

Multi-Sided Platform Strategy, Taxation, and Regulation: A Quantitative Model and Application to Facebook

Seth G. Benzell*, Avinash Collis^{†‡}

This draft: July, 2020
First draft: July, 2019

Abstract

Digital platforms, such as Facebook, Uber, and AirBnB, create value by connecting users, vendors, and contractors. Their supply and demand economies of scale make them natural monopolies, and have led to increasing calls for special regulations and taxes. We construct and illustrate an approach for modeling digital platforms. The model allows for heterogeneity in demand elasticity, disutility from advertising, and network effects across users. We parameterize our model using a survey of over 57,000 US internet users on their demand for Facebook. Facebook creates \$14 billion in social value per month, with consumer surplus concentrated among female and older users of Facebook. The most valuable friendships on Facebook are worth in excess of 50 cents per month, but most are worth far less, with connections to men more valuable on a per-friend basis. We find Facebook has too low a level of advertising relative to their short-term revenue maximizing strategy, confirming that they also value maintaining a large user base. We estimate that the welfare lost from Facebook's monopoly power, compared to the first best of a social welfare maximizing Facebook, is 9.6% of current social value, or \$1.3 Billion per month. We then simulate six proposed policies for government management of digital platforms, taking Facebook's optimal response into account. Taxes are mostly incident on Facebook profit and properly targeted taxes can even raise consumer surplus. Achieving perfect competition in social media would raise social surplus from Facebook by 4.8% of current value. But a botched regulation that left the US with two smaller, non-competitive social media monopolies would decrease social surplus by 84.7%.

*Argyros School of Business and Economics, Chapman University & MIT Initiative on the Digital Economy, sbenzell@mit.edu

[†]McCombs School of Business, University of Texas at Austin & MIT Initiative on the Digital Economy, avinash.collis@mcombs.utexas.edu

[‡]We thank Sinan Aral, 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 Chris Forman for a fantastic discussion. We additionally thank the MIT Initiative on the Digital Economy for research funding. Our study was determined to be exempt by MIT's IRB.

1 Introduction

Much of the value of digital platform businesses comes from “network effects”. A network effect is an externality that one participant in a market, digital platform, or similar system provides to others. But how exactly can one measure and exploit the value of network effects for any particular business or industry? In this paper we propose and implement a flexible strategy for the measurement and optimal harnessing of network effects. We then use the model to simulate the effect of several proposed or recently implemented digital platform regulations and taxes.

We make three main contributions. First, we provide a tractable framework for optimal pricing strategy on multi-sided platforms. This approach builds on traditional price discrimination models by taking into account network externalities. Second, we implement this model using data we collected on Facebook, introducing a novel methodology for the estimation of network effects. Using the calibrated model, we provide the first simulations of Facebook revenue, participation, and social surplus under counter-factual pricing policies. Finally, we use the model and data to estimate the social gains from proposals by politicians and academics to tax and regulate “Big Tech.”

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 value they get from other types of users using the platform, and the disutility they receive from advertising. It is a model of an n-sided platform in the sense that each individual or market segment can be thought of as a node of the network.¹ We show that the optimal pricing strategy for an n-sided platform entails decreasing fees or advertising for users who elastically demand the platform (the direct effect) and who create high amounts of network value for other profitable users who themselves demand the platform elastically (the network effect).

After introducing and analyzing our model, we proceed to an empirical illustration. We collected information on US internet users’ demand for Facebook across over 57,000 surveys conducted through Google Surveys. We categorize the surveyed into twelve demographic groups by their age and gender. To collect information on demand for and network effects on the platform, we use an experimental choice approach in the spirit of (Brynjolfsson et al., 2019) and (Allcott et al., 2019). These papers measure the consumer surplus generated by digital goods by conducting discrete choice experiments

¹When conceived of this way, any platform, including a one-sided platform, can be thought of as an n-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 vs. personal use, demographics, regional location, heterogeneity in activity (frequent users or not) and type of activity (always callers, callers and receivers, always receivers).

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 social connections) as well as by using information from the full distribution of responses to fit a demand curve (in our parameterization, logistic) rather than focusing on median and average responses. We adapt this approach to our case by giving consumers the choice to give up access to a subset of their network in exchange for monetary compensation.

Using this information about demand for Facebook, we estimate the parameters of a logistic demand curve for each demographic group, 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 current advertising revenues by demographic group.

With this model of individual participation, we can calculate the effects of counterfactual pricing policies, government policies and demand shocks. We begin by simulating Facebook’s revenue maximizing strategy. We find that Facebook could raise revenues by 2.38 billion dollars 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%. We infer that in addition to maximizing current revenues, Facebook values maintaining a large and happy 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 proceed to calculating the impact of changes in government policy on Facebook revenues, participation, and consumer welfare. We estimate that a first-best social welfare maximizing Facebook would subsidize usage, running a large deficit, and raising social surplus by 9.6%, or \$1.3 Billion a month. This result assumes that Facebook’s motivation to maintain a large user base is a publicly useful one (e.g. it represents the non-immediately monetizable value of data collected) rather than an only privately useful one (i.e. as a barrier to entry), as well as other assumptions about the marginal ability of the social planner to attract users.

We simulate three taxation and redistributive policies. 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 advertising revenues would raise consumer

surplus by 1.3%, and raise 2.4% of current Facebook advertising revenues in taxes. A tax on the number of users of Facebook, which raised the same amount of revenue, would lower consumer surplus by -.1%. Another proposed policy for redistributing the wealth from Facebook is Weyl’s “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 (which create less than an a dollar of direct disutility per dollar of revenue) 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 taking steps to enhance the competitiveness of the social media industry, by 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). We model this policy as creating perfect competition, and lowering the price of the platform to its marginal cost – i.e. forcing the elimination of advertising and other fees. Next we simulate the results two a ‘botched’ Facebook breakups. The first, a botched horizontal breakup, leaves America with two monopolies over half of the population each. The second, a botched vertical breakup, ends Facebook usage by 5% of the population with no offsetting increase in social media competitiveness. We predict that perfect competition would raise social welfare by 4.8%. Breaking Facebook into two non-competitive ‘baby Facebooks’ would be disastrous, lowering social welfare by 84.7%. A botched vertical breakup would have more limited effects, lowering social welfare by only about 10.1%.

2 Related Literature

Following the seminal work of Parker and Van Alstyne (2005) and Rochet and Tirole (2003), 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 for the other side (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 heterogeneous 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)).

There are several recent papers which model pricing in the presence of multi-dimensional network effects. For example, Bernstein and Winter (2012) determines a mechanism for optimally renting storefronts in a shopping mall where stores have heterogeneous externalities on other stores. Candogan et al. (2012) and Fainmesser and Galeotti (2015) 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 an algorithm guaranteeing 88% of the maximum. Fainmesser and Galeotti (2015) considers a similar model but assumes 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. The paper in this literature with a model most similar to ours is Weyl (2010). That paper, like ours, considers an indivisible platform good with network effects. It also, like this paper, allows for groups to vary in both their network effect on other groups and in their opportunity cost for using the platform. It finds that a wedge exists between the profit maximizing and social welfare maximizing pricing strategy.²

Our paper builds on these prior papers along several 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) which do not allow for such variation, and Weyl (2010) which features an unrealistic pricing scheme, where users are charged based on the level of participation of other users (i.e. an ‘insulating

²The exact nature of this wedge – as a marginal, not an average distortion – was clarified in a published comment (Tan and Wright, 2018).

tariff’). Weyl (2010) 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 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.³ The most important contribution of our model is that it is the first one 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 Facebook, a platform primarily monetized through advertising. 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., 2018a)). 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. (2018b) 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 modeled and users generating different amounts of network externalities end up seeing the same number of ads.⁴ 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). That paper has a model of an analog two-sided platform: the yellow pages. It uses instruments to find the spillover effects of additional advertisements on phone-book quality. Rysman that small *decreases* in competition might increase welfare, as there would be fewer better phone-books with more utilituous advertisements.

Our paper also speaks to the growing literature on the optimal regulation of digital platform monopolies. [Cite Fiona and Marshall, cite that spanish project]

³The 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. (2011) 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.

⁴Based 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.

3 A Model of an N-Sided Platform

The foundational element of a model of network effects is a stance on how agents connect to and gain welfare from the network. In our model, individuals with heterogeneous characteristics decide whether or not to participate on a platform. Their desire to participate in the network is a function of their expectation of which other individuals will participate. For example, Jane Doe’s desire to use Instagram is a function of which of her friends are also using Instagram. The key term in the model is the externality that users gain from others. Unlike other models of platforms, we allow for individuals of different characteristics to gain different amounts of value from the participation of others on the network. These market segments are the different sides of the platform.

Our focal example is a social network, because our illustration takes place in that setting. Therefore, in our baseline model, other incidental network characteristics mimic that of a social network. Once two users are using the network, there is no additional cost for them to form a connection. All connections where both users gain weakly positive value are immediately formed.⁵

The platform’s monetization is modeled as a unpleasant platform attribute (e.g. advertising), fee or subsidy faced by each participating user: A binary function of their decision to participate on the network.⁶ Users face disutility depending on how intensely they are monetized by the platform. The relationship between the platform’s degree of monetization and the user’s disutility need not be one-to-one. For example, the platform may raise revenues from users through unpleasant advertisements or by collecting personal data (potentially causing users the disutility of knowing one’s data will be harvested and resold). Alternatively, it may correspond to an explicit participation charge, such as WhatsApp’s original \$1 subscription cost.

⁵This abstracts from the reality that Facebook usage is not a binary choice, but rather a continuing decision about how often to post and view other’s posts. A model with an additional decision about how intensely to use Facebook might have slightly weaker total network effects, if the most prolific posters within a demographic group (who presumably create the strongest local network effects), are the least likely to leave the platform as a result of a policy or participation shock. The opposite might hold if the elasticity of posting with respect to others’ participation was larger than the elasticity of participation. We thank Xiang Hui for this insight.

⁶While some papers have found advertising to provide benefits to platform users (e.g. Rysman (2004)), in the context of a digital platform it is safe to assume advertisements cause disutility at the margin – it is effectively costless to increase their supply, and so, if this were not the case, all social media platforms would simply be an infinite feed of advertisements.

3.1 Consumers

A consumer i chooses whether to participate in the platform ($\mathbf{P}_i = 1$) or not ($\mathbf{P}_i = 0$).⁷ If the consumer i uses the platform ($\mathbf{P}_i = 1$), they expect to receive

$$E[U_i(\mathbf{P}_i = 1)] = \mu_i(P_1, \dots, P_I, -\phi_i) \quad (1)$$

where P_j is the probability individual j participates on the platform. ϕ_i is the revenue the platform raises from individual i . A firm which monetizes using advertising might raise \$1 in revenue by displaying additional ads which create \$.20 in additional disutility (i.e. $\frac{\partial \mu_i}{\partial \phi_i} = .2$). Other platforms, like local telephone and pre-2016 WhatsApp, monetize by charging a flat fee for participation (i.e. $\frac{\partial \mu_i}{\partial \phi_i} = \1).

$\frac{\partial \mu_i}{\partial P_j}$ is the marginal utility given to i from j being on the network to i (if i participates). These partial derivatives capture our model’s network effects. In our theoretical analysis, our only assumption is that μ_i be continuously differentiable. In our calibration, we further assume that utility from the platform is linearly additive in the network effect from friends and disutility from ϕ . In other words, the parametric analysis assumes that $\frac{\partial \mu_i}{\partial P_j}$ ($U_i(j)$ for short) and $\frac{\partial \mu_i}{\partial \phi_i}$ (written as a_i) are constant.⁸

The value to a consumer of not using the platform, their ‘opportunity cost’, is an ex-ante unknown random variable.

$$U_i(\mathbf{P}_i = 0) = \epsilon_i \quad (2)$$

where ϵ_i are independent random variables (not necessarily symmetrical or mean 0). ϵ ’s distribution may vary by individual. This means that the probability of participating on a network, P , conditional on a given level of utility from the network good $U(\mathbf{P} = 1)$ is consumer specific.⁹

⁷Note that while demand functions are here defined at the individual level, as a practical matter firms may estimate them at the level of a demographic or social group. We consider an example with twelve market segments in our illustration.

⁸The assumption that the value of platform connections are linearly additive is not a harmless one, despite being made in all of the most similar papers extant ((Candogan et al., 2012), (Fainmesser and Galeotti, 2015), and (Weyl, 2010) all make this assumption). It means, for example, the additional value that Jane Doe gets from James Smith joining Instagram isn’t a function of whether any third person is already on Instagram. This is a useful simplification in the context of social networks, but it is likely unrealistic. Taking a food delivery platform as an example, it is likely the case that the 10th pizza delivery service joining the platform provides less platform value to the typical user than the 1st. A related simplification is the assumption that the value of a connection is only a function of the characteristics of the connected individuals. In general, the value of a connection to one individual may be a function of that individuals’ connections to other individuals. We abstract from these possibilities in the calibration. The measurement of non-linearly additive network effects introduces large measurement challenges beyond the scope of this paper’s illustration, but is something we plan to explore in future work.

⁹By adding a negative sign, this term can also be interpreted as the value or disutility of Facebook use in the absence of any friends or advertisements.

The distribution of ϵ_i determines how elastic i will be to changes in the platforms' attractiveness. Consider the case where ϵ_i is expected to be approximately equal to the utility of participation $U_i(\mathbf{P}_i = 1)$ – in other words, that it is likely that the user is 'on the fence' about using the platform. In this case, changes in ϕ_i or other consumers' participation will be highly likely to change i 's participation. On the other hand, if ϵ_i is two-peaked, with half of users miserable without the platform and half who are very happy without it, use of the platform will be inelastic to changes in platform quality.

Each consumer gets to see the resolution of their private outside option ϵ_i before participating, but not the resolution of anyone else's. Therefore, they base their decision to participate on the platform based on their beliefs in the likelihood of others participating. The *ex-post* consumer demand function is

$$\begin{cases} \mathbf{P}_i = 1 & \text{if } E[U_i(\mathbf{P}_i = 1)] > \epsilon_i \\ \mathbf{P}_i = 0 & \text{otherwise} \end{cases}$$

Note that P_i 's are independent because ϵ_i 's are independent.

We can write the *ex-ante* demand function (i.e. expected demand before ϵ_i is known) as:

$$P_i = \text{Prob}[E[(U_i(\mathbf{P}_i = 1))] > \epsilon_i] = \Omega_i(\mu_i) \quad (3)$$

for more useful notation, define

$$\mu_i \equiv \mathbf{U}_i \equiv E[U_i(\mathbf{P}_i = 1)] \quad (4)$$

The network is in equilibrium when individuals' decisions to participate are optimal responses to their beliefs about every other individuals' decision to participate. In our empirical illustration, we calculate the new equilibrium as a response to a shock through evaluating a series of 'cascades'.

For example, if the firm were to raise ϕ_i we would first calculate the direct impact of only this change in price on user i . This is the first cascade. We would then calculate all individuals' decision to participate taking i 's new participation rate as given – the second cascade. Additional cascades estimate every groups' rate of participation, taking the previous cascades' rate of participation as an input. We calculate 1000 cascades in all of our simulations, but as a practical matter, the importance of cascades beyond the third or fourth is minimal for a typical network in a stable equilibrium.

For the symmetric network (i.e. where all individuals have the same distribution of opportunity costs ϵ , disutility from advertising A , and network externality $\mu_i(P_j)$), where utility is linearly additive in the network effects and disutility from advertising, an equilibrium is stable so long as

$$1 > \frac{\partial \Omega}{\partial U} U(i)(I - 1) \quad (5)$$

where $U(i) = \frac{\partial \mu_i}{\partial P_i}$ is the value from any consumer participating in the network to any other consumer, and I is the number of friends each user has. Intuitively, the network is unstable when users are very elastic and care a lot about the participation of others on the network. When a network is in an unstable equilibrium, small changes in platform quality can lead to unravelling (i.e. the partial derivative of participation with respect to platform quality can be infinite). The derivation of this equation is in appendix B.

4 Optimal Platform Strategy

There are many questions you can ask about optimal platform strategy in this setting. Here we focus on the managerial implications for a profit maximizing social network.

4.1 Pricing Strategy

Consider a social network which can price discriminate among its users taking their demand functions (as well as the actions of competitors) as given. Platforms in this setting can price discriminate either by directly charging or subsidizing some users, or by giving some subset of users more or less advertisements.

Firms maximize expected total profits. After uncertainty is resolved, the firm's revenues are

$$\Phi = \sum_i^I \phi_i P_i - F \quad (6)$$

Where ϕ_i is the revenue collected from or distributed to consumer i if they participate in the network. It is a choice variable from the perspective of the firm. P_i is a binary indicator of whether the consumer participates. F is the fixed cost of the platform firms' operation.¹⁰

P_i 's are independent random variables, so firms maximize

$$E[\Phi] = \sum_i^I \phi_i P_i - F \quad (7)$$

where

$$P_i = E[P_i] = \Omega_i(U_i) = \Omega_i(\mu_i(\phi_1, \phi_2, \dots)) \quad (8)$$

¹⁰We assume the platform faces no marginal costs, but adding a marginal cost does not change the qualitative results.

the probability of a consumer participating P_i is an individual specific function of \mathbf{U}_i . Ω_i is the effective individual specific demand function. Ultimately the equilibrium level of participation is a function of preference parameters and the vector of ϕ 's, and there are no variable costs, so the monopolist social media platforms' problem is to select the level of ϕ 's that maximizes revenues.¹¹

The firm seeks to maximize profits

$$\max_{\phi_i} E[\Phi] = \sum_i^I \phi_i P_i - F \quad (9)$$

s.t.

$$P_i = \Omega_i(\mu_i) \quad (10)$$

This yields the following first order condition

$$\frac{\partial \Phi}{\partial \phi_i} = P_i + \phi_i \frac{\partial P_i}{\partial \phi_i} + \sum_{j \neq i}^J \phi_j \frac{\partial P_j}{\partial \phi_i} \quad (11)$$

where

$$\frac{\partial P_i}{\partial \phi_i} = \frac{\partial \Omega_i}{\partial \mu_i} \left(-\frac{\partial \mu_i}{\partial \phi_i} + \sum_j^J \left(\frac{\partial \mu_i}{\partial P_j} \frac{\partial P_j}{\partial \phi_i} \right) \right) \quad (12)$$

and,

$$\frac{\partial P_j}{\partial \phi_i} = \frac{\partial \Omega_j}{\partial \mu_j} \left(\frac{\partial \mu_j}{\partial P_i} \frac{\partial P_i}{\partial \phi_i} + \sum_{k \neq i}^K \frac{\partial \mu_j}{\partial P_k} \frac{\partial P_k}{\partial \phi_i} \right) \quad (13)$$

This recursion is natural as P_i is a function of P_j , which is a function of P_i , etc. Equation (13) will converge to a finite value so long as each recursion of the network effect “dampens out”. This will occur so long as the equilibrium is stable.

4.2 Profit Maximization Problem Simplified

Equation 11 gives conditions for the optimal schedule of fees (or other revenue raising monetization strategies) and subsidies for the general case. Even if not enough is known about the entire curve of functions to find the optimum, knowing the first derivative of the goal with respect to the choice parameters is useful. *An experimenting firm can simply use these equations to inch towards a local maximum via gradient decent.*

For simplicity in interpreting the first order condition, we retain only the first term in brackets in 12 and 13. In other words, the following equations only take into account only one cascade of network effects. In the empirical calibration below we will show

¹¹Although the function $a_i(\phi_i)$ which relates user disutility from monetization to platform revenue might be thought of as being net of this fixed cost.

that the first cascade of network effects provides a good approximation of the total effect of the shock.

$$\frac{\partial P_i}{\partial \phi_i} = \frac{\partial \Omega_i}{\partial \mu_i} \left(-\frac{\partial \mu_i}{\partial \phi_i} + \sum_{j \neq i}^J \left(\frac{\partial \mu_i}{\partial P_j} \frac{\partial P_j}{\partial \phi_i} \right) \right) \quad (14)$$

and,

$$\frac{\partial P_j}{\partial \phi_i} = \frac{\partial \Omega_j}{\partial \mu_j} \left(\frac{\partial \mu_j}{\partial P_i} \frac{\partial P_i}{\partial \phi_i} + \sum_{k \neq i}^K \left(\frac{\partial \mu_j}{\partial P_k} \frac{\partial P_k}{\partial \phi_i} \right) \right) \quad (15)$$

This, substituting into 11, yields a new simplified approximate first order condition

$$\frac{\partial \Phi}{\partial \phi_i} = \underbrace{P_i - \phi_i \frac{\partial \mu_i}{\partial \phi_i} \frac{\partial \Omega_i}{\partial \mu_i}}_{\text{Direct Effect}} - \underbrace{\frac{\partial \mu_i}{\partial \phi_i} \frac{\partial \Omega_i}{\partial \mu_i} \sum_{j \neq i}^J \phi_j \frac{\partial \Omega_j}{\partial \mu_j} \frac{\partial \mu_j}{\partial P_i}}_{\text{Network Effect}} \quad (16)$$

The simplified first order condition consists of two sets of terms. The first two terms report the direct effect of raising the amount of advertising on individual i by one dollar. This will raise revenue, based on that individual's current likelihood of participation, and lose revenue based on how elastic that individual's participation is. The two direct effect terms are what normal firms have to consider when pricing their products (note that when $\frac{\partial \Omega_j}{\partial P_i} = 0 \forall i, j$, i.e. when no network effects are present at the current margin, 16 reduces to this pair of terms).

The last term in equation 16 is the network effect of an advertising increase. The increase in advertising makes i less likely to participate (in this approximation, by an amount $\frac{\partial \mu_i}{\partial \phi_i} \frac{\partial \Omega_i}{\partial \mu_i}$) which leads others to stop participating (by an amount $\frac{\partial \Omega_j}{\partial \mu_j} \frac{\partial \mu_j}{\partial P_i}$). When these third parties stop participating, the platform loses on the current revenues that they were paying ϕ_j . In reality, most pairs of individuals will receive no network effect from the participation of another. In our empirical calibration we assume that $\frac{\partial \Omega_j}{\partial P_i} = 0$ whenever i and j are not Facebook friends.

In other words, 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 (second term) plus the decreased revenue from all the charged person's friends potentially dropping out (third term).

This simplified first order condition can be made more precise by taking into account additional cascades of the network effect. In other words, because user i 's fee increasing causes j to be less likely to participate, all those connected to j should be less likely to participate as well.

Unsurprisingly, the firms' profit maximizing strategy deviates from social welfare maximizing pricing. Appendix C reports the social welfare maximization problem. In-

tuitively, the wedge between the revenue and social welfare maximizing strategies arises from the fact that the platform only cares about monetization disutility and network spillovers that effect marginal users of the platform, whereas a social planner takes into account welfare changes for infra-marginal individuals who will use the platform in both scenarios.

5 Empirical Illustration – Facebook

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 demonstrated that many value it highly.

To illustrate how our method can be used by firms to price discriminate, we collected survey data to estimate our model. We conducted 57,195 surveys on a representative sample of US internet population. Google Surveys provides information on a survey participants’ gender and age group, so we distinguish market segments based on those characteristics. We divided 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
- Age: 18-24; 25-34; 35-44; 45-54; 55-64; and 65+

Individuals under the 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 1 and the full list of questions and possible responses is documented in Appendix section D.

Figure 1 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 1 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 1 month	7,021

Table 1: Surveys conducted and number of responses. More detail on the survey instruments can be found in appendix D



Figure 1: Google survey interface example. Note that each respondent only receives a single survey question.

5.1 Calibrating the Model for Facebook

The very general utility function analyzed in section 4 is tractable enough to lead to some analytic results, some of which we have already elucidated. However, for the purposes of quantitative estimates, we need a more restrictive functional form for the utility function. We also need to make modifications to the model to account for the fact that we are estimating it over market segments, not individuals, and for the fact that not all individuals are friends.

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

$$Y = \beta_0 + \beta_1 X \quad (17)$$

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

Figure 2 reports the willingness to accept (WTA) demand curve for giving up

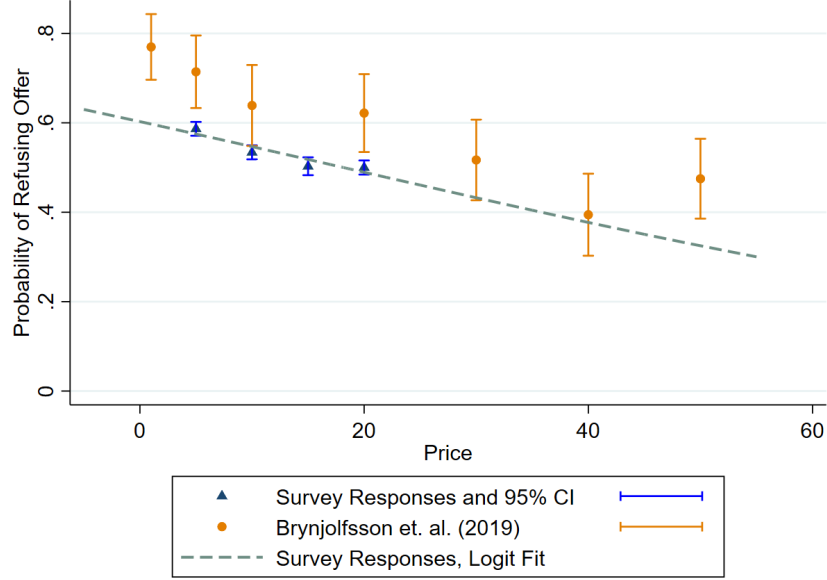


Figure 2: Probability of rejecting an offer to give up Facebook for one month for price listed. Mean responses to the question “Would you give up Facebook for 1 month in exchange for \$X? Choose Yes if you do not use Facebook.” 95% confidence intervals are based on binomial statistics. This survey’s findings plotted with blue triangles. Responses are contrasted with results from a directly analogous question from Brynjolfsson et al. (2019) (orange circles.). The logistic line of best fit for the new results is plotted with a dashed line.

Facebook for one month combining together all demographic groups in our sample. This constitutes a representative survey of the US internet using population population. 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) which asked a directly comparable questions. Our results are broadly in line with that previous paper, which had some methodological differences, but indicate a slightly lower valuations.

Figures A1 through A12 plot mean WTA responses and demand curves for various subgroups of the population. Table A1 reports the estimates underlying these curves.

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

$$p(\epsilon_i) \sim \frac{e^{-\frac{\epsilon_i - \eta_i}{s_i}}}{s_i \left(1 + e^{-\frac{\epsilon_i - \eta_i}{s_i}}\right)^2} \quad (18)$$

where

$$s_i = (\beta_{1,i})^{-1} \quad (19)$$

and

$$\eta_i = (-\beta_{0,i})s_i \quad (20)$$

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 and in disutility from advertising, i.e.

$$\mu_i = \sum^J U_i(j)P_j z_i(j)D_j - a_i\phi_i \quad (21)$$

where $U_i(j)$ is the (linear) 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. $\frac{\partial \mu_i}{\partial \phi_i}$).

We estimate the parameters of 21 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 current revenues 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.¹² To calculate initial revenue per user $\bar{\phi}_i$ we take in data on the cost of advertising to users of different types from Facebook's advertisement API. After selecting which demographic group to target, Facebook 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 to be the relative value of a demographic to Facebook's ad revenue (When a range is provided, we use the mean). 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

$$\bar{\phi}_i = z \text{Relative Value}_i \quad (22)$$

and

$$11.62 = q \frac{\sum^I \text{Relative Value}_i \bar{P}_i D_i}{\sum^I \bar{P}_i D_i} \quad (23)$$

where q is a scaling term, \bar{P}_i is the estimate of the initial participation rate on Facebook

¹²This is derived from Facebook's 2019 Q1 annual report, where they report \$ 34.86 in revenues per North American user per quarter.

by the demographic group (taken as our estimate of $\Omega_i(\overline{\mu_i}) - \mu_i = 0$), 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 type i is friends with, we combine the results of two sets of survey questions. First, we ask questions to solicit the total number of friends each demographic has on average. We then ask questions to solicit what percentage of their friends of of each demographic. We re-balance these responses to add to 100 percent (including a catchall category for individuals under age 18, who are not directly modeled). Figure 3 presents our estimate of the average number of friends by type for each demographic.

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

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 rescale these responses by the estimated number of friends each ethnic 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 .

Figure 4 graphically represents 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 4 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. Figures 5 restrict 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 elderly women, with middle aged women next. They value individual connections of all types only slightly, but amongst these most value connections to Males age 55-64.

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 platform changes its advertising level as $t = 1$. The participation rate on the platform for a demographic group after cascade t is

$$P_{i,t} = \Omega_i \left(\sum^J U_i(j) z_i(j) D_j P_{j,t-1} - a_i \phi_{i,t} \right) \quad (24)$$

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

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

$$\int_0^{P_{i,t}} ((\mu_i(\vec{P}_{j,t-1}, \phi_i) - e_i(\rho_i))) d\rho_i \quad (25)$$

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

$$e_i = -s_i \log\left(\frac{1 - p_i}{p_i}\right) + \mu_i \quad (26)$$

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.

The revenue to Facebook from user participation of a given demographic after t



Figure 4: A graphical representation of 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 $U_i(j)$. 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.

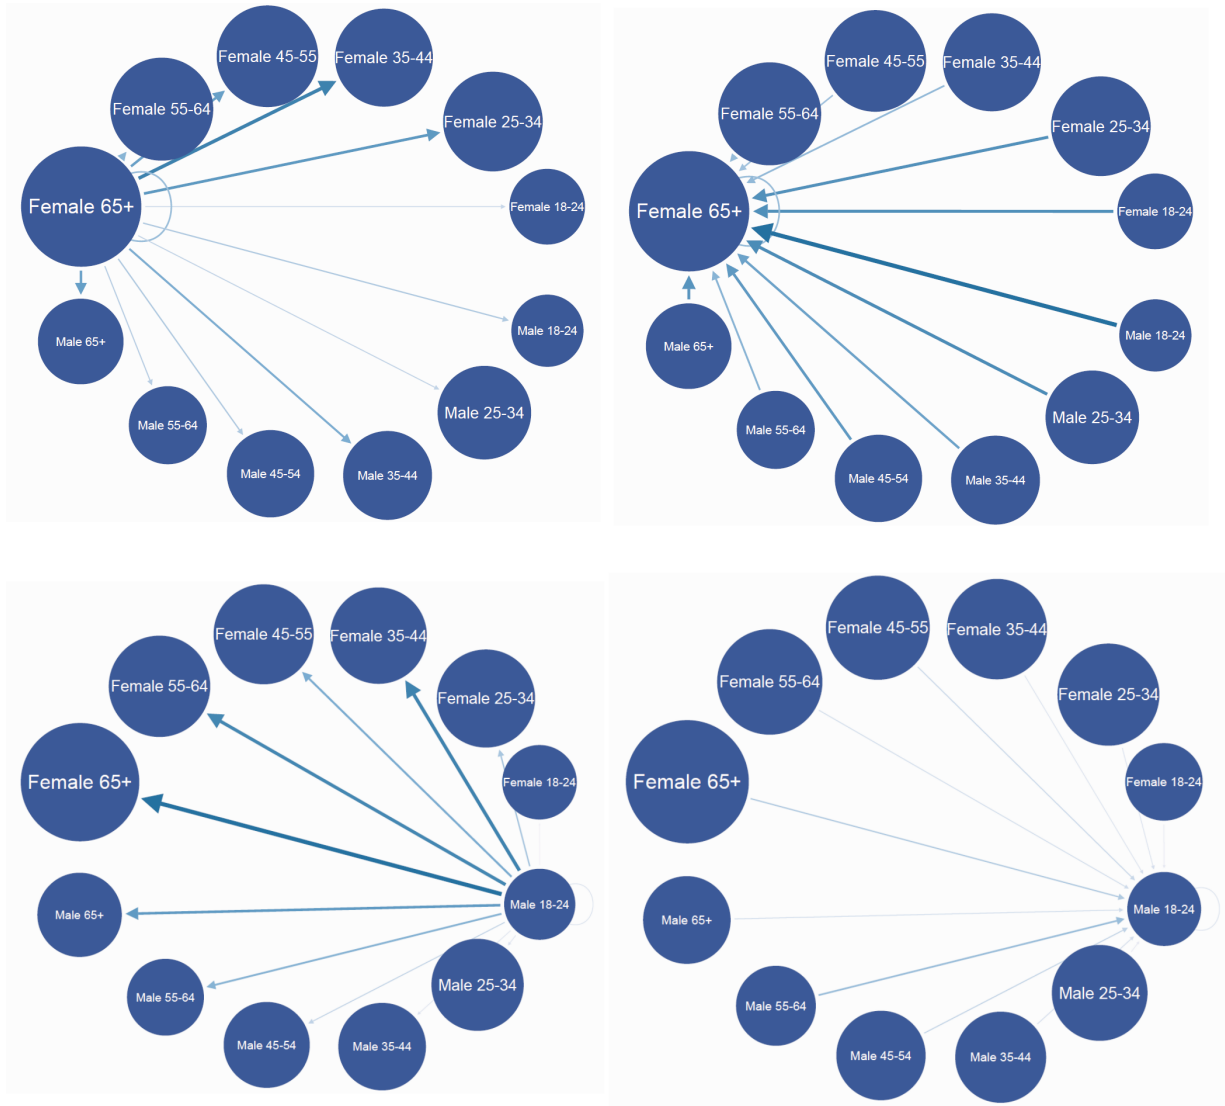


Figure 5: A graphical representation of Facebook usage and network externalities by market segment. The size of each node represents the relative current size of the Facebook user base by demographic. The thickness of the arrows corresponds to the relative value received by a Facebook user of the demographic the arrow is pointing towards from a friendship with a user of the source demographic $U_i(j)$. 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.

cascades is

$$\Phi_{i,t} = \phi_{i,t} D_i P_{i,t} \tag{27}$$

we calculate 1000 cascades of the network effect but, as will be seen, most of the action occurs in the first few cascades.

6 Illustration Results

With the parameterized model in hand, we can now proceed to simulating counterfactual pricing strategies and potential government policies. We will begin by estimating Facebook’s profit maximizing strategy. We then calculate the non-monetary value Facebook places on users which justifies their current monetization strategy as optimal – this is taken into account in the subsequent policy simulations.

6.1 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 we assume that the marginal disutility from increased monetization $\frac{\partial \mu_i}{\partial p h_i} a_i$ is equal to 1 for each group.¹³ While our model and simulation program allows for price discrimination among different demographic groups, we have no evidence on how Facebook currently 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 6 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

¹³ $a_i = 1$ 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.

		Ad Increase						
		Initial	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 6: Changes in consumer surplus and Facebook profit after N cascades in billions of dollars per month, after Facebook implements its profit maximizing monetization strategy.

to remaining users. Cascade 1 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 3 cascades of the network effect, Facebook reaches its new equilibrium to three significant digits of precision.

Why is Facebook leaving so much money on the table? One possibility is that Facebook 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 (for example, 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.

6.2 The First Best

Before evaluating the impact of different possible reforms, 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 table 2. Because of 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

	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 2: 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. 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.

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 .¹⁴ Under this policy, 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.

6.3 Evaluating Tax and Redistribution Policies

We next simulate the consequences of three tax and redistribution policies. 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

¹⁴There 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 an infinite subsidy. Second, for reasons explained more below in the ‘Data as Labor’ example, it is unrealistic to believe that Facebook usage could be frictionlessly subsidized, because it would induce scammers to enter with fake accounts.

a 10 dollar per user tax.¹⁵ A more radical proposal is the “Data as Labor” framework proposed in Weyl (2010). 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 demographic group the full revenue it collects from displaying them ads.

Before we proceed to simulations, our model has a novel theoretical point to make about the incidence of taxes on digital platforms. So long as a tax is flatly applied to all platform sources of revenue and utility, it will not distort the platform’s optimal vector of ϕ_i ’s. To see this, consider the maximization problem 7. A tax that is equally applied to all ϕ_i revenues would add a multiplier term to this equation. It would not change the firm’s first order maximization equation. In other words, under the assumptions of our model, including the assumption that platform’s only source of revenues are advertisements and the platform faces no marginal cost, a flat tax on advertising revenues is fully incident on the platform’s profits.¹⁶ We have found however, that Facebook derives non-monetary utility from maintaining a large user base. Therefore a tax on ad-revenues will cause it to shift between it’s two tasks from making revenues from selling advertising to increasing utility by cultivating a large user base. 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.

Table 3 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 it’s 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 the value that they create on social media is a clear win-win from a perspective of social welfare. The question is whether such a policy is implementable. Doubtless, were a cash payment for using Facebook to be implemented, the website would be swamped

¹⁵Professor Grauwe proposes this as an annual levy, but we consider a per-month tax which would raise the same amount as a 3% of revenues tax.

¹⁶In 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 3: Evaluating tax and redistribution policies. Current and simulated % Change in Facebook advertisement 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.

by scammers creating phony accounts.

6.4 Evaluating Digital Platform Regulations

The final set of policies we simulate are regulatory. Many proposals have been made for regulation of online 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, one assumed to be implemented perfectly, and two ‘worst case scenarios’ for botched reforms. The positive reforms are a move to increase the competitiveness of social media.

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. Perhaps the most important reason increasing platform competition is not an obvious win is that it has the potential to destroy network effects by splitting the market. 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 draw and purpose of digital platforms. To resolve this last concern, a recent study of anti-trust 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 Facebook to share posts and other communications with competitor social networks, who would then be allowed

to display them on their platforms. We consider ‘perfect competition’ as entailing this interoperability, and model it as the elimination of all advertisements on Facebook.¹⁷

One component of many plans to increase platform competition includes mandatory ‘breakups’. For example, an essay by a leading presidential candidate calls for, among other things, Facebook to be split from Instagram and Whatsapp (Warren, 2019). To the 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 botched 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 ‘worst case scenarios’. The first is a botched horizontal Facebook breakup resulting in the creation of two Facebook monopolies each serving half of the US population.¹⁸ The final scenario we simulate is a botched ‘vertical breakup’ that results in five percent of the US population completely losing interest in using Facebook. This represents the percentage of the population that only uses Facebook only because of its synergies with Instagram, and would stop using it if these products were completely disconnected.

Results from these three simulations are summarized in table 4. 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 botched breakups, which destroy network effects without an offsetting increase in platform quality, are unsurprisingly disastrous of social welfare. The worse of these is the botched horizontal breakup, which would lower social welfare by as much as 84.7%.

¹⁷There 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.

¹⁸Such 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.

	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 4: 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 failed regulation that reduces Facebook quality for a subset 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.

7 Discussion

Building on Rochet and Tirole (2003), Parker and Van Alstyne (2005) and Weyl (2010) we construct and illustrate an approach for firms to incorporate network effects in their monetization strategies. Taking the first order condition for profit maximization with respect to the advertising 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), don’t have much disutility from advertisements, and don’t 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 Facebook elastically (the first cascade of the network effect).

We then build a simulation tool for evaluating 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 that appear 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.

We find that Facebook’s monopoly power reduces social welfare by 9.6% relative

to the first best, and 4.8% relative to perfect competition. In addition to the assumptions embedded in our parameterization of the model, these results incorporate two additional key assumptions. The first is the assumption Facebook’s shadow value of 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 the assumption 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).¹⁹ Botched Facebook breakups, which reduce the quality of Facebook or the ability of users to connect on social media without increasing competition, are unambiguously bad.

Our approach is not without weaknesses. One important issue is trickiness 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. Both are clearly simplifications. 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. Another limitation of the current approach is that advertisers are treated as price setters, rather than as a side of the market. A more complete model would treat advertisers as a heterogenous mix of agents as well. Finally, our model conceives of

¹⁹These results hold so long as taxes are *more* incident on monetary value than shadow values, which seems plausible.

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.

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.
- 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, E, A Collis, and F Eggers**, “Using massive online choice experiments to measure changes in well-being.,” *Proceedings of the National Academy of Sciences of the United States of America*, 2019.
- 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**, “Co-creation of value in a platform ecosystem: The case of enterprise software,” *MIS Quarterly*, *Forthcoming*, 2011.
- 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.

- Hagi, Andrei**, “Two-sided platforms: Product variety and pricing structures,” *Journal of Economics & Management Strategy*, 2009, 18 (4), 1011–1043.
- 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, Ms Shan, Sinan Aral, Jeffrey Yu Hu, and Erik Brynjolfsson**, “Social Advertising Effectiveness Across Products: A Large-Scale Field Experiment,” 2018.
- Jackson, Matthew O**, *Social and economic networks*, Princeton university press, 2010.
- Morton, Fiona Scott, 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.
- 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.
- 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.
- 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.

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.

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.

A Additional Tables and Figures




Figure A1: 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 A2: 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 A3: 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 A4: 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 A5: 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 A6: 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 A7: 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 A8: 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 A9: 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 A10: 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 A11: 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.

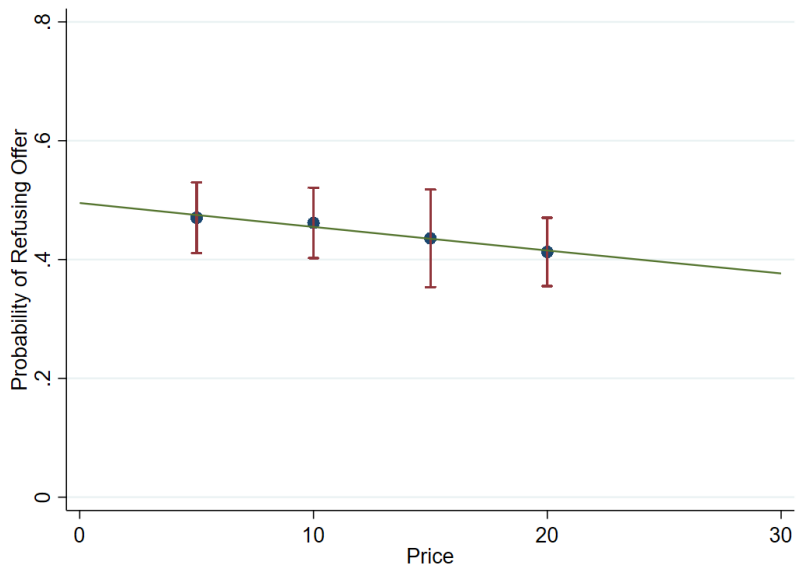


Figure A12: Underlying data and estimate of the the demand curve (Ω_i) for men age 65 or older. The points are the mean response to the question “Would you give up Facebook for 1 month in exchange for \$X? Choose Yes if you do not use Facebook.” for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.

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 A1: 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).

B Network Stability

B.1 Stability of Equilibria

An important first question is whether the network just described is stable. We define a network as stable at equilibrium \vec{P} if the derivative of a connected individual's best response function with response to these probabilities is less than 1.²⁰ This is a version of a 'trembling hand' perfect equilibrium, meaning that the equilibrium is robust to small fluctuations in each individual's likelihood of participation.

For a symmetric network (i.e. every individuals' demand function Ω is identical), assuming that utility is linearly additive in the network effects and disutility from advertisements, the probability of participation for any individual is

$$P = \Omega\left(\sum^I U(i)P - a\phi\right) \quad (28)$$

where $U(i)$ is the value of any connection.²¹ Then the best response function is

$$\frac{\partial \Omega}{\partial P} = \frac{\partial \Omega}{\partial U} \left(\sum^I U(i) - a\phi \right) \quad (29)$$

And so a network equilibrium is stable so long as

$$1 > \frac{\partial \Omega}{\partial U} U(i)(I - 1) \quad (30)$$

In other words, a network equilibrium is stable so long as the average user doesn't have too many connections, is too elastic in their individual participation, or gains too much value from every additional connection. If the inequality is violated, small deviations from an equilibrium are liable to send participation to a boundary condition of 100% participation or zero participation.

B.2 Stability of Equilibrium to Demand Shock

Relatedly, we can also consider the resilience of a network equilibrium to a shock in preferences.

Theorem 1. *Consider a symmetric network where Ω is continuously differentiable and utility is linearly additive in network effects and the disutility from advertisement. Then for any stable equilibrium (as defined above) $\frac{P_i}{\phi_j}$ and $\frac{P_j}{\phi_i}$ are finite*

²⁰This concept of equilibrium stability borrows from Jackson (2010) section (9.7.2). In that model, only some individuals are connected in the network, but in our model all are connected. In that model p corresponds to the percentage of neighbors who participate, but in our model it corresponds to the likelihood of anyone who participates.

²¹the value of a 'connection to oneself' is assumed to be 0

Proof. Rewriting equation 13 with the assumption all nodes are identical, before i gets hit with a fee, yields:

$$\frac{\partial P_j}{\partial \phi_i} = \frac{\partial \Omega}{\partial \mathbf{U}} \left(U(i) \frac{\partial P_i}{\partial \phi_i} + \sum_{k \neq i}^K U(i) \frac{\partial P_k}{\partial \phi_i} \right) \quad (31)$$

Substituting in 12 and summing yields

$$\frac{\partial P_j}{\partial \phi_i} = \frac{\partial \Omega}{\partial \mathbf{U}} \left((I-2) \frac{\partial P_j}{\partial \phi_i} U(i) + \frac{\partial \Omega}{\partial \mathbf{U}} U(i) \left(U(i)(I-1) \frac{\partial P_j}{\partial \phi_i} - \frac{\partial A}{\partial \phi_i} \right) \right) \quad (32)$$

Solving for $\frac{\partial P_j}{\partial \phi_i}$ yields

$$\frac{\partial P_j}{\partial \phi_i} = \frac{-\left(\frac{\partial \Omega}{\partial \mathbf{U}}\right)^2 U(i) \frac{\partial A}{\partial \phi_i}}{1 - \left(\frac{\partial \Omega}{\partial \mathbf{U}} U(i)(I-2) + \left(\frac{\partial \Omega}{\partial \mathbf{U}}\right)^2 (U(i))^2 (I-1)\right)} \quad (33)$$

The network will not unravel due to a welfare change so long as 33 is not infinite. This is equivalent to showing that the denominator is not equal to zero (as all other terms are finite).

However, the denominator never takes the value 0 when the network stability criteria is satisfied. Rearranging terms, the denominator can be written as

$$1 - \frac{\partial \Omega}{\partial \mathbf{U}} U(i) \left((I-2) + \left(\frac{\partial \Omega}{\partial \mathbf{U}}\right) (U(i))(I-1) \right) \quad (34)$$

From the assumption that the network is stable, we have

$$1 > \frac{\partial \Omega}{\partial \mathbf{U}} U(i)(I-1) \quad (35)$$

This implies

$$I-1 > (I-2) + \frac{\partial \Omega}{\partial \mathbf{U}} U(i)(I-1) \quad (36)$$

and applying B again implies

$$1 > \frac{\partial \Omega}{\partial \mathbf{U}} U(i) \left((I-2) + \left(\frac{\partial \Omega}{\partial \mathbf{U}}\right) (U(i))(I-1) \right) \quad (37)$$

And if $\frac{\partial P_j}{\partial \phi_i}$ is finite, clearly so to is $\frac{\partial P_i}{\partial \phi_i}$. So long as the network is stable in the normal sense, it is stable to welfare shocks. □

Lemma 2. *In a symmetric network, $\frac{\partial P_i}{\partial \phi_j} = 0$ if $\left(\frac{\partial \Omega}{\partial \mathbf{U}}\right)^2 \frac{\partial A}{\partial T} U(j) = 0$*

Proof. Directly from (33) □

C Social Welfare Maximization

The increase in welfare due to the platform's existence for a given individual i is

$$W_i = P_i E[\mu_i(\vec{P}, \phi_i) - \epsilon_i | U_i > \epsilon_i] \quad (38)$$

in other words, welfare from the platform is the odds an individual participates on the platform, multiplied by their expected surplus from platform use. This expected surplus is equal to the value of platform use less opportunity cost.

Evaluating this equation yields

$$W_i = P_i \int_{-\infty}^{U_i} \frac{\mu_i(\vec{P}, \phi_i) - \epsilon_i}{\text{Prob}(U_i > \epsilon_i)} f(\epsilon_i) d\epsilon_i \quad (39)$$

where $f(\epsilon_i)$ is the pdf of ϵ_i . There is an upper bound on the integral, because an individual only participates – and pays the opportunity cost – if the value of participation exceeds the opportunity cost. Now, μ_i is a constant with reference to the integral, so this reduces to

$$W_i = P_i \frac{\mu_i(\vec{P}, \phi_i) F(\epsilon_i)}{\text{Prob}(U_i > \epsilon_i)} \Big|_{-\infty}^{U_i} - P_i \int_{-\infty}^{U_i} \frac{\epsilon_i}{\text{Prob}(U_i > \epsilon_i)} f(\epsilon_i) d\epsilon_i \quad (40)$$

where $F(\epsilon_i)$ is the CDF of ϵ_i . Now, $\text{Prob}(U_i > \epsilon_i) = F(U_i) = P_i$ so

$$W_i = P_i \mu_i(\vec{P}, \phi_i) - \int_{-\infty}^{U_i} \epsilon_i f(\epsilon_i) d\epsilon_i \quad (41)$$

Social welfare maximization needs to take into account both consumer surplus and platform surplus. Using the same equation for platform profits as used above, this means social welfare maximization entails

$$\max_{\phi_i} \sum_i^I [P_i(\phi_i + \mu_i(\vec{P}, \phi_i)) - Q_i(\epsilon_i) \Big|_{-\infty}^{U_i}] - F \quad (42)$$

s.t.

$$P_i = \Omega_i(\mathbf{U}_i) \quad (43)$$

Where $Q_i(\epsilon_i)$ stands for the indefinite expectation integral $\int \epsilon_i f(\epsilon_i) d\epsilon_i$.

The first term is the utility from participation to users of the platform μ_i and to the firm ϕ_i . These are both multiplied by the odds of participation. The next term is the expected opportunity cost to an individual from participating.

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