allows for price discrimination among different demographic groups, we have no evidence on how or if Facebook does so. This is critical for calibrating the model to allow for price discrimination. Therefore, our simulations throughout this paper only allow Facebook to vary the overall intensity of advertising.

Figure 4 displays the change in Facebook ad revenues and consumer welfare after N cascades in billions of dollars per month. We find that Facebook's profit maximizing strategy entails increasing fees substantially.

Implementing this strategy would increase Facebook revenues by \$2.38 billion dollars per month (from a baseline of \$1.79 billion) at the cost of decreasing its user base by 49.1% and lowering consumer surplus by 42.1% (from a baseline of 12.2 billion). In other words, this strategy entails squeezing Facebook's most inelastic users for a much higher share of their surplus. These results also show the importance of network effects. If no network effects were present, Facebook would be able to raise revenues by 211.2% by implementing this fee increase (column "Cascade 0") at the cost of decreasing its user base by 32.3%. However, the first wave of users leaving reduce Facebook's quality to remaining users. Cascade one displays the effect of the first cascade of this network effect – an additional 11.3% of Facebook's initial user-base leaves, reducing Facebook revenue and further degrading platform quality. After three cascades of the network effect, Facebook reaches its new equilibrium to three significant digits of precision.

Why is Meta is leaving so much money on the table? One possibility is that it

[24] $a_i = 1$ is a logical upper bound, because Facebook could always simply charge a fee for use. In the policy simulations, which generally entail a reduction in advertising rates, we use our estimated a_i throughout.

		Initial	Ad Increase Cascade 0	Cascade 1	Cascade 2	Cascade 3	Cascade 20	Cascade 1000
Consumer Surplus (Billions of Dollars a Month)	Female 18-24	0.37B	-34%	-47%	-51%	-53%	-53%	-53%
	Female 25-34	1.76B	-15%	-31%	-36%	-38%	-38%	-38%
	Female 35-44	2.12B	-14%	-31%	-36%	-38%	-38%	-38%
	Female 45-54	1.05B	-35%	-49%	-54%	-56%	-56%	-56%
	Female 55-64	1.38B	-28%	-42%	-46%	-48%	-48%	-48%
	Female 65+	1.34B	-30%	-43%	-47%	-49%	-49%	-49%
	Male 18-24	0.52B	-22%	-34%	-38%	-40%	-40%	-40%
	Male 25-34	0.84B	-21%	-34%	-38%	-40%	-40%	-40%
	Male 35-44	0.65B	-27%	-40%	-44%	-46%	-46%	-46%
	Male 45-54	0.60B	-28%	-40%	-44%	-46%	-46%	-46%
	Male 55-64	0.60B	-11%	-24%	-28%	-30%	-30%	-30%
	Male 65+	0.99B	-5%	-20%	-24%	-26%	-26%	-26%
Tot. Consumer Surplus (Billions)		12.22B	-21.3%	-35.4%	-39.9%	-42.1%	-42.1%	-42.1%
Ad Revenue (Billions)		1.79B	211.2%	158.5%	141.2%	133.1%	133.1%	133.1%
Participation (Millions)		154.11M	-32.3%	-43.6%	-47.3%	-49.1%	-49.1%	-49.1%

Figure 4: Change in consumer surplus and Facebook profit after N cascades in billions of dollars per month, after Facebook implements its profit maximizing monetization strategy.

values having a large and happy user base. This could be because they value the data produced by a large user base (either for resale or for internal development), because they plan to monetize the user base further in the future (e.g., keeping a marginal user on Facebook might increase the odds that they use Libra or some Oculus product in the future), to give the platform some buffer between its current level of usage and a lower level where unravelling might become possible, or because having a large user base deters the entry of competitors. The first two motivations imply that a large user base is socially valuable as well as valuable to Facebook (the latter two are ambiguous). This is the interpretation we focus on below.

5.2 Social Welfare Maximization

Before evaluating the impact of different possible government actions, it makes sense to evaluate how far the current regime is from the first best – a nationalized Facebook which sets prices to maximize social value. The results of this policy are listed in Appendix Table E2. Due to the positive externality of Facebook participation, the social welfare maximizing policy entails a subsidy for Facebook use. The size of this subsidy is -381.7% of initial revenues (the first 100% of which would be achieved by eliminating advertising). To be consistent with other results below, we assume that the nationalized Facebook can subsidize usage (alternatively, raise Facebook participation quality) at the same rate at which it raises revenues: i.e. a_i.[25] Under this policy,

[25]There are additional reasons for restricting attention to a frictional subsidy. First, because Facebook usage carries a positive externality, if Facebook usage could be subsidized at \$1 creating an additional dollar of utility for users (and if there was no distortion from taxation) the optimal policy would be infinite advertisements and an infinite subsidy. In our 'Data as Labor' scenario we make the alternate assumption that a subsidy of \$1 does raise Facebook use utility by the same amount, but limit the size of this rebate to

consumer surplus increases by about 23.9%, but this is offset by the decrease in profits. The third row reports the change in social welfare from the increase in consumer surplus and decrease in profits alone. This would actually correspond to a 27.9% decrease in social welfare. Intuitively, ads on Facebook are in some sense productive, because they create less than $1 of disutility for every dollar raised. The fourth row in the table includes the "shadow value" of maintaining a large user base in calculating the change in social surplus. This shadow value is accounted for only in the margin (i.e., the percentage change is from the sum of the consumer surplus and the monetary producer surplus). Counting this shadow value, a first best nationalization of Facebook would raise social surplus by 9.6%, or about $1.3 billion dollars per month.

5.3 Evaluating Tax and Redistribution Policies

We next simulate the consequences of three tax and redistribution policies: an ad-revenue tax, a per-user tax, and "Data as Labor". The two taxes we simulate are a tax on advertising revenues and a per-user tax. A tax on ad revenues has been proposed by leading economists, such as Paul Romer (Romer, 2019). A three percent tax on sales of ads by large online platforms has recently been passed by France, but has not yet been implemented (CNBC, 2019). Grauwe (2017) proposes a 10 dollar per-user tax.[26] A third proposal is the "Data as Labor" framework proposed in Posner and Weyl (2018). In this framework, perhaps through a collective bargaining process, users would be compensated for their 'labor' in providing data and viewing advertisements. We operationalize this last policy as Facebook maintaining its current level of advertising, but rebating to each user the net revenue it collects from displaying them ads.

Before we proceed to the simulations, we have a novel theoretical point to make about the incidence of taxes on digital platforms. As we find in theorem 3.4, in the absence of a a non-revenue per-person incentive to acquire users, a flat tax does not distort the platform's optimal vector of ϕ_i's.

In our calibrated model, we find that Meta does derive non-monetary utility from maintaining a large user base. Therefore, a tax on ad revenues will cause it to shift between its two tasks. It will shift from making revenues from selling advertising to increasing utility by cultivating a large user base. Thus will raise social welfare (see Theorem 3.5) On the other hand, a tax on the amount of users will lead the platform to adopt a strategy that tries to squeeze a smaller share of users for more of their surplus. [27] Theorem 3.6 tells us that rebating ad revenues to users will be social

the current total net-revenue from advertising.

[26]Professor Grauwe proposes this as an annual levy, but we consider a per-month tax that would raise the same amount as a 3% of revenues tax.

[27]In the case of an ad tax that only applied to certain jurisdictions or demographic groups, there would be an incentive for the firm to increase monetization of users who provide value to the taxed group.

	Current (millions)	3% Tax	Per Capita Tax	Data as Labor
Net Ad Revenue	1,790.8	-21.4%	-1.6%	-100.0%
Consumer Surplus	12,219.8	1.3%	-0.1%	17.8%
Social Welfare (No SV)	14,010.6	-1.3%	0.1%	2.7%
Social Welfare (Including SV)		1.1%	0.0%	30.3%
Tax Revenue (in % of initial revenue)		2.4%	2.4%	0.0%
Number of Users	154.1	1.1%	0.0%	12.1%

Table 1: Evaluating tax and redistribution policies. Current and simulated % change in Facebook ad revenue, consumer surplus, social surplus, tax revenue and number of users after three policies are implemented. 3% tax indicates a tax on ad revenues. Per Capita tax indicates a tax on the number of users that raises the same amount of revenue as a 3% tax. Data as labor indicates a policy of rebating to users 100% of the revenue from the advertisements they see, keeping the level of advertisement constant. Baseline social welfare is the sum of Facebook ad revenues and consumer surplus. Percentage increase in social welfare "with SV" includes the non-monetary "shadow value" of maintaining a larger user base in calculating the change in social welfare, while "no SV" excludes this value (i.e., $\omega = 1$ or 0, respectively).

welfare enhancing as well.

While theory can tell us the direction of effects, we need to use simulations to understand their magnitudes. Table 1 summarizes the results of these three simulation experiments. As our theory suggested, a per-user tax slightly decreases the number of users and consumer surplus. On the other hand, a 3 percent tax slightly boosts consumer surplus and participation rates. However, it has a disproportionately negative effect on Facebook net revenues, because Facebook reduces its level of advertising in response. The "Data as Labor" policy has the most positive implications. Advertising, which is productive in the sense that it raises more revenue than the direct disutility it causes, is used to fuel a transfer to users. This directly makes users better off, and has a knock on effect of attracting additional users to the Facebook platform, who themselves provide positive spillovers to inframarginal users. About 58% of the welfare increase is due to the direct transfer to current users with the remainder due to new users who join the platform, consuming more ads and providing more value to other users. Rebating users for the value that they create on social media is a clear win-win from the perspective of social welfare. The question is whether such a policy is implementable, given the clear incentive to create fake accounts.

5.4 Evaluating Platform Regulation Policies

The final set of policies we simulate are regulatory. Many proposals have been made for the regulation of digital platforms and social media, some of which (especially those regarding "Fake News" and political manipulation) are beyond the scope of this current study. Here we consider a set of three potential reforms. The first simulates achieving perfect competition in social media, while preserving network effects for users of different Facebook-like social media services. We also simulate two Facebook break-up scenarios that fail to increase competition.

In principle, it is not obvious whether decreasing the market power of a digital platform is a good or bad thing for social welfare. On the positive side, completely eliminating market power would force platforms to 'price' at their marginal cost – here assumed to be zero. It might also have positive political implications. However, even if network effects were preserved across users of different Facebook like social media services, increasing competition may have negative effects. We show in Theorem 3.3 that increases in monetization from the competitive level have ambiguous effects, and may even increase social welfare. Still, the conditions for a decrease in monetization from the monopoly level to raise social welfare outlined in Theorem 3.2 apply, and so, in this setting, the reduction in price due to greater competition will raise social welfare.

In our first simulation we assume that the decrease in Meta's market power over Facebook-like social media services preserves network effects. But a change that decreases market power without preserving network effects across platforms is less likely to increase welfare. If multi-homing is costly and network effects do not spillover across platforms, then increasing the number of platforms may decrease the positive network effects that are the main source of value on digital platforms. To resolve this last concern, a recent study of antitrust and regulation in the context of digital platforms, (Scott Morton et al., 2019), has called for mandated 'interoperability' alongside other policy changes that would lower barriers to entry. Interoperability would require, for example, Facebook to share posts and other communiques with competitor social networks, who would then be allowed to display them on their platforms. Our 'perfect competition' scenario assumes this interoperability.[28]

One component of many plans to increase platform competition include mandatory breakups. For example, leading politicians and regulators have called for, among other things, Facebook to be split from Instagram and Whatsapp (Warren, 2019). To the

[28] There are also other reasons increasing competition could be bad. First, and most theoretically interesting, a monopolist can cross-subsidize different sides of a market in a way that a competitive firm cannot. In the same way that a government might subsidize an infant industry for the good of the total economy in the long-run, a monopolist platform is a sort of 'stationary bandit' who has an interest in taking into account at least some network effects. This incentive differs from the social planners' interest in that the monopolist only cares about the network effect on marginal platform users (rather than on all platform users). Another reason market power might be good in this setting in particular is that it might prevent 'production' through advertising. Because advertisements raise more revenue than the disutility they directly cause, the social welfare optimum may include a positive, rather than zero, amount of advertising. Of course, this argument is null if advertising revenues can be rebated (as we assumed in the "Data is Labor" case above), but one can imagine several frictions that might cause this.

	Current (millions)	Perfect Competition	Horizontal Breakup	Vertical Breakup
Net Ad Revenue	1,790.8	-100.0%	-49.5%	-15.1%
Consumer Surplus	12,219.8	6.6%	-33.0%	-3.9%
Social Welfare (No SV)	14,010.6	-7.0%	-35.1%	-5.3%
Social Welfare (Including SV)		4.8%	-84.7%	-10.1%
Number of Users	154.1	5.2%	-21.8%	-2.1%

Table 2: Evaluating regulatory policies. Current and simulated % Change in Facebook advertisement revenue, consumer surplus, social surplus and number of users after three policies are implemented. Perfect competition entails an optimally implemented regulatory policy that drives the price (i.e. advertising level) of social media services to zero, but does not split user bases in a way that reduces network effects. Horizontal breakup simulates a failed regulation that left the US with two smaller non-competitive Facebook-like social media platforms. Vertical breakup simulates a regulation that reduces Facebook quality for 5% of users without increasing competition. Percentage increase in social welfare 'with SV' includes the non-monetary 'shadow value' of maintaining a larger user base in calculating the change in social welfare, while 'no SV' excludes this value (i.e. $\omega = 1$ or 0, respectively).

extent that these are separate platforms that do not allow for network effects across them, such a breakup is sensible. But one can imagine a breakup of Facebook that both destroyed network effects and failed to increase competition (e.g., by dictating that users must use only one of the two platforms). We model two such breakup scenarios. The first is a horizontal Facebook breakup resulting in the creation of two Facebook monopolies each serving half of the US population.[29] The final scenario we simulate is a 'vertical breakup' that results in five percent of the US population losing interest in using Facebook, without any increase in competition. This represents a guess of the percentage of the population that uses Facebook only because of its synergies with Instagram, and would stop using it if these products were completely disconnected.

Results from these simulations are summarized in Table 2. We find that perfect competition would raise consumer surplus by 6.6%, at the cost of eliminating all monetary profits. Taking only Facebook's monetary revenues into account, perfect competition actually lowers social surplus (-7%), because the reduction in ad revenues is larger than the reduction in consumer welfare. However, if a large user base is still assumed to create social surplus at the same rate as for Facebook today, the policy creates a clear social welfare improvement of 4.8%. Framed differently, roughly half of the 9% increase in social welfare that could be achieved through a nationalized Facebook would be gained through perfect competition. The two other breakup scenarios, which fail to preserve cross-platform network effects without producing an offsetting increase in platform competition, are negative for social welfare. The worse of these is

[29] Such a scenario is not that far-fetched. The breakup of 'Ma' Bell Telephone led to the creation of several regional monopolies and one 'long-distance' monopoly, with regulated prices.

the horizontal breakup, which would lower social welfare by as much as 84.7%.

6 Discussion

Building on Rochet and Tirole (2003), Parker and Van Alstyne (2005), and Weyl (2010), we contribute a tractable model of a network good. Taking the first order condition for profit maximization with respect to the monetization schedule yields a recursive equation that can be evaluated to the desired degree of precision. The managerial insight is that platform owners should increase advertising on market segments which inelastically demand the platform (the direct effect), do not have much disutility from advertisements, and do not create much network value for others. Platforms should decrease advertisements on those who elastically demand the platform and create high amounts of network value for other profitable users who demand the platform elastically (the first cascade of the network effect).

We build a simulation tool to evaluate the consequences of different firm strategies and government interventions. This model is implementable in the sense that there is a clear strategy for measuring all the terms in the model. It is scalable in the sense that these terms can be measured with as much precision and for as small a market segment as desired. When calibrating the model, we must make additional assumptions about the functional form of user demand for the platform. A platform firm itself, with much more information about its users, would be able to implement a more precise and detailed version of this simulation tool.

Previous papers, including Weyl and White (2014) have made the theoretical point that digital platforms with market power introduce two distortions: a classical markup above marginal cost, and a failure to internalize network effects provided to inframarginal users (the Spence distortion). We find that Facebook's market power reduces social welfare by 9.6% relative to the first best, and 4.8% relative to perfect competition. In other words, we find that the Spence distortion is quantitatively just as important as the classical market power distortion.

In addition to the assumptions embedded in our parameterization of the model, these results incorporate two additional key assumptions. The first is the assumption that Meta's shadow value from maintaining a large user base should be incorporated into estimates of social welfare. If it should not be, then the implication of the model actually flips: Facebook's market power actually enables it to do "productive" advertising, that would be eliminated under perfect competition, lowering social welfare. The only policy evaluated in the model that is unambiguously positive despite what one assumes about the social value of this shadow price is Data as Labor. That is because this policy both preserves productive advertising and further subsidizes platform usage. If somehow Facebook were able to directly compensate users for a percentage of the value that they create, this would create clear increases in social surplus that could be used to make the policy a Pareto improvement.

The other key ancillary assumption we make in evaluating the welfare consequences of policies is that the shadow value of users is not impacted by tax or regulatory policy. This drives the result that profit taxes are good for consumer welfare (they cause the platform to reduce the amount of advertisements they show, as they shift from a smaller more intensely monetized user base to a larger one) and that per-capita taxes are bad for consumer welfare (the converse). The sign of the first result will hold so long as the tax is *more* incident on monetary value than shadow values, which is plausible. Facebook breakups which reduce the quality of Facebook or the propagation of network effects without increasing competition are unambiguously bad.

Our approach is not without weaknesses. One important issue is difficulty in soliciting the necessary data to estimate the model. Consumers may not fully understand or reliably answer questions about their valuations for different friend groups. Poor memory may also be an obstacle. There may also be important differences between short and long-term elasticities of demand. Similarly, if individuals have very high variance or skewness in their platform valuations, network effects, or number of friends, the average of these values within a group may be a poor summary statistic – especially if these measures are correlated within a side of the market/demographic group. Relatedly, in our parameterization, we currently assume that the value from friends is linearly additive and that the disutility from advertising revenues is linear. While these are fine assumptions as local approximations, they become more problematic as the size of the proposed interventions increase. The most extreme interventions are far from the support of our survey data. However, with a larger budget, incentive compatible experiments, smaller market segments or within-platform proprietary data, each of these concerns could be addressed, and the nature of utility functions measured more precisely.

On the theoretical side, one limitation of the current approach is that demand for advertising is flat, rather than as an elastic side of the market. A more complete model would treat advertisers as a heterogenous mix of agents as well. Finally, our model conceives of consumers as atomistic price takers. This ignores the possibility that highly valuable users with market power might bargain with the platform or that users might unionize to demand a better equilibrium. The implications of such a scenario could be estimated in an extension of the model. In future work, we look forward to further refining and testing this model, and the estimates of the utility functions of platform stakeholders that underlie it.

While this particular study has important limitations, particularly in the estimation of key parameters, we encourage both regulators and platform designers to design and release their own structural models. Currently, the conflict over platform taxation and regulation is being conducted at a very conceptual level. While this is important, actual policymaking entails quantitative estimates of harms, benefits, and incidences. The FTC, or the 'Digital Authority' in Scott Morton et al. (2019)'s vision, should publicly release their own estimates of the impact of potential policies on different platforms to inform the public debate. Platforms would be welcome to generate rebuttal models, and would ultimately be incentivized to work with regulators. If necessary, regulators should be given powers to compel the platforms themselves to privately share the data needed to properly calibrate their models. This paper is not intended to end a conversation, but rather we hope that it is the beginning of a new dialogue that leads the digital economy to create greater benefits for all.

References

Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow, "The welfare effects of social media," Technical Report, National Bureau of Economic Research 2019.

Aral, Sinan and Christos Nicolaides, "Exercise contagion in a global social network," *Nature communications*, 2017, *8*, 14753.

— **and Dylan Walker**, "Identifying influential and susceptible members of social networks," *Science*, 2012, *337* (6092), 337–341.

Åström, Karl J and Björn Wittenmark, *Computer-controlled systems: theory and design*, Courier Corporation, 1997.

Ballester, Coralio, Antoni Calvó-Armengol, and Yves Zenou, "Who's who in networks. Wanted: The key player," *Econometrica*, 2006, *74* (5), 1403–1417.

Bernstein, Shai and Eyal Winter, "Contracting with heterogeneous externalities," *American Economic Journal: Microeconomics*, 2012, *4* (2), 50–76.

Boudreau, Kevin, "Open platform strategies and innovation: Granting access vs. devolving control," *Management science*, 2010, *56* (10), 1849–1872.

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

Candogan, Ozan, Kostas Bimpikis, and Asuman Ozdaglar, "Optimal pricing in networks with externalities," *Operations Research*, 2012, *60* (4), 883–905.

Ceccagnoli, Marco, Chris Forman, Peng Huang, and DJ Wu, "Cocreation of value in a platform ecosystem! The case of enterprise software," *MIS quarterly*, 2012, pp. 263–290.

CNBC, "France targets Google, Amazon and Facebook with 3% digital tax," Mar 2019.

Eisenmann, Thomas, Geoffrey Parker, and Marshall W Van Alstyne, "Strategies for two-sided markets," *Harvard business review*, 2006, *84* (10), 92.

Evans, David S and Richard Schmalensee, "Failure to launch: Critical mass in platform businesses," *Review of Network Economics*, 2010, *9* (4).

Fainmesser, Itay P and Andrea Galeotti, "Pricing network effects," *The Review of Economic Studies*, 2015, *83* (1), 165–198.

Grauwe, Paul De, "Why Facebook Should Be Taxed And How To Do It – Paul De Grauwe," Oct 2017.

Hagiu, Andrei, "Two-sided platforms: Product variety and pricing structures," *Journal of Economics & Management Strategy*, 2009, *18* (4), 1011–1043.

Hałaburda, Hanna, "Unravelling in two-sided matching markets and similarity of preferences," *Games and Economic Behavior*, 2010, *69* (2), 365–393.

Hohnhold, Henning, Deirdre O'Brien, and Diane Tang, "Focusing on the Long-term: It's Good for Users and Business," in "Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining" ACM 2015, pp. 1849–1858.

Huang, Jason, David Reiley, and Nick Riabov, "Measuring Consumer Sensitivity to Audio Advertising: A Field Experiment on Pandora Internet Radio," *Available at SSRN 3166676*, 2018.

Huang, Shan, Sinan Aral, Yu Jeffrey Hu, and Erik Brynjolfsson, "Social advertising effectiveness across products: A large-scale field experiment," *Marketing Science*, 2020, *39* (6), 1142–1165.

Jeon, Doh-Shin, Yassine Lefouili, Yaxin Li, and Tim Simcoe, "Complementary Multi-Sided Platforms," *Mimeo*, 2021.

Kawaguchi, Kohei, Toshifumi Kuroda, and Susumu Sato, "Merger Analysis in the App Economy: An Empirical Model of Ad-Sponsored Media," in "TPRC48: The 48th Research Conference on Communication, Information and Internet Policy" 2021.

Knill, Oliver, *Linear Algebra And Vector Calculus; Chapter 22: Stability*, Vol. II, Public Domain, 2019.

Meta, "Meta Reports Fourth Quarter and Full Year 2021 Results," 2022.

Parker, Geoffrey G and Marshall W Van Alstyne, "Two-sided network effects: A theory of information product design," *Management science*, 2005, *51* (10), 1494–1504.

Petersen, Kaare Brandt, Michael Syskind Pedersen et al., "The matrix cookbook," *Technical University of Denmark*, 2008, *7* (15), 510.

Posner, Eric A and E Glen Weyl, *Radical markets: Uprooting capitalism and democracy for a just society*, Princeton University Press, 2018.

Rochet, Jean-Charles and Jean Tirole, "Platform competition in two-sided markets," *Journal of the european economic association*, 2003, *1* (4), 990–1029.

Romer, Paul, "A Tax that Could Fix Big Tech," *The New York Times*, May 2019.

Rysman, Marc, "Competition between networks: A study of the market for yellow pages," *The Review of Economic Studies*, 2004, *71* (2), 483–512.

Scott Morton, Fiona, Pascal Bouvier, Ariel Ezrachi, Bruno Jullien, Roberta Katz, Gene Kimmelman, A Douglas Melamed, and Jamie Morgenstern, "Committee for the Study of Digital Platforms: Market Structure and Antitrust Subcommittee Report," 2019.

Segal, Ilya, "Coordination and discrimination in contracting with externalities: Divide and conquer?," *Journal of Economic Theory*, 2003, *113* (2), 147–181.

Stanković, Nikola, Sorin Olaru, and Silviu-Iulian Niculescu, "On stability of discrete-time delay-difference equations for arbitrary delay variations," *IFAC Proceedings Volumes*, 2013, *46* (3), 242–247.

Sun, Leping, "Stability analysis for delay differential equations with multidelays and numerical examples," *Mathematics of computation*, 2006, *75* (253), 151–165.

Tan, Hongru and Julian Wright, "A Price Theory of Multi-Sided Platforms: Comment," *American Economic Review*, 2018, *108* (9), 2758–60.

Tucker, Catherine, "Identifying formal and informal influence in technology adoption with network externalities," *Management Science*, 2008, *54* (12), 2024–2038.

Veiga, André, E Glen Weyl, and Alexander White, "Multidimensional platform design," *American Economic Review*, 2017, *107* (5), 191–95.

Warren, Elizabeth, "Here's how we can break up Big Tech," Mar 2019.

Weyl, E Glen, "A price theory of multi-sided platforms," *American Economic Review*, 2010, *100* (4), 1642–72.

— **and Alexander White**, "Let the Right'One'Win: Policy Lessons from the New Economics of Platforms," *University of Chicago Coase-Sandor Institute for Law & Economics Research Paper*, 2014, (709).

Yan, Jinyun, Birjodh Tiwana, Souvik Ghosh, Haishan Liu, and Shaunak Chatterjee, "Measuring Long-term Impact of Ads on LinkedIn Feed," *arXiv preprint arXiv:1902.03098*, 2019.

Appendix for Online Publication

A Derivation of the Monopolist's First Order Conditions and Locally Optimal Price

Plugging equation (17) into equation (16) yields

$$\Pi = (1-\tau_1)\vec{\Phi}'(I-\sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1}(\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}\vec{\Phi}+\vec{P}_0)+(\lambda-\tau_2)'(I-\sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1}(\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}\vec{\Phi}+\vec{P}_0)+F. \tag{27}$$

Taking the first derivative with respect to $\vec{\Phi}$ [30] yields the following:

$$\frac{\partial \Pi}{\partial \vec{\Phi}} = (1-\tau_1)(I-\sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1}(\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}\vec{\Phi}+\vec{P}_0)+\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}'(I-\sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1\prime}\vec{\Phi}-((\lambda-\tau_2)'(I-\sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1}\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t})'. \tag{28}$$

Multiplying through yields:

$$\begin{aligned}\frac{\partial \Pi}{\partial \vec{\Phi}} =&(1-\tau_1)(I-\sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1}(\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}\vec{\Phi}) + (1-\tau_1)(I-\sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1}(\vec{P}_0)\\ &+(1-\tau_1)\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}'(I-\sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1\prime}\vec{\Phi}+((\lambda-\tau_2)'(I-\sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1}\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t})'.\end{aligned} \tag{29}$$

This equation can be simplified by noting that $\frac{\partial \vec{P}_{ss}}{\partial \vec{\Phi}} = (I-\sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1}\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}$ (i.e., equation (4)). Plugging this in simplifies the equation to:

$$\begin{aligned}\frac{\partial \Pi}{\partial \vec{\Phi}} =&(1-\tau_1)\Big(\frac{\partial P_{ss}}{\partial \vec{\Phi}}\vec{\Phi}+(I-\sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1}\vec{P}_0+\frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}'(I-\sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})^{-1\prime}\vec{\Phi}\Big)\\ &+\frac{\partial P_{ss}}{\partial \vec{\Phi}}'(\lambda-\tau_2).\end{aligned} \tag{30}$$

In addition to deriving first order conditions, we can also solve explicitly for the monopolist's optimal price assuming the linear approximation holds. For notational convenience, define $B := (I-\sum_{y=1}^{Y}\frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})$ and $A = \frac{\partial \vec{P}_t}{\partial \vec{\Phi}_t}$. Rearranging equation equation 29 to solve for $\vec{\Phi}$ yields:

$$\vec{\Phi} = -(1-\tau_1)^{-1}(B^{-1}A+(AB^{-1})')^{-1}((1-\tau_1)B^{-1}\vec{P}_0+((\vec{\lambda}-\vec{\tau}_2)'B^{-1}A)') \tag{31}$$

[30] To do so we use the equation (78) from Petersen et al. (2008).

multiplying through yields

$$\vec{\Phi} = -(B^{-1}A + (AB^{-1})')^{-1}(B^{-1}\vec{P_0}) - (1-\tau_1)^{-1}(B^{-1}A + (AB^{-1})')^{-1}((\vec{\lambda} - \vec{\tau_2})'B^{-1}A)' \tag{32}$$

combining B^{-1} into the main term yields [31]

$$\vec{\Phi} = -(BB^{-1}A + B(AB^{-1})')^{-1}\vec{P_0} - (1-\tau_1)^{-1}(B^{-1}A + (AB^{-1})')^{-1}(A'B^{-1\prime}(\vec{\lambda} - \vec{\tau_2})) \tag{33}$$

And, by the definition of an inverse, this can be reduced to

$$\vec{\Phi} = -(A + B(AB^{-1})')^{-1}\vec{P_0} - (1-\tau_1)^{-1}(B^{-1}A + (AB^{-1})')^{-1}(A'B^{-1\prime}(\vec{\lambda} - \vec{\tau_2})) \tag{34}$$

Taking the special case where $A = -I$ (i.e. one unit of increased monetization per user leads to one unit of disutility for that user, or that the fee is simply a price) and $(\vec{\lambda} - \vec{\tau_2}) = \vec{0}$ (i.e. the platform has no marginal incentive to acquire users above their revenues from fees) this can be simplified further. Making these substitutions yields:

$$\vec{\Phi} = (I + B(B^{-1})')^{-1}\vec{P_0}. \tag{35}$$

Note that τ_1 drops out of the equation here – that is because the optimal price is not impacted by a marginal tax, so long as there is no net non-fee derived revenue/utility from acquiring users. Substituting in $BB^{-1\prime} = I + (B - B')B'^{-1}$ yields:

$$\vec{\Phi} = (2I + (B - B')B^{-1\prime})^{-1}\vec{P_0}. \tag{36}$$

Rearranging using (Petersen et al., 2008), equation (157) yields:

$$\vec{\Phi} = \Big(1/2I - 1/4(B'(B - B')^{-1} + 1/2I)^{-1}\Big)\vec{P_0}. \tag{37}$$

Bringing into the second term 2^{-1} yields:

$$\vec{\Phi} = \Big(1/2I - 1/2(2B'(B - B')^{-1} + I)^{-1}\Big)\vec{P_0}. \tag{38}$$

Plugging in $I = (B - B')^{-1}(B - B')$ for the latter I and rearranging yields:

$$\vec{\Phi} = \Big(1/2I - 1/2((B - B' + 2B')(B - B')^{-1})^{-1}\Big)\vec{P_0}. \tag{39}$$

Adding yields:

$$\vec{\Phi} = \Big(1/2I - 1/2((B + B')(B - B')^{-1})^{-1}\Big)\vec{P_0}. \tag{40}$$

[31] Noting that taking the inverse of a product of matrix changes the order of the multiplication, (Petersen et al., 2008), equation (1).

Resolving the inverse on the outer term, which reverses the order of inner terms, yields:

$$\vec{\Phi} = \Big(1/2I - 1/2(B - B')(B + B')^{-1}\Big)\vec{P}_0. \tag{41}$$

Finally, substituting back in $B = I - \sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}}$ yields:

$$\vec{\Phi} = \Big(1/2I - 1/2((\sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})' - (\sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}}))(2I - (\sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})' + (\sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}}))^{-1}\Big)\vec{P}_0, \tag{42}$$

which can be rearranged to yield:

$$\vec{\Phi} = \Big(1/2I - 1/4((\sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})' - (\sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}}))(I - \frac{(\sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})' + (\sum_{y=1}^{Y} \frac{\partial \vec{P}_t}{\partial \vec{P}_{t-y}})}{2})^{-1}\Big)\vec{P}_0. \tag{43}$$

A special case of this solution, for the $Y = 0$ case, can be found as the optimal pricing strategy in Jeon et al. (2021).

B Additional Derivation Details on Consumer Surplus

In general, we have:

$$CU_i = \int_0^{P_i} \mu_i - \epsilon_i(\mathrm{i}) d_\iota. \tag{44}$$

Here is how to derive equation (23) from that point. First, note that, by taking the linear approximation for demand, we can write consumer surplus as the following triangular area:

$$CU_i = \frac{1}{2}(\mu_i - \epsilon_i(0))P_i, \tag{45}$$

where P_i is the base of the triangle, and $\mu_i - \epsilon_i(0)$ – i.e., the consumer surplus for the lowest opportunity cost person – is the height of the triangle.

Noting that the linearized demand equation is:

$$P_i = \frac{\partial P_i}{\partial \mu_i}\mu_i + P_0, \tag{46}$$

solving for μ_i yields:

$$\mu_i = (P - P_0)(\frac{\partial P_i}{\partial \mu_i})^{-1}, \tag{47}$$

and solving for P_0, or the amount of people who would participate if $\mu_i = 0$ yields:

$$P_0 = -\epsilon_i(0)\frac{\partial P_i}{\partial \mu_i}. \tag{48}$$

Plugging both of these into the above yields equation (23).

The change in consumer welfare with respect to a change in Φ_i can also be derived by taking the derivative of equation (46) with respect to $\vec{\Phi}_i$ which yields:

$$\frac{\partial CU_i}{\partial \Phi_i} = \frac{1}{2}\frac{\partial \mu_{ss,i}}{\partial \Phi_i}P_i + \frac{1}{2}(\mu_i - \epsilon_i(0))\frac{\partial P_{ss,i}}{\partial \Phi_i}, \tag{49}$$

which by making the same substitutions and noting $\frac{\partial P_{ss,i}}{\partial \Phi_i} = \frac{\partial P_i}{\partial \mu_i}\frac{\partial \mu_{ss,i}}{\partial \Phi_i}$ can be plugged in to yield equation (24).

C Additional Simulation Model Details

Google Surveys provides information on survey participants' gender and age, so we distinguish market segments based on those characteristics. We divide Facebook users into twelve market segments. These are a pair of genders and six age brackets. The market segments we consider are:

- Gender: Male or Female, and
- Age: 18-24; 25-34; 35-44; 45-54; 55-64; and 65+.

Individuals under age 13 are not allowed to have Facebook accounts.

We asked the following sets of questions about individuals' demand for Facebook, combining responses within the twelve market segments described. The list of surveys conducted is documented in Table C1 and the full list of questions and possible responses in Appendix D.

Figure C1 gives examples of how the surveys appeared to respondents. Respondents answered these surveys either as part of Google Rewards or to access premium content on websites.

Survey	Number of Responses
Number of friends on Facebook	3,509
Composition of friends by demographic group	15,660
Willingness to accept to give up Facebook for one month	17,649
Willingness to accept to give up a friend group on Facebook for one month	13,356
Willingness to pay to not see any advertisements on Facebook for one month	7,021

Table C1: Surveys conducted and number of responses. More detail on the survey instruments is in Appendix D

.

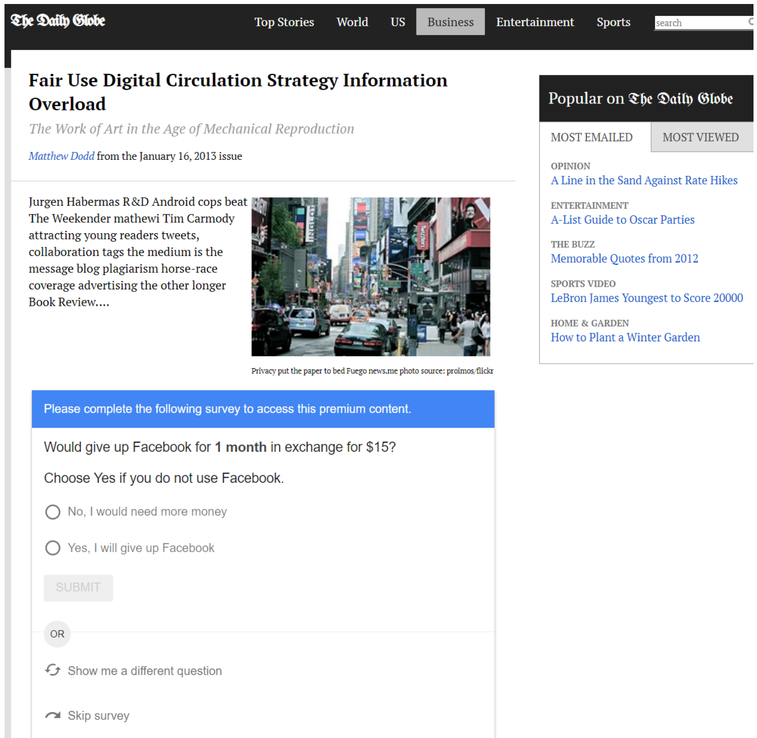

Figure C1: Google Surveys interface example. Note that each respondent only receives a single survey question.

We assume that the opportunity cost for using Facebook is distributed such that demand for Facebook, Ω_i, follows a logistic distribution. We estimate the parameters of Ω_i by running a logistic regression on responses to the question "Would you give up Facebook for 1 month in exchange for $X? Choose Yes if you do not use Facebook.". Regressions are separately estimated for each population group i. The logistic regression takes the form:

$$Y = \beta_0 + \beta_1 X, \tag{50}$$

where Y is an indicator for whether the offer is accepted, and X is the amount offered.

Figure C2 reports the willingness to accept (WTA) demand curve for giving up Facebook for one month combining together all demographic groups in our sample. The figure plots the mean response to this question for different offers, 95% confidence intervals, and the logistic line of best fit. The median WTA is $18.16. These findings are juxtaposed with those of Brynjolfsson et al. (2019), who asked a directly comparable question. Our results are broadly in line with that paper, but indicate slightly lower valuations.

Appendix Figures E2 through E13 plot mean WTA responses and demand curves for various subgroups of the population. Appendix Table E1 reports the estimates underlying these curves.

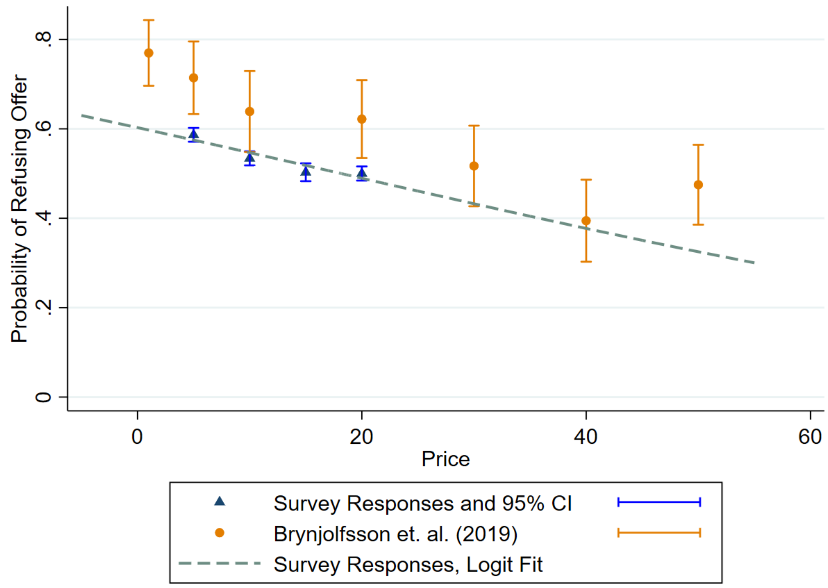

Figure C2: Probability of rejecting an offer to give up Facebook for one month for price listed. Mean responses to the question "Would you give up Facebook for 1 month in exchange for \$X? Choose Yes if you do not use Facebook." 95% confidence intervals are based on binomial statistics. The findings from our survey are plotted with blue triangles. Results from a directly analogous question from Brynjolfsson et al. (2019) are plotted with orange circles. The logistic line of best fit for the new results is plotted with a dashed line.

We convert from estimates of the CDF logistic equation to the PDF of the distribution of ϵ_i's using the equation:

$$p(\epsilon_i) \sim \frac{e^{-\frac{\epsilon_i - \eta_i}{s_i}}}{s_i \left(1 + e^{-\frac{\epsilon_i - \eta_i}{s_i}}\right)^2}, \tag{51}$$

where

$$s_i = (\beta_{1,i})^{-1} \tag{52}$$

and

$$\eta_i = (-\beta_{0,i}) s_i. \tag{53}$$

Ω_i, the probability of participating on a platform as a function of the fee and all other's participation, is a function of both the distribution of i's opportunity costs and how the value of Facebook participation changes as these fluctuate. The parametric model of consumer utility we calibrate for each market segment i is linear in the number of friends of each type j and in disutility from advertising:

$$\mu_i = \sum^{J} U_i(j) p_j z_i(j) D_j - a_i \phi_i, \tag{54}$$

where $U_i(j)$ is the utility an individual i receives from having a friend in market segment j (i.e. $\frac{\partial \mu_i}{\partial P_j}$), P_j is the percentage of Americans in group j who use Facebook, $z_i(j)$ is the percentage of users of type j who i is friends with, D_j is the population of demographic group j, and a_i is the disutility caused by a level of advertising ϕ_i (i.e., $\frac{\partial \mu_i}{\partial \phi_i}$). In terms of the microfoundation already presented, all we are doing is decomposing $\vec{P}$ into the total eligible population and the percentage who decide to participate $D_j p_j$, and decomposing B into the value of bilateral network effects per friend of a type $U_i(j)$ and the share of alters of a type who are friends $z_i(j)$.

We estimate the parameters of equation (54) through a combination of survey questions, government sources and information publicly available through Facebook's ad API and quarterly reports. D_j is taken from US Census reports for 2019. Our estimate of the revenue that Facebook make from users by demographic begins by noting that Facebook raises $11.62 dollars a month in revenue from US users through displaying them advertisements.[32] To understand how this varies across demographic groups, we gathered data on the cost of advertising to users of different types from Facebook's Ad API. After selecting which demographic group to target, Facebook's Ad API reports a range of how many impressions you are estimated to receive per dollar of spending. We take the inverse of this measure as the relative value of a demographic to Facebook's ad revenue (when a range is provided, we use the midpoint). By taking as given that the average value of a user per month is $11.62, we can then calculate the revenue per user of a demographic using the following equations:

$$\overline{\phi_i} = z \text{Relative Value}_i \tag{55}$$

and

$$11.62 = q \frac{\sum^I \text{Relative Value}_i \overline{P_i} D_i}{\sum^I \overline{P_i} D_i}, \tag{56}$$

where q is a scaling term, $\overline{p_i}$ is the estimate of the initial participation rate on Facebook by the demographic group, and D_i is the total population of the group in the US.

To estimate the share of users by type that a user of demographic group i is friends with, we combine the results of two sets of survey questions. We ask questions to solicit the average total number of friends by ego and alter demographic. We then ask questions to solicit what percentage of their friends are of each demographic. We re-balance these responses to sum to 100 percent (including a catchall category for

[32] This is derived from Facebook's 2019 Q1 annual report, where they report $ 34.86 in revenues per North American user per quarter.

individuals under age 18, who are not directly modeled). Appendix figure E1 presents our estimate of the average number of friends by type for each demographic.[33]

To estimate the value of friends by demographic group, we begin by asking users "On Facebook, would you unfriend all your friends who are [gender] between ages [age bracket] for $X? Choose Yes if you do not use Facebook." We then re-scale these responses by the estimated number of friends each demographic group has, and our estimate of initial average welfare from Facebook (derived from our estimates of Ω_i) so that the sum of all friend network effects is equal to our estimate of the average initial utility per user from the platform. Finally, to estimate the disutility from advertising a_i, we ask "What is the maximum amount of money (in US $) you would pay to personally not see any advertisements on Facebook for 1 month? Select 0 if you do not use Facebook." We divide this number by our estimates of initial revenues per user $\overline{\phi_i}$ to estimate a_i.[34]

We calculate the impact of a change in advertising strategy, or some other change in Facebook's environment, over the course of multiple cascades. We denote the period when the platform changes its advertising level as $t = 1$. The participation rate on the platform for a demographic group after cascade t is:

$$P_{i,t} = \Omega_i(\sum^{J} U_i(j)z_i(j)D_j p_{j,t-1} - a_i\phi_{i,t}), \tag{57}$$

where $P_{i,0} = \overline{P_i}$, the initial rate of platform participation for the market segment.

We calculate the perceived welfare to a user of a demographic group i from the existence of Facebook after cascade t as:

$$\int_0^{P_{i,t}} \big((\mu_i(\vec{p}_{t-1}, \phi_i) - e_i(\rho_i))\big) d\rho_i, \tag{58}$$

where e_i is the inverse of Ω_i, giving the implied opportunity cost of Facebook use for every percentile of the population:

$$e_i = -s_i log(\frac{1-p_i}{p_i}) + \mu_i. \tag{59}$$

[33] We reached out to Facebook to collaborate on calibrating our model, including by running internal experiments to estimate parameters. We presented our work in progress to the Facebook core data science team. While Facebook chose not to help us to the extent requested, they did release to us information on friendship shares by ego and alter demographic as part of their 2020 Social Cohesion Conference. Appendix Figure E14, contrasts our survey-based estimates of these rates against official measures. We look forward to revising our estimates if additional internal Facebook data becomes available.

[34] This parameter, a user's direct marginal disutility from experiencing a certain dollar amount in revenue of advertising, is one with very few estimates in the literature. The only alternate estimate we are aware of appears in (Kawaguchi et al., 2021) which finds the disutility caused by advertising on Google Play in Japan is about 6% of total revenues. Our equivalent estimate is approximately 20% overall, with significant heterogeneity across groups.

The total welfare to a demographic group from the existence of Facebook is the above amount times the number of users of that demographic group.

We calculate 1000 cascades of the network effect, but the steady-state is approximately reached after the first few cascades. The revenue to Facebook from user participation of a given demographic after t cascades is

$$\Phi_{i,t} = \phi_{i,t} D_i p_{i,t}. \tag{60}$$

D Survey Instruments

We conducted 57,195 surveys on Google Surveys to collect our data and quantify our parameters of interest. Each survey has several variations and each respondent answers only one variation of a survey.

The instrument for each survey is listed below, organized by question type.

D.1 Number of Friends on Facebook (n = 3,509)

Question text: How many friends do you have on Facebook?

Survey variations: We conducted two different surveys to get more coverage of users with very low and very high number of friends

Possible responses:

- Survey 1 (n = 2,507): 0-100; 100-200; 200-300; 300-400; 400-500; More than 500; I do not use Facebook
- Survey 2 (n = 1,002): 0-50; 50-100; 100-500; 500-700; 700-900; More than 900; I do not use Facebook

D.2 Composition of Friends by Demographic Group (n = 15,660)

Question text: What percentage of your friends on Facebook are *[demographic group]*?

Demographic groups:

- under age 18 (n = 1,200)
- men between age 18 and 24 (n = 1,207)
- men between age 25 and 34 (n = 1,206)
- men between age 35 and 44 (n = 1,203)

- men between age 45 and 54 (n = 1,201)
- men between age 55 and 64 (n = 1,208)
- men aged 65 or over (n = 1,203)
- women between age 18 and 24 (n = 1,206)
- women between age 25 and 34 (n = 1,204)
- women between age 35 and 44 (n = 1,207)
- women between age 45 and 54 (n = 1,201)
- women between age 55 and 64 (n = 1,208)
- women aged 65 or over (n = 1,206)

Possible responses: 0-10%; 10%-20%; 20%-40%; 40%-60%; 60%-80%; 80%-100%; I do not use Facebook

D.3 Willingness to Accept (WTA) Money to Give up Facebook for One Month (n = 17,649)

Question text: Would give up Facebook for 1 month in exchange for $[X]? Choose Yes if you do not use Facebook.

Price levels: [X] = 5 (n = 4,917), 10 (n = 4,912), 15 (n = 2,917), 20 (n = 4,903)

Possible responses:

- Yes, I will give up Facebook
- No, I would need more money

D.4 Willingness to Accept (WTA) Money to Give up a Friend Group on Facebook for One Month (n = 13,356)

Question text: On Facebook, would you unfriend all your friends who are [demographic group] for 1 month in exchange for $[X]? Choose Yes if you do not use Facebook.

Demographic groups:

- men between age 18 and 24 (n = 1,114)
- men between age 25 and 34 (n = 1,110)
- men between age 35 and 44 (n = 1,113)

- men between age 45 and 54 (n = 1,109)
- men between age 55 and 64 (n = 1,116)
- men aged 65 or over (n = 1,110)
- women between age 18 and 24 (n = 1,115)
- women between age 25 and 34 (n = 1,115)
- women between age 35 and 44 (n = 1,111)
- women between age 45 and 54 (n = 1,120)
- women between age 55 and 64 (n = 1,112)
- women aged 65 or over (n = 1,111)

Price levels: [X] = 5 (n = 3,655), 10 (n = 3,636), 15 (n = 2,431), 20 (n = 3,634)

Possible responses:

- Yes, I will unfriend all these friends
- No, I would need more money

D.5 Willingness to Pay (WTP) to Not See Any Advertisements on Facebook for One Month (n = 7,021)

Question text: What is the maximum amount of money (in US $) you would pay to personally not see any advertisements on Facebook for 1 month? Select 0 if you do not use Facebook.

Survey variations: We conducted two different surveys to get more coverage of users with a low WTP to not see any ads.

Possible responses:

- Survey 1 (n = 1,000): 0; $1-$5; $5-$10; $10-$15; $15-$20; More than $20
- Survey 2 (n = 6,021): 0; $1-$3; $3-$5; $5-$10; $10-$15; $15-$20; More than $20

E Additional Tables and Figures

Intercept	Coefficient on Cost	Demo Group
.576306	.0405408	Female 18-24
.578603	.0130975	Female 25-34
.753299	.0111212	Female 35-44
.856898	.0243794	Female 45-54
.798354	.0185573	Female 55-64
.570240	.0221589	Female 65+
.260814	.0248347	Male 18-24
.270967	.0231238	Male 25-34
.378010	.0282876	Male 35-44
.319277	.0294697	Male 45-54
-.03944	.0230791	Male 55-64
-.01939	.0161522	Male 65+

Table E1: Coefficient estimates from a logit regression of willingness to stop using Facebook on cost of Facebook proposed (equal to negative of the Price offered to stop using Facebook).

	Current (millions monthly)	First Best
Net Ad Revenue	$1,790.8	-381.7%
Consumer Surplus	$12,219.8	23.9%
Social Welfare (No SV)	$14,010.6	-27.9%
Social Welfare (With SV)		9.6%
Number of Users	154.1	16.5%

Table E2: Current and % change in Facebook advertisement revenue, consumer surplus, social surplus and number of users after Facebook is nationalized and the first best social welfare maximizing policy is implemented. Baseline social welfare is the sum of Facebook ad revenues and consumer surplus. The percentage increase in social welfare 'with SV' includes the non-monetary 'shadow value' of maintaining a larger user base in calculating the change in social welfare, while 'no SV' excludes this value (i.e. $\omega = 1$ or 0, respectively).

	Female 18-24	Female 25-34	Female 35-44	Female 45-54	Female 55-64	Female 65+	Male 18-24	Male 25-34	Male 35-44	Male 45-54	Male 55-64	Male 65+
Female 18-24	46.9	43.4	20.0	17.9	14.7	16.7	59.0	46.4	37.8	28.0	23.2	16.8
Female 25-34	36.9	64.2	34.9	18.3	19.7	15.3	40.2	71.0	47.2	43.2	24.0	21.5
Female 35-44	16.0	25.0	31.3	18.6	15.1	10.6	22.3	50.8	52.6	41.1	25.0	19.2
Female 45-54	14.7	21.7	23.6	25.4	23.7	10.7	15.4	42.3	38.7	34.0	29.1	14.6
Female 55-64	10.3	12.2	17.4	20.6	14.9	11.9	18.2	22.8	29.3	34.4	37.2	21.5
Female 65+	6.1	8.3	13.0	11.1	16.6	11.0	9.5	10.5	18.9	20.8	20.1	16.7
Male 18-24	32.3	44.9	27.3	19.1	15.1	22.5	45.1	32.6	25.9	26.6	26.7	17.4
Male 25-34	37.8	51.6	31.0	31.0	21.4	14.7	46.7	46.1	41.0	12.6	25.9	23.8
Male 35-44	27.5	37.4	39.0	31.4	22.9	19.3	31.9	46.3	46.2	33.3	39.1	9.6
Male 45-54	17.2	22.1	27.4	22.1	20.6	13.4	18.3	20.7	32.0	30.3	20.9	16.5
Male 55-64	9.3	18.5	20.0	27.2	19.5	17.2	14.1	21.3	20.4	27.1	22.8	19.9
Male 65+	9.6	10.4	7.0	20.2	17.7	14.1	12.4	12.7	14.1	17.8	20.6	12.3

Figure E1: Average number of friends someone in Y-axis market segment has of the type in the X-axis market segment.

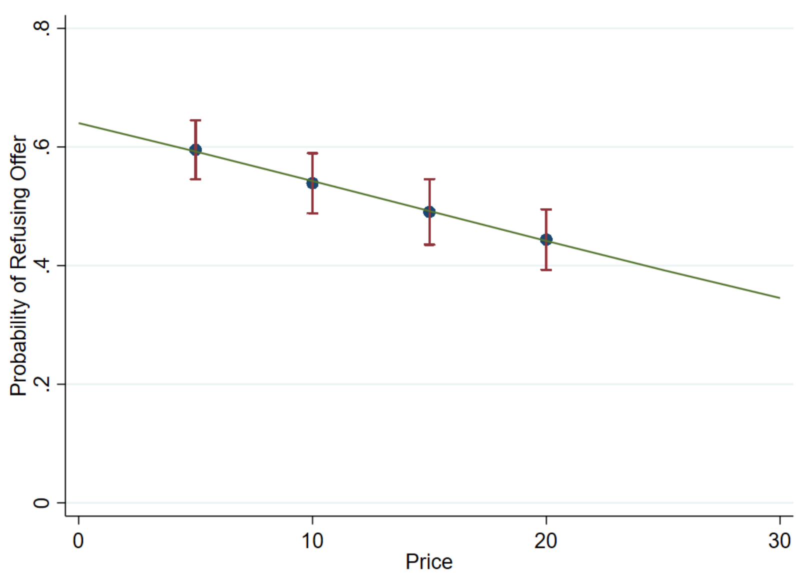

Figure E2: Underlying data and estimate of the the demand curve (Ω_i) for women age 18-24. The points are the mean response to the question "Would you give up Facebook for 1 month in exchange for $X? Choose Yes if you do not use Facebook." for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.

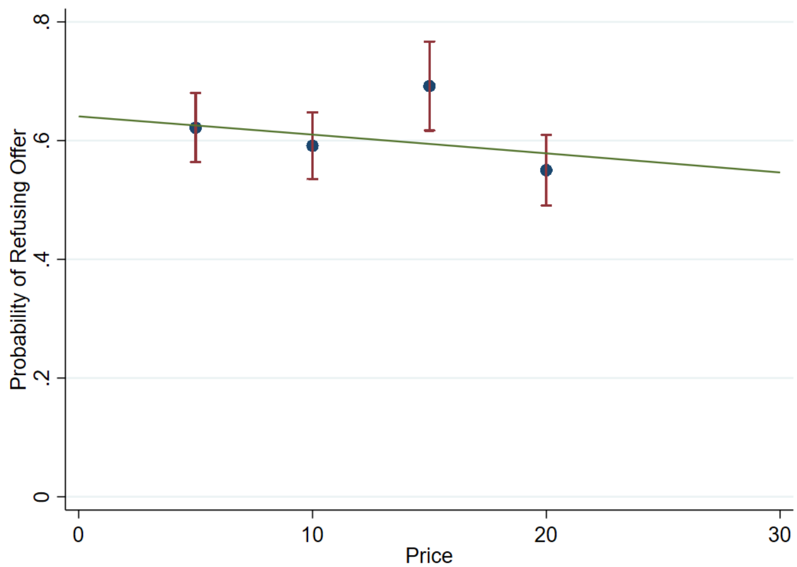

Figure E3: Underlying data and estimate of the the demand curve (Ω_i) for women age 25-34. The points are the mean response to the question "Would you give up Facebook for 1 month in exchange for \$X? Choose Yes if you do not use Facebook." for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.

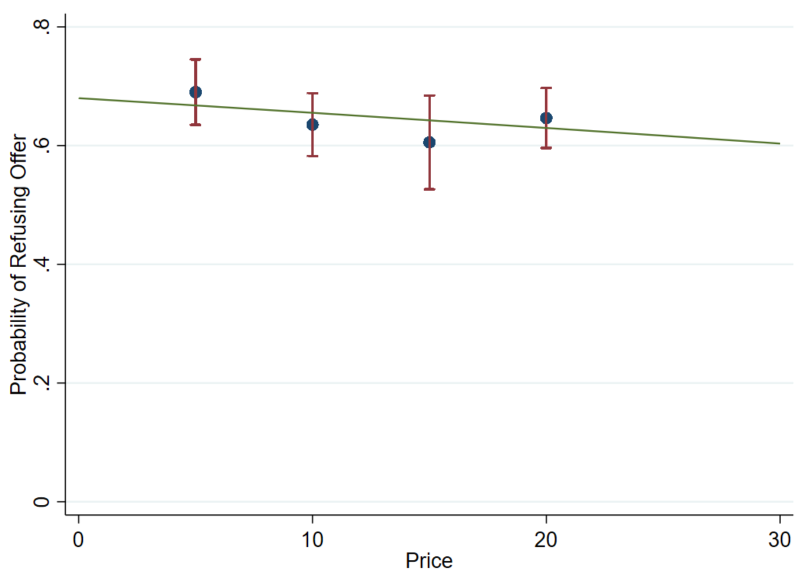

Figure E4: Underlying data and estimate of the the demand curve (Ω_i) for women age 35-44. The points are the mean response to the question "Would you give up Facebook for 1 month in exchange for \$X? Choose Yes if you do not use Facebook." for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.

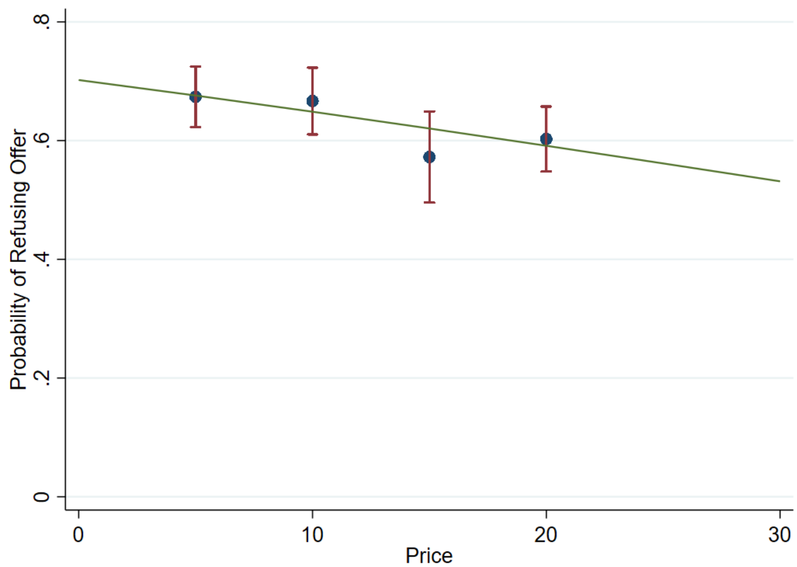

Figure E5: Underlying data and estimate of the the demand curve (Ω_i) for women age 45-54. The points are the mean response to the question "Would you give up Facebook for 1 month in exchange for \$X? Choose Yes if you do not use Facebook." for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.

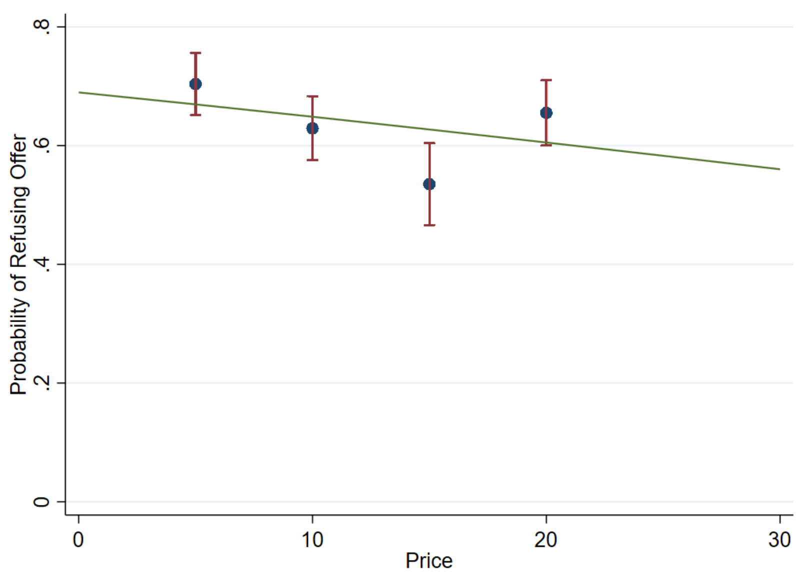

Figure E6: Underlying data and estimate of the the demand curve (Ω_i) for women age 55-64. The points are the mean response to the question "Would you give up Facebook for 1 month in exchange for \$X? Choose Yes if you do not use Facebook." for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.

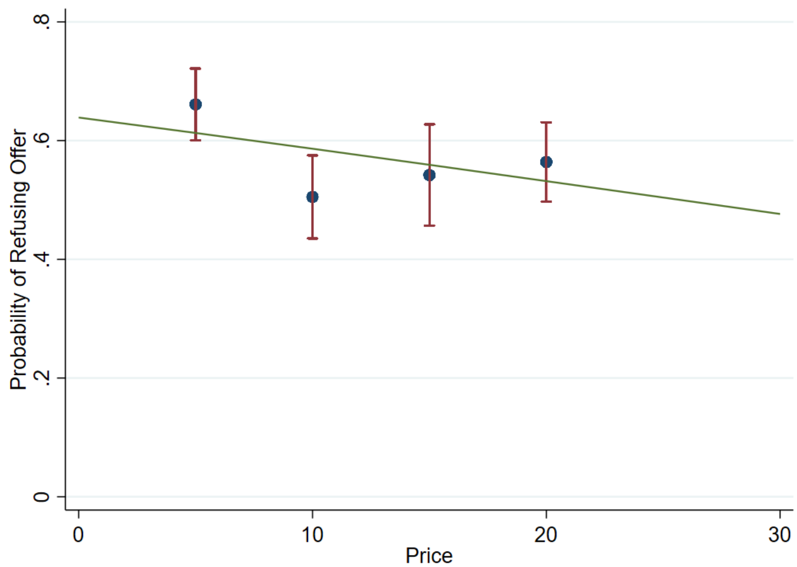

Figure E7: Underlying data and estimate of the the demand curve (Ω_i) for women age 65 or older. The points are the mean response to the question "Would you give up Facebook for 1 month in exchange for $X? Choose Yes if you do not use Facebook." for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.

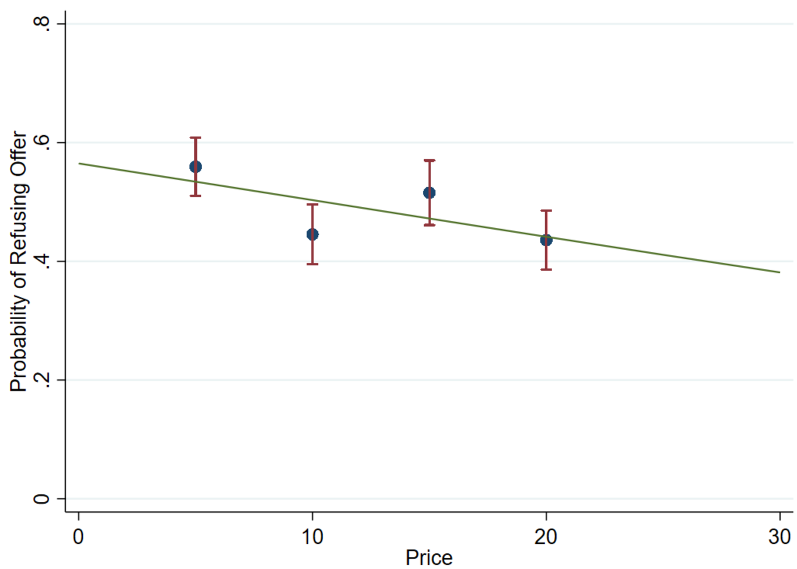

Figure E8: Underlying data and estimate of the the demand curve (Ω_i) for men age 18-24. The points are the mean response to the question "Would you give up Facebook for 1 month in exchange for $X? Choose Yes if you do not use Facebook." for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.

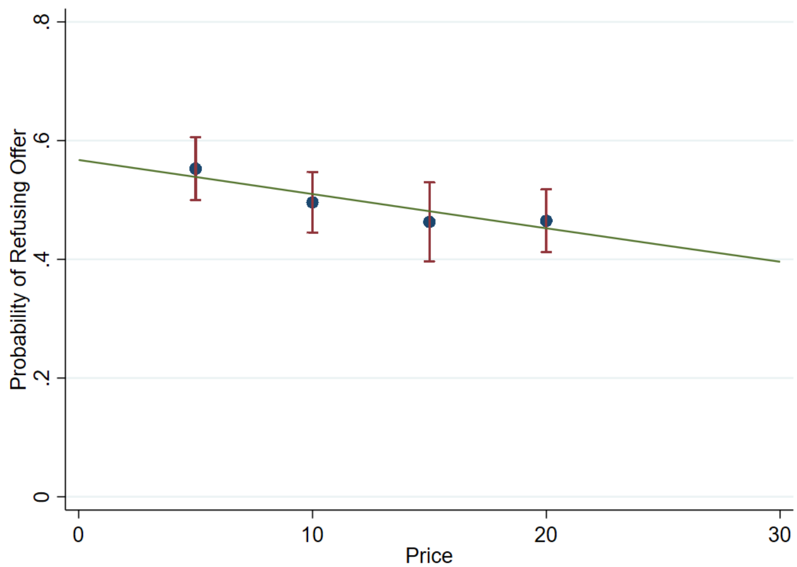

Figure E9: Underlying data and estimate of the the demand curve (Ω_i) for men age 25-34. The points are the mean response to the question "Would you give up Facebook for 1 month in exchange for $X? Choose Yes if you do not use Facebook." for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.

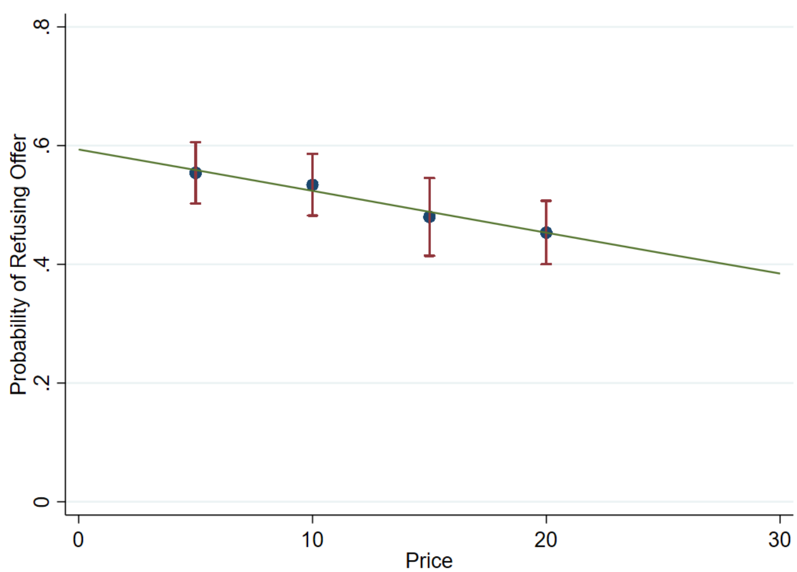

Figure E10: Underlying data and estimate of the the demand curve (Ω_i) for men age 35-44. The points are the mean response to the question "Would you give up Facebook for 1 month in exchange for $X? Choose Yes if you do not use Facebook." for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.

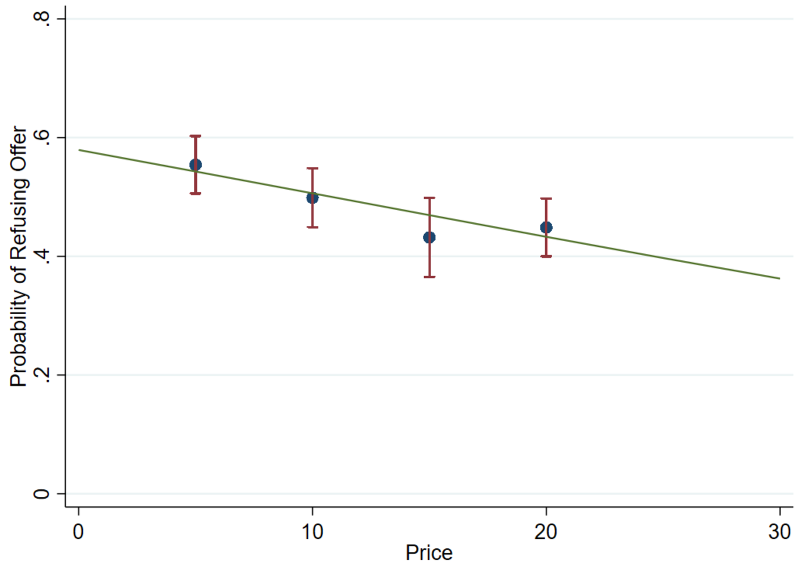

Figure E11: Underlying data and estimate of the the demand curve (Ω_i) for men age 45-54. The points are the mean response to the question "Would you give up Facebook for 1 month in exchange for $X? Choose Yes if you do not use Facebook." for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.

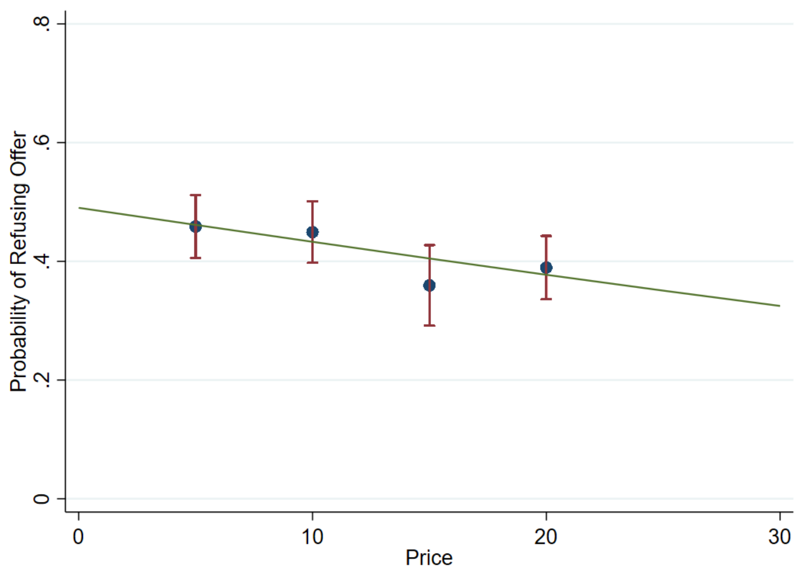

Figure E12: Underlying data and estimate of the the demand curve (Ω_i) for men age 55-64. The points are the mean response to the question "Would you give up Facebook for 1 month in exchange for $X? Choose Yes if you do not use Facebook." for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.

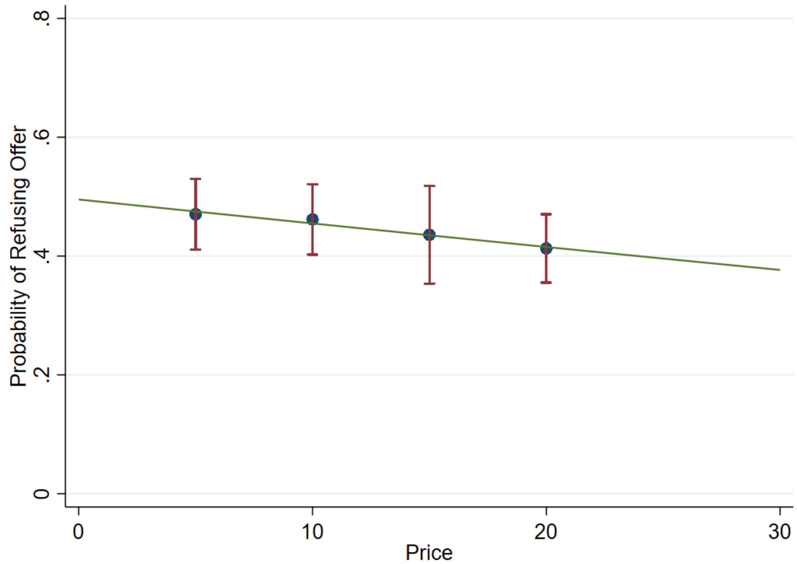

Figure E13: Underlying data and estimate of the the demand curve (Ω_i) for men age 65 or older. The points are the mean response to the question "Would you give up Facebook for 1 month in exchange for $X? Choose Yes if you do not use Facebook." for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.

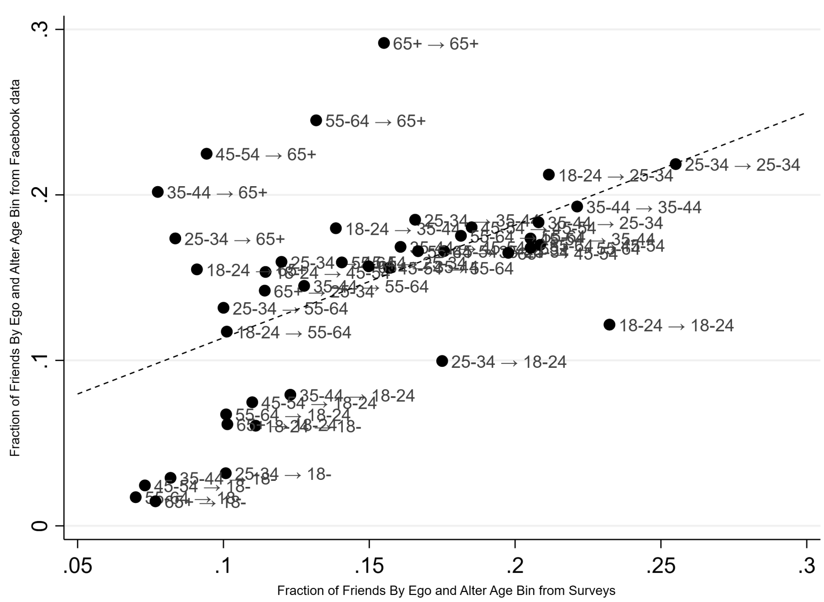

Figure E14: Comparison of fraction of ego's friends by ego and alter age bins. X axis gives our estimates, and Y axis those from Facebook's social cohesion data on connections within county. Pearson's correlation between both probabilities is 0.53.

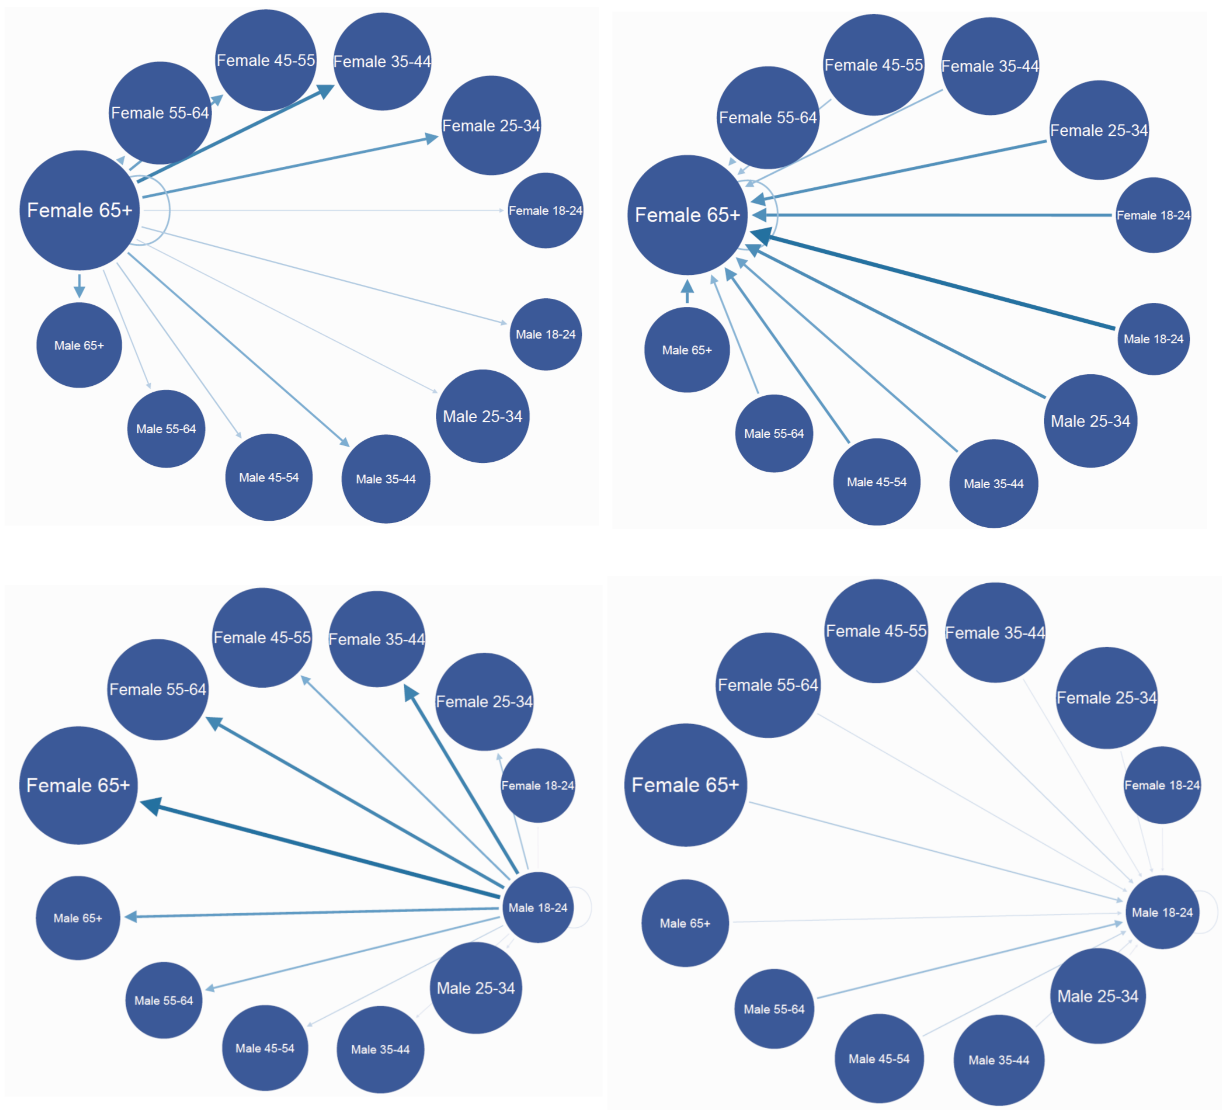

Figure E15: Facebook usage and network externalities by market segment. The size of each node represents the relative current size of the Facebook user base by demographic. The thickness of the arrows corresponds to the relative value received by a Facebook user of the demographic the arrow is pointing towards from a friendship with a user of the source demographic (i.e., $U_i(j)$ with i being where the arrow is pointing towards, and j being the source of the arrow). Only some friend valuations displayed. Clockwise from the top left, the friend values displayed are the value from women aged 65+ to Facebook users of different demographics, to women aged 65+, to men aged 18-24, and from men aged 18-24.

.