Insulated Platform Competition*

Alexander White† E. Glen Weyl‡

April 2016

Abstract

Does competition promote efficient platform pricing and market structures? We model imperfect platform competition using a new approach, *Insulated Equilibrium* (IE). Building on the observation that platforms often charge low prices to build a "critical mass", IE assumes platforms use "penetration pricing" to simplify the role of user beliefs. We show that competition's impact on efficiency depends crucially on heterogeneity in users' valuations for network effects. The standard Flat Pricing approach (FP) cannot tractably incorporate such heterogeneity. We show that IE's sharp, general predictions are economically similar to FP's, when FP is tractable, and we show how the nature of product differentiation determines the welfare consequences of competition in calibrated models of the video game and newspaper industries.

Keywords: platform competition, insulating tariffs, Spence distortion, user coordination, penetration pricing, excessive entry

JEL Codes: D21, D43, D85, L13

*We are grateful to many people for helpful comments and discussions, especially Simon Anderson, Eduardo Azevedo, Jacques Crémer, Bruno Jullien, Jon Levin, Ariel Pakes, Phil Reny, Bill Rogerson, Larry Samuelson, Jean Tirole and Julian Wright as well as Rabah Amir, Joyee Deb, Joe Farrell, Renato Gomes, Chiu Yu Ko, John Morgan and In-Uck Park for their conference discussions. Tesary Lin, Fengshi Niu, Toan Phan, Rui Wang, Will Weingarten, Stephen Wolff, Anthony Zhang and especially Yucheng Liang provided superb research assistance. White gratefully acknowledges the NET Institute (www.NETinst.org) and National Natural Science Foundation of China Project 71203113 and Weyl the Ewing Marion Kauffman Foundation, Nokia, Microsoft Corporation and the Alfred P. Sloan foundation for their financial support. Prior circulating drafts of this work, which is based on a chapter of White's Ph.D. dissertation at Toulouse School of Economics, were entitled 'Imperfect Platform Competition: A General Framework'. All errors are our own.

†School of Economics and Management and National Institute for Fiscal Studies, Tsinghua University, Weilun Building, 100084 Beijing, China: awhite@sem.tsinghua.edu.cn

‡Microsoft Research New England, One Memorial Drive, Cambridge, MA 02142 USA and Department of Economics, University of Chicago: glenweyl@microsoft.com.

1 Introduction

When it first enters new cities, Uber offers discounted prices to riders while paying drivers the full fare, helping to spur both riders and drivers to join its service. Other platforms' strategies mirror this approach, subsidizing usage during the "start-up" phase to overcome the "chicken-and-egg" problem whereby each side of the market is willing to participate only if it expects the other side to as well. The literature on dynamic platform competition often predicts some form of this "penetration pricing". However, the static literature on platforms used in policy-oriented work has, for tractability, largely assumed pricing to be independent of participation. In this paper, we propose a method for incorporating the most prominent features of penetration pricing into a static model of platform competition amenable to policy analysis.

Surprisingly, explicitly modeling such pricing strategies makes general analysis of the basic environment of interest simpler. This allows us to study policy-relevant cases that have not been tackled using the standard approach. This simplicity stems from the fact that, when platforms adopt such strategies, they do not need to worry about user (mis-) coordination, removing a complication from the model. Consequently, our approach sheds new light on the important question of the extent to which platforms efficiently provide users with network effects and how the competitive environment influences this. In a monopoly model, Weyl (2010) showed that heterogeneity in users' values for network effects is the key driver leading to distortion in their provision. Thus, in studying the provision of network effects under competition, allowing for rich user heterogeneity is crucial to reaching plausible policy conclusions. A central motivation for our work is to provide a framework that accommodates these realistic distortions while remaining tractable for policy analysis.

To illustrate the conclusions such an analysis could yield, we calibrate our model to two of the platform industries more widely studied in the empirical literature: video games and newspapers. In the former, we find competition can substantially improve the availability of games because consoles are differentiated by their inherent features, such as the controllers they use. Given this form of differentiation, competition, which leads to greater attention to "switching" users who are indifferent between the platforms, draws firms to cater more closely to inframarginal consumers' strong taste for greater game availability. In the latter, on the other hand, we find that competition may in fact further distort the amount of advertisement, as both high- and low-brow papers compete for readers with moderate taste for advertisements and who poorly represent the average readership of either paper.

Our approach centers around a new solution concept for platform competition called *Insulated Equilibrium* (IE). It is an alternative to the "Flat Pricing" (FP) assumption typically used to refine the large multiplicity of Nash equilibria in prices that depend on the participation rates on the other side of the market, as proposed in an important special case by Armstrong (2006). This concept draws motivation from a large literature in which firms, principals, or designers favor mechanisms that improve the reliability or "robustness" with which they induce consumers or agents to take the desired action.

In particular, it builds on Dybvig and Spatt's (1983) argument that in designing the pricing of a public good with network externalities, the authority should try to minimize the need for users to coordinate with one another.[1] Such strategies are most effective (Green and Laffont, 1977; Wilson, 1987; Bergemann and Morris, 2005; Chung and Ely, 2007) when prices give all users dominant strategies independent of other users' behavior. Take the case of Uber again. When it first enters a city, it should set prices just aggressively enough to perfectly offset the inconvenience created by its initially thin set of drivers. Then, as it matures, its improving quality will perfectly offset its rising prices as more drivers enter the market, insulating its user base. We thus refer to such a pricing strategy as an *insulating tariff*.

However, such perfect compensation is possible if and only if all users place the same value on network effects, so that the rate of compensation is the same for all users, as we assume in the next two sections. On the other hand, when users are heterogeneous in their valuation of network effects, the best a platform can do is to insulate the representative user's (average marginal) valuation of network effects. Weyl (2010) first proposed such insulation, in the context of Rochet and Tirole's (2006) monopoly model, as an expositional tool.

In this paper, we introduce insulation as a strategic element, examining its methodological and substantive implications when platforms face imperfect competition. Compared to the standard approach of non-adjusting FP, insulation greatly reduces the complexity of each platform's problem by obviating the need to account for potential indirect effects of one platform's strategy on its profits through a change in user beliefs about the network effects it and its competitors provide. By ruling out these effects, we obtain a much richer characterization of leading aspects of platform competition of concern to policy-makers than was feasible in previous literature. We now outline our results.

First, in Section 2, we focus on the homogeneous case. With these preferences, IE implies that each platform literally implements in dominant strategies its optimal allocation, given

[1] See also Becker (1991), who makes a similar argument about the pricing strategies of trendy restaurants. Another intriguing resolution to these problems is proposed by Ambrus and Argenziano (2009), who assume users coordinate themselves in their collective interests. Our focus on firm coordination is motivated both by the difficulty we perceive in highly disparate users coordinating with each other, especially when they have conflicting interests because of horizontal differentiation of products which Ambrus and Argenziano abstract from, and by the resulting intractability of the Ambrus and Argenziano approach in models that allow for such differentiation along dimensions orthogonal to network effects.

the residual demand it faces. We derive a first-order condition for equilibrium prices under the restrictions that users have a common valuation of network effects and join at most one platform. This simple pricing formula intuitively embeds that of a traditional differentiated Bertrand model (as valuations for network effects go to zero) while adding the central platform intuition that the benefits derived by the other side of the market are passed through as a subsidy to the side of the market that creates this value. Using the example of logit demand, we compare IE with FP and a notion of Cournot competition extended to incorporate network effects, showing that only under IE does this simple logic of "Bertrand with network effect pass-through" apply.

We then give the first conditions, as far as we know, for the existence of pure-strategy equilibrium in platform competition applying to a broad range of value distributions, as in Caplin and Nalebuff (1991)'s analysis of standard imperfect competition. In addition to Caplin and Nalebuff's conditions, which limit demand convexity, in the platform context, existence of IE also requires sufficient differentiation of platforms relative to the strength of network effects. Armstrong (2006), in his two-sided Hotelling model with FP,[2] gives a condition of this latter form to ensure the existence and uniqueness of an equilibrium in which platforms *split* the market, rather one in which the market *tips* to a single dominant platform.

In Section 3, we further explore the comparison between IE and FP, specializing to the aforementioned model of Armstrong, but studying the previously intractable parameter values for which, instead of splitting, the market may tip in favor of a single dominant platform, both under IE and under FP. We find that, when comparable, IE and FP lead to qualitatively similar results. In particular, both give rise to excess entry compared to the social optimum, and under both, network effects alone cannot lead consumers to be "locked in" to an inefficient platform. Thus, using either approach, we find the main message of economic substance seems to be at odds with two commonly expressed views: (a) antitrust authorities should be especially vigilant in preventing one platform in a particular industry from becoming too dominant, and (b) as argued in classic works by David (1985) and Arthur (1989), such markets are highly path dependent, so technologically inferior firms can persistently stay on top.

However, we also find that, on a technical level, IE is much less demanding than FP.[3] Because, under IE, in the final stage of the game, all users have dominant strategies, standard backward induction allows us to solve for the full set of pure-strategy equilbria. By contrast, constructing equilibria under FP is much more complicated, because it requires considering all possible arrangements of user beliefs in all possible subgames. Thus, we

[2] See also Spulber (1999), Chapter 3, for an early pass at Hotelling competition among platforms.

[3] Note that, in this discussion, by "equilibrium under IE", or "equilibrium under FP", we formally mean a Nash equilibrium in pricing functions satisfying the respective refinement.

do not know of any systematic algorithm for characterizing the complete set of equilibria under FP. Therefore, we believe IE can be very useful in further exploring these complex questions regarding market structure.

Next, in Section 4, we consider the full model in which users heterogeneously value network effects. We state a general definition of IE amenable to such rich environments. Here, each firm, instead of necessarily giving a dominant strategy to each of its users, sets prices that adjust to hold fixed its aggregate demand from each group of users (as a best response, given its residual demand). At IE, it thus holds that the *representative consumer* of each group has a dominant strategy. We then derive a first-order condition for platform pricing, which exhibits the key point that platforms distort their network-effect provision passing through the preferences of marginal rather than average consumers. These distorted incentives are analogous to those for monopoly quality provision studied by Spence (1975) and Sheshinski (1976). Evidently, this *Spence distortion* emerges only when users are heterogeneous in their valuations of network effects, and thus the valuations of marginal and average users may differ. Unlike the standard approach, IE allows this issue to be tractably analyzed under competition.

In Section 5, we further pursue this analysis. In particular, we focus on a prominent special case studied by both Anderson and Coate (2005) and Armstrong (2006), which we generalize to allow heterogeneous valuation of network effects. We develop two stylized versions of this model, one tailored to the video gaming console industry and the other to newspapers and in both cases we study a leading policy question: the social value of competition. In the former model, the two gaming consoles differentiate themselves by offering different features from one another that are orthogonal to network effects, whereas, in the latter model, a "high-brow" newspaper differentiates itself from a "low-brow" tabloid by devoting less space to ads but charging readers a higher price. We then calibrate these models to match ballpark quantities reported by Lee (2013) and Gentzkow (2007), respectively, in their empirical work on these industries, and analyze the impact of increasing competition on the Spence distortion and other predictions of interests.

In the case of gaming consoles, "switching" users, indifferent between the two consoles but with potentially high willingness to pay for either, are more representative of average users on both platforms than are "exiting" users, who are indifferent between buying one of the consoles and buying neither. This differentiation pattern implies that enhancing competition by making switching easier mitigates the Spence distortion and improves social welfare. This effect is substantial in our calibrations.

On the other hand, in the newspaper context, switching readers have tastes in between those of the loyal readers of either paper, as in the analysis of Hotelling (1929). Thus, increased competition may distort the provision of network effects by exerting pressure for the papers to cater to middle-brow readers. However, because exiting users are also imperfectly representative of average users, these effects are somewhat mild and non monotonic throughout the full parameter range. Thus, in view of the fact that, apart from its impact on the distortion of network effects, competition also reduces platforms' market power, our calibrations suggest the conditions necessary for competition to be, on balance, welfare-diminishing are quite strong.

A principal appeal of IE as a solution concept is the fact that it opens the door to addressing issues such as these that are inherently important to platform competition but were previously intractable for applied analysis. To do so, it borrows the demand-stabilizing motivation behind penetration pricing, various forms of which appear commonly in platform industries[4] and which are predicted by many dynamic models of platform competition (e.g., Katz and Shapiro, 1986; Mitchell and Skrzypacz, 2006; Cabral, 2011*a*,*b*; Veiga, 2014). IE is also grounded in the theoretically attractive notion of implementation in dominant strategies, although this connection is stronger in the case of homogeneous valuations than it is under rich heterogeneity.

Because cases of rich heterogeneity are the most important to understand, and, so far, no other approach has been proposed that can address them, IE offers the natural benchmark for studying platform competition in these environments. Furthermore, in the simpler cases in which FP does yield tractable predictions, these predictions do not seem to be qualitatively at odds with those of IE. As we discuss in a separate legal policy paper (Weyl and White, 2014), which informally surveys both the broader literature on platforms and the results of this paper, IE serves to identify, clarify, and expand upon the seemingly complicated lessons of prior work on platforms for antitrust and regulation; it does not dramatically overturn them.

Thus, in terms of their substantive economic implications, IE and FP are broadly consistent. Nevertheless, a large grain of salt is still required before applying their conclusions to the real world. For example, in the two-sided Hotelling model, under both IE and FP, there is neither excessive concentration nor a failure of more efficient entrants to displace incumbents purely as a consequence of network effects. However, such issues could plausibly arise in richer environments, for example, with asymmetric information, even if existing models do not cleanly capture them. In our conclusion, we discuss some of the more promising directions research is taking to model information and expectations in a richer way that incorporates such effects and may indeed, as a result, lead insulation to be

[4]A recent, notable example occured in China in 2014, when the two major taxi-hailing app services, *Didi Dache* and *Kuaidi Dache* each charged negative per-ride prices to both riders and drivers over a nine-month period. Note, however, that stabler forms of contingent pricing, such as video games' per-unit royalties, and the Apple App Store's *ad valorem* taxation of transactions between users and developers are also broadly consistent with IE, as they make the total payment from a user to the platfrom contingent on other users' demand. To the extent platforms adopt such pricing arrangements, compared to, for instance, fixed yearly fees, IE seems to us more empirically plausible than FP.

undesirable for a platform. As such research proceeds, we believe IE to be more natural than FP as a static, complete information benchmark.

We collect extended proofs and details of our applications into an appendix following the main text. The most involved calculations are in an online appendix.[5]

2 Homogeneous Interaction Values

There are two *platforms*, indexed by j and k, each of which serves two *sides* of the market, indexed by A and B; whenever the corresponding results can be inferred by symmetry, we state only results for platform j, side A.[6] The set of users on each side of the market is continuous and of mass one. In this section, we assume users *single-home* in the sense that they demand at most one platform.

Users have heterogeneous standalone or *membership* values for each platform, denoted by $\theta = \left(\theta_j, \theta_k\right) \in \mathbb{R}^2$, and they perceive externalities from the presence of opposite-side users that join the same platform. Let $(n_A, n_B) = \left(\left(n_{j,A}, n_{k,A}\right), \left(n_{j,B}, n_{k,B}\right)\right) \in [0,1]^4$ denote (endogenized-below) fractions of users that join each platform on each side of the market, and let $p_A = \left(p_{j,A}, p_{k,A}\right) \in \mathbb{R}^2$ denote the platforms' side A prices, whose determination we discuss below. User θ on side A that joins platform j receives payoff

$$u_j = \theta_j + \gamma_A n_{j,B} - p_{j,A}.$$

γ_A denotes the side A *interaction* value. This parameter, common to all side A users and constant across platforms, measures the constant marginal utility such an agent receives from the presence of an additional side B user on the same platform. Membership values of side A users are distributed according to density function f_A, which we assume to be continuously differentiable and strictly positive on $\mathbb{R}^2$; that is it has full support. The payoff from the outside option is normalized to zero.

Platform j's profits are

$$\left(p_{j,A} - c_{j,A}\right) n_{j,A} + \left(p_{j,B} - c_{j,B}\right) n_{j,B},$$

where $c_{j,A}, c_{j,B} \geq 0$ denote j's constant marginal costs of serving users on each side.

[5]This is available at http://ssrn.com/abstract=2601836.

[6]In a previous draft we considered an arbitrary number of platforms. Doing so changes little qualitatively in our analysis, but requires significant additional notation. We therefore omit this extension here.

2.1 Equilibrium

The game has two stages. First, platforms simultaneously choose their strategies; second, users observe these strategies and choose which, if any, platform to join. Regarding the specification of platform strategies, Rochet and Tirole (2003) observed that platforms typically set prices that depend on participation rates on the other side of the market, such as royalties for video game developers and per-reader fees in newspapers; see also Hagiu (2006). As highlighted above, prices may also dynamically adjust to market conditions, as with penetration pricing.

To capture both of these phenomena, we follow Armstrong (2006) and we allow platforms to compete in functions of the participation rate on the other side of the market.[7] Specifically, let $\sigma_{j,A}(n_B)$, a mapping from $(0,1)^2$ to $\mathbb{R}$, denote platform j's pricing strategy on side A, and let $\sigma_A(n_B)$ denote a profile of platforms' side A strategies. Thus, given j's strategy and the reaction of side B users, $p_{j,A} = \sigma_{j,A}(n_B)$ is the realized side A price to join platform j. We now briefly motivate the particular strategies on which we focus and then formally define our solution concept, Insulated Equilibrium (IE).

Armstrong (2006, Proposition 3) shows that when platforms' strategies are arbitrary functions that depend on opposite-side participation, the Nash Equilibrium solution concept has no predictive power. As he points out, the logic from the well-known result of Klemperer and Meyer (1989) on multiplicity of supply function equilibria applies here: just as the different Cournot and Bertrand reactions from rivals induce dramatically different pricing incentives, so too can different reactions to opposite-side participation shape equilibrium prices. In fact, he shows that one can construct appropriately sloped functions to support any outcome as a Nash Equilibrium, provided it gives both platforms nonnegative profits.

Armstrong and most subsequent literature resolve this issue by assuming that platforms compete in *flat prices* that are independent of opposite-side demand.[8] However, if platforms follow such strategies, the continuation game users play may have multiple Nash Equilibria, because the attractiveness of each platform depends on opposite-side participation. Platform profits, in turn, depend on which of these equilibria prevails.

Our approach instead assumes platforms actively avoid staking their profits on "lucky" user coordination. It builds on Dybvig and Spatt's (1983) mechanism for efficient public

[7]Armstrong restricts attention to affine functions that depend only on a platform's own opposite-side participation level, and such functions suffice for insulation with the homogeneous linear interaction values we consider in this section. However, it is simpler to immediately allow the generality we will use in the next section here.

[8]Reisinger (2014) takes an alternative approach, using variability in the number of interactions each individual engages in to tie down the slope of tariffs as a price discriminatory device. Because such optimal discrimination has no reason to coincide with insulation, such tariffs create similar tractability issues to FP and thus much of our discussion of FP applies to them. We therefore do not discuss them in greater detail.

good provision and, more broadly, follows the large literature (Green and Laffont, 1977; Wilson, 1987; Bergemann and Morris, 2005; Chung and Ely, 2007) arguing that a mechanism designer should, where possible, avoid relying on details of agents' beliefs by making desired behavior a dominant strategy.

We formalize this preference for robust implementation by supposing platforms use the *insulating tariffs*, proposed by Weyl (2010) in the monopoly context, which ensure, whenever possible, that users have a dominant strategy.

Definition 1 (Insulated Equilibrium)**.** *In the special case of homogeneous interaction values, an* Insulated Equilibrium *is a Nash Equilibrium in which, after platforms have moved, each user has a dominant strategy.*

As we will see in Section 4, when users have heterogeneous interaction values, dominant strategy implementation is not feasible. There we state a general definition of insulation, appropriate for such heterogeneous environments, requiring that *representative* users have dominant strategies. In that context, insulation calls for each platform's side A price to be allowed to vary with the entire vector of opposite-side participation, n_B. However, restricted to our present, homogeneous interaction values setting, this more general definition is equivalent to dominant strategy implementation, and as we now show platform j's equilibrium strategy on side A, $\sigma_{j,A}$, varies only with $n_{j,B}$.

2.2 Insulating Tariff Systems

As a preliminary step, we now characterize the shape of the price functions that arise under IE. Observe, first, that regardless of its competitor's strategy, j can guarantee each user a fixed payoff from joining that is independent of the opposite-side users' actions. To make such a guarantee to side A users, platform j simply needs to set

$$\sigma_{j,A}(n_B) = t + \gamma_A n_{j,B}, \tag{1}$$

where t is an arbitrary constant. This strategy offers user θ on side A a payoff of $u_j = \theta_j - t$. Second, note that if j, alone, adopts such a pricing strategy, users do not necessarily have a dominant strategy, because their payoff from joining platform k may still vary. However, if k adopts such a strategy, then u_k for each user is fixed. Thus, when each platform adopts such a strategy in response to the other, the two platforms' strategies jointly guarantee every user a dominant strategy to join a particular platform or to join neither. We now define an *Insulating Tariff System* (ITS) in this environment.

Definition 2 (Insulating Tariff Systems)**.** *In the special case of homogeneous interaction values, an* Insulating Tariff System *on side A is a pair of price functions that jointly give all side A users a dominant strategy.*

To formally characterize an ITS, we first define the side A demand function, $N_A : \mathbb{R}^2 \times [0,1]^2 \to (0,1)^2$. The demand for platform j on side A is

$$N_{j,A}(p_A, n_B) = \int_{\{\theta: u_j \geq \max\{u_k, 0\}\}} f_A(\theta)\, d\theta.$$

Proposition 1 establishes the existence and uniqueness of the ITS that implements (gives rise to) a given demand profile.[9]

Proposition 1. *Consider n_A, p_A, and n_B such that $N_A(p_A, n_B) = n_A$. The unique ITS that implements n_A is $\sigma_{j,A}(x) = p_{j,A} + \gamma_A \left(x_j - n_{j,B}\right)$.*

Proof. For any given value of n_B, our full support assumption on membership values implies the distribution of $u = \left(u_j, u_k\right)$ has full support. This, combined with quasi-linearity of preferences, is sufficient to satisfy Azevedo et al. (2013)'s conditions for a unique equilibrium price vector to exist in the Walrasian market with aggregate supply n_A. This unique price is p_A in the theorem. Furthermore, note that the distribution of u is invariant in n_B given σ_A; thus, σ_A is the unique price vector supporting n_A for any value of n_B. □

Thus, all insulating tariffs by platform j on side A have the same variable component, $-\gamma_A n_{j,B}$, and differ only in their fixed terms. Let $t_A \in \mathbb{R}^2$ denote the vector of fixed terms, and let $\widetilde{N}_A : \mathbb{R}^2 \to (0,1)^2$ denote side A demand, *given insulation*, where

$$\widetilde{N}_{j,A}(t_A) = \int_{t_{j,A}}^{\infty} \int_{-\infty}^{\theta_j - t_{j,A} + t_{k,A}} f_A(\theta)\, d\theta_k d\theta_j.$$

Proposition 1 then implies that platform j's total side A price is $P_{j,A} \equiv t_{j,A} + \gamma_A \widetilde{N}_{j,B}(t_B)$.

2.3 Pricing

Before analyzing the predictions of IE, we first consider the benchmark of socially optimal pricing. When a side A user joins platform j, she imposes on j a direct cost of $c_{j,A}$, but she creates a benefit of γ_B for each of platform j's opposite-side users. Hence, a total

[9] Note that full support plays an important role in defining IE. If support were more restricted, more tariffs could exist that create dominant strategies for this restricted set of users. For example, if users had to choose one platform (the market was assumed covered), as in our example in the next section, then any pair of tariffs with $\sigma_{j,A}(x) - \sigma_{k,A}(x) = a + \gamma_A \left(x_j - x_k\right)$ would be insulating, because only the difference between the tariffs would influence user choices. In the next section we resolve this indeterminacy by focusing on the unique insulating tariff in the limit of a full support distribution approaching the fully covered market model.

surplus-maximizing planner sets (t_A, t_B) to satisfy[10]

$$P_{j,A} = c_{j,A} - \gamma_B N_{j,B}. \quad (2)$$

Intuitively, cost is marked-down by passing through to side A users the benefits $(\gamma_B N_{j,B})$ that their presence generates for side B.

Turning to equilibrium, let

$$\pi_j(t_A, t_B) = \left(t_{j,A} + \widetilde{N}_{j,B}(t_B)\gamma_A - c_{j,A}\right)\widetilde{N}_{j,A}(t_A) + \left(t_{j,B} + \widetilde{N}_{j,A}(t_A)\gamma_B - c_{j,B}\right)\widetilde{N}_{j,B}(t_B)$$

denote platform j's profits when both platforms use insulating strategies. Maximizing this with respect to $t_{j,A}$ gives the IE pricing formula, which we now state in Proposition 2, henceforth denoting $\mu_{j,A} = N_{j,A}/\left(-\frac{\partial N_{j,A}}{\partial p_{j,A}}\right)$ as platform j's (Bertrand) *market power* on side A.

Proposition 2. *At IE, platform j's side A price satisfies*

$$P_{j,A} = c_{j,A} + \frac{N_{j,A}}{-\frac{\partial N_{j,A}}{\partial p_{j,A}}} - \gamma_B N_{j,B}. \quad (3)$$

Proof. The FOC of platform j on side A is

$$\frac{\partial \pi_j}{\partial t_{j,A}} = \widetilde{N}_{j,A} + \left(t_{j,A} + \widetilde{N}_{j,B}(t_B)\gamma_A - c_{j,A}\right)\frac{\partial \widetilde{N}_{j,A}}{\partial t_{j,A}} + \frac{\partial \widetilde{N}_{j,A}}{\partial t_{j,A}}\gamma_B \widetilde{N}_{j,B}(t_B) = 0.$$

Dividing by $\frac{\partial \widetilde{N}_{j,A}}{\partial t_{j,A}}$ and using Proposition 1 we have Equation 3. □

Comparing Equation 3 with Equation 2 reveals that, at IE, only one source of distortion is present: the market power term, $N_{j,A}/\left(-\frac{\partial N_{j,A}}{\partial p_{j,A}}\right)$, familiar from conventional "one-sided" oligopoly competition. Importantly, this term is analogous to the markup term that would arise under differentiated Nash-in-prices competition, in a one-sided model because platform j regards $t_{k,A}$ as fixed. Thus, here, platforms have efficient incentives to internalize and pass-through the externalities created by users, because as platform j increases its side A demand, the willingness to pay of all side B users to join j increases at a rate of γ_B. Platform j thus fully internalizes the externalities through the insulating tariff and treats them exactly as if they were marginal cost economies of scope. The heterogeneous valuation of network effects that we allow in Section 4 makes such precise internalization impossible.

[10]Proposition 7 derives this result formally in a more general setting. Note that this is a necessary condition for social optimization only; it might be satisfied, say, at a market splitting outcome even if an outcome in which the market mostly tips is socially superior, though our full support assumption rules out full tipping being optimal. We examine the structure of the social optimum beyond first-order conditions (FOCs henceforth) in greater detail, but in a more special case, in Section 3 below.

Before considering such heterogeneity, it is useful to first examine, in the homogeneous case, the way different solution concepts predict that network effects will shape competition. Two natural alternatives to IE, which we now explore in detail, are flat pricing (FP), described above and what could be called "platform Cournot" (PC), in which platforms each commit to serving a fixed fraction of the population on each side of the market. The former is the most typical approach in the two-sided markets literature, particularly following Armstrong (2006), whose specific models we analyze below, whereas the multi-sided platform literature has seldom considered the latter.[11]

2.4 Example with Logit Demand

A concrete and canonical example useful in making this comparison is the "Logit" demand system studied extensively by Anderson et al. (1992) among others. Suppose side A users' membership values, θ, are independently and identically distributed across platforms according to the "Type I Extreme Value" distribution with dispersion parameter λ. This assumption implies

$$N_{j,A}\left(p_A, n_B\right) = \frac{\exp\left\{\left(\gamma_A n_{j,B} - p_{j,A}\right)/\lambda\right\}}{1 + \exp\left\{\left(\gamma_A n_{j,B} - p_{j,A}\right)/\lambda\right\} + \exp\left\{\left(\gamma_A n_{k,B} - p_{k,A}\right)/\lambda\right\}}.$$

Thus, the IE pricing formula on side A becomes

$$P_{j,A} = c_{j,A} + \frac{\lambda}{1 - N_{j,A}} - \gamma_B N_{j,B},$$

which differs from the standard Bertrand logit pricing formula only due to the presence of the final term (Anderson and de Palma, 1992; Berry, 1994).

Now consider flat pricing under the same distributional assumption. Furthermore, suppose that some selection rule guarantees a unique equilibrium in the second stage, for any prices that platforms could set. Under these conditions, the simplest way of which we are aware to express platforms' FOCs is Filistrucchi and Klein's (2013) and Song's (2015) technique involving both a system of FOCs of the form

$$\left(p_{j,A} - c_{j,A}\right)\frac{dN_{j,A}}{dp_{j,A}} + \left(p_{j,A} - c_{j,A}\right)\frac{dN_{j,B}}{dp_{j,A}} = 0,$$

and a set of equations that implicitly determine the relationship between the above total derivates of demand, with respect to price, and the underlying demand system. In general this does not lead to a simple pass-through representation of the impact of network effects,

[11] See Gabszewicz and Wauthy (2014) for a recent exception.

given that the full equilibrium must be accounted for. We will see a particular quantitative example of this violation of the pass-through principle in the next section.

To see why arriving at such a simple intuition under FP is difficult, suppose γ_A, $\gamma_B > 0$. Then, the FP solution concept implies that, when platform j lowers its side A price to attract more users on that side of the market, not only does this perturbation attract some side A users away from platform k, but it also triggers a feedback loop: the shift of demand toward j on side A causes side B users to switch from k to j, which further boosts j's side A demand while diminishing k's, and so on. This exploitation by j of k's passivity in pricing implies that through aggressive competition, each platform can gain great advantage by pushing its rivals into a downward spiral.

By contrast, under IE, when j lowers its price on side A, platform k's strategy calls for a compensating price decrease on side B that preserves the payoff of its side B users and thus its market share on that side. Consequently, the incentives j faces when lowering its price on side A boil down to the standard marginal versus infra-marginal tradeoff and the additional pricing power it obtains on side B, thanks to the increase in *its own* participation level on side A.

Finally consider the Platform Cournot (PC) solution concept, whereby each platform chooses a demand level on each side of the market, and then prices adjust so that markets clear. Here, platform j's side A price satisfies

$$P_{j,A} = c_{j,A} + \frac{\lambda}{1 - \frac{N_{j,A}}{1-N_{k,A}}} - \gamma_B N_{j,B}.$$

Like the IE formula, this expression differs from its one-sided analog only in the presence of the final term, $\gamma_B N_{j,B}$, reflecting the absence of a feedback loop. However, the relevant one-sided analog differs, as the mark-up term becomes $\lambda / \left(1 - N_{j,A} / \left(1 - N_{k,A}\right)\right)$; platform j takes platform k's A side quantity fixed, believing k will adjust its *side A* price to hold this quantity fixed, not merely that it will compensate side B users for their lost partners.[12] The applied literature usually considers such assumptions of Cournot conduct unrealistic. Thus, while both PC and IE cut off feedback loops between the two sides, IE does so in a

[12] Although we focus in this paper on the common case in which all network effects flow across the two sides of the platform, in a previous draft, we considered a model in which a platform might have just a single side within which network effects flow. Even there, however, PC and IE do not coincide; the analysis is essentially the same, except that the pass-through term is given on the basis of own-side participation, and thus the mark-up term in the two models continues to differ. Crucially, under IE, if platform j aggressively decides to drop its fixed price t_j, k compensates its users for changes in network effects resulting from lost participation but does not lower its fixed fee, t_k, as would be necessary to hold fixed participation on platform k constant. As a result, the one-sided platform FOCs take the same form as in two-sided markets in both cases, except that in the pass-through term the aggregate participation rather than the opposite-side participation is the relevant N value.

manner that allows flexibility in platforms' "within-side" conduct.[13]

2.5 Existence of IE

To formalize the tractability of IE, we give as far as we know the first conditions for the existence of pure-strategy equilibrium in platform competition that apply to a broad range of value distributions as in Caplin and Nalebuff's (1991) analysis of standard imperfect competition. Our approach builds closely on theirs. In addition to conditions similar to theirs on the concavity of individual product profit functions, we impose conditions that ensure that network effects do not lead to a violation of concavity through interactions across the two sides of the market. In particular, and closely following Armstrong's (2006) analysis for FP in the special case of a Hotelling model that we discuss in the next section, we restrict jointly, along with convexity in the value distribution, the minimum differentiation of platforms and the size of network effects.

Intuitively, as Becker (1991) argued for a single-sided platform, if demand on the two sides of the markets, multiplied by network effects, is too elastic, concavity of profits in quantity fails as aggregate demand (under FP) is backward-bending. By inducing more users to enter side A, a platform can sufficiently increase the number of users on side B to generate network effects that will allow the price on side A to be *raised* rather than lowered. Under IE, although such a feedback loop does not literally arise, by an analogous mechanism, network effects that are too strong, combined with insufficient demand concavity on each side and too much elasticity (a lack of what Armstrong termed "sufficient platform differentiation"), prevent profits on the two sides from being jointly concave in the vector of fixed prices $\left(t_{j,A}, t_{j,B}\right)$. We now formally state conditions that rule out such a phenomenon and are sufficient to establish existence.

Because the proof involves some detailed notation and calculations, we have relegated it to the appendix. But the basic approach is straightforward: we calculate the Hessian of each platform's profit function with respect to its two fixed price components. The original Caplin and Nalebuff conditions, embodied in the following assumption, ensure the diagonal elements of this matrix are strictly negative globally.

Assumption 1. *The joint distributions of side A and side B users' membership values,* $f_A\left(\theta_j, \theta_k\right)$ *and* $f_B\left(\theta_j, \theta_k\right)$*, are respectively* $\left(-\frac{a}{1+2a}\right)$*-concave and* $\left(-\frac{b}{1+2b}\right)$*-concave, where* $-\frac{1}{2} < a, b < 1$*.*

To ensure the determinant of the Hessian is globally positive requires a bound on the product of (a) the degree of concentration of values (the inverse of platform differentiation),

[13] Indeed, one can interpret the PC solution concept as a special case of a generalized notion of IE that incorporates a within-side conduct parameter. We follow Nevo (1998)'s suggestion in focusing on the most canonical specific models, such as Nash-in-prices, in order to avoid the indeterminacy and identification problems such conduct parameters may create.

(b) the square of the sum of network effects on the two sides of the market, and (c) a term monotone increasing in the convexity bound from on the two sides of the market from Assumption 1, as stated in the following assumption.

Assumption 2. *For all θ_j and θ_k, the conditional density functions $f_{A,j|k}\left(\theta_j|\theta_k\right)$ and $f_{A,k|j}\left(\theta_k|\theta_j\right)$ are upper bounded by $M_A \in \mathbb{R}_+$, and $f_{B,j|k}\left(\theta_j|\theta_k\right)$ and $f_{B,k|j}\left(\theta_k|\theta_j\right)$ are upper bounded by $M_B \in \mathbb{R}_+$. M_A and M_B satisfy $M_A M_B < \frac{(1-a)(1-b)}{4(\gamma_A+\gamma_B)^2}$, where a and b are defined in Assumption 1.*

Condition (b) is easy to interpret and condition (c) is standard from Caplin and Nalebuff. Condition (a) can be seen as a quantitative strengthening of our full support assumption and can be satisfied, for example, in most standard statistical discrete choice models by ensuring sufficient variance or other measures of dispersion, as such assumptions place a global bound on densities and conditional densities as required by Assumption 2.

Note that although we now show these assumptions are sufficient to establish the existence of a pure-strategy equilibrium, they are not necessary. In fact, they ensure the existence of a particular kind of equilibrium in which all platforms always have a globally quasi-concave profit function that is only useful in describing a limited set of circumstances in a platform industry. Typically, when network effects are strong enough relative to platform differentiation to make the market "tip" toward most individuals using a single dominant platform, profit functions will not be globally quasi-concave, because it will be optimal, socially and privately, for each platform either to be small or to dominate the market. In fact, in the special symmetric Hotelling case we consider in the next subsection, these conditions collapse to those guaranteeing the existence of an equilibrium in which the two platforms perfectly split the market. However, we are not aware of any general conditions guaranteeing the existence of pure-strategy equilibrium in imperfect competition that do not rely on global concavity of profits, and proposing such conditions is beyond the scope of this paper. We instead focus here on extending the Caplin and Nalebuff approach to the platform setting.

Theorem 1. *In the homogeneous interaction values model if Assumptions 1 and 2 hold there exists a (pure-strategy) Insulated Equilibrium.*

Because IE eliminates the indirect interactions between platforms through the feedback across sides of the market, it makes the analysis of equilibrium tightly analogous to that in a standard market.[14]

Proof. See Subappendix A.1. □

[14]It thus seems plausible, but well beyond our scope here, that Caplin and Nalebuff's conditions for equilibrium uniqueness, which involve strengthening their existence conditions, could be used to obtain conditions for equilibrium uniqueness here.

3 The Structure of Equilibrium under IE and FP

We now further illustrate IE and its relation to FP by comparing their predictions about the structure of equilibrium, in the context of Armstrong's (2006) canonical *two-sided single-homing model*.[15] Armstrong's model specializes the one introduced above in two ways. First, all users must join one platform or the other. Second, their membership preferences between the two platforms take the standard Hotelling (1929) form, with users uniformly distributed on the unit interval between the two platforms. Here, we further simplify Armstrong's setup by assuming symmetry across the two sides of the market. Thus, let γ denote all users' interaction value, and let τ denote the *transport cost* parameter. We begin by assuming platforms share a common marginal cost which we normalize to zero.[16] Whenever possible, we drop platform and side subscripts. The sometimes elaborate calculations necessary to derive the result below appear in Section 2 of our online appendix.

A key distinction we emphasize is between splitting and tipping equilibria. A splitting equilibrium is a symmetric equilibrium in which each platform serves half the users on each side of the market. A tipping equilibrium is an asymmetric equilibrium in which, on both sides of the market, the same one platform serves all users and the other serves none.[17]

3.1 Price Levels

In the literature on platform competition, some form of splitting has been the overwhelmingly focal case. Proposition 3 reports price levels that must prevail in any splitting equilibrium, under each solution concept.

Proposition 3. *(a) (Armstrong, 2006) Under FP, at any splitting equilibrium,* $P = \tau - \gamma$*.*

(b) Under IE, at any splitting equilibrium, $P = \tau - \frac{\gamma}{2}$*.*

[15]See Section 4 of that paper.

[16]This normalization is essentially without loss of generality because the full market coverage assumption implies common marginal costs are simply passed-through to users and cause no change in allocation. In addition to these substantive simplifications, this model, and the applications we consider in Section 5 below, are not literally special cases of the models they "specialize" in several ways including:

1. The model in this section and some elements in Section 5 specialize by collapsing heterogeneity to a single dimension in ways that violate our full support assumption and
2. In Section 5, to make calibrations more intuitive, we frequently assume a non-unit mass of users.

We do not separately develop general theory for these special cases here, because they are so tightly analogous to the general cases we developed in the previous section and develop in the next section. However, in some cases, the general results above do not literally apply in this case, as they rely on full support, the lack of which may lead to non-uniqueness of insulating tariffs. In these cases we import the sharper results from the general case, viewing the specialization as a limit of the general model.

[17]The following propositions ignore the knife-edge transition cases in the parameter space. As one would expect, at these transition points, the behaviors that occur on either side can both occur.

The reason prices are $\frac{\gamma}{2}$ lower under FP is the feedback loop discussed in Subsection 2.4 above: under FP, when a platform attracts a side A user from its rival, it receives twice the per-side B user benefit (2γ) that it does under IE because it simultaneously undermines its rival's offering. This benefit multiplies the number of users ($\frac{1}{2}$) it has on the opposite side of the market. Thus the formula under FP involves twice the intuitive pass-through effect.

Because this environment involves full market coverage, only the structure of the equilibrium, and not prevailing prices, affect social efficiency; thus, our investigation of structure below is of greater direct import than are prices, given splitting. However, a natural benchmark for pricing is the unique socially optimal prices in a model with market expansion that converges to the Hotelling limit, where efficiency requires that $P = -\frac{\gamma}{2}$. As we show below in Proposition 4, a necessary condition for splitting equilibrium to exist under both FP and IE is that $\gamma < \tau$. Therefore, for any comparable set of parameter values, FP prices under splitting are closer to being socially optimal. IE prices, on the other hand, have a greater structural similarity to socially optimal prices,[18] as they are marked up by the standard market power term, τ.

3.2 Tipping versus Splitting

The prices under splitting, in turn, have an important impact on when the market tips.

Proposition 4. *(a) It is socially optimal for the market to split if $\gamma < \frac{\tau}{2}$ and to tip if $\gamma > \frac{\tau}{2}$.*

(b) Under FP, if $\gamma < \tau$, then there exists a unique equilibrium with splitting; if $\gamma > \tau$, then there exist tipping equilibria and no splitting equilibria.

(c) Under IE, if $\gamma < \frac{3}{2}\tau$, there exists a unique equilibrium with splitting; if $\frac{3}{2}\tau < \gamma < 2\tau$, then there exist only a single splitting equilibrium and tipping equilibria; if $\gamma > 2\tau$, then there exist only tipping equilibria.

Note, first, that under both IE and FP, excess fragmentation can occur, but never excess tipping. This tendency is related to the tendency toward excessive entry discussed by Spence (1976) and Mankiw and Whinston (1986). A fragmenting entrant can profit from diverting sales away from a dominant platform, because the latter cannot simultaneously profit from its position and deter entry. Here, however, the social cost of entry is the dissolution of network effects rather than the duplication of fixed costs.

Because network effects influence marginal costs and not just fixed costs, excess entry may, unlike in the analysis of Mankiw and Whinston, lower *user* and not just social surplus. For example, when $\gamma = \frac{3}{2}\tau$, the lower bound of the region where both the tipping and splitting equilibria exist, under IE, user surplus under splitting is only $\frac{1}{4}\tau$, but under

[18] See the discussion following Proposition 2.

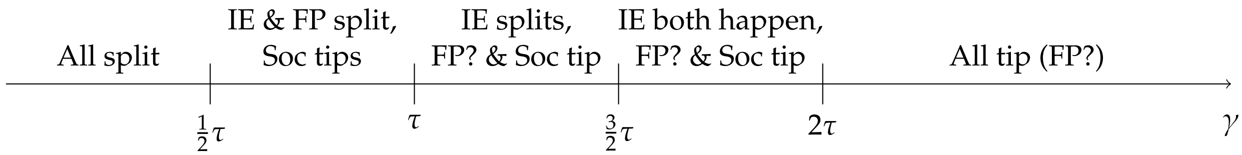

Figure 1: The structure of equilibrium for different parameter values under social optimization (Soc), Insulated Equilibrium (IE) and fixed pricing (FP). The "?"s after FP indicate that we cannot completely characterize FP equilibria in this case.

tipping it is twice as high at $\frac{1}{2}\tau$; when $\gamma = 2$, the upper bound of this region, user surplus depends on which of a continuum of tipping equilibria occurs. At the most competitive one, however, user surplus is again twice, under tipping, what it is under splitting.

In general settings, the above "business stealing" effect may be smaller or larger than the opposing force, stemming from the entrant's inability to capture infra-marginal surplus, that pushes towards insufficient entry. In this Hotelling setup, the infra-marginal user surplus from product differentiation is relatively small, and thus excessive entry occurs. Moreover, IE has greater scope than FP for excess fragmentation. Under FP, inefficient splitting occurs when $\gamma \in \left(\frac{1}{2}\tau, \tau\right)$, whereas, under IE, it always occurs when $\gamma \in \left(\frac{1}{2}\tau, \frac{3}{2}\tau\right)$, and it continues to be an equilibrium until γ reaches 2τ. IE's greater susceptibility than FP to inefficient fragmentation stems from the former's higher prices under splitting. These results are shown graphically in Figure 1.

On a technical level, our characterization of FP is much less complete than of IE; for $\gamma > \tau$, we know that no splitting equilibrium exists and that a tipping equilibrium does exist, but we cannot characterize whether other equilibria exist or how many tipping equilibria exist, though we believe a great many do. This type of concern motivated Ambrus and Argenziano (2009)'s analysis, which applies to the case in which there is no horizontal differentiation ($\tau = 0$), but provides a complete characterization of equilibria by assuming that users can act as coalitions that coordinate in their common interests.

By contrast under IE, and without further refinement of user behavior, we can fully characterize the equilibrium set in the cases considered throughout this section. The relative ease of characterizing IE versus FP stems from the fact that, unlike under IE, under FP, the set of equilibria depends sensitively on the nature of user beliefs, about which we have a great deal of flexibility, but of which careful track must be kept to establish the existence of an equilibrium.

Finally, note that the condition under which market splitting is the only equilibrium under IE, when $\gamma < \frac{3}{2}\tau$, is precisely the relevant special case of Assumption 2 that we used to ensure existence of IE in the previous subsection. In particular, in this context the maximum density of marginal users $M = \frac{1}{2\tau}$ on both sides of the market and, because no exiting margin exists, the factor of 4 dividing the right-hand side of the inequality in Assumption 2 drops out, which we discuss in detail in Subappendix A.1. Furthermore, $a = b = -1$ because demand is linear and thus weakly 1 concave. Thus, here, Assumption 2 simplifies to

$$\frac{1}{4\tau^2} < \frac{4}{(2\gamma)^2} \implies \gamma^2 < 4\tau^2 \implies \gamma < 2\tau,$$

that is, our conditions for the existence of a splitting equilibrium. This finding is consistent with our argument in the previous section that our conditions for existence of IE are actually conditions for the existence of a special type of equilibrium analogous to a splitting IE in the simple model we consider here.

3.3 Inefficient Lock-In?

A serious policy concern that motivated much of the literature on platforms is the risk, highlighted by David (1985) and Arthur (1989), that mis-coordination in favor of a dominant incumbent platform may prevent a more efficient competitor from entering. Given that these concerns are most preeminent in cases where tipping is assured, we focus now on the case in which $\gamma > 2\tau$ so that, under both IE and FP, tipping is guaranteed. We augment the model by now assuming one platform, the *entrant*, has Δc lower cost than the inefficient *incumbent* platform on both sides of the market.

Proposition 5. *Assume $\gamma > 2\tau$. Under both FP and IE, tipping occurs. Under both solution concepts, the market can tip to the incumbent only if $\Delta c \leq \tau$. Under IE, if $\Delta c \leq \tau$, then an equilibrium exists with tipping to the incumbent.*

Thus, inefficient tipping occurs not when network effects γ are sufficiently large, but rather when differentiation τ is sufficiently great. Inefficient tipping thus cannot be driven (solely) by network effects (it requires differentiation to be large) under either model and can occur only when the difference in efficiency between the two platforms is not too large. In the two cases, the mechanisms preventing inefficient tipping are different. Under IE, platforms' insulating tariffs eliminate user mis-coordination.

Under FP, an entrant can undermine mis-coordination by using a "divide-and-conquer" strategy (Jullien, 2011) of charging a low-enough price on one side of the market to ensure that for users on that side, joining it is a dominant strategy. It then offsets this subsidy by charging a correspondingly higher price on the other side of the market. Despite this difference, the main finding from both IE and FP appears to be that, at least in this stylized model, network effects cannot significantly drive inefficient tipping. Note, however, that matters would be very different in a one-sided platform where all network effects flow within a given side. In that case, there would be no way to divide and conquer; charging a low enough price to ensure coordination would also ensure that the entrant is forced to lose money in equilibrium. Consequently, in such a setting, inefficient tipping would be much more plausible under FP than IE.[19]

Thus, in this setting, IE and FP lead to moderately but not radically different conclusions. When the market splits, prices are higher and less responsive to the size of γ under IE than under FP, and this leads to a greater chance for inefficient fragmentation under the former. Because IE does not require keeping track of any user beliefs, it is straightforward to characterize, whereas FP is quite challenging. These results support our view that IE's conclusions are comparably, if not more realistic than those of FP. At the same time, they are easier to analyze, and have a more straightforward relationship to conventional models in industrial organization. These analytic advantages of IE are amplified in environments with rich user heterogeneity, to which we turn in the next two sections.

However, it is important to note that IE and FP share an assumption of complete information by platforms (and consumers) about market primitives that may be unrealistic and whose violation may make issues related to coordination and expectations important to platform competition, as suggested by David and Arthur. We return to these issues in our conclusion yet note here that, even when such issues are critical, a complete information benchmark is a useful point against which to judge the role of beliefs. We believe IE is a more useful such benchmark than FP, because it more explicitly embodies the complete information assumption, is more tractable, and does not leave out important coordination dynamics in FP that significantly change the equilibrium outcomes, to the extent that equilibrium outcomes in FP can be tractably analyzed.

4 The General Model

In the previous section, the pass-through term was always proportional to the homogeneous value all users had for network effects. Such homogeneity eliminates an important determinant of social efficiency Weyl (2010) highlighted in the monopoly context, namely *whose* preferences platforms internalize. In particular, he argues in the spirit of Spence and Sheshinski that monopoly platforms cater to the preferences of marginal rather than average users. In this section, we extend our model to allow for heterogeneous value for network effects and study the extent to competitive pressure mitigates this *Spence distortion*.

As before, two platforms serve users on two sides of the market. All assumptions from Section 2 remain unchanged, except for those pertaining to user preferences, which we

[19]Thus, if one believes that penetration pricing and other strategies for entering against an incumbent with network effects are relevant predictions, the value-added of IE may be even greater in a one-sided model. Conversely, it will lead to much different substantive conclusions in that context and thus would be unattractive if one believes the possibility of excessive standardization, as is possible in some equilibria of the one-sided, undifferentiated platforms analysis of Farrell and Saloner (1986), is important in explaining industry behavior.

now specify.

Users can choose to join any subset of platforms, $S \in 2^{\{j,k\}}$. The utility function of users on side A is now indexed by two parameters, a *membership value* $\theta \in \mathbb{R}^3$ and an *interaction value* $\gamma \in \mathbb{R}^L$:

$$\theta_S + v_A(S, n_B, \gamma) - \sum_{l \in S} p_{l,A}.$$

$\theta_\emptyset \equiv 0$ for all individuals, and θ_S is the membership value of individual of type θ joining the subset S of platforms. Denote the set of interaction values by $\Gamma = \mathbb{R}^L$, where $L \in \mathbb{N}_+$. The function $v_A : 2^{\{j,k\}} \times [0,1] \times \Gamma \to \mathbb{R}$ is thus a map to a user's willingness to pay for interactions from each possible consumption choice, the consumption pattern on side B, and the individual user's type.

The crucial substantive feature of this set-up is the heterogeneity it allows in users' interaction values, that is, their marginal valuations for additional interaction with side B users, through changes in n_B. These interaction values can differ across users and be non-linear. Side A users' types, (θ, γ), are distributed according to (smooth) density $f_A : \Gamma \times \mathbb{R}^3 \to \mathbb{R}$. Assumption 3 further characterizes side A users.

Assumption 3. *The functions v_A and f_A, jointly with their domains, have the following properties.*

(a) Full Support: *For all γ, the conditional distribution $f_A(\cdot|\gamma)$ of θ on $\mathbb{R}^3$ has full support.*[20]

(b) Smoothness: *$\forall \gamma$, $f_A(\cdot|\gamma)$ is twice continuously differentiable . $v_A(\cdot,\cdot,\cdot)$ is twice continuously differentiable in its second and third arguments with bounded derivatives.*

(c) No Externalities to Outsiders: *If $j \notin S$, then $v_A(S, n_B, \theta)$ is independent of $n_{j,B}$.*[21]

(d) Normalization: *For all $\gamma \in \Gamma$ and for all $n_B \in [0,1]^2$, $v_A(\emptyset, n_B, \gamma) = 0$.*

Lemma 1 characterizes side A demand, which we continue to denote by $N_A(p_A, n_B)$.

Lemma 1. *Side A demand, $N_A(p_A, n_B)$, exists and has the following properties.*

(a) Invertibility. *Holding fixed a particular n_B, for any interior side A consumption pattern, $n_A \in (0,1)^2$, there exists a unique price vector, p_A, such that $N_A(p_A, n_B) = n_A$.*

(b) Differentiability. *N_A is twice continuously differentiable.*

Proof. See Subappendix A.2. □

[20]This assumption is necessary to directly prove part (a) of Lemma 1 using the technique borrowed from Azevedo, Weyl and White (2013), which is useful for brevity. However, this assumption can be substantially relaxed (e.g., to one of convex support) by applying a more involved variant on the same proof. As this issue is orthogonal to the focus of this paper, we omit such a relaxation here.

[21]This assumption is not essential for our basic approach but is realistic in many industries and leads to sharper pricing results. Relaxing it could be an interesting avenue to studying "open" versus "closed" platforms as well as issues of externalities on nontraders, as studied by Segal (1999).

Thus, a side A inverse demand function, $P_A(n_A, n_B)$, exists and is twice continuously differentiable over the domain $(0,1)^2 \times [0,1]^2$ by the inverse function theorem.[22]

Finally, let $V_A(n_A, n_B) : (0,1)^2 \times [0,1]^2 \to \mathbb{R}$ denote gross user surplus on side A. Specifically, given an endowment, n_A, of available "slots" on each platform on side A and the aggregate behavior of side B users, summarized by n_B, $V_A(n_A, n_B)$ denotes the level of user surplus attained when side A users are optimally allocated to platforms.[23] Lemma 2 characterizes this function.

Lemma 2. *The side A gross user surplus, $V_A(n_A, n_B)$, exists and has the following properties.*

(a) Differentiability. *V_A is differentiable with respect to both arguments, $\frac{\partial V_A(n_A, n_B)}{\partial n_{j,A}} = P_{j,A}(n_A, n_B)$ and*

$$\frac{\partial V_A(n_A, n_B)}{\partial n_{j,B}} = \overline{\gamma}_{j,A} n_{j,A} \equiv \int_\gamma \int_{\theta : j \in \operatorname{argmax}_S v_A(S, n_B, \gamma) + \theta_S - \sum_{l \in S} P_{l,A}} \frac{\partial v_A}{\partial n_{j,B}}(n_A, n_B) f_A(\gamma, \theta) d\theta d\gamma.$$

We refer to $\overline{\gamma}_{j,B}$ as the average interaction value *of side B users on platform j.*

(b) Concavity. *V_A is strictly concave in n_A for any given n_B.*

Proof. See Subappendix A.2. □

4.1 Equilibrium

As described in Section 2, platforms move first, setting price functions. Users observe these price functions and then choose which platform(s) to join.

For the same reasons described in the previous sections, it is natural to suppose platforms seek to avoid mis-coordination by dominant strategy implementation. However, perfect dominant strategy implementation is typically impossible with heterogeneous interaction values. Suppose that there are two types of potential purchasers of a video game console. One, perhaps a family with young children, plans to use it occasionally for the basic games that come with the system and is close to indifferent between two major competing consoles; this family has no value for additional games beyond those that come with the console. The other is an avid gamer in his twenties, and availability of the latest games is the sole driver of his demand. He considers the games that come with the console too childish, and thus the console is less important to him than which games he can play. To

[22]This argument requires a slightly stronger condition, namely, that the Jacobian of $N_A(p_A, n_B)$ with respect to p_A is never singular. However, this condition is implied by full support using the same arguments we use in the proof of Lemma 1 and that Azevedo et al. (2013) use in their proof of invertibility. It also follows from the strict concavity of V_A shown in Lemma 2.

[23]Given quasi-linear utility, this is equivalent to the gross surplus on side A achieved at the unique Walrasian equilibrium with an aggregate supply of n_A slots on side A and with participation on side B exogenously set at n_B (Azevedo et al., 2013).

provide a dominant strategy to the second type of (marginal at a symmetric equilibrium) user, platforms would have to charge a tariff that rises steeply with the availability of games. But such a tariff would ensure that the first type of user would not have a dominant strategy; they would actually want to purchase the console with the fewest games, because it would be the cheapest.

Thus, dominant strategy implementation would require the platforms to offer a different price schedule for each user. While these schedules would agree at the equilibrium, they would require the platform to price discriminate in an individualized fashion if more games were suddenly available on one platform: the price to the avid gamer would have to rise, whereas the price for the family would have to stay flat. We know of no such discriminatory schemes in platform markets, and, presumably, if the platform could commit to such discrimination off the equilibrium path, it would also use it on the equilibrium path as well to raise profits, as in the classical analysis of third-degree monopoly price discrimination.

Therefore, rather than follow the motivation of our solution concept into this *reducio ad absurdum*, we relax our assumption about the platforms' ambitions. Just as platforms that cannot perfectly price discriminate choose prices to extract maximum surplus from the aggregate market, we assume platforms choose their price schedules to ensure a dominant strategy for the aggregate market. We thus define an aggregate *representative* user, in the spirit of Anderson et al. (1992), who can be given a dominant strategy.

Definition 3. *The side A* Representative User *chooses a vector of participation levels, n_A, so as to maximize $V_A(n_A, n_B) - n_A \cdot \sigma_A(n_B)$, where "$\cdot$" denotes the inner-product operator.*

We then define IE as follows.

Definition 4 (Insulated Equilibrium)**.** *An* Insulated Equilibrium *is a Nash Equilibrium of the overall game in which, after platforms have moved, each side's representative user has a dominant strategy.*

4.2 Insulating Tariff Systems

Definition 5. *An* Insulating Tariff System *(ITS) on side A is a pair of price functions that jointly give the side A representative user a dominant strategy.*

Note, first, that Definitions 4 and 5 are generalizations that subsume Definitions 1 and 2 in this more general setting. In the homogeneous environment of Section 2, the ITS as defined here necessarily gives all users a dominant strategy, because the utility of all users changes in the same way in reaction to changes in opposite-side participation.

Second, note that, at IE, each platform acts in a way that can most intuitively be described as follows.

1. Given the strategy of its competitor, platform j selects its profit maximizing demand vector, $\left(n_{j,A}, n_{j,B}\right)$.

2. Platform j selects the "residually insulating tariff", given its competitor's strategy, that robustly gives rise to $\left(n_{j,A}, n_{j,B}\right)$.

IE occurs precisely when both platforms behave this way as a best response to one another.

Regarding the shape of the ITS, note that even when each user has a constant marginal value for externalities, ITSs do not typically have a constant slope with respect to each platform's demand on the other side of the market when interaction values are heterogeneous. Returning to our stylized example above, the marginal gamer for a console with many games will typically be an avid gamer, whereas the marginal users of a platform with few games will typically be a family with small children. These marginal gamers are the ones the console platforms must compensate to give the representative gamer a dominant strategy.

Thus, as a platform's strategy changes, so too will its marginal users and thus the slope of its component of the ITS. However, the relative strategy of the two platforms will determine the identity of each platform's marginal users, because many of these marginal users are choosing between the two platforms. Consequently, each platform's component of the ITS will depend on the demand profile of all platforms and not just on its own user base on the other side of the market, as formally characterized in Proposition 6.[24,25]

[24]While, in practice, such "marketwide" contingent pricing is not typcially explicitly quoted, comparisons between platforms' relative network effects seem to play a (sometimes significant) role in the way their prices evolve, and, in this vein, we regard this form of ITS arising under rich user heterogeneity to be realistic. For instance, in a recent Amazon advertisement for the "fireTV stick", its competitor to the "Apple TV" platform, the former claims to offer "18x more apps and games than Apple TV", clearly referencing relative and not just absolute developer participation.

[25]Consider the following formal example illustrating the logic described above; side A represents developers and side B represents TV viewers. Platform j is Apple and k is Amazon. Viewers are single-homing. The type of a viewer is her interaction value γ, which has distribution F with full support on $\mathbb{R}$, and effort cost e of watching TV, which has distribution H on $[0, 1]$. The net utility a type (γ, e) viewer gets from joining Apple is

$$u_{j,B} = \gamma\left(n_{j,A} - e\right) - p_{j,B},$$

where $n_{j,A}$ and $p_{j,B}$ are, respectively, Apple's number of apps and the price it charges to viewers. Consider the limit case under our full support assumption in which viewers see platforms as differentiated only in their level of network effects.

Suppose that, in line with the example, $n_{k,A} > n_{j,A}$ and $p_{k,B} > p_{j,B}$. The interaction value of switching viewer is

$$\gamma^s = \frac{p_{k,B} - p_{j,B}}{n_{k,A} - n_{j,A}}.$$

All buyers who join Amazon have interaction values $\gamma \leq \gamma^s$, and all who join Apple have interaction values $\gamma \geq \gamma^s$.

Suppose an exogenous increase in apps were to occur for Apple. $p_{j,B}$ should respond to this increase with a slope smaller than γ^s, because the weighted average of γ of Apple's exiting margin and switching margin is smaller than γ^s. But then viewers would migrate from Amazon to Apple because the increase of $p_{j,B}$ does not fully offset the benefit of higher $n_{j,A}$ for the switching viewers. To keep $n_{k,B}$ fixed, Amazon must decrease $p_{k,B}$.

Proposition 6. *The unique ITS that implements n_A is characterized by the differential equation* $\left[\frac{\partial \sigma_A}{\partial n_B}\right] = \left[\frac{\partial N_A}{\partial p_A}\right]^{-1}\left[\frac{\partial N_A}{\partial n_B}\right]$*, and the boundary condition $\sigma_A(n_B) = P_A(n_A, n_B)$ for any fixed n_B.*

Proof. By definition, in an ITS, the side A representative user's choice of n_A does not depend on n_B. Therefore, the Jacobian of the ITS satisfies

$$\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} = \left[\frac{\partial N_A}{\partial n_B}\right] + \left[\frac{\partial N_A}{\partial p_A}\right]\left[\frac{\partial \sigma_A}{\partial n_B}\right] \quad \Leftrightarrow \quad \left[\frac{\partial \sigma_A}{\partial n_B}\right] = \left[-\frac{\partial N_A}{\partial p_A}\right]^{-1}\left[\frac{\partial N_A}{\partial n_B}\right]. \tag{4}$$

The boundary condition is obvious. □

The ITS holds fixed the aggregate demand on side A as the quantity on side B changes. Holding demand fixed in this way requires the vector of prices on side A to move in response to n_B so as to undo the direct impact of n_B on side A demand, N_A. If both demand profiles were scalars, this movement would just be the average marginal willingness of side A marginal users to pay for an additional side B user (i.e., the average marginal rate of substitution among marginal side A users between side B users and money). Because each side of the market has a full vector of demand, however, the slope of the ITS is a multidimensional generalization of this idea: the product of the inverse of the derivatives of N_A with respect to price, and the matrix of derivatives of N_A with respect to n_B. We will develop the interpretation of this matrix in greater detail using examples in the next section.

4.3 Pricing

We first consider, as a benchmark, prices that support a socially optimal allocation. Let $c_j = \left(c_{j,A}, c_{j,B}\right)$ and $n_j = \left(n_{j,A}, n_{j,B}\right)$ denote, respectively, the vectors of platform j's marginal costs and participation levels. Total surplus is maximized by solving

$$\max_{n_A, n_B} V_A(n_A, n_B) + V_B(n_B, n_A) - n_j \cdot c_j - n_k \cdot c_k. \tag{5}$$

Differentiating Expression 5 with respect to $n_{j,A}$ gives rise to Proposition 7.

Proposition 7. *At a socially optimal allocation, platform j's side A price satisfies*

$$P_{j,A} = c_{j,A} - \overline{\gamma}_{j,B} N_{j,B}.$$

This formula extends the Pigouvian logic of the Formula 2 for socially optimal pricing derived in Section 2 to account for heterogeneous interaction values. Rather than the number of users on the other side of the market being multiplied by their common interaction

Therefore, $p_{k,B}$ depends both on $n_{k,A}$ and on $n_{j,A}$.

value, the scaling factor is determined by averaging over the marginal utility for network effects of all the (optimally participating) users. As we will see shortly, platforms will not typically have the incentive to pass through the preferences of these average users.

Proposition 8 states the general IE pricing formula. Henceforth let $D_{jk,A} = \frac{\partial N_{k,A}}{\partial p_{j,A}} / \left(-\frac{\partial N_{j,A}}{\partial p_{j,A}}\right)$ denote platform j's *diversion ratio* on side A, that is, the fraction of side A users abandoning platform j, in response to an small increase in $p_{j,A}$, who turn to platform k.[26] Recall that $\mu_{j,A} = N_{j,A} / \left(-\frac{\partial N_{j,A}}{\partial p_{j,A}}\right)$.

Proposition 8. *At any IE, platform j's price on side A satisfies*

$$P_{j,A} = c_{j,A} + \mu_{j,A} - \left(\left[\frac{\partial N_B}{\partial p_B}\right]^{-1}\left[\frac{\partial N_B}{\partial n_A}\right]\right)_{j,\cdot}\begin{pmatrix}1\\ -D_{jk,A}\end{pmatrix} N_{j,B}. \tag{6}$$

Proof. Platform j's objective is to choose $n_{j,A}$ and $n_{j,B}$ to maximize its profit

$$P_{j,A}\left(n_A, n_B\right) n_{j,A} + P_{j,B}\left(n_B, n_A\right) n_{j,B} - n_j \cdot c_j,$$

given both platforms use insulating strategies. The FOC on side A is

$$P_{j,A} + n_{j,A}\frac{\partial P_{j,A}}{\partial n_{j,A}} - c_{j,A} + N_{j,B}\left(\frac{\partial \sigma_{j,B}}{\partial n_{j,A}} + \frac{\partial \sigma_{j,B}}{\partial n_{k,A}}\frac{\frac{\partial N_{k,A}}{\partial p_{j,A}}}{\frac{\partial N_{j,A}}{\partial p_{j,A}}}\right) = 0 \tag{7}$$

$$\Leftrightarrow P_{j,A} = c_{j,A} + \mu_{j,A} - \left[\frac{\partial \sigma_B}{\partial n_A}\right]_{j,\cdot}\begin{pmatrix}1\\ -D_{jk,A}\end{pmatrix} N_{j,B}.$$

Proposition 6 thus implies Equation 6. □

The new element in Equation 6 is the factor multiplying $N_{j,B}$ in the final pass-through term. We find this term most easily understood by analogy to the Spence-Sheshinski theory of optimal quality choice by a monopoly. Spence argues that monopolistic platforms distort the quality of their products relative to the social optimum by internalizing the preferences of the marginal rather than average user. Because, as Katz and Shapiro (1985) observe, the number of users determines the quality of a platform, this *Spence distortion* of the pass-through term from side B influences the price charged to side A. Weyl (2010) showed that, when side A users are heterogeneous in multiple dimensions and thus marginal users are heterogeneous, it is the average marginal user's preference that platforms internalize.

[26]Note that this diversion ratio need not be positive as we have not assumed that platforms must be substitutes for one another. The numerator of the diversion ratio is thus the net of the number of users that flow to or simultaneously leave platform k as they leave platform j. However, all of our applications concern weakly substitutable platforms, the title of our paper suggests, and thus our language below focuses on this case.

Proposition 8 takes this logic one step further by incorporating competition. With competition, a platform must account not just for the direct price impact of the users it attracts (the "1" in the top row of the column vector at the far right of Equation 6), but also for the price impact of the users the platform diverts from its rivals (the $-D_{jk,A}$ in the bottom row). These are both filtered through their price impact (that is, the willingness to pay of average marginal users) through the ITS described above. It is these preferences of marginal users, embedded in this matrix, that are passed through to the users on the other side of the market. We now make this abstract interpretation more concrete and draw sharper welfare implications from our analysis by considering more specific models where they can be signed, quantified, and calibrated to features of real industries.[27]

5 Platform Competition and Efficient Network-Effect Provision

In this section, we analyze the impact of competition on the Spence distortion and thus on the efficiency of network-effect provision and eventually on welfare. To investigate this issue, we focus on a prominent special case in the two-sided markets literature, studied by both Anderson and Coate (2005) and Armstrong (2006)[28] and labelled the *competitive bottlenecks model* by the latter. We develop two examples exhibiting, respectively, the potential benefits and harms of competition. In the first, calibrated to fit a stylized version of the video game industry, inherent features of gaming consoles drive differentiation – not differing availability of games – implying that competition improves welfare. In the second, calibrated to fit a local market for newspapers, the different amounts of space devoted to ads is an important driver of differentiation, implying that increasing competition may be harmful. We thus argue that the impact of competition is crucially related to the nature of product differentiation.[29]

[27]Note that even in the most general case, the condition suggests equilibrium is generically locally unique, because the system of equations defined by (6) will not typically have a singular Jacobian with respect to elements of the vector (n_A, n_B). Absent such a singular Jacboian, these conditions will not locally be satisfied at more than a single point.

However, we do not explicitly discuss issues of existence and uniqueness in the general case here, because, with rich heterogeneity, the formulas for and primitives determining the relevant second-order conditions are quite baroque. Instead, we focus on characterizing properties of and simulating equilibria based on the first-order necessary conditions in (6). This approach is common in the literature, following Berry et al. (1995), on standard, "non-platform" oligopolistic competition in which consumer preferences exhibit rich heterogeneity (i.e., "random coefficients"), which typically appeals to the results of Caplin and Nalebuff (1991), proved in a less heterogeneous environment. To provide additional assurance beyond what is standard in that literature, where failures of necessary conditions to be sufficient are sometimes a challenge, we check numerically that the points thus implied are global equilibria by simulating the profit of deviations away from this point.

[28]See Section 5 of both of these papers.

[29]Throughout the remainder of this section, we consider a series of "specializations" of the general model that in some instances are not strict special cases, as discussed in Section 3, but that obey the same logic as either limit cases or based on slight alterations of preceding logic.

5.1 The Competitive Bottlenecks Model

We now let A and B denote specific sides of the market. On side B, *buyers*, for example, gamers or readers, view the platforms as substitutes and join at most one. On side A, *providers*, e.g., game developers or advertisers, receive benefits that depend on the total number of users they gain access to via either platform. Because no buyer joins both platforms, providers view the decision of whether or not to join platform j independently from the analogous decision regarding k. Formally, a side B buyer of type $\left(\theta_j, \theta_k, \gamma\right)$ receives utility $\theta_j + \gamma n_{j,A} - p_{j,B}$ from joining only platform j and an arbitrarily large and negative payoff from joining both; that is, buyers have unit demand and thus "single-home". On side A, a provider of type $\theta = \left(\theta_j, \theta_k, \gamma\right)$ receives utility $\theta_j + \gamma n_{j,B} - p_{j,A}$ from joining only j, and the additive sum $\theta_j + \theta_k + \gamma \cdot \left(n_{j,B} + n_{k,B}\right) - p_{j,A} - p_{k,A}$ from joining both platforms, implying independence in their participation decisions across platforms.

The canonical analyses of this model by Anderson and Coate and Armstrong assume buyers have homogeneous interaction values. However, allowing buyer-side heterogeneity in this model is crucial both for realism, given strong empirical evidence that, for example, avid video gamers place much greater value than less enthusiastic gamers on the availability of a wide range of games, and for policy conclusions based on the model, as we analyze extensively in the next two subsections. However, once buyers are heterogeneous, an increase in the number of games available for platform j is *not* equivalent to a reduction in its price under FP, making it impossible to use the envelope conditions for platform j to define the appropriate amount to pass through to providers.[30] This makes analysis of FP with buyer-side heterogeneity in this model very challenging, and, thus, in what follows, we focus exclusively on IE.

Let $\widetilde{\gamma}_B^s = \mathbb{E}\left[\gamma | \theta_j + \gamma n_{j,A} - p_{B,j} = \theta_k + \gamma n_{k,A} - p_{B,k} \geq 0\right]$ denote the average interaction value of buyers along the *switching margin*, that is, the average among buyers wishing to join either platform but indifferent between the two. Let $\widetilde{\gamma}_{j,B}^{\emptyset} = \mathbb{E}\left[\gamma | 0 = \theta_j + \gamma n_{j,A} - p_{B,j} \geq \theta_k + \gamma n_{k,A} - p_{B,k}\right]$ denote the average interaction value of buyers along the exiting margin, that is, the average among those indifferent between joining j and joining neither platform, but who prefer platform j to platform k. We denote the *size* of these margins as $f_B^s = \frac{\partial N_{k,B}}{\partial p_{j,B}} = \frac{\partial N_{j,B}}{\partial p_{k,B}}$ (by Slut-

[30] To be more precise, after platform j increases its number of providers, thereby inducing under FP a reduction in buyer participation on platform k, the resulting decrease in providers on platform k does *not* have the same impact on platform j's profits (from buyers) that a change in the buyers' price that attracted in the same number of buyers to join platform j would have. The reason is that the initial increase in the number of platform j providers changes the composition of marginal buyers for platform j, who then react differently from the initial marginal buyers to the drop in platform k providers. In the gaming example, if platform j raises the games available, it will tend to selectively attract those most avid gamers, leaving its marginal gamers less avid and thus less attracted by the subsequent fall in games available for platform k, dampening (relative to a price reduction) the advantage platform j gains from stealing platform k buyers. This logic leads to a complex implicit definition of the pass-through term under FP.

sky symmetry) and $f^{\emptyset}_{j,B} = -\frac{\partial N_{j,B}}{\partial p_{j,B}} - f^s_B$, respectively. Finally, let $\omega_j = \left(1 + \frac{f^{\emptyset}_{j,B}}{f^s_B} + \frac{f^{\emptyset}_{j,B}}{f^{\emptyset}_{k,B}}\right)^{-1} \in (0,1)$ denote platform j's *margin weighting*, as it measures the relative thickness of j's switching and exiting margins. Note that ω_j increases in f^s_B, which rises if the platforms become closer substitutes, that is, as the market becomes more competitive. Proposition 9 characterizes pricing under IE.

Proposition 9. *At IE of the competitive bottlenecks model, platform j's price for providers satisfies*

$$P_{j,A} = c_{j,A} + \mu_{j,A} - \left(\omega_j \widetilde{\gamma}^s_B + \left(1 - \omega_j\right) \widetilde{\gamma}^{\emptyset}_{j,B}\right) N_{j,B}, \tag{8}$$

and its price for buyers satisfies

$$P_{j,B} = c_{j,B} + \mu_{j,B} - \widetilde{\gamma}_{j,A} N_{j,A}, \tag{9}$$

where $\widetilde{\gamma}_{j,A} = \mathbb{E}\left[\gamma | \theta_j + \gamma n_{j,B} = p_{j,A}\right]$ *denotes the average interaction value of j's* marginal *providers.*[31]

Proof. See Subappendix A.2. □

Because there is no competition for providers, any impact platform j increasing the number of buyers it attracts has on platform k's attractiveness to providers is irrelevant to platform j's profits. Thus, the amount passed through to buyers, internalizing value from providers, is exactly as in Weyl's monopoly analysis: platforms pass through the average interaction value of marginal providers $\widetilde{\gamma}_{j,A}$. The competitive bottlenecks model thus isolates competition, which was absent from his analysis, to the buyer side of the market and, therefore isolates its effects to the pass-through term to the provider side.

Furthermore and for the same reason, even in the buyer's side first-order condition that internalizes the value to providers of an additional buyer, the extra term in Proposition 8 involving the provider-side diversion ratio drops out. The amount passed through to providers is therefore now just a weighted average of the average value delivered to marginal switching and marginal exiting buyers on platform j, with the weight on the switchers increasing in the extent of competition. Thus, competition now impacts not only the market power that distorts buyer participation, but also the Spence distortion, because it affects the divergence between $\omega \widetilde{\gamma}^s_B + (1-\omega) \widetilde{\gamma}^{\emptyset}_{j,B}$ and $\overline{\gamma}_{j,B}$. This effect obviously can only be studied in a model with heterogeneous buyer interaction values; if γ is homogeneous, the pass-through term will always incorporate the homogeneous interaction value γ, which is why the Spence distortion is absent from the analyses of Anderson and Coate and Armstrong.

[31] Following Einav et al. (2010), this conditional (on a zero-measure set) expectation operator here is defined according to the standard average-marginal relationship as the derivative of $\mathbb{P}\left(\theta_j \geq p_{j,A} - \gamma n_{j,B}\right) \mathbb{E}\left[\gamma | \theta_j \geq p_{j,A} - \gamma n_{j,B}\right]$ with respect to $p_{j,A}$.

Although competition will unambiguously raise the extent to which the platforms attend to switching buyers (via ω_j), the social value of this effect is ambiguous. If switching buyers are more representative of average buyers than are exiting buyers, then increased catering to switchers is a good thing; but the reverse may be the case. Beyond even this ambiguity, the actual comparative static effects of competition are complex. Competition also directly influences buyer prices and, hence, the numbers and types of buyers participating in the market. Therefore it also indirectly affects the number of participating providers. Given these complications, we now turn away from further general analytic results and instead calibrate more specific models designed to illustrate these possibilities and match salient features of the video game and newspaper industries respectively.

The following stylized calibrations illustrate possible effects of competition on welfare. In line with this paper's focus, we attempt to fit the models' parameters in a way most relevant to determining the coarse-grained issues of aggregate network effect provision. In doing so, we abstract from many important features of the respective industries and do not, in any way, intend to make actual claims about welfare in these industries. Though most of our broad conclusions are robust over reasonable calibration ranges about the values we adopt, we selected these focal values both because of the sharp conclusions they allow and for their realism. Except where explicit closed forms are displayed or general results stated, we derive all results in these examples from numerical simulations.

5.2 Video Games and the Benefits of Competition

5.2.1 Model

The first example involves two symmetrically differentiated gaming platforms. Side A users are "developers" and side B users are "gamers". We follow Rochet and Stole (2002)'s analysis of competition in non-linear pricing in assuming gamers have heterogeneous vertical (interaction) values as well as horizontal (membership) tastes for one platform's console over the other that are distributed independently of one another. However, we slightly tweak the Hotelling preferences Rochet and Stole, Anderson and Coate, and Armstrong assume to allow a market that is not fully covered, while maintaining the same simplifying structure that Rochet and Stole exploit. Specifically, gamers are uniformly distributed over the unit interval, of which platforms $j \in \{0, 1\}$ are located at each end. If a gamer joins the platform closest to her, she incurs no transport cost, whereas, if she joins the platform that is further away, she incurs cost $\delta\tau$, where $\delta = 2\left(\left|j - x\right| - \frac{1}{2}\right)$ measures the *additional* distance required to travel to the less-favored platform and τ denotes the standard transport cost parameter.[32] Formally, a gamer of type (x, γ) who joins platform j receives utility

[32] For instance, a gamer located at $x = 0.1$ incurs zero transport cost to join platform 0 and, since $\delta = 0.9 - 0.1 = 0.8$, a transport cost of 0.8τ to join platform 1; a gamer located at $x = 0.9$ incurs zero transport cost to join

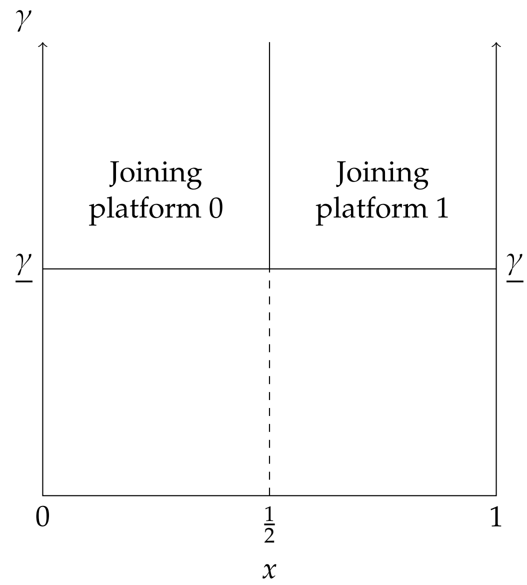

Figure 2: The type space at a symmetric equilibrium of our video game model.

$\gamma n_{j,A} - \max\{0, \delta\}\,\tau - p_{j,B}$.

Modifying Hotelling preferences in this way implies two useful simplifications at any symmetric equilibrium:

(1) As in Rochet and Stole's analysis, the average interaction value among switching gamers, $\widetilde{\gamma}_B^s$, is equal to the average interaction value among each platform's active gamers, $\overline{\gamma}_B$.

(2) In a natural extension of the Rochet and Stole set-up allowing for market expansion, exiting gamers share a common interaction value, $\widetilde{\gamma}_B^{\emptyset}$, that is also the minimum interaction value among purchasing gamers.

Figure 2 represents these "orthogonality" properties graphically. In this context, it is natural to expect competition to mitigate the Spence distortion, further increasing the benefits of competition, because competition increases the weight on the switching margin, which, unlike the exiting margin, is representative of average active gamers.

5.2.2 Calibration to the Sixth Generation of the US Video Game Industry

Where possible, we set magnitudes of parameters to match those reported or empirically estimated by Lee (2013), for the "sixth generation" of the video game industry in the United States, which took place between roughly 2000 and 2005. Though we continue to use a

platform 1 and 0.8τ to join platform 0; a gamer at $x = 0.5$ incurs zero transport cost to join either platform.

duopoly setup, three major gaming platforms in fact existed: Sony PlayStation 2, Microsoft XBox, and Nintendo GameCube. We make the following calibrations:

- Recruiting developers imposes no physical cost on the platforms.
- Consoles cost $400 to produce. This amount is in the ballpark of figures Lee (2012) and many media accounts report, though true costs are proprietary.
- There is a mass of 2,000 potential developers and 65 million potential gamers. The former rounds up to the first significant digit the number of software titles Lee observes, while the latter geometrically splits the households Lee observes owning consoles and the number of Nielsen households that he takes as the potential market, because not all households likeley to be potential gamers, even at very low prices.
- Developers have homogeneous interaction values of 50 cents per gamer. This homogeneity eliminates any Spence distortion among developers and is consistent with the roughly 10 games Lee finds a typical gamer buys and $50 per game (which is a bit higher than the $20-$30 he reports).
- Developers have heterogeneous fixed development costs that are log-normally distributed with a mean of about $10 million and a standard deviation of about $1 million. These amounts are generally consistent with press accounts, though they may be slightly overestimated.
- Gamers have interaction values that are log-normally distributed with a mean of approximately 50 cents per game available and a standard deviation of 25 cents. This implies an even split of per-game surplus between the two sides of the market, on average, and inequality in value across users, consistent with Lee's empirical conclusions, though the degree of inequality we assume is probably a bit low.
- We allow the platform differentiation parameter τ to range from $500 to $6500, because a τ value of $3500 generates platform profits of approximately $2.5 billion in IE, roughly consistent with the life-cycle profits reported in the press for the leading platforms.

5.2.3 Results

We focus on symmetric equilibria given the symmetric setup and solve for a point satisfying the necessary FOCs, which we calculate numerically using quadrature for all expectations and margins. We then search using Newton's method to solve for a symmetric solution, beginning at seed values calibrated to Lee's reported equilibrium prices and participation rates. Finally, we check that there are no profitable global deviations from our proposed

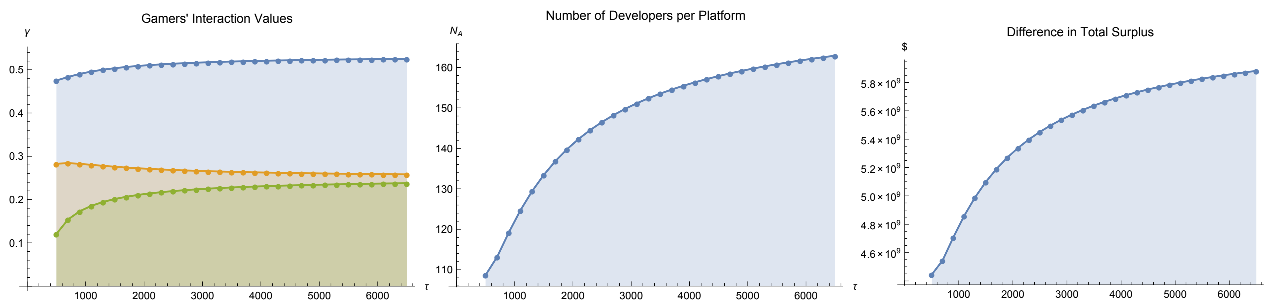

Figure 3: The difference between outcomes at IE and in a second-best scenario where the Spence distortion is absent, as a function of τ, the market power, in our calibrated video game model. On the left, the divergence between the average interaction values internalized by the platform (middle curve) and that internalized by the social planner (and thus in the second-best, as well as the average interaction value of switchers, pictured as the top curve), as well as that of exiting users (the bottom curve). In the center, the difference between the number of developers active. On the right, the difference in social surplus.

equilibrium by graphing the profit function of one platform holding fixed the other platform's equilibrium behavior. We discuss further details in Subappendix B.2.

The overall results we obtain are unsurprising: increasing τ lowers the equilibrium number of games created, gamers served and social welfare created. Of greater interest is what happens to the Spence distortion, which tends to lower the number of games available, because infra-marginal avid gamers value games more than marginal gamers under our modeling assumptions.

To study this, we compare the equilibrium we obtain to one where platforms mechanically use the average gamer value (multiplied by the number of gamers) to mark down the price they charged developers, an outcome we call the (no Spence distortion) *second-best*. This follows the approach of Ordover and Saloner (1989) and Besanko et al. (2014), who use it to analogously analyze the welfare impacts of various predatory incentives of firms. Figure 3 displays the results of our analysis in three forms.

First, on the left, we consider the size of the Spence distortion (that is, the divergence between the interaction value the platform internalizes and the interaction value it would optimally internalize), as a function of τ, at IE. To do so, we chart the interaction value of the average active gamer (the highest curve, which is also the interaction value of the average switching gamer), the average exiting gamer (the lowest curve), and the gamer whose utility the platform internalizes (the middle curve). The Spence distortion, which is the gap between the middle and top curve, steadily rises in market power τ. It is quite large in magnitude, cutting in half the amount passed through to developers, and rises by a bit over a third from $\tau = 500$ to $\tau = 6500$.

Second, in the center, we consider the impact of this distortion on the number of games developed for the platforms, plotting the divergence between the games developed under the two scenarios. To put these figures in context, the number of games developed in total under both scenarios falls, as τ increases, from ≈ 1425 to ≈ 1375 even under the second-best as the increased prices created by market power reduces the gamer base. The center panel shows that the Spence distortion leads to 110-160 fewer games being developed, and this distortion increases by about half as τ increases over our range. Thus, the increase in the Spence distortion accounts for about half of the reduction in games developed as market power increases under IE and distorts this development downward by about 10% overall relative to second-best.

Finally, in the right panel, we consider the impact of this distortion on social welfare. Again, to get a sense of magnitudes, second-best social welfare is in the range of $44-$48 billion for all values of τ we consider and declines in market power. The reduction in social surplus resulting from the Spence distortion is in the range of about $4.5 billion to $5.9 billion, a bit larger than the full impact of competition on second-best social welfare, and it increases by roughly $1.5 billion as τ increases, accounting for about a quarter of the loss in social welfare from the rise in market power. Most of this effect comes from the impact on gamer welfare, though we do not display this here.

To summarize, in this model, the Spence distortion is quite large in both absolute terms and relative to the total impact of competition over the range we consider. It grows significantly as market power rises. It accounts for a significant part of the reduction in welfare from rising market power, as a result of the pathway we discussed above: by reducing the amount passed through to developers, the Spence distortion lowers the number of games available and harms gamers and developers, and thus lowers social welfare.

5.3 Newspapers and the Costs of Competition

5.3.1 Model

The second example features two competing newspapers that serve advertisers on side A and readers on side B. One of the newspapers, h, is "high-brow", producing content that readers view to be of higher quality than the content of its "low-brow" competitor, l. We thus use vertical product differentiation in the spirit of Shaked and Sutton (1982) and Bresnahan (1987) to model the degree of competition between the papers.

Readers subscribe to at most one newspaper, while advertisers choose whether or not to place ads independently across papers.[33]

[33] While we use the term "subscribe", we follow Gentzkow (2007) in not distinguishing explicitly between newsstand and subscription sales, effectively averaging the two together to form our empirical magnitudes.

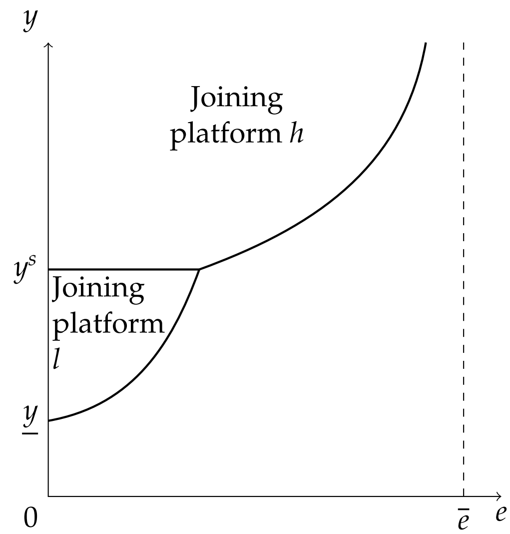

Figure 4: The structure of demand at equilibrium in our newspaper model. Hassle costs of reading are on the horizontal axis, and income is on the vertical axis. Those above the curve read and above the horizontal line read the high-brow paper.

Readers are heterogeneous in their income, y, and in the hassle or effort, e, they require to read a newspaper. Formally, a reader of type $\theta = (y, e)$ who subscribes to newspaper $j \in \{l, h\}$ receives utility

$$u_j = y\left(q_j\left(1 - n_{j,A}\right)^{\phi} - e\right) - p_{j,B}. \tag{10}$$

Here, q_j denotes the (exogenous) quality-level of j's content, $\left(1 - n_{j,A}\right)$ is the size of what is sometimes referred to as the "newshole" – the fraction of newspaper j's space devoted to content rather than ads, and $\phi > 0$ is a parameter governing the rate at which ads provoke nuisance.

Following Gentzkow et al. (2014), we assume all advertisers are identical. Thus, no direct market power distortion of the number of advertisements occurs, allowing us to focus only on the distortions from the way in which reader-side preferences are internalized.[34] The number of advertisers is normalized to be unity if the newspaper is completely filled with ads, and thus we maintain the assumption of a unit mass of advertisers. We measure the willingness of advertisers to pay, per reader, for filling the full paper with adds by a parameter a.

Let $\Delta q \equiv q_h - q_l > 0$ denote the difference in the respective quality parameters of the two newspapers, which measures differentiation and thus market power analogously to

[34] In fact, this makes our analysis essentially an extension of Spence's model to allow competition through IE, because the lack of a distortion on the advertiser's side of the market makes the choice of a number of advertisers essentially equivalent to a quality choice.

the way τ measured market power in the previous subsection, where differentiation was horizontal. As we will see below, the inherent quality advantage of newspaper h leads it also to show fewer ads than l. When this occurs, reader demand takes the form depicted in Figure 4. Those readers lying above the kinked curve buy one of the two papers; those above the horizontal line purchase the high-brow paper, while those below the horizontal line purchase the low-brow paper.

As this figure shows, a threshold income level, y^s, exists above which readers prefer h to l, and vice versa. Each newspaper also has an exiting margin of individuals who, for a given income, find the effort of reading the newspaper barely offsets the value they derive from it.

In Subappendix B.3, we formally define various average reader types (in terms of their income, which determines their preference for advertising by Equation 10), but for our purposes, we note what they are, informally:

(1) $\widetilde{y}_l^{\emptyset}$ and $\widetilde{y}_h^{\emptyset}$ are the average incomes of the low- and high-brow papers' respective exiting readers, and

(2) $\overline{y}_l$ and $\overline{y}_h$ are the average incomes of all of the low- and high-brow papers' respective readers.

As noted above, the newspapers cater to a weighted average of the first two reader income types, whereas the second-best outcome is for them to cater to the third.[35] Because switchers have higher income than readers purchasing the low-brow paper and lower income than readers purchasing reading the high-brow paper, newspapers catering to them represents undesirable convergence toward the center ground, as in Hotelling (1929). Given that average exiting readers for each paper will likely have incomes more representative of average readers of the two papers respectively, more competition (a fall in Δq) may be harmful, as it causes the papers to cater more heavily to unrepresentative switching readers, rather than to the fairly representative exiting readers. We might thus expect competition to actually exacerbate the Spence distortion in this case and thus to have an ambiguous impact on social welfare.

5.3.2 Calibration to US Daily Newspapers

We now attempt to generally match the US daily newspaper market studied by Gentzkow (2007) among others, with all quantities in annual terms:

[35]More formally, by "cater to" we mean that if the newspaper acts in equilibrium and for purposes of determining advertising levels, as if all readers were homogeneous and have the "catered to" income. This income in turn determines the disutility of advertising and thus the level of advertising the newspaper finds it optimal to provide.

- We set $c_{l,B} = c_{h,B} = 150 \approx 0.4 \cdot 365$ as Gentzkow estimates papers cost approximately $0.40 to print and distribute.

- We have in mind a mid-sized US metropolitan area and thus assume a mass of one million potential readers.

- We calibrate
$$a = \frac{\text{Total Ad Revenue}}{\text{Circulation}} / \frac{\text{Ad Space}}{\text{Total Space}}.$$
Available numbers for these quantities are, unfortunately, all over the map; Gentzkow and representative data from the Newspaper Association of America yield numbers in the $100-$1500 range. However, numbers much above $300 lead to negative prices for readers, which are not observed, and which impair convergence. We thus used $a = 300$.

- Readers' income and effort parameters are independently distributed. Income is log-normally distributed with log mean 10.86 and log standard deviation 0.5, a lower-variance fit to the median of US income distribution to account for the fact that a significant part of inequality is across metropolitan regions.

- We assume $\phi = 0.2$, because it leads to reasonable model outputs; such strong concavity seems reasonable to reflect the diminishing marginal appetite of readers for news and the increasing obtrusiveness of advertising as it comes to dominate the paper.

- Our comparative static exercise involves changing Δq, while holding fixed $\overline{q} \equiv (q_h + q_l)/2$. We assume $\overline{q} = 0.02$; that is, for the "average" paper with no advertising, the most enthusiastic newspaper reader would be willing to pay 2% of her annual income, which implies, reasonably, that a median income household that is maximally enthusiastic about newspapers would be willing to pay about $1000 a year for them if they had no advertising.

- We consider exogenous quality differentiation in the range $\Delta q \in [0.0005, 0.012]$; that is, the difference in quality (before factoring in the difference in endogenous advertising) is worth between 0.05% of income (viz. about $26 for a median income household) and 1.2% of income (viz. about $620 for a median income household) to readers.

- e is uniformly distributed on $[0, 2\overline{q}]$. This upper bound is the minimum necessary to ensure the high-brow paper has non-subscribers at all income levels and thus that the exiting margin is smooth.

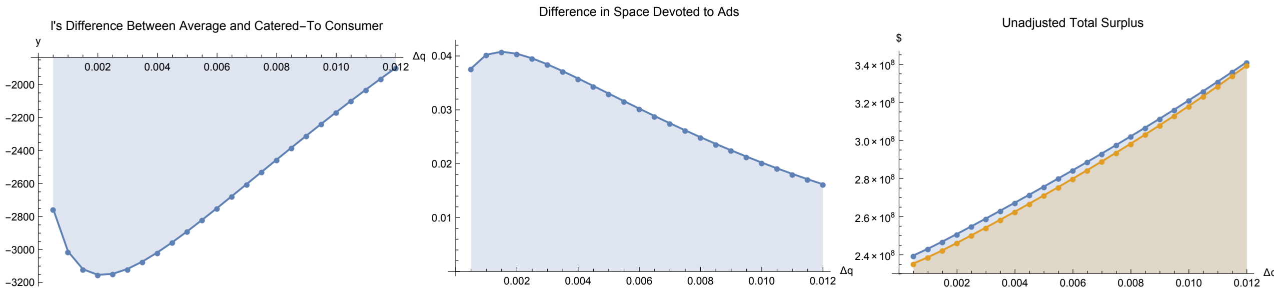

Figure 5: The difference between outcomes at IE and in the no-Spence-distortion second-best as a function of Δq, the difference in newspaper quality and thus market power, in our calibrated newspaper model. On the left is the divergence between the average income of the catered-to readers under the two scenarios. In the center is the difference between the IE and second-best newshole for the low-brow paper. On the right is the second-best (top) and IE (bottom) aggregate surpluses.

5.3.3 Results

The asymmetric set-up implies we should not expect a symmetric equilibrium. We therefore search (in a higher-dimensional space) for an asymmetric equilibrium using Newton's method on the system of now four FOCs and again verify second-order conditions through computational exploration.

Figure 5 is analogous to Figure 3 in that it compares outcomes under the no-Spence-distortion second-best to IE. However we only picture our intermediate outcomes for the low-brow paper for spatial economy; the results for the high-brow paper are basically similar, though they are both more dramatic and have greater failures of monotonicity and thus are somewhat messier to exposit.

In the left panel, the figure shows the difference between the average income of the readers to whom the low-brow newspaper caters and its average reader $\overline{y}_l$, which is now negative because the low-brow paper caters partly to the switchers between it and the high-brow newspaper who have higher income than its own average readers. To fix ideas about magnitudes, the typical average reader of the low-brow paper has income in the \$30-35k range, with this figure steadily falling as Δq rises. The Spence distortion is thus fairly modest here, never more than 10%. Except for very small values of Δq, the Spence distortion of the catered-to reader steadily declines in magnitude as Δq increases, falling by a third in total and accounting for about a fifth of the total fall in the income of the reader to whom the low-brow paper caters.

The central panel shows the impact of the Spence distortion on the size of the newshole. To fix magnitudes, the newshole shrinks from about 38% to about 22% at the second-best over the range of Δq we consider, shrinking steadily as Δq increases. The Spence distortion distorts the newshole upward by about 2-4 percentage points, with the distortion declining monotonically except, as above, for very small Δq.

The right panel shows the impact of the Spence distortion on total welfare, picturing in the top curve second-best total surplus and on the bottom that which occurs at IE. Social welfare in aggregate is on the order of \$200-400 million. The (aggregate) surplus foregone because of the Spence distortion is only \$2-\$5 million and declines from the high end to the low end of this range as Δq increases, again except for very small Δq. Note that aggregate social surplus rises in Δq mechanically, because raising Δq increases the quality of the paper read by those who most value quality, and we make no attempt to adjust our metrics for this effect.

To summarize, in this context, the Spence distortion is quite small and has a limited impact relative to other factors. However, market power typically makes this small distortion even smaller to a significant degree in relative terms, though this effect is not entirely monotonic. Of all the specifications we tried, this one showed the clearest effects in this direction. These exercises lead us to tentatively view the excess convergence result (Hotelling, 1929; de Palma et al., 1985) as more fragile and less common than the tendency of competition, when differentiation is orthogonal to the quality choice variable, to reduce the Spence distortion.[36] However, these issues demand further investigation, and we hope these examples will spur broader research on this topic.

6 Conclusion

In this paper, we propose Insulated Equilibrium (IE), a novel refinement of Nash Equilibrium in pricing functions, giving crisp predictions in a general static model of platform competition amenable to applied policy analysis. The basic idea behind IE is that each platform, while best-responding to its competitor, offers prices that compensate users for variation in the level of network effects provided, thus eliminating possible issues of (mis-) coordination and multiple equilibria among users.

We argue that, in some respects, IE is more realistic than the standard approach of "flat pricing", because it encapsulates in a static model a reduced form of the "penetration pricing" frequently used by firms in network industries and predicted in dynamic models of platform competition (Katz and Shapiro, 1986; Mitchell and Skrzypacz, 2006; Cabral, 2011*a*,*b*; Veiga, 2014). Moreover, when comparison between the predictions of IE and those of flat pricing is possible, they are qualitatively similar, but IE can be used to analyze a much broader range of the model space. In particular, it allows us to study the effects of

[36]This conclusion is broadly consistent with the follow-on literature about how endogenous differentiation compares to the social optimum (D'Aspremont et al., 1979; Economides, 1986; Irmen and Thisse, 1998). However, in this literature, all dimensions of differentiation are endogenous, and thus there is no way to analyze the impact we emphasize of competition *per se* on endogenous quality provision.

competition when users are heterogeneous in their *valuations of network effects*. Only when this form of user heterogeneity is present does a key incentive arise for platforms to distort their prices in a way that leads to inefficient network-effect provision. Nevertheless, prior approaches to modeling platform competition do not easily accommodate this heterogeneity, and thus the literature has typically assumed it away. We illustrate IE's usefulness in this regard by developing calibrated models of video game and newspaper industries in which heightened competition, respectively, reduces and exacerbates this distortion.

The most natural direction for future research is the use of our solution concept to study a variety of practical antitrust and regulatory policy issues that were intractable under previous analytic approaches. Antitrust authorities have reviewed numerous mergers in platform industries in recent years and have struggled to build a coherent two-sided platform model due to the intractability of the standard approach. Many of these, such as the Google/DoubleClick merger, featured competition between incumbent firms and potentially disruptive competitors. Given the concern such cases have recently raised about the elimination of dynamic competition, a salient extension of our analysis of competition would be to a framework for merger analysis applicable to such settings.

Other leading antitrust concerns in these industries include predation (of which Uber and other ride-sharing platforms have been accused by incumbent taxi providers) and exclusive vertical contracting (which has played an important role in the development of the video game industry). By providing an account of optimal pricing consistent with the patterns of user heterogeneity that empirical analysis has shown to be critical, our approach provides a natural pathway to applied analysis of these policy challenges.

Another interesting direction for future research regards the micro-foundations of IE. In particular, understanding in what explicit dynamic model, if any, exactly or approximately insulating tariffs are uniquely optimal for platforms would be helpful. Preliminary works by Cabral (2011*a*) and Veiga (2014) investigate this question, but much more work remains to be done. Conversely, we showed that both FP and IE effectively rule out the concerns about expectations and coordination emphasized by David (1985) and Arthur (1989). This suggests that, in models of complete information about fundamentals, mis-coordination is unlikely to be an important source of inefficiency.

However, in models with incomplete information about primitives, such as the value of a particular platform, such concerns seem likely to be important. In such cases, adopting strategies like insulation could, in contrast to our setting, be undesirable for platforms, because, for example, users may know better than the firm whether any prices exist at which (under complete information) the platform brings positive value to the market relative to existing alternatives. In such a case, insulation could expose the platform to losses if the latter turns out to be of low value. Indeed, users' decisions whether to coordinate on a platform might be informative to the platform about its potential future profitability. Allowing the kind of "coordination failures" that insulation prevents might thus be optimal for the platform. Jullien and Pavan (2014) consider a platform model in which the demand distribution is not common knowledge and find that, unless the platform shares the beliefs of the marginal consumers, charging insulating tariffs is typically not optimal. However, their model does not allow coordination failures stemming purely from the lack of insulation, as their structure has sufficient differentiation to ensure a unique equilibrium in which platforms split the market. Further investigation on the role of such incomplete information, its relationship to insulation, and its impact on our benchmark results is an exciting topic for future research.

References

Ambrus, Attila and Argenziano, Rossella. (2009). 'Asymmetric Networks in Two-Sided Markets', *American Economic Journal: Microeconomics* 1(1), 17–52.

Anderson, Simon P. and Coate, Stephen. (2005). 'Market Provision of Broadcasting: A Welfare Analysis', *The Review of Economic Studies* 72(4), 947–972.

Anderson, Simon P. and de Palma, Andre. (1992). 'The Logit as a Model of Product Differentiation', *Oxford Economic Papers* 44(1), 51–67.

Anderson, Simon P., de Palma, André and Thisse, Jacques-François. (1992), *Discrete Choice Theory of Product Differentiation*, MIT Press, Cambridge, Massachusetts.

Armstrong, Mark. (2006). 'Competition in Two-Sided Markets', *RAND Journal of Economics* 37(3), 668–691.

Arthur, W. Brian. (1989). 'Competing Technologies, Increasing Returns, and Lock-In by Historical Events', *Economic Journal* 99(394), 116–131.

Azevedo, Eduardo, Weyl, E. Glen and White, Alexander. (2013). 'Walrasian Equilibrium in Large, Quasi-Linear Markets', *Theoretical Economics* 8(2), 281—290.

Becker, Gary S. (1991). 'A Note on Restaurant Pricing and Other Examples of Social Influence on Price', *Journal of Political Economy* 99(5), 1109–1116.

Bergemann, Dirk and Morris, Stephen. (2005). 'Robust Mechanism Design', *Econometrica* 73(6), 1771–1813.

Berry, Steven, Levinsohn, James and Pakes, Ariel. (1995). 'Automobile Prices in Market Equilibrium', *Econometrica* 63(4), 841–890.

Berry, Steven T. (1994). 'Estimating Discrete-Choice Models of Product Differentiation', *RAND Journal of Economics* 25(2), 242–262.

Besanko, David, Doraszelski, Ulrich and Kryukov, Yaroslav. (2014). 'The Economics of Predation: What Drives Pricing When There is Learning-by-Doing?', *American Economic Review* 104(3), 868–897.

Bresnahan, Timothy F. (1987). 'Competition and Collusion in the American Automobile Industry: the 1955 Price War', *Journal of Industrial Economics* 35(4), 457–482.

Cabral, Luís. (2011*a*), A Dynamic Theory of Two-Sided Markets. http://goo.gl/rONRby.

Cabral, Luís. (2011*b*). 'Dynamic Price Competition with Network Effects', *Review of Economic Studies* 78(1), 83–111.

Caplin, Andrew and Nalebuff, Barry. (1991). 'Aggregation and Imperfect Competition: On the Existence of Equilibrium', *Econometrica* 59(1), 25–59.

Chung, Kim-Sau and Ely, J. C. (2007). 'Foundations of Dominant-Strategy Mechanisms', *Review of Economic Studies* 74(2), 447–476.

D'Aspremont, C., Gabszewicz, J. Jaskold and Thisse, J.-F. (1979). 'On Hotelling's "Stability in Competition"', *Econometrica* 47(5), 1145–1150.

David, Paul A. (1985). 'Clio and the Economics of QWERTY', *American Economic Review* 75(2), 332–337.

de Palma, A., Ginsburgh, V., Papageorgiou, Y. Y. and Thisse, J.-F. (1985). 'The Principle of Minimum Differentaiton Holds under Sufficient Heterogeneity', *Econometrica* 53(4), 767–781.

Dybvig, Philip H. and Spatt, Chester S. (1983). 'Adoption Externalities as Public Goods', *Journal of Public Economics* 20(2), 231–247.

Economides, Nicholas. (1986). 'Minimal and Maximal Product Differentiation in Hotelling's Duopoly', *Economics Letters* 21(1), 67–71.

Einav, Liran, Finkelstein, Amy and Cullen, Mark R. (2010). 'Estimating Welfare in Insurance Markets Using Variation in Prices', *Quarterly Journal of Economics* 125(3), 877–921.

Farrell, Joseph and Saloner, Garth. (1986). 'Standardization and Variety', *Economics Letters* 20, 71–74.

Filistrucchi, Lapo and Klein, Tobias J. (2013), Price Competition in Two-Sided Markets with Heterogeneous Consumers and Network Effects. http://papers.ssrn.com/abstract=2336411.

Gabszewicz, Jean J. and Wauthy, Xavier Y. (2014). 'Vertical Product Differentiation and Two-Sided Markets', *Economics Letters* 123(1), 58—61.

Gentzkow, Matthew. (2007). 'Valuing New Goods in a Model with Complementarity: Online Newspapers', *The American Economic Review* 97(3), 713–744.

Gentzkow, Matthew, Shapiro, Jesse M. and Sinkinson, Michael. (2014). 'Competition and Ideological Diversity: Historical Evidence from US Newspapers', *American Economic Review* 104(10), 3073–3114.

Green, Jerry and Laffont, Jean-Jacques. (1977). 'Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods', *Econometrica* 45(2), 427–438.

Hagiu, Andrei. (2006). 'Pricing and Commitment by Two-Sided Platforms', *RAND Journal of Economics* 37(3), 720–737.

Hotelling, Harold. (1929). 'Stability in Competition', *The Economic Journal* 39(153), 41–57.

Irmen, Andreas and Thisse, Jacques-François. (1998). 'Competition in Multi-Characteristics Spaces: Hotelling Was Almost Right', *Journal of Economic Theory* 78(1), 76–102.

Jullien, Bruno. (2011). 'Competition in Multi-Sided Markets: Divide and Conquer', *American Economic Journal: Microeconomics* 3(4), 186–219.

Jullien, Bruno and Pavan, Alessandro. (2014), Platform Pricing under Dispersed Information. http://goo.gl/AQ2D9l.

Katz, Michael L. and Shapiro, Carl. (1985). 'Network Externalities, Competition, and Compatibility', *The American Economic Review* 75(3), 424–440.

Katz, Michael L. and Shapiro, Carl. (1986). 'Technology Adoption in the Presence of Network Externalities', *Journal of Political Economy* 94(4), 822–841.

Klemperer, Paul D. and Meyer, Margaret A. (1989). 'Supply Function Equilibria in Oligopoly under Uncertainty', *Econometrica* 57(6), 1243–1277.

Lee, Robin S. (2012), Home Videogame Platforms, *in* **Martin Peitz and Joel Waldfogel**., eds, 'The Oxford Handbook of the Digital Economy', Oxford University Press.

Lee, Robin S. (2013). 'Vertical Integration and Exclusivity in Platform and Two-Sided Markets', *American Economic Review* 103(7), 2960–3000.

Mankiw, N. Gregory and Whinston, Michael D. (1986). 'Free Entry and Social Inefficiency', *RAND Journal of Economics* 17(1), 48–58.

Milgrom, Paul and Segal, Ilya. (2002). 'Envelope Theorems for Arbitrary Choice Sets', *Econometrica* 70(2), 583–601.

Mitchell, Matthew F. and Skrzypacz, Andrzej. (2006). 'Network Externalities and Long-Run Market Shares', *Economic Theory* 29(3), 621–648.

Nevo, Aviv. (1998). 'Identification of the Oligopoly Solution Concept in a Differentiated-Products Industry', *Economics Letters* 59(3), 391–395.

Ordover, Janusz A. and Saloner, Garth. (1989), Predation, Monopolization and Antitrust, *in* **Richard Schmaleness and Robert D. Willig**., eds, 'Handbook of Industrial Organization', Vol. I, Amsterdam: Elsevier.

Reisinger, Markus. (2014). 'Two-Part Tariff Competition between Two-Sided Platforms', *European Economic Review* 68(C), 168–180.

Rochet, Jean-Charles and Stole, Lars A. (2002). 'Nonlinear Pricing with Random Participation', *Review of Economic Studies* 69(1), 277–311.

Rochet, Jean-Charles and Tirole, Jean. (2003). 'Platform Competition in Two-Sided Markets', *Journal of the European Economic Association* 1(4), 990–1029.

Rochet, Jean-Charles and Tirole, Jean. (2006). 'Two-Sided Markets: A Progress Report', *RAND Journal of Economics* 37(3), 645–667.

Segal, Ilya. (1999). 'Contracting with Externalities', *The Quarterly Journal of Economics* 114(2), 337–388.

Shaked, Avner and Sutton, John. (1982). 'Relaxing Price Competition Through Product Differentiation', *The Review of Economic Studies* 49(1), 3–13.

Sheshinski, Eytan. (1976). 'Price, Quality and Quantity Regulation in Monopoly Situations', *Economica* 43(170), 127–137.

Song, Minjae. (2015), Estimating Platform Market Power in Two-sided Markets with an Application to Magazine Advertising. http://ssrn.com/abstract=1908621.

Spence, A. Michael. (1975). 'Monopoly, Quality, and Regulation', *Bell Journal of Economics* 6(2), 417–429.

Spence, Michael. (1976). 'Product Differentiation and Welfare', *American Economic Review* 66(2), 407—-414.

Spulber, Daniel F. (1999), *Market Microstructure and the Theory of the Firm*, Cambridge University Press, Cambridge, U.K.

Veiga, André. (2014), Dynamic Platform Design. http://ssrn.com/abstract=2506461.

Weyl, E. Glen. (2010). 'A Price Theory of Multi-Sided Platforms', *American Economic Review* 100(4), 1642–1672.

Weyl, E. Glen and White, Alexander. (2014). 'Let the Right "One" Win: Policy Lessons from the New Economics of Platforms', *Competition Policy International* 10(2), 29–51.

Wilson, Robert. (1987), Game Theoretic Analysis of Trading Processes, *in* **Truman F. Bewley**., ed., 'Advances in Economic Theory', Cambridge University Press, Cambridge, U.K.

Appendices

A Proofs

A.1 Proof of Theorem 1

We establish the results using a series of lemmata that we now state and prove.

Lemma 3. *Under Assumptions 1 and 2, $\pi_j(t_A, t_B)$ is strictly quasi-concave in $\left(t_{j,A}, t_{j,B}\right)$ given any $\left(t_{k,A}, t_{k,B}\right)$.*

Proof. For any t_k, consider the profit function $\pi_j\left(t_j; t_k\right)$; given that t_k is held fixed, from now on we suppress it entirely, considering only residual profit and demand functions, but it should be understood as implicitly present. We can change variables to n_j according to the transformation $n_{j,A} = \widetilde{N}_{j,A}\left(t_{j,A}\right)$. Note that this transformation is globally strictly monotone as by full support $\widetilde{N}'_{j,A} < 0$ globally. Thus, it suffices to prove that $\widetilde{\pi}_j\left(n_j\right) \equiv \pi_j\left(\widetilde{N}^{-1}_{j,A}\left(n_{j,A}\right), \widetilde{N}^{-1}_{j,B}\left(n_{j,B}\right)\right)$ is strictly concave.

By the chain rule and the inverse function theorem,

$$g_{j,A}\left(t_j\right) \equiv \frac{\partial \tilde{\pi}_j}{\partial n_{j,A}}\left(\widetilde{N}_{j,A}\left(t_{j,A}\right), \widetilde{N}_{j,B}\left(t_{j,B}\right)\right) = \frac{\frac{\partial \pi_j}{\partial t_{j,A}}}{N'_{j,A}}$$

so that the Hessian of $\widetilde{\pi}_j$ is

$$\begin{bmatrix} \frac{\partial g_{j,A}}{\partial t_{j,A}}\frac{1}{N'_{j,A}} & \frac{\partial g_{j,A}}{\partial t_{j,B}}\frac{1}{N'_{j,B}} \\ \frac{\partial g_{j,B}}{\partial t_{j,A}}\frac{1}{N'_{j,A}} & \frac{\partial g_{j,B}}{\partial t_{j,B}}\frac{1}{N'_{j,B}} \end{bmatrix}.$$

To show this is globally negative definite, it suffices to show that globally

$$\frac{\partial g_{j,A}}{\partial t_{j,A}}, \frac{\partial g_{j,B}}{\partial t_{j,B}}, \left(\frac{\partial g_{j,A}}{\partial t_{j,A}}\frac{\partial g_{j,B}}{\partial t_{j,B}} - \frac{\partial g_{j,A}}{\partial t_{j,B}}\frac{\partial g_{j,B}}{\partial t_{j,A}}\right)\frac{1}{N'_{j,A}N'_{j,B}} > 0;$$

and this in turn implies the concavity of $\widetilde{\pi}_j$ and thus the result.

We first consider $\frac{\partial g_{j,A}}{\partial t_{j,A}}$. By direct differentiation of π_j from the text,

$$g_{j,A} = t_{j,A} + \gamma_A \widetilde{N}_{j,B} - \mu_{j,A} - c_{j,A} + \gamma_B \widetilde{N}_{j,B}. \tag{11}$$

Thus, $\frac{\partial g_{j,A}}{\partial t_{j,A}} = 1 + \psi$ where

$$\psi \equiv \frac{\partial \frac{\widetilde{N}_{j,A}}{\widetilde{N}'_{j,A}}}{\partial t_{j,A}}$$

But by definition of ρ-concavity, if $\widetilde{N}_A(t_A)$ is ρ-concave in $t_{j,A}$ then $\psi \geq \rho$. Note that side A user preferences satisfy the requirements of Caplin and Nalebuff's Assumption A1 for the case of 2 dimensions of heterogeneity. Thus, Caplin and Nalebuff's Theorem 1 and Assumption 1 implies that $\widetilde{N}_A(t_A)$ is $-a$ concave. Thus, $\frac{\partial g_{j,A}}{\partial t_{j,A}} \geq 1 - a > 0$ globally.

Next note that $\frac{\frac{\partial g_{j,B}}{\partial t_{j,A}}}{N'_{j,A}} = \gamma_A + \gamma_B$. Thus,

$$\left(\frac{\partial g_{j,A}}{\partial t_{j,B}}\frac{\partial g_{j,B}}{\partial t_{j,A}}\right)\frac{1}{N'_{j,A}N'_{j,B}} > 0 \iff \frac{\partial g_{j,A}}{\partial t_{j,A}}\frac{\partial g_{j,B}}{\partial t_{j,B}} > (\gamma_A + \gamma_B)^2 \widetilde{N}'_{j,A}\widetilde{N}'_{j,B}.$$

But $\widetilde{N}'_{j,A}\widetilde{N}'_{j,B} < 4M_AM_B$ because

$$\widetilde{N}'_{j,A} = -\left(\int_{-\infty}^{\theta_j - t_{j,A}+t_{k,A}} f_A\left(t_{j,A},\theta_k\right)d\theta_k + \int_{t_{j,A}}^{\infty} f_{A,j}\left(\theta_j\right)f_{A,k|j}\left(\theta_j - t_{j,A} + t_{k,A}|\theta_j\right)d\theta_j\right),$$

$$\int_{-\infty}^{\theta_j - t_{j,A}+t_{k,A}} f_A\left(t_{j,A},\theta_k\right)d\theta_k < \int_{-\infty}^{\infty} f_A\left(t_{j,A},\theta_k\right)d\theta_k = f_{A,j}\left(t_{A,j}\right) \leq M_A$$

and

$$\int_{t_{j,A}}^{\infty} f_{A,j}\left(\theta_j\right)f_{A,k|j}\left(\theta_j - t_{j,A} + t_{k,A}|\theta_j\right)d\theta_j < \int_{-\infty}^{\infty} f_{A,j}\left(\theta_j\right)f_{A,k|j}\left(\theta_j - t_{j,A} + t_{k,A}|\theta_j\right)d\theta_j \leq$$

$$M_A\int_{-\infty}^{\infty} f_{A,j}\left(\theta_j\right)d\theta_j = M_A$$

so that $\widetilde{N}'_{j,A} > -2M_A$. This proves the result as it implies that

$$\frac{\partial g_{j,A}}{\partial t_{j,A}}\frac{\partial g_{j,B}}{\partial t_{j,B}} > (\gamma_A + \gamma_B)^2 \widetilde{N}'_{j,A}\widetilde{N}'_{j,B} \impliedby (1-a)(1-b) > 4\left(\gamma_A + \gamma_B\right)^2 M_AM_B.$$

□

Note that the factor of four above can clearly be eliminated if there is only one margin, as it is an artifact of the two terms entering the sum defining $\widetilde{N}'_{j,A}$. This justifies our eliminating the factor of four when comparing the current result to those of the Armstrong single-homing model in Subsection 3.2.

The following lemma is a generic fixed point argument for a function with an open domain. It is perhaps canonical but we have not found the right reference and thus state the result and prove it in our online appendix. Note that the upper-bound of the interval stated in point 2 is closed, whereas the upper-bound of the interval stated in point 3 is open. The closed upper bound in point 2 reflects the intuition, proved formally in Lemma 5, that when its competitor charges an infinitely high price, a given firm essentially becomes a monopolist.

Lemma 4. *Let $h : \mathbb{R}^n \to \mathbb{R}^n$ be continuous and have generic elements $h_i(x_i, x_{-i})$ satisfying*

1. $\frac{\partial h_i}{\partial x_i} > 0$ *globally,*
2. $\forall \hat{x}_{-i} \in (-\infty,\infty]^{n-1}, \lim_{x_i\to\infty}\lim_{x_{-i}\to\hat{x}_{-i}} h_i(x) > 0$
3. $\forall \hat{x}_{-i} \in (-\infty,\infty)^{n-1}, \lim_{x_i\to-\infty}\lim_{x_{-i}\to\hat{x}_{-i}} h_i(x) < 0.$

Then there exists $x^\star \in \mathbb{R}^n$ such that $h(x^\star) = 0$.

Proof. See our online appendix Section 1. □

We also require the follow technical lemma.

Lemma 5. *For any $\hat{t}_{k,A} \in (-\infty, \infty]$, $\lim_{t_{j,A}\to\infty} \lim_{t_{k,A}\to\hat{t}_{k,A}} t_{j,A} - \mu_{j,A}(t_A) = \infty$.*

Proof. For notational compactness in this proof we define $x \equiv t_{j,A}$ and $y \equiv t_{k,A}$.

$$x - \mu_{j,A}(x, y) = x - \frac{\int_x^\infty \int_{-\infty}^{a-x+y} f_A(a,b)dbda}{\int_{-\infty}^{y} f_A(x,b)\,db + \int_x^\infty f_A(a, a-x+y)\,da}. \tag{12}$$

When $x \in \mathbb{R}$, for all $\hat{y} \in (-\infty, \infty]$, $\lim_{y\to\hat{y}} x - \mu_{j,A}(x, y)$ is finite as both the numerator and the first term of the denominator are finite and the second term of the denominator approaches 0. As is implied by the $-a$ concavity established in the proof of Lemma 3, as x increases, $x - \mu_{j,A}(x, y)$ increases at a rate lower-bounded by a strictly positive number. Therefore, for sufficiently large x, $\lim_{y\to\hat{y}} x - \mu_{j,A}(x, y)$ is positive. □

Lemma 6. *Under Assumption 1 and 2, there exists a "critical tariff" $t^\star$ at which both platforms' first-order optimization conditions on both sides of the market are satisfied.*

Proof. Note that if, at some tariff vector $t^\star$, $g(t^\star) = 0$ then this is a critical point as the numerators of the entries of g are the first-order conditions of the platforms. To establish that such a point exists, we verify that the conditions of Lemma 4 hold with the diagonal elements being associated in the natural manner with the same platform-side pair.

The continuity of g follows immediately from its differentiability established in our proof of quasi-concavity in Lemma 3. We directly showed that $\frac{\partial g_{A,j}}{\partial t_{A,j}} > 0$ there as well. We now establish the two limit conditions.

First, let us consider the limit as $t_{j,A} \to -\infty$. Note that all terms of Equation 11 defining $g_{j,A}$ are either bounded (both externality terms and cost), strictly negative ($-\mu_{j,A}$) or approach $-\infty$ as $t_{j,A}$ does (namely the term $t_{j,A}$). Thus $\lim_{t_{j,A}\to-\infty} \lim_{t_{-j,A}\to\hat{t}_{-j,A}} g_{j,A}(t) = -\infty < 0$ for any $\hat{t}_{-j,A} \in \mathbb{R}^3$.

Finally, all terms in $g_{j,A}$ are bounded in t other than $t_{j,A} - \mu_{j,A}$. Thus, by Lemma 5, for any $\hat{t}_{-j,A} \in (-\infty, \infty]^3$, $\lim_{t_{j,A}\to\infty} \lim_{t_{-j,A}\to\hat{t}_{-j,A}} g_{j,A}(t) = \infty > 0$.

□

Proof of Theorem 1. The critical tariff identified in Lemma 6 must be an IE as, by Lemma 3, it is a global optimum for each platforms given the rival's strategy . □

A.2 Other proofs

Proof of Lemma 1 . We establish the two parts of the result in turn. For the first, note that for a given n_B, Assumption 3 (a) implies the distribution of willingness to pay for each bundle on side A has full support. To see this, note that it holds for each γ as adding any constant vector to a full support random variable yields another full support random variable. Integrating over γ implies that the support of the total value distribution is also full. Therefore, the result follows by the argument in the proof of Proposition 1.

For the second, note that by definition,

$$N_{j,A} = \int_{\gamma} \int_{\theta : j \in \operatorname{argmax}_S v_A(S, n_B, \gamma) + \theta_S - \sum_{j \in S} p_{j,A}} f_A(\gamma, \theta) d\theta d\gamma =$$

$$\int_{\gamma} \int_{\theta_{\{k\}}} \left(\int_{\theta_{\{j\}}=-\infty}^{\underline{\theta_{\{j\}}}} \int_{\theta_{\{j,k\}}=\underline{\theta_{\{j,k\}}}}^{\infty} f_{A,(\theta_{\{j\}},\theta_{\{j,k\}})|(\gamma,\theta_{\{k\}})} d\theta_{\{j,k\}} d\theta_{\{j\}} + \int_{\theta_{\{j\}}=\underline{\theta_{\{j\}}}}^{\infty} f_{A,\theta_{\{j\}}|(\gamma,\theta_{\{k\}})} d\theta_{\{j\}} \right) f_{A,(\gamma,\theta_{\{k\}})} d\left(\gamma, \theta_{\{k\}}\right),$$

where $f_{A,x}$ is the marginal density of x under f_A, $f_{A,x|y}$ is the (marginal) conditional distribution of x under A given y, arguments are dropped wherever possible for spatial economy, $\underline{\theta_{\{j\}}}$ denotes the cut-off type preferring $\{j\}$ over either $\{k\}$ or the empty set

$$\underline{\theta_{\{j\}}}\left(\gamma, \theta_{\{k\}}, p_A, n_B\right) \equiv \left[\theta_{\{k\}} + v_A\left(\{k\}, n_B, \gamma\right) - p_{k,A}\right]^+ + p_{j,A} - v_A\left(\{j\}, n_B, \gamma\right)$$

and $\underline{\theta_{\{j,k\}}}$ denotes the cut-off type preferring $\{j,k\}$ over either $\{k\}$ or the empty set

$$\underline{\theta_{\{j,k\}}}\left(\gamma, \theta_{\{k\}}, p_A, n_B\right) \equiv \left[\theta_{\{k\}} + v_A\left(\{k\}, n_B, \gamma\right) - p_{k,A}\right]^+ + p_{j,A} + p_{k,A} - v_A\left(\{j,k\}, n_B, \gamma\right).$$

Note that $\underline{\theta_{\{j\}}}$ and $\underline{\theta_{\{j,k\}}}$ are twice continuously differentiable *except* where $\theta_{\{k\}} = p_{k,A} - v_A\left(\{k\}, n_B, \gamma\right)$; that is, at precisely one point in the $\theta_{\{k\}}$ space. By the assumption that $\theta|\gamma$ has a smooth density function, this point is of measure 0 in the integral defining $N_{j,A}$ and thus $N_{j,A}$ is twice continuously differentiable so long as the Leibniz rule derivatives defining its relevant derivatives are everywhere finite. To see that this holds, consider a representative case, the partial derivative with respect to $n_{j,B}$: $\frac{\partial N_{j,A}}{\partial n_{j,B}} =$

$$\int_{\gamma} \int_{\theta_{\{k\}}} \left(\int_{\theta_{\{j\}}=-\infty}^{\underline{\theta_{\{j\}}}} \frac{\partial v_A}{\partial n_{j,B}}(\{j,k\}) f_{A,\left(\theta_{\{j\}},\theta_{\{j,k\}}\right)|\left(\gamma,\theta_{\{k\}}\right)}\left(\theta_{\{j\}}, \underline{\theta_{\{j,k\}}}\right) d\theta_{\{j\}} + \int_{\theta_{\{j,k\}}=-\infty}^{\underline{\theta_{\{j,k\}}}} \frac{\partial v_A}{\partial n_{j,B}}(\{j\}) f_{A,\left(\theta_{\{j\}},\theta_{\{j,k\}}\right)|\left(\gamma,\theta_{\{k\}}\right)}\left(\underline{\theta_{\{j\}}}, \theta_{\{j,k\}}\right) d\theta_{\{j,k\}} \right) f_{A,\left(\gamma,\theta_{\{k\}}\right)} d\left(\gamma, \theta_{\{k\}}\right),$$

that is, the sum of expectations of bounded quantities (partial derivatives of v_A) multiplied by densities of these events, both of which are finite as a result. Thus, the derivative exists. Similar arguments show the result for the other derivatives, first and second.

□

Proof of Lemma 2. Everything but the second part of claim (a) follows again by applying Proposition 1 of Azevedo et al. (2013), given full support is satisfied for any fixed n_B. To establish the second part of claim (a), note that by definition,

$$V_A\left(n_A, n_B\right) = \int_{\gamma,\theta} \left[\max_S v_A\left(S, n_B, \gamma\right) + \theta_S - \sum_{l \in S} P_{l,A} \right] f_A(\gamma, \theta) d\theta d\gamma + P_{j,A} n_{j,A} + P_{k,A} n_{k,A}.$$

When taking the derivative with respect to n_B we can ignore

1. Any changes in the value of S that achieves the maximum for each (γ, θ) by the envelope theorem for general choice sets of Milgrom and Segal (2002).

2. The effects on the price terms, as these are defined to match those of V_A (market clearing).

3. Effects on any $(\gamma, \theta) : j \notin \text{argmax}_S\, v_A(S, n_B, \gamma) + \theta_S - \sum_{l \in S} P_{l,A}(n_A, n_B)$ by Assumption 3 part (c).

The result then follows directly by differentiating V_A.

□

Proof of Proposition 9. Because the providers multi-home, $\frac{\partial N_{k,A}}{\partial p_{j,A}} = 0$, it holds that $-D_{jk,A} = 0$. Hence,

$$
\begin{aligned}
& N_{j,B}\left(\left[-\frac{\partial N_B}{\partial p_B}\right]^{-1}\left[\frac{\partial N_B}{\partial n_A}\right]\right)_{j,j} \\
= \quad & N_{j,B}\left(\begin{bmatrix} f^{\emptyset}_{j,B}+f^s_B & -f^s_B \\ -f^s_B & f^{\emptyset}_{k,B}+f^s_B \end{bmatrix}^{-1}\begin{bmatrix} \widetilde{\gamma}^{\emptyset}_{j,B}f^{\emptyset}_{j,B}+\widetilde{\gamma}^s_B f^s_B & -\widetilde{\gamma}^s_B f^s_B \\ -\widetilde{\gamma}^s_B f^s_B & \widetilde{\gamma}^{\emptyset}_{k,B}f^{\emptyset}_{k,B}+\widetilde{\gamma}^s_B f^s_B \end{bmatrix}\right)_{j,j} \\
= \quad & N_{j,B}\left(\frac{\begin{bmatrix} f^{\emptyset}_{k,B}+f^s_B & f^s_B \\ f^s_B & f^{\emptyset}_{j,B}+f^s_B \end{bmatrix}}{f^{\emptyset}_{j,B}f^{\emptyset}_{k,B}+f^{\emptyset}_{j,B}f^s_B+f^{\emptyset}_{k,B}f^s_B}\begin{bmatrix} \widetilde{\gamma}^{\emptyset}_{j,B}f^{\emptyset}_{j,B}+\widetilde{\gamma}^s_B f^s_B & -\widetilde{\gamma}^s_B f^s_B \\ -\widetilde{\gamma}^s_B f^s_B & \widetilde{\gamma}^{\emptyset}_{k,B}f^{\emptyset}_{k,B}+\widetilde{\gamma}^s_B f^s_B \end{bmatrix}\right)_{j,j} \\
= \quad & N_{j,B}\frac{\begin{bmatrix} \widetilde{\gamma}^s_B f^{\emptyset}_{k,B}f^s_B+\widetilde{\gamma}^{\emptyset}_{j,B}f^{\emptyset}_{j,B}\left(f^{\emptyset}_{k,B}+f^s_B\right) & \left(\widetilde{\gamma}^{\emptyset}_{k,B}-\widetilde{\gamma}^s_B\right)f^{\emptyset}_{k,B}f^s_B \\ \left(\widetilde{\gamma}^{\emptyset}_{j,B}-\widetilde{\gamma}^s_B\right)f^{\emptyset}_{j,B}f^s_B & \widetilde{\gamma}^s_B f^{\emptyset}_{j,B}f^s_B+\widetilde{\gamma}^{\emptyset}_{k,B}f^{\emptyset}_{k,B}\left(f^{\emptyset}_{j,B}+f^s_B\right) \end{bmatrix}_{j,j}}{f^s_B f^{\emptyset}_{k,B}+f^{\emptyset}_{j,B}\left(f^{\emptyset}_{k,B}+f^s_B\right)} \\
= \quad & N_{j,B}\frac{\widetilde{\gamma}^s_B f^{\emptyset}_{k,B}f^s_B+\widetilde{\gamma}^{\emptyset}_{j,B}f^{\emptyset}_{j,B}\left(f^{\emptyset}_{k,B}+f^s_B\right)}{f^s_B f^{\emptyset}_{k,B}+f^{\emptyset}_{j,B}\left(f^{\emptyset}_{k,B}+f^s_B\right)} = N_{j,B}\left(\omega_j\widetilde{\gamma}^s_B+\left(1-\omega_j\right)\widetilde{\gamma}^{\emptyset}_{j,B}\right).
\end{aligned}
$$

Because providers are multi-homing, we have $\frac{\partial \sigma_{j,A}}{\partial n_{k,B}} = 0$. Thus, Equation 7 implies that platform j's FOC on side B is

$$
P_{j,B} = c_{j,B} + \mu_{j,B} - \frac{\partial \sigma_{j,A}}{\partial n_{j,B}} N_{j,A}.
$$

The definition of $\widetilde{\gamma}_{j,A}$ directly implies Equation 9. □

B Details of Applications

We reserve the most detailed calculations, derivations and numerical methods related to our applications for our online appendix. However, for readers interested in peering just a bit underneath the hood, we briefly summarize some of the most salient additional details in this appendix. All proofs and derivations are contained in the online appendix, Sections 2-4.

B.1 Structure of IE in the Armstrong single-homing model

We begin by discussing the full set of equilibria under IE. We then turn to details behind our user surplus results. Finally we characterize one tipping equilibrium under FP completely.

For the following proposition, we assume costs are symmetrically equal to 0 on both sides for both platforms.

Proposition 10. *If $\gamma < 1.5\tau$, the unique IE is the splitting IE described by Proposition 3. When $1.5\tau \leq \gamma \leq 2\tau$, the set of equilibria are the unique splitting equilibrium described by Proposition 3 and a set of all "tipping equilibria" in which all users on both sides adopt a single platform, without loss of generality j, and fixed tariff components are defined by*

1. $t_{k,A} - t_{j,A} = \tau$,
2. $t_{k,A} \leq 0$,
3. $t_{k,A} + t_{k,B} \leq 2\tau - 2\gamma$ *and*
4. $t_{k,A} \geq 3\tau - 2\gamma$.

This shows that there exists bounded, two-dimensional continuous space of tipping equilibria. In the case when platform j has a cost advantage or disadvantage of Δc, we focus on characterizing the set of tipping equilibria only.

Proposition 11. *Suppose platform j has a cost of Δc on both sides of the market and platform k has 0 cost on both sides. Then there is an equilibrium tipping to platform j if and only if $\Delta c \leq \min\{\tau, 2\gamma - 3\tau\}$ and the set of such equilibria is given by*

1. $t_{k,A} - t_{j,A} = \tau$,
2. $t_{k,A} \leq 0$,
3. $t_{k,A} + t_{k,B} \leq 4\tau - 2\gamma$,
4. $t_{k,A} \geq \Delta c + 3\tau - 2\gamma$ *and*
5. $t_{k,A} + t_{k,B} \geq 2(\Delta c + \tau - \gamma)$.

This set of equilibria is depicted visually in Figure 6. The region marked with T is the set of tipping equilibria; these are restrictions on the fixed component of the tariffs and are portrayed in the first quadrant by considering negative values of these fixed components (and thus positive user utility values).

The following result characterizes user surplus in the symmetric, zero-cost case.

Proposition 12. *At any splitting equilibrium average user surplus is $\gamma - \frac{5}{4}\tau = \gamma - \tau - \frac{1}{4}\tau$, as $\frac{1}{4}\tau$ is the average transport cost born by users. Average user surplus at tipping equilibria equal the set $\left[\frac{1}{2}\tau, 2\gamma - \frac{5}{2\tau}\right]$.*

The results quoted in the paper about user surplus follow directly from this proposition.

Proposition 13. *Under FP, when $\gamma > \tau$, there always exists an equilibrium where the market tips to platform j, $p_{j,A} = 0$ and $p_{k,A} = \tau - \gamma$.*

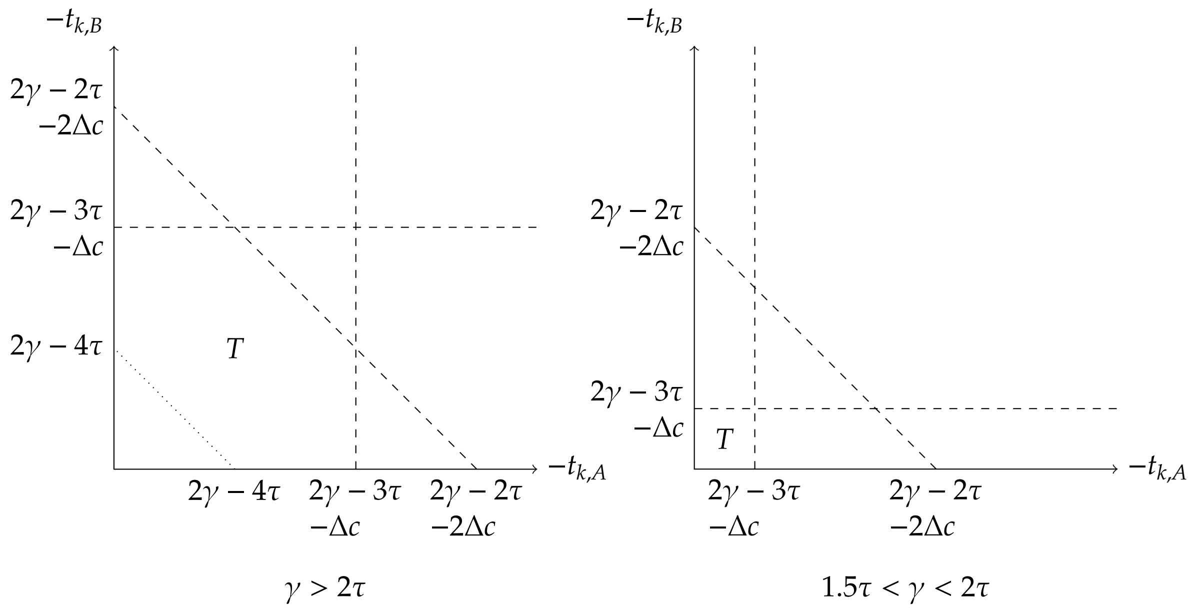

Figure 6: The set of tipping equilibria (labeled with T) for various parameter values in the Armstrong single-homing model.

B.2 Video games

In this subappendix, we only formally state the first-order equilibrium conditions and second-best conditions that we solve and provide some additional figures displaying results we could not fit in the text.

The market has two platforms located at either end of a Hotelling line; we refer to them as platforms 0 and 1. The mass of developers is m_A and of gamers is m_B. Developers have homogeneous interaction values a and heterogeneous development costs $d \sim f_A$. The set of developers for platform j is therefore characterized by a cut-off $\overline{d}_j$ below which developers will develop; at a symmetric equilibrium, this will just be $\overline{d}$ for both platforms.

Gamers have a two-dimensional type: each has a value per developer $\gamma \sim f_B$ and a horizontal position x distributed uniformly on the unit interval. Given our set-up, at a symmetric equilibrium there is a cut-off type $\underline{\gamma}$ such that any gamer with $\gamma > \underline{\gamma}$ buys some gaming platform, and all those who join a platform choose the one closest to them. Note that $\underline{\gamma}$ is the interaction value of all exiters, and thus clearly their average interaction value, while

$$\overline{\gamma}\left(\underline{\gamma}\right) \equiv \frac{\int_{\underline{\gamma}}^{\infty} \gamma f_B(\gamma) d\gamma}{1 - F_B\left(\underline{\gamma}\right)}$$

is both the average interaction value of switchers and that of the average user. Here, ω, the

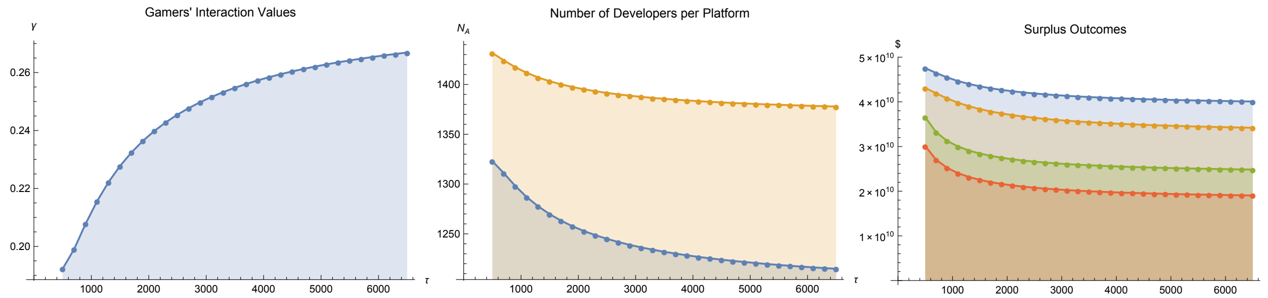

Figure 7: Additional outcomes in our calibrated video game model for various values of τ. On the left is the Spence distortion gap between average and catered-to gamer interaction values. In the center is the number of games developed at the second-best and IE. On the right are total (top two curves) and gamer (bottom two curves) surplus at the second-best (top of each pair of curves) and IE (bottom of each pair of curves.)

symmetric equilibrium weight on the switchers, is given by

$$\omega = \frac{1}{2 + \frac{\frac{f_B\left(\underline{\gamma}\right)}{m_A F_A\left(\overline{d}\right)}}{\frac{1-F_B\left(\underline{\gamma}\right)}{\tau}}},$$

as the two platforms are symmetric, gamer value heterogeneity is in per-equilibrium developer ($m_A F_A\left(\overline{d}\right)$) units and the number of exiters is the number of those who buy a console, divided by their transport cost τ, by the standard Hotelling logic.

Insulation involves extracting from the developers all their (assumed homogeneous) revenue from selling the games and then offering them a subsidy of $\overline{d}$. The first-order condition for this subsidy at a symmetric equilibrium, dropping as many arguments as possible (all are evaluated at their putative equilibrium values), is

$$\overline{d} = \underbrace{\left[\omega\overline{\gamma} + (1-\omega)\underline{\gamma}\right]}_{\text{internalized gamer}} \underbrace{m_B\left(\frac{1-F_B}{2}\right)}_{\text{\#of gamers}} - c_A \underbrace{-\frac{F_A}{f_A}}_{\text{developer monopsony power}}.$$

The second-best with no Spence distortion is characterized identically to the equilibrium, except that the fixed component of the developer subsidy is given by

$$\overline{d} = \underbrace{\overline{\gamma}}_{\text{average gamer}} \cdot m_B\left(\frac{1-F_B}{2}\right) - c_A - \frac{F_A}{f_A}.$$

That is, the only change is that the internalized gamer is now the average gamer, rather than a weighted average of the average and the marginal gamer. This leads to a strictly greater subsidy *holding fixed all the relevant cut-offs*, though, of course, these cut-offs are not actually held fixed in the re-equilibration process. However, this gives an intuition for why the Spence distortion leads to fewer developers in equilibrium.

Figure 7 shows some of the results that we used to fixed magnitudes or alluded to in the text but omitted there for brevity. The left panel pictures the Spence distortion gap between average active gamer interaction values and catered-to user interaction values. The central panel shows the total number of developers at the second-best (the higher of the two curves) and at IE. The right panel shows the total (top two curves) and gamer (bottom two curves) surplus. In each pair of curves the top one is second-best and the bottom one is IE. As we mention in the text, most of the change in total surplus, and most of the gap between the second-best and the IE outcome, results from gamer surplus.

B.3 Newspapers

Here we provide only a formal definition for the various catered-to types, though again, this requires setting up the model a bit, and then we display a few additional results, as in the previous subappendix.

The threshold switching income can easily be solved for as

$$y^s = \frac{p_{h,B} - p_{l,B}}{q_h\left(1 - n_{h,A}\right)^{\phi} - q_l\left(1 - n_{l,A}\right)^{\phi}}.$$

All these switching readers have the same income so we do not need to take averages. Readers with lower incomes will read the low-brow paper, if any. They will be indifferent to reading the low-brow paper and none at all if

$$e_h(y) = q_l\left(1 - n_{l,A}\right)^{\phi} - \frac{p_{l,B}}{y};$$

thus the lowest income of a reader buying any paper is

$$\underline{y} = \frac{p_{l,B}}{q_l\left(1 - n_{l,A}\right)}.$$

Similarly an individual who prefers the high-brow paper to the low-brow one will be indifferent between the high-brow paper and none at all if

$$e_l(y) = q_h\left(1 - n_{h,A}\right)^{\phi} - \frac{p_{h,B}}{y}.$$

We assume that e's distribution is uniform on $[0, \bar{e}]$ and that $y \sim f$, so the number of readers of the high-brow paper is

$$n_{h,B} = \int_{y^s}^{\infty} \frac{e_h}{\bar{e}} f(y)dy$$

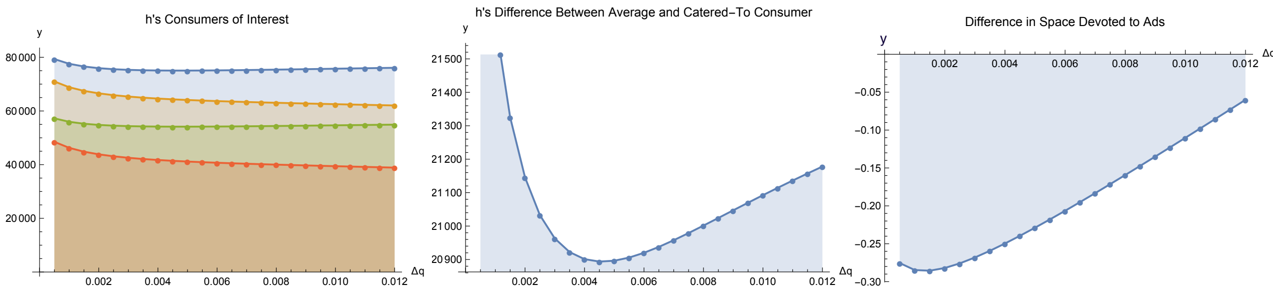

Figure 8: Outcomes for the high-brow paper in our calibrated newspaper model. On the left are various average reader incomes; the top is for the average readers; next is among exiting readers; next is among the readers catered-to; and the bottom is among switching readers. In the center is the Spence-distortion gap between average and catered-to readers. On the right is the impact of the Spence distortion on the number of advertisements; it is negative, indicating more advertisements under IE than at the second-best.

and the number of readers of the low-brow paper is

$$n_{l,B} = \int_{\underline{y}}^{y^s} \frac{e_l}{e} f(y)dy,$$

given that, as discussed in the text, our choice of $\bar{e} = 2\bar{q}$ guarantees that, at each income, some readers do not buy any paper.

Their average values are thus defined by

$$\overline{y_h} = \frac{\int_{y^s}^{\infty} y \frac{e_h}{\bar{e}} f(y)dy}{n_{h,B}}$$

and

$$\overline{y_l} = \frac{\int_{\underline{y}}^{y^s} y \frac{e_l}{\bar{e}} f(y)dy}{n_{l,B}}.$$

The density of individuals on the exiting margins are

$$f^{\emptyset}_{h,B} = \int_{y^s}^{\infty} \frac{\frac{\partial e_h}{\partial p_{h,B}}}{e} f(y)dy = \int_{y^s}^{\infty} \frac{1}{y\bar{e}} f(y)dy$$

and similarly for $f^{\emptyset}_{l,B}$. The average income of an exiter from the low-brow paper is then

$$\widetilde{y}^{\emptyset}_l = \frac{\int_{\underline{y}}^{y^s} \frac{1}{\bar{e}} f(y)dy}{f_{l,B}} = \frac{F(y^s) - F\left(\underline{y}\right)}{\int_{\underline{y}}^{y^s} \frac{1}{y} f(y)dy},$$

and similarly for $\widetilde{y}^{\emptyset}_h$.

Figure 8 shows outcomes similar to those shown in the text for the low-brow paper but for the high-brow paper. On the left, we see that the average purchasing readers (the highest curve) are consistently richer than exiting readers (the second-highest curve), who are in turn richer than the catered-to readers (the third curve) and the switching readers (the bottom curve). The Spence-distortion gap between the incomes of the average and catered-to readers is highly non-monotonic, but declines fairly steeply between as Δq rises from its lowest to intermediate values, though it eventually rises again. Despite this lack of monotonicity, the Spence distortion is fairly large here, around a quarter of the average income. The reason for the non-monotonicity is likely, as we see on the left, that despite the level of income to which the paper caters moving decisively toward exiting readers and away from switching readers, the income of both declines as Δq rises. Therefore, the high-brow paper becomes relatively more popular even at lower incomes, even while maintaining a high average income readership because of the skewness of the income distribution.

More monotone therefore, and similarly large, is the impact of the Spence distortion on space devoted to ads, pictured on the right. This distortion always leads to too many ads, because of the greater distaste average readers, relative to catered-to readers, have for ads. It falls dramatically as Δq (and thus the market power of the high-brow paper) rises, from nearly 30% of the paper to less than 10%. This is presumably largely driven by the increase in the readership of the high-brow paper making the externality tax on the advertisers larger, despite the non-monotone behavior of the Spence distortion. As we discuss in the text, these outcomes are more complicated and mixed, though in some ways more dramatic, than those that occur with the low-brow paper and are generally consistent with our summary message in the text.

Figure 9 shows some additional results for the low-brow paper and aggregate reader surplus. On the left are the various average readers for the low-brow paper: switchers are richer than catered-to readers who in turn are richer than average readers and exiting readers are the poorest. In the center is the space devoted to ads under the second-best (top) and IE (bottom). At right is reader surplus. The Spence distortion only ever slightly reduces reader surplus (no more than 10%), though, in relative terms, reader surplus falls significantly with Δq, tracking social surplus fairly closely in proportional terms.

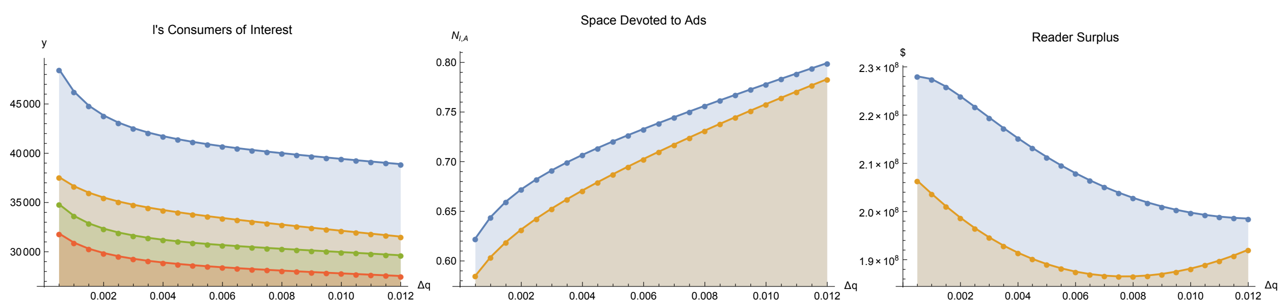

Figure 9: Additional outcomes for the low-brow paper (left two panels) and for all readers (right panel) in our calibrated newspaper model. On the left are various average readers for the low-brow paper; the top is switching readers, as in Figure 8; next is catered-to readers; next is average readers; finally exiting readers. In the center is the space devoted to ads at the second-best (higher) and IE (lower). On the right is total reader surplus at second-best (top) and IE (bottom).