

Imperfect Platform Competition:

A General Framework*

Alexander White[†]

E. Glen Weyl[‡]

November 2010

Job Market Paper

Abstract

The externalities advertisers receive from newspaper readers and that operating system users receive from software developers are among the leading features of those “platform” industries. However, they are rarely incorporated into applied models of imperfect competition. We argue this omission is due to a basic theoretical indeterminacy created by these externalities and propose the solution concept of *Insulated Equilibrium* to resolve it. At such equilibrium, each platform’s price on one side of the market adjusts to participation on the other side so as to insulate its own allocation, eliminating both the necessity for consumer coordination and the multiplicity of platform best replies. This allows us to solve a model of oligopoly without the unrealistic restrictions typically imposed for tractability and to demonstrate that the fundamental additional distortion created by consumption externalities is analogous to that identified by Spence (1975)’s analysis of a quality-choosing monopolist.

Keywords: Two-Sided Markets, Multi-Sided Platforms, Quality Competition, Oligopoly, Insulated Equilibrium, Antitrust and Mergers in Network Industries

JEL Codes: D21, D43, D85, L13

*We are grateful to Alexei Alexandrov, Mark Armstrong, Gary Biglaiser, Ed Glaeser, Renato Gomes, Sonia Jaffe, Bruno Jullien, Scott Kominers, Thomas Mariotti, Ariel Pakes, Patrick Rey, François Salanié, Larry Samuelson, Paul Seabright, Yossi Spiegel, Jean Tirole, Catherine Tucker, André Veiga, Julian Wright and especially Jacques Crémer for very helpful comments and discussions. We thank Rui Wang and Will Weingarten for their superb research assistance. White gratefully acknowledges the NET Institute (www.NETinst.org) and Weyl the Microsoft Corporation for their financial support. All errors are our own.

[†]Harvard University Department of Economics; alexanderwhite@fas.harvard.edu

[‡]Harvard University Society of Fellows & Toulouse School of Economics; weyl@fas.harvard.edu

1 Introduction

Apple and Microsoft's competing operating systems confront a set of challenges that also faces newspapers, credit cards, Internet service providers and search engines. However, these challenges are not reflected in the "canonical model" of imperfect competition. Such firms all sell multiple products, and the challenges in question arise from the *externalities* between the consumers of these different products. For instance, a particular operating system is more appealing to end users if the number of applications it boasts is greater; meanwhile, it is more appealing to application developers if it features a greater number of users. This paper offers a way to incorporate these consumption externalities into the canonical model and to analyze their implications for industrial policy.

In doing so, we build on a recent theoretical literature on such "two-sided markets" or "multi-sided platforms".¹ This literature has highlighted the pervasiveness of these consumption externalities throughout different industries and the fact that their influence on pricing can be of first-order importance. For example, it has offered a convincing explanation for otherwise puzzling negative prices we observe, drawing a clear link between phenomena such as operating system subsidies to application development, credit card point systems and the free availability of almost all websites.

Policymakers have clearly expressed interest in the effects of these externalities. For example, many have claimed that network neutrality regulation benefits consumers by expanding their choice of websites. Quantitative evaluation of these claims requires a model flexible enough to incorporate rich structures of consumer preferences and firm heterogeneity. However, rich models of "one-sided" competition, such as Berry, Levinsohn, and Pakes (1995) (BLP), have been considered intractable in the platform context.²

We argue that this apparent intractability stems from a basic indeterminacy in the current theory. In view of this, the paper aspires to make three contributions. First, it builds a model whose generality is comparable to that of standard one-sided models of competition that also includes consumption externalities, and it uses this model to illustrate this indeterminacy. Second, it proposes the solution concept of *Insulated Equilibrium*, which restores full predictive power to the model and whose criterion, we argue, is motivated by sound economic logic. Third, it identifies a fundamental force, which, in addition to classical market power can lead prices, under imperfect platform competition, to differ from their socially efficient levels.

¹An excellent recent survey of this literature, pioneered by Caillaud and Jullien (2003), Rochet and Tirole (2003), Evans (2003) and Parker and Van Alstyne (2005), is given by Rysman (2009).

²Armstrong (2006) says of the extension of a rich two-sided monopoly model, "A full analysis of this case is technically challenging in the case of competing platforms" (p. 671).

In order to understand both the indeterminacy we refer to above and our proposed solution, consider two issues: *Consumer Coordination* and *Armstrong's Paradox*. Consumer Coordination refers to the fact that, when there are consumption externalities, the optimal decision for any given consumer of which platform(s) to patronize, or, as we say hereafter, to "join", depends on the choices made by consumers on the other "sides of the market". Thus, for a given set of prices charged by platforms, there can be multiple equilibria. This has been a well known issue in the economics of networks since at least Katz and Shapiro (1985), and it can arise when there is only one platform as well under competition. Its crucial implication is that, when considering the price setting decisions of platforms, *demand is not necessarily well defined*.

By contrast, Armstrong's Paradox, first observed by Armstrong (2006), arises only when there are at least two platforms. It says that when platforms can charge prices for their products that are *functions* of the participation of consumers on the other side of the market, *virtually any outcome can be supported as a Subgame Perfect Equilibrium*.³ The assumption that platforms can charge participation-contingent, as opposed to just flat, prices seems very reasonable in view, for example, of the per transaction pricing of credit cards, the per game pricing of video games and the per click pricing of search engines.

Insulated Equilibrium resolves these two issues by positing that platforms will engage in a form of "robust implementation", which can be summarized as follows.

1. Holding fixed the strategies of all other platforms, each platform identifies its optimal feasible allocation on each side of the market.
2. From among the many price functions that weakly implement this desired allocation, each platform selects *Residually Insulating Tariffs*, which are special in that they remove any scope for problems of Consumer Coordination and thus guarantee that the chosen allocation will be realized.

When all platforms do this as a best response to one another, it is an Insulated Equilibrium. It is straightforward to see that their doing so eliminates issues of Consumer Coordination.⁴ We show that, for any allocation, a unique set of such tariffs exists, and thus Insulated Equilibrium completely eliminates the indeterminacy brought on by Armstrong's Paradox. An implication of this is that the demand system and observation of equilibrium prices suffice to identify firms' marginal costs, returning us, from an identification perspective, to the familiar confines of standard static industrial organization.

³The argument for why this is true is quite subtle and thus we defer it to the body of the paper (see Section 5.3). Note, however, that it is very closely related to the argument of Klemperer and Meyer (1989).

⁴In this regard, we build on Weyl (2010), which introduces Insulating Tariffs in a monopoly setting.

This Myersonian (1981) approach whereby platforms are seen to “select” allocations and then implement them is underpinned in our setting by a classical result from quasi-linear general equilibrium theory, which we extend to our smooth, large, two-sided economy. We show that for a given equilibrium allocation of consumers to platforms, there is a unique supporting price vector. This result ensures the validity of this way of framing the problem, as it guarantees that the mapping from allocation to profits is indeed a function. Note, however, that despite this *framing* in terms of allocation, platforms’ *conduct* is Nash-in-prices, and that our model generalizes the differentiated Bertrand (and not Cournot) model of competition to multi-sided markets.⁵

Armed with these technical tools, we then analyze first-order conditions characterizing Insulated Equilibria. We show that there are two fundamental forces governing the relationship between the equilibrium allocation and the optimum. One of these forces is the classical Cournot (1838) *market power distortion* and the other is the *Spence distortion*, owing its name to the seminal analysis in Spence (1975) of a monopolist’s choice of quality. While, as a general matter, the effect that an intensification of competition has on the market power distortion is well known, the effect of such an intensification on the Spence distortion depends crucially on the structure of consumer heterogeneity and the distribution of the demand.

Importantly, our model accommodates such issues. We make no specific assumptions on (i) functional forms for firm costs or distribution of user preferences, (ii) the dimensions of heterogeneity of consumer preferences, (iii) the dimensions of heterogeneity of consumer preferences, (iv) the number and symmetry of platforms or (v) consumption patterns (i.e., single versus multi-homing). Instead we assume only mild “regularity” conditions in these dimensions.

While the model we consider throughout most of the paper has exactly two sides and no externalities within sides, we show at the end how these restrictions can be easily relaxed. It should thus be possible to use our framework to evaluate models of competition among firms in markets *with* consumption externalities that are no more restrictive than the models typically used to study competition in markets *without* such externalities. We therefore believe that our approach has the potential to enrich the applied analysis of platform competition and to significantly inform regulatory policy in such markets.

Section 2 frames our argument, previewing the payoffs of our approach and relating it to other work. Section 3 develops the formal model. Following this, we derive our main technical results in Sections 4 and 5, though longer and less instructive proofs are

⁵The conduct we assume *on a given side of the market* can easily be adapted. Since issues of one-sided conduct are not our focus here, and due to its use in applied work, we stick throughout to Bertrand.

left to an appendix. We discuss first-order conditions in Section 6 and conditions for the stability, uniqueness and existence of equilibrium in Section 7. Section 8 considers several applications and extensions: 8.1 models mergers between platforms, 8.2 covers the aforementioned generalizations, and 8.3 sketches a way forward for using our approach to perform structural estimation in multi-sided industries. Section 9 concludes.

2 This Paper’s Contribution in Context

In this section, we first preview the payoffs delivered by the model and solution concept that we develop in the subsequent sections of the paper. We then describe the ways in which our results enrich previous literature on multi-sided platforms.

2.1 Platform Pricing

Our model and solution concept provide precise and intuitive, but general, first-order conditions characterizing the market equilibrium of competing multi-sided platforms. These generalize the classical conditions for Nash-in-prices equilibrium in a differentiated products industry, to a multi-sided setting. They also nest, as a special case, the optimality conditions for a multi-sided monopolist of Weyl (2010) (W10).

Let j denote a particular firm, \mathcal{I} denote a side of the market, P denote price, N denote the fraction of consumers participating, $C_{\mathcal{I}}^j$ denote marginal cost to platform j of serving side \mathcal{I} , μ denote the inverse (partial) hazard rate of demand (the standard market power distortion often denoted by $P'Q$). Let \mathbf{D} represent the *diversion ratio matrix* with j, k^{th} element $\frac{\partial N^k}{\partial p^j} / \left(-\frac{\partial N^j}{\partial p^j}\right)$, the fraction of sales lost by platform j in response to an increase in its price increase that are recouped by platform k . Finally, let $\mathbf{M}_{j,\cdot}$ and $\mathbf{M}_{\cdot,j}$ denote, respectively, the j^{th} row and column of a matrix \mathbf{M} . The first-order condition for insulated equilibrium pricing is that, for each firm j , on each side of the market \mathcal{I} ,

$$\underbrace{P^{\mathcal{I},j} = C_{\mathcal{I}}^j + \mu^{\mathcal{I},j}}_{\text{Exactly as in a standard market}} - N^{\mathcal{J},j} \cdot \underbrace{\left(\left[\begin{array}{c} -\frac{\partial N^{\mathcal{J}}}{\partial \mathbf{P}^{\mathcal{J}}} \\ \frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}} \end{array} \right]^{-1} \left[\begin{array}{c} \frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}} \end{array} \right] \right)_{j,\cdot}}_{\approx \text{Average value to marginal opposite-side consumers}} \cdot [-\mathbf{D}_{\cdot,j}^{\mathcal{I}}], \quad (1)$$

where $\mathcal{J} \neq \mathcal{I}$. Note that the first terms come directly from classical industrial organization theory: price equals marginal cost plus the optimal differentiated Bertrand mark-up, the inverse partial hazard rate of demand. To interpret the additional “two-sided markets”

term, it is useful to compare it to that arising in the monopoly setting of W10 where $\left(\left[-\frac{\partial N^{\mathcal{J}}}{\partial P^{\mathcal{J}}} \right]^{-1} \left[\frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}} \right] \right)_{j,\cdot} \cdot -\mathbf{D}^{\mathcal{I}}_{\cdot,j}$ collapses to the *average* willingness of a *marginal* consumer on side \mathcal{J} to pay for the participation of a *marginal* consumer on side \mathcal{I} . This is the part of the externality created by this marginal side \mathcal{I} consumer that the platform can extract per consumer on side \mathcal{J} .

As we discuss extensively in Section 6.2, the broader expression that we show is valid under oligopoly is a natural extension of this same notion. The j^{th} diagonal entry of $\frac{\partial N^{\mathcal{J}}}{\partial P^{\mathcal{J}}}$ is the density of side \mathcal{J} users just indifferent between consuming a bundle including platform j and consuming a bundling excluding it: the mass of j 's marginal users. This matrix's j, k^{th} entry for $j \neq k$ is the mass of users indifferent between consuming a bundle including platform j but not platform k and consuming a bundle including platform k but not platform j : the mass of "switching" users marginal between j and k . Thus, $\frac{\partial N^{\mathcal{J}}}{\partial P^{\mathcal{J}}}$ is a natural multi-product extension of the "mass of marginal users". Similarly, we show that the j^{th} diagonal entry of $\frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}}$ is the product of the density of j 's marginal users and the average value these place on a marginal side \mathcal{I} user, while its j, k^{th} entry for $j \neq k$ is the density of j, k switching users multiplied by the average value such users would place on a marginal side \mathcal{I} user joining platform k , if they were to join k . Thus this matrix is a natural extension of the product of the mass of marginal users and their average marginal valuations for users on the other side.

Therefore $\left(\left[-\frac{\partial N^{\mathcal{J}}}{\partial P^{\mathcal{J}}} \right]^{-1} \left[\frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}} \right] \right)_{j,\cdot} \cdot [-\mathbf{D}^{\mathcal{I}}_{\cdot,j}]$ generalizes W10's monopoly pricing rule to the oligopoly setting, in the same way that, for example, the matrix equation for a multivariate regression generalizes the ratio of the covariance to the variance of the regressor. For example, in the case considered by Armstrong (2006), when all marginal values are constant and homogeneous across all individual-platform pairs, this quantity collapses to exactly that marginal value.

This allows us to consider the impact of intensified competition on the relationship between social and private objectives. While it is well known that, in standard markets, intensified competition will reduce incentives for distortionary above-cost pricing, in platform settings, market power introduces a second Spence (1975) distortion into pricing, as firms have an incentive to focus on externalities perceived by marginal consumers, rather than those perceived by all consumers. As we argue in Section 6.3, whether competition is likely to alleviate or exacerbate the Spence distortion depends on the nature of heterogeneity *among platforms*.

If platforms differ along horizontal or vertical dimensions orthogonal to consumer valuations of externalities, then competition is likely to ameliorate the Spence distortion as it leads platforms to attend switching rather than exiting users' valuation of externalities,

which are more likely to be representative of the full population of participating users. However, if platforms differentiate themselves vertically in the number of users they have on the other side of the market, users switching between the platforms are likely to have valuations for users on the other side that are below those of the “high quality” and that are above those of the “low quality” platform.

A canonical issue in competition policy is the evaluation of the impacts on consumer welfare of a potential merger. Evaluating a merger between two multi-sided platforms requires extending standard merger evaluation techniques to accommodate both the multi-product nature of multi-sided platforms and, more importantly, the additional presence of Spencian welfare effects. To illustrate how our model enables this, in Section 8.1 we extend Jaffe and Weyl (2010b) (JW)’s quantification of the standards embodied in the US government’s recently released merger guidelines to the context of platform competition (U.S. Department of Justice and the Federal Trade Commission, 2010).

In the multi-sided extension of the JW formula, the marginal opportunity costs of sales created by the merger, often called *Upward Pricing Pressure* or “UPP” (Farrell and Shapiro, 2010), are multiplied by pass-through rates to obtain estimates of price effects and then by quantities to obtain a local approximation to the effect on consumer welfare. In our setting, two additional forces emerge. First, the marginal opportunity cost of a sale now incorporates not only the standard value of diverted sales that determine UPP, but also the marginal harm a competitor would incur by offering decreased externalities to consumers on the opposite side, as a result of these diverted sales.

Second, the effect on consumer welfare is not only through the direct harm brought by the incentive for firms to raise prices; changes in the levels of externalities due to changes in participation on a given side also affect consumer welfare on the opposite side of the market. Following Spence’s logic, these harms are proportional to the change in the number of consumers on the other side of the market multiplied by (a certain version of) the difference between the value that marginal consumers on the other side of the market place on those externalities (which is extracted by the platform) and the value placed on the externalities by average consumers on the other side. Our model can also be used for other standard comparative static exercises. In particular, we look forward, in section 8.3, to the future development of special cases of our model that can easily be estimated.

2.2 Context

Why were these simple and general results not feasible in prior work? The two issues of multiplicity we discussed in the introduction, Consumer Coordination and Armstrong’s

Paradox, stymied the tractability of a general model of multi-sided platforms. We now discuss these two challenges as well as the dimensions along which we generalize, and fail to generalize, with respect to the existing literature.

Caillaud and Jullien (2003), Ellison and Fudenberg (2003), Ellison, Fudenberg, and Möbius (2004), Hagiu (2006), Ambrus and Argenziano (2009), Lee (2010) and Anderson, Ellison, and Fudenberg (2010) study the coordination of consumers given prices. Many equilibria are possible in these settings and the payoffs in the game played by platforms depend sensitively on which equilibrium consumers are assumed to coordinate on by refinements in the second stage. Our approach instead attributes the role of coordination to *platforms*, who have a large stake in the matter and significant powers over the outcome, rather than to consumers, who are multitudinous and dispersed.⁶

Armstrong's Paradox, whereby infinitely many allocations can be supported as equilibria, or not, among competing platforms, stems from Proposition 3 of Armstrong (2006). Armstrong argues that firms' best responses are determined by the exact degree to which competitors' prices on one side of the market respond to changes in the number of consumers on the other side, but that only the level of prices, and not the slope of such responses, is tied down by equilibrium. We discuss this issue in more detail in Section 5.3. While the motivation for Insulated Equilibrium is the view that it is reasonable to expect platforms to pin down consumer behavior, if platforms indeed do this, then Armstrong's Paradox is resolved.⁷

Regarding the ways in which our model generalizes with respect to existing literature,⁸ a crucial aspect is its accommodation of arbitrary preference heterogeneity among consumers. W10 shows that the comparative statics of a model of a two-sided monopolist depend crucially on whether consumers differ primarily in their valuations for *membership* or in their valuations for *interaction* with other consumers. However, with little or no empirical basis for these assumptions, in prominent theoretical models in which platforms compete for consumers, such as those in Anderson and Coate (2005), Armstrong (2006), Armstrong and Wright (2007) and Peitz and Valletti (2008), and in econometric works, such as those of Rysman (2004) and Kaiser and Wright (2006), consumers are assumed to

⁶Dybvig and Spatt (1983) discuss a seemingly similar but in most cases quite different notion of insurance.

⁷In a recent paper, Reisinger (2010) proposes an alternative approach to getting around Armstrong's Paradox, in a setting with the particular assumptions of Armstrong's model and abstracting from the issue of multiplicity of Consumer Equilibria. In essence, this approach points out that, under two-part tariffs, introducing heterogeneity in consumers' interaction behavior is equivalent to allowing them to price discriminate in a regime with flat pricing. This, in turn, ties down platforms' competitive responses to one another.

⁸The best-known model of a monopoly platform is perhaps that of Rochet and Tirole (2006), which is generalized by W10, while the best-known model of competing platforms is likely Armstrong (2006).

be homogenous in their interaction values. A crucial implication of such a setup is that it rules out, *ex hypothesi*, the Spence distortion (discussed in the introduction and in Section 6). Our model provides a framework for analyzing the interaction between this distortion and variation in the competitive environment.⁹

Our approach also does not require making assumptions on functional forms of, for instance, the distribution of consumer preferences or the platforms' cost curves. In contrast, a common assumption in models of competition, following Armstrong (2006), has been that of a two-sided Hotelling (1929) setup giving rise to linear demand. Several benefits come from relaxing this assumption, including compatibility with the approach taken in the empirical industrial organization literature, which we discuss in Section 8.3, reduced vulnerability to the forms of criticism given in Werden, Froeb, and Scheffman (2004) to using such models as bases for arguments in antitrust cases. Furthermore, Jaffe and Weyl (2010a) have recently shown that with more than two firms, it is impossible for a discrete choice model to generate linear demand.

This paper's framework does not restrict the number of firms that can compete nor does it require them to be symmetric, making the model more realistic. In addition, not requiring symmetry among platforms protects against the possibility of making unusual-seeming findings that may be driven by this assumption (see, for instance, Amir and Lambson (2000) as well as the criticism in BLP of the substitution patterns in the logit model). This is particularly true in models of competing platforms, in which equilibria can be sensitive to "tipping", as discussed in Sun and Tse (2007). Moreover, our model is amenable to merger analysis, which cannot be performed using models in the style of Armstrong (2006), due to their setup with two platforms and non-market-expanding demand.¹⁰

Our approach gives consumers free reign over their consumption choices, as they can select any bundle of platforms they find optimal. Existing models in which consumers "multi-home" (Caillaud and Jullien, 2003; Rochet and Tirole, 2003; Rysman, 2004; Armstrong, 2006; Doğanoglu and Wright, 2006; Armstrong and Wright, 2007), consider cases with just two platforms and/or exogenously impose single-homing on one side of the market. In the one-sided discrete choice literature, works such as Hendel (1999) and Gentzkow (2007) have moved towards incorporating such flexibility into consumers' choice set, and our approach follows in this spirit.

⁹In Bedre-Defolie and Calvano (2009) and White (2009), consumers are heterogenous in both dimensions, but they do not learn their interaction benefits until *after* they have selected a platform.

¹⁰The aforementioned survey, Rysman (2009), speaks of the lack of such a framework, "Naturally, if we were to analyze the merger between two platform firms, we would need to account for complex two-sided issues that arise" (p. 137).

The restrictions that have so far been present in the theory of platform competition have made carrying out empirical studies of such industries more difficult, forcing authors to adapt to the circumstances of their studies in somewhat constrained ways. For instance, in Cantillon and Yin (2008), the authors lack a model to predict platforms' equilibrium prices and instead take them as exogenous, while Argentesi and Filistrucchi (2007) and Wilbur (2008) use a reduced form inverse demand functions to model one side of the market. Our model, we hope, can serve as a basis for applied studies and can thus help to solve such difficulties.

While our model generalizes in the dimensions listed above, it retains two important assumptions that are typical of models in the literature on multi-sided platforms. First, we employ what Economides (1996) refers to as the “macro approach” to modeling networks, taking as exogenous the interaction among the consumers on different sides, once they join platforms and assuming consumer payoffs from joining a set of platforms depends only on the number of consumers participating and the payment to the platform(s). This approach brings useful generality when the interventions one considers are unlikely to affect the microstructure of interactions. However, if one's focus is on policies aimed at microstructure, an explicit model of such is crucial, as in Nocke et al. (2007), Hagiu (2009b), White (2009) and Weyl and Tirole (2010).

Second, we assume all consumers on a given side are homogenous in the externalities they cause. That is, a consumer from one group cares about how *many* consumers of another group join each platform, but not *which* consumers these are. This restrictive and unrealistic assumption has been relaxed in a few specific contexts (Chandra and Collard-Wexler, 2009; Hagiu, 2009a; Rochet, 2010; Gomes, 2009; Athey et al., 2010), but work in progress by Veiga and Weyl (2010) provides the first general approach to incorporating heterogeneous externalities. They show that heterogeneity of externalities matter in pricing to the extent that valuation of participation on the other side covaries (on the margin) with the value of externalities brought by a consumer.¹¹

Other issues that we do not consider include dynamics and price discrimination within sides. Anderson and Coate (2005), Hagiu (2006), Sun and Tse (2007), Lee (2010) all include consideration of the former, while Gomes (2009), Doğanoglu and Wright (2010) and Hagiu and Lee (forthcoming) deal with the latter.

¹¹We are optimistic that such an extension can be incorporated without great difficulty into our framework, but given the early stage of this research on heterogenous externalities, we do not include it in the current version of this paper.

3 The Model

There is a set, $\mathcal{M} = \{1, \dots, m\}$, of “two-sided platforms”, with elements indexed by j . These firms serve two separate groups of “consumers” or “users”, of measure 1, said to be on opposite “sides of the market”, \mathcal{A} and \mathcal{B} , indexed by \mathcal{I} . For concreteness, consider as examples

- payment card issuers, whose two groups of consumers are shoppers that carry the card and use it to make purchases and merchants that accept the card
- publishers of newspapers that sell their final product to readers and that sell space in their pages to advertisers
- jobs-listing websites, catering to both job seekers and employers.

Consumers on each side of the market can choose to “join” any combination of platforms, i.e., they pick an element in the power set of the set of platforms, $\wp(\mathcal{M})$. We denote the particular subset or “bundle” of platforms that consumer i on side \mathcal{I} chooses by $\mathcal{M}_i^{\mathcal{I}} \in \wp(\mathcal{M})$. We sometimes refer to the non-empty elements of $\wp(\mathcal{M})$ in an order, from 1 to $2^m - 1$, which can be arbitrarily chosen but which, once established, we refer to as the *bundle labeling*.

The unifying and distinguishing feature of the types of firms we might refer to as “platforms” is the presence of some form of externality across their different groups of consumers. To capture such externalities, or “cross-network effects”, we assume that the payoff to a consumer on side \mathcal{I} from joining a given set of platforms depends, in some way, on the number of consumers of the opposite side of the market, $\mathcal{J} \equiv -\mathcal{I}$, that join each of the platforms in this bundle.¹² Intuitively, one may think of the number of side \mathcal{J} consumers participating on each platform in a bundle as, from the standpoint of a consumer on side \mathcal{I} , a *characteristic* of that bundle, partially determining its perceived quality. We now introduce a statistic that keeps track of these characteristics.

Definition 1. A Coarse Allocation, $\mathbf{N} \equiv (\mathbf{N}^{\mathcal{A}}, \mathbf{N}^{\mathcal{B}}) \in [0, 1]^{2^m}$, specifies the total measure or “number” of consumers participating on each side of each platform. We denote a generic element by $N^{\mathcal{I},j}$.

Demand. Consumers have quasi-linear utility, and their optimization problem takes the form of a discrete choice over bundles of platforms. We write the payoff to user i on

¹²As we discuss in Section 2.2, we do not explicitly model the “interaction” that may take place among users on opposite sides.

side \mathcal{I} from joining bundle of platforms \mathcal{X} as

$$v^{\mathcal{I}}(\mathcal{X}, \mathbf{N}^{\mathcal{J}}, \boldsymbol{\theta}_i^{\mathcal{I}}) - \sum_{j \in \mathcal{X}} p^{\mathcal{I}, j},$$

where $\boldsymbol{\theta}_i^{\mathcal{I}} \in \Theta^{\mathcal{I}}$ denotes consumer i on side \mathcal{I} 's "type". The set of side \mathcal{I} types, $\Theta^{\mathcal{I}} = \mathbb{R}^{L^{\mathcal{I}}}$, $2^m - 1 \leq L^{\mathcal{I}} \in \mathbb{N}$, does not impose any particular restrictions on the dimensions in which consumers can be heterogenous.¹³ The function $v^{\mathcal{I}} : \wp(\mathcal{M}) \times [0, 1]^m \times \Theta^{\mathcal{I}} \rightarrow \mathbb{R}$ is thus a map to a consumer's willingness to pay from each possible consumption choice, the characteristics of the available goods and the user's individual characteristics. $p^{\mathcal{I}, j}$ denotes the total price a user on side \mathcal{I} must pay to join platform j , the details of which we discuss below, when defining platforms' strategies. Assumption 1 further characterizes the demand system.

Assumption 1. *The functions $v^{\mathcal{I}}$, $\mathcal{I} = \mathcal{A}, \mathcal{B}$, jointly with their domains, have the following properties:*

1. *Smoothness: $v^{\mathcal{I}}$ is C^2 in all dimensions of its second and third arguments.*
2. *Gross Substitutes: For all $N^{\mathcal{J}} \in [0, 1]$, if platform j is in a side \mathcal{I} consumer's optimal bundle and the price of some other platform in this bundle increases, then j remains in the consumer's optimal bundle.*

This is the weakest condition known in the literature that assures the existence of an equilibrium price vector in the one-sided analogue to our model. In the context of our model, we believe this assumption to be very realistic.¹⁴

3. *Full Support: For all $N^{\mathcal{J}} \in [0, 1]$ and utility profiles $\mathbf{u}^{\mathcal{I}} \in \mathbb{R}^{2^m - 1}$ over all bundles satisfying gross substitutes, $\exists \boldsymbol{\theta} \in \Theta^{\mathcal{I}}$ such that $v^{\mathcal{I}}(\cdot, N^{\mathcal{J}}, \boldsymbol{\theta})$ takes on the value $\mathbf{u}^{\mathcal{I}}$.*
4. *No Externalities to Outsiders: if $j \notin \mathcal{X}$ then $v^{\mathcal{I}}(\mathcal{X}, N^{\mathcal{J}}, \boldsymbol{\theta})$ is independent of $N^{\mathcal{J}, j}$. This is an intuitive assumption reflecting the idea that consumers on opposite sides of the market do not "interact", unless they join at least one common platform. This affords a clean*

¹³For example, the natural extension of the primary preferences discussed in W10 (see section I.B.) would include, for each bundle of platforms, a *membership* benefit, or dummy variable, and, for each platform within a given bundle, an *interaction* coefficient multiplying the number of side \mathcal{J} consumers on that platform.

¹⁴We are confident that this assumption can be relaxed and have sketched an argument of how to do so. However, this is an active issue in general equilibrium and matching theory that is largely orthogonal to our focus in this paper. On this issue, see Kelso and Crawford (1982), Gul and Stacchetti (1999) and Hatfield, Kominers, Nichifor, Ostrovsky, and Westkamp (2010).

interpretation of our results, although somewhat modified versions of them can be derived without this assumption.¹⁵

5. Normalization: for all $\theta \in \Theta^I$, $v^I(\emptyset, N^I, \theta) = 0$. For all consumers the “outside option” gives a payoff normalized to zero.

Let $f^I : \Theta^I \rightarrow \mathbb{R}$ be the probability density function of user types on side $I = \mathcal{A}, \mathcal{B}$, satisfying $\int_{\Theta^I} f^I(\theta) d\theta = 1$. We assume that each function f^I is C^1 and has full support.

Supply. Platform j ’s profits are given by

$$\Pi^j \equiv P^{\mathcal{A},j} N^{\mathcal{A},j} + P^{\mathcal{B},j} N^{\mathcal{B},j} - C^j(N^{\mathcal{A},j}, N^{\mathcal{B},j}),$$

where $C^j(N^{\mathcal{A},j}, N^{\mathcal{B},j})$ denotes platform j ’s costs as a function of the number of users on each side and is assumed to be C^2 in both arguments.

Timing. Platforms move first, simultaneously. Then, having observed the platforms’ moves, all consumers simultaneously choose which platforms to join.

Strategies. As we mention in the introduction, we employ an *allocation approach* to solve the game. Two features of this approach that we regard to be particularly appealing are, first, that it allows for *consideration* of a very large strategy space for platforms, and, second, that while doing so, it does *not* require explicitly keeping track of platforms’ potentially very complicated pricing functions or solving an optimization problem using calculus of variations. In Section 4, we explain this approach in detail and relate it to previous work.

We thus allow for each platform to charge a tariff to consumers on side I that is a function of the *entire coarse allocation* on side \mathcal{J} . A (pure) strategy for platform j , $\sigma^j \equiv (\sigma^{\mathcal{A},j}(N^{\mathcal{B}}), \sigma^{\mathcal{B},j}(N^{\mathcal{A}}))$, is a pair of such functions. Previous work on platform competition of which we are aware all restricts the prices that a given platform charges its consumers on one side to be, at most, a function of the number of opposite-side consumers that *it* serves. Moreover, typically (see, e.g., Armstrong (2006)) these prices are assumed to be affine in this argument. When using our approach to solve the game, however, imposing such restrictions does not buy us anything. Moreover, in practice, platforms may have an incentive to charge quite sophisticated, market-dependent tariffs.¹⁶

Formally, $\sigma^{I,j} : [0, 1]^m \rightarrow \mathbb{R}$. In order to ensure differentiability of each platform’s residual profits, we assume that all platforms’ price functions, $\sigma^{I,j}$, are C^2 . While this as-

¹⁵In this regard, our model contrasts with that of Segal (1999), which studies the effects of positive versus negative “externalities on nontraders”.

¹⁶For instance, the prices that television networks charge for ads shown during a given program frequently depend on that program’s success *versus other programs* in drawing an audience. “Penetration pricing”, and other practices that are observed in network industries seem to reflect such complexity of pricing strategies.

sumption facilitates our approach, it will become clear that, under Insulated Equilibrium, platforms never have an incentive to deviate to charging tariffs that violate this assumption. Let Σ denote the set of all pairs of C^2 functions, and let Σ^m denote the m^{th} cartesian power of this set. We denote the profile of strategies of the entire set of platforms by $\sigma \in \Sigma^m$. A profile of platform strategies is a function, $\sigma : [0, 1]^{2m} \rightarrow \mathbb{R}^{2m}$, which we assume can be written $\sigma(N) \equiv (\sigma^{\mathcal{A}}(N^{\mathcal{B}}), \sigma^{\mathcal{B}}(N^{\mathcal{A}}))$. Under this notation, $\sigma^{\mathcal{I}}(N^{\mathcal{J}}) : [0, 1]^m \rightarrow \mathbb{R}^m$ maps from the coarse allocation on side \mathcal{J} to the vector of prices charged by all platforms on side \mathcal{I} .

Consumers react to platforms' announcement of price functions. Thus, a pure strategy for consumer i on side \mathcal{I} , is a *functional*,¹⁷ which we denote by $\mathcal{M}_i^{\mathcal{I}}[\sigma]$, where $\mathcal{M}_i^{\mathcal{I}} : \Sigma^m \rightarrow \wp(\mathcal{M})$. To denote a *Side Strategy Profile*, for the set of consumers on side \mathcal{I} , we define the correspondence $\mathcal{M}^{\mathcal{I}}(\theta^{\mathcal{I}}, [\sigma])$. To avoid having to distinguish between economically equivalent outcomes where sets of consumers of measure zero can select different bundles of platforms among which they are indifferent, we impose Assumption 2.

Assumption 2. *Strategy profiles adopted by consumers satisfy the following properties:*

1. *Purity: In every subgame, each consumer takes some action with probability 1.*
2. *Symmetry: All agents sharing a common type adopt the same strategy.*
3. *Tie-breaking: When indifferent, all agents choose to join the set of platforms that comes first in the established bundle labeling.*

It follows from the Purity and Symmetry components of Assumption 2 that $\mathcal{M}^{\mathcal{I}}$ is a functional, where $\mathcal{M}^{\mathcal{I}} : \Theta^{\mathcal{I}} \times \Sigma^m \rightarrow \wp(\mathcal{M})$ identifies all side \mathcal{I} consumers' behavior in response to all $\sigma \in \Sigma^m$. We denote the *Marketwide* consumer strategy profile by $\widehat{\mathcal{M}}(\theta, [\sigma])$, where $\widehat{\mathcal{M}} : \{\Theta^{\mathcal{A}} \times \Theta^{\mathcal{B}}\} \times \Sigma^m \rightarrow \wp(\mathcal{M})$.

Above, we defined a coarse allocation to be the number of consumers participating on each platform, on each side of the market. We can now derive this statistic as a function of the strategy profiles of consumers and platforms. To do so, we define the functional $N : \{\widehat{\mathcal{M}}\} \times \Sigma^m \rightarrow [0, 1]^{2m}$, mapping from marketwide consumer strategy profile and platform strategy profile to coarse allocation. $N[\widehat{\mathcal{M}}, \sigma]$ has generic elements

$$N^{I,j}[\widehat{\mathcal{M}}, \sigma] = \int_{\{\theta^{\mathcal{I}} \in \Theta^{\mathcal{I}} : j \in \widehat{\mathcal{M}}(\theta^{\mathcal{I}}, [\sigma])\}} f^{\mathcal{I}}(\theta) d\theta.$$

¹⁷By "functional", we mean a function that takes a function as at least one of its input-arguments. Hereafter, we surround arguments of functionals with square brackets when the entire function is to be taken as the input. E.g., $z[f]$ indicates that z depends on the entire shape of the function f . In contrast, when the input to a function is a function *evaluated* at a particular point, we surround the argument with ordinary parentheses. E.g., $Z(f(x))$ indicates that Z depends on the value $f(\cdot)$ takes when evaluated at x .

4 Allocation \implies Price

In this section, we prove Theorem 1, which provides the foundation for what we call the *allocation approach* to analyzing the game. We refer to this approach as such because, despite the fact that we assume platforms' *conduct* vis-à-vis one another, on each side of the market, to be Bertrand (Nash-in-prices), we *represent* their profit maximization problem as a choice of quantity or allocation on each side of the market.¹⁸ We now explain the motivation for doing this.

Consider the *Consumer Game* that takes place in the second stage, after platforms have announced their strategies. It is well known¹⁹ that such games can have multiple Nash Equilibria, since the optimal bundle for consumers on one side of the market can depend on the actions of consumers on the other side. (See Figure 1.) This implies that, in order to evaluate platforms' profits as functions of their strategies, one must have some criteria for selecting which Nash Equilibrium prevails in the subsequent Consumer Game.

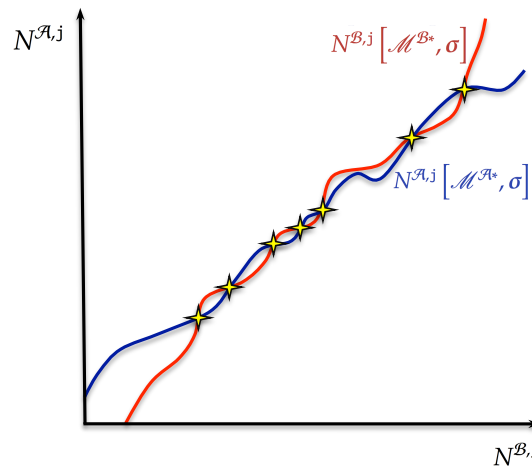


Figure 1: A hypothetical case with multiple equilibria in the second stage, for a given profile of platform strategies, σ .

The need for such criteria potentially adds significant technical complication, and it is not at all obvious how to begin asking which criteria are the most apt. Even if this issue were resolved and “demand” was defined as a function of platforms' strategies, there would be the additional difficulty that platforms' profits are functionals whose inputs are price functions. Consequently, solving for platforms' optimal strategies would be a potentially intractable problem.

Fortunately, thanks to the result of Theorem 1, we can safely ignore these issues. This result says that, for any coarse allocation that may arise as a Nash Equilibrium among consumers in the last stage of the game, there is a unique vector of total platform prices

¹⁸Indeed, the merits of such a representation are the same regardless of platforms' “intra-side” conduct.

¹⁹We cite papers that discuss this at the beginning of Section 2.2.

that is consistent with – i.e. that can support as a Nash Equilibrium – this coarse allocation. The crucial implication of this is that platform profits can be written as *functions* of the coarse allocation, and the game can be fully analyzed in this simpler manner.

This result allows for the generalization to oligopoly of the allocation approach, which W10 employs in the simpler setting of a monopoly platform. Note that the technique we use here is quite similar to those used by Myerson (1981) and Riley and Samuelson (1981) to analyze properties of optimal auctions. In particular, it makes use of restrictions implied by the equilibrium behavior on the part of consumers (bidders) in order to steer clear of extraneous aspects of the pricing mechanisms used by platforms (auctioneers).

We now formally describe the second stage of the game, in order to arrive at this result. Let a Consumer Game be defined as a subgame that takes place once the platforms' strategy profile σ has been determined. Taking σ as a parameter, let $\widehat{P^{I,\mathcal{X}}}(\mathbf{N}^{\mathcal{J}}[\mathcal{M}^{\mathcal{J}};\sigma]) \equiv \sum_{j \in \mathcal{X}} \sigma^{I,j}(\mathbf{N}^{\mathcal{J}}[\mathcal{M}^{\mathcal{J}};\sigma])$ denote the sum of prices charged to a consumer on side I , joining the set of platforms $\mathcal{X} \in \wp(\mathcal{M})$. Let U_i^I denote the net payoff to a consumer of type θ_i^I , joining this bundle, where

$$U_i^I(\mathcal{X}, \mathbf{N}^{\mathcal{J}}; \sigma) \equiv v^I(\mathcal{X}, \mathbf{N}^{\mathcal{J}}[\mathcal{M}^{\mathcal{J}}; \sigma], \theta_i^I) - \widehat{P^{I,\mathcal{X}}}(\mathbf{N}^{\mathcal{J}}[\mathcal{M}^{\mathcal{J}}; \sigma]).$$

Finally, let $\mathcal{M}^{I*} : \Theta^I \times [0, 1]^m \times \Sigma^m \rightarrow \wp(\mathcal{M})$ denote the *Best Response Correspondence* for side I , where $\mathcal{M}^{I*}(\theta_i^I, \mathbf{N}^{\mathcal{J}}; \sigma) \in \arg \max_{\mathcal{X} \in \wp(\mathcal{M})} U_i^I(\mathcal{X}, \mathbf{N}^{\mathcal{J}}; \sigma)$, $\forall \theta_i^I \in \Theta^I$. Note that by the Tie-breaking component of Assumption 2, \mathcal{M}^{I*} is a function. We can now state our solution concept for a Consumer Game.

Definition 2. In Consumer Game, σ , a marketwide consumer strategy profile, $\widehat{\mathcal{M}}$, forms a Consumer Nash Equilibrium (CNE) if the associated side strategy profiles, $\{\mathcal{M}^I\}_{I=\mathcal{A},\mathcal{B}}$, satisfy $\mathcal{M}^I(\theta^I; [\sigma]) = \mathcal{M}^{I*}(\theta^I, \mathbf{N}^{\mathcal{J}}[\widehat{\mathcal{M}}^{\mathcal{J}}; \sigma]; [\sigma])$, $\forall \theta^I \in \Theta^I$.

To proceed to Theorem 1, we first define *Gross Consumer Surplus* on side I , as a function of side I consumers' best response strategy profile and the coarse allocation on side \mathcal{J} . We denote this by $V^I : \{\mathcal{M}^I\} \times [0, 1]^m \rightarrow \mathbb{R}$, where

$$V^I([\mathcal{M}^{I*}], \mathbf{N}^{\mathcal{J}}) \equiv \sum_{\mathcal{X} \in \wp(\mathcal{M})} \int_{\theta^I : \mathcal{M}^{I*}(\theta^I, \mathbf{N}^{\mathcal{J}}) = \mathcal{X}} v^I(\mathcal{X}, \mathbf{N}^{\mathcal{J}}, \theta) f(\theta) d\theta. \quad (2)$$

The right-hand side of expression (2) is the sum over bundles of platforms of side I consumers' gross payoffs, given that they are best-responding both to platform prices and to the allocation of consumers on side \mathcal{J} .

Since, in a CNE, consumers maximize their utility given the prevailing prices, it must

be the case that allocation of bundles to consumers is Pareto optimal, given the quantity constraints imposed by the induced side \mathcal{I} coarse allocation. Since consumers have quasi-linear utility, a Pareto optimal allocation also maximizes the sum of consumer utility, subject to the same quantity constraints. Therefore, we can write Gross Consumer Surplus as a function, $V^I : (0, 1)^{2m} \rightarrow \mathbb{R}$, of the side \mathcal{I} coarse allocation, given by the solution to the following constrained maximization problem

$$V^I(\widetilde{N}^I, N^J) \equiv \max_{\mathcal{M}^I \in \{\mathcal{M}^I : N^I[\mathcal{M}^I] = \widetilde{N}^I\}} \sum_{\mathcal{X} \in \varphi(\mathcal{M})} \int_{\theta^I : \mathcal{M}^I(\theta^I) = \mathcal{X}} v^I(\mathcal{X}, N^J, \theta) f(\theta) d\theta. \quad (3)$$

Lemma 1 establishes the formal relationship between these objects.

Lemma 1 (Equilibrium \iff Maximization). *If $\widehat{\mathcal{M}}$ forms a CNE of a Consumer Game, σ , then the associated side strategy profiles $\{\mathcal{M}^I\}_{I=\mathcal{A}, \mathcal{B}}$ satisfy*

$$\mathcal{M}^I \in \arg \max_{\mathcal{M}^I \in \{\mathcal{M}^I : N^I[\mathcal{M}^I] = \widetilde{N}^I\}} \sum_{\mathcal{X} \in \varphi(\mathcal{M})} \int_{\theta^I : \mathcal{M}^I(\theta^I) = \mathcal{X}} v^I(\mathcal{X}, N^J[\mathcal{M}^J], \theta) f(\theta) d\theta.$$

Conversely, for any allocation $N \in (0, 1)^{2m}$ there exists some $\widehat{\mathcal{M}}^$ and some Consumer Game σ in which $\widehat{\mathcal{M}}^*$ forms a CNE such that $V^I([\mathcal{M}^{I*}], N^J) = V^I(N)$, for $I = \mathcal{A}, \mathcal{B}$.*

Proof. See Appendix A.1. □

Lemma 2 shows that the gross surplus function is differentiable in the elements of the coarse allocation; that is, the marginal utilities of consumption of the “representative consumer” exist.

Lemma 2 (Differentiability). *For any $N \in (0, 1)^{2m}$, $\frac{\partial V^I}{\partial N^{I,j}} \Big|_N$ exists. We denote it by $P^{I,j}(N)$.*

Proof. See Appendix A.2. □

We now use these results to show that every allocation implies exactly one price. A price vector supporting the allocation exists because preferences exhibit gross substitutes (Kelso and Crawford, 1982; Gul and Stacchetti, 1999). Only one such vector may exist because of the rich heterogeneity of consumers: at any price vector there is an *exactly* marginal consumer of zero mass whose utility it equated to the marginal gross surplus from an additional unit of that good. Because this marginal gross surplus is unique, by Lemma 2, any prices supporting a given allocation must be identical.

Theorem 1 (Allocation \implies Price). *If $\widehat{\mathcal{M}}$ is a CNE of consumer game, σ , then $\sigma(N[\widehat{\mathcal{M}}]) = P(N[\widehat{\mathcal{M}}])$. Conversely, if a consumer game σ has $\sigma(\widetilde{N}) = P(\widetilde{N})$ for some $\widetilde{N} \in (0, 1)^{2m}$ then there exists a CNE of σ , call it $\widehat{\mathcal{M}}^*$, with $N[\widehat{\mathcal{M}}^*] = \widetilde{N}$.*

Proof. First take the forward direction. By Lemma 1 we know that $\sigma(N[\widehat{\mathcal{M}}])$ must be a general equilibrium price vector for the economy with endowment $N[\widehat{\mathcal{M}}]$. But by the proof of Lemma 2 we know any such vector must be equated to $P(\widetilde{N})$.

In the reverse direction, simply apply the construction from the proof of the reverse direction of Lemma 1. □

5 Insulated Equilibrium

In this section, we consider the first stage of the game, in which platforms move. The previous section suggests that it is natural to think of platforms' problem as a choice of *allocation*. We briefly expound upon this idea in abstract terms, before defining our solution concept of *Insulated Equilibrium*.

An individual platform j , taking as given the strategies of other platforms, can identify a set of coarse allocations that can feasibly occur as CNE in the subsequent Consumer Game. Moreover, since each platform's profits can be written as a direct function of the CNE coarse allocation, platform j can identify the subset of feasible coarse allocations that, for it, are profit-maximizing. For the sake of exposition, let us suppose that this subset is single-valued and call its member N^j . Clearly, given the non-cooperative nature of the game, platform j fares strictly better if it is able to implement N^j than it does if some other coarse allocation arises. In light of this observation, the obvious issues are whether and how platform j can "robustly" implement its desired coarse allocation.

Holding fixed the strategies of the other platforms, there is an infinite set of strategies platform j can choose in order to implement N^j as *one of the* CNE in the subsequent Consumer Game. However, an arbitrary strategy that weakly implements N^j can also lead to other, less profitable allocations for platform j . This multiplicity of CNE arises from the fact that a coordination game takes place between consumers on opposite sides of the market. In the rest of this section, we show that by shaping its tariffs properly, a platform can completely eliminate the coordination game among consumers on opposite side of the market.

5.1 Definition

We now formally define *Residual Insulating Tariffs*,²⁰ which W10 introduces in the case of a monopoly platform. In the context of competing platforms, they work in the following way. In order to “pick” a coarse allocation on side \mathcal{J} , a platform charges an insulating tariff on side \mathcal{I} and thus guarantees that, on the latter side, a particular coarse allocation prevails.

Definition 3. *Given a profile of strategies of other platforms, σ^{-j} , platform j is said to charge a Residual Insulating Tariff on side \mathcal{I} if, $\forall N^{\mathcal{J}}, \widetilde{N}^{\mathcal{J}} \in [0, 1]$,*

$$N^{\mathcal{I},j} \left[\mathcal{M}^{\mathcal{I}*} \left(\theta^{\mathcal{I}}, N^{\mathcal{J}}, [\sigma] \right), \sigma \right] = N^{\mathcal{I},j} \left[\mathcal{M}^{\mathcal{I}*} \left(\theta^{\mathcal{I}}, \widetilde{N}^{\mathcal{J}}, [\sigma] \right), \sigma \right].$$

For a firm j to charge an insulating tariff on side \mathcal{I} , it must choose a price function, $\sigma^{\mathcal{I},j}(N^{\mathcal{J}})$, that, given the strategies of the other platforms, preserves the coarse allocation on side \mathcal{I} , regardless of the strategy profile adopted by side \mathcal{J} consumers. To see how such a function operates, consider the demand for platform j among side \mathcal{I} consumers, $N^{\mathcal{I},j}$. It can be written

$$N^{\mathcal{I},j} = N^{\mathcal{I},j} \left(\sigma^{\mathcal{I},j} \left(N^{\mathcal{J}} \right), \sigma^{\mathcal{I},-j} \left(N^{\mathcal{J}} \right), N^{\mathcal{J}} \right).$$

An insulating tariff, charged by firm j on side \mathcal{I} is thus a function, $\sigma^{\mathcal{I},j}(\cdot)$, that takes into account the shape of $N^{\mathcal{I},j}(\cdot, \cdot, \cdot)$ and the shape of other firms' side \mathcal{I} price functions, denoted by the vector $\sigma^{\mathcal{I},-j}(\cdot)$, in order to ensure that the output of $N^{\mathcal{I},j}$ is constant. Lemma 3 establishes the existence and uniqueness of an insulating tariff for firm j on side \mathcal{I} , given the side \mathcal{I} price functions announced by other firms and the coarse allocation on the other side of the market, $N^{\mathcal{J}}$.

Lemma 3 (Existence and Uniqueness of a Residual Insulating Tariff). *There exists a unique function, $\overline{P^{\mathcal{I},j}} \left(N^{\mathcal{J}}; \widetilde{N}, \left[\sigma^{\mathcal{I},-j} \left(N^{\mathcal{J}} \right) \right] \right)$, such that, $\forall N^{\mathcal{J}}, \forall \widetilde{N} \in (0, 1), \forall \sigma^{\mathcal{I},-j}$*

$$N^{\mathcal{I},j} \left(\overline{P^{\mathcal{I},j}} \left(N^{\mathcal{J}}; \widetilde{N}, \left[\sigma^{\mathcal{I},-j} \left(N^{\mathcal{J}} \right) \right] \right), \sigma^{\mathcal{I},-j} \left(N^{\mathcal{J}} \right), N^{\mathcal{J}} \right) = \widetilde{N}.$$

Moreover, $\overline{P^{\mathcal{I},j}}$ is C^2 in all dimensions of its first argument.

Proof. For existence, note that (i) $N^{\mathcal{I},j}(\cdot, \cdot, \cdot)$ is continuous in its first argument, since it is the integral of a smooth set, and (ii) $\forall N^{\mathcal{J}}, \forall \sigma^{\mathcal{I},-j}, \lim_{P^{\mathcal{I},j} \rightarrow -\infty} N^{\mathcal{I},j} \left(P^{\mathcal{I},j}, \sigma^{\mathcal{I},-j}, N^{\mathcal{J}} \right) = 1$ (and

²⁰Hereafter, when discussing them informally, we typically drop the term “residual” when speaking of such tariffs.

$\lim_{p^{I,j} \rightarrow \infty} N^{I,j} (p^{I,j}, p^{I,-j}, N^J) = 0$), since $\forall \theta^I, \forall N^J, \forall \sigma^{I,-j}, \exists p^{I,j}$ such that

$$\max_{\mathcal{X}: j \in \mathcal{X}} \left\{ v^I(\mathcal{X}, N^J, \theta^I) - \widehat{p^{I,\mathcal{X}}} \right\} > (<) \max_{\mathcal{Y}: j \notin \mathcal{Y}} \left\{ v^I(\mathcal{Y}, N^J, \theta^I) - \widehat{p^{I,\mathcal{Y}}} \right\}.$$

For uniqueness, note that $N^{I,j}(\cdot, \cdot, \cdot)$ is nonincreasing in its first argument, since it is the sum of a set of nonincreasing functions. To see that it is in fact strictly decreasing, note that by our full support assumption, strictly positive density must always exist on the set of marginal consumers for whom the above relationship holds with equality.

To show that $\overline{p^{I,j}}$ is C^2 in all dimensions of its first argument, we note that, in response to a change in the value of an arbitrary element of $N^J, N^{J,k}$, in order to be insulating $\overline{p^{I,j}}$ must be differentiable by the inverse function theorem and have derivative equal to

$$\frac{\sum_{l \neq j} \frac{\partial N^{I,j}}{\partial p^{I,l}} \frac{\partial \sigma^{I,l}}{\partial N^{J,k}} + \frac{\partial N^{I,j}}{\partial N^{J,k}}}{-\frac{\partial N^{I,j}}{\partial p^{I,j}}}$$

so long as the denominator of this is bounded away from zero, which it is by our argument above that demand is strictly declining in own-price. Furthermore, this expression is, itself, differentiable in all elements of N^J by the smoothness assumptions we have imposed. \square

We now introduce vocabulary to describe the case when all platforms charge insulating tariffs.

Definition 4. An Insulating Tariff System (ITS) on side $I, \overline{P^I}(\mathbf{N}^J; \widetilde{\mathbf{N}}^I)$, is a profile of insulating tariffs, parameterized by the coarse allocation it induces, $\widetilde{\mathbf{N}}^I$. We say that $\overline{P^I}$ is “anchored” at Reference Allocation $\widetilde{\mathbf{N}}^I$. We denote a marketwide ITS by $\overline{P}(\widetilde{\mathbf{N}}) \equiv \left(\overline{P^A}(\mathbf{N}^B; \widetilde{\mathbf{N}}^A), \overline{P^B}(\mathbf{N}^A; \widetilde{\mathbf{N}}^B) \right)$.

Note that the ITS, at any anchor allocation, like residual insulating tariffs, exists and is unique directly from Theorem 1: it is exactly the unique set of price consistent with the anchor allocation and the allocation at which it is evaluated. It is also C^2 by the smoothness of the demand system.

We can now define our solution concept. Insulated Equilibria are particular Subgame Perfect Equilibria. We first define the latter in the context of our game and then we state the definition of IE. Given a consumer strategy profile, $\widehat{\mathcal{M}}$, and a profile of strategies adopted by other firms, σ^{-j} , denote firm j 's profits by

$$\Pi^j \left[\sigma^j, \sigma^{-j}; \widehat{\mathcal{M}} \right] \equiv \sum_{I=A,B} \sigma^{I,j} \left(N^{J,j} \left[\widehat{\mathcal{M}}, \sigma \right] \right) N^{I,j} \left[\widehat{\mathcal{M}}, \sigma \right] - C^j \left(N^{A,j} \left[\widehat{\mathcal{M}}, \sigma \right], N^{B,j} \left[\widehat{\mathcal{M}}, \sigma \right] \right). \quad (4)$$

Definition 5. In a particular platform game, defined by $\widehat{\mathcal{M}}$, a platform strategy profile, σ , forms a Platform Nash Equilibrium (PNE) if $\sigma^j \in \arg \max_{x \in \Sigma} \Pi^j(x, \sigma^{-j}; \widehat{\mathcal{M}})$, $\forall j \in \mathcal{M}$.

Definition 6. A set containing a profile of strategies for platforms and for consumers on each side, $\{\sigma^*, \{\mathcal{M}^I\}_{I=\mathcal{A}, \mathcal{B}}\}$, forms a Subgame Perfect Equilibrium (SPE) if σ^* forms a PNE given $\{\mathcal{M}^I\}_{I=\mathcal{A}, \mathcal{B}}$ and $\mathcal{M}^I = \mathcal{M}^{I*}(\theta^I, N^{\mathcal{J}}[\mathcal{M}^{\mathcal{J}}, x]; [x])$, $\forall \theta^I \in \Theta^I$, $\forall x \in \Sigma^m$.

We now state the definition of an Insulated Equilibrium.

Definition 7. Let $\{\sigma^*, \widehat{\mathcal{M}}^*\}$ be an SPE with coarse allocation $N^* = (N^{\mathcal{A}*}, N^{\mathcal{B}*})$. The SPE $\{\sigma^*, \widehat{\mathcal{M}}^*\}$ is an Insulated Equilibrium (IE) if platforms' strategy profile is the Insulating Tariff System anchored at N^* , i.e. if $\sigma^* = \overline{P}(N^*)$.

In a Subgame Perfect Equilibrium, platforms select their strategies as if they had complete certainty of the outcome of the continuation Consumer Game, even when the particular consumer game that they induce has multiple Consumer Nash Equilibria. Thus, one must speak of platforms' profits as functions of both platforms' strategies and of consumers' strategies. *Under Insulated Equilibrium, on the other hand, the particular strategy profile adopted by consumers is of no consequence*, since, when the platforms' strategy profile amounts to an Insulating Tariff System, in the subsequent Consumer Game, there is a unique Consumer Nash Equilibrium.

5.2 A Special Property of Insulating Tariff Systems

One way to look at the Insulating Tariff System is in terms of a *Representative Consumer (RC)*.²¹ Suppose that on side \mathcal{I} there is a single agent in charge of choosing quantities, or "slots" on platforms, for his constituent consumers on side \mathcal{I} to efficiently allocate among themselves, and that the RC's objective is to maximize the sum of constituents' utility. The RC's objective function can thus be written as

$$V^{\mathcal{I}}(N_{RC}^{\mathcal{I}}, N^{\mathcal{J}}) - \sigma^{\mathcal{I}}(N^{\mathcal{J}}) \cdot N_{RC}^{\mathcal{I}}.$$

where " \cdot " denotes the inner-product operator and where, as defined in (3), $V^{\mathcal{I}}$ denotes Gross Consumer Surplus on side \mathcal{I} , which, here, can be interpreted as the gross payoff to the representative consumer. Consider the following *Representative Consumer Game*, defined by the strategy profile, σ , announced by platforms. On side \mathcal{I} , the RC chooses coarse allocation $N_{RC}^{\mathcal{I}}$; activity among side \mathcal{J} consumers occurs as before. We can now state Theorem 2.

²¹See Anderson et al. (1992), particularly chapter 3, for foundations of the representative consumer approach in a one-sided discrete choice setting.

Theorem 2. In an RC game, σ , it is a strictly dominant strategy for the Representative Consumer to select $N_{RC}^I = N^{I*}$ if and only if the platforms' side \mathcal{I} strategy profile is the Insulating Tariff System anchored at N^{I*} . Formally, it holds that

$$V^I(N^{I*}, N^J) - \sigma^I(N^J) \cdot N^{I*} > V^I(\widetilde{N}^I, N^J) - \sigma^I(N^J) \cdot \widetilde{N}^I, \quad \forall N^J, \forall \widetilde{N}^I \neq N^{I*}$$

$$\Leftrightarrow \sigma = \left(\overline{P}^I(N^J; N^{I*}), \sigma^J(N^J) \right).$$

Proof. First note that

$$\begin{aligned} V^I(N_{RC}^I, N^J) - \sigma^I(N^J) \cdot N_{RC}^I &= \\ & \max_{\mathcal{M}^I \in \{\mathcal{M}^I: N^I[\mathcal{M}^I, \sigma] = N_{RC}^I\}} \sum_{\mathcal{X} \in \varphi(\mathcal{M}^I)} \int_{\theta^I: \mathcal{M}^I(\theta^I, [\sigma]) = \mathcal{X}} \left(v^I(\mathcal{X}, N^J, \theta) - \widehat{P}^I, \mathcal{X}(N^J) \right) f(\theta) d\theta \\ & \leq \sum_{\mathcal{X} \in \varphi(\mathcal{M}^I)} \int_{\theta^I: \mathcal{M}^{I*}(\theta^I, N^J, [\sigma]) = \mathcal{X}} \left(v^I(\mathcal{X}, N^J, \theta) - \widehat{P}^I, \mathcal{X}(N^J) \right) f(\theta) d\theta. \quad (5) \end{aligned}$$

where, by revealed preference, the inequality in (5) is strict if and only if $N_{RC}^I \neq N^I[\mathcal{M}^{I*}(\theta^I, N^J; [\sigma]), \sigma]$ as net surplus has a unique maximizer by its concavity from gross substitutes. Second, note that, $N^I[\mathcal{M}^{I*}(\theta^I, N^J; [\sigma]), \sigma] = N^{I*}$, $\forall N^J$, if and only if $\sigma^I(\cdot) = \overline{P}^I(\cdot; N^{I*})$, by the definition of the Insulating Tariff System. This establishes our claim. \square

5.3 Marginal Costs Are Identified Under Insulated Equilibrium

We propose Insulated Equilibrium as a refinement of Subgame Perfect Equilibrium. In doing so, we claim, as motivation for *this particular refinement*, that platforms can reasonably be expected to charge insulating tariffs. Independently, however, of the issue of multiplicity of CNE, there is another, perhaps more damning problem with SPE as a solution concept for our class of games, namely that it is largely vacuous. This issue of multiplicity of *Platform Equilibria*, holding fixed consumers' strategy profile, is discussed by Armstrong (2006), in Proposition 3 and in the discussion thereafter. The basic issue, which we refer to as *Armstrong's Paradox*, follows from the multiplicity of supply function equilibria in a deterministic setting, analyzed by Klemperer and Meyer (1989).

This can be understood by observing expression (4) of the profits of a platform j . Suppose there is some allocation, $(N^{\mathcal{A}, j*}, N^{\mathcal{B}, j*})$, that uniquely maximizes j 's profits, given the strategies of the other platforms and given the consumers' strategy profile. Then, j 's equilibrium strategy, σ^j , must, in one way or another, include this set. However,

the functions $\sigma^{\mathcal{I},j}(\cdot)$ are not pinned down when evaluated at non-equilibrium quantities, $N^{\mathcal{J},j} \neq N^{\mathcal{J},j*}$. Figure 2 illustrates this aspect of platforms' best responses.




Figure 2: The best-response criterion does not tie down a platform's *strategy*.

In particular, even with respect to those parts of the pricing functions directly relevant to first-order incentives, the two full $m \times m$ Jacobian matrices of $\sigma^{\mathcal{I}}$ and $\sigma^{\mathcal{J}}$ are free. In order to satisfy the first-order equilibrium conditions at any point these need only satisfy $2m < 2m^2$ for $m > 1$ first-order conditions at the conjectured equilibrium point. As a result, we conjecture that it is possible to construct a set of platform strategies that support, as an SPE, any coarse allocation in which all platforms make positive profits.²² However, regardless of whether this is precisely true, the set of SPEs is very large. One clear manifestation of this problem is the fact that it is impossible, based solely on first-order conditions, as is commonly done in standard one-sided markets (Rosse, 1970), to identify firms' marginal cost from a measurement of the demand system and an observation of prices.

Thus, if a solution concept for this class of games is to have significant predictive power, it must be stronger than SPE. Here we show that, under the IE solution concept, the issue of multiplicity of equilibria is reduced to the point where it takes the same form as in traditional "one-sided" models of imperfect competition typically used in industrial organization. Theorem 3 states this from the perspective of the standard Rosse exercise: under IE, just as in one-sided differentiated products Bertrand competition, marginal production costs are identified by observing prices, quantities and the demand system.

Theorem 3 (Under Insulated Equilibrium, Marginal Cost is Identified). *Suppose $\{\widehat{\mathcal{M}}^*, \sigma^*\}$ is an IE with coarse allocation N^* , with generic elements $N^{\mathcal{I},j*}$. Then, the vector of platform*

²²We are working on a formal proof of this conjecture; the difficulty is showing that global optimality and stability conditions are also satisfied.

marginal costs is identified jointly by the vector of prices, $\{\mathbf{P}^I\}_{I=\mathcal{A},\mathcal{B}}$, the coarse allocation, the payoff functions $\{v^I\}_{I=\mathcal{A},\mathcal{B}}$ and the distribution of types $\{f^I\}_{I=\mathcal{A},\mathcal{B}}$.

Proof. Since $\{\widehat{\mathcal{M}}^*, \sigma^*\}$ is an IE with coarse allocation \mathbf{N}^* , the equilibrium profile of platform strategies is $\sigma^* = \overline{\mathbf{P}}(\mathbf{N}^*)$. Thus, platform j 's profit maximization problem can be written

$$\max_{\{N^{\mathcal{A},j}, N^{\mathcal{B},j}\}} \sum_{I=\mathcal{A},\mathcal{B}} N^{I,j} \cdot P^{I,j}(N^{I,j}, N^{\mathcal{J},j}) - C^j(N^{\mathcal{A},j}, N^{\mathcal{B},j}), \quad (6)$$

where

$$P^{I,j}(N^{I,j}, N^{\mathcal{J},j}) = \overline{P}^{I,j}(N^{\mathcal{J}}; N^{I,j}, [\overline{P}^{I,-j}(N^{\mathcal{J}}, N^{I*})])$$

and

$$N^{\mathcal{J}} = N^{\mathcal{J}}(\overline{P}^{\mathcal{J},j}(N^I; N^{\mathcal{J},j}, [\overline{P}^{\mathcal{J},-j}(N^I; N^{\mathcal{J}*})]), \overline{P}^{\mathcal{J},-j}(N^I; N^{\mathcal{J}*}), N^I).$$

The values that maximize (6), $N^{\mathcal{A},j*}$ and $N^{\mathcal{B},j*}$, satisfy first-order condition

$$\begin{aligned} P^{I,j} + N^{I,j*} \frac{\partial P^{I,j}}{\partial N^{I,j}} + N^{\mathcal{J},j*} \frac{\partial P^{\mathcal{J},j}}{\partial N^{I,j}} \\ = \overline{P}^{I,j} + N^{I,j*} \frac{\partial \overline{P}^{I,j}}{\partial N^{I,j}} + N^{\mathcal{J},j*} \frac{\partial \overline{P}^{\mathcal{J},j}}{\partial N^I} \cdot \frac{\partial N^I}{\partial P^{I,j}} \frac{\partial P^{I,j}}{\partial N^{I,j}} = \frac{\partial C^j}{\partial N^{I,j}}. \end{aligned} \quad (7)$$

All of these quantities are well-defined based on the demand and cost systems and the observed allocation by Theorem 1, and thus a unique vector of marginal costs is consistent with a given Insulated Equilibrium. \square

Figure 3 illustrates the intuition behind Theorem 3.




Figure 3: The shape of the Insulating Tariff System is tied down by the demand system.

6 Pricing Under Insulated Equilibrium

In the previous sections, we have shown how to solve for Insulated Equilibrium and explained its motivation. The rest of the paper focuses on analyzing the economic predictions of our model, using this solution concept. In this section, we study the prices that arise under Insulated Equilibrium and compare them with those that correspond to a socially optimal allocation.

It is divided into three parts: in Sections 6.1 and 6.2, we continue in the general environment that we have assumed thus far. In 6.1, we first consider the benchmark of socially optimal pricing and then derive the Insulated Equilibrium pricing formula previewed in equation (11). In 6.2, we examine, in detail, the components of the “two-sided externality” term in this formula. Section 6.3 then specializes the model in several different ways, illustrating the relationship of our results to existing literature and paying special attention to the differing impacts of different forms of consumer heterogeneity.

6.1 General Pricing

Socially Optimal

The utilitarian social welfare corresponding to an allocation is equal to

$$\sum_{I=\mathcal{A},\mathcal{B}} V^I(N^I, N^{\mathcal{J}}) - \sum_{j \in \mathcal{M}} C^j(N^{\mathcal{A},j}, N^{\mathcal{B},j}), \quad (8)$$

where $V^I(N^I, N^{\mathcal{J}})$ denotes gross consumer surplus on side I , as defined in equation (3). In Proposition 1 we state the pricing formula for maximizing this quantity. To do so, let us denote by

$$\overline{v}_j^{I,j} \equiv \frac{\int_{\theta^I: j \in \mathcal{M}^{I*}(\theta^I)} \frac{\partial v^I(\mathcal{M}^{I*}(\theta), N^{\mathcal{J}}, \theta^I)}{\partial N^{\mathcal{J},j}} f(\theta) d\theta}{N^{I,j}}$$

the average valuation, among *all* of platform j 's side I consumers, for an additional side \mathcal{J} consumer to join platform j .

Proposition 1. *At a socially optimal allocation, the total price charged to side I consumers to join platform j satisfies*

$$P^{I,j} = C_I^j - N^{\mathcal{J},j} \overline{v}_j^{\mathcal{J},j}. \quad (9)$$

Proof. By Lemma 2, we have $\frac{\partial V^I}{\partial N^{I,j}} = P^{I,j}$. The other terms in (9) are straightforward. \square

Equation (9) affords two complementary interpretations. The first, “Pigouvian” interpretation is that social efficiency requires the consumers that join platform j to be

those whose willingness to pay (holding fixed the set of other platforms in their optimal bundle) exceeds the net social cost they impose by joining. This net social cost consists of the “physical” cost incurred by the platform minus the externality they impose on opposite-side consumers.

The second, “Spenceian” interpretation can best be appreciated by rearranging (9) as

$$\frac{C_I^j - P^{I,j}}{N^{J,j}} = \overline{v_j^{J,j}}. \quad (10)$$

In the one-sided model of Spence (1975), it is socially optimal for the firm choose its quality level so that the marginal cost, per consumer, of an improvement in quality equals the average valuation of all its consumers for such an improvement. Equation (10) says precisely this, since it prescribes that the net social marginal cost, per side \mathcal{J} consumer, of providing such a quality increase²³ be equated with this average valuation.

Under Insulated Equilibrium

We now state the Insulated Equilibrium pricing formula in Proposition 2.

Proposition 2. *At an IE allocation, the total price platform j charges to side \mathcal{I} consumers satisfies*

$$P^{I,j} = C_I^j + \mu^{I,j} - N^{J,j} \left(\left[-\frac{\partial N^J}{\partial P^J} \right]^{-1} \left[\frac{\partial N^J}{\partial N^I} \right] \right)_{j,j} \cdot [-D_{\cdot,j}^I]. \quad (11)$$

Proof. In view of equation (7), it remains for us to show that $\overline{\frac{\partial P^{J,j}}{\partial N^I}} = \left(\left[-\frac{\partial N^J}{\partial P^J} \right]^{-1} \left[\frac{\partial N^J}{\partial N^I} \right] \right)_{j,j}$. In an IE, platforms’ tariffs constitute an Insulating Tariff System. Thus, by definition, in response to changes in the side \mathcal{I} coarse allocation, prices on side \mathcal{J} adjust so as to hold the side \mathcal{J} coarse allocation unchanged. Therefore, the Jacobian of the Insulating Tariff System satisfies

$$\underbrace{\mathbf{0}}_{m \times m} = \begin{bmatrix} \frac{\partial N^J}{\partial N^I} \end{bmatrix} + \begin{bmatrix} \frac{\partial N^J}{\partial \overline{P^J}} \end{bmatrix} \begin{bmatrix} \overline{\frac{\partial P^J}{\partial N^I}} \end{bmatrix} \Leftrightarrow \begin{bmatrix} \overline{\frac{\partial P^J}{\partial N^I}} \end{bmatrix} = \begin{bmatrix} -\frac{\partial N^J}{\partial \overline{P^J}} \end{bmatrix}^{-1} \begin{bmatrix} \frac{\partial N^J}{\partial N^I} \end{bmatrix}, \quad (12)$$

and we have our result. □

Two forces govern the relationship between the optimal pricing formula of equation

²³For simplicity of exposition, we refer to this change in platform j ’s side \mathcal{J} quality as an “increase”, even though it could, in principle, be negative; in some applications such as ad-based media, we would expect it to be so.

(10) and the equilibrium pricing rule of equation (11). The first is the classical market power distortion, captured by $\mu^{I,j} \equiv -N^{I,j} / \left(-\frac{\partial N^{I,j}}{\partial P^{I,j}}\right)$. As is well known from classical industrial organization theory, this term decreases as competition intensifies, through an increase in the number of platforms and/or an increase in their substitutability.

The second force is what we refer to as the *Spence distortion*. As discussed above, the allocation that a platform chooses on side \mathcal{I} determines the quality of the platform for consumers on side \mathcal{J} . In Spence’s model, the quality that a one-sided monopolist provides to consumers is distorted because it depends on the willingness to pay for quality of its *marginal* consumers, rather than the average of all its consumers. As established in Proposition 1, the socially efficient quality level for a platform to provide to its side \mathcal{J} consumers depends on the average preferences of such consumers. Analogously to the result in Spence’s model, the quality a platform chooses in equilibrium depends on the appreciation for quality of its *marginal* side \mathcal{J} consumers. We now examine, in more detail, the precise sense in which this is true.

6.2 Decomposing the *Two-Sided Externality* Term

Under competition, platforms have multiple margins on each side of the market. In an Insulated Equilibrium, the amount by which the presence of an additional side \mathcal{I} consumer allows platform j to increase its revenue from its current set of side \mathcal{J} consumers is measured by $N^{\mathcal{J},j} \left(\left[-\frac{\partial N^{\mathcal{J}}}{\partial P^{\mathcal{J}}} \right]^{-1} \left[\frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}} \right] \right)_{j,\cdot} \cdot [-\mathbf{D}^{\mathcal{I}}]_{\cdot,j}$, the “two-sided” or “cross-externality” term in the pricing formula of Proposition 2. The middle and right-hand factors in this term reflect the role of these various margins in determining the size of such an increase.

The right-hand factor in this term is the negative of the j^{th} column of the side \mathcal{I} diversion ratio matrix. This vector measures the displacement of *side \mathcal{I} consumers* from other platforms to platform j that occurs when j lowers its price on side \mathcal{I} and all other platforms keep their side \mathcal{I} prices fixed.²⁴ This switching by side \mathcal{I} consumers determines the extent to which *the characteristics* of all platforms change, from the perspective of side \mathcal{J} consumers, in response to platform j ’s lowering its price on side \mathcal{I} .

The middle factor in this term is the j^{th} row of the Jacobian of the side \mathcal{J} Insulating Tariff System. Each element in this vector measures the amount by which platform j must change its price on side \mathcal{J} , per unit of change in the number of side \mathcal{I} consumers participating on a particular platform, in order to hold fixed the level of its demand on side \mathcal{J} . It is thus clear why the size of the cross-externality term depends, on the rate at which *marginal consumers* on side \mathcal{J} are willing to trade off money for additional “interactions”

²⁴Recall that we assume Bertrand conduct among firms within each side.

or “quality”.

Since, in equilibrium, all platforms charge *residual* insulating tariffs, the amount by which each platform’s side \mathcal{J} price changes in response to a shift in the side \mathcal{I} allocation is somewhat subtle. These changes, however, are tied down by the underlying demand system in a way that, after reflection, becomes quite intuitive. As W10 shows, in the case of a monopoly platform, the side \mathcal{J} insulating tariff responds to a change in side \mathcal{I} allocation by exactly compensating the platform’s average marginal consumer on side \mathcal{J} . In other words, it responds to a quality change by “overcompensating” a number of previously excluded consumers equal to the number of previously included consumers that it “undercompensates”, thus leaving its level of side \mathcal{J} demand unchanged.

In the general case of our model, each platform’s insulating tariff behaves in precisely the same way; only, when there’s competition, these notions of “compensation” must be thought of as net of the price changes of *other* platforms. These compensation levels turn out to be weighted averages, with different weights put on different margins, of the average valuations of consumers within each margin for an additional interaction with an opposite-side consumer. A fundamental issue, then, for understanding the effect of changes in competition on prices and welfare in two-sided industries is understanding the effects of such changes in the environment on the equilibrium *weights* that platforms place on different margins. To investigate this, we first take formal inventory of these marginal valuations and their weights, in the general case, and then make use of them in Section 6.3, in which we consider examples with different competitive environments.

The Insulating Tariff System

We now derive the relationship between the Jacobian of the Insulating Tariff System, as defined by equation (12), and the underlying demand system. Two sorts of quantities comprise the matrix $\left[-\frac{\partial N^{\mathcal{J}}}{\partial p^{\mathcal{J}}}\right]^{-1} \left[\frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}}\right]$. One sort are densities of consumers that are *on the margin* between two bundles of platforms. The other sort are the aforementioned quality or *interaction values*—valuations for “interacting”, on a given platform, with an additional consumer from the opposite side, averaged over sets of such marginal consumers. In order to express these latter quantities, we first define the relevant sets of marginal consumers.

First, let $\widetilde{\Theta}_j^{\mathcal{I}} \equiv \left\{ \theta^{\mathcal{I}} \in \Theta^{\mathcal{I}} : \exists \mathcal{X}, \mathcal{Y} \in \arg \max_{\mathcal{Z} \in \varphi(\mathcal{M})} \left\{ v^{\mathcal{I}}(\mathcal{Z}, N^{\mathcal{J}}, \theta^{\mathcal{I}}) - \widehat{P}^{\mathcal{I}, \mathcal{Z}} \right\} \text{ s.t. } j \in \mathcal{X}, j \notin \mathcal{Y} \right\}$ denote the entire set of consumers on side \mathcal{I} that are indifferent between consuming some bundle of platforms, \mathcal{X} , containing platform j , and consuming some other bundle, \mathcal{Y} , not containing platform j . Second, let $\widetilde{\Theta}_{j,k}^{\mathcal{I}} \equiv \left\{ \theta^{\mathcal{I}} \in \Theta^{\mathcal{I}} : \exists \mathcal{X}, \mathcal{Y} \in \right.$

$\arg \max_{\mathcal{X} \in \varphi(\mathcal{N})} \left\{ v^I(\mathcal{X}, N^J, \theta^I) - \widehat{P^{I, \mathcal{X}}} \right\}$ s.t. $j \in \mathcal{X}, j \notin \mathcal{Y}, k \in \mathcal{Y}, k \notin \mathcal{X}$ } denote the set of consumers on side I that are indifferent between consuming some bundle of platforms, \mathcal{X} , containing platform j and not containing platform k , and consuming some other bundle \mathcal{Y} , containing platform k and not containing platform j . The matrix $\left[-\frac{\partial N^I}{\partial P^I} \right]$ is simply the negative of the Slutsky matrix of the side I demand system that arises, given the coarse allocation on the opposite side, N^J . Let us denote the density of a set of marginal consumers, $\widetilde{\Theta}$, by $F[\widetilde{\Theta}] \equiv \int_{\widetilde{\Theta}} f^I(\widetilde{\Theta}) d\widetilde{\Theta}$, where $\widetilde{\theta}$ is an index of dimension $L^I - 1$ tracing out $\widetilde{\Theta}$. The elements of the Slutsky matrix are thus

$$\frac{\partial N^{I,j}}{\partial P^{I,k}} = \begin{cases} -F[\widetilde{\Theta}_j^I], & \text{if } j = k \\ F[\widetilde{\Theta}_{j,k}^I], & \text{if } j \neq k \end{cases}.$$

Terms on the diagonal of this matrix capture the number of consumers a platform loses when it increases its price by a small amount, and terms on the off-diagonal capture the number of consumers that switch to a bundle containing platform j when another platform, k , increases its price.

The *Interaction Matrix*, $\left[\frac{\partial N^I}{\partial N^J} \right]$, mirrors the Slutsky matrix except that it is weighted by the average over *marginal consumers' valuations for additional interaction, on a particular platform*, with consumers on the opposite side of the market. For a marginal set of consumers, $\widetilde{\Theta}$, defined in terms of bundle of platforms, \mathcal{X} , we denote the average, over $\widetilde{\Theta}$, of such interaction values by $v_k^{I, \mathcal{X}}[\widetilde{\Theta}]$, where

$$v_k^{I, \mathcal{X}}[\widetilde{\Theta}] \equiv \frac{\int_{\widetilde{\Theta}} \frac{\partial v^I(\mathcal{X}, N^J, \widetilde{\Theta})}{\partial N^{J,k}} f^I(\theta) d\theta}{F[\widetilde{\Theta}]}.$$

The elements of the interaction matrix are thus

$$\frac{\partial N^{I,j}}{\partial N^{J,k}} = \begin{cases} v_k^{I, \mathcal{X}}[\widetilde{\Theta}_j^I] \cdot F[\widetilde{\Theta}_j^I], & \text{if } j = k \\ -v_k^{I, \mathcal{Y}}[\widetilde{\Theta}_{j,k}^I] \cdot F[\widetilde{\Theta}_{j,k}^I], & \text{if } j \neq k \end{cases}.$$

Note, first, that the signs of the terms in this matrix correspond to the signs of marginal consumers' average "interaction valuations". Thus, in the case where consumers have positive interaction values, the signs of the terms in this matrix are the reverse of those in the Slutsky matrix. Second, note that the first argument of $\frac{\partial v^I(\cdot, N^J, \theta^I)}{\partial N^{J,k}}$ in the various terms corresponds to the subset to which the platform on which there is a change in allocation belongs. In the case of the set $\widetilde{\Theta}_j^I$, the change in coarse allocation being contemplated,

$\partial N^{\mathcal{J},k}$, occurs on a platform that forms part of the bundle, \mathcal{X} , of which platform j is a member. In contrast, in the case of the set $\widetilde{\Theta}^{\mathcal{I}}_{j,k}$, the change under consideration occurs on a platform that is part of a bundle, \mathcal{Y} , to which platform j does not belong.

6.3 Examples and Intuition

We now draw on the above discussion of the underlying demand system to consider various special cases of IE pricing. We begin by stating Proposition 3, pertaining to the case where consumers on one of the two sides of the market have independent demand for each platform. Numerous articles in the literature, such as Rysman (2004), Anderson and Coate (2005), Armstrong and Wright (2007), as well as Armstrong (2006) in the section on multi-homing, have studied such scenarios and argued for their relevance in particular contexts, such as the markets for advertisement in Yellow Page directories and broadcast media.

Proposition 3 (Independent Demand). *Suppose that demand on side \mathcal{J} is independent across platforms (i.e., consider the limit case as $\frac{\partial N^{\mathcal{J},k}}{\partial P^{\mathcal{J},j}}$ and $\frac{\partial N^{\mathcal{J},k}}{\partial N^{\mathcal{I},j}}$ approach zero, for $k \neq j$). Then, IE pricing on side \mathcal{I} collapses to the monopoly formula of W10, given by*

$$P^{\mathcal{I},j} = C^j_{\mathcal{I}} + \mu^j_{\mathcal{I}} - N^{\mathcal{J},j} \sigma^{\mathcal{I},\mathcal{X}}_j \left[\widetilde{\Theta}^{\mathcal{I}}_j \right].$$

Proof. When demand on side \mathcal{J} is independent across platforms, the inverse of the side \mathcal{J} Slutsky matrix is diagonal with j^{th} entry that is the inverse of platform j 's marginal mass of side \mathcal{J} consumers. This leads to our result. \square

The simplicity of this case comes from the fact that, on side \mathcal{J} , each platform has only a market expansion margin. As a result, each platform's insulating tariff charged to side \mathcal{J} consumers need only vary as a function of *its own* allocation on side \mathcal{I} . Consequently, when platform j increases its allocation on side \mathcal{I} , it does not need to take into account any response by other platforms in order to keep its own allocation on side \mathcal{J} fixed, leaving the pricing formula to be the same as that of a monopolist.

As we discuss in section 2.2, much of the literature on two-sided markets has proceeded by considering extensions of Armstrong (2006)'s two-sided single-homing model, which adopts a Hotelling setup. A key assumption of these models is that all consumers on a given side have the same interaction values. Proposition 4 states the IE pricing formula that emerges under this homogeneity assumption.

Proposition 4 (Generalized Armstrong Pricing). *When all side \mathcal{J} consumers have a common, constant valuation $\overline{v^{\mathcal{J}}}$ for interacting with additional side \mathcal{I} consumer, firm j 's IE price on side \mathcal{I} satisfies*

$$P^{\mathcal{I},j} = C_I^j + \mu_I^j - N^{\mathcal{J},j} \overline{v^{\mathcal{J}}}.$$

Proof. When all side \mathcal{J} consumers have the same interaction valuation, $\overline{v^{\mathcal{J}}}$, this term can be factored out of the interaction matrix, leaving the inverse Slutsky and Slutsky matrices to cancel each other out, so that the Jacobian of the Insulating Tariff System becomes $\overline{v^{\mathcal{J}}}\mathbf{I}$. Thus the right hand side of expression 2 becomes

$$C_I^j + \mu_I^j - N^{\mathcal{J},j} \overline{v^{\mathcal{J}}} [\mathbf{I}]_{j,j} \cdot [-\mathbf{D}_{:,j}^{\mathcal{I}}] = C_I^j + \mu_I^j - N^{\mathcal{J},j} \overline{v^{\mathcal{J}}},$$

because the j^{th} entry of $[-\mathbf{D}_{:,j}^{\mathcal{I}}]$ is equal to one. □

As in the case of Proposition 3, when side \mathcal{J} consumers have homogenous interaction valuations, each platform's insulating tariff depends only on its own side \mathcal{I} allocation.²⁵ However, the reason for this is different in this case, since, here, there can be arbitrary substitution patterns among side \mathcal{J} consumers among platforms. Instead, when all side \mathcal{J} consumers have the same interaction valuations and the side \mathcal{I} allocation changes on one platform, the adjustment of its own insulating tariff alone preserves the entire coarse allocation on side \mathcal{J} . This stems from the fact that, in this special case, insulating tariffs provide *full insurance* to all consumers against variation in the opposite side allocation.

Two Platforms

We now suppose there are two platforms ($m = 2$) and first look at a pair of instructive special cases before stating a general two platform pricing formula. Suppose that demand on side \mathcal{I} is independent across platforms. Then, expression (11), of platform j 's price on side \mathcal{I} becomes

$$P^{\mathcal{I},j} = C_I^j + \mu^{\mathcal{I},j} - N^{\mathcal{J},j} \frac{\partial P^{\mathcal{J},j}}{\partial N^{\mathcal{I},j}}. \quad (13)$$

In equation (13), the right-hand factor in the cross-externality term is the partial derivative of firm j 's side \mathcal{J} insulating tariff, with respect to its own coarse allocation. Further suppose that the number of consumers on side \mathcal{J} that multi-home is negligible. Under these assumptions, on side \mathcal{J} , there are three margins – one “cannibalization margin” between firms 1 and 2 and one “market expansion margin” between each firm and \emptyset . This

²⁵We thank Julian Wright for drawing our attention to this point.

factor then simplifies to

$$\frac{\partial \overline{P^{\mathcal{J},j}}}{\partial N^{\mathcal{I},j}} = \left((1 - \kappa) \cdot v_j^{\mathcal{J},\mathcal{X}} \left[\overline{\Theta^{\mathcal{J}}_{j,\emptyset}} \right] + \kappa \cdot v_j^{\mathcal{J},\mathcal{X}} \left[\overline{\Theta^{\mathcal{J}}_{j,k}} \right] \right), \quad (14)$$

where

$$\kappa \equiv 1 / \left(1 + F \left[\overline{\Theta^{\mathcal{J}}_{j,\emptyset}} \right] \left(\frac{1}{F \left[\overline{\Theta^{\mathcal{J}}_{j,k}} \right]} + \frac{1}{F \left[\overline{\Theta^{\mathcal{J}}_{k,\emptyset}} \right]} \right) \right).$$

The term $\frac{\partial \overline{P^{\mathcal{J},j}}}{\partial N^{\mathcal{I},j}}$ is thus a weighted average of the average interaction values for an additional interaction on platform j of side \mathcal{J} consumers along its own two margins. This weighting, governed by κ , depends on the relative measures of consumers on each of the three side \mathcal{J} margins.

When firm j 's market expansion margin is more crowded, then κ is small and firm j behaves similarly to a monopoly. In particular, it is analogous to a monopoly in that it sets its quality on side \mathcal{J} to cater to consumers on the market expansion margin, on which the consumers would likely be similar to those on the margin of a monopolist.

On the other hand, when the cannibalization margin is heavier, then κ is larger, and platform j caters more to consumers on this margin. Consumers on the cannibalization margin are quite different from those on the market expansion margin. Crucially, with respect to the overall decision of whether or not to join *some* platform, they are infra-marginal – and to all different degrees. As a result, it is natural to suppose that the average interaction value of consumers on the cannibalization margin, $v_j^{\mathcal{J},\mathcal{X}} \left[\overline{\Theta^{\mathcal{J}}_{j,k}} \right]$, will be closer than the average interaction value of consumers on j 's expansion margin, $v_j^{\mathcal{J},\mathcal{X}} \left[\overline{\Theta^{\mathcal{J}}_{j,\emptyset}} \right]$, to the average interaction value among *all* of platform j 's consumers on side \mathcal{J} . As Figure 4 illustrates, under circumstances such as those where the two platforms' primary dimension of differentiation, on side \mathcal{J} , is *horizontal* in membership benefits, the former group of consumers on the *cannibalization margin* constitutes a more representative sample than the latter group of consumers on the *market expansion* margin.

This scenario thus represents a mechanism through which competition among platforms can reduce the Spence distortion. This need not be the case, however. For instance, when platforms are vertically differentiated from one another in a manner analogous to that of Shaked and Sutton (1982), then an increase in competition can have the opposite effect. We now sketch such an example.

To fix ideas, assume that demand system and platform cost functions lead to an equilibrium allocation on side \mathcal{I} such that $N^{\mathcal{I},j} > N^{\mathcal{I},k}$. Furthermore, suppose that consumers on \mathcal{J} differ significantly from one another in both the membership and interaction ben-




Figure 4: On the left, the thin diagonal line represents an “unpopulated” margin between platforms j and k on side \mathcal{J} and thus a low value of κ ; on the right, the thick diagonal line represents a “crowded” margin between platforms and thus a high value of κ .

efits they perceive but that these preferences are, for most consumers, very stable across platforms. Formally, such preferences can be straightforwardly represented by the utility function giving a payoff

$$B_i^{\mathcal{J}} + b_i^{\mathcal{J}} N^{\mathcal{I},j} + \epsilon_i^{\{j\}} - P^{\mathcal{J},j}$$

to consumer i on side \mathcal{J} , when he joins the set containing only platform j , where $\epsilon_i^{\mathcal{X}}$ is a bundle-specific idiosyncratic term, $B_i^{\mathcal{J}}$ denotes consumer i 's membership value and $b_i^{\mathcal{J}}$ denotes consumer i 's interaction value. As implied by the description above, suppose that consumers' values of $\epsilon^{\mathcal{X}}$ are heavily concentrated around some value, normalized to zero.

In this setup, provided appropriate cost functions, under Insulated Equilibrium, both platform j and platform k attract a significant number of side \mathcal{J} consumers. Platform j charges its consumers a higher price than does platform k , while also allowing for interaction with a larger number of side \mathcal{I} consumers. Accordingly, (ignoring noise term, $\epsilon^{\mathcal{X}}$), we can define a threshold interaction value, $\widetilde{b}_{j,k}^{\mathcal{J}} \equiv \frac{p^{\mathcal{J},j} - p^{\mathcal{J},k}}{N^{\mathcal{I},j} - N^{\mathcal{I},k}}$, which represents the interaction value of all side \mathcal{J} consumers that lie on the cannibalization margin between platforms j and k .

Recall the first-order condition in expression (14), and note that as the mass of consumers with an interaction value of $\widetilde{b}_{j,k}^{\mathcal{J}}$ increases, so does κ . As a result, if such an increase were to occur, each platform would have an incentive to adjust its allocation on side \mathcal{I} so as to cater more to consumers on this cannibalization margin. In contrast to the previous example, however, this *exacerbates* the Spence distortion on side \mathcal{J} inflicted by each of the

two platforms. As Figure 5 illustrates, this is because, on the one hand, (except for an arbitrarily small measure) all of platform j 's side \mathcal{J} consumers have interaction values greater than $\widetilde{b}_{j,k}^{\mathcal{J}}$, while all of platform k 's side \mathcal{J} consumers have interaction values less than $\widetilde{b}_{j,k}^{\mathcal{J}}$.




Figure 5: Platforms are vertically differentiated on side \mathcal{J} . The thick blue line on the right illustrates mechanism by which more intense competition leads to a larger Spence distortion, since consumers on this margin have far from average interaction values.

Thus far in this section, we have “turned off” the competition among platforms on side \mathcal{I} by assuming that the cannibalization margin on this side is negligible. We now activate this feature of the model and examine the general case of competition between two platforms. We have

$$P^{\mathcal{J},j} = C_{\mathcal{I}}^j + \mu^{\mathcal{I},j} - N^{\mathcal{J},j} \left((1-\kappa) \cdot v_j^{\mathcal{J},\mathcal{X}} [\widetilde{\Theta}_{j,\emptyset}^{\mathcal{J}}] + \kappa \left(v_j^{\mathcal{J},\mathcal{X}} [\widetilde{\Theta}_{j,k}^{\mathcal{J}}] + \overbrace{D_{k,j}^{\mathcal{I}} \left(v_k^{\mathcal{J},\mathcal{X}} [\widetilde{\Theta}_{k,j}^{\mathcal{J}}] - v_k^{\mathcal{J},\mathcal{X}} [\widetilde{\Theta}_{k,\emptyset}^{\mathcal{J}}] \right)}^{\varphi} \right) \right) \right). \quad (15)$$

Note, first, the term in (15), indicated by φ , that does not enter the prior first-order condition, (14). This appears, since, when there is competition on side \mathcal{I} and j changes its quantity on this side, this affects the number of consumers that k serves as well. The diversion ratio on side \mathcal{I} represents the significance of this side \mathcal{I} reallocation. This overall reallocation of side \mathcal{I} consumers influences the perceived quality by side \mathcal{J} consumers of not only platform j but also of platform k . As a result, in order to hold fixed its quantity on side \mathcal{J} , j must take into account the interaction values of k 's marginal consumers for an additional interaction on that platform.

In particular, the relevant quantity for this purpose is $v_k^{\mathcal{J},\mathcal{X}} [\widetilde{\Theta}_{k,j}^{\mathcal{J}}] - v_k^{\mathcal{J},\mathcal{X}} [\widetilde{\Theta}_{k,\emptyset}^{\mathcal{J}}]$, the

difference between the value for an additional interaction on platform k of the side \mathcal{J} consumers on k 's cannibalization margin and those consumers on the market expansion margin. Thus, the extent to which j distorts the quality it provides to its side \mathcal{J} consumers depends not only on the divergence between the interaction values of its own marginal versus average consumers but also on the distribution of such valuations among consumers on other platforms. As competition on side \mathcal{I} toughens through an increase in $D_{jk}^{\mathcal{I}}$, in determining its quality on side \mathcal{J} , platform j puts more weight on the preferences of consumers on the cannibalization margin. This can bring the platform's incentives closer to or further from the social planners', according to the heterogeneity issues we discuss above.

m Symmetric Platforms

As a final example, we consider a symmetric Insulated Equilibrium among m identical platforms. Let $F_k^{\mathcal{J}}$ and $v_k^{\mathcal{J}}$ denote, respectively, the mass and average interaction value of side \mathcal{J} consumers on a given platform's cannibalization margin, and let $F_{\emptyset}^{\mathcal{J}}$ and $v_{\emptyset}^{\mathcal{J}}$ denote the mass and interaction value of \mathcal{J} consumers on a given platform's market expansion margin. The side \mathcal{I} first-order condition for platform j is given by

$$P^{\mathcal{I},j} = C_{\mathcal{I}}^j + \mu^{\mathcal{I},j} - N^{\mathcal{J},j} \left((1 - \kappa^{sym}) v_{\emptyset}^{\mathcal{J}} + \kappa^{sym} v_k^{\mathcal{J}} \right), \quad (16)$$

where

$$\kappa^{sym} \equiv \frac{\frac{(m-1)F_k^{\mathcal{J}}}{F_{\emptyset}^{\mathcal{J}} + mF_k^{\mathcal{J}}}}{1 - \frac{F_k^{\mathcal{I}}}{F_{\emptyset}^{\mathcal{I}} + mF_k^{\mathcal{I}}}}.$$

Expression (16) reinforces the themes we discuss in the previous examples. As in expression (14), the extent to which the quality provided to consumers on side \mathcal{J} depends on the characteristics of consumers on the two types of margins and on the weight the platform attributes to each of these margins. In particular, since κ^{sym} is increasing in $F_k^{\mathcal{J}}$, the mass of consumers on the side \mathcal{J} cannibalization margins, such an increase in competition on side \mathcal{J} reduces the Spence distortion experienced by consumers on that side if and only if the average interaction value of consumers on the cannibalization margin is closer than that of the consumers on the market expansion margin to the average interaction value of all consumers.

7 Stability, Uniqueness and Existence

Our extensive discussion above of the first-order conditions characterizing an Insulated Equilibrium are obviously only necessary and not sufficient for such an equilibrium to prevail. To analyze the conditions for existence, stability and uniqueness of equilibrium in such models, it has been recognized since at least the work of Samuelson (1941) that the gradient of the vector of first-order derivatives is fundamental. We now define the gradient matrix that is relevant for such analysis as well as for that of Section 8.1 on mergers. In debt to Samuelson, and to distinguish it from the better-known and related (but more restrictive) Hessian matrix of second partial derivatives of a single objective function, we refer to this matrix as the *Samuelsonian*.

Let ψ represent the vector of first-order derivatives of profits with respect to quantity which, in equation 11, are equated to 0. Given the allocation approach we have been employing, it is most natural to take the gradient of ψ with respect to quantity, as we will denote by $\nabla\psi$. However when studying the stability of equilibria in our price-choosing game, as well as the comparative statics of prices, this matrix must be transformed so as to, effectively, represent gradients with respect to prices rather than quantities.

There are two steps to this transformation of $\nabla\psi$ into the Samuelsonian matrix. One step involves changing the units of the matrix's entries from quantity to price. To do this, we pre-multiply $\nabla\psi$ by a matrix whose diagonal is made up of the diagonal terms of the Slutsky matrix on each side of the market, $\frac{\partial N^{I,j}}{\partial p^{I,j}}$, and whose off-diagonal terms are zeros. We denote this *unit-transforming* matrix by \mathbf{T}_U .

The other step in this transformation modifies $\nabla\psi$ to match platforms' conduct in the game. Specifically, it allows for each entry in the Samuelsonian to correspond to a change *in a given price, holding fixed all other prices on the same side of the market*, as dictated by the game's Bertrand conduct, and holding fixed quantities on the other side of the market, as dictated by the Insulating Tariff System. To do this, we post-multiply $\nabla\psi$ by a block matrix, whose diagonal blocks are negative diversion ratio matrices and whose off-diagonal blocks are zeros. We denote this *conduct-transforming* matrix by \mathbf{T}_C , where

$$\mathbf{T}_C \equiv \begin{bmatrix} -\mathbf{D}^A & \mathbf{0} \\ \mathbf{0} & -\mathbf{D}^B \end{bmatrix}.$$

Letting \mathfrak{S} denote the Samuelsonian matrix, we thus have

$$\mathfrak{S} \equiv [\mathbf{T}_U][\nabla\psi][\mathbf{T}_C].$$

Some intuition for \mathfrak{S} comes from noting its relationship to the game's (cost) pass-through matrix (Weyl and Fabinger, 2009). To see this, suppose that there were a specific tax levied against each platform, per consumer that it serves on each side of the market, and let the vector of such taxes be denoted by t . The necessary condition for Insulated Equilibrium can be written $\psi = t$. Implicitly differentiating this system of first-order conditions with respect to each of the taxes then yields

$$\left[\frac{\partial \psi}{\partial P} \right] \left[\frac{\partial P}{\partial t} \right] = \mathbf{I}_{2m} \quad \Leftrightarrow \quad [\nabla \psi] [\mathbf{T}_C] \left[\frac{\partial P}{\partial t} \right] = \mathbf{I}_{2m},$$

where \mathbf{I}_{2m} denotes the $2m$ -dimensional identity matrix, and $\frac{\partial P}{\partial t}$ is the Jacobian matrix of equilibrium price changes in response to changes in the specific taxes. Rearranging this equation gives

$$\frac{\partial P}{\partial t} = [\mathfrak{S}]^{-1} [\mathbf{T}_U]. \quad (17)$$

Equation 17 thus shows that the Samuelsonian, transformed into units of quantity, is the inverse of the game's pass-through matrix.

The necessary conditions for equilibrium, discussed above, would also be sufficient, provided that platforms' objective functions are quasi-concave, given the residual inverse demand defined by other platforms' insulating tariffs. In terms of the Samuelsonian, a common such condition (i.e. concavity of profits) is that the principal submatrix of \mathfrak{S} , formed by the two rows corresponding to a particular platform j , be negative definite for every price pair. Alternative such conditions typically involve the negative definiteness of some (possibly price-dependent) positive-diagonal transformation of this submatrix.

If one were interested in a stronger notion of sufficiency, such as local stability, independent of adjustment speed, in the sense of Enthoven and Arrow (1956), the matrix \mathfrak{S} would have to be (local to the conjectured equilibrium) D-stable, a standard generalization of negative definiteness to non-symmetric matrices. Furthermore, it is well-known that if such conditions hold globally, equilibrium is unique (if it exists).

Existence of an insulated equilibrium also relies on conditions placed on the demand system entirely analogous to those in standard markets. If the marginal externalities to users on each side of the market are bounded uniformly regardless of the number of users on the other side, the additional "marginal costs or benefits" to the platform arising from two-sidedness are similarly bounded, as the latter is a simple linear transformation of the former. Thus if Bertrand equilibrium exists in a sufficiently wide range of cases on each side independently, so too will an Insulated Equilibrium: each allocation on side \mathcal{A} will lead to a Bertrand equilibrium allocation on side \mathcal{B} and this in turn will induce a Bertrand

equilibrium allocation on side \mathcal{A} . A fixed point of this process must exist, given standard fixed point theorems, as the allocations on each side are in the compact set $[0, 1]^m$.

An analysis of conditions on the primitives of our model ensuring such standard conditions on endogenous variables is beyond the scope of our analysis here and thus we do not belabor these points common to all general standard industrial organization models, but readily acknowledge their potential importance. In future work we may consider these issues in greater detail.

8 Applications and Extensions

8.1 First-Order Merger Analysis

In this section, we extend the techniques of Jaffe and Weyl (2010b), hereafter “JW”, to consider the effect on consumer surplus of a potential merger of platforms. The key to this extension is that we must take into account not only the potential harms or benefits to consumers from the changes *in* the (insulating) tariffs charged due to the merger, but also the welfare effects of movements *along* these insulating tariffs caused by changes in the degree of externalities generated due to the change in the rate of consumer participation induced by these price changes. As JW argue, so long as the induced change in price is small, the first effect may be measured using the standard Jevons (1871)-Hotelling (1938) rule: $-\Delta p \cdot q$. The second effect consists of two parts: the harms caused by the increased prices charged for increased externalities and the benefits brought by these increased externalities themselves. Because the former depend on the benefits derived only by marginal users and the latter depend on the benefits delivered only to average users, the difference between these closely resembles the Spence distortion. The total local approximation may thus be written as a sum of Jevons-Hotelling effects and Spence effects, multiplied by the number of consumers experiencing these:

$$\begin{bmatrix} \left(\overline{\mathbf{V}}_{\mathcal{B}}^{\mathcal{A}} - \frac{\partial \overline{\mathbf{P}}^{\mathcal{A}}}{\partial N^{\mathcal{B}}} \right) \frac{\partial N^{\mathcal{B}}}{\partial \mathbf{P}^{\mathcal{B}}} \Delta \mathbf{P}^{\mathcal{B}} - \Delta \mathbf{P}^{\mathcal{A}} \\ \left(\overline{\mathbf{V}}_{\mathcal{A}}^{\mathcal{B}} - \frac{\partial \overline{\mathbf{P}}^{\mathcal{B}}}{\partial N^{\mathcal{A}}} \right) \frac{\partial N^{\mathcal{A}}}{\partial \mathbf{P}^{\mathcal{A}}} \Delta \mathbf{P}^{\mathcal{A}} - \Delta \mathbf{P}^{\mathcal{B}} \end{bmatrix}^T \cdot \begin{bmatrix} N^{\mathcal{A}} \\ N^{\mathcal{B}} \end{bmatrix}, \quad (18)$$

where $\Delta \mathbf{X}$ denotes the (small) difference between the pre- and post-merger value of vector \mathbf{X} , under Insulated Equilibrium. The change in prices counted here is only the *directly* induced change, that is change *in* the insulating tariff, not *along* it. The matrix $\overline{\mathbf{V}}_{\mathcal{I}}^{\mathcal{J}}$ is diagonal and has generic diagonal element $\overline{v}_j^{\mathcal{I},j}$, which, recall from Section 6, denote the average valuation for an additional interaction among the set of *all* side \mathcal{I} consumers on

platform j . From our discussion in Section 6, it should not be surprising that the Jacobian matrix of the insulating tariff $\frac{\partial P^I}{\partial N^J}$ plays a role, under oligopoly, that is analogous to that of the average value of marginal consumers, under monopoly.

Calculating the appropriate extension of Farrell and Shapiro (2010) (FS)'s *Upward Pricing Pressure* (UPP) to this context is relatively straightforward.²⁶ The first term is exactly as in FS, the value (in terms of profits, that is the mark-up) of sales of platform k diverted as a result of one more slot on platform j being filled. The second, novel term arises from the fact that, post-merger, platform j must now consider not only how increasing its participation positively impacts the externalities for which it can charge consumers on the other side of the market but also how it negatively impacts the externalities for which the *merger partner* can charge on the other side. Without loss of generality, we assume the merger occurs between platforms 1 and 2; in this case the UPP vector is given by

$$\tau^{I,j} = \begin{cases} \mathbf{D}_{k,j}^I (P^{I,k} - C_I^k) + N^{J,k} \left[\frac{\partial P^{J,k}}{\partial N^I} \right] \times [\mathbf{D}_{\cdot,j}^I], & \text{if } j, k = 1, 2, k \neq j \\ 0, & \text{if } j \neq 1, 2 \end{cases}.$$

JW show that if sufficient technical conditions are satisfied (for example, the pre- and post-merger equilibria must be stable in a strong sense) and the product of the pass-through matrix and UPP, $\mathfrak{S}^{-1} [T_U] \tau$, is sufficiently small, then

$$\Delta \mathbf{P} \approx \left([T_U]^{-1} \mathfrak{S} - \nabla_P \tau \right)^{-1} \tau,$$

where all quantities are evaluated at the pre-merger allocation. This formula is perfectly analogous to the one pertaining to the standard markets that JW consider, with the exception of what enters into the determination of τ . These local approximations of price changes may then be inserted into expression (18), where all other terms in that expression are also evaluated at the pre-merger allocation, to obtain a first-order approximation to the full effect of the merger on consumer welfare. Note that we may also obtain independent approximations of the effect of a merger on each side's welfare by simply evaluating each side independently, rather than summing over the two.

To summarize, we extend JW's formula (quantities multiplied by pass-through, multiplied by the value of diverted sales) in two ways:

1. The value of diverted sales is extended to include the *full opportunity cost* of those diverted sales in a two-sided setting. This value takes into account both the direct mark-up that diverted sales bring on the side of the market in question and their

²⁶Note that because we assume Bertrand conduct, we need consider only UPP and not JW's generalization of it, GePP, which allows more general conduct.

impact on each of the merging platform's ability to extract value from externalities, as perceived by marginal consumers, on the opposite side of the market.

2. The effects of the predicted price changes are also accounted for via their impact on participation and, consequently, on the externalities experienced by users on the other side.

8.2 Generalizations

Many Sides of the Market

Thus far, for expository purposes, we have focused on market configurations with two "sides" or groups of consumers. The model easily extends to accommodate an arbitrary number of sides. To see this, suppose there are S groups of consumers, indexed by $\mathcal{I} = \mathcal{A}, \mathcal{B}, \mathcal{C}, \dots$, and let the gross payoff of joining a bundle of platforms, \mathcal{X} , to a consumer of type $\theta^{\mathcal{I}}$ on side \mathcal{I} be $v^{\mathcal{I}}(\mathcal{X}, N^{-\mathcal{I}}, \theta^{\mathcal{I}})$, where $v^{\mathcal{I}} : \wp(\mathcal{M}) \times [0, 1]^{m(S-1)} \times \Theta^{\mathcal{I}} \rightarrow \mathbb{R}$ now depends on $N^{-\mathcal{I}} \in [0, 1]^{m(S-1)}$, the coarse allocation on the $S - 1$ other sides of the market apart from side \mathcal{I} . Also, let platform j 's strategy now be given by $\sigma^j \equiv (\sigma^{\mathcal{A},j}(N^{-\mathcal{A}}), \sigma^{\mathcal{B},j}(N^{-\mathcal{B}}), \sigma^{\mathcal{C},j}(N^{-\mathcal{C}}), \dots)$, where $\sigma^{\mathcal{I},j} : [0, 1]^{m(S-1)} \rightarrow \mathbb{R}$ maps from $N^{-\mathcal{I}} \in [0, 1]^{m(S-1)}$ to a total price that side \mathcal{I} consumers pay to join platform j .

It is straightforward to see that, when the model is extended in this way, none of the arguments made thus far in the paper depend on the presence of only two sides. In particular, the result of Theorem 1, that a CNE coarse allocation implies a price vector, continues to hold. Thus, the simplest way to consider a platform's profit maximization problem continues to be as a choice of allocation, holding fixed the strategies of the other platforms

$$\max_{\{N^{\mathcal{A},j}, N^{\mathcal{B},j}, N^{\mathcal{C},j}, \dots\}} \sum_{\mathcal{I}=\mathcal{A}, \mathcal{B}, \mathcal{C}, \dots} N^{\mathcal{I},j} P^{\mathcal{I},j}(N^{\mathcal{I},j}, N^{-\mathcal{I}}) - C^j(N^{\mathcal{A},j}, N^{\mathcal{B},j}, N^{\mathcal{C},j}, \dots). \quad (19)$$

Analogously to the results of Section 6.1, the prices that implement the socially optimal allocation satisfy

$$P^{\mathcal{I},j} = C^j_{\mathcal{I}} - \sum_{\mathcal{J} \neq \mathcal{I}} N^{\mathcal{J},j} \overline{v^{\mathcal{J},j}}, \quad (20)$$

and the platforms' prices under Insulated Equilibrium satisfy

$$P^{\mathcal{I},j} = C^j_{\mathcal{I}} + \mu^{\mathcal{I},j} - \sum_{\mathcal{J} \neq \mathcal{I}} N^{\mathcal{J},j} \left(\left[-\frac{\partial N^{\mathcal{J}}}{\partial P^{\mathcal{J}}} \right]^{-1} \left[\frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}} \right] \right)_{\mathcal{J}} \cdot [-\mathbf{D}^{\mathcal{I}}_{\cdot j}]. \quad (21)$$

The only difference between these expressions and those discussed in Section 6 is that here, since there are S sides of the market, the number of consumers on side \mathcal{I} affects the payoffs of consumers on all of the $S - 1$ other sides. Consequently, the prices charged to side \mathcal{I} consumers under both the socially optimal allocation and the Insulated Equilibrium allocation take into account the sum of such externalities, with the latter still subject to both the market power and Spence distortions.

Within-Side Externalities

Until now, we have also assumed that consumers' preferences over platforms are independent of the number of consumers *of the same group* that join each platform. This section extends the model to allow for such within-side network effects, which play a significant role in many industries, such as the provision of mobile phone service and social networking websites. Note that, while our focus is indeed on competition among *multi-sided* platforms, embedded in this generalization is the case, where $S = 1$, of competition among one-sided network providers, as in the literature stemming from the seminal paper of Katz and Shapiro (1985).

When joining a bundle of platforms, \mathcal{X} , a consumer of type $\theta^{\mathcal{I}}$ on side \mathcal{I} receives gross payoff $v^{\mathcal{I}}(\mathcal{X}, N, \theta^{\mathcal{I}})$, where $v^{\mathcal{I}} : \wp(\mathcal{M}) \times [0, 1]^{m^S} \times \Theta^{\mathcal{I}} \rightarrow \mathbb{R}$ depends on N , the *entire* coarse allocation.

We extend platforms' strategy space in the way that allows for the solution concept of Insulated Equilibrium to be most naturally preserved. Let platform j 's strategy be given by $\sigma^j \equiv (\sigma^{\mathcal{A},j}(N), \sigma^{\mathcal{B},j}(N), \sigma^{\mathcal{C},j}(N), \dots)$, where $\sigma^{\mathcal{I},j} : [0, 1]^{m^S} \rightarrow \mathbb{R}$ maps from N , the entire coarse allocation of consumers, including on side \mathcal{I} , to a total price.

When there are within-side network externalities, in order to speak of Insulating Tariffs, it becomes convenient to introduce the notion of consumers' *beliefs*, as discussed by Katz and Shapiro (1985),²⁷ about the strategies of other consumers on the same side. Suppose that prior to choosing their actions, side \mathcal{I} consumers form beliefs about one another's strategies, which, for our purposes, it is not restrictive to assume are common among all consumers. Formally, let $\overset{\dots}{N}^{\mathcal{I}}$ denote the coarse allocation on side \mathcal{I} that side \mathcal{I} consumers believe will prevail in the given consumer game. It is apparent that the best-response strategy profile of side \mathcal{I} consumers depends on $\overset{\dots}{N}^{\mathcal{I}}$. Residual Insulating Tariffs, whose definition we now restate, adapted to this context, pin down a unique value of $\overset{\dots}{N}^{\mathcal{I}}$ and thus a unique value of $N^{\mathcal{I}}$ that is consistent with Consumer Nash Equilibrium.

²⁷The object that we refer to as "beliefs" is, in fact, called "expectations" by Katz and Shapiro (1985) but is given the former name in more recent literature such as Caillaud and Jullien (2003).

Definition 8. Given a profile of strategies of other platforms, σ^{-j} , platform j is said to charge a Residual Insulating Tariff on side \mathcal{I} if $\forall \ddot{N}^{\mathcal{I}}, \widetilde{\ddot{N}}^{\mathcal{I}} \in [0, 1]^m$ and $\forall N^{-\mathcal{I}}, \widetilde{N}^{-\mathcal{I}} \in [0, 1]^{m(S-1)}$

$$N^{\mathcal{I},j} \left[\mathcal{M}^{\mathcal{I}*} \left(\theta^{\mathcal{I}}, \left(\ddot{N}^{\mathcal{I}}, N^{-\mathcal{I}} \right), [\sigma] \right), \sigma \right] = N^{\mathcal{I},j} \left[\mathcal{M}^{\mathcal{I}*} \left(\theta^{\mathcal{I}}, \left(\widetilde{\ddot{N}}^{\mathcal{I}}, \widetilde{N}^{-\mathcal{I}} \right), [\sigma] \right), \sigma \right].$$

As before, all platforms announcing Residual Insulating Tariffs on all sides of the market, gives rise to an Insulating Tariff System, $\overline{P}(N)$, anchored at a reference allocation. Insulated Equilibrium thus continues to be defined as in Definition 7, and the shape of the Insulating Tariff System, in response to variation in the own side coarse allocation is pinned down by the equation

$$0 = \left[\frac{\partial N^{\mathcal{I}}}{\partial \overline{P}^{\mathcal{I}}} \right] \left[\frac{\partial \overline{P}^{\mathcal{I}}}{\partial \ddot{N}^{\mathcal{I}}} \right] + \left[\frac{\partial N^{\mathcal{I}}}{\partial \ddot{N}^{\mathcal{I}}} \right] \Leftrightarrow \left[\frac{\partial \overline{P}^{\mathcal{I}}}{\partial \ddot{N}^{\mathcal{I}}} \right] = \left[-\frac{\partial N^{\mathcal{I}}}{\partial \overline{P}^{\mathcal{I}}} \right]^{-1} \left[\frac{\partial N^{\mathcal{I}}}{\partial \ddot{N}^{\mathcal{I}}} \right]. \quad (22)$$

When platforms' tariffs satisfy equation (22), for *any* beliefs that side \mathcal{I} consumers might have, prices adjust to maintain a given CNE coarse allocation. Therefore, there is a unique coarse allocation that consumers can consistently believe will occur in equilibrium.

Platform j 's profit maximization problem continues to be given by expression (19), and the prices that arise under the socially optimal allocation and the Insulated Equilibrium allocation are, respectively,

$$P^{\mathcal{I},j} = C_I^j - \sum_{\mathcal{J}=\mathcal{A},\mathcal{B},\mathcal{C}\dots} N^{\mathcal{J},j} \overline{v}_j^{\mathcal{J},j} \quad (23)$$

and

$$P^{\mathcal{I},j} = C_I^j + \mu^{\mathcal{I},j} - \sum_{\mathcal{J}=\mathcal{A},\mathcal{B},\mathcal{C}\dots} N^{\mathcal{J},j} \left(\left[-\frac{\partial N^{\mathcal{J}}}{\partial \overline{P}^{\mathcal{J}}} \right]^{-1} \left[\frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}} \right] \right)_{j,\cdot} \cdot \left[-\mathbf{D}_{\cdot,j}^{\mathcal{I}} \right]. \quad (24)$$

Notice that the only difference between equations (23) and (24), corresponding to the case where there are within-side externalities, and equations (20) and (21), corresponding to the case without such effects, is in the final term. When there are within-side externalities and firm j changes the number of consumers it serves on side \mathcal{I} , this affects the quality of j as perceived by side \mathcal{I} consumers in addition to consumers of all other sides.

8.3 Empirical Application of an Affine Discrete Choice Model

An important motivation for our work this is to help extend the tools for empirical research in industrial organization developed during the past two decades to allow for consump-

tion externalities. While an extensive treatment of how to apply these tools in our context is beyond the scope of this paper, it is a particularly promising area for future research. In view of this potential, here, we briefly speculate on the possible road ahead in this dimension. We draw attention to two points that seem particularly notable: first, the connection between “characteristics space” representations and “random coefficients” models on the one hand and the estimation of externalities and the “Spence distortion” in our setting, and, second, the possibility, under the solution concept of Insulated Equilibrium, of jointly estimating a demand system using both demand and supply equations.

Regarding the first point, there has been an increasing focus in the last two decades, stimulated particularly by Berry (1994) and BLP, on using characteristics-based representations of demand systems to reduce the dimensionality of demand estimation. In our setting, such representations have an additional relative benefit compared to product-based demand systems, since consumer valuations of network effects are of direct interest, not merely indirectly useful for estimation of demand for products. Furthermore, the concurrent development of *random* rather than simple logit models, originally stimulated by a desire to accommodate more realistic substitution patterns, is highly complementary with applications of our framework. Only a random coefficient model allows for the heterogeneity among consumers in their valuation for network effects which generates the possibility of Spence distortions.²⁸ Since, as we have shown, the presence of network externalities and particularly the Spence distortion are important forces that can exacerbate or counteract the ill effects of market power, in a two-sided setting, there is an additional argument in favor of the popular random coefficient, characteristic-based approach.

To see how such an approach could proceed, consider the simplest specification of demand in our model. This is a combination of the model of the affine preference specification of Rochet and Tirole (2006) with Armstrong (2006)’s assumption in Section 4 (standard in the discrete choice demand estimation literature), that consumers must “single-home” (purchase at most one product) and the standard characteristics assumption of Berry (1994) that valuation of product characteristics is common across products (here platforms). In particular, consumer i on side I ’s utility from consuming singleton bundle j would be

$$v_i^{I,j} = \beta_i^I N^{\mathcal{J},j} + \eta_i^{I,j} - \alpha_i^I P^{I,j}$$

where β_i^I represents the firm-homogeneous random coefficient of consumer i for network externalities from side \mathcal{J} , $\eta_i^{I,j}$ represents all non-network-generated value to i from platform j including mean utility, idiosyncratic valuation of non-network characteristics and

²⁸Note that this is true as well in one-sided models with endogenous characteristic choices.

good-consumer idiosyncratic errors (typically assumed Type I Extreme Value distributed for tractability) and α_i^T is the distaste for price (usually assumed to be income-related). The Spence distortion would then be a function of the correlation across consumers between $\frac{\beta_i^T}{\alpha_i^T}$ and factors entering into $\frac{\eta_i^{T,j}}{\alpha_i^T}$ and thus leading users to be marginal or infra-marginal. A particularly natural such relationship is obviously income heterogeneity (heterogeneity in α_i^T), but others might be coefficients of the same (or opposite) sign in the regression of β_i^T and $\eta_i^{T,j}$ on the same demographic characteristics or correlation between demographic characteristics, with systematically related coefficient signs in these two regressions.

One difficulty in estimating β_i^T is the likely correlation of $N^{\mathcal{J},j}$ with unobserved platform characteristics. In one-sided demand estimation, this issue is typically thought to arise with respect to prices, and the customary (Nevo, 2000; Akerberg, Benkard, Berry, and Pakes, 2007) approach to dealing with it is to use instruments for price. The natural extension of this approach to our context would be to use additional instruments for the number of opposite-side consumers. It appears to us that the criteria for evaluating the appropriateness of such instruments for opposite side participation should be similar to those for price instruments *on the other side of the market*. If this is correct, the same instruments used for $P^{\mathcal{J},j}$ could be used for $N^{\mathcal{J},j}$ (as price affects demand) in the estimation of the demand system on the other side of the market. Regardless, this seems like an important area for further research.

The second point mentioned above is the possibility, under Insulated Equilibrium, of jointly estimating a parameterized demand system, using both demand and supply equations. An implication of Armstrong's Paradox is that Subgame Perfect Equilibrium does not yield platform first-order conditions that could be used for such purposes, since they depend on platforms' off-equilibrium beliefs. As Theorem 3 shows, however, Under Insulated Equilibrium, the first-order condition is expressed only as a function of market-level observables, marginal costs and the demand system. Thus, while the pricing equations have an additional "two-sided" term, the only substantive requirement for imposing these equations, compared to the one-sided case, is to make use of the derivatives of each platform's market share, not only with respect to own and other platforms' prices, but also with respect to opposite side *participation*. Evaluating such a derivative is a substantively, and thus we suspect computationally, analogous exercise: the price derivative effectively involves computing average value of α_i^T along the set of marginal consumers while the participation derivative involves computing the average value of β_i^T (both involve also computing the size/density of the marginal set).

9 Conclusion

This paper aspires to make three contributions. First it develops, for the first time, a model with generality comparable to that of standard static industrial organization models, but incorporating the “multi-sided platforms” features of multiple goods and consumption externalities. Second, it develops a conceptual approach, extending the notion of the allocation approach to oligopoly and proposing the solution concept of *Insulated Equilibrium*, that allow this broad model to be analyzed. Finally, it shows how a natural extension of the logic of Spence (1975) can be used to understand both the distortions created by oligopolistic market power and the capacity of competition to remedy these.

While we believe this constitutes one important step forward in the literature on multi-sided platforms, it is certainly no more than that: much remains to be done, both for us and others. We therefore now briefly discuss both some extensions we plan to develop going forward, as well as some of what we consider the most promising directions for future research by others.

We hope to generalize and extend our analysis in a number of ways. First, given the current limitations of the theory of general equilibrium with indivisible goods, we limited our model to consumption patterns exhibiting gross substitutes. We do not believe this is necessary and in particular suspect that a continuum of richly heterogeneous agents, as we assume, may be sufficient to ensure the existence and uniqueness of a price vector supporting a given allocation. If this conjecture is correct, it should be straightforward to allow arbitrary substitution patterns in our model. On a similar technical level, a more detailed analysis of conditions for existence, stability and uniqueness of IE would be useful.

More substantive generalization would also be useful, most importantly dealing with multi-homing and other sources of heterogeneity of externalities across users within a side of the market. If third-degree price discrimination is possible to all groups bringing different externalities either exogenously or endogenously through their choice of platforms (it may be less effective to advertise to a reader who has already seen the advertisement in another paper), it is relatively straightforward to extend our model to allow such externality or bundle-contingent contracts by simply increasing the number of sides. However we did not explicitly discuss this above because it is not very realistic in many settings. More promising, therefore, is the prospect of combining into our model the analysis of Veiga and Weyl (2010) which allows within-side heterogeneity while still permitting rich preference heterogeneity.

Finally, we would like to analyze a number of other substantive issues using the framework. These include regulation, such as price and quantity controls that are relevant in, for example, the analysis of network neutrality policies and a more general characterization of the cases in which intensifying competition helps alleviate, or exacerbate, the Spence distortion. Perhaps most importantly, we would like to build a workhorse parametric version of the model in the spirit of BLP, as outlined in Section 8.3, that could be applied in a range of empirical settings.

Beyond our work here, our paper suggests many natural directions for future research. Most clearly, relaxing the assumptions (the “macroness” of the model, homogeneous quality, no price discrimination) we discussed in Section 2 is important for the literature to progress. Our solution concept also seems naturally connected to a number of other problems in economics; elucidating these connections would help unify these areas. Most clearly, White is currently constructing a model, with Germain Gaudin, that builds on the techniques developed in this paper to study the effect of competition on the quality provision by one-sided firms. Similarly by the Bulow and Roberts (1989) equivalence, W10 is equivalent to Segal (1999)’s general model of contracting with externalities with asymmetric information. Thus it seems natural that our model should be closely related to common agency with multiple agents, externalities and asymmetric information. It would therefore be interesting to consider whether Insulated Equilibrium has a natural analogy to solution concepts invoked in that literature, or whether it offers a potential alternative concept.

At a deeper theoretical level, it would be interesting to understand more clearly the dynamic incentives of multi-sided platforms, in the spirit of Chen et al. (2009) and Cabral (forthcoming), and whether these lead to price paths resembling Insulating Tariff Systems. Also the intersection of profit maximization and matching market design (Roth, 2002) is conspicuously limited but very promising; see Gomes (2009) for an exception proving the rule.

On the applied side, we believe our paper offers a number of tools that make possible a range of interesting empirical analyses of multi-sided platforms, measuring market power and Spence distortions and predicting counter-factual effects of policy interventions, which we hope will develop in coming years. Making the theory of multi-sided platforms useful to policy makers will also require enriching our model to consider issues that are beyond the scope of this paper such as interconnection, vertical restraints, bundling, predation and regulatory design. We are thus hopeful that the theory and measurement of multi-sided platform industries will be increasingly put to use in helping to clarify an important and often ideologically-driven set of industrial policy debates.

Appendices

A Omitted Proofs

A.1 Proof of Lemma 1

Proof. First consider the forward direction, supposing the result failed. Note that if the maximization condition fails, then there exists another marketwide consumer strategy profile $\widehat{\mathcal{M}}$ yielding the same coarse allocation, at least as great gross consumer surplus on each side and strictly greater gross consumer surplus on the other side. Consider an arrangement in which each consumer made a payment of her gross utility under $\widehat{\mathcal{M}}$ and received a payment of her gross utility under $\widehat{\mathcal{M}}$, then $\widehat{\mathcal{M}}$ is implemented and the net revenue obtained is distributed equally among all consumers. Note that this is clearly a Pareto improvement for consumers over the CNE as they receive the same utility before the distribution of net revenue as before the change and the net revenue is strictly positive and thus improves their utility.

However, in CNE consumers maximize their utility taking prices as given and thus a CNE inducing coarse allocation N must be a general equilibrium of an economy with endowment N . But by the first fundamental theorem of welfare economics, such an equilibrium must be Pareto efficient, contradicting our premise and establishing the result.

In the reverse direction, note that by the Kelso and Crawford (1982); Gul and Stacchetti (1999) theorems, for any allocation N an general equilibrium price vector P^* exists. Let

$$\mathcal{M}^{I*}(\theta^I) \equiv \arg \max_{\mathcal{X} \in (\mathcal{M})} v^I(\mathcal{X}, N^J, \theta) - \sum_{j \in \mathcal{X}} P^{I,j},$$

with Assumption 2 breaking any ties so this is a function not a correspondence. This is by construction a CNE of the game where $\sigma \equiv P^*$. Thus, by the forward direction of the proof, it is an gross-surplus-maximizing strategy profile given the allocation N . \square

A.2 Proof of Lemma 2.

Proof. Consider an allocation N . By Lemma 1 there exists some general equilibrium price vector P^* “supporting” this allocation in the sense that CNE strategies given these prices (in the sense of the proof of Lemma 1 above and denoted by $\widehat{\mathcal{M}}^*$) induce this allocation.

We will show that $\left. \frac{\partial V^I}{\partial N^{I,j}} \right|_N$ exists and equals $P^{I,j}$.

Smoothness, Full Support and Gross Substitutes imply that, given N^J there is a smooth, full support distribution over all gross utility profiles, \mathbf{U}^I , in \mathbb{G} , where \mathbb{G} is the subset of \mathbb{R}^{2^m-1} exhibiting gross substitutes. Denote this distribution by $g^{I^*}(\mathbf{U}^I)$ and denote the measure of set under g by $\gamma(\Omega)$ where $\Omega \subset \mathbb{G}$.

Let $\Omega^* \equiv \{\mathbf{U} \in \mathbb{G} : \forall \mathcal{X} \neq \{j\}, \sum_{k \in \mathcal{X}} P^{I,k} > U_{\mathcal{X}}\}$; this is the *j-isolated set* of consumers who find all bundles other than the singleton j bundle less attractive than the null bundle. Clearly a consumer in Ω^* with $U_{\{j\}}$ weakly greater than $P^{I,j}$ consume the singleton j bundle at the associated CNE and those with $U_{\{j\}}$ strictly below $P^{I,j}$ consume the null bundle. Furthermore the marginal distribution over $U_{\{j\}}$ clearly has full support by the same logic as above.

Let $\Omega_{+\delta}^* \equiv \{\mathbf{U} \in \Omega^* : P^{I,j} + \delta > U_{\mathcal{X}} > P^{I,j}\}$ and similarly for $\Omega_{-\delta}^*$. For any $\delta > 0$, let $\epsilon_+(\delta) \equiv \gamma(\Omega_{+\delta}^*)$ and similarly for ϵ_- . Note that this is well-defined by full support and the fact that the distribution is non-atomic, and for the same reasons is a strictly positive, strictly increasing function over its domain.

To establish the existence of $\left. \frac{\partial V^I}{\partial N^{I,j}} \right|_N$ and its equality to $P^{I,j}$ we bound all of the Dini partial derivatives appropriate to demonstrate their equality to one another and to $P^{I,j}$. In particular, using the standard convention we will denote the Dini analog of $\left. \frac{\partial V^I}{\partial N^{I,j}} \right|_N$ by $\left. \frac{\partial V^I}{\partial N^{I,j}} \right|_N^+$ upper from the right, $\left. \frac{\partial V^I}{\partial N^{I,j}} \right|_N^-$ lower from the right, $\left. \frac{\partial V^I}{\partial N^{I,j}} \right|_{+,N}$ upper from the left and finally $\left. \frac{\partial V^I}{\partial N^{I,j}} \right|_{-,N}$ lower from the left. Clearly to establish existence and equality it suffices to show that both upper Dini derivatives are bounded above by $P^{I,j}$ and both lower Dini derivatives are bounded below by $P^{I,j}$. We will formally derive the bounds on the lower right Dini derivatives, somewhat formally consider the upper right and then informally describe the logic for the left Dini derivatives, as this logic is closely analogous.

First consider the lower right Dini derivative:

$$\left. \frac{\partial V^I}{\partial N^{I,j}} \right|_N^- \equiv \liminf_{\epsilon \rightarrow 0^+} \frac{V^I(N^I + \mathbf{1}_j \epsilon, N^J) - V^I(N^I, N^J)}{\epsilon}$$

where $\mathbf{1}_j$ is a vector of all 0's except that it has a 1 in the j^{th} entry. To show this is bounded below by $P^{I,j}$ it is sufficiently to demonstrate that for any $\delta > 0$ it is possible to find a sufficiently small $\epsilon > 0$ such that

$$V^I(N^I + \mathbf{1}_j \epsilon, N^J) > V^I(N^I, N^J) + \epsilon(P^{I,j} - \delta)$$

Suppose we are challenged with some $\delta > 0$; we claim this bound must hold if we select any $\epsilon < \epsilon_-(\delta)$. To see this note that one feasible assignment under $(N^I + \mathbf{1}_j \epsilon, N^J)$ is one

identical to $\widehat{\mathcal{M}}^*$ except that all consumers in $\Omega_{-\epsilon^{-1}(\epsilon)}^*$ are assigned to consume $\{j\}$; note that ϵ_- is clearly invertible because it is strictly increasing. Call this assignment $\widehat{\mathcal{M}}_{-\epsilon}^*$. This is clearly feasible because, by construction, these consumers represent a mass ϵ . Furthermore

$$V^I \left(\left[\widehat{\mathcal{M}}_{-\epsilon}^* \right], N^{\mathcal{J}} \right) \geq \left(P^{I,j} - \epsilon^{-1}(\epsilon) \right) \epsilon + V^I \left(\left[\widehat{\mathcal{M}}^* \right], N^{\mathcal{J}} \right) = \left(P^{I,j} - \epsilon^{-1}(\epsilon) \right) \epsilon + V^I(N)$$

by construction. But given that this assignment is feasible, the optimal assignment must do at least as well as it and given that ϵ^{-1} is strictly increasing our claim is established given that if $\epsilon^{-1} < \delta$ clearly $\epsilon \left(P^{I,j} - \delta \right) < \left(P^{I,j} - \epsilon^{-1}(\epsilon) \right) \epsilon$.

Now consider the upper right Dini derivatives:

$$\left. \frac{\partial V^I}{\partial N^{I,j}} \right|_N^- \equiv \limsup_{\epsilon \rightarrow 0^+} \frac{V^I(N^I + \mathbf{1}_j \epsilon, N^{\mathcal{J}}) - V^I(N^I, N^{\mathcal{J}})}{\epsilon}$$

Suppose, contrary to what we want to demonstrate, that this were strictly greater than $P^{I,j}$. Then

$$\forall \bar{\epsilon} > 0, \exists \epsilon < \bar{\epsilon} : \frac{V^I(N^I + \mathbf{1}_j \epsilon, N^{\mathcal{J}}) - V^I(N^I, N^{\mathcal{J}})}{\epsilon} > \frac{\left. \frac{\partial V^I}{\partial N^{I,j}} \right|_N^- + P^{I,j}}{2}$$

or

$$V^I(N^I + \mathbf{1}_j \epsilon, N^{\mathcal{J}}) - \epsilon \left(\frac{\left. \frac{\partial V^I}{\partial N^{I,j}} \right|_N^- + P^{I,j}}{2} \right) > V^I(N^I, N^{\mathcal{J}})$$

But note that we can construct an isolated set just as above at the allocation $(N^I + \mathbf{1}_j \epsilon, N^{\mathcal{J}})$ and that by Gross Substitutes (which implies the law of demand) any general equilibrium price vector at this allocation must have lower $P^{I,j}$. Thus we could always achieve a utility close to $V^I(N^I + \mathbf{1}_j \epsilon, N^{\mathcal{J}})$ under the allocation N by using the optimal assignment at that allocation and removing a group of isolated near-marginal consumers. Thus a contradiction can easily be established to $\left. \frac{\partial V^I}{\partial N^{I,j}} \right|_N^- \neq P^{I,j}$.

Similar arguments may then be applied for the left lower and left upper Dini derivatives. If the left upper were too large we could always have evicted from consumption the near-marginal consumers in the isolated set and this would have been a less costly way of removing consumers. If the left lower were too small then we could have added more near-marginal types at the lower allocation in and have achieved a higher utility at the original allocation that was posited. Thus all necessary bounds can be established to show differentiability and equality of the the derivative to $P^{I,j}$. \square

References

- Akerberg, Daniel, C. Lanier Benkard, Steven Berry, and Ariel Pakes (2007), "Econometric Tools for Analyzing Market Outcomes." *Handbook of Econometrics*, 6A, 4171–4276.
- Ambrus, Attila and Rossella Argenziano (2009), "Asymmetric Networks in Two-Sided Markets." *American Economic Journal: Microeconomics*, 1, 17–52.
- Amir, Rabah and Val E. Lambson (2000), "On the Effects of Entry in Cournot Markets." *The Review of Economic Studies*, 67, 235–254.
- Anderson, Robert M., Glenn Ellison, and Drew Fudenberg (2010), "Location Choice in Two-Sided Markets with Indivisible Agents." *Games and Economic Behavior*, 69, 2–23.
- Anderson, Simon P. and Stephen Coate (2005), "Market Provision of Broadcasting: A Welfare Analysis." *The Review of Economic Studies*, 72, 947–972.
- Anderson, Simon P., André de Palma, and Jacques-François Thisse (1992), *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, Massachusetts.
- Argentesi, Elena and Lapo Filistrucchi (2007), "Estimating Market Power in a Two-Sided Market: The Case of Newspapers." *Journal of Applied Econometrics*, 22, 1247–1266.
- Armstrong, Mark (2006), "Competition in Two-Sided Markets." *The RAND Journal of Economics*, 37, 668–691.
- Armstrong, Mark and Julian Wright (2007), "Two-sided Markets, Competitive Bottlenecks and Exclusive Contracts." *Economic Theory*, 32, 353–380.
- Athey, Susan, Emilio Calvano, and Joshua Gans (2010), "Will the Internet Destroy News Media?" *Mimeo*.
- Bedre-Defolie, Özlem and Emilio Calvano (2009), "Pricing Payment Cards." *Mimeo*.
- Berry, Steven, James Levinsohn, and Ariel Pakes (1995), "Automobile Prices in Market Equilibrium." *Econometrica*, 63, 841–890.
- Berry, Steven T. (1994), "Estimating discrete-choice models of product differentiation." *The RAND Journal of Economics*, 25, pp. 242–262.
- Bulow, Jeremy and John Roberts (1989), "The Simple Economics of Optimal Auctions." *The Journal of Political Economy*, 97, 1060–1090.
- Cabral, Luís (forthcoming), "Dynamic Price Competition with Network Effects." *Review of Economic Studies*.
- Caillaud, Bernard and Bruno Jullien (2003), "Chicken & Egg: Competition among Intermediation Service Providers." *The RAND Journal of Economics*, 34, 309–328.

- Cantillon, Estelle and Pai-Ling Yin (2008), "Competition between Exchanges: Lessons from the Battle of the Bund." *Mimeo*.
- Chandra, Ambarish and Allan Collard-Wexler (2009), "Mergers in Two-Sided Markets: An Application to the Canadian Newspaper Industry." *Journal of Economics & Management Strategy*, 18, 1045–1070.
- Chen, Jiawei, Ulrich Draszelski, and Jr. Joseph E. Harrington (2009), "Avoiding Market Dominance: Product Compatibility in Markets with Network Effects." *RAND Journal of Economics*, 40, 455–485.
- Cournot, Augustin (1838), *Researches into the Mathematical Principles of the Theory of Wealth* (N.T. Bacon trans.), 1971 edition. Augustus M. Kelley, New York.
- Doğanoglu, Toker and Julian Wright (2006), "Multihoming and Compatibility." *International Journal of Industrial Organization*, 24, 45–67.
- Doğanoglu, Toker and Julian Wright (2010), "Exclusive Dealing with Network Effects." *International Journal of Industrial Organization*, 28, 145–154.
- Dybvig, Philip H. and Chester S. Spatt (1983), "Adoption Externalities as Public Goods." *Journal of Public Economics*, 20, 231–247.
- Economides, Nicholas (1996), "The Economics of Networks." *International Journal of Industrial Organization*, 14, 673 – 699.
- Ellison, Glenn and Drew Fudenberg (2003), "Knife-Edge or Plateau: When Do Market Models Tip?" *The Quarterly Journal of Economics*, 118, 1249–1278.
- Ellison, Glenn, Drew Fudenberg, and Markus Möbius (2004), "Competing Auctions." *Journal of the European Economic Association*, 2, 30 – 66.
- Enthoven, Alain C. and Kenneth J. Arrow (1956), "A Theorem on Expectations and the Stability of Equilibrium." *Econometrica*, 24, pp. 288–293.
- Evans, David S. (2003), "The Antitrust Economics of Multi-Sided Platform Markets." *Yale Journal of Regulation*, 20, 325–431.
- Farrell, Joseph and Carl Shapiro (2010), "Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition." *Berkeley Electronic Journal of Theoretical Economics: Policies and Perspectives*, 10.
- Gentzkow, Matthew (2007), "Valuing New Goods in a Model with Complementarity: Online Newspapers." *The American Economic Review*, 97, 713–744.
- Gomes, Renato (2009), "Mechanism Design in Two-Sided Markets: Auctioning Users." *Mimeo*.
- Gul, Faruk and Ennio Stacchetti (1999), "Walrasian Equilibrium with Gross Substitutes." *Journal of Economic Theory*, 87, 95 – 124.

- Hagiu, Andrei (2006), "Pricing and Commitment by Two-Sided Platforms." *The RAND Journal of Economics*, 37, 720–737.
- Hagiu, Andrei (2009a), "Quantity vs. Quality and Exclusion by Two-Sided Platforms." *Mimeo*.
- Hagiu, Andrei (2009b), "Two-Sided Platforms: Product Variety and Pricing Structures." *Journal of Economics & Management Strategy*, 18, 1011–1043.
- Hagiu, Andrei and Robin S. Lee (forthcoming), "Exclusivity and Control." *Journal of Economics & Management Strategy*.
- Hatfield, John William, Scott Duke Kominers, Alexandru Nichifor, Michael Ostrovsky, and Alexander Westkamp (2010), "Stability and Competitive Equilibrium in Trading Networks." *Mimeo*.
- Hendel, Igal (1999), "Estimating Multiple-Discrete Choice Models: An Application to Computerization Returns." *The Review of Economic Studies*, 66, 423–446.
- Hotelling, Harold (1929), "Stability in Competition." *The Economic Journal*, 39, 41–57.
- Hotelling, Harold (1938), "The General Welfare in Relation to Problems of Taxation and of Railway and Utility Rates." *Econometrica*, 6, 242–269.
- Jaffe, Sonia and E. Glen Weyl (2010a), "Linear Demand Systems are Inconsistent with Discrete Choice." *Mimeo*.
- Jaffe, Sonia and E. Glen Weyl (2010b), "The First-Order Approach to Merger Analysis." This paper is currently under preparation. Contact Weyl at weyl@fas.harvard.edu for notes.
- Jevons, William Stanley (1871), *The Theory of Political Economy*. London: MacMillan.
- Kaiser, Ulrich and Julian Wright (2006), "Price Structure in Two-Sided Markets: Evidence from the Magazine Industry." *International Journal of Industrial Organization*, 24, 1–28.
- Katz, Michael L. and Carl Shapiro (1985), "Network Externalities, Competition, and Compatibility." *The American Economic Review*, 75, 424–440.
- Kelso, Jr., Alexander S. and Vincent P. Crawford (1982), "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica*, 50, pp. 1483–1504.
- Klemperer, Paul D. and Margaret A. Meyer (1989), "Supply Function Equilibria in Oligopoly under Uncertainty." *Econometrica*, 57, 1243–1277.
- Lee, Robin S. (2010), "Dynamic Demand Estimation in Platform and Two-Sided Markets." *Mimeo*.
- Myerson, Roger B. (1981), "Optimal Auction Design." *Mathematics of Operations Research*, 6, 58–73.

- Nevo, Aviv (2000), "A Practitioners Guide to Estimation of Random-Coefficients Logit Models of Demand." *Journal of Economics & Management Strategy*, 9, 513–548.
- Nocke, Volker, Martin Peitz, and Konrad Stahl (2007), "Platform Ownership." *Journal of the European Economic Association*, 5, 1130–1160.
- Parker, Geoffrey G. and Marshall W. Van Alstyne (2005), "Two-Sided Network Effects: A Theory of Information Product Design." *Management Science*, 51, 1494–1504.
- Peitz, Martin and Tommaso M. Valletti (2008), "Content and Advertising in the Media: Pay-TV versus Free-to-Air." *International Journal of Industrial Organization*, 26, 949–965.
- Reisinger, Markus (2010), "Unique Equilibrium in Two-Part Tariff Competition between Two-Sided Platforms." *Mimeo*.
- Riley, John G. and William F. Samuelson (1981), "Optimal Auctions." *The American Economic Review*, 71, pp. 381–392.
- Rochet, Jean-Charles (2010), "The Pricing of Academic Journals: A Two-Sided Market Perspective." *American Economic Journal: Microeconomics*, 2, 222–55.
- Rochet, Jean-Charles and Jean Tirole (2003), "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association*, 1, 990–1029.
- Rochet, Jean-Charles and Jean Tirole (2006), "Two-Sided Markets: A Progress Report." *RAND Journal of Economics*, 37, 645–667.
- Rosse, James N. (1970), "Estimating Cost Function Parameters Without Using Cost Data: Illustrated Methodology." *Econometrica*, 38, 256–275.
- Roth, Alvin E. (2002), "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics." *Econometrica*, 70, 1341–1378.
- Rysman, Marc (2004), "Competition between Networks: A Study of the Market for Yellow Pages." *The Review of Economic Studies*, 71, 483–512.
- Rysman, Marc (2009), "The Economics of Two-Sided Markets." *Journal of Economic Perspectives*, 23, 125–143.
- Samuelson, Paul A. (1941), "The Stability of Equilibrium: Comparative Statics and Dynamics." *Econometrica*, 9, pp. 97–120.
- Segal, Ilya (1999), "Contracting with Externalities." *The Quarterly Journal of Economics*, 114, 337–388.
- Shaked, Avner and John Sutton (1982), "Relaxing Price Competition Through Product Differentiation." *The Review of Economic Studies*, 49, pp. 3–13.
- Spence, A. Michael (1975), "Monopoly, Quality, and Regulation." *The Bell Journal of Economics*, 6, 417–429.

- Sun, Mingchun and Edison Tse (2007), "When Does the Winner Take All in Two-Sided Markets?" *Review of Network Economics*, 6.
- U.S. Department of Justice and the Federal Trade Commission (2010), "Horizontal Merger Guidelines." URL <http://www.ftc.gov/os/2010/08/100819hmg.pdf>.
- Veiga, Andre and E. Glen Weyl (2010), "Multi-Sided Platforms with Heterogeneous Externalities." This work is currently in progress. Contact Weyl at weyl@fas.harvard.edu for notes.
- Werden, Gregory J., Luke M. Froeb, and David T. Scheffman (2004), "A Daubert Discipline for Merger Simulation." *Antitrust*, 18, 89–95.
- Weyl, E. Glen (2010), "A Price Theory of Multi-Sided Platforms." *American Economic Review*, 100, 1642–1672.
- Weyl, E. Glen and Jean Tirole (2010), "Materialistic Genius and Market Power: Uncovering the Best Innovations." *Mimeo*.
- Weyl, Glen and Michal Fabinger (2009), "Pass-Through as an Economic Tool." *Mimeo*.
- White, Alexander (2009), "Search Engines: Left Side Quality versus Right Side Profits." *Mimeo*.
- Wilbur, Kenneth C. (2008), "A Two-Sided, Empirical Model of Television Advertising and Viewing Markets." *Marketing Science*, 27, 356–378.