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

Link to Latest Draft

Seth G. Benzell and Avinash Collis*

November 4, 2019

Abstract

Digital platforms, such as Facebook, Uber, and AirBnB, create value by connecting users, creators, and contractors of different types. Their rapid growth, nontraditional business models, and disruptive nature presents challenges for managers and asset pricers. These features also, arguably, make them natural monopolies, leading to increasing calls for special regulations and taxes. We construct and illustrate an approach for modeling digital platforms. The model allows for heterogeneity in elasticity of demand and heterogeneous network effects across different users. We paramaterize our model using a survey of over 40,000 US internet users on their demand for Facebook. Facebook creates about 11.2 billion dollars in consumer surplus a month for US users age 25 or over, in line with previous estimates. We find Facebook has too low a level of advertising relative to their revenue maximizing strategy, suggesting that they also value maintaining a large user base. We simulate six proposed government policies for digital platforms, taking Facebook's optimal response into account. Taxes are mostly incident on Facebook profits, only slightly changing consumer surplus. Three more radical proposals, including 'data as labor' and nationalization, have the potential to raise consumer surplus by up to 42%. But a botched regulation that left the US with two smaller, non-competitive social media monopolies would decrease consumer surplus by 44%.

1   Introduction

Much of the value of many digital platform businesses comes from what are known as "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 regulations and taxes on Facebook profits and consumer surplus.

We make three main contributions. First, we provide a tractable framework for third-degree price discrimination on multi-sided platforms. This approach builds on traditional price discrimination models for price discrimination in multi-market consumer population by taking into account the externalities among them. 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 and participation 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 network in the sense that each individual or market segment can be thought of as a side of the network.[1] We show that in this setting, the optimal pricing strategy 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 40,000 surveys conducted through Google Surveys. We categorize the surveyed into ten 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

---

[1]When 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).

the spirit of (Brynjolfsson et al., 2019) and (Allcott et al., 2019).[2] 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 ten by ten 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 2.65 billion dollars a month (from a baseline of 1.57 billion) if it optimally price discriminated. It could raise revenues by only 2.08 billion dollars a month if it increased monetization equally across users. This begs the question of why Facebook is 'leaving so much money on the table'. Implementing their revenue maximizing strategy entails squeezing value from their most inelastic users, reducing Facebook usage by 55.6% and lowering total consumer surplus by 82.8%. We infer that in addition to maximizing current revenues, Facebook values maintaining a large and happy user base. We impute the 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 non-monetary 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 consider three taxation and redistributive policies. We show theoretically that a flatly applied 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 2% and that a ten dollar per month tax on users would lower consumer surplus by 3%. 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

---

[2]These papers measure the consumer surplus generated by digital goods by conducting discrete choice experiments where they offer consumers the choice to give up access to the good in exchange for monetary compensation. We build on these papers by asking about a new type of free good (the value of social connections) as well as by using information from the full distribution of responses to fit a demand curve (in our case logistic) rather than focusing on median and average responses.

targeted advertisements (Posner and Weyl, 2018). We conceive of this policy as a rebate of Facebook's current advertising revenues to users. We find this policy would boost consumer welfare by 24%, about 58% of which 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.

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. A second policy we evaluate is the nationalization of Facebook for the purpose of maximizing its social welfare. If transfers can be frictionlessly distributed to Facebook users, the optimal policy is to implement an infinite subsidy for Facebook use – this is because transfers don't change social surplus, and every user of Facebook creates positive externalities for other users. In our simulation, we assume that a nationalized Facebook can create value for users as the inverse of how it creates disutility through monetization.[3] Finally we simulate the results of a 'botched' Facebook breakup which leaves America with two monopolies over half of the population each. We predict that perfect competition would raise consumer surplus by 9%, at the cost of eliminating all monetary profits, and that a social welfare maximizing Facebook would raise consumer surplus by 42%, at the cost of Facebook needing to go -255% into the red. Breaking Facebook into two non-competitive 'baby Facebooks' would be disastrous, lowering consumer surplus by 44%. It would also lower combined ad revenues by 93% as the baby Facebooks lowered advertising rates to retain even 82% of their original combined user base.

The paper concludes with a discussion of the strengths and weaknesses of this approach to modelling platform businesses and contemplates future work.

2 Related Literature

A rich stream of theoretical literature studying network effects in the context of platform businesses has evolved over the past decade and a half. 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 vari-

---

[3]In other words, if we estimate that a group experiences 20 cents in disutility from a dollar's worth of revenue in advertising, we assume that a nationalized Facebook can only increase the desirability of Facebook to a demographic by one dollar by spending five dollars.

ous 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 different types of users may provide different network externalities. 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 in-between fares). However, a large literature suggests that even within a 'side' of a one or two-sided platform, users are heterogenous in the effect their actions have on the network. The empirical literature on network effects uses several techniques for their estimation, including studying exogenous shocks to the network (e.g. Tucker (2008)), using an instrumental variable approach (e.g. Aral and Nicolaides (2017)) and conducting field experiments (e.g. Aral and Walker (2012)).

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.[4]

---

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

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 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. [5] 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.[6] 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 instru-

---

[5]The fact that platforms cannot fully first-degree price discriminate is testified to by papers which show that users benefit considerably on average from joining a platform. For example, Ceccagnoli et al. (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.

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

ments 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 utilitous advertisements.

3 Analytic Model

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 in a network. 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.

We use the example of a social network, because our implementation section takes place in that setting. Therefore, in our baseline model, other incidental network characteristics mimic that of an internet 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. We assume that the fee or subsidy faced by each participating network user is a binary function of their decision to participate on the network. This assumption is easy to modify for other contexts where fees are a function of the number or type of connections or interactions.[7]

The platform's monetization is also modeled. Users face disutility depending on how intensely they are monetized by the platform. This may correspond to the unpleasantness of advertisements or 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.

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. As a first pass, a platform might distinguish between the network externalities and demand characteristics of broad user groups such as women and men. A more sophisticated platform with a larger research budget might estimate and incorporate into their optimization network externalities at the individual level. When calibrating

---

[7]An example of an online platform with network effects and a non-binary fee structure would be an online auction house like eBay. eBay's main source of revenue is a progressive fee on the value of every transaction.

the model, we make additional assumptions about the functional form of user demand for the platform.

3.1   Consumers

A consumer i chooses whether to participate in the platform ($\boldsymbol{P_i} = 1$) or not ($\boldsymbol{P_i} = 0$).[8] If the consumer i uses the platform ($\boldsymbol{P_i} = 1$)), they expect to receive

$$E[U_i(\boldsymbol{P_i} = 1)] = \mu_i(P_1, .., P_I, -\phi_i) \tag{1}$$

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. $\frac{\partial \mu_i}{\partial \phi_i} = .2$). Local telephone calls and pre-2016 WhatsApp monetized by charging a flat fee for participation (i.e. $\frac{\partial \mu_i}{\partial \phi_i} = \$1$). [9] Note that users do not directly care about what other users are charged, but it is indirectly important to them insofar as it causes other users to participate on the network.

$\frac{\partial \mu_i}{\partial P_j}$ is the marginal utility of j being on the network to i (if i participates). For convenience, we will sometimes write the marginal value of a user j to a user i as

$$U_i(j) = \frac{\partial \mu_i}{\partial P_j} \tag{2}$$

and the marginal disutility of advertising as

$$a_i = \frac{\partial \mu_i}{\partial \phi_i} \tag{3}$$

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 $U_i(j)$ and a_i are constant.[10]

---

[8]Note 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 ten market segments in our calibration.

[9]In general, platforms monetize in many different ways. Some monetize by charging fees for transactions (Ebay, AirBnB, etc), some subsidize one side while charging others (Credit Cards), some by charging a flat fee for participation (Local telephone calls, pre-2106 WhatsApp), and some monetize by charging advertisers or selling advertisements (social networks). Our baseline model is best suited for evaluating the latter two approaches, but can be straightforwardly modified to handle other monetization methods.

[10]The 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 in the case of other networks it is

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

$$U_i(\boldsymbol{P_i} = 0) = \epsilon_i \tag{4}$$

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 $U(\boldsymbol{P} = 1)$ is consumer specific.[11]

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 $U_i(\boldsymbol{P_i} = 1)$ – 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.

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

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

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:

$$P_i = Prob[E[(U_i(\boldsymbol{P_i} = 1)] > \epsilon_i] = \Omega_i(\mu_i) \tag{5}$$

for more useful notation, define

$$\boldsymbol{U_i} \equiv E[U_i(\boldsymbol{P_i} = 1)] = \mu_i \tag{6}$$

The network is in equilibrium when individuals' decisions to participate are optimal

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.

[11]By 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.

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 the symmetric network (i.e. where all individuals have the same ϵ distribution, A, and network externality), where utility is linearly additive in the network effects and disutility from advertising, the equilibrium is stable so long as

$$1 > \frac{\partial \Omega}{\partial \boldsymbol{U}} U(i)(I - 1) \tag{7}$$

where $U(i)$ 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. 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 revenue maximizing monopoly social network.

4.1   Monopoly Firm Price Setting

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

$$\Phi = \sum_i^I \phi_i \boldsymbol{P_i} - F \tag{8}$$

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. $\boldsymbol{P_i}$ is a binary

indicator of whether the consumer participates. F is the fixed cost of the platform firms' operation.[12]

P_i's are independent random variables, so firms maximize

$$E[\Phi] = \sum_i^I \phi_i P_i - F \tag{9}$$

where

$$P_i = E[\boldsymbol{P_i}] = \Omega_i(\boldsymbol{U_i}) = \Omega_i(\phi_1, \phi_2, ...) \tag{10}$$

the probability of a consumer participating P_i is an individual specific function of $\boldsymbol{U_i}$. 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.[13]

The firm seeks to maximize revenues

$$\max_{\phi_i} E[\Phi] = \sum_i^I \phi_i P_i \tag{11}$$

s.t.

$$P_i = \Omega_i(\boldsymbol{U_i}) \tag{12}$$

This yields the following first order condition

$$\frac{\partial \Phi}{\partial \phi_i} = P_i + \phi_i \frac{\partial P_i}{\partial \phi_i} + \sum_{j \neq i}^J \phi_j \frac{\partial P_j}{\partial \phi_i} \tag{13}$$

where

$$\frac{\partial P_i}{\partial \phi_i} = \frac{\partial \Omega_i}{\partial \boldsymbol{U_i}} \left(-\frac{\partial \mu_i}{\partial \phi_i} + \sum_j^J \left(\frac{\partial \mu_i}{\partial P_j} \frac{\partial P_j}{\partial \phi_i} \right) \right) \tag{14}$$

and,

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

This recursion is natural as P_i is a function of P_j, which is a function of P_i, etc. Equation (15) 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. In our calibrated example evaluating only the first two recursions, or cascades, of this function

---

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

[13]Although hypothetically 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.

tends to yield a decent approximation of the total change in revenues from a pricing change.

4.2   Profit Maximization Problem Simplified

Equation 13 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 first term in brackets in 14 and 15. In other words, the following equations only take into account one cascade of network effects.[14] For clarity and parsimony, we also make the substitutions from equations 2 and 3

$$\frac{\partial P_i}{\partial \phi_i} = \frac{\partial \Omega_i}{\partial \boldsymbol{U_i}} \left(-a_i + \sum_{j \neq i}^{J} \left(\cancel{U_i(j) \frac{\partial P_j}{\partial \phi_i}} \right) \right) \tag{16}$$

and,

$$\frac{\partial P_j}{\partial \phi_i} = \frac{\partial \Omega_j}{\partial \boldsymbol{U_j}} \left(U_j(i) \frac{\partial P_i}{\partial \phi_i} + \sum_{k \neq i}^{K} \cancel{U_j(k) \frac{\partial P_k}{\partial \phi_i}} \right) \tag{17}$$

Then, substituting into 13, yields a new simplified first order condition

$$\frac{\partial \Phi}{\partial \phi_i} = \underbrace{P_i - \phi_i a_i \frac{\partial \Omega_i}{\partial \boldsymbol{U_i}}}_{\text{Direct Effect}} - \underbrace{a_i \frac{\partial \Omega_i}{\partial \boldsymbol{U_i}} \sum_{j \neq i}^{J} \phi_j \frac{\partial \Omega_j}{\partial \boldsymbol{U_j}} U_j(i)}_{\text{Network Effect}} \tag{18}$$

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 $U_j(i) = 0 \; \forall \; i, j$, i.e. when no network effects are present, 18 reduces to this pair of terms).

The last term in equation 18 is the network effect of an advertising increase. The increase in advertising makes i less likely to participate (in this approximation, by an amount $a_i \frac{\partial \Omega_i}{\partial U_i}$) which leads others to stop participating (by an amount $\frac{\partial \Omega_j}{\partial U_j} U_j(i)$). When these third parties stop participating, the platform loses on the current revenues

---

[14]In the parametric section we will show that the first cascade of network effects is quantitatively much more important than subsequent cascades for a reasonable parameterization.

that they were paying ϕ_j.

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 equation is a powerful tool for managers to think about monetizing their platform. While a similar equation might be able to be derived from the model in (Weyl, 2010) (i.e. that this result is latent in that model), a major contribution of this paper is a recasting of the firm's maximization problem in terms of elasticities of demand and other more interpretable terms.

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. Intuitively, 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, wheras 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 approximately 40,000 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 ten market segments. These are a pair of genders and five age brackets. The market segments we consider are

- Gender: Male or Female

- Age: 25-34; 35-44; 45-54; 55-64; and 65+

we also intended to include individuals age 18-24 in our analysis, but we found it

13

difficult to get a sufficient number of survey responses for this group. Individuals under the age 13 are not formally allowed to have Facebook accounts.

We asked the following sets of questions about individuals' demand for Facebook, combining responses within the ten market segments described. The full list of surveys conducted is documented in figure 1.

Question Text	Possible Responses	Number of Responses
How many friends do you have on *Facebook*?	Survey 1: 0-100; 100-200; 200-300; 300-400; 400-500; More than 500; I do not use Facebook Survey 2: 0-50; 50-100; 100-500; 500-700; 700-900; More than 900; I do not use Facebook	2509
Would you give up Facebook for *1 month* in exchange for $[X]? Choose Yes if you do not use Facebook.	Yes, I will give up Facebook; No, I would need more money	20050
"On Facebook, would you unfriend all your friends who are *[demo_group]* for 1 month in exchange for $[X]? Choose Yes if you do not use Facebook."	Yes, I will unfriend all these friends; No, I would need more money	13272
What *percentage* of your *friends on Facebook* are *[demo_group]*?	0-10%; 10%-20%; 20%-40%; 40%-60%; 60%-80%; 80%-100%; I do not use Facebook	13037
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.	0; $1-$5; $5-$10; $10-$15; $15-$20; $20+	1001

Figure 1: List of surveys conducted

Figure 2 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.

5.1   Calibrating the Model for Facebook Market Segments

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 to select 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
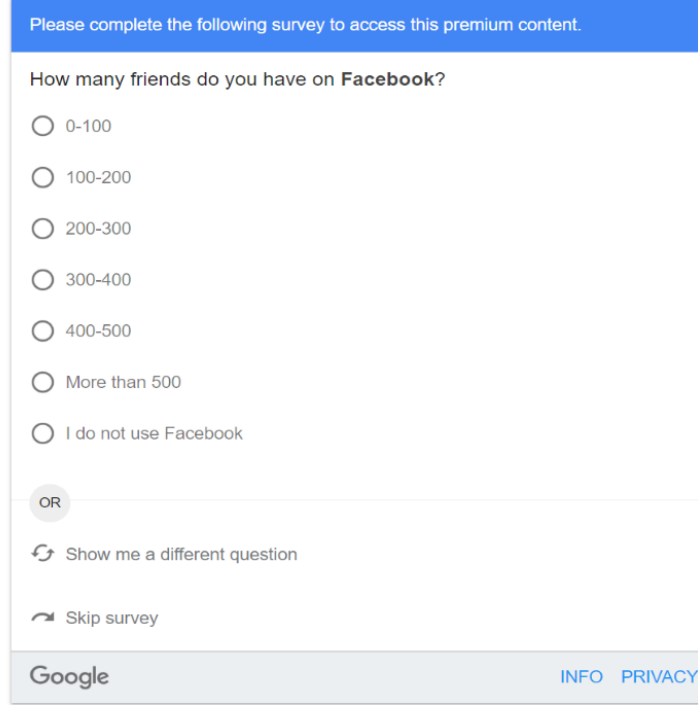
14

Figure 2: Google survey interface example. Note that each respondent only receives a single survey question, and that responses are limited to seven multiple choice options.

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.". Figures 9 through 18 report responses to these questions, and the logistic line of best fit. Table 1 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 equations

$$p(\epsilon_i) \sim \frac{e^{-\frac{\epsilon_i - \eta_i}{s_i}}}{s_i \left(1 + e^{-\frac{\epsilon_i - \eta_i}{s_i}}\right)^2} \tag{19}$$

where

$$s_i = (\text{Coef. on Cost}_i)^{-1} \tag{20}$$

and

$$\eta_i = (-\text{Intercept}_i)s_i \tag{21}$$

As an argument the function Ω_i takes the change in the value of Facebook due to an individual gaining or losing friends, or from experiencing a direct change in their advertising level ϕ_i. 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.

$$\mu_i = \sum^{J} U_i(j)P_j z_i(j)D_j - a_i\phi_i \tag{22}$$

where $U_i(j)$ is the (linear) utility an individual i receives from having a friend in market segment 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.

We estimate the parameters of 22 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.[15] To calculate initial revenue per user $\overline{\phi_i}$ we take in data on the cost of advertising to users of different types from Facebook's advertisement API. After selecting which demographic group you would like to target, Facebook tells you 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. By taking as given that the average value of a user per month is \$11.62, we can then calculate the revenue per user of a demographic using the following equations

$$\overline{\phi_i} = z\text{Relative Value}_i \tag{23}$$

and

$$11.62 = q\frac{\sum^{I} \text{Relative Value}_i\overline{P_i}D_i}{\sum^{I} \overline{P_i}D_i} \tag{24}$$

where q is a scaling term, $\overline{P_i}$ is the estimate of the initial participation rate on Facebook 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

---

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

age 25, who are not directly modeled). Figure 3 presents our estimate of the average number of friends by type for each demographic.

GENDER		Female	Female	Female	Female	Female	Male	Male	Male	Male	Male
	AGE	25-34	35-44	45-54	55-64	65+	25-34	35-44	45-54	55-64	65+
Female	25-34	71.0	47.2	43.2	24.0	21.5	64.2	34.9	18.3	19.7	15.3
Female	35-44	50.8	52.6	41.1	25.0	19.2	25.0	31.3	18.6	15.1	10.6
Female	45-54	42.3	38.7	34.0	29.1	14.6	21.7	23.6	25.4	23.7	10.7
Female	55-64	22.8	29.3	34.4	37.2	21.5	12.2	17.4	20.6	14.9	11.9
Female	65+	10.5	18.9	20.8	20.1	16.7	8.3	13.0	11.1	16.6	11.0
Male	25-34	46.1	41.0	12.6	25.9	23.8	51.6	31.0	31.0	21.4	14.7
Male	35-44	46.3	46.2	33.3	39.1	9.6	37.4	39.0	31.4	22.9	19.3
Male	45-54	20.7	32.0	30.3	20.9	16.5	22.1	27.4	22.1	20.6	13.4
Male	55-64	21.3	20.4	27.1	22.8	19.9	18.5	20.0	27.2	19.5	17.2
Male	65+	12.7	14.1	17.8	20.6	12.3	10.4	7.0	20.2	17.7	14.1

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 relative value received by a Facebook user of the demographic the arrow is pointing towards from an additional Facebook user of the source demographic (i.e. the product of $z_i(j)$ and $U_i(j)$). 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. This is due in part due to older users having high valuations for connections to relatively non-abundant Facebook participants. The high value that older users glean from younger users is even more clear when restricting attention to the ten most valuable network effects, as figure 5 does. Appendix figures 19 through 22 restrict attention to the network effects experienced by and caused by other nodes of interest, displaying the rich heterogeneity of externalities on Facebook.

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

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 an additional Facebook user of the source demographic (i.e. $z_i(j)U_i(j)$)

.

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 an additional Facebook user of the source demographic (i.e. $z_i(j)U_i(j)$). Only the ten edges with the largest network externalities displayed.

platform for a demographic group after cascade t is

$$P_{i,t} = \Omega_i(\sum_{j}^{J} U_i(j)z_i(j)D_jP_{j,t-1} - a_i\phi_{i,t}) \tag{25}$$

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

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

$$\int_0^{P_{i,t}} \left((\mu_i(\vec{P}_{j,t-1}, \phi_i) - e_i(\rho_i))\right)d\rho_i \tag{26}$$

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.

$$e_i = -s_i log(\frac{1 - p_i}{p_i}) + \mu_i \tag{27}$$

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 cascades is

$$\Phi_{i,t} = \phi_{i,t}D_iP_{i,t} \tag{28}$$

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. We then calculate, taking into account where appropriate Facebook's optimal response, the consumer welfare and Facebook revenue consequences of three tax and redistributive policies and three regulatory policies.

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

	Cascade 1	Cascade 2	Cascade 3	Cascade 10	Cascade 20	Cascade 1000
Women 25-34	-0.94	-1.18	-1.25	-1.27	-1.27	-1.27
Women 35-44	-1.17	-1.47	-1.56	-1.59	-1.59	-1.59
Women 45-54	-0.50	-0.68	-0.73	-0.75	-0.75	-0.75
Women 55-64	-0.67	-1.03	-1.14	-1.18	-1.18	-1.18
Women 65+	-0.60	-1.17	-1.34	-1.41	-1.41	-1.41
Men 25-34	-0.29	-0.46	-0.51	-0.52	-0.52	-0.52
Men 35-44	-0.24	-0.37	-0.41	-0.43	-0.43	-0.43
Men 45-54	-0.21	-0.44	-0.51	-0.54	-0.54	-0.54
Men 55-64	-0.21	-0.47	-0.54	-0.57	-0.57	-0.57
Men 65+	-0.45	-0.87	-1.00	-1.05	-1.05	-1.05
Total CW Change	-5.28	-8.15	-8.99	-9.30	-9.31	-9.31
Profit Change	3.60	2.90	2.72	2.65	2.65	2.65

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.

entails a large increase in the level of monetization. Therefore for this analysis we assume that the marginal disutility from increased monetization a_i is equal to 1 for each group.[16]

We find that Facebook's profit maximizing strategy entails increasing fees substantially across all groups. Both on a per-user level and in aggregate, the negative incidence of price increases falls mostly on women. This is because they demand Facebook more inelastically, and because they provide lower positive network externalities on average. They also currently provide less advertising revenue on average. These factors combine to make them relatively attractive targets for increased monetization. Figure 6 displays the change in Facebook ad revenues and consumer welfare after N cascades in billions of dollars per month.

Implementing this strategy would increase Facebook revenues by 2.65 billion dollars per month (from a baseline of 1.57 billion) at the cost of decreasing its user base by 55.6% and lowering consumer surplus by 82.8% (from a baseline of 11.2 billion). In other words, this strategy entails squeezing Facebook's most inelastic users for a much higher share of their surplus. If we restrict attention only to non-price discriminating strategies (i.e. requiring Facebook to raise ϕ_i for all groups proportionately), we find that the profit maximizing flat increase in ϕ is $47.48 a month, which leads to an increase in Facebook revenues of only 2.078 billion dollars a month, with even larger associated decreases in consumer surplus. This last result gives a sense of the importance of price discrimination to Facebook profit maximization, even when applied to potentially less important market segments such as age and gender.[17]

---

[16]$a_i = 1$ 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.

[17]Likely, even more gains from price discrimination could be achieved by price discriminating over more

Why do our results imply that Facebook is leaving so much money on the table? There are two possible sets of answers. The first set of answers 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), or because it deters the entry of competitors. A second set of possibilities is that this is due to our model missing something important. For example, by taking into account only US users over age 25, we are potentially missing important network spillovers to and from other users of Facebook. If US users provide lots of value to users abroad, then it makes less sense to monetize them so intensely. Another possibility is that our surveys are soliciting a short-term demand elasticity for Facebook, whereas long-term demand for Facebook is more elastic.

6.2 The Impact of Tax and Redistribution Policies on Facebook Revenues and Social Welfare

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 a 10 dollar per user tax.[18] 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 9. 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

---

important attributes for demand like occupation or income.

[18]Professor Grauwe proposes this as an annual levy, but we consider a per-month tax.

	Initial	$10 Per User Tax	3% Ad Revenue Tax	"Data as Labor"
Net Ad Revenue (millions)	$ 1,565.7	-58%	-30%	-100%
Consumer Suplus (millions)	$ 11,225.5	-3%	2%	24%
Number of Users (millions)	134.741	-2%	1%	13%
Addl Tax Revenue (millions)	$ -	$ 1,326.1	$ 32.7	$ -

Figure 7: Estimated Facebook revenue, usage and consumer surplus changes as a result of three different redistributive reforms: a three percent tax on ad revenues, a ten dollar per a month per user tax, and a rebate of ad revenue to users in the spirit of Posner and Weyl's 'Data as Labor' proposal. The first two policies assume Facebook responds optimally to the tax change. The final proposal assumes advertising levels for all demographics are kept fixed at their current level.

revenues is fully incident on the platform's profits.[19] 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.

Figure 8 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, while raising a large amount of revenues. On the other hand, a 3 percent tax slightly boosts consumer surplus and participation rates. However, it does not raise much revenue, and it has a disproportionate 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.

6.3 The Impact of Regulations on Facebook Revenues and Social Welfare

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

---

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

study.

Here we consider a set of three potential reforms, two assumed to be implemented perfectly, and one a 'worst case scenario' for a botched reform. The two positive reforms are a move to increase the competitiveness of social media and nationalization by a benevolent social planner.

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.

Increasing competition might be bad for a few reasons. 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.

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 communiques 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.

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).

	Initial	Perfect Competition	Facebook Breakup	Welfare Maximization
Net Ad Revenue (millions)	$ 1,565.7	-100%	-93%	-255%
Consumer Suplus (millions)	$ 11,225.5	9%	-44%	42%
Number of Users (millions)	134.741	5%	-18%	20%
Change in Non-Rev Value (millions)	$ -	-	-$8,155	$ 5,019

Figure 8: Estimated Facebook revenue, usage and consumer surplus changes as a result of three different regulatory reforms: mandating interoperability and lowering barriers to entry to create perfect competition, a botched Facebook breakup that creates two uncompetitive monopolies for half of the US, and nationalization by a benevolent social planner.

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 such a 'worst case scenario' Facebook breakup as the creation of two Facebook monopolies each serving half of the US population.[20]

The final scenario we simulate is one in which a benevolent social planner takes over Facebook, and runs it to maximize social welfare. Such a planner would internalize all network effects. Here, we also model the planner as taking into account the platforms' desire to have a large user base (as this might represent the future value of data collection). This simulation entails the platform running 'negative advertisements' (i.e. expending money to boost the welfare of users on the platform). [21]

Results from these three simulations are summarized in figure 8. We find that perfect competition would raise consumer surplus by 9%, at the cost of eliminating all monetary profits. Taking only Facebook's monetary revenues into account, perfect competition actually lowers social surplus, 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. A social welfare maximizing Facebook would raise consumer surplus by 42%, at the cost of Facebook needing to go -255% into the red. However, breaking Facebook into two non-competitive 'baby Facebooks' would be disastrous, lowering consumer surplus by 44%. It would also lower combined ad revenues by 93% as the baby Facebooks lowered advertising rates to retain even 82% of their original combined user base.

---

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

[21]We assume that these negative advertisements symetrically create utility at the rate a_i. If we assume that $a_i \geq 1$, then the social welfare optimum comes at a transfer of negative infinity.

7   Conclusion and Managerial Implications

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. The specific example we emphasize is a firm which can discriminate in its advertising to profit maximize. 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 decree 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 use this model to estimate, in the case of Facebook, the revenue and welfare consequences of different pricing strategies, taxes, regulations and market structures. As far as we know this is the first paper to produce such predictions. Hopefully these findings will be useful in guiding policy makers, and will serve as one approach among many for projecting the impact of policy changes.

That being said, 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 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. However, the implications of such a scenario could be estimated in an extension of the model. In

any case an intriguing area for future investigation is to actually conduct experiments on platforms to see how well real world phenomena match our predictions.

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.

Hagiu, 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.

Figure 9: 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.

A   Additional Tables and Figures

Figure 10: 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 11: 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 12: 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 13: 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.

33

Figure 14: 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 15: 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 16: 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 17: 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 18: Underlying data and estimate of the the demand curve (Ω_i) for men age 65 or older. The points are the mean response to the question "Would you give up Facebook for 1 month in exchange for \$X? Choose Yes if you do not use Facebook." for individuals of the group. Confidence intervals are based on binomial statistics. The curve, in green, is the logistic line of best fit.
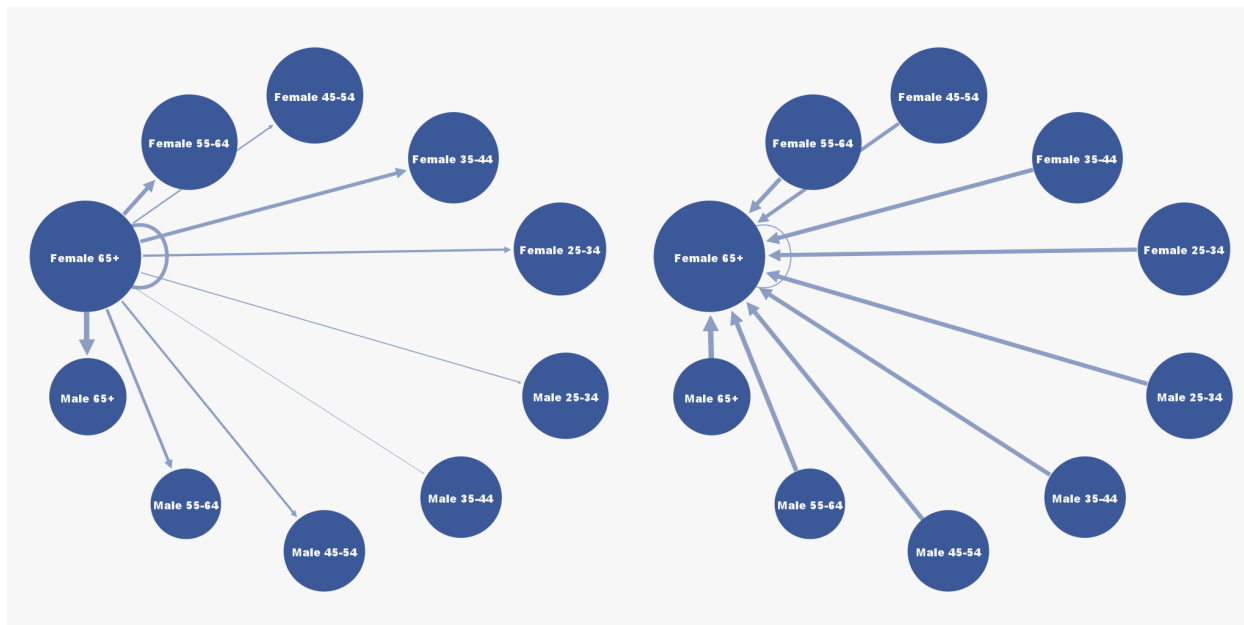

Figure 19: A graphical representation of Facebook usage and network effects. Relative value of Females 65+ to users of different demographics displayed.

Figure 20: A graphical representation of Facebook usage and network effects. Relative value of other users to Females 65+ displayed.

Figure 21: A graphical representation of Facebook usage and network effects. Relative value of Males age 25-34 to users of different demographics displayed.

Figure 22: A graphical representation of Facebook usage and network effects. Relative value of other users to Males age 25-34 displayed.

Intercept	Coefficient on Cost	Demo Group
.578603	.0130975	Female 25-34
.753299	.0111212	Female 35-44
.856898	.0243794	Female 45-54
.798354	.0185573	Female 55-64
.570240	.0221589	Female 65+
.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 1: 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.[22] 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

$$P = \Omega\Big(\sum^{I} U(i)P - a\phi \Big) \tag{29}$$

where $U(i)$ is the value of any connection.[23] Then the best response function is

$$\frac{\partial \Omega}{\partial P} = \frac{\partial \Omega}{\partial \boldsymbol{U}} \Big(\sum^{I} U(i) - a\phi \Big) \tag{30}$$

And so a network equilibrium is stable so long as

$$1 > \frac{\partial \Omega}{\partial \boldsymbol{U}} U(i)(I-1) \tag{31}$$

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 $\frac{P_i}{\phi_j}$ are finite*

---

[22]This 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.

[23]the value of a 'connection to oneself' is assumed to be 0

Proof. Rewriting equation 15 with the assumption all nodes are identical, before i gets hit with a fee, yields:

$$\frac{\partial P_j}{\partial \phi_i} = \frac{\partial \Omega}{\partial \boldsymbol{U}}\left(U(i)\frac{\partial P_i}{\partial \phi_i} + \sum_{k \neq i}^{K} U(i)\frac{\partial P_k}{\partial \phi_i}\right) \tag{32}$$

Substituting in 14 and summing yields

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

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

$$\frac{\partial P_j}{\partial \phi_i} = \frac{-\left(\frac{\partial \Omega}{\partial \boldsymbol{U}}\right)^2 U(i)\frac{\partial A}{\partial \phi_i}}{1 - \left(\frac{\partial \Omega}{\partial \boldsymbol{U}}U(i)(I-2) + \left(\frac{\partial \Omega}{\partial \boldsymbol{U}}\right)^2 \left(U(i)\right)^2(I-1)\right)} \tag{34}$$

The network will not unravel due to a welfare change so long as 34 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

$$1 - \frac{\partial \Omega}{\partial \boldsymbol{U}}U(i)\left((I-2) + \left(\frac{\partial \Omega}{\partial \boldsymbol{U}}\right)\left(U(i)\right)(I-1)\right) \tag{35}$$

From the assumption that the network is stable, we have

$$1 > \frac{\partial \Omega}{\partial \boldsymbol{U}}U(i)(I-1) \tag{36}$$

This implies

$$I - 1 > (I-2) + \frac{\partial \Omega}{\partial \boldsymbol{U}}U(i)(I-1) \tag{37}$$

and applying B again implies

$$1 > \frac{\partial \Omega}{\partial \boldsymbol{U}}U(i)\left((I-2) + \left(\frac{\partial \Omega}{\partial \boldsymbol{U}}\right)\left(U(i)\right)(I-1)\right) \tag{38}$$

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.

$\square$

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

Proof. Directly from (34) $\square$

C Social Welfare Maximization

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

$$W_i = P_i E[\mu_i(\vec{P}, \phi_i) - \epsilon_i | U_i > \epsilon_i] \tag{39}$$

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

$$W_i = P_i \int_{-\infty}^{U_i} \frac{\mu_i(\vec{P}, \phi_i) - \epsilon_i}{Prob(U_i > \epsilon_i)} f(\epsilon_i) d\epsilon_i \tag{40}$$

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

$$W_i = P_i \frac{\mu_i(\vec{P}, \phi_i) F(\epsilon_i)}{Prob(U_i > \epsilon_i)} \Big|_{-\infty}^{U_i} - P_i \int_{-\infty}^{U_i} \frac{\epsilon_i}{Prob(U_i > \epsilon_i)} f(\epsilon_i) d\epsilon_i \tag{41}$$

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

$$W_i = P_i \mu_i(\vec{P}, \phi_i) - \int_{-\infty}^{U_i} \epsilon_i f(\epsilon_i) d\epsilon_i \tag{42}$$

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 maximizaiton entails

$$\max_{\phi_i} \sum_i^I [P_i(\phi_i + \mu_i(\vec{P}, \phi_i)) - Q_i(\epsilon_i) \Big|_{-\infty}^{U_i}] - F \tag{43}$$

$$\text{s.t.}$$

$$P_i = \Omega_i(\boldsymbol{U_i}) \tag{44}$$

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.