

TWELFTH CEPR/JIE CONFERENCE ON APPLIED INDUSTRIAL ORGANIZATION

Hosted by

The Eitan Berglas School of Economics, Recanati School of Business, Department of Public Policy at Tel Aviv University

Supported by

The Pinhas Sapir Center for Development, Tel Aviv University
Recanati School of Business, Tel Aviv University
Department of Public Policy, Tel Aviv University
Journal of Industrial Economics (JIE)
CEPR

Tel Aviv; 24-27 May 2011

Imperfect Platform Competition: A General Framework

***Alexander White (Harvard University)**

E. Glen Weyl (Harvard University and Toulouse School of Economics)

Imperfect Platform Competition:

A General Framework^{*}

Alexander White[†]

E. Glen Weyl[‡]

May 17, 2011

Abstract

The externalities that operating system users receive from software developers are among the leading features of those ‘platform’ industries, but are rarely incorporated into applied models of imperfect competition. We argue this omission arises from the difficulty of collapsing the dynamic pricing characterizing such industries into a static policy analysis model. Given the role these pricing strategies play in coordinating consumer behavior, a theory ignoring them quickly becomes intractable and indeterminate. Postulating that platforms *identify* and then *robustly implement* best response allocations, we show platforms play an *Insulated Equilibrium* that eliminates the need for consumers to coordinate their behavior. This facilitates the analysis of an oligopoly model without unrealistic restrictions imposed for tractability. We use this to illustrate the additional distortion, analogous to that identified by Spence (1975)’s study of a quality-choosing monopolist, arising when platforms determine both their prices and their (externality-driven) level of quality.

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 thank Alexei Alexandrov, Rabah Amir, Mark Armstrong, Susan Athey, Gary Biglaiser, Joyee Debb, Ed Glaeser, Michal Fabinger, Renato Gomes, Sonia Jaffe, Bruno Jullien, Scott Kominers, Greg Lewis, Thomas Mariotti, Ariel Pakes, In-Uck Park, Eric Rasmusen, Patrick Rey, François Salanié, Larry Samuelson, Paul Seabright, Yossi Spiegel, Jean Tirole, Catherine Tucker, André Veiga, Ken Wilbur, Julian Wright and especially Andrew Clausen and Jacques Crémer for very helpful comments and discussions. We thank Toan Phan, 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

A prominent aspect of competitive strategies in industries with consumption externalities is the tendency to charge low prices when participation is low, cashing in only once the service has matured, in order to avoid a ‘chicken-and-egg’ problem. The recent theoretical literature on ‘two-sided markets’ or ‘multi-sided platforms’ suggests, often implicitly, a technique to model this aspect of such industries: to allow platforms to charge prices that *depend on realized levels of demand*. This paper explicitly takes this idea to its logical conclusion, collapsing such coordination dynamics into a canonical static model of imperfect competition amenable to policy analysis.¹

The two-sided markets literature has both highlighted the pervasiveness of consumption externalities throughout different industries and argued 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. Meanwhile, policymakers have 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. However, quantitative evaluation of these claims requires a model flexible enough to incorporate rich structures of consumer preferences and firm heterogeneity, and rich models of ‘one-sided’ competition, such as Berry, Levinsohn, and Pakes (1995) (BLP), have been considered intractable in the platform context.²

We build a model whose generality is comparable to that of standard one-sided models of competition. In doing so, we illustrate that, in the absence of platform pricing policies that coordinate consumers, the possibility of multiple equilibria in which consumers either succeed or fail to coordinate on joining a platform, makes general analysis intractable. At the same time, as observed by Armstrong (2006), allowing platforms arbitrary conditional (reduced-form dynamic) strategies creates a different form of indeterminacy, in the spirit of Klemperer and Meyer (1989).

We argue that both of these problems are resolved by observing that platforms’ pricing strategies are *designed* to coordinate consumers on their desired outcome. Thus, we explicitly assume that each platform, taking as given other platform’s pricing strategies, *identifies* and then seeks to *robustly implement* its best response levels of consumer participation. Our approach preserves the simplicity of a static model, yet, through two crucial assumptions, it captures, in a reduced form manner, the mechanics of coordination that might otherwise require the specification of an elaborate dynamic process.

The first of these assumptions, *No Sunspots*, developed in Section 4, ensures that platforms can always identify their best response allocation. It posits that only changes in platform strategies which affect the *level of participation of consumers on that platform*, which we refer to as its *coarse*

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).

allocation or ‘allocation’ for short, can impact the coarse allocations of other platforms. Thus we rule out platforms’ strategies playing a pure ‘sunspot’ coordination role among consumers participating on other platforms, in the spirit of Cass and Shell (1983). Furthermore because the participation levels of users on one side of the market tie down the utility profiles of users on the other side we can apply the results from our recent work on exchange economies with a continuum of consumers (Azevedo, Weyl, and White, 2011), which guarantee that the coarse allocation ties down the prices of all platforms and thus their profits. Consequently, platforms need only identify their best response coarse allocation.

The second of these assumptions, *Insulation*, developed in Section 5, is inspired jointly by observations about dynamic pricing discussed above and the recent theoretical literature on robust implementation. Platforms’ dynamic pricing strategies, and sometimes even their static tariffs in the case of advertising pricing, typically adjust to shifts in the attractiveness of the platform, lowering or raising them as the platform becomes less or more attractive because of changes in participation on the other side of the market. We model this statically as an attempt to robustly implement (Bergemann and Morris, 2005) the platform’s desired allocation by providing consumers with, as near as possible, a dominant strategy (Chung and Ely, 2007) to either participate or not in the platform. When users are only heterogeneous along a single dimension, this is literally possible and in fact is the modeling strategy adopted by one of the foundational papers in the literature, Rochet and Tirole (2003) (RT2003). However, as argued by Weyl (2010), when consumers are richly heterogeneous a uniform price cannot induce dominant strategies. We thus, instead, assume that platforms adopt the ‘next best thing’, an extension to competition of Weyl’s notion of an *Insulating Tariff*.

A firm charging a *Residual Insulating Tariff* (RIT) ensures that, regardless of what happens on the opposite side of the market, the equilibrium allocation on this side of the market is insulated and thus remains at the level the platform desires to implement. Whenever it is feasible to implement in dominant strategies, the RIT does this; when it is not, the RIT makes the desired allocation a dominant strategy for the representative consumer on each side to choose this allocation. When all platforms use RITs, platforms are said to implement an *Insulating Tariff System* (ITS), ‘anchored’ at whatever allocation it implements. We call an equilibrium in which all platforms optimize, given the behavior of other platforms, by charging their part of an ITS an *Insulated Equilibrium* (IE). Because such a system of tariffs, anchored at any given allocation, are unique, platform incentives are well-defined at an IE. Thus, whether a given allocation constitutes an IE can be checked based on primitives or, equivalently, marginal costs may be recovered from the demand system à la Rosse (1970).

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 ef-

fect that an intensification of competition has on the market power distortion is well-known, the effect of such an intensification on the Spence distortion, and the consequences of this distortion for social welfare, depend crucially the nature of differentiation between platforms, of consumer heterogeneity and the distribution of demand. A preview of these findings appears in Subsection 2.1.

Crucially, though, our model, unlike those previously appearing in the literature, can accommodate 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 number and symmetry of platforms or (iv) 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 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 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_I^j denote marginal cost to platform j of serving side 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 I ,

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

where $J \neq I$. 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 $\left(\left[-\frac{\partial N^J}{\partial P^J} \right]^{-1} \left[\frac{\partial N^J}{\partial N^I} \right] \right)_{j,\cdot} \cdot [-\mathbf{D}^I_{\cdot,j}]$ collapses to the *average* willingness of a *marginal* consumer on side J to pay for the participation of a *marginal* consumer on side I . This is the part of the externality created by this marginal side I consumer that the platform can extract per consumer on side 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 $\frac{\partial N^J}{\partial P^J}$ is the density of side 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, $\frac{\partial N^J}{\partial P^J}$ is a natural multi-product extension of the ‘mass of marginal users’. Similarly, we show that the j^{th} diagonal entry of $\frac{\partial N^J}{\partial N^I}$ is the product of the density of j ’s marginal users and the average value these place on a marginal side 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 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 $\left(\left[-\frac{\partial N^J}{\partial P^J} \right]^{-1} \left[\frac{\partial N^J}{\partial N^I} \right] \right)_{j,\cdot} \cdot [-\mathbf{D}^I_{\cdot,j}]$ 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

Coordination and Pricing to Coordinate

Since the work of Rohlfs (1974), the literature on imperfect competition with consumption (or ‘network’) externalities has lacked a consensus view on how to model the process by which consumers coordinate their actions.³ For instance, in Katz and Shapiro (1985), platforms take consumer beliefs as given, while Caillaud and Jullien (2003) mainly assume that consumers coordinate on an ‘incumbent’ platform. Ellison and Fudenberg (2003), Ellison, Fudenberg, and Möbius (2004), Hagiu (2006) and Anderson, Ellison, and Fudenberg (2010) study the coordination problem in detail and find a large multiplicity of equilibria. Ambrus and Argenziano (2009) and Lee (2010) offer alternative notions of consumers coordinating always in their collective best interest to refine these equilibria. While each offers its own formal justification, a potential criticism of such approaches is their apparent attempt to predict the way in which a large and diffuse population of consumers will coordinate with one another.

A particularly influential approach is that adopted by RT2003. This model exploits the fact that, if all consumers’ values are *proportional* to the number of consumers on the other side, differing only in the constant of proportionality and prices are denominated *per-interaction*, then every consumer has a dominant strategy to join one platform or another, given prices. This eliminates the scope for mis-coordination and squares with the intuition we discuss in the introduction about the use of prices as an instrument to coordinate demand. It relies, however, on the particular preference structure assumed in this model and is not obviously generalizable.

The monopoly model in Armstrong (2006) follows a strategy that is similar in this respect. It assumes that consumers have homogeneous valuation for externalities, which allows for platforms to promise *utility* levels as in Dybvig and Spatt (1983) and Armstrong and Vickers (2001). W10 generalizes this, in the context of monopoly, to *Insulating Tariffs*, which apply in settings with rich heterogeneity. However, Armstrong also points out that in the context of competition, allowing platforms to charge prices that vary with the number of users on the other side of the market creates a paradox in the spirit of Klemperer and Meyer (1989). If one platform believes its rivals will lower prices dramatically when it loses consumers, this acts as a deterrent to intense competition. Thus, as we discuss in Section 5, *some* structure of tariffs can support nearly any outcome as an equilibrium among platforms. Approaches to refine these equilibria, such as assuming platforms prices do not depend on the number of consumers on the other side, which Armstrong considers, or by reference to price discriminatory motives (Reisinger, 2011) threaten to revive the consumer coordination problem.

Our approach thus stems from the following claim: the motivation behind the RT2003 tariffs

³We prefer the phrase ‘consumption externality’ over ‘network effect’ as the latter is easily confused with a number of other related but not equivalent concepts in the literature on airline networks (Borenstein, 1989) and economies on graphs (Jackson, 2008).

was not to give firms, *per se*, the strategic flexibility that is well-known to eliminate the predictive power of theory. Rather the goal was to capture the ways in which platform strategies act to coordinate consumers. Not all conditional pricing schemes achieve this goal; in fact, in the RT2003 context it is very specifically proportional pricing that does it, while in the Armstrong context it is determined by the shape of consumer's homogeneous benefits. What these share is that they both eliminate the scope for coordination failures by giving consumers dominant strategies. In Section 5 we maintain that platforms will generally use generalizations of these strategies, and of W10's notion of Insulating Tariffs, to avoid mis-coordination. This eliminates both the scope for coordination among consumers and ties down tariffs to resolve Armstrong's Paradox.

The Model's Scope

Regarding the ways in which our model generalizes with respect to existing literature, 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 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.⁴

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

⁴In 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.

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.⁵

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.

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; Jeon and 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.⁶

⁵The 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).

⁶We are optimistic that such an extension can be incorporated without great difficulty into our framework, but given the early stage of this research on heterogeneous externalities, we do not include it in the current version of this paper.

Other issues that we do not consider include explicit 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.

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} . 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 $\mathcal{M}_i^{\mathcal{I}} \in \wp(\mathcal{M})$.

To capture consumption 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.⁷ 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, $N \equiv (N^{\mathcal{A}}, N^{\mathcal{B}}) \in [0, 1]^{2m}$, specifies the total fraction 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 side \mathcal{I} from joining bundle of platforms \mathcal{X} as

$$v^{\mathcal{I}}(\mathcal{X}, N^{\mathcal{J}}, \theta_i^{\mathcal{I}}) - \sum_{j \in \mathcal{X}} P^{\mathcal{I},j},$$

where $\theta_i^{\mathcal{I}} \in \Theta^{\mathcal{I}}$ denotes consumer i on side \mathcal{I} ’s ‘type’ or ‘characteristics’. 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.⁸ 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

⁷As we discuss in Section 2.2, we do not explicitly model the ‘interaction’ that may take place among users on opposite sides.

⁸For example, the natural extension of the primary preferences assumed by Rochet and Tirole (2006) 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.

\mathcal{I} must pay to join platform j , the details of which we discuss below, when defining platforms' strategies. Let $f^{\mathcal{I}} : \Theta^{\mathcal{I}} \rightarrow \mathbb{R}$ denote the probability density function of user types on side $\mathcal{I} = \mathcal{A}, \mathcal{B}$, satisfying $\int_{\Theta^{\mathcal{I}}} f^{\mathcal{I}}(\theta) d\theta = 1$.

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. Full Support: *For all $N^{\mathcal{J}} \in [0, 1]$ and utility profiles $\mathbf{u}^{\mathcal{I}} \in \mathbb{R}^{2^m-1}$, there exists $\theta \in \Theta^{\mathcal{I}}$ such that $v^{\mathcal{I}}(\cdot, N^{\mathcal{J}}, \theta)$ takes the value $\mathbf{u}^{\mathcal{I}}$.*
2. Smoothness: *For all $N^{\mathcal{J}} \in [0, 1]$, the corresponding (2^m-1) -dimensional distribution of gross utility profiles is twice continuously differentiable. Moreover, $v^{\mathcal{I}}(\cdot, \cdot, \cdot)$ is twice continuously differentiable in all dimensions of its second argument.*
3. Normalization: *For all $\theta \in \Theta^{\mathcal{I}}$, $v^{\mathcal{I}}(\emptyset, N^{\mathcal{J}}, \theta) = 0$. For all consumers the 'outside option' gives a payoff normalized to zero.*

Before further specifying the strategic aspects of the game, we turn off the 'two-sided' feature of the model in order to establish useful properties that such a demand system exhibits in a conventional one-sided setting.

Lemma 1 (Invertibility). *Hold fixed the coarse allocation on the opposite side of the market, $N^{\mathcal{J}}$, and consider the demand system on side \mathcal{I} . For any interior side \mathcal{I} coarse allocation, $N^{\mathcal{I}} \in (0, 1)^m$, there exists a unique vector, $\mathbf{P}^{\mathcal{I}} \in \mathbb{R}^m$, of platforms' total prices, that supports this allocation.*

Proof. Holding fixed $N^{\mathcal{J}}$, the demand system in this model is a special case of the demand system in Azevedo, Weyl, and White's (2011) exchange economy with a continuum of agents. Therefore, Theorem 1 of that paper implies this result. \square

Lemma 1 ensures that demand on side \mathcal{I} , $N^{\mathcal{I}}(\mathbf{P}^{\mathcal{I}}, N^{\mathcal{J}})$, can be inverted with respect to the vector of total prices on side \mathcal{I} and that this inverse demand function, $\mathbf{P}^{\mathcal{I}}(N^{\mathcal{I}}, N^{\mathcal{J}})$, is well-defined over the domain $(0, 1)^m \times [0, 1]^m$. Also note that our assumptions guarantee that the functions $N^{\mathcal{I}}$ and $\mathbf{P}^{\mathcal{I}}$ are twice continuously differentiable with respect to all of their arguments.

Lemma 2 (Consumers' Marginal Value of Supply). *Let $V^{\mathcal{I}}(N^{\mathcal{I}}, N^{\mathcal{J}})$ denote the gross utilitarian surplus to consumers on side \mathcal{I} associated with coarse allocation $(N^{\mathcal{I}}, N^{\mathcal{J}})$. It holds that*

$$\frac{\partial V^{\mathcal{I}}}{\partial N^{\mathcal{I},j}} = P^{\mathcal{I},j}. \quad (2)$$

Moreover, $V^{\mathcal{I}}(N^{\mathcal{I}}, N^{\mathcal{J}})$ is concave in all elements of $N^{\mathcal{I}}$.

Proof. This follows from Proposition 1 of Azevedo et al. *op. cit.* \square

Lemma 2 formalizes, in a setting with rich complementary preferences, the familiar intuition that a good's marginal consumers have valuations for it that are equal to its price.

Supply

Turning to platforms, 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 twice continuously differentiable in both arguments.

Equilibrium

Platforms move first, simultaneously. Then, having observed the platforms' moves, all consumers simultaneously choose which platforms to join. We allow for each platform to charge a tariff to consumers on side \mathcal{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. Formally, $\sigma^{\mathcal{I},j} : [0, 1]^m \rightarrow \mathbb{R}$. In order to ensure differentiability of each platform's residual profits, we assume that all platforms' price functions, $\sigma^{\mathcal{I},j}$, are twice continuously differentiable. While this assumption 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*,⁹ 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 $\mathcal{M}^{\mathcal{I}}(\theta^{\mathcal{I}}, [\sigma])$. With our atomless continuum of consumers, there will always be sets of consumers of measure zero who are indifferent between bundles. For definiteness, we tie down the behavior of such agents by assuming that consumers' strategy profiles are symmetric – in every subgame, each consumer takes some action with probability one – and pure – all agents sharing a common type adopt the same strategy. It thus follows 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 $N : \{\widehat{\mathcal{M}}\} \times \Sigma^m \rightarrow [0, 1]^{2m}$,

⁹By '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 .

mapping from marketwide consumer strategy profile and platform strategy profile to coarse allocation. $N[\widehat{\mathcal{M}}, \sigma]$ has generic elements

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

Finally, we denote by $\mathcal{M}^{\mathcal{I}*}(\theta^{\mathcal{I}}, N^{\mathcal{J}}, [\sigma])$ a *best response* strategy profile for consumers on side \mathcal{I} , where $\mathcal{M}^{\mathcal{I}*} : \Theta^{\mathcal{I}} \times [0, 1]^m \times \Sigma^m \rightarrow \wp(\mathcal{M})$ specifies an optimal bundle for every consumer on side \mathcal{I} , given a coarse allocation on side \mathcal{J} and platforms' strategy profile.

4 The Allocation Approach

In this section, we establish a simple framing of platforms' profit maximization problem, in the spirit of Myerson (1981) and Riley and Samuelson's (1981) approach to solving for optimal auctions based on allocation and implied revenue. First, we explain the motivation for such an approach in our oligopoly setting with consumption externalities. We then define a *No Sunspots* property of consumers' strategies', which, if satisfied, ensures the validity of this approach.

In the standard analysis of Nash-in-prices equilibrium of a differentiated products industry, each firm takes as given other firms' prices and chooses the price(s) for its own good(s) that maximizes profits. The presence of consumption externalities complicates matters. Consider the *Consumer Game* that takes place in the second stage, after platforms have announced their strategies. As we discuss in Section 2.2, such games can have multiple Nash Equilibria, since the optimal bundle for consumers on one side of the market depends on the actions of consumers




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

on the other side. (See Figure 1.) Thus, a platform's optimal strategy may depend sensitively on the particular way it expects consumers to coordinate. This potential for multiple equilibria among consumers makes each platform's optimization problem *over prices* potentially extremely complex.

Rather than considering the large space of potential price functions, the platform (and we, as modelers) can focus on the much smaller space of residual allocations. In order for this framing of the problem to be fully useful, it must be the case that, given the strategies of the other platforms,

there is a function that maps from the number of consumers that platform j serves on each side of the market, $(N^{\mathcal{A},j}, N^{\mathcal{B},j})$, to the profit that j realizes. Such a mapping exists, so long as consumers' strategy exhibits the No Sunspots property, the term for which we borrow from Cass and Shell (1983) and which we now define.

Definition 2 (No Sunspots). *Given the strategy profile of the other platforms, σ^{-j} , a consumer strategy profile, $\widehat{\mathcal{M}}$, perceives No Sunspots from platform j if, for all $\sigma^j, \widehat{\sigma}^j$,*

$$N^j[\widehat{\mathcal{M}}, (\sigma^j, \sigma^{-j})] = N^j[\widehat{\mathcal{M}}, (\widehat{\sigma}^j, \sigma^{-j})] \text{ implies } N[\widehat{\mathcal{M}}, (\sigma^j, \sigma^{-j})] = N[\widehat{\mathcal{M}}, (\widehat{\sigma}^j, \sigma^{-j})].$$

To appreciate the meaning of this condition, suppose that it is violated. Then, platform j could, through changes in its policies, affect *other platforms'* allocations without affecting its own. In the context of our model, we believe this condition to be very weak, because it seems highly implausible that the prices of one platform would serve purely as a coordination device mediating the decisions of consumers to join *other* platforms, purely by influencing their beliefs. We thus posit that this condition is fulfilled in Assumption 2.

Assumption 2. *We restrict attention to consumer strategy profiles that perceive No Sunspots from any platform.*

Under Assumption 2, it is possible to assign a unique profit level, for platform j , to every (interior) level of demand that it serves on each side of the market. We state this formally in Lemma 3

Lemma 3. *There exists a profit function for platform j , $\Pi^j(N^{\mathcal{A},j}, N^{\mathcal{B},j}, \sigma^{-j})$, that is well-defined over the domain $(0, 1)^2 \times \Sigma^{m-1}$.*

Proof. When No Sunspots holds for platform j , each value of the coarse allocation for platform j , $N^j = (N^{\mathcal{A},j}, N^{\mathcal{B},j})$, implies a unique value for the entire coarse allocation, $N = (N^{\mathcal{A}}, N^{\mathcal{B}})$. By Lemma 1, it holds that P^I is a function of $N^{\mathcal{A}}$ and $N^{\mathcal{B}}$, for $I = \mathcal{A}, \mathcal{B}$. Therefore, platform j 's prices, $(p^{\mathcal{A},j}, p^{\mathcal{B},j})$ are uniquely determined by $N^{\mathcal{A},j}, N^{\mathcal{B},j}$ and σ^{-j} . \square

Lemma 3 says that, in the absence of sunspots, platform j has a coherent residual inverse demand system. Thus, given the strategies of other platforms, each can identify a *best response allocation*. As in any optimization problem in a rich environment, a best-response is generically unique. Once this best response allocation is identified, the platform must determine how to implement this allocation using its available strategies. This question forms the basis of the following section.

5 Robust Implementation and Insulation

Section 4 casts platforms' problem as the choice of a best response allocation. Here, we address the question of how each platform induces or 'implements' this. We identify strategies called *Residual Insulating Tariffs* which, in a sense to be precisely specified, most robustly do so. We then posit

our crucial assumption – that platforms use RITs – forming the basis for our solution concept of *Insulated Equilibrium*, which generates testable predictions of prices based on cost and demand primitives.

Suppose that, given the strategy profile of other firms, N^{j*} uniquely maximizes platform j 's profits. Then, in order to be a best response to σ^{-j} , platform j 's strategy, σ^j , must 'weakly implement' N^{j*} . More formally, for σ^j to be a best response to σ^{-j} , there must be a Nash Equilibrium in the Consumer Game defined by (σ^j, σ^{-j}) that features N^{j*} .

There are, however, *infinitely many* strategies for platform j that fulfill this criterion. To see this, let N^* denote the entire coarse allocation corresponding to N^{j*} . It is straightforward to see that σ^j is a best response to σ^{-j} if and only if $\sigma^j(N^*) = (P^{A,j}(N^{B*}), P^{B,j}(N^{A*}))$. In other words, the value of $\sigma^j(\cdot)$ is tied down by the demand system *only when evaluated at coarse allocation* N^* . As Figure 2 illustrates, the value $\sigma^j(\cdot)$ takes on at other allocations is unconstrained, and thus the slope j 's tariff with respect to the allocation on the other side of the market is unconstrained.




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

As pointed out by Armstrong's (2006) Proposition 3, this leads to a great multiplicity of *Platform Equilibria*, holding fixed consumers' strategy profile. This issue, which we refer to as *Armstrong's Paradox*, mirrors the multiplicity of supply function equilibria in a deterministic setting, analyzed by Klemperer and Meyer (1989). 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 σ^I and σ^J are free, where, recall, m denotes the number of platforms. In order to satisfy the first-order equilibrium conditions at any point these need only satisfy $2m \leq 2m^2$ 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.¹⁰ 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-

¹⁰We are working on a formal proof of this conjecture; the difficulty is showing that global optimality and stability conditions are also satisfied.

order conditions, as is commonly done in standard one-sided markets (following 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.

Our solution to Armstrong's paradox relies on the following observation: the *entire motivation* for introducing allocation-dependent tariffs into models with consumption externalities was to allow firms to solve coordination problems among consumers. After all, despite the analysis of Klemperer and Meyer, most static empirical analysis in industrial organization uses the Nash-in-prices solution concept. Thus the motivation for allowing richer tariffs was not simply to expand the generality of the analysis, but rather, it was to allow firms to condition their prices in order to prevent the mis-coordination among consumers discussed in Section 4. As we now discuss, not all tariffs do this equally well, providing a rationale for our assumption that platforms choose tariffs that minimize the scope for such mis-coordination.

5.1 Mis-coordination and robust implementation.

Many tariffs available to platform j implement N^{j*} as a unique Nash Equilibrium. However, the platform might wish to achieve something stronger than this form of uniqueness. At least since the work of Wilson (1987), a large literature in mechanism design has stressed the importance of implementing desired outcomes in ways robust to players' higher-order beliefs (Bergemann and Morris, 2005; Chung and Ely, 2007). In a quasi-linear, private values environment like the one considered here, doing so would require platforms to make it a dominant strategy for each consumer either to join or not join. Such robust implementation of its best response allocation seems a compelling strategy for a platform to adopt.

In fact this is precisely the technique that Rochet and Tirole (2003) (RT2003) use to solve their model of competition and that Armstrong (2006) uses to solve his general model of monopoly. In RT2003, the joint assumption that both consumers' utility and platforms' prices are proportional to the number of users on the other side of the market implies that each user has a dominant strategy. In Armstrong's monopoly model, the assumption that consumers are heterogeneous only in the value of their outside option plays the same role in this regard: following Dybvig and Spatt (1983) and Armstrong and Vickers (2001), he argues that platforms can commit to offer consumers a particular level of utility by adjusting their price in response to changes in the number of users on the other side, fully insuring users against any change in their utility, thus offering them a dominant strategy.

However, note that both of these modeling techniques rely on a special assumption about consumer heterogeneity, namely that it is one-dimensional. When this is not the case, marginal consumers are heterogeneous and thus it is not possible, using a uniform price, to give all consumers a dominant strategy. We now illustrate this point with a simple example.

A monopolist platform faces consumers on side \mathcal{A} that have heterogeneous 'participation benefits', $B_i^{\mathcal{A}}$, and heterogeneous 'interaction benefits', $b_i^{\mathcal{A}}$, per consumer on side \mathcal{B} , thus deriving

gross utility from joining the platform given by $B_i^{\mathcal{A}} + b_i^{\mathcal{A}} N^{\mathcal{B}}$. Suppose the platform wishes to implement $N^{\mathcal{A}} = N^{\mathcal{B}} = \frac{1}{2}$. In the absence of price discrimination, there must be some threshold, $P^{\mathcal{A}}$, such that all consumers i with $B_i^{\mathcal{A}} + \frac{1}{2} b_i^{\mathcal{A}} > P^{\mathcal{A}}$ participate. For concreteness, suppose the threshold above which half of the population satisfies this is $P^{\mathcal{A}} = 5$. Then, any implementation in dominant strategies $\sigma^{\mathcal{A}}(N^{\mathcal{B}})$ must to provide an incentive for all users with $B_i^{\mathcal{A}} + \frac{1}{2} b_i^{\mathcal{A}} > 5$ to participate for every $N^{\mathcal{B}} \in (0, 1)$ and those with $B_i^{\mathcal{A}} + \frac{1}{2} b_i^{\mathcal{A}} < 5$ a similar incentive not to participate. In particular a $(B^{\mathcal{A}}, b^{\mathcal{A}}) = (5.1, 0)$ consumer would have to have a dominant strategy to participate, while a $(4.4, 1)$ consumer would need to have a dominant strategy not to participate. Among other things, this would require that

$$5.1 > \sigma^{\mathcal{A}}(.9) > 4.4 + .9 \cdot 1 = 5.3,$$

an obvious contradiction.

5.2 Residual Insulating Tariffs

Given that dominant strategy implementation is infeasible with rich heterogeneity, one may consider a weaker notion of robust implementation that encompasses implementation in dominant strategies whenever it is feasible. To this end, we now define *Residual Insulating Tariffs* (RITs).

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, for all $N^{\mathcal{J}}, \widetilde{N^{\mathcal{J}}} \in [0, 1]$,

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

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 its 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

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

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, $\sigma^{\mathcal{I},-j}(\cdot)$, in order to ensure that the output of $N^{\mathcal{I},j}$ is constant. Note that, whenever dominant strategy implementation is feasible, it is exactly the insulating tariff: any tariff that leads all consumer's choices to be independent of others' choices will, *a fortiori*, lead the aggregate quantity to be independent of the allocation on the other side. Thus, the RT2003 and Armstrong dominant strategy tariffs are the special cases of the insulating tariff, under the preference structures assumed in those settings. More generally, as Lemma 4 establishes, a unique insulating tariff exists, given the strategies of other platforms.

Lemma 4 (Existence and Uniqueness of a Residual Insulating Tariff). *There exists a unique function, $\overline{P^{I,j}}(N^{\mathcal{J}}; \tilde{N}, [\sigma^{I,-j}(N^{\mathcal{J}})])$, such that, $\forall N^{\mathcal{J}}, \forall \tilde{N} \in (0, 1), \forall \sigma^{I,-j}$,*

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

where \tilde{N} denotes the ‘anchor allocation’ that $\overline{P^{I,j}}$ implements. Moreover, $\overline{P^{I,j}}$ is C^2 in all dimensions of its first argument.

Proof. For existence, note that (i) $N^{I,j}(\cdot, \cdot, \cdot)$ is continuous in its first argument, since it is the integral of a smooth set, and (ii) $\forall N^{\mathcal{J}}, \forall \sigma^{I,-j}, \lim_{P^{I,j} \rightarrow -\infty} N^{I,j}(P^{I,j}, \sigma^{I,-j}, N^{\mathcal{J}}) = 1$ (and $\lim_{P^{I,j} \rightarrow \infty} N^{I,j}(P^{I,j}, \sigma^{I,-j}, N^{\mathcal{J}}) = 0$), since $\forall \theta^I, \forall N^{\mathcal{J}}, \forall \sigma^{I,-j}, \exists P^{I,j}$ such that

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

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 $N^{\mathcal{J}}, N^{\mathcal{J},k}$, in order to be insulating $\overline{P^{I,j}}$ must be differentiable by the inverse function theorem and have derivative equal to

$$\frac{\sum_{1 \neq j} \frac{\partial N^{I,j}}{\partial P^{I,1}} \frac{\partial \sigma^{I,1}}{\partial N^{\mathcal{J},k}} + \frac{\partial N^{I,j}}{\partial N^{\mathcal{J},k}}}{-\frac{\partial N^{I,j}}{\partial P^{I,j}}}$$

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

Another perspective from which to view RITs is by comparison to Armstrong and Vickers (2001)’s notion of competition in utility space. When consumers are homogeneous, it is feasible, with a uniform price, for a platform to make their utility independent of the number of users on the other side of the market, as in Armstrong’s monopoly model. Even in the case when consumers are heterogeneous, but the *marginal* consumers are homogenous, the marginal users may be perfect compensated for changes in the allocation on the other side of the market, as in RT2003. On the other hand, when even the marginal consumers are heterogeneous, the most a firm can hope to accomplish is to compensate the *average marginal* consumer for changes in the allocation on the other side. This is equivalent to making the platform’s desired allocation the dominant strategy for the *Representative Consumer* (RC) on a given side of the market, who aggregates together the

where ‘ \cdot ’ denotes the inner-product operator and where, as defined in Lemma 2, V^I denotes gross consumer surplus on side 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 I , the RC chooses coarse allocation N_{RC}^I ; activity among side J consumers occurs as before. We can now state Theorem 1.

$$\begin{aligned} V^I(\overline{N^I}, N^J) - \sigma^I(N^J) \cdot \overline{N^I} &> V^I(N^I, N^J) - \sigma^I(N^J) \cdot N^I, \\ &\Leftrightarrow \sigma^I = \left(\overline{P^{I,j}}(\cdot; N^{I,j*}, [\sigma^{I,-j}]), \sigma^{I,-j} \right). \end{aligned}$$

F95 138 7.9701 Tf 5.451 0-4.505 Td [(;)]TJ 112 10.9091 Tf 4.546 0 Td [(N)]TJ F143 7.9701 Tf 9.043 4.505 Td [(I

$$\begin{aligned} V^I(N_{RC}^I, N^{\mathcal{I}}) - \sigma^I(N^{\mathcal{I}}) \cdot N_{RC}^I &= \\ \max_{\mathcal{M}^I \in \{\mathcal{M}^I: N^I[\mathcal{M}^I, \sigma] = N_{RC}^I\}} &\sum_{\mathcal{X} \in \wp(\mathcal{M})} \int_{\theta^I: \mathcal{M}^I(\theta^I, [\sigma]) = \mathcal{X}} \left(v^I(\mathcal{X}, N^{\mathcal{I}}, \theta) - \widehat{P_{I, \mathcal{X}}}(N^{\mathcal{I}}) \right) f(\theta) d\theta \\ &\leq \sum_{\mathcal{X} \in \wp(\mathcal{M})} \int_{\theta^I: \mathcal{M}^{I*}(\theta^I, N^{\mathcal{I}}; [\sigma]) = \mathcal{X}} \left(v^I(\mathcal{X}, N^{\mathcal{I}}, \theta) - \widehat{P_{I, \mathcal{X}}}(N^{\mathcal{I}}) \right) f(\theta) d\theta. \quad (3) \end{aligned}$$

$$\left\{ \left(\Sigma'' \left(\Gamma' \right)^T \left[\left(\Gamma' \right)^T \left(\Gamma' \right) \left(\Sigma \right) \right] \right) \right\}$$

5.3 Key Assumption: Insulation

The crucial substantive assumption from which we derive our solution concept posits that, in order to minimize the scope for mis-coordination, all platforms use Residual Insulating Tariffs. Because there always exists a (unique) RIT anchored at a platform's best response allocation, this merely refines platforms choice space without restricting it; that is, platforms can never do better by unilaterally deviating to *any* other tariff.

Assumption 3 (Insulation). *All platforms employ Residual Insulating Tariffs.*

While it is common for advertising rates to depend on the number of consumers viewing those advertisements and even the advertisements of rivals, static prices in platform markets almost never *literally* resemble insulating tariffs. On the other hand, a qualitatively obvious and seemingly very common practice of platform businesses is to set prices low when the platform's (endogenous) quality is low to have them rise, over time as quality increases, in order to generate profits. Our insistence that platforms will, as far as possible, robustly implement their best response allocation is a static reduced form for this inherently dynamic process: the platform aims, relentlessly and consistently, at implementing its best response allocation, regardless of the transitory market conditions that may impede this.

5.4 Insulated Equilibrium

We now introduce vocabulary to describe the world implied by Assumption 3, in which all platforms charge insulating tariffs.

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

Note that, at any anchor allocation, the ITS exists and is unique directly from Lemma 1: it is exactly the unique set of price consistent with the anchor allocation, given the coarse allocation on the opposite side 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

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

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, on each side of the market, \mathcal{M}^I is a best response strategy profile for all $\sigma \in \Sigma^m$.

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 $\sigma^* = \bar{P}(N^*)$.

Remark. By construction, IEs are the only SPEs consistent with Insulation.

In a Subgame Perfect Equilibrium, platforms select their strategies as if they had complete certainty of the outcome of the ensuing Consumer Game, even when the particular Consumer Game that they induce has multiple 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 Nash Equilibrium.

5.5 Insulated Equilibrium is Identified

Because the ITS anchored at any reference allocation is unique, it defines a unique inverse demand system for each platform at each candidate equilibrium, thus tying down each platform's incentives. Intuitively, as shown in Figure 3, assumption 3 ties down each platform's tariff, resolving Armstrong's Paradox. Thus if both costs and demand are known, any candidate allocation either is, or is not, an IE, without any further specification of strategies. Similarly, if costs are unknown, marginal costs may now be recovered à la Rosse. Thus the predictive power of IE is equivalent to that of Nash-in-prices equilibrium in a standard industry without consumption externalities. We formalize this statement in Theorem 2.

Theorem 2 (Under Insulated Equilibrium, Marginal Cost is Identified). Suppose $\{\widehat{\mathcal{M}}^*, \sigma^*\}$ is an IE with coarse allocation N^* , with generic elements $N^{I,j*}$. Then, the vector of platform marginal costs is identified jointly by the vector of prices, $\{P^I\}_{I=\mathcal{A},\mathcal{B}}$, 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 N^* , the equilibrium profile of platform strategies is $\sigma^* = \bar{P}(N^*)$. Thus, platform j 's profit maximization problem can be written

$$\max_{\{N^{\mathcal{A},j}, N^{\mathcal{B},j}\}} \sum_{I=\mathcal{A},\mathcal{B}} N^{I,j} \cdot P^{I,j}(N^{I,j}, N^{\mathcal{J},j}) - C^j(N^{\mathcal{A},j}, N^{\mathcal{B},j}), \quad (5)$$

where

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




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

and

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

The values that maximize (5), $N^{\mathcal{A},j*}$ and $N^{\mathcal{B},j*}$, satisfy first-order condition

$$p^{\mathcal{I},j} + N^{\mathcal{I},j*} \frac{\partial p^{\mathcal{I},j}}{\partial N^{\mathcal{I},j}} + N^{\mathcal{J},j*} \frac{\partial p^{\mathcal{J},j}}{\partial N^{\mathcal{I},j}} = \overline{p^{\mathcal{I},j}} + \frac{N^{\mathcal{I},j*}}{\frac{\partial N^{\mathcal{I},j}}{\partial p^{\mathcal{I},j}}} + N^{\mathcal{J},j*} \frac{\partial \overline{p^{\mathcal{J},j}}}{\partial N^{\mathcal{I}}} \cdot \frac{\frac{\partial N^{\mathcal{I}}}{\partial p^{\mathcal{I},j}}}{\frac{\partial N^{\mathcal{I},j}}{\partial p^{\mathcal{I},j}}} = \frac{\partial C^j}{\partial N^{\mathcal{I},j}}. \quad (6)$$

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

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 (10). 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.

In order to afford a clean interpretation of the pricing formulae, we impose the following No

Externalities to Outsiders assumption. 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.¹² For instance, it implies that the gross utility a user derives from reading *only* one newspaper does not depend on the number of advertisements in other newspapers.

Assumption 4 (No Externalities to Outsiders). *If $j \notin \mathcal{X}$ then $v^I(\mathcal{X}, N^{\mathcal{J}}, \theta)$ is independent of $N^{\mathcal{J},j}$.*

6.1 General Pricing

Socially Optimal

The utilitarian social welfare corresponding to an allocation is equal to

$$\sum_{I=\mathcal{A},\mathcal{B}} v^I(N^I, N^{\mathcal{J}}) - \sum_{j \in \mathcal{M}} C^j(N^{\mathcal{A},j}, N^{\mathcal{B},j}), \quad (7)$$

where $v^I(N^I, N^{\mathcal{J}})$ denotes gross consumer surplus on side I , as defined in Lemma (2). In Proposition 1 we state the pricing formula for maximizing this quantity. To do so, let us denote by

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

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*

$$p^{\mathcal{I},j} = C_I^j - N^{\mathcal{J},j} \overline{v_j^{\mathcal{J},j}}. \quad (8)$$

Proof. By Lemma 2, we have $\frac{\partial v^I}{\partial N^{\mathcal{I},j}} = p^{\mathcal{I},j}$. The other terms in (8) are straightforward. \square

Equation (8) 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 (8) as

$$\frac{C_I^j - p^{\mathcal{I},j}}{N^{\mathcal{J},j}} = \overline{v_j^{\mathcal{J},j}}. \quad (9)$$

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

¹²In this regard, the scope of these results differs from that of Segal (1999), which studies the effects of positive versus negative ‘externalities on nontraders’.

of all its consumers for such an improvement. Equation (9) says precisely this, since it prescribes that the net social marginal cost, per side \mathcal{J} consumer, of providing such a quality increase¹³ 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*

$$P^{\mathcal{I},j} = C_{\mathcal{I}}^j + \mu^{\mathcal{I},j} - N^{\mathcal{J},j} \left(\left[-\frac{\partial N^{\mathcal{J}}}{\partial P^{\mathcal{J}}} \right]^{-1} \left[\frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}} \right] \right)_{j,\cdot} \cdot [-D^{\mathcal{I}}_{\cdot,j}]. \quad (10)$$

Proof. In view of equation (6), it remains for us to show that $\frac{\partial \overline{P^{\mathcal{J},j}}}{\partial N^{\mathcal{I}}} = \left(\left[-\frac{\partial N^{\mathcal{J}}}{\partial P^{\mathcal{J}}} \right]^{-1} \left[\frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}} \right] \right)_{j,\cdot}$. 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

$$\underbrace{\mathbf{0}}_{m \times m} = \left[\frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}} \right] + \left[\frac{\partial N^{\mathcal{J}}}{\partial P^{\mathcal{J}}} \right] \left[\frac{\partial \overline{P^{\mathcal{J}}}}{\partial N^{\mathcal{I}}} \right] \Leftrightarrow \left[\frac{\partial \overline{P^{\mathcal{J}}}}{\partial N^{\mathcal{I}}} \right] = \left[-\frac{\partial N^{\mathcal{J}}}{\partial P^{\mathcal{J}}} \right]^{-1} \left[\frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}} \right], \quad (11)$$

and we have our result. \square

Two forces govern the relationship between the optimal pricing formula of equation (9) and the equilibrium pricing rule of equation (10). The first is the classical market power distortion, captured by $\mu^{l,j} \equiv -N^{\mathcal{I},j} / \left(-\frac{\partial N^{\mathcal{I},j}}{\partial P^{\mathcal{I},j}} \right)$. 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.

¹³For 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.

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.¹⁴ 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’ 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

¹⁴Recall that we assume Bertrand conduct among firms within each side.

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 (11), and the underlying demand system. Two sorts of quantities comprise the matrix $\left[-\frac{\partial N^I}{\partial P^I}\right]^{-1} \left[\frac{\partial N^I}{\partial N^J}\right]$. 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 $\widetilde{\Theta}_j^I \equiv \left\{ \theta^I \in \Theta^I : \exists \mathcal{X}, \mathcal{Y} \in \arg \max_{\mathcal{Z} \in \varphi(\mathcal{M})} \left\{ v^I(\mathcal{Z}, N^J, \theta^I) - \widehat{P}^{I, \mathcal{Z}} \right\} \text{ s.t. } j \in \mathcal{X}, j \notin \mathcal{Y} \right\}$ denote the entire set of consumers on side 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 $\widetilde{\Theta}_{j,k}^I \equiv \left\{ \theta^I \in \Theta^I : \exists \mathcal{X}, \mathcal{Y} \in \arg \max_{\mathcal{Z} \in \varphi(\mathcal{M})} \left\{ v^I(\mathcal{Z}, N^J, \theta^I) - \widehat{P}^{I, \mathcal{Z}} \right\} \text{ s.t. } j \in \mathcal{X}, j \notin \mathcal{Y}, k \in \mathcal{Y}, k \notin \mathcal{X} \right\}$ 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

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

The elements of the interaction matrix are thus

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

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 $\frac{\partial v^{\mathcal{I}}(\cdot, N^{\mathcal{J}}, \theta^{\mathcal{I}})}{\partial N^{\mathcal{J},k}}$ 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 $\widetilde{\Theta}_j^{\mathcal{I}}$, 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 $\widetilde{\Theta}_{j,k}^{\mathcal{I}}$, 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*

$$p^{\mathcal{I},j} = C_I^j + \mu_I^j - N^{\mathcal{J},j} v_j^{\mathcal{I},\mathcal{X}} \left[\widetilde{\Theta}_j^{\mathcal{I}} \right].$$

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

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

because the j^{th} entry of $[-\mathbf{D}_{\cdot,j}^{\mathcal{I}}]$ 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.¹⁵ 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 (10), of platform j 's price on side \mathcal{I} becomes

$$P^{\mathcal{I},j} = C_I^j + \mu_I^j - N^{\mathcal{J},j} \frac{\overline{\partial P^{\mathcal{J},j}}}{\partial N^{\mathcal{I},j}}. \quad (12)$$

In equation (12), 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

¹⁵We thank Julian Wright for drawing our attention to this point.

expansion margin" between each firm and \emptyset . This factor then simplifies to

$$\frac{\partial \overline{P^{\mathcal{J},j}}}{\partial N^{\mathcal{I},j}} = \left((1 - \kappa) \cdot v_j^{\mathcal{J},\mathcal{X}} \left[\overline{\Theta^{\mathcal{J}}_{j,\emptyset}} \right] + \kappa \cdot v_j^{\mathcal{J},\mathcal{X}} \left[\overline{\Theta^{\mathcal{J}}_{j,k}} \right] \right), \quad (13)$$

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 benefits 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

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




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 .

to consumer i on side \mathcal{J} , when he joins the set containing only platform j , where $\epsilon_i^{\mathcal{J}}$ 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 $\epsilon^{\mathcal{J}}$ 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, $\epsilon^{\mathcal{J}}$), 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 (13), 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 $\widetilde{b}_{j,k}^{\mathcal{J}}$.

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




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.

of the model and examine the general case of competition between two platforms. We have

$$P^{I,j} = C_I^j + \mu^{I,j} - N^{\mathcal{J},j} \left((1-\kappa) \cdot v_j^{\mathcal{J},\mathcal{I}} [\widetilde{\Theta}_{j,\emptyset}] + \kappa \left(v_j^{\mathcal{J},\mathcal{I}} [\widetilde{\Theta}_{j,k}] + \overbrace{D_{k,j}^I (v_k^{\mathcal{J},\mathcal{I}} [\widetilde{\Theta}_{k,j}] - v_k^{\mathcal{J},\mathcal{I}} [\widetilde{\Theta}_{k,\emptyset}])}^{\varphi} \right) \right). \quad (14)$$

Note, first, the term in (14), indicated by φ , that does not enter the prior first-order condition, (13). 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 $v_k^{\mathcal{J},\mathcal{I}} [\widetilde{\Theta}_{k,j}] - v_k^{\mathcal{J},\mathcal{I}} [\widetilde{\Theta}_{k,\emptyset}]$, 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}^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

$$P^{\mathcal{I},j} = C_I^j + \mu^{\mathcal{I},j} - N^{\mathcal{J},j} \left((1 - \kappa^{\text{sym}}) v_\emptyset^{\mathcal{J}} + \kappa^{\text{sym}} v_k^{\mathcal{J}} \right), \quad (15)$$

where

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

Expression (15) reinforces the themes we discuss in the previous examples. As in expression (13), 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 10, 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 \mathbf{t} . The necessary condition for Insulated Equilibrium can be written $\psi = \mathbf{t}$. Implicitly differentiating this system of first-order conditions with respect to each of the taxes then yields

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

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

$$\frac{\partial \mathbf{P}}{\partial \mathbf{t}} = [\mathfrak{S}]^{-1} [\mathbf{T}_U]. \tag{16}$$

Equation 16 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:

$$\begin{bmatrix} \left(\overline{\mathbf{V}}_{\mathcal{B}}^{\mathcal{A}} - \frac{\partial \overline{\mathbf{P}}^{\mathcal{A}}}{\partial \mathbf{N}^{\mathcal{B}}} \right) \frac{\partial \mathbf{N}^{\mathcal{B}}}{\partial \mathbf{P}^{\mathcal{B}}} \Delta \mathbf{P}^{\mathcal{B}} - \Delta \mathbf{P}^{\mathcal{A}} \\ \left(\overline{\mathbf{V}}_{\mathcal{A}}^{\mathcal{B}} - \frac{\partial \overline{\mathbf{P}}^{\mathcal{B}}}{\partial \mathbf{N}^{\mathcal{A}}} \right) \frac{\partial \mathbf{N}^{\mathcal{A}}}{\partial \mathbf{P}^{\mathcal{A}}} \Delta \mathbf{P}^{\mathcal{A}} - \Delta \mathbf{P}^{\mathcal{B}} \end{bmatrix}^T \cdot \begin{bmatrix} \mathbf{N}^{\mathcal{A}} \\ \mathbf{N}^{\mathcal{B}} \end{bmatrix}, \quad (17)$$

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 $\overline{v}_{\mathcal{I},j}^{\mathcal{J}}$, 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 $\frac{\partial \overline{\mathbf{P}}^{\mathcal{I}}}{\partial \mathbf{N}^{\mathcal{J}}}$ 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.¹⁶ 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

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

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

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

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 (17), 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

¹⁶Note that because we assume Bertrand conduct, we need consider only UPP and not JW's generalization of it, GePP, which allows more general conduct.

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 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 θ^I on side \mathcal{I} be $v^I(\mathcal{X}, N^{-I}, \theta^I)$, where $v^I : \wp(\mathcal{M}) \times [0, 1]^{m(S-1)} \times \Theta^I \rightarrow \mathbb{R}$ now depends on $N^{-I} \in [0, 1]^{m(S-1)}$, 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 $N^{-I} \in [0, 1]^{m(S-1)}$ 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 ??, 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

$$\max_{\{N^{\mathcal{A},j}, N^{\mathcal{B},j}, N^{\mathcal{C},j}, \dots\}} \sum_{\mathcal{I}=\mathcal{A}, \mathcal{B}, \mathcal{C}, \dots} N^{\mathcal{I},j} P^{\mathcal{I},j}(N^{\mathcal{I},j}, N^{-I}) - C^j(N^{\mathcal{A},j}, N^{\mathcal{B},j}, N^{\mathcal{C},j}, \dots). \quad (18)$$

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

$$P^{\mathcal{I},j} = C^j_I - \sum_{\mathcal{J} \neq \mathcal{I}} N^{\mathcal{J},j} \overline{v^{\mathcal{J},j}}, \quad (19)$$

and the platforms' prices under Insulated Equilibrium satisfy

$$P^{\mathcal{I},j} = C^j_I + \mu^{\mathcal{I},j} - \sum_{\mathcal{J} \neq \mathcal{I}} N^{\mathcal{J},j} \left(\left[-\frac{\partial N^{\mathcal{J}}}{\partial P^{\mathcal{J}}} \right]^{-1} \left[\frac{\partial N^{\mathcal{J}}}{\partial N^{\mathcal{I}}} \right]_{j,} \right) \cdot [-D^{\mathcal{I},j}]. \quad (20)$$

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 I affects the payoffs of consumers on all of the $S - 1$ other sides. Consequently, the prices charged to side 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 θ^I on side I receives gross payoff $v^I(\mathcal{X}, N, \theta^I)$, where $v^I : \wp(\mathcal{M}) \times [0, 1]^{mS} \times \Theta^I \rightarrow \mathbb{R}$ 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 $\sigma^{I,j} : [0, 1]^{mS} \rightarrow \mathbb{R}$ maps from N , the entire coarse allocation of consumers, including on side 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),¹⁷ about the strategies of other consumers on the same side. Suppose that prior to choosing their actions, side 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 $\overset{\dots}{N}^I$ denote the coarse allocation on side I that side I consumers believe will prevail in the given consumer game. It is apparent that the best-response strategy profile of side I consumers depends on $\overset{\dots}{N}^I$. Residual Insulating Tariffs, whose definition we now restate, adapted to this context, pin down a unique value of $\overset{\dots}{N}^I$ and thus a unique value of N^I that is consistent with Consumer Nash Equilibrium.

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

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

As before, all platforms announcing Residual Insulating Tariffs on all sides of the market, gives

¹⁷The 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).

rise to an Insulating Tariff System, $\bar{P}(N)$, 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

$$0 = \left[\frac{\partial N^I}{\partial P^I} \right] \left[\frac{\partial \bar{P}^I}{\partial N^I} \right] + \left[\frac{\partial N^I}{\partial N^I} \right] \Leftrightarrow \left[\frac{\partial \bar{P}^I}{\partial N^I} \right] = \left[-\frac{\partial N^I}{\partial P^I} \right]^{-1} \left[\frac{\partial N^I}{\partial N^I} \right]. \quad (21)$$

When platforms' tariffs satisfy equation (21), for *any* beliefs that side 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 (18), and the prices that arise under the socially optimal allocation and the Insulated Equilibrium allocation are, respectively,

$$P^{I,j} = C_I^j - \sum_{\mathcal{J}=\mathcal{A},\mathcal{B},\mathcal{C}\dots} N^{\mathcal{J},j} \bar{v}_j^{\mathcal{J},j} \quad (22)$$

and

$$P^{I,j} = C_I^j + \mu^{I,j} - \sum_{\mathcal{J}=\mathcal{A},\mathcal{B},\mathcal{C}\dots} N^{\mathcal{J},j} \left(\left[-\frac{\partial N^{\mathcal{J}}}{\partial P^{\mathcal{J}}} \right]^{-1} \left[\frac{\partial N^{\mathcal{J}}}{\partial N^I} \right] \right)_{j,\cdot} \cdot [-D^I_{\cdot,j}]. \quad (23)$$

Notice that the only difference between equations (22) and (23), corresponding to the case where there are within-side externalities, and equations (19) and (20), 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 I , this affects the quality of j as perceived by side 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 consumption 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.¹⁸ 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 good-consumer idiosyncratic errors (typically assumed Type I Extreme Value distributed for tractability) and α_i^I 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 $\frac{\beta_i^I}{\alpha_i^I}$ and factors entering into $\frac{\eta_i^{I,j}}{\alpha_i^I}$ and thus leading users to be marginal or infra-marginal. A particularly natural such relationship is obviously income heterogeneity (heterogeneity in α_i^I), but others might be coefficients of the same (or opposite) sign in the regression of β_i^I and $\eta_i^{I,j}$ 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^I 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 $P^{I,j}$ could be used for $N^{\mathcal{J},j}$ (as price affects demand)

¹⁸Note that this is true as well in one-sided models with endogenous characteristic choices.

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 2 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 of the extensions that we consider to be most promising. One interesting avenue would be more detailed treatment of 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.

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. On a more technical level, a more detailed analysis of conditions for existence, stability and uniqueness of IE would be useful.

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, as Bulow and Roberts (1989) shows, Weyl’s (2010) model is equivalent to Segal’s (1999) 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, Draszelski, and Harrington (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.

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." *RAND Journal of Economics*, 37, 668–691.
- Armstrong, Mark and John Vickers (2001), "Competitive Price Discrimination." *RAND Journal of Economics*, 32, 579–605.
- 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*.
- Azevedo, Eduardo, E. Glen Weyl, and Alexander White (2011), "Walrasian Equilibrium without Gross Substitutes." *Mimeo*.
- Bedre-Defolie, Özlem and Emilio Calvano (2009), "Pricing Payment Cards." *Mimeo*.
- Bergemann, Dirk and Stephen Morris (2005), "Robust Mechanism Design." *Econometrica*, 73, 1771–1813.
- 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." *RAND Journal of Economics*, 25, 242–262.
- Borenstein, Severin (1989), "Hubs and High Fares: Dominance and Market Power in the U.S. Airline Industry." *RAND Journal of Economics*, 20, 344–365.
- 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." *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*.
- Cass, David and Karl Shell (1983), "Do Sunspots Matter?" *Journal of Political Economy*, 91, 193–227.
- 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 Joseph E. Jr. Harrington (2009), "Avoiding Market Dominance: Product Compatibility in Markets with Network Effects." *RAND Journal of Economics*, 40, 455–485.
- Chung, Kim-Sau and J. C. Ely (2007), "Foundations of Dominant-Strategy Mechanisms." *Review of Economic Studies*, 74, 447–476.
- 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, 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*.
- Hagiu, Andrei (2006), "Pricing and Commitment by Two-Sided Platforms." *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*.
- 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.
- Jackson, Matthew O. (2008), *Social and Economic Networks*. Princeton University Press, Princeton, New Jersey.
- Jaffe, Sonia and E. Glen Weyl (2010a), "Linear Demand Systems are Inconsistent with Discrete Choice." *B. E. Journal of Theoretical Economics*, 10.
- Jaffe, Sonia and E. Glen Weyl (2010b), "The First-Order Approach to Merger Analysis." *Mimeo*.
- Jeon, Doh-Shin and Jean-Charles Rochet (2010), "The Pricing of Academic Journals: A Two-Sided Market Perspective." *American Economic Journal: Microeconomics*, 2, 222–55.
- Jevons, William Stanley (1871), *The Theory of Political Economy*. MacMillan, London.
- 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.
- 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 Practitioner's 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 (2011), "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, 381–392.
- 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.
- Rohlf, Jeffrey (1974), "A Theory of Interdependent Demand for Communications Service." *Bell Journal of Economics and Management Science*, 5, 16–37.
- 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, 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, 3–13.
- Spence, A. Michael (1975), "Monopoly, Quality, and Regulation." *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.
- Wilson, Robert (1987), "Game Theoretic Analysis of Trading Processes." In *Advances in Economic Theory* (Truman F. Bewley, ed.), Cambridge University Press, Cambridge, UK.