

Two-sided Markets, Pricing, and Network Effects*

Bruno Jullien
Toulouse School of Economics

Alessandro Pavan
Northwestern University and CEPR

Marc Rysman
Boston University

July 2021

Abstract

The chapter has 9 sections, covering the theory of two-sided markets and related empirical work. Section 1 introduces the reader to the literature. Section 2 covers the case of markets dominated by a single monopolistic firm. Section 3 discusses the theoretical literature on competition for the market, focusing on pricing strategies that firms may follow to prevent entry. Section 4 discusses pricing in markets in which multiple platforms are active and serve both sides. Section 5 presents alternative models of platform competition. Section 6 discusses richer matching protocols whereby platforms price-discriminate by granting access only to a subset of the participating agents from the other side and discusses the related literature on matching design. Section 7 discusses identification in empirical work. Section 8 discusses estimation in empirical work. Finally, Section 9 concludes.

1 Introduction

Consider a consumer that starts the day by searching for news on a tablet computer and then ordering a car ride from a ride-sharing company over a mobile telephone. At each of these points, an intermediary connects the consumer to various providers. The consumer's search terms for news are processed by a web search engine that mediates between content consumers and content providers. Simultaneously with this search, a separate intermediary system connects the consumer (and the consumer's search behavior) to advertisers, leading to advertisements. The ride-sharing service is an intermediary between the consumer and

*Chapter book to appear in the Handbook of Industrial Organization. For comments and suggestions, we thank the editors, anonymous referees, as well as Mark Armstrong, Jay-Pil Choi, Avi Goldfarb, Hanna Halaburda, Igal Hendel, Chuqing Jin, Marshall Van Alstyne, Shuang Wang, Stefan Weiergraeber, Julian Wright, and Junjie Zhou. Special thanks to Martin Peitz for very detailed comments. Jishan Du, Tyler Haas, and Shuang Wang provided excellent research assistance. Bruno Jullien acknowledges funding from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program) and ANITI (ANR grant 3IA), and from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 670494).

potential ride providers. Payment for the ride-sharing service is processed through a payment network, an intermediary that facilitates payments between buyers and sellers. These interactions are enabled by an internet service provider that is an intermediary between consumers and a complex market of telecommunication systems. In addition, the devices in this story, the tablet and phone, can be regarded as intermediaries themselves, connecting consumers to software applications via the devices' operating systems. And most likely, at least one of them was purchased through an online marketplace that connects consumers to sellers.

Economics regards these intermediaries as *platforms*. Platforms are firms, or services of firms, that connect market participants and allow them to interact or transact. Platforms are prevalent in many markets: health insurance companies, for example, mediate between consumers and care providers, stock exchanges mediate between buyers and sellers (or liquidity providers and liquidity takers) and cable television companies mediate between content providers and viewers. The rise of platforms, especially on the internet, has changed the way vast parts of the economy are organized and has led to intense reexamination of the nature of market power, competition policy, and regulation. Platform economics is central to understanding a wide variety of recent policy debates, such as net neutrality, financial market reforms, antitrust policy, as well as related topics such as privacy, consumer protection, and media diversity.

A central aspect of platform economics is the role of network effects, which apply when a product is valued based on the extent to which other market participants adopt or use the same product. Of particular interest are indirect network effects, which emerge when the adoption and use of a product leads to increased provision of complementary products and services, with the value of adopting the original product increasing with the provision of such complementary goods. For instance, as more consumers adopt a video game console such as the Sony Playstation, more game-developers invest in developing games for that platform, raising the value of the console to consumers. In this sense, indirect network effects lead the platform firm to take into account the various interdependencies between the two sides of the market, and the pertinent literature studying such interdependencies is often termed the study of *two-sided markets*. Naturally, one can also consider multi-sided markets consisting of more than two sides. For instance, in producing the Windows operating system, Microsoft manages a three-sided market between hardware providers, software providers, and consumers. One might even distinguish between different types of hardware (CPUs, printers, screen devices, etc.) and types of software (productivity apps, games, etc.), in which case Windows mediates a large many-sided market.

As a product with network effects diffuses into the market, it becomes more valuable and drives further adoption. Indirect network effects thus lead to a feedback loop as more participants on each side of the platform find it more valuable to adopt and use the platform when they expect the other side to attract more users. This phenomenon leads to efficiencies as more market participants are able to interact with each other but also, in some circumstances, market power, as network effects can protect platform owners from entry. In markets with low marginal costs, as is the case for many digital markets, platforms with strong network effects can grow to be enormous and eventually dominate the market.¹

¹See also Goldfarb and Tucker (2019) for an extended analysis of digital markets

The most obvious tool that a platform has to manage and expand its use is price. Pricing decisions in the face of indirect network effects are complex because raising the price on one side of the market affects demand not only on that side of the market but also on the other, as each side responds to changes in the participation on the other side. Finding the correct approach to pricing is key to the success of a platform. For policy-makers attempting to evaluate the performance of a market, understanding the complex determinants of prices is also crucial.

In this chapter, we provide an overview of the literature on platform economics, focusing on indirect network effects and pricing. We first provide a unified theoretical treatment of a wide variety of approaches to these issues. We derive equilibrium and welfare-maximizing prices and compare the factors that lead to deviations from optimality. Standard intuition about how prices should respond to costs and competition often do not hold, and we lay out the intuition required to understand these markets. We discuss different treatments of how network effects may enter into consumer decision-making. For instance, whether agents have heterogeneous valuations for the number of agents on the platform turns out to have important implications. In addition, the platforms' ability to price-discriminate is critical, leading to the so-called *divide-and-conquer* strategies. Modeling platform firms may also require new equilibrium concepts as the relevant models are rife with multiple equilibria and one may wish to focus on the most relevant ones.

We contrast monopoly platforms with oligopoly platforms. Whether each agent *multi-homes* across several platforms or *single-homes* on one platform plays an important role. One may think that if agents are divided into two sides of the market, such as buyers and sellers, and one set of agents multi-homes, then irrespective of whether this set comprises buyers or sellers, all agents can interact with all agents on the other side and this guarantees a level of efficiency. However, a robust result in the literature is that which side multi-homes has important implications for pricing and the efficiency of the ensuing allocation. When only one side multi-homes, platforms compete on the single-homing side but behave like monopolists on the side that multi-homes, selling access to their single-homing agents. Total network benefits increase, but prices on each side may increase or decrease as a result of multi-homing.

We also take on some more challenging concepts to model, such as consumer information and beliefs, and dynamics. Furthermore, in many contexts, a platform plays the central role of a matchmaker. Matching markets, and more generally market design, play an important role in platform economics. We provide an overview of the matching-design literature, once again focusing on pricing. Although we show that prices respond to this wide variety of modeling assumptions, we are able to draw out some consistent themes about pricing in these markets.

Next, we turn to empirical work in this area. A starting point and a central issue in this literature is estimating the magnitude of network effects. We view network effects as a form of “social spillover” or “neighborhood effect”; the choice of an agent is influenced by the choices of the agent's peers. Identifying network effects is fraught with problems of simultaneity and omitted variables. These problems are similar to those in the literature on peer effects in social economics, such as identifying how the performance of one student affects the performance of classmates, or how out-of-wedlock childbirth and criminal arrest are influenced by the behavior and outcomes of peers. The social economics literature has

made great strides in understanding the identification of these peer effects, exemplified by the work of Manski (1993, 1995) on the *reflection problem* and the literature that followed. We view the relationship between the social economics and network effects literatures as under-explored. In this chapter, we provide a review of empirical approaches to estimate network effects from the perspective of the literature on social economics. Doing so provides a valuable connection between the network effects literature and the wider economics literature and provides new insights into the identification and estimation of network effects.

With the estimation of network effects in hand, we turn to a review of the estimation of pricing in settings with network effects. We argue that empirical papers have mostly taken one of two approaches. The first approach is to obtain data on many platforms, often separated by geography, and to analyze how endogenous variables such as prices respond to exogenous market characteristics. These papers often use reduced-form econometrics and attempt to use theory to evaluate the effects that are expected to be found. The other approach is to obtain data on a single platform and study the interaction between agents on the platform, taking the platform pricing as exogenous. These papers are often quite structural and utilize detailed models of agent behavior on the platform. These papers then calculate various counterfactual outcomes associated with different price changes and evaluate why firms choose the prices they do and what would happen if prices were to change. These methods provide different insights into pricing by platforms, but we are able to draw out some general lessons and connect them to the theoretical results in the first part of the chapter. For instance, in the theoretical, part we discuss the *seesaw effect*, namely the property that factors that depress price on one side of the market tend to raise price on the other side; in the empirical part, we then discuss papers that support the existence of the seesaw effect across a variety of industries using a variety of methods.

1.1 Terminology and background

The literature on network effects and two-sided markets has early antecedents in Rohlfs (1974), Baxter (1983), and, on the empirical side, Rosse (1970). The modern incarnation of the literature on network effects is typically associated with Arthur (1989), Katz and Shapiro (1985), and Farrell and Saloner (1986). The first theoretical models of indirect network effects are Church and Gandal (1992), as well as Chou and Shy (1990). An important overview of this earlier work is Farrell and Klemperer (2007). This earlier research on network effects and particularly indirect network effects typically focuses on consumer adoption decisions and does not treat pricing or other decisions by the platform as central. The focus changed with the launch of the literature on two-sided markets, pioneered by Caillaud and Jullien (2001), Rochet and Tirole (2002), and Parker and Van Alstyne (2005). Important contributions to the development of the literature are Rochet and Tirole (2006) and Armstrong (2006). Rysman (2009) provides an introduction, and a recent overview of various issues related to platform economics is Belleflamme and Peitz (2021).

In this section, we define what we mean by a platform. Hereafter, we use the terms “two-sided market” and “platform market” interchangeably. Several papers offer definitions of two-sidedness, such as Rochet and Tirole (2006), Weyl (2010), and Hagiu and Wright (2015a,b). We do not wish to adjudicate among these different definitions, but rather offer some sense of what economists normally mean with these terms. We generally consider a

two-sided market to be one in which at least two distinct sets of agents (or sides) interact through an intermediary—the platform—and in which the behavior of each set of agents directly impacts the utility, or the profit, of the other set of agents. The impact of one set of agents on the other, and the resulting feedback to the first set of agents, is an indirect network effect. To the extent that such an effect is not perfectly internalized by the platform’s pricing, it represents a *network externality*.

For example, suppose buyers and sellers appear at an open marketplace where they may engage in trade. The marketplace organizer is the intermediary, or platform, and the sellers naturally care about how successful the intermediary is in attracting buyers, and vice versa, so the firm is two-sided. In contrast, suppose the marketplace organizer decides to purchase the sellers’ products wholesale and then retail the inventory to buyers. The sellers receive the wholesale price whether buyers materialize or not, and so the sellers no longer have an interest in the buyers’ behavior. In this case, we refer to the firm as “traditional,” or one-sided.²

In reality, almost every real-world firm has some elements of two-sidedness to it, and it makes little sense to classify firms with some binary distinction as being a platform or not. Rather, one should see the platform nature of a firm as a continuous dimension.³ Platform economics may be more important for understanding some phenomena than others, even within a given firm or market. In an antitrust or in a regulatory setting, it may be necessary to categorize firms as one- or two-sided. In this chapter, instead, we do not attempt to do so, and instead favor the continuous nature of two-sidedness.

It is common in overview articles on platforms to list the largest firms in the world and identify which are platforms and which are not, highlighting the growth and importance of platform firms in the organization of modern business activity. However, we eschew this because we do not view the distinction between platform and non-platform firms as binary. For example, it is typical to identify Amazon as a platform firm and automobile manufacturers as traditional one-sided firms. However, Amazon has always sold books in a one-sided manner and, to the extent that it is expanding the sales of its own branded products, it is moving further in that direction. A significant source of operating profits at Amazon currently comes from its cloud-computing services, of which a substantial share is provided directly by Amazon.⁴ In contrast, the Ford Motors distribution network can be seen as a platform mediating the interactions between consumers and dealerships. Ford cannot attract dealers to sell its products unless it has interested consumers, and consumers cannot purchase Ford products unless dealers offer them. In addition to setting a wholesale

²This distinction is similar to the one in Hagiu and Wright (2015b), who locate the definition of a platform market in the control the platform exerts over the transactions. The market is one-sided if the intermediary controls all of the features of trade, particularly the price that the buyer pays to the seller.

³To push this point further, consider again the example of a marketplace organizer buying a seller’s product wholesale and retailing the inventory to consumers. We argued above that such instances are typically associated with traditional one-sided markets due to the fact that the seller does not care about the success of the organizer. In practice, however, we often observe buy-back clauses whereby the marketplace organizer can force the seller to buy inventory back if the organizer sells poorly. Even in the absence of such formal clauses, if the seller has a long-term relationship with the organizer, the seller may have an interest in the success of the organizer with consumers and this breaks down the strict one-sided nature of the example.

⁴In fact, Amazon also has an AWS (Amazon Web Services) Marketplace that connects cloud users to third-party sellers through Amazon’s cloud services offerings. To the extent that the marketplace is driving Amazon’s cloud services revenue and profits, Amazon is a platform for cloud services.

price to dealers, Ford manages a system of promotions to consumers and dealers to achieve successful outcomes. In practice, whether a firm is best viewed as a platform or a traditional firm can depend on which questions one seeks to answer about the firm's behavior.

Taking a broader perspective, whether a firm is two-sided (or multi-sided) is often endogenous to the firm. Firms can integrate into one side of the market to reduce the platform nature of the firm. For example, as pointed out above, Microsoft produces the Windows operating system and thus relies on a three-sided market between consumers, hardware manufacturers, and software developers. Adjacent to that, Microsoft produces both the operating system and the hardware for the Xbox gaming system, leaving only a two-sided market between consumers and software developers. Because of the endogenous nature of two-sidedness, it may be best to talk about two-sided strategies rather than two-sided markets to emphasize the role of a firm's choice.⁵ The example above about Amazon expanding its own brand of products is an example of a firm endogenously choosing to be less two-sided in specific product markets, which ties together these ideas about two-sidedness being endogenous and continuous. Similarly, internet search is thought of as a canonical two-sided market that connects content consumers with content providers. However, to the extent that Google integrates into content (i.e., videos, maps, shopping services, music, etc.), it appears to be choosing a less two-sided approach.⁶

An idea related to the endogeneity of two-sidedness is that two-sidedness typically describes firms, not markets. There are many examples of platform firms competing against one-sided firms in the same market. For instance, brokers that connect shippers with trucks compete in the same market against shippers that own their own trucks and may not need for-hire trucks. Media that is supported by advertising (and so is two-sided) competes against media that does not rely on advertising. Examples of media that do not use advertising are Consumer Reports, Wirecutter.com, and *Le Canard Enchaîné*. We refer the reader to Hagiu and Wright (2015b) for an overview of the recent literature addressing the endogeneity of two-sidedness.

Thus, from our perspective, calling this literature "two-sided markets" is not ideal. Two-sidedness is not a binary outcome endowed by a market but is typically rather a choice made by firms about what ways to be two-sided. Having clarified all of this, we follow in this chapter the common practice of referring to many of the issues we study as pertaining to two-sided and platform markets, as is clear from the chapter's title. After all, this is the way this literature is referred to, and, given our focus on pricing (which is the most well-studied phenomenon in this literature), there is little lost by treating two-sidedness as a market concept.

Our focus on pricing leads us to leave out large and important swaths of the litera-

⁵A famous definition of two-sidedness appears in Rochet and Tirole (2006), where a market is defined as two-sided if the price structure of the transactions matters beyond the price level. As they point out, this often hinges on whether the platform charges participation fees or not. However, we often think of a firm's price structure as endogenous and not as a natural way to define a market type. One could argue that their definition is a useful way to determine whether a firm adopts a two-sided strategy or not, rather than whether the market is two-sided or not.

⁶Some readers will recall that America Online provided internet search in the early days of the World Wide Web and, for some consumers, essentially all of their internet content. To the extent that America Online adopted a one-sided approach to internet search, it is an example of how markets are not inherently two-sided or not.

ture, that we briefly review here. First, whereas the concepts we describe can be applied to all varieties of platform markets, we do not attempt to address the details of specific industries, such as finance and advertising-supported media. For overviews of the literature on two-sided media, see Anderson and Jullien (2015) and Peitz and Reisinger (2015). Rysman and Wright (2014) review the two-sided markets literature in the context of payment cards. Many industry-specific applications of this literature exist, such as voluntary standard setting organizations (see Lerner and Tirole, 2006).

Further, whereas we discuss deviations of pricing from social optimality and the effects of competition on outcomes, we do not review the competition policy literature here. See, for example, Jullien and Sand-Zantman (2020) and Evans and Schmalensee (2015) for an overview of this literature. The effects of platform economics on competition policy are profound, with numerous commissions around the world addressing issues related to platforms and laying out strategies for approaching digital platforms.⁷ Platform economics has pointedly affected policy-making. For instance, in the United States, in the context of a monopolization case, the Supreme Court has recently integrated two-sided markets into its approach to market definition, and the Securities and Exchange Commission explicitly recognizes the role of platform economics in regulating securities exchanges.⁸ Indeed, platform economics appears to be leading to new forms of regulation, such as the recently proposed Digital Markets Act in the EU and the Digital Markets Unit in the United Kingdom.⁹ Platform economics has also important implications for privacy regulation, which we do not review here.¹⁰

Finally, we stop short of covering platform design. A successful platform requires more than just appropriate pricing.¹¹ Platforms also invest in quality that affects one side or the other, and choose the rules and nature of engagement between market participants, as well as the information that each side gets to observe. Ratings systems are a familiar feature of many digital platforms, and they bring up many design questions, as described in Tadelis (2016). The question of whether to be two-sided or not, as in Hagiu and Wright (2015b) is also largely an example of platform design. While these issues come up in the chapter, particularly when we discuss matching markets, we focus on them only to the extent that they interact with our questions about pricing. Along these lines, we restrict ourselves to

⁷Examples include (a) the report of the Stigler Committee on Digital Platforms from the Stigler Center for the Study of the Economy and the State, University of Chicago Booth School of Business, (b) “Unlocking Digital Competition” from the Digital Competition Expert Panel for the United Kingdom, and (c) “Competition policy for the digital area” by Jacques Cr  mer, Yves-Alexandre de Montjoye and Heike Schweitzer, for the European Commission, all released in 2019.

⁸The SEC released its “Staff Guidance on SRO Rule Filings Relating to Fees” in May 2019 which states “The platform theory of competition (“a total platform theory”) also provides a potential pathway to demonstrating a competitive environment. Total platform theory generally asserts that when a business offers facilities that bring together two or more distinct types of customers, it is the overall return of the platform, rather than the return of any particular fees charged to a type of customer, that should be used to assess the competitiveness of the platform’s market.” See also *Ohio vs. American Express Co.*, 585 US 2018.

⁹See European Commission (2020), “Proposal for a regulation of the European Parliament and the Council on contestable and fair markets in the digital sector (Digital Markets Act)” and United Kingdom (2020), “Government response to CMA digital advertising market study”.

¹⁰See Acquisti, Taylor, and Wagman (2016) for an overview of some of the privacy literature.

¹¹Teh (2021) and Choi and Jeon (2020) provide interesting analysis of the link between pricing instruments and platform design.

matching markets with payments, although there is a substantial literature on matching without payments (see, e.g., Roth, 2018).

In peer-to-peer networks, such as WhatsApp and initially Facebook and LinkedIn, content creators and content consumers are one and the same. Their dynamics are governed by a direct network effect rather than indirect network effects, and so their analysis falls outside of the scope of this chapter. Einav, Farronato, and Levin (2016) reviews peer-to-peer markets. To the extent that these peer-to-peer networks also sell advertising that reaches network users, these networks fall within our scope of analysis, although our focus is not on media. In particular, we do not discuss situations in which spillovers run in only one direction and we largely ignore negative spillovers. For instance, advertisers value newspaper readers positively but readers may not value, or negatively value, advertisers. The articles on media that we cite above cover the case of negative spillovers in detail.

It is also worth keeping in mind that terms such as “platforms,” “networks,” “network effects,” and “network externalities,” are used in many other areas in economics and across different literatures. For instance, in the strategic management literature, Cusumano, Gawer, and Yoffie (2019) distinguish between *innovation platforms* and *transaction platforms*. Innovation platforms are systems or general-purpose tools that facilitate new products. In this view, Java is a platform for creating internet applications and cloud computing is a platform for delivering services or infrastructure over the internet. Transaction platforms facilitate transactions. From our perspective, there is nothing particularly two-sided or connective about what this literature calls innovation platforms and our concept of platform and two-sidedness fits within their concept of transaction platforms. Similarly, there is a large literature on social networks and network formation that models exactly which agent is connected to which other agent (see, for example, Jackson, 2008). We discuss how platforms use payments to influence specific matches and hence shape the network of the interactions. Furthermore, some of the issues we review when discussing the empirical literature on two-sided markets also appear relevant for the broader literature on network formation. However, we do not discuss the connection between these two literatures at length in this chapter.

2 Monopoly

In this section we review the literature that examines pricing in two-sided markets where the interactions between the sides are mediated by a single platform. We first introduce the setting and some useful notation and review the key results. We then discuss distortions compared to welfare maximization. Finally we discuss various issues such as the chicken & egg problem, zero prices, congestion, transaction fees or dynamic pricing.

2.1 Basic framework and notation

A single firm (hereafter the platform) matches agents from two different sides of the market, denoted by $i = 1, 2$. Each side is populated by a unit-mass continuum of agents. Each agent derives utility from the product offered by the platform (e.g., Amazon Retail or the iPhone) but also by interacting with agents from the other side of the market who also join the platform. Formally, each side- i agent has a type $\theta_i = (v_i, \gamma_i) \in \Theta_i \subset \mathbb{R}^2$ and derives a gross

payoff

$$u_i(\theta_i, q_j) = v_i + \gamma_i q_j$$

from interacting with a mass of agents of size q_j from side $j \neq i$. These gross payoffs are net of any total payment P_i that the agents make to (or receive from) the platform. Preferences are quasi-linear in payments so that the net payoff is equal to $u_i(\theta_i, q_j) - P_i$. Each agent who does not join the platform obtains a reservation payoff normalized to zero. Each type θ_i is an independent draw from some distribution F_i and is the agent's private information. The component v_i in the agent's payoff is a proxy for the utility that the agent derives from the platform's product or service, whereas the component γ_i is a proxy for the importance the individual assigns to interacting with agents from the other side of the market.¹² Hereafter, we refer to v_i as the agent's "stand-alone value" and to the term γ_i as the agent's "interaction benefit." Following the pertinent literature, interactions with the other side will also be referred to as "usage."

Notice that the above formalization implies that different agents from the same side may assign different value to interacting with agents from the opposite side, but each agent values all agents from the other side equally.¹³ Consistently with the rest of the literature surveyed in the first few sections of this chapter, we also assume that platforms do not engage in price discrimination within sides. Any pair of agents from the same side who joins the platform is granted access to all agents on-board from the opposite side and is charged the same total payment by the platform. We relax these assumptions in Section 6 where we consider richer preference specifications and allow for discriminatory prices.

The total payment P_i to the platform may have two components, an access fee p_i and a transaction fee proportional to the number of interactions with agents from side $j \neq i$. Assuming that the number of interactions is proportional to the mass of agents on side j and letting t_i denote the expected transaction fee per agent on side j , we have that the total payment by each agent joining from side i is equal to $P_i = p_i + t_i q_j$. The above representation assumes that agents do not exchange money directly with agents from the other side or do so at a fixed price. It also assumes that the only relevant margin is the extensive one: agents choose whether or not to join the platform but do not control the intensity of their interactions with agents from the other side of the market. Finally, it assumes that every agent who joins from one side expects the same number of interactions with the agents participating from the opposite side. We discuss the implications of relaxing some of these assumptions later in the chapter.

Fixing the measure of agents q_j joining the platform from side j , we then have that the total demand from side i when the side- i 's total payment is equal to P_i is given by

$$q_i = D_i(P_i; q_j) \equiv \Pr \{(v_i, \gamma_i) \in \Theta_i : v_i + \gamma_i q_j \geq P_i\}. \quad (1)$$

¹²Given our focus on pricing, in most of the chapter, we treat v_i and γ_i as exogenous parameters, although in certain markets they may be influenced by non-price decisions that we do not consider here. As most of the literature, we also assume that network benefits are linear in the other side's participation. Linearity is relaxed in Section 4.1 and 6. Sections 2.5 and 6 endogenize network benefits through taxation and matching.

¹³This assumption only holds before the agents make their participation decisions, and reflects the lack of information the agents have at the contracting stage about the identity of the participating agents from the other side of the market. This assumption will be relaxed later in the chapter when considering price discrimination

The measure of agents q_1 and q_2 from the two sides of the market is then the solution to the two conditions $q_i = D_i(P_i; q_j)$, $j \in \{1, 2\}, j \neq i$. This formulation highlights that, from side- i 's perspective, q_j is a quality dimension that is determined endogenously through prices. The platform incurs a cost c_i for each side- i agent it brings on board and a cost σ for each interaction between the two sides. The platform's profits are thus given by

$$\Pi = \sum_{i,j=1,2,j \neq i} (P_i - c_i) q_i - \sigma q_1 q_2.$$

Hereafter, we refer to the term

$$\eta_i(P_i; q_j) \equiv -\frac{D_i(P_i; q_j)}{\frac{\partial D_i(P_i; q_j)}{\partial P_i}} = \frac{P_i}{\epsilon_i(P_i; q_j)}$$

as the side- i inverse semi-price elasticity of the demand, where

$$\epsilon_i(P_i; q_j) \equiv -\frac{\partial D_i(P_i; q_j)}{\partial P_i} \frac{P_i}{D_i(P_i; q_j)}$$

is the standard elasticity of the side- i demand with respect to its own price.

Note that the sale of the platform's stand-alone product is bundled with the interaction with the opposite side.

The environment described above encompasses most of the models considered in the literature. In particular, it admits the following specifications as special cases.

- Caillaud and Jullien (2001, 2003): F_i is a Dirac measure at some $(v_i^0, \gamma_i^0) \in \mathbb{R}^2$, $i = 1, 2$, meaning that agents' preferences are homogeneous within sides;
- Rochet and Tirole (2003): for each $i = 1, 2$, $c_i = p_i = 0$; furthermore, F_i assigns measure one to the set $\{(v_i, \gamma_i) \in \Theta_i : v_i = 0\}$, meaning that all agents derive utility only from interacting with other agents from the opposite side of the market (zero stand-alone values);
- Armstrong (2006): $\sigma \equiv 0$ and $t_i = 0$, $i = 1, 2$; furthermore, for each $i = 1, 2$ there exists γ_i^0 such that F_i assigns measure one to the set $\{(v_i, \gamma_i) \in \Theta_i : \gamma_i = \gamma_i^0\}$, meaning that all agents assign the same value to interacting with agents from the opposite side of the market (homogeneous interaction benefits).
- Rochet and Tirole (2006) and Weyl (2010): F_i is absolutely continuous over \mathbb{R}^2 , meaning that agents are heterogeneous in both their stand-alone values and their interaction benefits.

2.2 Profit-maximizing prices

Assuming a smooth demand and no coordination problems¹⁴, the profit-maximizing prices (P_1^*, P_2^*) , along with the demand/participation (q_1^*, q_2^*) they induce, solve¹⁵

$$P_i^* - [c_i + q_j^* (\sigma - \tilde{\gamma}_j(P_1^*, P_2^*))] = \eta_i(P_i^*; q_j^*) \quad (2)$$

for $i, j = a, b, j \neq i$, with $q_j^* = D_j(P_j^*; q_i^*)$, and with

$$\tilde{\gamma}_j(P_1, P_2) \equiv -\frac{\frac{\partial D_j(P_j; q_i)}{\partial q_i}}{\frac{\partial D_j(P_j; q_i)}{\partial P_j}} = \mathbb{E}[\gamma_j | v_j + \gamma_j q_i = P_j]. \quad (3)$$

The above formula is the two-sided analog of the familiar monopoly price formula in one-sided markets. The right-hand side is the familiar effect by which a monopolist's ability to price above marginal cost is inversely related to the elasticity of the demand. The only novelty relative to one-sided markets is that such elasticity is now computed by accounting for the fact that it depends on the size of the side- j 's participation, q_j^* . The term in square parenthesis on the left-hand-side is the cost of bringing a marginal agent on board, c_i , augmented by the cost of matching the agent to all agents on board from the opposite side, σq_j^* , and reduced by the product between the average interaction benefit $\tilde{\gamma}_j(P_1^*, P_2^*)$ experienced by the side- j marginal agents and the size of the side- j demand q_j^* . To understand this last term, note that when the platform brings on board an extra agent from side i , it can then increase its side- j price by $\tilde{\gamma}_j(P_1^*, P_2^*)$ while keeping constant the demand on side j at q_j^* . Holding the side j 's participation constant, the marginal effect on the side j 's profits is thus equal to $\tilde{\gamma}_j(P_1^*, P_2^*) q_j^*$. When the average network effect $\tilde{\gamma}_j(P_1^*, P_2^*)$ among the side- j marginal agents is positive, this last effect thus contributes to a reduction in the marginal cost of bringing new agents on board from side i .

2.2.1 Homogeneous interaction benefits (Armstrong, 2006)

Consider the case in which all agents from the same side have the same interaction benefit, with the latter equal to γ_1 on side 1 and γ_2 on side 2. Further assume that the platform incurs no cost in matching agents, with the only cost being the cost of getting agents on board: $\sigma = 0$ and $c_1, c_2 > 0$. In this case, the profit-maximizing prices are given by

$$P_i^* = c_i - \gamma_j q_j^* + \eta_i(P_i^*; q_j^*)$$

with $\eta_i(P_i^*; q_j^*) = [1 - F_i^v(P_i^* - \gamma_i q_j^*)] / f_i^v(P_i^* - \gamma_i q_j^*)$ and $q_j^* = 1 - F_j^v(P_j^* - \gamma_j q_i^*)$, $i = 1, 2$, $j \neq i$, where F_i^v and f_i^v are the cumulative distribution function and the density of the marginal distribution of F_i over the v -dimension, respectively.

The novel term relative to the formula for the optimal price in a one-sided market is the term $-\gamma_j q_j^*$. This term can be interpreted as an opportunity cost of excluding a marginal

¹⁴There may be multiple demand configurations compatible with the same prices. When this is the case, we assume here that the monopolist can coordinate the market on the allocation most favorable to her.

¹⁵See Weyl (2010). This first-order condition corresponds to the optimal choice of price P_i holding the quantity q_j constant on the other side, which requires adjusting P_j accordingly.

consumer on side i . Indeed suppose the platform increases the side- i participation by $\varepsilon_i > 0$. It can then maintain the participation on side j constant at q_j^* by increasing its side- j price by $\gamma_j \varepsilon_i$, thereby generating an additional revenue of $\gamma_j q_j^* \varepsilon_i$.

2.2.2 Homogeneous stand-alone benefits (Rochet and Tirole, 2003)

Next, consider the case where agents are homogeneous in their stand-alone values, with the latter normalized to zero, but have heterogeneous interaction benefits, on both sides. Further assume that the platform incurs no cost to bring an agent on board, but incurs a cost for each agent it matches: $c_1 = c_2 = 0$ and $\sigma > 0$. Finally assume that the platform charges each side- i agent t_i per transaction, i.e., per each match with the other side. The total payment collected from each participating agent from side i is thus equal to $P_i = t_i q_j$, $i, j = 1, 2, j \neq i$. Let F_i^γ and f_i^γ denote the cumulative distribution function and the density of the marginal distribution of F_i over the γ -dimension, respectively. Hence, in this case, $q_i^* = 1 - F_i^\gamma(P_i^*/q_j^*)$, which implies that

$$\eta_i(P_i^*; q_j^*) = \frac{[1 - F_i^\gamma(P_i^*/q_j^*)] q_j^*}{f_i^\gamma(P_i^*/q_j^*)}.$$

The profit-maximizing payments can then be obtained from (2) which in this case yields

$$P_i^* - \sigma q_j^* + \frac{P_j^*}{q_i^*} q_j^* = \eta_i(P_i^*; q_j^*).$$

Letting $t_i^* = P_i^*/q_j^*$, $i = 1, 2, j \neq i$, be the transaction fees associated with the profit-maximizing total payments, we then have that the optimal transaction fees t_1^* and t_2^* are given by the solution to

$$t_1^* + t_2^* - \sigma = \frac{1 - F_1^\gamma(t_1^*)}{f_1^\gamma(t_1^*)} = \frac{1 - F_2^\gamma(t_2^*)}{f_2^\gamma(t_2^*)}, \quad (4)$$

which is the price formula identified by Rochet and Tirole (2003). Again, to interpret this price formula, note that $[1 - F_i^\gamma(t_i^*)]/f_i^\gamma(t_i^*)$ is the inverse semi-transaction fee-elasticity of the side- i demand (only agents with interaction benefit above t_i^* participate). The term $\sigma - t_j^*$, instead, is the opportunity cost (per agent joining from side j) of getting on board one extra agent from side i . The latter combines the cost to the platform of matching the agent with any of the agents joining from side j along with the fact that any side- j agent pays t_j^* per each match with any of the side- i agents. The optimal transaction fees thus reflect the general cross-subsidization pattern identified by the optimality conditions in (2).

2.3 Welfare-maximizing prices

Consider next the problem of a planner maximizing the sum of all agents' utilities and the platform's profit. For any vector of prices P_1 and P_2 , total welfare is given by

$$W \equiv \sum_{i=1,2} \int_{\{\theta_i: v_i + \gamma_i q_j \geq P_i\}} (v_i + \gamma_i q_j - c_i) dF_i(\theta_i) - \sigma q_1 q_2,$$

with q_i given by (1), $i = 1, 2$. One can then show that the welfare-maximizing prices (P_1^e, P_2^e) , along with the efficient participation profile (q_1^e, q_2^e) they induce, solve

$$P_i^e = c_i + q_j^e (\sigma - \bar{\gamma}_j(P_1^e, P_2^e)) \quad (5)$$

for $i, j = 1, 2, j \neq i$, with $q_i^e = D_i(P_i^e; q_j^e)$, where the demand functions D_i are given by the same functions as in (1), and where

$$\bar{\gamma}_j(P_1^e, P_2^e) \equiv \mathbb{E} [\gamma_j | v_j + \gamma_j q_i^e \geq P_j^e] \quad (6)$$

is the *average interaction benefit* of the participating agents from side j .

The efficient price on each side is thus equal to the marginal cost $c_i + \sigma q_j^e$ that the platform incurs to get a marginal agent on board from side i and then match him to all participating agents from side j , discounted by the network externality that the marginal agent exerts on all the participating agents from the other side of the market. The key difference with respect to profit maximization is that the externality accounts for the benefit that all agents from side j derive from interacting with the marginal agent from side i , whereas, in case of profit maximization, the only relevant externality accounted for by the monopolist is the one exerted on the *marginal* agent from the opposite side of the market (as in Spence, 1981, in a one-sided setting with endogenous quality).

2.4 Distortions

Using (2) and (5), one can express the difference between the profit-maximizing and the efficient prices as follows (see Tan and Wright, 2018 and Gomes and Pavan, 2019):

$$\begin{aligned} P_i^* - P_i^e &= \underbrace{\eta_i(P_i^*; q_j^*)}_{\text{markup}} + \underbrace{q_j^e (\bar{\gamma}_j(P_1^e, P_2^e) - \bar{\gamma}_j(P_1^e, P_2^e))}_{\text{Spence distortion}} \\ &+ \underbrace{q_j^e (\bar{\gamma}_j(P_1^e, P_2^e) - \bar{\gamma}_j(P_1^*, P_2^*))}_{\text{displacement distortion}} + \underbrace{(q_j^e - q_j^*) (\bar{\gamma}_j(P_1^*, P_2^*) - \sigma)}_{\text{scale distortion}}. \end{aligned} \quad (7)$$

The first term in the right-hand side of (7) is the usual distortion originating in the platform's market power, according to which the firm reduces output to extract more surplus from the infra-marginal agents. The only novelty relative to one-sided markets is the dependence of the markup on the measure of the participating agents from the opposite side, q_j^* .

The second term in the right-hand side of (7) is a Spencian distortion. It originates in the fact that a profit-maximizing monopolist internalizes only the effect of expanding the side- i participation on the side- j marginal agents, whereas a welfare-maximizing platform internalizes the effect that such an expansion has on all participating agents from side j .

The third term in the right-hand side of (7) is a distortion that accounts for the fact that the marginal agents under the profit-maximizing prices are not the same as under the welfare-maximizing ones and hence may have different interaction benefits.

Lastly, the the fourth term in the right-hand side of (7) is a distortion that originates in the different size of the participation by side j under profit and welfare maximization:

the average benefit $\tilde{\gamma}_j(P_1^*, P_2^*)$ enjoyed by the side- j marginal agents when the platform expands the participation on side i (net of the platform's matching cost σ) applies to a measure of agents equal to q_j^e under welfare maximization, whereas it applies to a measure of agents of size q_j^* under profit maximization.

In general, signing the net effect of the interaction among the above four distortions is difficult. One can identify markets in which the profit-maximizing prices are higher than their efficient counterparts on one side but lower on the other, as well as markets in which they are higher on both sides (see Tan and Wright, 2018 for examples).

To obtain further insights, Gomes and Pavan (2019) suggest to express prices on a *per-unit* basis, i.e., by normalizing them by the size of the participating population on the opposite side of the market. To see this, assume that the costs of getting more agents on board on each side is equal to zero, that is, $c_1 = c_2 = 0$. Furthermore, and without loss of optimality (both in the case of profit and in the case of welfare maximization), assume that the participation fees are equal to zero (i.e., $p_i^* = p_i^e$, $i = 1, 2$) so that the total payments can be expressed entirely in terms of transaction fees, with $t_i^* \equiv P_i^*/q_j^*$ and $t_i^e \equiv P_i^e/q_j^e$, $i, j = 1, 2$, $j \neq i$.¹⁶ We then have that

$$\begin{aligned}
 t_i^* - t_i^e &= \underbrace{\frac{1}{q_j} \eta_i(P_i^*; q_j^*)}_{\text{markup}} + \underbrace{(\tilde{\gamma}_j(P_1^e, P_2^e) - \tilde{\gamma}_j(P_1^e, P_2^*))}_{\text{Spence distortion}} \\
 &\quad + \underbrace{(\tilde{\gamma}_j(P_1^e, P_2^e) - \tilde{\gamma}_j(P_1^*, P_2^*))}_{\text{displacement distortion}}. \tag{8}
 \end{aligned}$$

Note that both the displacement and the Spence distortions are equal to zero in the Armstrong (2006) model (due to the fact that the interaction benefits are homogeneous, on both sides). In this case, the efficient transaction fees are unambiguously lower under welfare maximization than under profit maximization, on both sides of the market: $t_i^e < t_i^*$, $i = 1, 2$.

When, instead, preferences are as in the Rochet and Tirole (2003) model, one can show that the Spence distortion is always positive on both sides, whereas the displacement distortion is negative on at least one side. While, in general, transaction fees can be either smaller or larger under welfare maximization than under profit maximization on either side of the market, the sum (across sides) of the transaction fees is always higher under profit maximization than under welfare maximization: $t_1^* + t_2^* > t_1^e + t_2^e$. The possibility that prices are lower under profit maximization than under welfare maximization, though, is a consequence of the impossibility for the platform to engage in second-degree price discrimination. As we show in Section 6, when platforms discriminate among agents from the same side of the market, prices are always higher under profit than under welfare maximization (see also Gomes and Pavan, 2016).

¹⁶Clearly, the results below extend to arbitrary payments P_i by interpreting t_i as the “per-unit” payment (i.e., the payment per transaction).

2.5 Miscellaneous: chicken & egg problem, non-negative prices, distortionary taxation

2.5.1 Chicken & egg

The analysis so far assumes that either there exists a unique demand configuration compatible with the proposed prices, or the monopoly can choose its preferred price-quantity pair on the demand schedule. This last assumption may not be appropriate in many markets. When the monopolist cannot pick its preferred allocation, a network may fail to launch despite the long-run efficiency of the activity (Rohlfs, 1974). In a nutshell, if all consumers have a negative stand-alone utility (joining is costly), then for any positive price, there is an equilibrium market allocation where no consumer joins, despite the fact that some positive participation would be efficient and profitable. In the case of platforms, this possibility translates into the following chicken & egg problem (Caillaud and Jullien, 2003)¹⁷ Consider a new platform with no users on board and suppose that the stand-alone utility is negative for all users on both sides, that is, $v_i \leq 0$ for all consumers, $i = 1, 2$. To induce the side- i agents to join, the platform must convince them that the side- j agents will join as well. In the worst-case scenario (see the discussion of consumers' pessimistic beliefs in Section 3), given any pair of positive prices p_1 and p_2 , no consumer joins from either side. To overcome such pessimistic beliefs and successfully launch, the platform can subsidize one side, through monetary payments or by offering additional services that raise the stand-alone values and make them positive.

Notice that the ability to charge different prices on the two sides is key to overcoming the chicken & egg problem, an important difference with respect to standard one-sided markets with network effects. To illustrate this last point, suppose that sides are symmetric, with demand equal to $D(p_i; q_j)$ on each side $i = 1, 2$ and let \bar{v} be the smallest price such that $0 = D(\bar{v}; 0)$. Consider a firm selling at a uniform price a good with one-sided network effects such that demand Q satisfies $Q = 2D(p; Q/2)$ — the demand thus coincides with the total symmetric demand of the platform at symmetric prices p . Because, for all $p \geq \bar{v}$, the demand is equal to zero, under the worst-case scenario, the one-sided platform cannot be active if its cost is above \bar{v} . By contrast, the platform can set prices such that $p_1 < \bar{v}$ and $\bar{v} < p_2 < \bar{v} + \gamma D(p_1; 0)$ which ensures positive sales on both sides (because $q_1 \geq D(p_1; 0) > 0$). The profit may then be positive if the margin on side 2 is large enough. Notice that a similar strategy could be adopted by the one-sided platform if it could price-discriminate.

2.5.2 Vertical integration

When subsidies are not possible, an alternative way to overcome the chicken & egg problem and convince consumers on one side of the market to join is for the platform to integrate vertically.¹⁸ Hagiu and Spulber (2013) argue that a platform mediating the interactions between content suppliers and consumers may offer some integrated content to appear

¹⁷See Evans and Schmalensee (2010) for a dynamic version emphasizing the need to reach a critical mass to succeed.

¹⁸Like any firm, platforms may have other efficiency motivations for vertical integration, such as eliminating double marginalization, creating investment incentives, and promoting diversity and stability of supply (see, for example, Perry, 1989, and Rey and Tirole, 2007).

more attractive to consumers.¹⁹ Offering content in addition to matching services allows the platform to raise the value that buyers derive from joining and secure their participation. However, it comes at the cost of creating congestion on the seller side if integration reduces the expected profit per consumer of non-integrated content suppliers.

Miao (2009) and De Cornière and Taylor (2014) study integration in complementary services used as input for the interactions between the two sides and point to the risk of foreclosure of competing services when the platform cannot appropriate all profits from interactions.²⁰

A recent empirical literature studies platform integration into the provision of complementary services, such as Google producing apps for Android (Wen and Zhu, 2019) and Amazon selling products under its own brand name (Zhu and Liu, 2016). This literature finds that the platform is attracted to product categories with relatively high sales and that the threat of entry leads to reduced innovation and investment on the part of producers of complementary products. These results suggest that platforms are not entirely driven by the goal of ensuring the consumer experience, at least in these cases. In contrast, Lee (2013) finds that producers of video game consoles sign exclusive deals with game developers in order to attract consumers and this practice particularly contributes to the success of smaller console producers.

2.5.3 Non-negative prices

When, on one side of the market, the marginal cost is small (which is often the case for digital platforms), the demand is very elastic, and/or the interaction benefits on the other side of the market are large, the platform may find it optimal to set a negative price on the side satisfying the aforementioned properties. Several contributions discuss why negative prices may attract opportunistic demand from users that are of no value to the platform, who join without any serious intent to interact with the other side. When this is the case, the lowest price that a platform may set is often zero.²¹ Alternatively, a platform may optimally set the price on one side to zero if it faces transaction costs in collecting monetary payments from that side. Examples are free-to-the-air TV and street markets.²²

Amelio and Jullien (2012) point out that, when the unconstrained monopoly price is negative, constraining the price to be non-negative has a detrimental effect not only on the platform but also on consumers. The reason is that the motivation for setting a negative price on side i is to boost side- i participation so as to raise the value perceived by the side- j users. As for any quality, as long as the pass-through of the interaction benefits into prices is less than one, the net utility of the side- j consumers increases when the platform is allowed to set a negative price on side i and so does the side- j participation. Consumers on side i then benefit from a larger participation on side j and a larger subsidy. By contrast, when a zero-price constraint is exogenously imposed, say because of high transaction costs, the

¹⁹See also Nocke, Peitz, and Stahl (2007).

²⁰Illustrating the arguments in these papers requires richer models that go beyond what is assumed in this chapter.

²¹Gans (2019) provides a formal justification of non-negative prices, based on a kinked demand curve at a zero price due to free disposal of the platform's good.

²²We thank Martin Peitz for suggesting these examples to us.

welfare effect on consumer surplus of a zero price on one side is ambiguous—as the optimal price could be positive or negative without the constraint.

Consider the optimal subscription price on side 2 when the platform is constrained to set a non-negative price on side 1 and the constraint binds. Profits in this case can be written as

$$\Pi_2 = (p_2 - c_2) D_2(p_2; q_1) - \sigma q_1 D_2(p_2; q_1) - c_1 q_1$$

with $q_1 = D_1(0; D_2(p_2; q_1))$. To ensure uniqueness and positivity of q_1 , assume that $\Delta \equiv \frac{\partial D_1}{\partial q_2} \frac{\partial D_2}{\partial q_1} < 1$ and that $D_1(0; 0) > 0$. Assuming also positive profits, we have that the first-order condition for the optimal price p_2^0 is given by (Here we use superscripts “0” to highlight the zero-lower-bound on the side-1 prices; We also use Δ^0 as a shortcut for Δ evaluated at p_2^0 and q_1^0):

$$p_2^0 - \left[c_2 + \sigma q_1^0 + (\sigma q_2^0 + c_1) \frac{\partial D_1(0; q_2^0)}{\partial q_2} \right] = (1 - \Delta^0) \eta_2(p_2^0; q_1^0). \quad (9)$$

Comparing these conditions to their counterparts in the absence of the zero-lower-bound (see Condition (2) above), we can highlight two differences (in addition to the change in the quantity q_1).²³ First, as the price p_1 is not adjusted to compensate for the change in q_2 and maintain the side-1 demand constant, the relevant opportunity cost of getting more agents on board from side 2 is larger: the extra benefit $\tilde{\gamma}_1 q_1$ on side 1 is now replaced by a cost $(\sigma q_2^0 + c_1) \frac{\partial D_1}{\partial q_2}$ corresponding to the increase in q_1 . Second, the demand on side 2 is more elastic at the constant price p_1 than at the constant sales q_1 , which is reflected in the term $1 - \Delta$. The global effect on the side-2 price of imposing a zero-lower bound on the side-1 price is thus a priori ambiguous.

For an illustration, consider a market with linear interior demands $D_i(p_i; q_j) = a_i - p_i + \gamma_i q_j$ and zero interaction costs, $\sigma = 0$. In this case, the optimal prices with and without the zero lower bound are related by $p_2^0 = p_2^* - \gamma_2 p_1^*$, $q_2^0 = q_2^*$ and $q_1^0 = q_1^* + p_1^*$.²⁴ Hence all prices increase and quantities decrease (weakly) when $p_1^* < 0$ whereas the reverse is true when $p_1^* > 0$.

²³Note that a similar condition holds for optimal prices in the absence of the zero-lower-bound. Such a condition can be obtained by maximizing profits by means of p_2 , holding p_1^* constant:

$$p_2^* - \left[c_2 + \sigma q_1^* + (\sigma q_2^* + c_1 - p_1^*) \frac{\partial D_1(p_1^*; q_2^*)}{\partial q_2} \right] = (1 - \Delta^*) \eta_2(p_2^*; q_1^*).$$

²⁴From the previous footnote, the condition

$$p_2 - [c_2 + (c_1 - p_1) \gamma_2] = (1 - \gamma_1 \gamma_2) q_2$$

holds both in the presence of the zero-lower-bound and in its absence. Treating p_1 as a parameter varying between 0 and p_1^* and differentiating yields

$$\frac{\partial p_2}{\partial p_1} = -\gamma_2; \quad \frac{\partial q_2}{\partial p_1} = 0; \quad \frac{\partial q_1}{\partial p_1} = -1.$$

Welfare with a zero-price on the side-1 price is maximal when

$$p_2 - \left[c_2 + \sigma q_1 + (\sigma q_2 + c_1) \frac{\partial D_1(0; q_2)}{\partial q_2} \right] = - \left[\bar{\gamma}_1(0, p_2) q_1 + \bar{\gamma}_2(0, p_2) q_2 \frac{\partial D_1(0; q_2)}{\partial q_2} \right],$$

where $\bar{\gamma}_j(0, p_2)$ is the average interaction benefit experienced by the side- j agents, as defined in (6). Because the right hand side is negative whereas it is positive in (9), this suggests that the profit-maximizing price p_2^0 is excessively high compared to its welfare-maximizing level.

2.5.4 Distortionary taxation and two-part tariffs

The demand representation in (1) assumes that the value that each agent from each side $i = 1, 2$ assigns to interacting with each agent from the other side does not depend on prices. As a result, tariffs affect “participation” but not “usage” conditional on participation. When this is the case, the structure of the tariffs is not important for the monopoly, so that the first-order conditions in (2) apply to any two-part tariff $P_i = p_i + t_i q_j$, $i, j = 1, 2$, $i \neq j$, irrespective of how the latter is decomposed into a participation fee p_i and a transaction fee t_i . Rochet and Tirole (2006) also show that their analysis of optimal transaction fees can be adjusted to the case of two-part tariffs.²⁵

However, in general, taxation of usage has a distortionary effect and in this case the choice of transaction fees matters for efficiency. For instance, consider the following situations.

- The interaction benefit s_i for a side- i consumer is random and learned only after the agent chooses to participate. In this case, a type θ_i participating trades only if the realization of s_i is above t_i .
- The interacting parties trade a variable quantity x at some price τ that are determined ex-post (i.e., after joining the platform) through efficient Nash bargaining or a posted price mechanism. In this case, the interaction surplus and volume traded depends on the total per-unit tax $t = t_1 + t_2$.²⁶

Under these circumstances, the composition of the tariffs in terms of participation fees p_i and transaction fees t_i matters for optimality. For instance, Bedre-Defolie and Calvano (2013) extend the credit-card model of Rochet and Tirole (2002) by considering two-part tariffs on both sides with random interaction benefits. In their model, participating merchants do not decide on the payment mode. As a consequence of the asymmetry between merchants and consumers, they find that the transaction fees are inefficiently distorted (from a total-welfare perspective) toward lower fees for consumers.²⁷

²⁵They show that their formula (4) applies to normalized transaction fees $\hat{t}_i = (p_i - c_i)/q_j + t_i$, by replacing $(1 - F_i^\gamma(t_i))/f_i^\gamma(t_i)$ with $\eta_i(P_i; q_j)/q_j$, where $\eta_i(P_i; q_j)$ is a generic expression for the side- i inverse-semi-price elasticity.

²⁶This is the total payment to the platform per unit of output transferred from side 2 to side 1, within each match.

²⁷In their credit card model, this means that the interchange fee (the fee paid by the merchant’s bank to the consumer’s bank) is too large.

Notice that, when interaction benefits are ex-ante homogeneous, the model can be rewritten in reduced form by expressing interaction benefits as a function of transaction fees, $\gamma_i \equiv \gamma_i(t_1, t_2)$, and introducing a measure of the expected fee paid by the side- i agent per interaction with each side- j agent, $\beta_i(t_1, t_2)$. Then the gross utility of each consumer joining from side i becomes $v_i + \gamma_i(\mathbf{t})q_j$, where $\mathbf{t} \equiv (t_1, t_2)$. The expected payment by each side- i agent to the platform is $p_i + \beta_i(\mathbf{t})q_j$. With this notation at hand, the demand on side i (i.e., the measure of side- i agents who join the platform) can be written as

$$q_i = 1 - F_i^v(p_i + (\beta_i(\mathbf{t}) - \gamma_i(\mathbf{t}))q_j).$$

The expected interaction cost to the platform $\sigma(\mathbf{t})$ must also account for the effect of \mathbf{t} on the probability of trade. One may then invert the demand to write the platform's total profits, using $p_i + \beta_i(\mathbf{t})q_j = (F_i^v)^{-1}(1 - q_i) + \gamma_i(\mathbf{t})q_j$, as

$$\sum_{i,j=1,2,j \neq i} \left((F_i^v)^{-1}(1 - q_i) - c_i \right) q_i + (\gamma_1(\mathbf{t}) + \gamma_2(\mathbf{t}) - \sigma(\mathbf{t})) q_1 q_2.$$

The first part is the profit net of trade surplus and takes a standard form. The second part is the total trade surplus generated by the platform. With this formulation we see that if the platform is free to set any participation fees $\mathbf{p} = (p_1, p_2)$ and any transaction fees $\mathbf{t} = (t_1, t_2)$, then it chooses the transaction fees that maximize the total interaction benefits between participants $\gamma_1(\mathbf{t}) + \gamma_2(\mathbf{t}) - \sigma(\mathbf{t})$. Participation fees can then be adjusted to induce the desired participation levels q_1 and q_2 .

The conclusion that transaction fees should aim at maximizing the surplus generated by the activity of all agents on board is in line with similar conclusions on two-part tariffs in other domains of IO (vertical wholesale contracts for instance). It hinges on the platform's ability to use participation fees p_i to extract surplus and on consumers being risk neutral (Jullien, 2012). If these conditions are not met, the transaction fees naturally affect rent extraction and risk sharing between the platform and consumers.

Moreover, with heterogeneous agents, a platform must account for the screening effect of different tariffs. In particular, if the trade benefits are ex-post heterogeneous and the agents have private information, then the platform should optimally use the transaction fees to implement a screening mechanism (Myerson and Satterthwaite, 1983). In this context, Halaburda and Yehezkel (2013) points to the interaction between ex-post screening and the chicken & egg problem, arguing that solving the chicken & egg problem implies distortions in the trade mechanism. In their model, the platform's users learn the trade benefits only after participating and, absent any chicken & egg issue, the platform would implement efficient trade. However, to solve the chicken & egg problem, the platform needs to subsidize participation of one side and charge only the other side. The platform then finds it profitable to bias the trade mechanism against the subsidized side, so as to raise the other side's surplus and capture it through the participation fee. The paper then shows that competition may exacerbate the trade distortion.

Hagiu (2006) argues that transaction fees matter when the platform cannot fully commit. In his model of a video-game platform, there are two stages. In a first stage, transaction fees are negotiated with independent game developers, who then invest in developing games. In a second stage, the platform sells the game console while developers sell their games²⁸.

²⁸As sellers choose the prices of games, interaction benefits are endogenous in this model.

Retail prices of the console and of the games are not contractible, and as the console and the games are complements, prices tend to be too high (Cournot, 1838). It would be efficient from an interaction surplus perspective—an interaction occurs when a consumer buys a game—to set low transactions fees, so as to avoid double marginalization on games. But then the platform would exploit its monopoly power by setting an excessively high price for the console and restricting sales. Increasing the transaction fees permits the platform to increase the expected revenue to the platform per gamer and this reduces its incentives to restrict sales. For a given mass of games available, a higher transaction fee implies more sales of consoles, which may attract more developers. Hence a high transaction fee works as a commitment device that helps the platform convincing the developers of its willingness to set low prices and high marketing efforts.

2.6 Dynamic pricing

Cabral (2019) offers a dynamic analysis of a monopoly platform with switching costs.²⁹ The model assumes that, at random times, consumers outside the platform decide whether to pay a subscription fee to join, and consumers inside the platform decide whether to renew their subscription to the platform. The platform is thus treated as a durable good with two-sided network externalities and random obsolescence. Consumers have fixed homogeneous interaction benefits, γ_i , heterogeneous and random stand-alone valuations (on each side, v_{it} is i.i.d. across agents and time t) and consumers are myopic³⁰—meaning that when deciding to pay the subscription fee, they assess the value of joining/staying at the platform based on current utility given prices and participation.³¹ Hence demand at time t depends on subscription prices (p_{1t}, p_{2t}) and the volume of consumers with active subscription (q_{1t}, q_{2t}). The platform price-discriminates between renewing insiders and joining outsiders and follows a Markov strategy with state variable (q_{1t}, q_{2t}).

The paper finds that, for small discount factors, the price on side i is increasing in q_{jt} and, under some regularity conditions, decreasing in q_{it} . At any time, the platform must find a balance between exploiting current market power (with high prices) and building the customer base (with low prices) so as to "harvest" in the future with higher prices. When q_{jt} increases (and thus side- i interaction benefits increase), and the firm is impatient, the market power effect dominates and the firm raises its side- i price. The effect of q_{it} on p_{it} is more specific to the dynamic setting, as in a static framework there would be no reason to link the price to the same-side participation, in the absence of within-side network externalities. The dynamic effect of a consumer on the participation of the other side, along with the complementarity between sides due to the two-sided network effects, induces a dynamic indirect network externality between consumers on the same side. Another interesting point is that, when consumers are homogeneous with non-random stand-alone valuations ($v_{it} \equiv 0$),

²⁹Peitz, Rady, and Trepper (2017) consider a continuous-time model of monopoly with no switching costs but imperfect information about demand, where the intertemporal leakage follows from learning and experimentation by the platform.

³⁰See Cabral, Salant, and Woroch (1999) for an analysis of one-sided dynamic monopoly pricing with forward-looking consumers.

³¹Each consumer thus also ignores the effect of her current decision on future participation levels. The analysis also rules out within-period coordination issues by assuming that, in any period, only one consumer is given the chance to move (as in Cabral, 2011, discussed in Section 4).

the trade-off between exploitation and building customer base disappears. In this case, the firm achieves monopoly profit with full participation by setting a price $p_{it} = \gamma_i q_{jt}$ at any date, which is the maximal price that any myopic consumer on side i is willing to pay to join at date t .

Simulations show that, when stand-alone values take on negative values, dynamics may have two absorbing states, one with no participation and one with full participation (this is in line with insights from Evans and Schmalensee, 2010, that a platform may fail to launch if it does not reach a critical mass). Results also suggest that heterogeneity in stand-alone values may help a successful launch of the platform (reaching the large participation absorbing state): as heterogeneity increases, more agents may be willing to join even if they expect little interactions on the platform, which helps initiating the process of participation.

The paper shows also that, in a pure intermediation model where $v_{it} < 0$ with probability 1, the initial subscription fees charged by a platform starting with zero participation are negative. Here, myopia plays a role similar to expectations in Caillaud and Jullien (2001, 2003)'s analysis of the role of subsidies as a means to overcome consumers' coordination problems. In both cases, those consumers expecting or contemplating no participation on the other side refuse to subscribe unless the fee is negative.³² The difference is that, in a dynamic setting, the subsidy vanishes and prices become positive as the customer base increases.

Peitz, Rady, and Trepper (2017)'s analysis of learning by experimentation also shows that a platform may benefit from lowering its prices at an early stage of development, relative to short-run profit maximization, but for a different reason. Lower prices speed up the learning process by raising the informativeness of the market signal. This effect is however countervailed by the fact that reducing the price on one side raises the marginal benefit of increasing the price on the other side (a "seesaw" effect that will be discussed in Section 4.3). Hence, in their model, early monopoly prices are below myopic level on one side but may be below or above the myopic level on the other side, depending on the nature of the uncertainty.

3 Competition for the market

We now consider markets in which multiple firms (the platforms) compete to match agents from various sides of the market. Specifically, suppose there are now two platforms $k = A, B$ and that each side- i agent has a type $\theta_i = (v_i^A, v_i^B, \gamma_i)$, where v_i^k is the agent's stand-alone valuation for firm k 's product. By joining platform k , the agent obtains a utility equal to

$$u_i^k(\theta_i, q_j^k) = v_i^k + \gamma_i q_j^k \tag{10}$$

where q_j^k is the measure of agents from side j joining platform k . Let $\mathbf{p}_i = (p_i^A, p_i^B)$ denote the side- i prices set by the two platforms and $\mathbf{q}_i = (q_i^A, q_i^B)$ the measure of agents joining

³²Note that, for an infant platform, consumers' myopic beliefs are disadvantageous. Indeed a forward-looking consumer expecting positive participation in the future may be willing to pay a positive price to join in the current period, if she fears that a long time will elapse before she is given again the opportunity to join. Indeed, in this case, not joining would result in foregone interaction benefits.

each of the two platforms on side i . The side- i demand for platform A is equal to

$$D_i^A(\mathbf{p}; \mathbf{q}_j) = \Pr(v_i^A + \gamma_i q_j^A - p_i^A \geq \max\{v_i^B + \gamma_i q_j^B - p_i^B, 0\}).$$

To simplify the formulas, in this section we ignore the interaction costs by assuming that $\sigma^k = 0$, $k = A, B$, and then denote by c_i^k the cost that firm k incurs to bring agents on-board on side $i = 1, 2$.

A key driver of the nature of competition in such markets is the magnitude of the interaction benefits vis-a-vis the degree of horizontal differentiation between the two platforms' products, as measured by the differential in the agents' stand-alone values. When the interaction benefits are small relative to the degree of differentiation between the two platforms' products, the two platforms can share the market. When, instead, the interaction benefits are large, a single platform is likely to dominate the entire market. In this section, we focus on the case of large interaction benefits. The case where interaction benefits are relatively small is examined in the next section.

3.1 Divide-and-conquer strategies

To understand how competition works in markets with large interaction benefits, Caillaud and Jullien (2001, 2003) consider a model in which an incumbent platform (firm A) and an entrant (firm B) set prices to protect and conquer the market, respectively. Their model is one in which platforms are pure matching intermediaries (agents derive no stand-alone value from the platforms' products) and where the agents' interaction benefits are homogeneous within sides and the same for both platforms (hence firm B is equally efficient as firm A). That is, there exists γ_1 and γ_2 such that F_i is a Dirac measure assigning probability one to $(0, 0, \gamma_i)$, $i = 1, 2$.

The paper analyzes a simultaneous-pricing game and characterizes the full set of equilibria (a continuum) under a mild regularity assumption on demand (referred to as monotonicity). In particular, the analysis does not assume any incumbency advantage. In this section, we discuss only the case where all agents single-home (i.e., choose at most one of the two platforms). The case where agents multi-home (i.e., use the two platforms) will be addressed in Section 3.3.

Caillaud and Jullien (2001) show that an incumbent platform may monopolize the market and make positive profits when facing an equally efficient competitor.³³ Because the equilibrium outcome yielding the highest profit to the incumbent firm coincides with the equilibrium outcome of a Stackelberg game and because arguments are easier to expose in this alternative game, hereafter we consider a Stackelberg game whereby firm A sets prices before firm B does.

In this simple setup, when the price differential $|p_i^A - p_i^B|$ between platforms is small on each side $i = 1, 2$, both tipping³⁴ on firm A (that is, all agents joining platform A) and tipping on firm B (all agents joining platform B) can be sustained in equilibrium. The multiplicity of equilibria in the network-formation game (that is, in the continuation game among the agents, for given access prices set by the firms) translates into a multiplicity of

³³This equilibrium is not unique; another equilibrium exists in which the competitor wins.

³⁴In the chapter, we say that there is tipping when all consumers join the same platform.

equilibria in the price-posting game. Whether or not the incumbent firm can have positive profits in equilibrium then depends on the extent of the coordination failure between the agents. If agents coordinate efficiently on joining the platform that can deliver them the highest surplus,³⁵ then a standard undercutting logic implies that no platform can enjoy strictly positive profits in equilibrium. By contrast, if firm B faces “unfavorable beliefs,” the incumbent firm may benefit from the consumers’ coordination failure and enjoy strictly positive profits. The notion of “unfavorable beliefs” in Caillaud and Jullien (2003) captures precisely the idea that consumers’ coordination failure disadvantages a firm.³⁶

Unfavorable beliefs A firm faces unfavorable beliefs at prices (p_1, p_2) if, in the continuation game that follows the selection of the prices (p_1, p_2) , consumers coordinate on the equilibrium allocation that minimizes the firm’s market share on both sides.³⁷

When B faces unfavorable beliefs, we say that platform A is *focal*. Suppose that initially each agent expects all other agents to join platform A . Then an equally efficient platform B facing unfavorable beliefs would have no demand if it sets positive prices. This is because any agent expects no match on platform B and such beliefs are self-enforcing. However, firm B could still attract both sides by convincing one of the two sides to join with a subsidy and charging a positive price on the other side. Such strategies are referred to in the literature as *divide-and-conquer* strategies.

Divide-and-conquer strategies: Strategies that set prices below cost for some consumers compensated with prices above cost on other consumers.

Divide-and-conquer strategies have been shown to be an effective instrument to steal consumers in markets with increasing return to scale (Innes and Sexton, 1994). When the entrant is expected to use such strategies, the incumbent cannot corner the market at a positive price if the entrant is equally or more efficient. However it can protect itself, preventing entry, by subsidizing one side.

Divide-and-conquer strategies in two-sided markets mirror introductory pricing in dynamic models of competition with interaction benefits, where a firm subsidizes early purchases to induce a bandwagon effect and recover the subsidy on late-comers (Farrell and Saloner, 1986, Katz and Shapiro, 1986, 1992). With introductory pricing, a firm subsidizes demand in one period and recovers the loss in future periods. In two-sided markets, the loss on one side is recovered on the other side within the same period. In some markets, there is a natural ordering, with one side joining after the other (see for instance Hagiu, 2006, on video games). In other cases, instead, there is no natural order and either side can be targeted with a subsidy. In this respect, divide-and-conquer strategies are more flexible than sponsoring strategies in dynamic setting.

³⁵That is, no agent joins platform k when all agents joining platform $-k$ (from both sides) yields a Pareto improvement vis-a-vis the situation where all agents (from both sides) join platform k .

³⁶See also Hagiu (2006), Halaburda and Yehezkel (2016) and the discussion in Halaburda, Jullien, and Yehezkel (2020).

³⁷Such an allocation exists because the network formation game between consumers is of a super-modular nature. Unfavorable beliefs amounts to assuming that consumers coordinate on a “minimal” equilibrium of the network-formation game, which always exists for supermodular games (Topkis, 1998). Unfortunately, this existence result does not immediately extend to markets with more than 2 networks.

When the entrant faces unfavorable beliefs, Caillaud and Jullien (2001, 2003) show that the incumbent can indeed obtain positive profits. This is achieved by setting a negative price on one side and a positive price on the other side. These prices must satisfy two conditions. First, the subsidy must be such that the cost for B of dividing the subsidized side is larger than the surplus from network effects on the other side. Second, the positive price must be low enough that dividing the non-subsidized side and competing with A 's subsidy on the subsidized side is not profitable. Interestingly, firm A can find such prices and make profits only if (i) both sides derive positive interaction benefits with the other side and (ii) there is an asymmetry in interaction benefits between the two sides. Firm A 's optimal strategy then subsidizes the side with less interaction benefits.

The above results rely on several modeling assumptions that were already discussed in the monopoly case.

Negative prices: If subsidies are limited then an efficient firm facing unfavorable beliefs may not be able to enter the market. To see this, suppose that firm A sets prices $0 \leq p_i^A \leq \gamma_i$ on both sides. Then firm B cannot sell when it faces unfavorable beliefs and negative prices are not feasible. Indeed, for any pair of non-negative prices by firm B , the network-formation game admits an equilibrium in which all consumers join only firm A . In this case, an equilibrium exists where firm A obtains monopoly profits, unless firm B finds a way to turn the agents' coordination in its favor.

Transaction fees: The positive-profit result discussed above relies on the firms offering simple prices that are not contingent on the participation of the other side. For instance, Caillaud and Jullien (2003) show that if the two firms can use non-distortionary transaction fees, then the equilibrium is efficient—i.e., all consumers buy from the same (most efficient) firm—and the active platform obtains zero profit if both platforms generate the same interaction benefits.

The analysis in Caillaud and Jullien (2001, 2003) focuses on pure matching intermediaries, where agents derive value only from the platforms' matching services. Jullien (2011) points to the importance of the relative magnitude of interaction benefits and stand-alone values for the nature of platform competition. The paper contrasts the case of pure intermediation with no stand-alone values and large interaction benefits to the case of platforms offering services generating large stand-alone values and relatively small interaction benefits. In the latter case, an incumbent firm cannot secure positive profits when facing a competitor at least as efficient as the incumbent. To see that, suppose that agents' stand-alone values are homogeneous within sides and across platforms with $v_i^k = v_i \gg \gamma_i$, $i = 1, 2$, $k = A, B$. Suppose that firm A sets prices $p_1^A < 0 < p_2^A$. Then firm B can undercut firm A on side 1 with price $p_1^B = p_1^A - \gamma_1$ and charge a premium on side 2 with price $p_2^B = p_2^A + \gamma_2$, generating a revenue $p_1^A + p_2^A + \gamma_2 - \gamma_1$ larger than A 's revenue.³⁸ The paper shows that, in the presence of multi-sided network effects, a Stackelberg leader loses the market when faced with an

³⁸By contrast, with a pure intermediation service, $v_i = 0$, $i = 1, 2$, dividing side 1 would allow to charge only γ_2 on side 2 because consumers' reservation value is zero on this side if side 1 joins B . Firm B 's revenue would then be only $p_1^A + \gamma_2 - \gamma_1$ which may be negative. This underlies the ability of firm A to generate positive profit with a positive price p_2 in this case.

equally efficient or even slightly less efficient competitor. When, instead, the leader is more efficient on some side, it may be able to win the market by choosing an adequate pattern of subsidies and taxes to the various sides.

3.2 Congestion within sides

The canonical model assumes that there are no network externalities within sides. Such externalities may originate from congestion effects or competition between agents on the same side. Belleflamme and Toulemonde (2009) show that intra-group negative externalities expand the set of divide-and-conquer strategies, by allowing platforms to target a fraction of agents from the congested side. This is because migration from A to B becomes less attractive as congestion is reduced at A . To illustrate, consider again a focal platform A and an entrant B . Assume that $v_i^k = 0$, $i = 1, 2, k = A, B$, and let the interaction benefits be homogeneous at γ_1 on side 1. Suppose that agents from side 2 are exposed to congestion. This can be captured by assuming that the side-2 interaction benefits depend on the measure q_2 of the side-2 participating agents, with the function $\tilde{\gamma}_2(q_2)$ decreasing in q_2 . When all agents join A , the side-1 consumers experience an utility equal to $\gamma_1 - p_1^A$ whereas the side-2 agents experience utility $\tilde{\gamma}_2(1) - p_2^A$. Faced with prices $p_1^A < 0$ and $p_2^A > 0$, platform B can now attract all consumers from both sides with the following strategy: choose $m_2 \in [0, 1]$ and set prices (slightly below) $p_2^B = p_2^A - \tilde{\gamma}_2(m_2)$ and $p_1^B = p_1^A + \gamma_1(1 - 2m_2)$. First notice that, at these prices, in any equilibrium demand configuration, at least $1 - m_2$ consumers from side 2 join B , because the (negative) price p_2^B makes B more attractive than A to the side-2 agents if congestion on platform A exceeds m_2 , irrespective of the side-1 participation. This in turn implies that all consumers on side 1 join B because p_1^B is such that they strictly prefer joining B than joining A when the former platform attracts strictly more than $1 - m_2$ side-2 agents and the latter platform attracts strictly less than m_2 side-2 agents. Finally, given that all side-1 agents join B , it must also be the case that all side-2 agents join B because no consumer on side 2 is willing to pay the positive price p_2^A for no interaction with the side-1 agents. Hence platform B has some flexibility in the choice of the subsidy it offers to the side-2 agents. A higher subsidy on the congested side 2 allows platform B to charge a higher price on the non-congested side 1.³⁹

Negative within-side network externalities may also prevent tipping, as agents may prefer to avoid negative externalities by joining a smaller platform.⁴⁰ For instance, Karle, Peitz, and Reisinger (2020) develop a model of e-commerce platforms where sellers can multi-home but competition with other sellers reduces their profit, inducing negative externalities between sellers on the same platform. When seller competition is intense, sellers have an incentive to join different platforms to avoid head-to-head competition. In this case, market fragmentation emerges even if platforms offer identical services and efficiency would require all trade to take place on the same platform.

³⁹The other divide-and-conquer strategy, subsidizing side 1, is unchanged if B wants to sell to all agents, with $\tilde{\gamma}_2(1)$ being the relevant interaction benefit on side 2.

⁴⁰See Ellison and Fudenberg (2003) and Ellison, Fudenberg, and Möbius (2004) for an analysis of tipping and market fragmentation in markets with indirect and direct network externalities but no prices charged by the platforms.

These effects can be important empirically. Augereau, Greenstein, and Rysman (2006) study the adoption of 56K modems in which the market was originally served by two essentially identical but incompatible technologies. Despite the benefits to consumers from having the market coordinate on a single technology, Augereau, Greenstein, and Rysman (2006) show that ISPs in the same local market adopted different technologies, presumably in order to differentiate themselves. They argue that this led not only to market fragmentation, but to overall adoption failure of 56K modems until a standard setting organization intervened.

Halaburda, Piskorski, and Yıldırım (2018) argue that platforms may choose to restrict agents' choices in order to limit within-side negative externalities (they study restrictions on each agent's available choice set under inefficient multi-lateral bargaining in a one-to-one matching process).⁴¹ The degree of restriction is then a competitive tool. With heterogeneous populations on both sides, offering different levels of restrictions allows platforms to differentiate their services, which prevents tipping equilibria to emerge. (See also the discussion of Ambrus and Argenziano, 2009, in Section 4.1).

3.3 Multi-homing

The analysis so far assumes that agents who join one platform do not join the other (single-homing). The opposite situation where agents may opt to join more than one firm (multi-homing) leads to slightly different results. For the discussion, assume that each agent cares only about the mass of agents she can interact with on the other side and not about the number of channels for interactions. The conclusion that firm A can corner the market and make strictly positive profit still holds but the nature of the divide-and-conquer strategies is different. When agents can multi-home at no cost, it is easier to convince a consumer to try a new product—the divide part of the strategy. This is because a consumer who decides to join firm B does not have to forgo interacting with those agents from the other side who join only firm A . However, inducing a consumer to join firm B does not generate the same competitive effects as in the case of single-homing. In particular, it does not trigger the same strong bandwagon effect as in the case of single-homing. This is because attracting users to firm B does not necessarily imply poaching them from firm A and hence does not necessarily make firm A weaker. This makes the conquest of the side where prices are positive less profitable. In the context of the simple example discussed above, when $p_j^A > 0$, firm B attract both sides by asking a price slightly below zero on side i and then setting a price slightly below p_j^A on side j . The maximal profit for firm B is thus equal to $p_j^A - c_1 - c_2$. Firm A can however prevent firm B from profiting from such a strategy by setting prices $p_1^A = p_2^A = c_1 + c_2$ and still make positive profits equal to $c_1 + c_2$.⁴² Thus incumbent firms still enjoy an incumbency advantage under multi-homing but their profits are smaller than under single-homing. In particular, Caillaud and Jullien (2003) show that the maximal incumbent's profit is bounded from above by the total cost $c_1 + c_2$. The incumbency advantage is thus small when costs are low.

⁴¹Casadesus-Masanell and Halaburda (2014) also argue that a platform may limit supply if the goods exchanged on the platform exhibit network effects.

⁴²Note that the same prices also prevent firm B from attracting agents under single-homing.

Multi-homing also alters the way complex tariffs, involving taxation of interactions, can be used to conquer a market. The role of taxation of usage is analyzed by Caillaud and Jullien (2003) who allows agents to multi-home and firms to charge fees for both participation and usage. In the case of single-homing, an optimal strategy for an entrant facing unfavorable beliefs is to subsidize participation and heavily tax interactions when taxes are not distortionary. This strategy relies on the fact that, once one side is attracted, agents from the other side must also join if they want to benefit from the network effects. However, under multi-homing, even if a subsidized side is attracted, agents from the other side can continue to enjoy interaction benefits by staying with the incumbent firm. Hence if the entrant taxes interactions too heavily, all users may stay with the incumbent firm and benefit from lower transaction fees. Moreover, multi-homing creates scope for a new type of entry strategies. The entrant does not need to convince consumers to stop buying the incumbent's product. It needs only to convince them to stop interacting with the other side using the incumbent's platform and instead use its own platform. The aim is to induce all agents to multi-home and then compete on transaction fees. As a result, when agents can multi-home, divide-and-conquer strategies typically involve low transaction fees. By setting low transaction fees, an incumbent firm can also raise its protection against entry. The paper shows that, in the equilibrium maximizing the incumbent's profits, the incumbent sets zero transaction fees and relies solely on participation fees. Moreover, the incumbent makes zero profits when faced with an equally efficient entrant, unless the volume of interactions increases when all agents multi-home compared with the case where they single-home on the same platform.⁴³

3.4 Dynamic competition

Most of the existing literature on the dynamics of competition is related to the irreversibility of the choice of a network by consumers, which may be due to the durability of equipment or to switching costs (adoption cost, learning, data transfer). It has been known since Arthur (1989) that in the presence of network externalities and irreversibility, firms with a large installed base benefit from a competitive hedge that may prevent entry of more efficient competitors. As discussed in Section 7.2.4, a number of empirical papers use installed base as a variable to explain technology diffusion in markets with network effects. An implication of the strategic value of an installed base is that firms supplying network goods gain from quickly building clientele (see Farrell and Saloner, 1986, and Katz and Shapiro, 1986, 1992). This allows them to induce a bandwagon effect that facilitates sales to future consumers. The dynamics of competition is thus shaped by the intensity of network effects, with incentives to preempt and intense competition at the earlier stages of competition.

Under imperfect competition, firms competing in the presence of network effects have to find a balance between maximizing short-run profits and investing in building their

⁴³Their model allows for an imperfect matching technology such that a pair of agents on each side using two platforms obtains higher interaction benefits than if they use only one platform. In this case, entry with all agents multi-homing is efficient if costs are low enough (this is referred to as global multi-homing). For larger costs, instead, the possibility of global multi-homing makes entry less profitable by increasing the cost of strategies that rely on one side single-homing.

network.⁴⁴ In this context, an important question for policy is whether there is increasing dominance of the largest network eventually resulting in tipping. Almost all contributions on the dynamics of competition focus on one-side network effects and uniform pricing by firms.⁴⁵ We highlight here some recent contributions that are particularly relevant for two-sided markets.

In a computational duopoly model with product differentiation, network effects and dynamic stochastic demand, Cabral (2011) offers a new perspective on tipping by showing that the market may exhibit long periods of dominance of one firm, alternating with transition periods of intense competition.⁴⁶ The general idea is that a firm that has acquired a large customer base will be tempted to exploit its position by setting a high price. Such a strategy permits another firm to survive with a low price and few customers until a demand shift allows it to challenge the large firm.⁴⁷

With small network effects, product differentiation dominates the equilibrium dynamics and the market converges to a situation where the two firms share the market. When network effects are large, instead, the equilibrium features relatively long periods in which a larger firm shares the market with a smaller firm, ending when a sufficiently large shift in the market demand favors the small firm and a new cycle starts.

A key point in the analysis of Cabral (2011) is that the expected delay before the small firm reaches a position where it becomes a credible challenger is long enough that the incumbent does not find it worthy to eliminate its competitor. Things are different when markets switch fast from one dominant firm to another. Fudenberg and Tirole (2000) point out that a monopoly firm may inflate its sales by reducing prices to deter entry and preserve its dominance. They build a model with overlapping generations of consumers with switching costs and show that, in equilibrium, when firms are patient, an incumbent facing the threat of entry in every period may sustain permanent low prices so as to maintain a large installed base that dissuades entry. Indeed, an entrant’s profitability decreases if it has to compete against a firm with a large installed base. The outcome is that of a contestable market where monopoly is persistent but prices remain low. In this case, the price reflects the shadow value of sales which includes entry deterrence.

Halaburda, Jullien, and Yehezkel (2020) and Biglaiser and Crémer (2020) argue that dynamic leakage may not be related to switching costs but to consumer inertia.

Halaburda, Jullien, and Yehezkel (2020) relies on the concept of pessimistic beliefs and a focal platform. The paper considers two competing platforms offering vertically-differentiated goods with large network effects and no switching cost. In a dynamic context, the paper assumes that “focality” is endogenous and history-dependent. More precisely, it assumes that the firm that wins the market in one period gains an incumbency advantage

⁴⁴The literature on dynamic competition with myopic consumers but forward-looking firms includes, among others, Doganoglu (2003), Laussel, de Montmarin, and Van Long (2004), Mitchell and Skrzypacz (2006), Driskill (2007), Markovich (2008), Markovich and Moenius (2009) and Laussel and Resende (2014).

⁴⁵ See the review by Farrell and Klemperer (2007). A contribution with a dynamic model of two-sided market is Zhu and Iansiti (2012)

⁴⁶See also Chen, Doraszelski, and Harrington (2009).

⁴⁷As all other contributions described in this section, this result is shown in the context of a market with one-sided network effects and uniform price instead of a two-sided market. But the insights are clearly relevant also for the dynamics of two-sided markets.

(see also Biglaiser and Crémer, 2020) by becoming focal in the next period. The prospect of winning focality affects incentives in a somewhat similar way as the prospect of building an installed base in the case of switching costs (Farrell and Saloner, 1986, Katz and Shapiro, 1986, 1992): firms are willing to subsidize sales in early periods to benefit from the incumbency advantage in later periods. The paper shows that, with a finite but long horizon, the more efficient platform wins the market in the first and all subsequent periods. The reason is that the value attached by a firm to gaining focality is larger for the higher-quality firm⁴⁸ and the value differential increases with the horizon, exceeding, for large horizons, any incumbency advantage that may initially favor the lower-quality firm.

Hence, dynamic considerations weaken concerns that an incumbency advantage creates barriers to entry, because prospective future benefits of a superior-quality entrant are larger than the prospective benefits of the lower-quality incumbent.⁴⁹

This conclusion must be qualified when the horizon is infinite. For large discount factors, both efficient and inefficient equilibria exist. In the case of inefficient equilibria, one firm is always aggressive and successful in gaining or preserving focality whereas the other does not attempt to gain focality, for it expects tough competition in case of success. Due to the advantage attached to incumbency, aggressiveness or softness in competition are self-enforcing. To see this point, suppose that the least efficient firm A is expected to win the market in the future, irrespective of which firm will be the incumbent. Then the other firm B has no reason to price below cost, since it expects no future sales. To the contrary, firm A is always willing to make a sacrifice by pricing below cost to become focal. Indeed, although firm A expects to win in the future irrespective of its focality position in the next period, it strictly prefers to be focal in the next period, because this permits it to set a higher price. Firm A may then win the current period both when it is already focal and when it is not. With a large discount factor and infinite horizon, this effect is strong enough that an equilibrium exists in which A always wins. In other words, when firms are patient, a firm's "reputation" for being aggressive in defending and restoring its incumbency position after a downturn may dissuade other firms from challenging it.

Biglaiser and Crémer (2020) analyze the effect of demand heterogeneity on market fragmentation in a dynamic model with no switching costs but an incumbency advantage. Their modeling of the incumbency advantage is based on a fictitious network formation game. Their pseudo-dynamics start from the allocation of consumers of the last period and assumes that consumers change their demand sequentially to a strictly preferred (myopic) choice, where the first to move are those with the highest benefit to do so. In the case of homogeneous network effects, the dynamics are the same as in Halaburda, Jullien, and Yehezkel (2020). Their concept applies to multiple platforms, which they exploit to study market fragmentation. In a model of price competition between firms offering homogeneous (one-sided) network goods to a heterogeneous population, they show that, for small discount factors and large heterogeneity of interaction benefits, an inefficient equilib-

⁴⁸To see that, suppose platforms A and B offer homogeneous products to a unit-mass population, with stand-alone values v^A and v^B and network effect γ per user such that $\gamma > v^B - v^A > 0$. Then, in a one-period game, A sells to all consumers if it is focal, with profit $\gamma + v^A - v^B < \gamma$, whereas B sells to all consumers if it is focal, with profit $\gamma + v^B - v^A > \gamma$.

⁴⁹However, this argument requires the entrant to be forward looking and able to sustain large losses in current competition with the incumbent.

rium with two active networks exists (where tipping would be efficient). This conclusion follows from the same logic that prevents full tipping in models with switching costs such as Cabral (2011), except that, as in Halaburda, Jullien, and Yehezkel (2020), an installed base doesn’t raise value but affects the consumer coordination process. An incumbent with a large legacy network, and thus a large incumbency advantage, prefers to “harvest” its position by charging a high price and keep only consumers with high interaction benefits, which allows a second firm to attract unsatisfied customers with low interaction benefits. Obviously such an equilibrium is sustainable only if the second network is small.

4 Competition on the market

Four forces that can prevent tipping in two-sided markets are: platform differentiation, multi-homing, compatibility⁵⁰, and, as discussed in Section 3.2, congestion.⁵¹ Many platforms offer differentiated products and services in addition to their intermediation services. For example, smartphones provide communication services, integrated apps, and photo and video recording, in addition to permitting users to download the apps that have been developed for their operating system and which are available “on-demand” (the other side of the market). By differentiating their stand-alone offerings, platforms alleviate the competition with other platforms. If the stand-alone service provided by a smaller platform is valuable enough, then the platform may be able to survive even if it attracts only a few consumers. In this section, we discuss some of the key results in the literature on competition between platforms, when the market has stabilized with more than one firm serving both sides of the market.

4.1 Single-homing

We first consider the case where all users single-home on both sides. Single-homing may be due to physical constraints or be the result of contractual restrictions set by the platforms. (Armstrong and Wright 2007). In this section, we treat single-homing as a technical constraint.

4.1.1 Homogeneous interaction benefits

Armstrong (2006) and Rochet and Tirole (2006) study platform competition under high degrees of product differentiation. In particular, Armstrong (2006) considers a duopoly in which platforms sell a horizontally-differentiated consumption good delivering different stand-alone values to consumers in addition to two-sided interaction benefits (through access to the other side). The analysis assumes that (i) consumers on both sides single-home, (ii) the market is covered (meaning that each agent from each side $i = 1, 2$ who does not join

⁵⁰We use the term compatibility for technologies that allow users of one platform to interact with users of other platforms. Examples include email and the connectivity of telecommunication networks.

⁵¹Another force against tipping may be dynamic pricing. Mitchell and Skrzypacz (2006) study a dynamic duopoly model of direct network effects (not a two-sided market) and identify settings in which a firm with a large lead raises its price to harvest its installed base instead of lowering its price and further reduce its rival’s market share.

platform k joins platform $-k$) and (iii) interaction benefits are homogeneous within each side.

Preferences take the additive form in (10) with γ_i identical for any pair of agents from the same side. Interaction costs are zero (i.e., $\sigma_i = 0$, $i = 1, 2$) and prices take the form of access fees so that $P_i^k = p_i^k$, $i = 1, 2$, $k = A, B$.

Platform differentiation on each side $i = 1, 2$ is captured by the dispersion of the stand-alone value differential $v_i^B - v_i^A$. When there is little differentiation, coordination failures may arise, and multiple allocations of consumers over the two platforms can be sustained in equilibrium. In this case, insights similar to those from the analysis of competition for the market apply. When, instead, the dispersion of stand-alone values is sufficiently large, a unique “stable” equilibrium allocation exists for *any* price vector.⁵² In this case, the demand for firm A is defined by

$$D_i^A(\mathbf{p}_i; \mathbf{q}_j) = 1 - D_i^B(\mathbf{p}_i; \mathbf{q}_j) = \Phi_i(p_i^B - p_i^A - \gamma_i(q_j^B - q_j^A)), \quad (11)$$

where Φ_i is the side- i cumulative distribution function of the differential in stand-alone values $v_i^B - v_i^A$, assumed to be smooth on a compact support, $\mathbf{p}_i = (p_i^A, p_i^B)$ is the vector of side- i prices, and $\mathbf{q}_j = (q_j^A, q_j^B)$ is the vector of side- j quantities (equivalently, of side- j participations).

Specifically, Armstrong (2006) assumes that stand-alone values take the Hotelling form $v_i^A = v - \tau_i x_i$ and $v_i^B = v - \tau_i(1 - x_i)$ where v is large enough to guarantee that the entire market is served in equilibrium, and where x_i is drawn from a uniform distribution over $[0, 1]$. In this case, the condition for uniqueness of the allocation of agents over the two platforms for all price vectors reduces to $\gamma_1 \gamma_2 < \tau_1 \tau_2$.⁵³

Under this condition, given any collection of prices $\mathbf{p}_1 = (p_1^A, p_1^B)$ and $\mathbf{p}_2 = (p_2^A, p_2^B)$, each platform’s profits are equal to

$$\Pi^k = \sum_{i=1,2} (p_i^k - c_i^k) q_i^k$$

where

$$q_i^k = D_i^k(\mathbf{p}_i; \mathbf{q}_j).$$

It is known from the earlier literature on one-sided markets that network effects can raise the elasticity of the demand by magnifying the demand response to a change in prices, as consumers that would not react to a price increase alone may react to the induced reduction of interaction benefits. The same logic applies to the elasticity of the residual demand of each platform in a two-sided market. Moreover, a higher elastic demand on one side means that increasing price on this side induces a higher loss of sales on both sides of the market.

⁵²The assumption of large differentiation bears some similarity with the role of heterogeneous beliefs in global games (see Jullien and Pavan, 2019 for a discussion).

⁵³To see this, consider an arbitrary allocation of consumers over the two platforms covering the entire market $(q_i^k)_{i=1,2,k=A,B}$ and assume that a small mass ε of side- i consumers moves from platform A to B . Then, the value of platform B relative to that of platform A for the side- j agents increases by $2\gamma_j \varepsilon$. This induces $\Phi_j' 2\gamma_j \varepsilon$ side- j agents to switch from A to B . By the same reasoning, this change in the side- j participation in turn induces an additional mass of agents from side i of measure $\Phi_i' 2\gamma_i (\Phi_j' 2\gamma_j) \varepsilon$ to migrate from A to B . The adjustment process converges if $4\Phi_i' \Phi_j' \gamma_i \gamma_j < 1$. In the Hotelling model, this condition reduces to $\gamma_1 \gamma_2 < \tau_1 \tau_2$.

The equilibrium pricing formulas thus reflect the enhanced price elasticity as discussed in the monopoly case, but accounting for the presence of a competing firm. To derive intuition, consider a small price change by firm k consisting in a reduction of its side- i price along with an increase in the firm's side- j price that maintains the platform's side- j demand constant at q_j^k . When the platform adds ε new consumers on side i , it can increase its side- j price p_j^k by $2\gamma_j\varepsilon$. This is because, in a covered market, the ε consumers that join the platform come from the competing firm. Attracting such consumers thus also means reducing the attractiveness of the other platform. That is, stealing a consumer from the competitor on side i not only raises the platform's value for those agents on the other side of the market who join the same platform but also reduces the value of joining the competing firm. By attracting ε more consumers from side i the platform can thus increase its side- j revenues by $2\gamma_j\varepsilon q_j^k$ without expanding the side- j participation and hence without incurring any extra cost. Holding constant the participation on side j at q_j^k , this means that the total opportunity cost of attracting ε more consumers on side i is equal to $(c_i^k - 2\gamma_j q_j^k) \varepsilon$. Once such opportunity costs are incorporated into the analysis, the optimality condition for equilibrium prices can be derived as in standard oligopoly games.

Under the above specification, equilibrium prices of an interior equilibrium take the form

$$p_i^k - [c_i^k - 2\gamma_j q_j^k] = \mu_i^k(\mathbf{p}_i; \mathbf{q}_j) \quad (12)$$

where

$$\mu_i^k(\mathbf{p}_i; \mathbf{q}_j) = -\frac{D_i^k(\mathbf{p}_i; \mathbf{q}_j)}{\frac{\partial D_i^k(\mathbf{p}_i; \mathbf{q}_j)}{\partial p_i}} \quad (13)$$

is the inverse semi-elasticity of firm k 's demand on side i , holding the participation of side j fixed at the level specified by $\mathbf{q}_j = (q_j^A, q_j^B)$ and fixing the other firm's prices at the level specified in $\mathbf{p}_i = (p_i^A, p_i^B)$.⁵⁴ In the case of the symmetric Hotelling model of Armstrong (2006), the equilibrium quantities on the other side are equal to $q_j^A = q_j^B = 1/2$ and the inverse-semi-elasticity of each firm's residual demand at the symmetric allocation is equal to $\mu_i^k(\mathbf{p}_i; \mathbf{q}_j) = \tau_i$. The equilibrium prices on each side are thus equal to⁵⁵

$$p_i^A = p_i^B = c_i + \tau_i - \gamma_j.$$

Extending these pricing formula to general oligopoly games turns out to be complex for two reasons. First, computing the opportunity cost is complicated if the intensity of the interaction benefits γ_j is heterogenous across agents and/or if there are more than 2 firms. Second, computing the price elasticity of the residual demands at constant opportunity cost may also be challenging, for neutralizing the effects of a change in the side- i demand on all the remaining sides is more involved. However, the idea of deriving the equilibrium prices on each side by holding quantities constant on the other sides has proved useful in richer

⁵⁴More precisely $\mu_i^A(\mathbf{p}_i; \mathbf{q}_j) = \frac{\Phi_i(p_i^B - p_i^A - \gamma_i(q_j^B - q_j^A))}{\Phi_i'(p_i^B - p_i^A - \gamma_i(q_j^B - q_j^A))}$ and $\mu_i^B(\mathbf{p}_i; \mathbf{q}_j) = \frac{1 - \Phi_i(p_i^B - p_i^A - \gamma_i(q_j^B - q_j^A))}{\Phi_i'(p_i^B - p_i^A - \gamma_i(q_j^B - q_j^A))}$.

⁵⁵Existence of the equilibrium requires that $\tau_1\tau_2 > \gamma_1\gamma_2 + \frac{1}{2}(\gamma_1 - \gamma_2)^2$ which is stronger than the previous condition for demand stability. Weyl and White (2016) show that, in the symmetric case with $\gamma_i = \gamma$ and $\tau_i = \tau$, there is a unique symmetric equilibrium if $\gamma < \tau$ while only tipping equilibria exist if $\gamma > \tau$.

settings too. The key assumption in the derivation of the equilibrium prices is that *the market is fully covered*.

For instance, when the demand on each side takes the general form $q_i^k = D_i^k(p_i^k, p_i^{-k}; q_j^k, q_j^{-k})$, changing price on side i while adjusting the price on side j so as to maintain the side- j participation constant and using the full-market-coverage condition $q_j^{-k} = 1 - q_i^k$ yields the following pricing formula on each side of the market and for each platform

$$p_i^k - c_i^k - q_j^k \left(\frac{\frac{\partial D_j^k}{\partial q_i^k} - \frac{\partial D_j^k}{\partial q_i^{-k}}}{\frac{\partial D_j^k}{\partial p_i^k}} \right) = \mu_i^k(\mathbf{p}_i; \mathbf{q}_j). \quad (14)$$

The above formula applies to arbitrary (asymmetric) duopoly markets.

In a recent paper, Tan and Zhou (2020) generalize the pricing formula derived by Armstrong (2006) to the case where there are more than two sides and more than two platforms, assuming a symmetric discrete choice model. The paper shows that, with two sides, no interaction costs (i.e., $\sigma^k = 0$, all k) and K firms covering the market, the first-order conditions for a symmetric equilibrium $(p_i^k, q_i^k) = (p_i, 1/K)$ all k is equal to

$$p_i - \left[c_i - \frac{\gamma_j}{K-1} \right] = \mu_i^k,$$

where μ_i^k is again the inverse semi-price elasticity of the side- i demand of firm k , holding constant the participation of side j and the other firms' prices. The interpretation is similar to the one discussed above. Attracting a small mass ε of consumers on side i allows firm k to raise its side- j price by $\gamma_j \left(\varepsilon + \frac{\varepsilon}{K-1} \right)$. The first term in parenthesis is the direct benefit of increasing the attractiveness of platform k for the side- j agents. The second term is the extra benefit of weakening the attractiveness of firm k 's competitors, again in the eyes of the side- j agents. In a symmetric equilibrium, the ε new agents the platform brings on board on side i come from a uniform reduction (by $\varepsilon/(K-1)$) in the side- i participation to each of the other $K-1$ firms. As in Armstrong (2006), the above conditions are valid only in markets with a sufficiently high degree of horizontal differentiation. With more than 2 sides of mass 1 each, the same reasoning shows that γ_j should be replaced by $\sum_{j \neq i} \gamma_{ij}$, where

γ_{ij} is the intensity of the interaction benefits from side i to side j (that is, the extent to which side j benefits from a marginal variation in the side- i participation).⁵⁶

The above pricing formulas can also be adjusted to accommodate for interaction costs. For example, in a symmetric market with K firms and two sides, it suffices to replace the side- i cost c_i in the above pricing formulas with $c_i + \sigma/K$, where the term σ/K is the cost of matching any new agent brought on board on side i to all the agents the platform attracts on the other side of the market (in a symmetric equilibrium, the measure of such agents is equal to $1/K$).

⁵⁶Tan and Zhou (2020) also derive first-order equilibrium conditions for the case where the market is not fully covered. In this case, reducing a price on one side expands the total demand and the simple additivity of the formulation above is lost.

While Tan and Zhou (2020) uncover important new insights, the literature still lacks a tractable model of platform competition in asymmetric and/or partially-covered markets.⁵⁷ As anticipated above, the difficulty is the complexity of identifying the price adjustments necessary to maintain participation on other sides constant when the participation of one side changes, as well as the complexity of identifying the elasticity of the relevant residual demands when maintaining the prices of the competing firms fixed. Section 9 discusses several papers that structurally estimate models of competing platforms with asymmetric and partial market coverage. As these papers are naturally computational in nature, they can utilize models that are intractable from a theoretical perspective.

4.1.2 Dispersed information

The possibility for multiple platforms to share the market hinges on the presence of heterogeneous stand-alone values, as interaction benefits alone tend to generate tipping, as explained above. However, in most markets of interest, it is difficult for an agent to predict whether many agents prefer one platform to another and hence how the platforms' prices will lead to different participation to each platform. This raises the question of how platforms should price access to each side when agents face idiosyncratic uncertainty about the distribution of preferences and hence about the ability of each platform to attract participation. This issue has been studied recently by Jullien and Pavan (2019) in a model that can be thought of as the extension of Armstrong (2006) to a market with dispersed information. The analysis accommodates for the possibility that each agent is uncertain about the distribution of preferences on each side of the market and uses his own appreciation for each platform's product to form beliefs about other agents' participation to each platform.

Formally, consider a two-sided market populated by two platforms and assume that interaction benefits are homogeneous within each side and given by γ_1 and γ_2 , respectively.⁵⁸ The distributions from which the agents' stand-alone values v_i^A and v_i^B for the two platforms are drawn is unknown both to the platforms and to the individual agents. Under suitable assumptions, the market is fully covered and each agent from each side $i = 1, 2$ uses the differential $v_i \equiv v_i^B - v_i^A$ in his stand-alone values for the two platforms to form beliefs over the distributions of preferences on each side.

Assume that the relevant "aggregate state" which is unknown to both the agents and the platforms is the pair of marginal distributions $\mathcal{F}^v \equiv (\mathcal{F}_1^v, \mathcal{F}_2^v)$ from which the agents' stand-alone differentials v_i are drawn. Given \mathcal{F}^v , each side- i agent's stand-alone differential v_i is drawn from \mathcal{F}_i^v independently across agents. Further assume that each agent's stand-alone differential v_i is the only information the agent has about \mathcal{F}^v . For simplicity, the platforms are assumed to possess no private information.

Given the firms' prices $(p_1^k, p_2^k)_{k=A,B}$ each agent from each side i then uses his type v_i to form beliefs about the participation to each platform on each side. Assuming that,

⁵⁷Some progress in this direction was recently made by Belleflamme, Peitz, and Toulemonde (2020) who extend Armstrong duopoly model by allowing for different network externalities on each platform. Among other things, the paper shows that the platform with the largest market shares need not be the most profitable one.

⁵⁸Consistently with the rest of the pertinent literature, Jullien and Pavan (2019) assume that the intensity of the network benefits γ_1 and γ_2 are symmetric across platforms. See Belleflamme, Peitz, and Toulemonde (2020) for a discussion of potential implications of dropping such an assumption.

as in Armstrong (2006), interaction benefits are not too large, we then have that, given any vector of prices, each side- i consumer chooses firm A if his stand-alone differential v_i is below a threshold \hat{v}_i , and firm B otherwise (with type \hat{v}_i being indifferent over which platform to join). When this is the case, the side- j participation to firm A expected by a side- i agent with type v_i is given by $\Pr(v_j \leq \hat{v}_j \mid v_i)$. Importantly, such participation need not coincide with the participation to platform A expected by the platforms, which is given by $q_j = \Pr(v_j \leq \hat{v}_j)$. This happens even if the platforms and all the agents share a common prior over the distribution from which the “aggregate state” \mathcal{F}^v is drawn.⁵⁹

Let $\hat{q}_{ji}^A = \Pr(v_j \leq \hat{v}_j \mid \hat{v}_i)$ denote the side- j participation to firm A expected by the “marginal” agent on side i — i.e., by the side- i agent with type \hat{v}_i implicitly defined by

$$\hat{v}_i = p_i^B - p_i^A - \gamma_i [\Pr(v_j \geq \hat{v}_j \mid \hat{v}_i) - \Pr(v_j \leq \hat{v}_j \mid \hat{v}_i)] \quad (15)$$

who is just indifferent between joining platform A and joining platform B .

Continue to assume that interaction costs are zero and that prices on each side take the form of participation fees so that $P_i^k = p_i^k$, $i = 1, 2$, $k = A, B$. Jullien and Pavan (2019) then show that platform A ’s equilibrium prices satisfy the following optimality conditions (similar conditions apply to platform B)

$$p_i^A - \left[c_i^A - 2\gamma_j \frac{\partial \hat{q}_{ij}^A}{\partial q_i^A} q_j^A \right] = \hat{\mu}_i^A(\mathbf{p}_i; \mathbf{q}_j), \quad (16)$$

where $\hat{\mu}_i^A(\mathbf{p}_i; \mathbf{q}_j)$ is the inverse-semi-price elasticity of the platform’s side- i expected demand q_i^A , holding the opponent’s prices and the side- j threshold \hat{v}_j constant, and where $\partial \hat{q}_{ij}^A / \partial q_i^A$ is the change in the side- i demand expected by the side- j marginal agent, when the platform increases the participation it expects from side i .

Private information affects pricing in two important ways. First, the opportunity cost of expanding the side- i participation (the term in square brackets in (16)) must now account for the fact that when the platform expects a change in the side- i price to bring on board ε more agents from side i , the marginal agent on side j expects the side- i demand to change by $\varepsilon \cdot \partial \hat{q}_{ij}^A / \partial q_i^A$. The difference in the side- i participation expected by the platform and by the side- j marginal agent reflects the difference in their beliefs over the “aggregate state” (i.e., over the cross sectional distribution of stand-alone differentials).⁶⁰

Second, the inverse-semi-price elasticity must now account for the fact that, when the platform expands the side- i participation, it also changes the beliefs of the side- i marginal agent over the side- j participation. This change in beliefs occurs even if, from the platform’s perspective, the side- j demand remains constant (that is, even if \hat{v}_j does not change). In

⁵⁹The analysis in Jullien and Pavan (2019) does not assume that the agents’ and the platforms’ beliefs be derived from a common prior. We assume a common prior here just to facilitate the exposition.

⁶⁰The platform and the side- j marginal agent agree that a decrease in the side- i price p_i^A by δ , when paired with a variation in the side- j price p_j^A that leaves the identity of the side- j marginal agent \hat{v}_j (and hence the side- j demand expected by the platform) unchanged triggers a change in the identity of the side- i marginal agent by $(\partial \hat{v}_i / \partial p_i^A) \delta$, with $\partial \hat{v}_i / \partial p_i^A$ obtained from (15). However, they disagree over the effect that such a change in the identity of the side- i marginal agent has on the side- i demand. In particular, while the platform expects the side- i demand to change at a rate $d \Pr(v_i \leq \hat{v}_i) / d \hat{v}_i$, the side- j marginal agent expects it to change at a rate $\partial \Pr(v_i \leq \hat{v}_i \mid \hat{v}_j) / \partial \hat{v}_i$. The term $\partial \hat{q}_{ij}^A / \partial q_i^A$ captures this discrepancy.

turn, such changes in the beliefs of the side- i marginal agent over the side- j 's participation affect the slope of the platform's side- i inverse residual demand which is now given by

$$\frac{\partial p_i^A}{\partial q_i^A} = \frac{\partial p_i^A}{\partial q_i^A} \Big|_{\hat{q}_{ji}^A = cst} + 2\gamma_i \frac{\partial \hat{q}_{ji}^A}{\partial q_i^A}.$$

The sign of the second term depends on whether preferences are “aligned” or “misaligned” between the two sides (that is, on the affiliation between the stand-alone differentials between the two sides).⁶¹ When preferences are aligned, agents with a higher appreciation for a platform's product also expect a higher appreciation by agents from the opposite side, whereas the opposite is true when preferences are misaligned.

When preferences are aligned, the new term is negative, thus contributing to a steeper inverse residual demand, whereas the opposite is true when preferences are misaligned. This effect has no counterpart under complete information. To understand this new effect, notice that increasing the side- i demand implies increasing \hat{v}_i which in turn means attracting marginal consumers with a lower appreciation for platform A 's product vis-a-vis platform B 's product. When consumers form beliefs over the distribution of preferences on the other side of the market using their own appreciation for the platforms' products, and preferences are aligned between sides, the new marginal agent from side i is thus less optimistic about the side- j participation to platform A than the old marginal agent (the one with a lower \hat{v}_i). Other things equal, this new effect tends to reduce the elasticity of the residual demand and thus contribute to higher prices when preferences are aligned, and to increase the elasticity of the residual demands and hence reduce the equilibrium prices when preferences are misaligned.

This finding reflects the fact that, when preferences are aligned, the new marginal agent is less optimistic than the infra-marginal agents about the participation of the other side. Therefore, the platform must cut the price more to trigger the same expansion in the side- i demand compared to the case where preferences are independent across sides (note that preferences are always independent when the cross-sectional distribution of stand-alone differentials on each side is common knowledge).

The above observations have various implications for the structure of the equilibrium prices. For example, the equilibrium price on each side depends directly on the magnitude of the interaction benefits on both sides whereas, with complete information, it depends only on the intensity of the interaction benefits on the opposite side, as shown in Armstrong (2006). More importantly, platforms may alter the intensity of competition by engaging in marketing activities and various other information-management policies geared at influencing the ability of each side to predict the other side's participation (through forums, showrooms, the dissemination of information about early adoption decisions and the like). Jullien and Pavan (2019) study the effects of the aforementioned policies on equilibrium prices, profits, and consumer surplus restricting attention primarily to symmetric markets that are fully covered. Extending the analysis to richer configurations remains an interesting line for future research.

⁶¹Formally, preferences are *aligned* if, for all \hat{v}_j , $\Pr(v_j \leq \hat{v}_j \mid v_i)$ is decreasing in v_i . They are *misaligned* if $\Pr(v_j \leq \hat{v}_j \mid v_i)$ is increasing in v_i .

4.1.3 Heterogeneous interaction benefits

Consider now a two-sided market in which preferences are as in (10), but where the interaction benefits are heterogenous across agents (that is, F_i is not restricted to assign probability one to types $\theta_i = (v_i^A, v_i^B, \gamma_i)$ that share the same γ_i).

To make things simple assume that types can be parametrized by a uni-dimensional variable $\omega_i \in \Omega_i \subset \mathbb{R}$, drawn from a distribution \mathcal{F}_i , $i = 1, 2$. That is, there exists a function $g_i : \Omega_i \rightarrow \Theta_i$ such that when ω_i is drawn from Ω_i according to \mathcal{F}_i , independently across agents, then the cross-sectional distribution of $g_i(\omega_i)$ coincides with the type distribution F_i on Θ_i . For any ω_i then let $\gamma_i(\omega_j)$ be the intensity of the interaction benefit for each side- i agent whose type θ_i is parametrized by ω_i (in short, for any ω_i -agent). Assuming for conciseness that the differential in stand-alone values $v_i(\omega_i) \equiv v_i^B(\omega_i) - v_i^A(\omega_i)$ is increasing in ω_i , under appropriate conditions, we then have that, for each price vector, there exist thresholds $\hat{\omega}_1$ and $\hat{\omega}_2$ such that each side- i agents joins firm A if $\omega_i < \hat{\omega}_i$ whereas he joins firm B if $\omega_i > \hat{\omega}_i$, $i = 1, 2$.⁶² Paralleling the analysis above, we then have that the price adjustment required to maintain the participation of side j constant when expanding the side- i participation can be identified by preserving the utility differential of the side- j marginal agent (the one whose type is indexed by $\hat{\omega}_j$). The equilibrium prices then solve the first-order conditions

$$p_i^A - [c_i^A - 2\gamma_j(\hat{\omega}_j)q_j^A] = \mu_i^A(\mathbf{p}_i; \mathbf{q}_j),$$

where the term $\mu_i^A(\mathbf{p}_i; \mathbf{q}_j)$ continues to denote the inverse-semi-price elasticity of platform A 's side- i demand with respect to the price p_i^A , for given prices by firm B and given participation \mathbf{q}_j on side j . As in the monopoly case, a Spence distortion then arises due to the fact that the firm accounts only for the effect of a variation in the side- i demand on the interaction benefit $\gamma_j(\hat{\omega}_j)$ of the side- j 's marginal agent, instead of all side- j agents on board (see Weyl, 2010). This distortion is in addition to the one originating in the firm's market power, as reflected in the inverse-semi-price elasticity term μ_i^A .

When the heterogeneity of the interaction benefits γ_i is large, asymmetric equilibria may arise. An alternative to product differentiation is for platforms to differentiate their business models. Introducing different price skewness to court different market segments on each side of the market is one way to alleviate competition. An illustration is provided in Ambrus and Argenziano (2009) who investigate the conditions for the existence of asymmetric networks, as for example, in the market for online job search in the US or in the credit-card industry. They assume that, on each side, there are two types of consumers, the first with large interaction benefits γ_i , the second with small interaction benefits. Assuming that the measure of the second type is large relative to that of the first type, Ambrus and Argenziano (2009) then show that each platform's optimal strategy is to set a low price on one side (say side 1), in order to attract the large mass of consumers with low interaction benefits and set a high price on the other side (side 2) targeted to those agents with large interaction benefits. When platform A sets low prices on side 1 and high prices on side 2, the best response for platform B is to do the opposite (set a low price on side 2 and a high price on side 1). This is because many agents are not served by platform A on side 2 and

⁶²This is the case for instance if $\gamma_i(\omega_i)$ are small, for all ω_i , $i = 1, 2$, and if the distribution of $v_i(\omega_i)$ is sufficiently diffused.

hence are easily attractable by firm B by setting a low price. Conversely, on side 1 those agents with low interaction benefits are already attracted by platform A in which case it is more profitable for platform B to target those side-1 agents with high interaction benefits who value being connected to many agents from the other side.

In the case of online job search, the result implies that a natural equilibrium configuration is one where two platforms asymmetrically split the market, with one platform attracting many job posters and few job seekers and the other platforms attracting many job seekers and few job posters.⁶³

4.2 Multi-homing

The previous contributions share the important assumption that each agent joins at most one platform. How robust are the ideas developed above to the possibility of multi-homing? As long as there are two active firms, and some agents choose to join only one platform, the agents on the other side will have incentives to subscribe to both platforms. Multi-homing allows agents to benefit from larger interaction benefits but also enjoy the differentiated products supplied by both firms.⁶⁴

Note that, for given horizontal differentiation in the stand-alone values, the incentives for each side- i agent to multi-home are inversely related to the measure of side- j agents who multi-home. Let m_1 and m_2 denote the measure of agents who multi-home on side 1 and 2, respectively. The incremental value of multi-homing relative to single-homing depends on the additional stand-alone value that each agent obtains by joining a second platform and the incremental interaction benefit that the agent obtains by reaching those agents from the opposite side that single-home. Specifically, consider a side- i agent of type $\theta_i = (v_i^A, v_i^B, \gamma_i)$ buying from firm A . The agent prefers to multi-home if

$$v_i^B - \kappa_i + \gamma_i (q_j^B - m_j) \geq p_i^B$$

where κ_i can be interpreted as the cost of multi-homing, which can be either a direct adoption cost or the utility loss from purchasing a product that is not fundamentally different from that provided by firm A . In other words, $v_i^B - \kappa_i$ is the incremental value from consuming the good provided by platform B when consuming already the good provided by firm A . The term $\gamma_i(q_j^B - m_j)$, instead, is the interaction benefit of reaching those side- j agents who are present on firm B and who do not multi-home. Everything else equal, it is easy to see that the mass of consumers willing to multi-home on side i decreases with the mass m_j of side- j consumers who multi-home. In particular, when interaction benefits are an important driver of the demand on each side (relative to stand-alone values), if all consumers multi-home on side- j , few consumers multi-home on side i . This led researchers to investigate asymmetric models with multi-homing on one side only.

⁶³In a media context, Calvano and Polo (2020) study how platforms may achieve product differentiation by choosing different business models, arguing that when consumers multi-home, an ad-financed business model becomes less attractive if the other platform adopts it.

⁶⁴There is a link between single-homing/multi-homing and substitutability/complementarity. But since the features of a platform depend on the number/type of agents joining the platform, the notion of substitutability/complementarity is endogenous in a two-sided context. If one side multi-homes, the other sides views the competing platforms are providing imperfectly substitutable goods. But if there is single-homing on the first side, the competing platforms are seen as providing complementary goods by the agents on the other side of the market.

4.2.1 Competitive bottleneck

As an example of a model where agents multi-home on one side only, consider again the setting in Armstrong (2006) but assume now that the side-2 agents multi-home whereas the side-1 agents continue to single-home (this is the structure considered in Armstrong and Wright, 2007⁶⁵).

The multi-homing behavior of the side-2 agents dramatically changes the competition between the two platforms. The side-2 agents de facto choose whether or not to join each platform as if it were a monopoly: given that choosing a platform does not preclude joining the other, there is no direct competition between the two platforms to attract the side-2 agents.

The lack of competition on side 2 de facto changes the bargaining power between the agents and the platform. Indeed, because the side-1 agents single-home, they are a scarce resource for which the platforms are ready to fight. If any of the side-2 agents wants to be connected to a specific agent from the other side, he has to join the platform this specific agent has joined. This situation is often referred to as a competitive bottleneck. Because of the lack of competition on side 2, platforms raise their side-2 price, inducing too few side-2 agents to join. The effects of multi-homing on the equilibrium side-1 prices are ambiguous. On the one hand, because side 2 multi-homes, the value of stealing a side-1 agent as a way to weaken the competitor is not present anymore. On the other hand, the value of attracting a side-1 agent may be larger under multi-homing because each platform may attract more agents on side 2 under multi-homing and hence may match the additional side-1 agent it attracts to a larger set of side-2 agents.

To illustrate the above effects assume that consumers on side 2 multi-home at no cost ($\kappa_2 = 0$), that all side-1 consumers single-home ($m_1 = 0$) and that interaction benefits are homogeneous within each side and equal to γ_1 and γ_2 for side 1 and side 2, respectively. Then for each consumer on side 2, the decision of whether or not to join platform k doesn't depend on the other platform's price or sales. The side-2 demand for each platform $k = A, B$ is given by

$$q_2^k = \Pr \left(v_2^k + \gamma_2 q_1^k - p_2^k \geq 0 \right) \geq D_2^k (\mathbf{p}_2; \mathbf{q}_1),$$

where $D_2^k (\mathbf{p}_2; \mathbf{q}_1)$ is firm k 's demand on side 2 under single-homing.

Many applications of the competitive bottleneck model assume that firms set quantities on the multi-homing side instead of prices. In this case, each firm k chooses a price p_1^k for the single-homers and the quantity q_2^k of multi-homers, with the quantity q_1^k and the price p_2^k obtained as residuals by market-clearing conditions.⁶⁶ In this case, most of the intuitions developed in the previous sections extend but multi-homing has two novel implications for the determination of the opportunity costs. First, as anticipated above, the value of stealing a side-1 user as a way to weaken the competitor on the other side of the market vanishes, as the two platforms no longer compete on the multi-homing side. Hence, in computing the opportunity cost of getting a marginal side-1 consumer on board, each firm internalizes only the value of increasing the attractiveness of its network for its side-2 agents. As a

⁶⁵In this paper the side multi-homing is endogenously derived

⁶⁶This is for instance the case for most models of advertising. See Anderson and Coate (2005), Armstrong (2006), and Crampes, Haritchabalet, and Jullien (2009).

consequence of this effect, the optimal prices on the single-homing side are now determined by the first-order condition⁶⁷

$$p_1^k - [c_1^k - \gamma_2 q_2^k] = \mu_1^k(\mathbf{p}_1; \mathbf{q}_2)$$

where $\mu_1^k(\mathbf{p}_1; \mathbf{q}_2)$ continues to denote the inverse-semi-price elasticity of the residual demand on the single-homing side.

Second, on the multi-homing side, attracting a marginal consumer does not imply stealing him from the competitor, if the marginal consumer multi-homes. This is because the attracted consumer continues to buy from the competing firm. Attracting dq_2^k new side-2 consumers while adjusting the side-1 price p_1^k by $\gamma_1 dq_2^k$ would leave \mathbf{q}_1 unchanged. The first-order optimality condition on the multi-homing side then becomes

$$p_2^k - [c_2^k - \gamma_1 q_1^k] = \eta_2^k(p_2^k; q_1^k) \quad (17)$$

where $\eta_2^k(p_2^k; q_1^k)$ is the inverse-semi-price elasticity of the monopoly demand.

Hence each firm de facto acts as a two-sided monopolist on the multi-homing side, with the opportunity cost adjusted to account for the effects of expanding the side-2 participation on the side-1 price (recall the discussion in Section 2). Notice that despite each firm acting as a monopoly on the multi-homing side, there is still an indirect price externality between the firms on this side. Any reduction of the side-2 price p_2^A induces a shift of the side-1 demand toward A and therefore a reduction of the demand for firm B on the single-homing side.

As pointed out in Armstrong (2006), the formula in (17) implies that the platforms choose their prices on the multi-homing side so as to maximize the joint surplus of the firm and of the single-homing side at constant volume of participation. In other words, p_2^k maximizes $(p_2^k - c_2^k) q_2^k + \gamma_1 q_1^k q_2^k$ holding q_1^k constant. General comparisons between the multi-homing and the single-homing prices is not easy because multi-homing changes the level and the elasticity of the demands. For given participation on side 2 and elasticity on side 1, the price each platform sets on side 1 is higher because the firm does not value stealing consumers on side 1 as a way of weakening its competitor on side 2. In this respect, multi-homing reduces the impact of network effects on prices by reducing the competitive edge generated by higher interaction benefits. On the single-homing side, this effect, however, must be balanced by the change in sales on side 2, as anticipated above. When multi-homing increases the side-2 demands, the value of attracting a marginal consumer on the single-homing side is higher because the firm may then match the consumer to a larger user base on the other side of the market. Thus, there are conflicting effects. For instance, in a model where all consumers on side 2 multi-home and have inelastic demands, Armstrong (2006) shows that the equilibrium price on the single-homing side is the same as when both sides single-home.

⁶⁷This condition is computed assuming that platforms choose quantity on side 2. If, instead, they choose prices, the interaction benefits between sides induce a 2-sided feedback loop on the competing firm's demand q_2^{-k} that requires tracking the change of q_2^{-k} in the computation of the inverse-semi-price elasticity μ_1^k . This happens even if the market is fully covered. See Armstrong (2006) and Crampes, Haritchabalet, and Jullien (2009).

The implications of multi-homing for consumer welfare are not straightforward either. First, *ceteris paribus*, allowing some agents to multi-home raises their options and thus their utility. However, prices may increase or decrease due to multi-homing, thus making the analysis of the effects of the latter on consumer surplus ambiguous. Moreover, the elasticity η_2^k of the monopoly demand differs in a non-trivial way from the elasticity of the residual demand μ_2^k in the duopoly case, which also makes the comparison of prices and hence of consumer welfare tedious. Still, intuition suggests that the single-homing side should benefit if its participation is very valuable to the other side (γ_2 large) and if multi-homing raises the volume of sales significantly. Similarly, the multi-homing side should benefit if the monopoly demand is elastic enough. Belleflamme and Peitz (2019a) confirm these insights in an extension of Armstrong (2006)'s Hotelling model, varying the cost κ_2 of multi-homing.

4.2.2 Multi-homing on both sides

The competitive bottleneck model assumes that agents multi-home only on one side. For example, when platforms are the owners of smartphone operating systems, it makes sense to expect most users to buy a single smartphone. This is because the cost of purchasing multiple handsets typically offsets the benefit of being able to use the apps that are developed for one smartphone operating system but not the others. On the other hand, developers typically develop apps for multiple operating systems.

In other markets, though, multi-homing may occur on both sides and this possibility brings new effects.⁶⁸ There are three types of reasons for multi-homing to occur.

1. Incremental stand-alone values: the two stand-alone products need not be perfect substitutes so that buying the second good may provide additional utility irrespective of whether it also provides access to more agents from the other side.
2. Incremental volume of interactions: when some agents single-home on the other side, multi-homing allows an agent to raise the volume of interactions.
3. Incremental value of interactions: interacting with the same agents on multiple platforms may be more valuable than interacting with the same agents only on one platform. In the context of advertising, multiple impressions may be more valuable than single impressions. More generally, interacting over multiple platforms may increase the likelihood that a final match is formed and/or its quality, as pointed out in Cailaud and Jullien (2003).

To see the implications of multi-homing on both sides on equilibrium prices, assume that, on each side, a consumer who multi-homes obtains a total stand-alone utility equal to $w_i = v_i^A + v_i^B - \kappa$, where, as before, κ denotes the cost of combining the two stand-alone products. Further assume that each multi-homing consumer obtains an additional interaction benefit equal to χ_i from interacting with an agent from the other side who multi-homes. Continue to let q_j^k denote the total participation to firm k on side j and m_j the mass of

⁶⁸See the recent analysis in Bakos and Halaburda (2020) and Jeitschko and Tremblay (2020).

side- j multi-homers. The total participation on side j is then equal to $q_j^A + q_j^B - m_j$. The utility of a multi-homing consumer on side i is thus equal to

$$w_i + \gamma_i (q_j^A + q_j^B - m_j) + \chi_i m_j - p_i^A - p_i^B.$$

Assuming full participation (i.e., $q_j^A + q_j^B - m_j = 1$), the consumer thus chooses to multi-home if

$$w_i - v_i^{-k} + \gamma_i (q_j^k - m_j) + \chi_i m_j > p_i^k \text{ for } k = A, B. \quad (18)$$

The term $w_i - v_i^{-k}$ captures the degree of substitutability between the two firms' stand-alone products, ranging from 0 in case of perfect substitutes to v_i^k in case of independent products. The second term captures the increase in the volume of interactions with the other side (recall that $q_j^k - m_j$ is the measure of side- j agents who are present only on platform k). The third term $\chi_i m_j$ captures the increase/decrease in interaction efficiency (i.e., the change in benefits originating from interacting with the same agents on both platforms as opposed to a single platform).

For our concerns, the interaction terms are the most interesting. Typically, the additional benefit of interacting with the same agent on multiple platforms is less than the benefit of interacting with the same agent on a single platform, that is, $\chi_i < \gamma_i$. This implies that, at fixed prices, increasing the volume of agents multi-homing on one side decreases the volume of agents who multi-home on the other side. Hence, when products are substitutes and the incremental gain χ_i of interacting with agents from the opposite side on multiple platforms is small, we may expect one side to single-home if the number of agents who multi-home on the other side is large. Gabszewicz and Wauthy (2014) illustrate this point. They find that only one side multi-homes in a model with heterogeneous interaction benefits. They suggest a link between multi-homing in two-sided markets and multi-purchases in vertical differentiation models. Interaction benefits constitute a quality dimension and multi-homing is similar to buying a bundle that raises quality compared to each single product in the bundle. Interaction benefits thus introduce a vertical dimension in the competition between the firms. As in models of standard vertical differentiation (Mussa and Rosen, 1978 and Shaked and Sutton, 1983), only consumers with a high value for quality (high interaction benefits) buy the high quality product (here multi-home).

Along this line, Doganoglu and Wright (2006) build a model of platforms where consumers are heterogeneous in their interaction benefits and both sides multi-home.⁶⁹ Suppose there are two types of consumers on each side: half of the consumers have a high interaction benefit $\bar{\gamma}_i$ and the other half has a low interaction benefit $\underline{\gamma}_i < \bar{\gamma}_i$. Products are independent and there is no efficiency from multi-homing ($\chi_i = 0$) so the sole motive for multi-homing is to expand the network interactions. Then a consumer multi-homes if $v_i^k + \gamma_i (q_j^k - m_j) - p_i^k > 0$ for $k = A, B$. Assume that in equilibrium the market is covered on each side and prices are at intermediate levels so that low types single-home whereas high types multi-home. The demand from low types is then as in the baseline model $D_i^k(\mathbf{p}_i; \mathbf{q}_j) / 2$ (accounting for the mass 1/2) while the demand from high types is as in the competitive bottleneck model and given by

$$\bar{D}_i^k(p_i^k; q_j^k - m_j) = \frac{1}{2} \Pr \left(v_i^k + \bar{\gamma}_i (q_j^k - m_j) - p_i^k > 0 \right),$$

⁶⁹See also Kim and Serfes (2006).

with the volume of multi-homers given by⁷⁰

$$m_j = \bar{D}_j^A(p_j^A; q_i^A - m_i) + \bar{D}_j^B(p_j^B; q_i^B - m_i) - \frac{1}{2}.$$

Compared to the competitive bottleneck model, there are two differences. First, each firm faces both single-homers and multi-homers and marginal consumers are spread over the two categories. Hence the situation is a hybrid between the single-homing case and the competitive bottleneck. Second, the value offered to marginal multi-homing consumers on one side decreases with the volume of multi-homing agents on the other side. This is the consequence of a more general principle that governs competition with multi-homing. When the marginal consumer is already buying from the other firm, a firm can only charge for the additional interaction benefits it brings to the consumer given the level of interactions at the other firm. This has been referred to as the “incremental pricing principle” in the advertising literature.⁷¹

When applied to each firm, this principle implies that each firm is constrained in its ability to extract rents from the consumers. For instance, if all consumers multi-home on side j (that is, $q_j^A = q_j^B = m_j$) and all marginal consumers multi-home on side i , then no firm can charge a mark-up for interaction benefits on side i . This incremental-pricing principle has been used to explain the decline of aggregate advertising revenue following the rise of the Internet and the declining ability by large advertising platforms to charge higher prices (Athey, Calvano, and Gans, 2018). Rochet and Tirole (2003) point to another aspect of competition with multi-homing that they refer to as “steering”. In their setup, firms are credit cards setting transaction fees. Sellers must decide whether to accept transactions on both platforms (multi-home) or only one.⁷² In case a seller accepts to trade on both platforms, the buyer chooses which platform to use. A seller choosing to trade only on the platform setting the lowest transaction fee reduces its total sales but increases its average margin by forcing some buyers to trade on their least-preferred platform. Reducing transaction fees in this context raises the volume of transactions by convincing some sellers to stop accepting transactions on the other more expensive platform..

4.3 Concentration, merger and collusion

4.3.1 Concentration and entry of platforms

It is well known that there may be excessive or insufficient entry in industries absent network effects.⁷³ The same holds true for goods involving network externalities, with the added feature that market fragmentation reduces interactions benefits on each platform (see Katz and Shapiro, 1994, for a discussion). Thus increasing concentration need not hurt consumers in markets with network externalities.

⁷⁰As the market is covered, every agent on side j is either a single-homer or a multi-homer. If x_j^A and x_j^B are the mass of high type single-homers then $\bar{D}_j^k(p_j^k; q_i^k - m_i) = x_j^k + m_j$ and $x_j^A + x_j^B + m_j = 1/2$.

⁷¹See Anderson, Foros, and Kind (2018), Ambrus and Reisinger (2006), Ambrus, Calvano, and Reisinger (2016) and Anderson and Peitz (2020a).

⁷²Liu, Teh, Wright, and Zhou (2020) extend the model to allow consumers to multi-home as well, showing that it induces some re-balancing of transaction fees in favor of consumers.

⁷³See for instance Anderson, de Palma, and Nesterov (1995) for conditions under which entry is excessive from a total-welfare perspective.

The literature based on incumbency advantage in one-period models confirms this insight for two-sided markets, despite barriers to entry. For instance, Jullien (2011) and Vasconcelos (2015) conclude that both excessive or insufficient entry may occur if network externalities are not too large. This is because an incumbent may find it too costly to deter all divide-and-conquer strategies by an entrant and decide to accommodate some (inefficient) market fragmentation. Amelio, Giardino-Karlinger, and Valletti (2020) characterize a limit-pricing equilibrium with inefficient entry deterrence for some intermediate range of entry costs.

In the context of competition on the market, Tan and Zhou (2019) analyze the effect of concentration in the discrete-choice model with K differentiated platforms and single-homing consumers on both sides discussed in Section 4.1.1. Reducing market fragmentation—reducing K —raises average participation in each platform on both sides and thus raises total interaction benefits. However, a lower elasticity of residual demands tends to increase price-cost margins, due to the familiar market-power effect. The market-power effect dominates if network effects are small. The paper shows that, depending on the specification of the demand function, increasing K may reduce or increase prices on any side of the market.^{74,75} In particular, prices may increase on one side and decrease on the other side. This property is a manifestation of the “seesaw” principle of Rochet and Tirole (2006), according to which *“a factor that is conducive to a higher price on one side, to the extent that it raises the platform’s margin on that side, tends also to call for a lower price on the other side, as attracting members on that other side becomes more profitable.”*

Correia-da Silva, Jullien, Lefouili, and Pinho (2019) study the effect of concentration in a Cournot model with K homogeneous platforms (see Section 5.2 for a discussion of the Cournot model). The paper finds that reducing K reduces quality-adjusted prices, $p_i^k - \gamma_i q_j^k$, $i \neq j$, on both sides $i = 1, 2$ if pre-merger quality-adjusted prices are below marginal costs on both sides, which occurs when interactions benefits are large. As lower quality-adjusted prices translate into higher consumer surplus in their model, consumer surplus increases on both sides with concentration if network effects are important. The reverse holds if interaction benefits are small and pre-merger quality-adjusted prices are above cost on both sides. The paper also identifies an intermediate range of interaction benefits such that increasing concentration harms consumers on one side and benefits consumers on the other side—a seesaw effect.

Anderson and Peitz (2020b) also analyze the seesaw effect in an aggregative-game formulation⁷⁶ of a media market. The paper focuses on a competitive bottleneck model and argues that the seesaw effect is more likely to occur when advertising is a nuisance and content is offered free of charge to consumers.⁷⁷

⁷⁴The possibility that increasing concentration may reduce prices also exists in one-sided setting with product differentiation (Chen and Riordan, 2008). We focus on platform-specific effects.

⁷⁵The paper also derives conditions for excessive entry to occur in a multi-sided market model.

⁷⁶Aggregative games simplify the equilibrium analysis by using an aggregate of players’ actions to reduce the dimensionality of strategic interactions (Selten, 1973). See Acemoglu and Jensen (2013) and Anderson, Erkal, and Piccinin (2020). This simplification allows Anderson and Peitz (2020b) to consider asymmetric platforms and partial market coverage

⁷⁷In a competitive bottleneck setup with ad nuisance, the attention to ads is the price to access content. A merger relaxes competition for consumers so that it raises the volume of advertising and hurts consumers (Anderson and Coate, 2005). This result may be reversed if consumers multi-home, because multi-homing softens competition for consumers (Anderson, Foros, and Kind, 2019).

4.3.2 Horizontal mergers

Unsurprisingly, similar conclusions emerge for mergers. Because of two-sidedness, a merger may not lead to higher prices and lower consumer surplus on both sides. This was first pointed out by Chandra and Collard-Wexler (2009). Considering two newspapers selling content to readers and space to advertisers, they notice that, due to imperfect screening, the reader subscription price may be excessive from the industry-profit perspective. Newspapers cannot discriminate between readers so that advertising revenue per reader is equal to the average value of a reader for advertisers. In order to sell to the readers most valuable to advertisers, a newspaper needs also to sell to readers not valuable to advertisers, who may contribute negatively to profit.⁷⁸ Then suppose that when a newspaper raises its subscription price, the marginal readers diverted toward the competing newspaper are those contributing negatively to profit, whereas the most attractive marginal consumers stop buying all together. In this case, raising one newspaper’s price has a negative externality on the other. Thus a monopoly owner of two newspapers may reduce subscription prices.⁷⁹

Tan and Zhou (2020) show, in a model with three platforms and $\gamma_2 = 0$, that, if γ_1 is intermediate, then a merger is profitable and results in lower prices and higher consumer surplus on side 2. Intuitively, if the pass-through of interaction benefits from side i to side j increases, it may be worthy boosting participation on side i by reducing the prices on that side. Higher participation on side i —which is a quality dimension for side- j consumers—may then benefit side j . Baranes, Cortade, and Cosnita-Langlais (2019) consider a model where four firms are located on the Salop circle so that competition is localized. When two adjacent firms merge without synergy, the market power effect dominates and all prices increase. They then show the possibility of a seesaw effect of cost synergies, with consumer surplus increasing on one side but decreasing on the other side when cost saving increases.

Finally, the owner of two platforms can raise compatibility between platforms and consequently boost the interaction benefits on the merged platforms (Leonello, 2010). Increasing compatibility may in turn incentivize the merged entity to reduce some price. To see that, suppose the owner of two platforms initially incompatible decides to make them compatible. The interaction benefits of a side- j consumer joining platform k that is merged with platform l becomes $\gamma_j (q_i^k + \xi q_i^l)$ instead of just $\gamma_j q_i^k$. As the merged entity coordinates prices, it internalizes the effect of side i ’s participation to platform k on the side- j ’s participation to platform l and may decrease the price p_i^k for $\xi > 0$ below the level chosen when $\xi = 0$.

4.3.3 Collusion

A seesaw effect is identified in Dewenter, Haucap, and Wenzel (2011)’s analysis of price coordination within a competitive bottleneck model: coordination reduces participation on the multi-homing side but may increase participation and consumer welfare on the single-homing side. Dewenter, Haucap, and Wenzel (2011) also introduce the possibility of price coordination on one side only (with competition on the other side), referred to as “semi-collusion.” They show that semi-collusion in the multi-homing side generates higher total

⁷⁸This occurs when the sum of the subscription price and the value of the consumer for advertisers is smaller than the newspaper’s marginal cost.

⁷⁹See also Foros, Kind, and Sørgaard (2015) for a discussion of the impact of mergers in media markets.

welfare than a full merger and may even generate higher total welfare than competition.⁸⁰

Lefouili and Pinho (2020) present a full-fledged analysis of tacit collusion in grim-trigger strategies (with Nash reversal) in Armstrong’s model. At the most profitable collusive equilibrium, incentive compatibility implies additional distortions of the price structure compared to competition, raising the price differential between the side with the larger degree of differentiation and the other side. In this model, the prices maximizing total profit are above competitive prices on both sides, but incentive compatibility may call for reducing the price below the competitive level on the least differentiated side, so as to reduce deviation profits. The paper then finds a seesaw effect of a different nature than in the merger analysis. A more standard seesaw effect occurs under semi-collusion, although the price may decrease on any side depending on the level of network effects. Under semi-collusion, firms have two possibilities: they may agree to raise margins on the collusive side with the consequence of reducing prices on the competitive side; or they may agree to boost demand on the collusive side by reducing margins, inducing higher interaction benefits and prices on the competitive side.⁸¹

4.4 Exclusivity and bundling

4.4.1 Exclusivity

Platforms use exclusive dealing as a way to prevent multi-homing.⁸² Requiring exclusivity on a multi-homing side j allows platforms to raise the price charged on side i for a given mass of side- j users on the platform. In a duopoly context, imposing exclusivity amounts to imposing that agents single-home. When there is little differentiation between platforms, exclusivity permits platforms to monopolize the side where exclusivity is imposed and gain a competitive hedge on the other side of the market. Hence exclusive contracts may destabilize equilibria where agents multi-home. Armstrong and Wright (2007) show that this is the case in a model of competitive bottleneck and characterize the resulting tipping equilibrium (discussed in Section 3) where all agents use the same platform that emerges when exclusive contracts are allowed and platforms are homogeneous. The tipping equilibrium is efficient but with lower consumer surplus than the competitive bottleneck without exclusive contracts.

The emergence of tipping equilibria when platforms can require exclusivity raises the question of strategic barriers to entry. The fact that exclusivity may be used by an incumbent platform to raise barriers to entry is confirmed by Doganoglu and Wright (2010). They first extend the naked-exclusion principle (Rasmusen, Ramseyer, and Wiley Jr, 1991, Segal and Whinston, 2000)⁸³ to markets with network effects.⁸⁴ They then focus on a two-sided

⁸⁰See Argentesi and Filistrucchi (2007) in Section 9 for empirical evidence.

⁸¹The paper also analyzes semi-collusion in a competitive bottleneck with similar insights (full collusion in their competitive bottleneck model has no effect on quantities and welfare because the market is fully covered on both sides).

⁸²There are of course other motives for exclusivity not specific to platforms, such as protecting specific investments or preventing free riding.

⁸³This principle states that, with increasing returns to scale, a monopoly can prevent efficient entry by offering discriminatory exclusive contracts, subsidizing a subset of consumers large enough to make trading with the entrant inefficient for other consumers.

⁸⁴In this case, subsidizing a subset of consumers for exclusive participation raises the value differential

market and assume that an incumbent platform A cannot discriminate within each side but may offer an exclusive contract to all agents on one side before an entrant B steps in and competes. Agents are assumed to be homogeneous on both sides with no stand-alone values ($v_i^k = 0$, $i = 1, 2$, $k = A, B$) and multi-homing. They assume that the entrant offers higher interaction benefits than the incumbent (i.e., $\gamma_i^B > \gamma_i^A$, $i = 1, 2$) and faces no chicken & egg problem (consumers coordinate on Pareto efficient allocations). Under these assumptions, absent exclusive contracts, the equilibrium is efficient and B wins the market.

When exclusive contracts are feasible, the incumbent can however prevent entry by offering exclusivity on one side of the market. To see that, suppose that all agents on side 1 grants exclusivity to platform A before B can make any offer. Then platform B cannot generate any value as it can only trade with one side. This allows A to charge a monopoly price γ_2^A to the non-exclusive side 2. Platform A thus offers exclusivity if there is a price $p_1^A > -\gamma_2^A$ that is accepted by the side-1 agents for their exclusive participation. Given that the surplus on side 1 is $\gamma_1^A - p_1^A$, this is the case if $\gamma_1^A + \gamma_2^A$ —the total value of A —is larger than the surplus accruing to side 1 under no exclusivity and competition between A and B . But without exclusivity, competition implies that consumers surplus, and a fortiori side-1 surplus, cannot exceed the maximal surplus that platform A can offer, which is precisely $\gamma_1^A + \gamma_2^A$.⁸⁵ As a consequence, the platform can profitably sell exclusive contracts and exclude its competitor. Notice that the argument requires no incumbency advantage but the ability to obtain exclusivity prior to the competition stage with the entrant. Clearly, if the entrant can also make such offers, there would be competition to sell exclusivity and the entrant would win the market.

The above argument reflects the view that exclusive dealing may protect incumbents. An alternative view is that exclusive contracts may help the entrant when the incumbent benefits from an incumbency advantage (such as favorable beliefs). In this case, exclusive offers may help a new platform overcome the chicken & egg problem by targeting key strategic players (Markovich and Yehezkel, 2021). Indeed, one issue faced by an entrant using a divide-and-conquer strategy when consumers can multi-home is that, if it subsidizes the participation of side i , it needs to prevent the possibility that side j stays with the incumbent and side i just multi-homes (collecting the subsidy but interacting with the other side on the incumbent's platform). This requires low prices on side j . Offering exclusivity to side i would then amount to switching to the single-homing divide-and-conquer strategy, which allows a platform to charge higher prices to side j but requires a larger subsidy to side i . When introductory exclusive contracts can be targeted in a discriminatory way, the entrant may then selectively target some agents on side i with a large subsidy for exclusive participation. This is in particular the case if some agents generate large interaction benefits, or if a subset of agents are part of an institution coordinating their choices (a user-group). By convincing these key agents to join exclusively, the entrant raises its relative value on the other side.

The view that exclusive deals can help smaller firms has some support empirically. As mentioned above, Lee (2013) studies exclusive deals between producers of video game consoles and video game developers. In his empirical results, without exclusivity, consumers and developers (who are allowed to multi-home) flow to a single console but exclusive deals

with the entrant for other consumers. See also Karlinger and Motta (2012).

⁸⁵This argument relies on efficient coordination by the consumers.

provide a way for smaller consoles to differentiate themselves and maintain market share.

An interesting feature pointed out in Carroni, Madio, and Shekhar (2020) is that exclusivity offered by a platform k to some key strategic agents on side i not only does it boost participation on the other side j , it may also help convincing other agents of the same side i to stop multi-homing and join exclusively platform k . The reason is that increasing platform k 's demand on side j to the expense of the other platform raises the value of platform k for agents of side i and reduces the incremental value of the other platform. To see that, recall from the analysis in the previous sections that the (gross) incremental value of platform $-k$ is the extra utility that a user of platform k obtains by joining platform $-k$ (in addition to joining platform k) and is given by

$$w_i - v_i^k + \gamma_i (q_j^{-k} - m_j) + \chi_i m_j = w_i - v_i^k + (\gamma_i - \chi_i) (q_j^{-k} - m_j) + \chi_i q_j^{-k},$$

where q_j^{-k} is total participation to platform $-k$ on side j and $q_j^{-k} - m_j$ is the exclusive participation to platform $-k$. Assume that $0 \leq \chi_i < \gamma_i$. Then given that both total and exclusive participation to platform $-k$ decrease when platform k becomes relatively more attractive, the incremental value of platform $-k$ decreases and more agents single-home. Hence, gaining exclusivity of some agents brings additional exclusive participation at no cost on the same side.⁸⁶ However, gaining exclusivity of these strategic agents is costly as they need to be compensated for giving up some interaction benefits. In the analysis of Carroni, Madio, and Shekhar (2020) of a competitive bottleneck, exclusivity arises when competition on the single-homing side is intense. As in Armstrong and Wright (2007), the equilibrium may be more efficient due to larger interaction benefits when concentration increases.

Hence, exclusivity raises equilibrium concentration when network effects are large and, depending on the context, may raise or reduce barriers to entry.⁸⁷

4.4.2 Bundling

Bundling is a widespread practice by platforms (which stands a bit outside the models discussed so far) that is motivated by specific features of platforms and leads to new economic insights. Choi (2010), and its corrigendum Choi, Jullien, and Lefouili (2017), argue that bundling may help a firm coordinating agents on its platform. The paper considers a situation where the platforms are pure complements of a monopoly base product. The monopoly provider A of the base product owns one platform, but cannot capture the revenue from an independent complementary platform—because it cannot directly charge the platform. When platform A is not tied to the base product, partial multi-homing on both sides of

⁸⁶Ishihara and Oki (2020) point out that, when both sides multi-home, increasing exclusive participation on one side raises the incremental value of each platform and thus the mass of multi-homing consumers on the other side. Put it differently, exclusivity on one side increases platform differentiation on the other side and may relax competition.

⁸⁷See also the literature on exclusivity in media, mostly focusing on premium content. For instance, Chowdhury and Martin (2017) focus on exclusive licensing. D'Annunzio (2017) and Stennek (2014) study the relation between exclusivity and quality of content. Also related is Kourandi, Krämer, and Valletti (2015), focusing on Internet fragmentation and net-neutrality.

the market dissipates profits. By tying the platform to the base product, platform A then convinces both sides to single-home on A , excluding the competing platform and raising total revenue.

Amelio and Jullien (2012) point out that non-negative price constraints create new incentives for bundling that are specific to platforms.⁸⁸ Indeed, as discussed in Section 2.2.3, a zero price in some contexts reveals that the platform would be willing to set negative prices to attract consumers but cannot do so. Bundling may then be used to boost adoption when the service is free. In a competitive context, relaxing the non-negative price constraint by bundling participation with goodies has two effects. First, it raises the attractiveness of the platform on the other side. Second, it relaxes the incentives to compete with other platforms as consumers on the profitable side become more costly (due to the feedback effect on subsidies). Bundling has thus both a direct effect on the attractiveness of the platform and a strategic effect that reduces the intensity of competition. While in the monopoly case bundling is welfare improving, the effects of bundling in a duopoly are ambiguous. In particular, in the symmetric model of Armstrong (2006), allowing symmetric bundling possibilities reduces consumer surplus. Choi and Jeon (2021) clarify that a key factor beyond new theories of bundling in two-sided markets is that a platform’s ability to cope with aggressive competition is limited by the non-negative price constraint. Hence, if one platform has exclusive access to some way of relaxing the non-negative price constraint and subsidizing the participation of one side (with bundling for instance), it can win the competition even if it is less efficient than its competitor and there are no consumer coordination failures.

Bundling is of great interest in the context of cable television. Cable systems are platforms connecting channels with consumers, and most cable systems offer bundles of channels organized in plans. Crawford and Yurukoglu (2012) estimate a structural model of consumer demand, pricing, and bargaining between cable systems and channels. They find that consumers would benefit from *à la carte* pricing if fees from cable systems to channels remained constant. However, the switch to *à la carte* pricing increases pricing power for channels and enables them to bargain for higher fees because these will be passed through to consumers more directly. When accounting for this effect on input prices, *à la carte* pricing hurts consumers relative to bundled pricing.⁸⁹

5 Alternative modeling of competition, coordination and beliefs

Equilibria in games of platform competition with network externalities typically depend on the assumptions involving (i) the set of feasible price structures, and (ii) the demand configuration given the firms’ prices. In this section, we discuss in more detail each of the above aspects.

⁸⁸Traditional motives for bundling include price-discrimination and exclusion of competitors (see Fumagalli, Motta, and Calcagno, 2018, for an overview).

⁸⁹While Crawford and Yurukoglu (2012) focus on bundling in the retail market, De Cornière and Taylor (2019) argue that a seller of multiple ad-financed products may improve its bargaining position in the negotiation with a distribution platform by bundling at the wholesale level.

5.1 Richer price structures

The standard Bertrand oligopoly model discussed above combines several assumptions typically encountered in the literature on competing platforms. First, platforms set a single price on each side that can take the form of a participation fee or a transaction fee, without transferability.⁹⁰ Second, demand on each side adjusts to changes in prices to always yield an equilibrium demand configuration in the continuation game that starts after the firms announce their prices.⁹¹

Many platforms can monitor both participation and usage. In this case, the choice of the price instruments is endogenous and the platforms can either charge fixed participation fees or individual transaction fees. As mentioned in Section 2, in the case of monopoly pricing, Rochet and Tirole (2006) allow for general two-part tariffs by focusing on the average price per transaction, assuming away the issue of multiplicity of continuation equilibria in the agents' choice of whether or not to join the platform. The generalization to oligopoly raises challenging problems and it seems difficult to find a single representation capturing the diversity and richness of oligopoly interactions with multi-sided network effects.

The difficulty in the oligopoly case is that the elasticity of the residual demand for any firm on any side depends on the structure of the tariffs used by the competing firms on all sides. Indeed, even if a firm can coordinate participation on both sides with its tariffs, the value of the agents' outside option (buying from the other firms) depends on the other firms' tariffs.⁹²

Suppose agents single-home on both sides and firm B charges a total price $P_i^B = p_i^B + t_i^B q_j^B$ on each side $i = 1, 2$, where p_i^B is a fixed fee and t_i^B is a fee per interaction. A deviation from the equilibrium tariffs by firm A that raises participation on side i and leaves the side- j participation constant requires an adjustment of the side- j price P_j^A equal to $\gamma_j + (\gamma_j - t_j^B) = 2\gamma_j - t_j^B$ per consumer attracted on side i , where γ_j is the extra value a side- j user of platform A derives from the increased participation to platform A on side i whereas the term $\gamma_j - t_j^B$ is the lost interaction value for a side- j user of platform B . The first-order condition for platform A 's prices (assuming it faces no chicken & egg problem) can be written as

$$P_i^A - [c_i^A + \sigma^A q_j^A - 2\gamma_j q_j^A + t_j^B q_j^A] = \mu_i^A(\mathbf{P}_i; \mathbf{q}_j). \quad (19)$$

where again μ_i^A is the inverse-semi-price elasticity of the residual demand of platform A on side i , holding q_j^A , q_j^B and P_i^B constant. In this approach to pricing, the tariff offered by the competing platform B determines the relevant opportunity cost of bringing a marginal consumer on board. This is another illustration of the point made by Rochet and Tirole (2003) that, in two-sided markets, the price structure matters as well as the price level.

The equilibrium allocations then depend on whether the firms set flat prices or two-part tariffs with a positive slope. For instance, when platforms compete in tariffs, Armstrong

⁹⁰We define transferability as in the matching literature as the possibility for a pair of agents interacting on a platform to negotiate a transfer between them to redistribute the fees.

⁹¹See Correia-da Silva, Jullien, Lefouili, and Pinho (2019) for a discussion of this assumption and a few issues related to it.

⁹²The complexity here is similar to the one noticed in the competing-principals literature – see, e.g., Martimort and Stole (2009).

(2006) shows the existence of a continuum of equilibria indexed by the level of the interaction fees t_i^k . To see that, suppose that demand is as in (11) and focus on the case of a symmetric equilibrium where quantities are $q_i^k = 1/2$, $k = A, B$, $i = 1, 2$. Consider a symmetric tariff in which the participation fees are $p_i^A = p_i^B = p_i$ and the interaction fees are $t_i^A = t_i^B = t_i$, with t_i not too large. In a symmetric equilibrium with full coverage, the total price paid by each side- i consumer is $P_i = p_i + t_i/2$. Under suitable concavity conditions, any symmetric tariff such that

$$p_i - c_i^A - \sigma^A/2 + \gamma_j + t_i/2 - t_j/2 = \frac{\Phi_i(0)}{\Phi_i'(0)}$$

can be sustained in a symmetric equilibrium, where Φ_i is the cdf of the side- i distribution of the stand-alone value differential between the two platforms (Armstrong, 2006).

Reisinger (2014) shows however that the multiplicity disappears if there is enough heterogeneity in the interaction benefits γ_i as well as in the stand-alone values. In his model, different types of consumers have different trade volume and thus pay a different total price for the same transaction fee. Each part of a two-part tariff then induces a different demand reaction, reflecting the mix of the different types of marginal consumers. Optimal subscription fees and transaction fees are tailored to these different elasticities, leading to unique equilibrium tariffs.

Weyl (2010) and White and Weyl (2016) argue that platforms have enough pricing instruments to overcome any coordination problem. They define the concept of *Insulated Equilibrium*, based on the idea that prices can be designed so that participation on each side is independent of the participation decisions on the other side. More specifically, they consider an *insulating tariff system* whereby the tariffs set by all firms make the choice of whether or not to join a platform invariant in the agents' beliefs about the participation of the other side. In other words, given the equilibrium tariffs, each agent's participation decision is determined by a dominant strategy. Formally, the notion of dominance applies to each agent in the case of homogeneous interaction benefits or to a representative consumer on each side in the case of heterogeneous interaction benefits. In either case, the insulating tariff system is given by $\mathbf{P}_i(\mathbf{q}_j) = (P_i^A(\mathbf{q}_j), P_i^B(\mathbf{q}_j))$ with

$$\frac{\partial D_i^k(\mathbf{P}_i(\mathbf{q}_j); \mathbf{q}_j)}{\partial q_j^l} = 0 \text{ for } k, l = A, B.$$

The concept of an insulating equilibrium however implicitly assumes coordination between the firms. Indeed whether the tariff of a firm is insulating or not depends on the tariff set by other firms. In general, a firm may not find it profitable to offer an insulating tariff if it expects the other firms to offer tariffs other than the insulating ones, which raises the usual question of equilibrium selection. Lastly, while insulating tariffs are quite simple in the case of homogeneous interaction benefits, these tariffs become highly complex in the case of heterogeneous interaction benefits. This is because an insulating tariff system must preserve for the marginal consumer the value differential between the platform and the consumer's outside opportunities. This requires controlling the value offered by the platform to the representative consumer in a manner that compensates for any change in outside options.⁹³

⁹³Hence, the firm must have enough instruments to implement optimal prices and these instruments

In the duopoly model with homogeneous interaction benefits discussed above, things are simpler and an insulated equilibrium arises when each firm selects a two-part tariff of the form $p_i^k + \gamma_i q_j^k$. When both firms offer such tariffs, all externalities are eliminated through taxation and the characterization of the insulated tariffs is given by the following first-order conditions:

$$P_i^k - [c_i^k + \sigma^k q_j^k - \gamma_j q_j^k] = \mu_i^k(\mathbf{P}_i; \mathbf{q}_j). \quad (20)$$

The difference with Condition (12) is that the term $\gamma_j q_j^k$ is counted only once instead of twice. Insulated tariffs eliminate the feedback loop between sides discussed above. When a firm attracts one more consumer on side i , the competitor's tariff responds by reducing the price on side j so as to preserve the competitor's value, which reduces the platform's ability to raise its own price on side j . Because the competitor's tariff taxes interaction benefits, demand is less elastic and prices are higher than under flat access fees. Hence, tariffs that solve the consumers' coordination problem may not benefit consumers. On the contrary, coordination failures may protect consumers against excessive use of market power.⁹⁴

Earlier versions of insulating tariffs appear in Caillaud and Jullien (2003). The paper shows that competition in two-part tariffs between homogeneous platforms leads to an efficient outcome with zero profits (see also Armstrong, 2006, for a similar point). With complex tariffs, a platform can tax all the surplus it creates for its members and redistribute it through flat registration fees. By doing so, the firm ensures its customers against the risk that participation does not meet expectations. When all firms offers such tariffs, network effects are internalized and coordination failures vanish.

The concept of insulating tariffs relies on the existence of multiple best-replies. In other words, it assumes that, given other firms' tariffs, each firm has several tariffs that would achieve the same revenue. The concept then implicitly assumes that each firm breaks indifference by choosing the tariff that minimizes the risk of consumers' mis-coordination. However, whether this multiplicity of best replies remains in cases involving more realistic assumptions than the base model is questionable. For instance, in Reisinger (2014) (discussed above), the platforms prefer to exploit the richness of tariffs to achieve second-degree price discrimination and there is a unique optimal tariff that may not be insulating.⁹⁵

must be sufficiently flexible (see the discussion in Veiga, Weyl, and White, 2017). White and Weyl (2016) and Veiga, Weyl, and White (2017) state that contingencies can be achieved with dynamic pricing, with firms adjusting their price to market conditions over time. As mentioned in Section 2, for the case of a monopoly platform, Cabral (2019) finds one particular instance (i.e. homogeneous consumers with $v_i = 0$) where optimal dynamic tariffs mimic insulating tariffs, although the tariffs in his analysis reflect exploitation motives rather than insurance motives. In the case of competition, dynamic strategic considerations must also be accounted for, which casts some doubts about the dynamic justification of insulated equilibria.

⁹⁴To illustrate this point, consider a monopoly platform facing "unfavorable beliefs". Consumers benefit if the platform must subsidize participation on one side to overcome the unfavorable beliefs. When, instead, consumers coordinate efficiently, the platform can extract the full surplus from the consumers by charging high prices on both sides.

⁹⁵See also Halaburda and Yehezkel (2013) who show how to use tariffs to screen agents with adequate choices of transaction fees.

5.2 Quantity competition

An alternative to the model of price competition considered thus far is the Cournot model in which platforms choose quantities instead of prices. Indeed, the assumption in Rochet and Tirole (2006) that the monopoly can select the desired demand configuration among all allocations consistent with the posted prices amounts to assuming that the monopoly chooses the participation on each side, which then uniquely determines the prices. This observation suggests that having platforms choosing quantity (i.e., the participation) on each side instead of prices may bypass some of the complexity of the pricing games discussed above. To our knowledge, Katz and Shapiro (1985) are the first to introduce network effects in the Cournot model of quantity competition. The Cournot equilibrium concept has also been used in a limited number of papers on two-sided markets. An early example is Schiff (2003) who finds, in a setting with homogeneous interaction benefits, that a monopoly is socially preferable to a duopoly because it permits each side to interact with a larger number of agents on the opposite side. Gabszewicz and Wauthy (2014) show that heterogeneous interaction benefits may lead to asymmetric market shares, despite the platforms being symmetric *ex ante*.⁹⁶ That paper also finds that the equilibrium participation and welfare are higher under the Cournot equilibrium than under price competition with passive beliefs (see also the discussion below). Correia-da Silva, Jullien, Lefouili, and Pinho (2019) analyze the impact of platform mergers in a Cournot setting with homogeneous interaction benefits within each side.

The Cournot model should be interpreted as imposing a specific form of firm’s *conduct*. Its key assumption is that each firm anticipates that its rivals will maintain their quantities at a given level when the platform under consideration changes its practices and that equilibrium prices will clear the market at targeted quantities.⁹⁷

Models of Cournot competition can be more tractable than models of price competition in the presence of cost asymmetries or if there are more than two platforms. In particular, the Cournot model generates predictions for externality-adjusted prices that make it possible to determine the welfare effects of a merger. To see that, consider a market with K firms and suppose that demand on each side depends only on the hedonic prices on that side, $p_i^k - \gamma_i q_j^k$, $k = A, B$. Invert the demand to obtain each platform’s hedonic price on side i as a function of participation to all platforms: $p_i^k - \gamma_i q_j^k = Z_i^k(\mathbf{q}_i)$. Then each platform’s profits can be written as

$$\sum_i \left(Z_i^k(\mathbf{q}_i) + \gamma_i q_j^k - c_i^k \right) q_i^k - \sigma^k q_1^k q_2^k.$$

⁹⁶The logic is similar to the one behind the emergence of vertical quality differentiation as a way of relaxing competition. The quantity chosen on one side is the quality perceived on the other side. As price competition would erode profit for equal quality, platforms are better off differentiating their offers, with one firm offering a high quantity catering to consumers with high interaction benefits and the other a low quantity catering to consumers with low interaction benefits.

⁹⁷As mentioned above, in media markets this may be due to pricing per user. In some two-sided markets, capacity is a key strategic variable (e.g., ad-finance media platforms committing to ad-space and auctioning ad inventories; other examples are nightclubs, shopping centers and exhibition halls). In such cases, Cournot competition is a natural way of modeling platforms’ behavior, at least on one side.

In the Cournot model of platform competition, each firm k chooses quantities (q_1^k, q_2^k) and prices adjust so that demand equals supply on each side and for each firm, $p_i^k = Z_i^k(\mathbf{q}_i) + \gamma_i q_j^k$. It is then easy to see that equilibrium quantities and prices must satisfy

$$p_i^k - \left[c_i^k + \sigma^k q_j^k - \gamma_j q_j^k \right] = \mu_i^k(\mathbf{p}_i; \mathbf{q}_j) \quad (21)$$

where $\mu_i^k(\mathbf{p}_i; \mathbf{q}_j)$ is as defined above. Comparing the above formula with the first-order conditions for insulated equilibria discussed above, we see that, under homogeneous interaction benefits and regularity conditions ensuring equilibrium existence and uniqueness, the outcome of Cournot equilibrium and Insulated Equilibrium coincide. However, this coincidence should not be expected to hold in more general situations.

The concepts of Insulated Equilibrium and of Cournot equilibrium are related insofar as they both reduce the effect of two-sidedness on prices by mitigating the internalization of demand feedback effects by the platforms. However, the two concepts differ when interaction benefits are heterogeneous, because the insulating tariffs insure only a “representative” consumer against the risk related to the other side’s participation, whereas the Cournot equilibrium entails insurance for all users.⁹⁸

It is worth pointing that many papers in the media literature assume that media platforms set quantities on the advertiser side.⁹⁹ The literature is surveyed by Anderson and Jullien (2015) and Peitz and Reisinger (2015). For instance, Kind, Nilssen, and Sørsgard (2007) build a model in which TV stations offer advertising space to advertisers and (free) content to viewers. They show that there is too little advertising when the channels’ programs are close substitutes and that the more viewers dislike ads, the more likely it is that social welfare is increasing in the number of channels. Crampes, Haritchabalet, and Jullien (2009) and Peitz and Valletti (2008) consider similar models but investigate also the case in which platforms charge viewers. This allows them to compare the free-to-air and pay-tv business models (see also Calvano and Polo, 2020). Rysman (2004) utilizes a quantity setting assumption in an empirical study of Yellow Pages directories in order to avoid the multiple equilibria that arise under price setting. A common feature of these papers is that consumers are assumed to single-home. Anderson, Foros, and Kind (2019) relax this assumption and show that a number of puzzles related to entry, merger and profitability identified in the previous literature on media can be resolved by allowing for multi-homing consumers.

5.3 Passive Beliefs

Coordination plays an important role in platform competition. As coordination largely depends on agents’ beliefs, the way beliefs are formed and modeled plays a crucial role in the analysis of platform markets.

All contributions using static models de facto assume that consumers’ expectations about participation decisions are fulfilled in equilibrium. The equilibrium concept must

⁹⁸As pointed out informally by Glen Weyl, this property impedes the existence of a Cournot equilibrium with multi-dimensional heterogeneity, which is one motivation for the introduction of the concept of insulated equilibria.

⁹⁹In a competitive bottleneck model, this is equivalent to setting an ad price per consumer.

then include a conduct assumption for firms, as discussed above, and an assumption about the way consumers coordinate when firms act out of equilibrium. The most prominent view (Caillaud and Jullien (2003), Rochet and Tirole (2003), Armstrong (2006), Weyl (2010)) is that, when prices change, consumers adjust their beliefs and correctly anticipate the adjustment in the participation level on the other side. This property is sometimes referred to as rational expectations, or reactive beliefs, although neither expression seems ideal. Some contributions take the alternative view that consumers hold fixed beliefs about the other side’s participation, that do not vary when prices change.¹⁰⁰ These beliefs are referred to as “passive” and their interaction with network effects was first examined in Katz and Shapiro (1985). With passive beliefs, because the side- j agents’ expectations of the side- i demand do not vary when a platform changes its side- i prices, then opportunity costs no longer incorporate network effects. Indeed profits can now be written as

$$\Pi^k = \sum_i \left(p_i^k - c_i^k \right) D_i^k \left(\mathbf{p}_i; \mathbf{q}_j^e \right)$$

where \mathbf{q}_j^e is the side- j participation expected by the side- i users and is not affected by the firm’s strategy. It follows that the first-order condition for the side- i prices under passive beliefs reduces to the standard oligopoly pricing rule

$$p_i^k - c_i^k = \mu_i^k \left(\mathbf{p}_i; \mathbf{q}_j \right).$$

In this case, two-sidedness affects prices only through the effect of interaction benefits on the elasticity of the residual demands, not through the firms’ incentives.

In the context of the competitive bottleneck model, Hagiu and Halaburda (2014) show that monopoly profits are lower under passive beliefs than under reactive beliefs. This is because, with passive beliefs, platforms de facto cannot credibly commit to the prices they set on the other side. Things, however, are different in a duopoly context. Duopoly profits in their competitive bottleneck model are higher under passive beliefs than under reactive beliefs because residual demands are less elastic in the absence of feedback effects (see also Belleflamme and Peitz, 2019b, for a model with single-homing). While interesting, a concern with the assumption of passive beliefs is that, because platforms take users’ expectations about participation as given, they do not account for the value that a user creates for the other side(s) of the market when setting prices or quantities. Thus, this approach ignores a key driver of the differences between one-sided and multi-sided markets (see the discussion in Rochet and Tirole, 2003, 2006): the (imperfect) internalization of network externalities.¹⁰¹

6 Matching design

6.1 Second-degree price discrimination and matching design

The literature surveyed in the previous sections is not concerned with the specific process by which agents on board are matched to agents from the other side of the market. In particu-

¹⁰⁰See Gabszewicz and Wauthy (2014) and Griva and Vettas (2011). Hurkens and López (2014) study the role of expectations in the context of mobile telephony and show that passive beliefs provide better modeling of actual behaviors.

¹⁰¹Belleflamme and Peitz (2019b) obtain partial internalization by considering a mix of passive and reactive beliefs reflecting different information of consumers on prices.

lar, the latter process is treated as exogenous to the platform. Some models are compatible with one-to-one matching in a buyer-seller context in which buyers are interested in one specific product version, provided by at most one seller. In these models, the probability by which buyers find their desired product on the platform is naturally increasing in the sellers’ participation on the other side of the market. However, in these models, the platform does not engage in any activity to favor the matching between buyers and sellers other than influencing the participation of each side through the choice of its prices.

Over the last few years the literature on platform markets has started considering models in which platforms engage in more sophisticated design, matching participating agents in a customized manner. This change in modeling reflects the observation that many platforms engage in discriminatory practices whereby they match different agents from the same side to different sets of agents from the opposite side of the market. For example, ad-exchanges match advertisers to only a subset of the content providers that join the exchange. Similarly, the packages offered by most cable TV providers (the platforms) are designed with a similar intent of price discriminating among the viewers by granting them access to different subsets of the channels that join the cable providers. In a similar vein, Business-to-Business platforms typically match vendors to a subset of the procurers in their network. Similar discriminatory practices are used by video-game consoles when mediating the interactions between gamers and video-game developers, and by employment agencies when mediating the interactions between employers and job seekers.¹⁰² For models of price discrimination in two-sided markets, see Halaburda and Yehezkel (2013), Reisinger (2014), Gomes and Pavan (2016, 2019), and Jeon, Kim, and Menicucci (2021).

The possibility of offering differential access to the other side of the market opens the door to *matching design*. In this subsection, we review some of the results in this literature, focusing on existing work where matching allocations are supported by payments.¹⁰³ In addition to accommodating for price discrimination, the literature on matching design also allows for richer preference structures and for the possibility that agents differ not only in their preferences but also in their attractiveness, that is, in the utility they bring to those agents they are matched with.

Consider the baseline model of Section 2, where a single platform mediates the interactions between two sides of a market, but now assume that the interaction benefit that a side- i agent of type θ_i derives from interacting with a side- j agent of type θ_j depends on the latter agent’s type.¹⁰⁴

Think of the platform as offering to each side of the market a sophisticated tariff specifying the price asked to each agent for each possible set of matches with the agents from

¹⁰²Most of the literature on two-sided markets assumes that platforms price discriminate across sides but not within side. Price discrimination within the same side has been considered in the networks literature. See, for example, Belleflamme and Peitz (2020), Bloch and Quérou (2013), Candogan, Bimpikis, and Ozdaglar (2012), Chen, Zenou, and Zhou (2018) and Fainmesser and Galeotti (2016).

¹⁰³This is in contrast to the rich literature examining the stability of matching allocations in environments without payments. See Roth and Sotomayor (1992) and Roth (2008) for a review of the earlier literature, and Kojima (2017) and Pathak (2017) for more recent developments. For a connection between the two literatures, see also the book “Market Design” by Haeringer (2018) and the chapter book “Two-sided Markets and Matching Design” by Gomes and Pavan (2021).

¹⁰⁴For models with competing matchmakers see Damiano and Li (2008) and Halaburda, Piskorski, and Yildirim (2018).

the other side of the market. The equilibrium outcomes induced by any pair of such tariffs can be conveniently described by a pair of *matching rules* $(s_i(\cdot))_{i=1,2}$ and a pair of *payment rules* $(p_i(\cdot))_{i=1,2}$, one for each side. The matching rule s_i specifies, for each type $\theta_i \in \Theta_i$, the set of types $s_i(\theta_i) \subset \Theta_j$ from the opposite side of the market that each side- i agent of type θ_i is matched to. The payment rule p_i specifies for each side- i type θ_i the total payment $P_i(\theta_i)$ that type θ_i makes to the platform, where the payment can be positive for some types and negative for others. As in other mechanism-design problems, such functions should be interpreted as describing the equilibrium allocations generated by the agents' choices on each side of the market.

A feasible pair of matching rules $(s_i(\cdot))_{i=1,2}$ must satisfy the *reciprocity condition* according to which type θ_j from side j is matched to type θ_i from side i if and only if type θ_i is in type θ_j 's matching set, $i, j = 1, 2, j \neq i$:

$$\theta_j \in s_i(\theta_i) \Rightarrow \theta_i \in s_j(\theta_j). \quad (22)$$

The utility that each side- i agent of type θ_i derives from paying $P_i(\theta_i)$ to the platform and being matched to a set of types $s_i(\theta_i)$ from the opposite side is equal to $u_i(s_i(\theta_i)|\theta_i) - P_i(\theta_i)$, with the real-valued function u_i describing the agent's gross payoff.¹⁰⁵

As types are the agents' private information, the matching and the payment rules must be *individually rational* and *incentive compatible*, meaning that each type θ_i must prefer the allocation $(s_i(\theta_i), P_i(\theta_i))$ to his outside option (assumed to yield a payoff equal to zero) and to any other available allocation, in particular to the equilibrium allocation $(s_i(\theta'_i), p_i(\theta'_i))$ of any other type θ'_i .

To simplify things, we assume here that the costs to the platform of getting agents on board are equal to zero, so that $c_i = 0, i = 1, 2$.

6.1.1 One-to-one matching

Consider first markets in which matching is one-to-one, meaning that each agent can be matched up to at most one agent from the opposite side. One-to-one matching typically arises in markets where the two sides are buyers and sellers and in which buyers have unit demands and sellers have unit supply. Also assume that the type of each agent is unidimensional and drawn from $\Theta_i = [\underline{\theta}_i, \bar{\theta}_i] \subset \mathbb{R}$ according to an atom-less distribution F_i , independently across agents.

The gross utility that each side- i agent of type θ_i derives from interacting with a side- j agent of type θ_j is equal to

$$\phi_i(\theta_i, \theta_j), \quad (23)$$

where the function ϕ_i is strictly increasing in both arguments and supermodular.¹⁰⁶ An example often considered in applications is $\phi_i(\theta_i, \theta_j) = \theta_i\theta_j$, with $\Theta_i \subset \mathbb{R}_+, i = 1, 2$.

To make things simple (but interesting), assume that the cost σ that the platform incurs for each match it creates is such that it is inefficient to match pairs of agents with

¹⁰⁵Importantly, the analysis accommodates for the case where the matching rule is stochastic, in which case $s_i(\theta_i)$ is a lottery over the collection of feasible subsets of Θ_j . We describe the rule as deterministic here to facilitate the exposition.

¹⁰⁶In addition, the function is typically assumed to be differentiable and equi-Lipschitz continuous which guarantees that the equilibrium payoffs are well behaved.

the lowest types, but it is efficient to match those with the highest types: $\sum_{i=1,2} \phi_i(\underline{\theta}_1, \underline{\theta}_2) < \sigma < \sum_{i=1,2} \phi_i(\bar{\theta}_1, \bar{\theta}_2)$.

A deterministic matching rule of special interest is the (*truncated*) *positive assortative* one. Formally, for any $U \in [0, 1]$, let $F_i^{-1}(U) \equiv \inf\{\theta_i \in \Theta_i : F_i(\theta_i) \geq U\}$ be the generalized inverse cdf of the type distribution on side i . Say that the pair of matching rules $\{s_i(\cdot)\}_{i=1,2}$ is truncated positive assortative if, for each $i = 1, 2$, there exists a threshold $\hat{\theta}_i \in \Theta_i$ such that

$$s_i(\theta_i) = \begin{cases} (F_j)^{-1}(F_i(\theta_i)) & \text{if } \theta_i \geq \hat{\theta}_i \\ \emptyset & \text{if } \theta_i < \hat{\theta}_i. \end{cases}$$

Hence, any agent from side i with a type below $\hat{\theta}_i$ is excluded (that is, he is matched to nobody), whereas any side- i agent with a type above $\hat{\theta}_i$ is matched in a *positively assortative* manner. That is, higher types from each side are matched to agents from the other side with higher types. The optimality of matching agents assortatively is a direct consequence of the gross payoffs ϕ_i being supermodular. Given the feasibility constraint imposed by one-to-one matching, positive assortativeness requires that an agent whose type θ_i occupies the n^{th} quantile on one side be matched with an agent from the opposite side whose type occupies the same quantile.¹⁰⁷

Next, let

$$\hat{\phi}_i(\theta_i, \theta_j) \equiv \phi_i(\theta_i, \theta_j) - \left(\frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \right) \frac{\partial \phi_i}{\partial \theta_i}(\theta_i, \theta_j)$$

denote the “virtual” utility that type θ_i derives from being matched to type θ_j . This is the true utility $\phi_i(\theta_i, \theta_j)$, adjusted by a term that controls for the cost to the platform of matching the pair θ_i and θ_j , due to the informational rents that the platform must leave to those agents from side i with a type above θ_i .

When the virtual utility functions $\hat{\phi}_i$, $i = 1, 2$, are also nondecreasing and supermodular, on both sides, one can show that both the welfare-maximizing and the profit-maximizing matching rules are truncated positive assortative (variants of this result appear in Damiano and Li, 2007, Johnson, 2013, and Galichon, 2018 – see also Gomes and Pavan, 2021 for an alternative and unifying proof). That is, when the virtual utility functions $\hat{\phi}_i$ satisfy the same properties as the true utility functions ϕ_i , the matching partner of any participating agent under profit maximization is the same as under welfare maximization. The set of participating agents under profit maximization though need not coincide with the efficient set: A profit-maximizing platform may exclude an inefficiently large set of agents from both sides of the market: $\hat{\theta}_i^* \geq \theta_i^e$, $i = 1, 2$, where $\hat{\theta}_i^*$ is the threshold type under profit maximization and θ_i^e is the threshold type under welfare maximization.¹⁰⁸

The above result, however, hinges on the virtual utility functions $\hat{\phi}_i$ being supermodular. When the true utility functions are supermodular but their virtual analogs are not, profit-maximization, in addition to excluding too many agents, may lead to inefficient matching of the participating agents.

Finally, one can show that, both in the case of profit maximization and in the case of

¹⁰⁷Feasibility also requires that $\hat{\theta}_j = F_j^{-1}(F_i(\hat{\theta}_i))$.

¹⁰⁸These threshold solve $\hat{\phi}_1(\hat{\theta}_1^*, \hat{\theta}_2^*) + \hat{\phi}_2(\hat{\theta}_1^*, \hat{\theta}_2^*) = \sigma$ and $\phi_1(\hat{\theta}_1^e, \hat{\theta}_2^e) + \phi_2(\hat{\theta}_1^e, \hat{\theta}_2^e) = \sigma$, respectively.

welfare maximization, the total payment asked to each side- i agent of type θ_i is equal to

$$P_i(\theta_i) = \phi_i(\theta_i, s_i(\theta_i)) - \int_{\hat{\theta}_i}^{\theta_i} \frac{\partial \phi_i}{\partial \theta_i}(\tilde{\theta}_i, s_i(\tilde{\theta}_i)) d\tilde{\theta}_i,$$

all $\theta_i \in \theta_i$, $i = 1, 2$.¹⁰⁹

6.1.2 Many-to-many matching

Now suppose that matching is many-to-many, meaning that each agent from each side may be matched to many agents from the other side of the market, with the only constraint being that the matching be reciprocal (that is, the resulting matching rules satisfy Condition (22)). As in the previous subsection, let $\Theta_i \equiv [\underline{\theta}_i, \bar{\theta}_i] \subseteq \mathbb{R}$. Assume that agents with a positive type like interacting with agents from the opposite side, whereas the opposite is true for agents with a negative type. To avoid trivial cases, assume that $\bar{\theta}_i > 0$ for some $i \in \{a, b\}$.

In addition to parametrizing an agent's preferences, suppose that each agent's type also parametrizes the utility the agent brings to the other side. To capture such a possibility in the simplest possible way, suppose that there exists a function $\lambda_i : \Theta_i \rightarrow \mathbb{R}_+$ describing the ‘‘salience’’ of each side- i agent of type θ_i . An agent of higher salience contributes more to the utility of those agents from the other side who like interacting with the opposite side and generates higher disutility to those who dislike interacting with the other side. For example, in the case of advertising, an ad of higher salience is one delivering a higher utility to those viewers or readers who like advertisement and a larger disutility to those who dislike it.

For any set of type $\mathbf{s} \subset \Theta_j$, let

$$|\mathbf{s}|_j \equiv \int_{\theta_j \in \mathbf{s}} \lambda_j(\theta_j) dF_j(\theta_j) \quad (24)$$

denote the ‘‘total salience’’ of the set \mathbf{s} . Assume that the gross utility $u_i(\mathbf{s}|\theta_i)$ that any side- i agent of type θ_i obtains from being matched to all agents from the opposite side whose type is in $\mathbf{s} \subset \Theta_j$ depends only on the agent's type and on the set's total salience and takes the form

$$u_i(\mathbf{s}|\theta_i) = \theta_i \cdot g_i(|\mathbf{s}|_j). \quad (25)$$

The function $g_i(\cdot)$ is positive, strictly increasing, continuously differentiable, and satisfies $g_i(0) = 0$. Note that the above specification admits as a special case the additive preference structure considered in most of the discussion in the previous sections, which is nested as $\theta_i = \gamma_i$, $\lambda_j(\theta_j) = 1$ all θ_j , and $g_i(x) = x$, $i, j = 1, 2$, $j \neq i$. In other words, the structure considered in the previous sections corresponds to a market in which salience is constant across types (in which case agents care only about the measure of agents they interact with and not the type of such agents) and where preferences are linear in the size of the matching set. The more flexible specification in (25) accommodates for the possibility that different agents contribute differently to the utility they bring to the other side, and that the incremental value of adding a marginal agent to an agent's matching set depends on

¹⁰⁹The result follows from usual envelope-type of arguments standard in mechanism design.

the salience of the latter agent's matching set. The case where the functions $g_i(\cdot)$, $i = 1, 2$, are concave corresponds to the case where agents experience decreasing returns to scale to match quality, whereas the case where these functions are convex corresponds to the case of increasing returns to scale.

However, even when the functions g_i are linear and salience is constant across types, an important point of departure from the analysis in the previous sections is that the platform is not restricted to matching each agent on board to all agents on board from the other side. Instead, the platform may engage in second-degree price discrimination by charging different payments for different matching sets (when preferences take the special structure as in the previous sections, this means that total payments need not be linear in the interactions with the other side).

Finally, note that the multiplicative specification in (25) permits one to retain the usual vertical-differentiation structure from the price-discrimination literature (e.g., Mussa and Rosen, 1978) and conveniently interpret an agent's type as the (positive or negative) value the agent assigns to interacting with agents from the opposite side.

To make things simple, assume that the virtual values $\theta_i - [1 - F_i(\theta_i)]/f_i(\theta_i)$ are nondecreasing in θ_i , as typically assumed in the monopolistic screening literature. Finally, assume that $c_i = 0 = \sigma$ so that there are no costs to get agents on board and/or to match any pair of agents.

Say that a matching rule is a *threshold rule* if there exists a pair of non-increasing functions $T_i : \theta_i \rightarrow \Theta_j$, along with threshold types $\hat{\theta}_i \in \Theta_i$, such that, for any $\theta_i \in \Theta_i$, $i = 1, 2$, the equilibrium matching set takes the form

$$s_i(\theta_i) = \begin{cases} [T_i(\theta_i), \bar{\theta}_j] & \text{if } \theta_i \geq \hat{\theta}_i \\ \emptyset & \text{otherwise,} \end{cases}$$

with $T_i(\theta_i) = \min\{\theta_j : T_j(\theta_j) \leq \theta_i\}$ for all $\theta_i \in [\hat{\theta}_i, \bar{\theta}_i]$.

Thus, under a threshold rule, any type below $\hat{\theta}_i$ is excluded, while any type $\theta_i > \hat{\theta}_i$ is matched to any agent from the other side whose type is above the threshold $T_i(\theta_i)$. Finally, the condition $T_i(\theta_i) = \min\{\theta_j : T_j(\theta_j) \leq \theta_i\}$ means that the pair of rules $s_i(\cdot)$, $i = 1, 2$, satisfies the reciprocity condition (22).

Gomes and Pavan (2016) show that, both in the case of profit maximization and in the case of welfare maximization, a threshold rule is optimal provided one of the following two sets of conditions holds:¹¹⁰

(a) the functions $g_i(\cdot)$ are weakly concave and the functions $\lambda_i(\cdot)$ are weakly increasing, $i = 1, 2$;

(b) the functions $g_i(\cdot)$ are weakly convex and the functions $\lambda_i(\cdot)$ are weakly decreasing, $i = 1, 2$.

In markets satisfying one of the above two sets of conditions, the tariffs that maximize profits (alternatively, welfare) thus induce a form of *negative assortativeness at the margin*: those agents with a lower value for interacting with the opposite side are matched only to those agents from the opposite side whose value for matching is large enough. Furthermore, the matching sets of any pair of agents from the same side are ranked, in the sense that one

¹¹⁰The role of the convexity/concavity of the g_i functions is to guarantee the optimality of threshold rules. It has nothing to do with the optimality/suboptimality of stochastic rules.

is a superset of the other. The above properties reflect the optimal way platforms subsidize interactions in markets in which they can engage in price discrimination (across and within sides) and agents' preferences are their own private information. Importantly, note that the above conditions are satisfied by the linear specification considered in the previous sections. Gomes and Pavan (2016) also show that the above conditions are “almost necessary” in the sense that, when violated, one can identify cases in which either the profit-maximizing or the welfare-maximizing tariffs fail to induce matching sets with a threshold structure.¹¹¹ The intuition for the above result is the following. Consider a welfare-maximizing platform (the problem for a profit-maximizing platform is similar, once values are replaced with their “virtual” analogs). Suppose that the market satisfies the conditions in scenario (a). Clearly, efficiency requires that any agent from side i whose type θ_i is positive be matched to any agent from the opposite side whose type is non-negative. However, efficiency may also require some *cross-subsidization*. That is, the platform may want to assign to a side- i agent with type $\theta_i > 0$ a matching set of total salience

$$|s_i(\theta_i)|_j > \int_{[0, \bar{\theta}_j]} \lambda_j(\theta_j) dF_j(\theta_j).$$

This requires matching the agent also to some of the side- j agents who dislikes interacting with the side- i agents (namely, those with $\theta_j < 0$). When $g_i(\cdot)$ are weakly concave and $\lambda_i(\cdot)$ are weakly increasing, $i = 1, 2$, and the agents' types are the agents' private information (all these assumptions are important), the *least costly* way to deliver such matching intensity is to match type θ_i to all agents from the opposite side whose type is large enough (i.e., above a threshold $T_i(\theta_i)$ satisfying $\int_{[T_i(\theta_i), \bar{\theta}_j]} \lambda_j(\theta_j) dF_j(\theta_j) = |s_i(\theta_i)|_j$). This is because the side- j agents with the highest type are the most attractive ones (by virtue of the monotonicity of the salience functions λ_j) and because using the same side- j agents *intensively* to deliver high match quality to the side- i agents who ask for high match quality is less costly than using different agents from side j (by virtue of the g_j functions being concave).

A similar logic applies to markets satisfying the conditions in scenario (b). That $g_i(\cdot)$ are weakly convex, and $\lambda_i(\cdot)$ weakly decreasing, on both sides, along with the fact that types are private information, implies that the most profitable way of using any agent from side i with a negative type $\theta_i < 0$ is to match the agent to those agents from side j with the highest positive type. This is because such side- j agents are those that benefit the most from interacting with the side- i agent and because such agents are the least salient ones and hence exert the smallest negative externalities on type θ_i .¹¹²

We conclude by discussing possible distortions in the provision of matching services due to market power. For this purpose, let $\nu_j(\theta_i) \equiv g'_j(|[\theta_i, \theta_i]_i) \cdot \lambda_i(\theta_i)$ be type θ_i 's marginal contribution to the total salience of a matching set that contains all side- i types above θ_i . That is, for any side- j type θ_j , $\theta_j \nu_j(\theta_i)$ is the marginal value of expanding the matching

¹¹¹The analysis in Gomes and Pavan (2016) is however more general and accommodates for the possibility that the agents' salience λ_i is a stochastic function of the agents' match values θ_i and that the agents have private information about both θ_i and λ_i .

¹¹²That the side- j agents with the highest type are those that benefit the most from interacting with the side- i agent follows from the fact that such agents necessarily have matching sets with the highest salience, as required by incentive compatibility. Because of the convexity of the functions $g_j(\cdot)$, such agents thus have the highest marginal utility for interacting with the side- i agent.

set of type θ_j starting from $[\theta_i, \bar{\theta}_i]$. Assume that the functions

$$\psi_i^W(\theta_i) \equiv \frac{\theta_i}{\nu_j(\theta_i)} \text{ and } \psi_i^P(\theta_i) \equiv \frac{\theta_i - \frac{1-F_i(\theta_i)}{f_i(\theta_i)}}{\nu_j(\theta_i)}$$

are strictly increasing, $i = 1, 2$. Gomes and Pavan (2016) refer to this property as “match regularity”. To interpret the condition, take the case of profit-maximization. The numerator in $\psi_i^P(\theta_i)$ is type θ_i ’s “virtual type”. This term captures the marginal effect on the platform’s profits of expanding the salience of type θ_i ’s matching set (accounting for the extra surplus that the platform must leave to all side- i agents with type above θ_i , as required by incentive compatibility). The denominator of $\psi_i^P(\theta_i)$, instead, captures the marginal effect on the platform’s profit of adding type θ_i to the matching set of any side- j agent θ_j whose matching set is $[\theta_i, \bar{\theta}_i]$ (i.e., whose threshold is given by $T_j(\theta_j) = \theta_i$). The condition then requires that the contribution of each side- i agent to the platform as a “consumer” increases faster with the agent’s type θ_i than his contribution as an “input.” This condition is thus the analog of Myerson (1981)’s *regularity condition* in a two-sided many-to-many matching environment. As in standard screening problems, the role of the match regularity condition is to rule out bunching. Under such a condition, in a matching environment, bunching occurs only at the bottom of the distribution where it takes the form of exclusion, or at the very top of the distribution where agents are matched to all agents on board from the other side of the market.

Gomes and Pavan (2016) then show that when, in addition to the conditions guaranteeing the optimality of a threshold rule, the agents’ utilities satisfy the above match regularity condition, then, relative to the welfare-maximizing rule, the profit-maximizing rule (a) always excludes a larger group of agents and (b) matches each agent who is on board to a subset of his efficient matching set.

The result follows directly from the fact that a profit-maximizing platform internalizes the effects of cross-subsidization on marginal revenues (which are proportional to virtual values $\theta_i - [1 - F_i(\theta_i)] / f_i(\theta_i)$), rather than their effects on marginal welfare (which are proportional to the true values θ_i). Because the former values are always smaller than the latter (reflecting the extra costs due to informational rents), the matching sets of all agents are a subset of the efficient sets.

Note that, contrary to other mechanism design problems, the matching sets of all types may be distorted, including those types $\bar{\theta}_i$ at the “top” of the distribution. The reason is that the cost of cross-subsidizing such types is strictly larger under profit maximization than under welfare maximization, due to the infra-marginal losses imposed by reciprocity on the opposite side.¹¹³

Strong market power in such markets may thus result in the complete exclusion of many agents from either side of the market, and/or to inefficient matching, whereby those agents on board are either matched to the wrong partners (in the case of one-to-one matching) or to a subset of their efficient matching set (in the case of many-to-many matching). Such distortions call for regulation and government interventions, an area that is receiving growing attention in recent years.

¹¹³A similar result appears in Jeon, Kim, and Menicucci (2021).

6.2 Targeting and third-degree price discrimination

The result about the optimality of threshold rules in Gomes and Pavan (2016) assumes that any pair of agents from side i may disagree on the relative attractiveness of any two agents from side j only when the first agent from side i likes interacting with the side- j agents (formally, has a positive type θ_i) whereas the second agent dislikes it (formally, has a negative type θ_i). This is a direct consequence of the assumed multiplicative structure. The above property is a good representation of markets in which preferences are differentiated only along a vertical dimension. In a recent paper, Gomes and Pavan (2019) show that the optimality of threshold rules extends to various models where agents' preferences are both vertically and horizontally differentiated, under appropriate regularity conditions.¹¹⁴

To see this, suppose each agent's type is now bi-dimensional $\theta_i = (x_i, v_i)$. The first component, x_i , is the agent's "location," whereas the second component, v_i , is the agent's "value for matching," that is, the overall importance the agent assigns to interacting with agents from the opposite side of the market. The component x_i thus parametrizes heterogeneity along a horizontal dimension, whereas the component v_i parametrizes heterogeneity along a vertical dimension.

For example, in the case of cable TV, the vertical dimension may capture the overall importance that a viewer assigns to cable TV, or the overall importance that a channel assigns to reaching viewers, reflecting the channel's expected advertising revenue as well as possible costs stemming from broadcasting rights. In turn, a viewer's location captures the viewer's tastes for different types of programming, whereas a channel's location proxies the type of content broadcasted by the channel.

Gomes and Pavan (2019) assume that the vertical dimensions are the agents' private information. As for the locations, they consider cases where they are publicly observable as well as cases where they are the agents' private information. In the cable TV application, for example, each viewer's ideal type of broadcasting is likely to be his own private information, whereas each channel's broadcasting profile is publicly observable. The analysis in Gomes and Pavan (2019) then shows that, under appropriate conditions, the optimal matching rule continues to have a *threshold structure*, but with a threshold $T_i(v_i; x_i, x_j)$ that depends on the involved agents' locations (a form of third-degree price discrimination). The paper then uses the result to study the effects on matching allocations and welfare of various regulations prohibiting platforms from engaging in third-degree price discrimination (see also Belleflamme and Peitz, 2020, for an alternative model of price discrimination in platform markets).¹¹⁵

Finally, the paper shows how the analysis of the effects of such regulations can also shed light on the effects of a transition from a centralized structure whereby matching is mediated by platforms to a decentralized structure whereby one side (typically the seller side) sets

¹¹⁴The analysis in the works cited above abstracts from within-side network externalities. In a recent paper, Valenzuela-Stookey (2020) shows that threshold rules are also optimal in certain markets combining across-side network effects with within-side network effects (e.g., congestion or, more generally, within-side agents' competition). See also Teh (2019) for an analysis of matching design with within-side network effects.

¹¹⁵The paper develops tools that one can use to derive the optimal matching plans in the presence of such regulatory restrictions. The difficulty is that these restrictions do not allow one to use standard results from mechanism design to represent the payments as a function of the matching allocations. The reason is that such regulations impose specific restrictions directly on the structure to the implementing tariffs.

prices and the other side then chooses the composition of the matching sets. The equivalence between such a decentralized structure and a centralized one subject to the aforementioned non-discriminatory restrictions follows from the fact that, in a decentralized market, sellers cannot bundle their products with other sellers' products. Because bundling is what permits platforms to condition the composition of the matching sets on locations, when the latter are private information, the matching allocations induced by a profit-maximizing platform when third-degree price discrimination (alternatively, bundling) is banned are similar to those in a decentralized market.

6.3 Dynamic arrivals and evolving private information

Another question that is receiving attention in the matching-design literature is how platforms respond to the *dynamic arrival* of information and/or of new agents by adjusting the matching allocations they propose to agents from different sides of the market.

In a couple of recent papers, Fershtman and Pavan (2017, 2021) consider dynamic matching markets in which agents arrive stochastically over time and experience shocks to their match values for specific agents from the other side of the market. Such shocks can be interpreted as reflecting the arrival of new information or variations in the environment that alter the attractiveness of the potential interactions. As a result of such changes, agents on both sides of the market are frequently re-matched. Both the agents' arrivals and the shocks the agents experience to their matching valuations are the agents' private information. Agents face capacity constraints that limit the number of interactions they can entertain in each period. In addition, the platform may also face constraints on the total number of interactions it can accommodate within each period. Such limits may reflect time, resource, or facility constraints, but also capture certain non-separabilities and non-linearities in the agents' preferences.

The combination of the agents' private information with the aforementioned capacity constraints suggests that platforms could benefit from using auctions to determine the relevant matches (e.g., Martens, 2016, and Pinker, Seidmann, and Vakrat, 2003, argue that search rankings and price auctions will soon become the main tools to facilitate online matching).¹¹⁶ However, standard auction formats used for the sale of physical goods or services (e.g., first-price, second-price, English, clock, and double auctions) do not appear appropriate for such matching markets in which agents play the double role of buyers and inputs and where the match values are expected to evolve frequently over time and, as a result, agents need to be re-matched.

Fershtman and Pavan (2021) show that, in many such markets, the profit-maximizing matches can be induced by using appropriate "matching auctions," specifically designed for such markets. Upon arrival, agents are asked to select a "membership status" which is then used to determine the weight assigned to the agents' bids. In each subsequent period, each agent is then asked to submit a bid for each possible agent on board from the other side of the market, with the identity of such agents disclosed by the platform. Each bilateral

¹¹⁶Online ad exchanges such as Google's DoubleClick and Microsoft's Exchange already use auctions to match advertisers with content providers. Advertisers bid repeatedly over time to place their ads on the website of multiple content providers and, over time, content providers ask different ad-specific prices to display the ads (see, e.g., Mansour, Muthukrishnan, and Nisan, 2012. See also the book "Market Design: Auctions and Matching" by Haeringer, 2018).

match is then assigned a “score” that depends on the involved agents’ reciprocal bids, their membership status, and the platform’s cost of implementing the match.¹¹⁷ The matches maximizing the sum of the bilateral scores are implemented in each period, and each agent is charged a total payment that reflects the “externality” the agent imposes on others (by altering the composition of other agents’ matching sets). Importantly, such externalities are adjusted to account for the platforms’ cost of leaving information rents to the agents.

When the arrival of new information is endogenous (as is the case when agents learn the attractiveness of their partners by interacting with them), the score assigned to each match resembles a Gittins “index,” and accounts for the option value of generating new information to be used in subsequent periods. Despite the complexity of the environment and the discrepancy between the agents’ and the platform’s value for experimentation, the proposed matching auctions admit simple profit-maximizing equilibria in which bidding is straightforward and myopic in each period.

Fershtman and Pavan (2021) also contrast matching dynamics under profit maximization to their counterparts under welfare maximization. The paper shows that, when all agents like interacting with all other agents from the opposite side of the market and only the aggregate capacity constraints possibly bind, profit maximization involves fewer and shorter interactions than welfare maximization. This conclusion, however, need not extend to environments where some agents dislike certain interactions and/or where agents’ individual capacity constraints limit the number of matches that each agent can be involved in at each period.

The research on mediated matching surveyed above is still in its infancy. Important work remains to be done to incorporate within-side network effects (namely, the possibility that an agent’s preferences for interacting with those agents from the other side they are matched with depend on who else from the same side interacts with the latter agents). Furthermore, the literature discussed above assumes that the platform can fully control the matching allocations (that is, can prevent agents from changing partners, for example by charging a large enough penalty in case agents do so). Most of the theoretical literature on matching of the last twenty years has abstracted from the role played by matchmakers in platform markets and from the possibility to incentivize matching with payments. It has studied, instead, the implications of appropriate stability notions and other desiderata (e.g., envy-freeness, and strategy proofness). Further bridging the Industrial Organization and the mechanism-design literatures briefly discussed above with the aforementioned literature on non-cooperative matching can shed more light on the properties of matching allocations induced by matchmakers with significant market power.

7 Identification of network effects

The previous sections built up theoretical models of indirect network effects to study pricing equilibria and the implications for profit and welfare. In this and the following sections, we focus on empirical analysis of these market phenomena. In this section, we discuss the identification of network effects. In Section 8, we discuss the implementation of these ideas

¹¹⁷These scores play a role similar to that played by the “compatibility scores” used by ad exchanges to select matches between advertisers and content providers (see, e.g., Moghaddam and Nof, 2017).

and estimation in practice. And in Section 9, we discuss empirical papers that study pricing in the context of platforms and indirect network effects.

For identification, we start with the question of how to estimate the strength of a direct network effect. The preceding sections focused on indirect network effects almost exclusively. However, from an empirical perspective, it will be useful to begin with direct network effects. Recall that direct network effects are when the value of a product depends on other consumers purchasing or using the same product whereas indirect network effects are when the value depends on the provision of some complementary good and that provision depends on other consumers purchasing or using the product. Because direct network effects study adoption of only one product, direct network effects will in some ways be an easier context in which to establish ideas about identification. Also, some empirical papers that study indirect network effects settings actually model them as direct network effects, often because of lack of data. For instance, consumers may value an operating system based on how many applications it has but a researcher that lacks data on the application market may model consumers as making choices based on the installed base of the operating system. The researcher relies on the close association between installed base and application availability.

Another reason to study direct network effects is that direct network effects are more easily comparable to other fields of economics. A central goal of this section is discuss identification of direct network effects from the perspective of the social economics literature, as in Durlauf and Ioannides (2010), Benhabib, Bisin, and Jackson (2011), Angrist (2014) and Graham (2018). This line of research has made substantial progress on the identification and estimation of neighborhood effects and social spillovers, which are modeled in similar ways to direct network effects. Our goal is to utilize the results from the social economics literature to gain insight into the estimation of network effects. Doing so provides perspective on estimating network effects and, in cases where the social economics literature has used approaches that the empirical network effects literature has not, provides new avenues for research. We are particularly interested in the reflection problem of Manski (1993). We discuss the identification of direct and indirect network effects from this perspective.¹¹⁸

In this section, we first present a basic model of direct network effects and we use it to highlight identification problems associated with social effects and network effects. We add consumer heterogeneity to the model and show how that contributes (or does not contribute) to identification, and we consider the concept of contextual effects. We then turn to a series of solutions to the identification problem drawing on ideas from the literature on network effects and social effects. Finally, we introduce a model of indirect network effects which enables us to tie these ideas back to the rest of the chapter.

The starting point for the empirical social economics literature is a *linear-in-means* model that is best represented by continuous choices. The theoretical treatment in the previous sections discussed primarily discrete choices by consumers, such as whether to adopt a technology or not. We return to this question below, but for now, we adjust the presentation of the theoretical model to generate a continuous choice.

¹¹⁸An early version of these ideas appears in Rysman (2019), which covered only Manski (1993) and Manski (1995). This chapter covers the large literature on the reflection problem that has followed, as well as a wider set of empirical papers.

7.1 Direct network effects

In this section, we present a simple model of network effects and discuss issues with identifying the network effect in empirical settings. We start with a model with no consumer heterogeneity. We add consumer heterogeneity to the model and then add a contextual effect, further discussed below.

7.1.1 Base model

We start with the case of direct network effects. Consider a partition of agents into markets $m = 1, \dots, M$. Each market has one network and agents choose how much to use the network in their market. Markets play the role of neighborhoods and determine which agents interact, so network effects operate only within a market. Each market can be thought of as a distinct network. Each market m has a continuum of agents as in the previous sections. Unlike the previous sections, each agent makes a continuous choice $a \in \mathbb{R}$ that incurs a constant marginal price of P_m . The optimal choice for each agent is:

$$a = \beta_0 - \beta_1 P_m + \beta_2 q_m + \zeta_m + \varepsilon \quad (26)$$

The variable $q_m = \bar{a}$ is the average choice of agents in the same market. In the parlance of the reflection problem, β_2 represents the *endogenous effect*, or what we refer to as the direct network effect. We may observe many markets, each with a different P_m and thus a different q_m . In this section, we use ε to represent an idiosyncratic match between the consumer and the platform and assume that $E[\varepsilon] = 0$ in each market m .

A primary concern in both network effects papers and social economics papers is addressing omitted variables or correlation in unobserved terms across agents in the same market. These may arise for a variety of reasons, such as unobserved shocks to a market or sorting by agents with similar unobserved terms into the same markets. These omitted variables are captured by ζ_m , and Manski (1993) terms ζ_m as *correlated effects*. These terms are typically treated as a nuisance to be controlled for or to be eliminated by experimental design. We argue in Section 7.3 that while this makes sense for direct network effects, common approaches to indirect network effects turn the focus of interest onto the correlated effects rather than the endogenous effect.

An equally difficult challenge is addressing simultaneity in choices. Because ε affects the agent's choice, and the agent's choice affects neighbors' choices, ε is correlated with q_m through simultaneity. However, in this simple model, a more significant problem arises than omitted variables or simultaneity. In particular, we recognize that $\bar{a} = q_m$ can be derived by taking the mean of Equation 26:

$$q_m = \beta_0 - \beta_1 P_m + \beta_2 q_m + \zeta_m. \quad (27)$$

We can rewrite this as:

$$q_m = \frac{\beta_0 - \beta_1 P_m + \zeta_m}{1 - \beta_2}.$$

Plugging back into Equation 26 leads to the reduced-form of our model:

$$a = \frac{\beta_0 - \beta_1 P_m + \zeta_m}{1 - \beta_2} + \varepsilon$$

We have more parameters than regressors and thus, the parameters are unidentified. As is well-known, estimating the regression in Equation 26 is guaranteed to find $\beta_2 = 1$.¹¹⁹ That negative result holds even if we eliminated correlated effects by assuming $\zeta_m = 0$ for all m .

What is the economic problem? Intuitively, suppose that we observe the price decrease on a platform. The reduced price leads agents to use the platform more as governed by β_1 , a standard demand effect. The increased use makes the platform more valuable to agents and thus leads to increased usage as governed by β_2 , the network effect. But all we observe in data is that price went down and quantity went up, and we cannot distinguish how much of that increase was due to shifts along the demand curve (β_1) versus the demand curve shifting out (β_2).

One might think that this relationship could be broken by calculating the mean of neighbors' choice only among other neighbors, so that it varies across each agent. However, note that the equation already includes a and q_m , which are sufficient (along with the number of agents) to construct the leave-one-out mean. In this sense, the leave-one-out mean does not introduce any new information. Angrist (2014) is particularly critical of using the leave-one-out mean to address this identification issue.

7.1.2 Consumer heterogeneity

If we observe individual-level explanatory variables, the situation is somewhat improved. Suppose that agents have a type ν that describes some demographic variables such as income or perhaps an individual-specific platform tax or subsidy. Our new version of Equation 26 is:

$$a = \beta_0 - \beta_1 P_m + \beta_2 q_m + \beta_3 \nu + \zeta_m + \varepsilon. \quad (28)$$

Then, q_m is:

$$q_m = \frac{\beta_0 - \beta_1 P_m + \beta_3 \bar{\nu}_m + \zeta_m}{1 - \beta_2}, \quad (29)$$

where $\bar{\nu}_m$ is the mean value of ν in market m . Thus, the reduced-form is:

$$a = \frac{\beta_0 - \beta_1 P_m + \beta_2 \beta_3 \bar{\nu}_m + \zeta_m}{1 - \beta_2} + \beta_3 \nu + \varepsilon.$$

Because β_3 appears by itself in front of ν , it could be identified, which gives the key to separating out the elements in the fraction. This is at least somewhat more promising than the result without ν .

Intuitively, consider the case of a discrete set of consumers in each market rather than a continuum. Equation 28 faces a simultaneity problem because ε affects q_m . Equation 29 provides a path to addressing this problem. Interpreting the equation as a first-stage prediction of the endogenous variable in a two-stage least squares estimator, we have that $\bar{\nu}_m$ predicts the endogenous variable but is otherwise excluded from the agent's decision. That

¹¹⁹Angrist (2014) argues that because we can estimate Equation 26 and find a parameter for β_2 , the parameter is identified (in some sense), but the estimand is not informative about the underlying model.

is, exogenous characteristics of other agents in the market provide instruments for their choices that can be used for identification.¹²⁰

This regression still faces the problem that P_m and $\bar{\nu}_m$ must be distinguished from the correlated effects ζ_m . If ζ_m are addressed with market fixed effects, we have a collinearity problem leading to a lack of identification. If we were willing to restrict the flexibility of these effects, identification follows. For instance, the presentation of Manski (1993) assumes that ζ_m is a linear function of observable market characteristics. Intuitively, unrestricted unobserved market-level shocks make interpreting the effect of q_m impossible, and so one approach is to restrict those shocks in some way.

7.1.3 Contextual effects

Relative to the social economics literature, it is important to keep in mind that we have so far ruled out one important effect. In the presentation of Manski (1993), $\bar{\nu}_m$ enters directly into Equation 28 as a variable of interest, i.e.:

$$a = \beta_0 - \beta_1 P_m + \beta_2 q_m + \beta_3 \nu + \beta_4 \bar{\nu}_m + \zeta_m + \varepsilon.$$

The parameter β_4 measures the *contextual effect* and represents the response of an agent directly to the characteristics of others in the market separately from the effect of their choices a . The contextual effect is often of major interest in social economics studies. For instance, researchers might be interested in the effect on the performance of students in a classroom from having a higher proportion of students from an underrepresented minority population, or from having very wealthy classmates. Research on racial segregation in friendship networks or in housing sorting sometimes focuses almost entirely on the contextual effect and largely ignores the endogenous effect. Graham (2018) discusses this line of research.

In contrast, we are not aware of network effects papers that allow for a contextual effect or attempt to address it. It certainly could be an issue. For instance, we might think that technology/platform adoption by a household depends not only on whether neighborhood households adopt (Goolsbee and Klenow, 2002, has a structure like this) but also on whether neighbors have high income or white-collar jobs. Neighbor demographics might affect the likelihood of hearing about a new technology or the value of the technology if it allowed the household to sell items to neighbors.

It is clear from the modeling above that adding the contextual effect will create significant problems. Going through the steps above leads to:

$$a = \frac{\beta_0 - \beta_1 P_m + (\beta_2 \beta_3 + \beta_4) \bar{\nu}_m + \zeta_m}{1 - \beta_2} + \beta_3 \nu + \varepsilon. \quad (30)$$

Again, this equation allows us to learn β_3 but that is not sufficient to unpack the parameters in the fraction because β_4 also sits in front of $\bar{\nu}_m$. Thus, we cannot distinguish between the endogenous effect and the contextual effect. This identification problem is the heart of the

¹²⁰Note that doing so in the problem without ν implies that q_m is perfectly collinear with P_m and so does not lead to identification.

reflection problem, which Manski (1993) describes as the difficulty in determining causality when the behavior of agents is reflected in their neighbors.¹²¹

Although we cannot generally separate between the contextual and endogenous effects, we can check if the composite term is equal to zero.¹²² If we find that it is not, we can conclude that there is some social effect, either endogenous or contextual. That is the approach of a number of social economics papers. For instance, Angrist and Lang (2004) measure the effect of the number of minority students in a classroom on other students without distinguishing whether the effect is driven by the presence of those students or by the performance of those students.

Because network effects papers do not seem to have considered the contextual effect, we do not allow for a contextual effect for the rest of this chapter, but it is worth considering whether the contextual effect might be important in some circumstances. The focus on contextual effects in the social economics literature suggests it might be relevant in the network effects literature.

7.2 Some solutions for direct network effects

The social economics literature has made progress on finding settings that lead to identification, and we review several approaches here. We also consider whether these approaches have been considered in the network effects literature and how they might be extended. In particular, we consider random assignment, heterogeneous networks, non-linear models, dynamic, and the use of variance of outcomes as a dependent variable.

7.2.1 Random assignment

Perhaps the most substantial focus has been on random assignment to markets. Reviews of these approaches are in Sacerdote (2014) and Graham (2018). A well-known example is Sacerdote (2001), which relies on random matching of college roommates to study the effects of roommate characteristics on GPA. Another example is Angrist and Lang (2004), which studies the effect of a program that assigns students from urban schools to suburban schools. The implementation of the program generates randomness in the number of urban students assigned to each class.

The basic idea of random assignment is that it should eliminate self-selection into the market. If selection is the only source of variation in ζ_m , random assignment can eliminate ζ_m . Note that if we are considering a contextual effect, eliminating ζ_m does not allow us to separately identify β_2 and β_4 in Equation 30. Sacerdote (2001) estimates the overall effect of student characteristics on roommate outcomes. Thus, Sacerdote estimates a reduced-form coefficient that could be positive either because of the endogenous effect or the contextual effect (β_2 or β_4), both of which represent social effects.

¹²¹Manski (1993) writes (pg. 532): “This paper examines the ‘reflection’ problem that arises when a researcher observing the distribution of behaviour in a population tries to infer whether the average behaviour in some group influences the behaviour of the individuals that comprise the group. The term reflection is appropriate because the problem is similar to that of interpreting the almost simultaneous movements of a person and his reflection in a mirror. Does the mirror image cause the person’s movements or reflect them? An observer who does not understand something of optics and human behaviour would not be able to tell.”

¹²²Blume, Brock, Durlauf, and Ioannides (2011) discuss how the contextual effect could be some alternative function of group characteristics other than the mean, which can lead to identification.

Unfortunately, random assignment does not obviously set $\zeta_m = 0$. Recall that ζ_m can represent not only sorting and self-selection but also any common market shocks. For instance, college roommates may be jointly affected by their residence counselor or room location. In a network effects context, random assignment to a network may eliminate selection concerns but an unobserved marketing campaign may still affect how much members use a network. As a result, Sacerdote focuses only on the social spillover effect of characteristics that are determined before the students start at college. For example, even in the case of random roommate matching, regressing GPA on the fraternity membership status of a roommate may be subject to omitted variable bias because GPA and fraternity membership for roommates may be affected by common shocks. However, regressing GPA on a roommate's parent income or distance to their household can be regarded as exogenous due to random roommate matching.

Similarly, Angrist and Lang (2004) focus on the effect of the number of students assigned to a class (a predetermined number) rather than the effect of student performance (an outcome) on one and another. We are not aware of papers in the network effects literature that have been implemented in the same way as Sacerdote (2001) or Angrist and Lang (2004). Given that contextual effects are not of primary interest in network effects models, the focus on contextual effects due to pre-determined heterogeneity is likely not particularly attractive. That is, randomly assigning agents to markets does not allow us to be sure an endogenous effect exists because of common market shocks, and in this sense, is of limited value for identifying a network effect.

However, random assignment can be utilized in a different way in a network effects context. Unlike typical social economics problems, network effects often operate through the size of the network rather than just the average behavior of the people in the network. As a result, even when randomly matching agents to markets of the same size (like the roommates or classroom examples) may be of limited value, randomly matching agents to networks of different sizes could be fruitful. Weiergraeber (2019) is an example of this approach. The paper estimates the network effect between adopters of a given mobile carrier. A cell phone is valuable in part because of how many phones it can call and at what price. The paper studies a period when carriers eliminated fees for calling users who subscribed to a different service. Eliminating fees led consumers to exogenously experience a wider and more accessible network, with the extent of this depending on the consumer's carrier and market shares of other carriers. Taking market structure as exogenous, eliminating out-of-network fees causes consumers to experience random assignment to a new network size. Weiergraeber exploits this exogenous variation to find a significant network effect.

7.2.2 Heterogeneous networks

The linear-in-means model implies that each agent responds to every other agent in the market in a symmetric manner. An alternative is to allow agents to respond to only a subset of agents in the same market. This approach leads us to consider which agents are *linked* to which.

Our presentation follows Bramoullé, Djebbari, and Fortin (2009). Let the number of agents be an integer and let \mathbf{a} with length equal to the number of agents be the vector of actions a , and similarly for $\boldsymbol{\nu}$, $\boldsymbol{\gamma}$ and \mathbf{u} . Let $\mathbf{\Lambda}$ be a square matrix with dimension equal to

the number of agents. Let element λ_{ij} be element $\{i, j\}$ of Λ and let $\lambda_{ij} = 1$ if i is affected by j and $\lambda_{ij} = 0$ otherwise.¹²³ There is no need to distinguish between markets in this approach as the existence of separate markets could be embedded in Λ . Also, Λ need not be symmetric, so some agents may respond to agents that do not respond to them.

In matrix notation, the equivalent to Equation 26 is:

$$\mathbf{a} = \beta_0 - \beta_1 \mathbf{P} + \beta_2 \Lambda \mathbf{a} + \beta_3 \boldsymbol{\nu} + \boldsymbol{\gamma} + \boldsymbol{\varepsilon}.$$

Here, the endogenous effect q from Equation 26 is replaced by $\Lambda \mathbf{a}$.

Solving for \mathbf{a} finds:

$$\mathbf{a} = (\mathbf{1} - \beta_2 \Lambda)^{-1} (\beta_0 - \beta_1 \mathbf{P} + \beta_3 \boldsymbol{\nu} + \boldsymbol{\gamma} + \boldsymbol{\varepsilon}).$$

Multiplying each side by Λ and taking an expectation over \mathbf{a} :

$$E[\Lambda \mathbf{a} | \mathbf{P}, \boldsymbol{\nu}] = (\mathbf{1} - \beta_2 \Lambda)^{-1} \Lambda (\beta_0 - \beta_1 \mathbf{P} + \beta_3 \boldsymbol{\nu} + \boldsymbol{\gamma}).$$

We can interpret this equation as a reduced-form for the endogenous effect and, as in Equation 29, it provides a potential instrument for $\Lambda \mathbf{a}$: $\Lambda \boldsymbol{\nu}$. In this model, an agent's choice is instrumented for by exogenous characteristics of connected agents as these characteristics affect the choices of neighbors but are otherwise excluded from the agents' decision problem.

Augereau, Greenstein, and Rysman (2006) follow an approach like this. As described in Section 3.2, they study the adoption decision of internet-service providers (ISPs) in the face of competing incompatible 56K modems. ISPs may agglomerate on the same standard due to indirect network effects with consumer adoption or ISPs may choose to differentiate from each other in order to avoid competition even if that means serving a smaller market. Because ISPs have overlapping service areas, the matrix indicating which ISPs compete with which (and how much) has a structure like Λ envisioned in Bramoullé et al. (2009). In addition to a two-stage least squares specification as suggested here, the main specification is a structural model of adoption in a Bayesian Nash equilibrium where firms take expectations over the choices of rivals, as suggested by Seim (2006). This structural model exploits exogenous characteristics of connected agents as an instrument.

Another example is Tucker (2008), who studies the adoption of a video-messaging technology within an international investment bank. She allows different employees to have different sets of connections and allows connections to be asymmetric; for instance, the choice of a manager may affect the choice of an employee differently than an employee's choice affects that of a manager. The technology could have alternative uses such as for viewing the World Cup soccer tournament. Thus, whether an employee is interested in soccer (perhaps explained by the country of the employee) is an exogenous shifter that plays the role of an instrument for the effect of the employee's choice on others in its network. Ryan and Tucker (2012) provide dynamic analysis of this setting in a structural model. In the developing country context of Rwanda, Björkegren (2019) estimates a model of adoption of mobile phones allowing for a consumer's network and calling patterns to influence adoption.

¹²³Alternatively, rows of Λ could be scaled so that the agent responded to the mean of actions by connected agents rather than the sum.

An even more straightforward approach arises if some agents are affected by agents that are not affected by any other agents. That is, for some i and $j \neq i$, $\lambda_{ij} = 1$ and $\lambda_{jk} = 0$ for all k . We can take these agents that are not affected by other agents as exogenous actors. There is no need to instrument for the choices of exogenous actors. We can regress the choices of agents on the choices of exogenous actors directly.

In this vein, Gowrisankaran and Stavins (2004) study the adoption of clearinghouse services by banks. The clearinghouse provides a very efficient method of moving funds across banks so Gowrisankaran and Stavins (2004) expect a network effect as more local banks (the ones a bank is most likely to trade with) adopt. Gowrisankaran and Stavins (2004) take several approaches including assuming that large banks are exogenous to small banks, so for instance, Citibank's choice can affect the Bank of Cape Cod but not vice versa.

Interestingly, the approach of Bramoullé, Djebbari, and Fortin (2009) can be used to distinguish between the endogenous and contextual effect. Letting the contextual effect have the same network structure as the endogenous effect, we can write the choice function as:

$$\mathbf{a} = \beta_0 - \beta_1 \mathbf{P} + \beta_2 \mathbf{\Lambda} \mathbf{a} + \beta_3 \boldsymbol{\nu} + \beta_4 \mathbf{\Lambda} \boldsymbol{\nu} + \boldsymbol{\gamma} + \boldsymbol{\varepsilon}.$$

Then, the function for the endogenous term is:

$$E[\mathbf{\Lambda} \mathbf{a} \mid \mathbf{P}, \boldsymbol{\nu}] = (\mathbf{1} - \beta_2 \mathbf{\Lambda})^{-1} \mathbf{\Lambda} (\beta_0 - \beta_1 \mathbf{P} + \beta_3 \boldsymbol{\nu} + \beta_4 \mathbf{\Lambda} \boldsymbol{\nu} + \boldsymbol{\gamma}).$$

Thus, $\mathbf{\Lambda}^2 \boldsymbol{\nu}$ provides a potential instrument for $\mathbf{\Lambda} \mathbf{a}$.¹²⁴ Instruments for the actions of connected agents are exogenous features of the agents that are connected to them. Blume, Brock, Durlauf, and Jayaraman (2015) extend this point in several ways, such as allowing different network structures to determine the endogenous and contextual effect. Bramoullé, Djebbari, and Fortin (2020) review recent developments in the literature on using network structure to identify peer effects.

With all of these approaches based on network structure, it is possible to find spurious results if error terms have the same correlation structure. That would be the case if the correlation of unobserved terms across agents were based on the network structure. Angrist (2014) constructs an example like this when discussing Bramoullé, Djebbari, and Fortin (2009). It could be reasonable to consider such a pattern in some situations. For instance, we might normally assume that error terms have a market-level component as in Equation 26, but if we are studying firms with a network of interaction across different markets, we should ask why firms exhibit their connection pattern, which might lead us to consider unobserved terms with a parallel structure.

7.2.3 Nonlinear models

Many of these problems in identification follow from the linear-in-means set-up. However, many if not most network effects are better thought of as binary or discrete choices, such as the choice to adopt or not. Binary outcome models are often nonlinear and that can lead

¹²⁴A requirement is that $\mathbf{\Lambda}$ and $\mathbf{\Lambda}^2$ are not perfectly collinear. If agents are broken up into markets and interact with everyone in their market, then $\mathbf{\Lambda}$ is a block-diagonal matrix of ones, which is perfectly collinear with $\mathbf{\Lambda}^2$. Thus, this approach requires some asymmetry in who is connected with who.

to new identification results. Important contributions on this are Brock and Durlauf (2001, 2007).

Consider the following model. Let $a \in \{-1, 1\}$ be the choice of a given agent. For instance, we can interpret $a = 1$ as adoption and $a = -1$ as non-adoption. As before, let q_m be the average choice in a market, so naturally, $q_m \in [-1, 1]$. Let a^* be the latent utility parameter defined as:

$$a^* = \beta_0 - \beta_1 P_m + \beta_2 q_m + \varepsilon.$$

Brock and Durlauf (2001) consider parametric assumptions on ε such that this model exhibits multiple equilibria. For instance, there may be one equilibrium with very high adoption and another with very low adoption.¹²⁵ Brock and Durlauf (2001) argue that this nonlinearity generates identification:

Why is there this difference between the binary choice and the linear-in-means frameworks? The answer is that the binary choice framework imposes a nonlinear relationship between group characteristics and group behaviours whereas the linear-in-means model (by definition, of course) does not. ... So, for example, if one moves across a sequence of richer and richer communities, the percentage of high school graduates cannot always increase proportionately with income. (p. 254)

Brock and Durlauf (2007) extend these ideas to consider minimal assumptions on the functional forms for ε and q_m . However, the existence of the identification problem in linear-in-means models means that this approach cannot be extended to a fully nonparametric estimator. All of the papers mentioned in this section up to now (such as Gowrisankaran and Stavins, 2004; Augereau et al., 2006; Tucker, 2008; Goolsbee and Klenow, 2002; Weiergraeber, 2019) and many more use discrete models of technology adoption. However, they tend to focus on sources of identification other than nonlinearity. Whether the network effects literature also provides applications for the ideas about nonlinearity as a source of identification for network effects remains an interesting question.

7.2.4 Dynamics

In general, *installed base* is a key determinant of future consumer decision-making in models of product diffusion under network effects. Section 3.4 discussed the importance of installed base in the theory of dynamic adoption of technologies characterized by network effects. Although this approach does not seem to have been widely explored in the social economics literature, a number of empirical papers on network effects adopt estimation strategies in which current decisions depend on accumulated past decisions of nearby agents. Past choices can sometimes be treated as exogenous to current choices, which makes this particularly attractive. Even in a model in which agents' utility depends on current adoption rather than past adoption, focusing on installed base might be justified by a model in which agents have imperfect information over each other's current choices and use installed base as a predictor of current choices.

¹²⁵Brock and Durlauf (2001) interpret ε as private information to the agent so the agents take expectations over other agents' choices and play a Bayesian Nash equilibrium. We suppress that notation here, as well as the market effect ζ_m .

A number of papers have adopted strategies like this. For instance, Ohashi (2003) and Park (2004) study adoption of VCRs. In this case, indirect network effects between console standards (i.e. VHS vs. Beta) and pre-recorded movies are likely at play but for data reasons, they treat the problem as one of direct network effects between purchasers of the same standard. They use installed base at the end of the previous period as the measure for the size of the network. It is the key variable for measuring the network effect. Ohashi (2003) experiments with different lags of the installed base such as the one year lead (which might capture expectations) and the two-year lag.¹²⁶ Naturally, relying on installed base or lagged sales as an exogenous variable requires a strong assumption on the independence of the error term over time. This may be supported with first-differencing strategies well-known in time series economics.¹²⁷ Kim, Newberry, Wagman, and Wolff (2021) is a recent example studying network effects through installed base, in this case in the context of online fantasy sports leagues. They use local team success as an exogenous instrument for the installed base.

7.2.5 Variance

Graham (2008) focuses on conditional variance rather than conditional means. Network effects are a social multiplier that magnify the response of a population of agents to choices of any one agent. Intuitively, if there is no network effect or social spillover, the variance of outcomes a will be the same regardless of how many agents are in the market. With network effects, the choices a respond to the average shock in the market and if the variance of the average shock increases, so should the variance of a . We can exploit this by noticing that as the number of agents rises, we average over more shocks so the variance of market-level shock decreases. If we observe that the variance of outcomes a differs across markets with different numbers of agents holding all else equal, it must be because of the existence of a social multiplier such as a network effect.

To see this, we return to the model in Section 7.1. This time, assume there is a finite, discrete set of consumers so the average of individual shocks, which we denote $\bar{\varepsilon}_m$, does not average out. That is, $\bar{\varepsilon}_m \neq 0$. Thus, Equation 27 is:

$$q_m = \beta_0 - \beta_1 P_m + \beta_2 q_m + \zeta_m + \bar{\varepsilon}_m$$

Solving for q_m and plugging into the previous equation:

$$a = \frac{\beta_0 - \beta_1 P_m + \zeta_m}{1 - \beta_2} + \frac{\beta_2}{1 - \beta_2} \bar{\varepsilon}_m + \varepsilon$$

We see that $\bar{\varepsilon}_m$ affects a only if the network effect parameter β_2 is non-zero. The variable $\bar{\varepsilon}_m$ is not observed so we cannot estimate this equation to learn about β_2 . However, Graham (2008) recognizes that we can work with the conditional variance rather than the conditional

¹²⁶Park (2004) uses time-format dummies for the main specification which captures all elements of the network effect, such as installed base, available complements (like the number of compatible pre-recorded videos), and expectations about the future. He uses installed base to decompose the elements of the network effect captured by the dummy variable.

¹²⁷Weiergraeber (2019), mentioned above in the context of random assignment, considers an alternative identification scheme utilizing lagged values of other consumers in the market.

mean. The variance of $\bar{\varepsilon}_m$ should fall as the number of agents in a market grows. If we assume that the variance of the first term and the variance of ε stay the same as the number of agents grows, then we can regress the variance of a on the number of agents in the market and interpret a negative coefficient as evidence of the existence of a network effect.

Graham (2008) outlines the assumptions that are necessary for this approach to work and allows for market characteristics to differ in a number of dimensions. He applies his method to an education setting where he can use institutional knowledge to argue that classes with different numbers of students are comparable. For instance, he needs to argue that the variance of teacher quality does not differ with classroom size.

We are not aware of an application of this method to a network effects setting. The requirements of the conditional variance method for the orthogonality of potentially confounding factors are significant for a typical industrial organization application, but it seems feasible. For instance, Tucker (2008) and Gowrisankaran and Stavins (2004) study markets (or offices in the case of Tucker) with different numbers of agents making simultaneous adoption decisions, which would seem to fit into the general approach of Graham (2008).

7.2.6 Other approaches

Naturally, not every paper fits into one of the categories we describe so far. Particularly in industry-specific settings, specialized arguments may be feasible. For instance, an early empirical paper on network effects is Gandal (1994), which shows that spreadsheet software that was compatible with the most popular brand at the time (Lotus 1-2-3) was more highly valued, presumably because it allowed users to more easily trade files. In addition to the approach described above, Gowrisankaran and Stavins (2004) also test whether more concentrated markets are more likely to adopt ACH banking as firms in these markets are better able to internalize network effects.

7.3 Indirect network effects

Many of the techniques we discuss for estimating direct network effects can be applied to indirect network effects but indirect network effects open up new possibilities. Suppose there are two sets of agents, $i = 1, 2$ in each market m . Let $q_{im} = \bar{a}_i$ in market m , the average choice of agents on side i in the market, and let P_{im} be the price of the platform to the set of agents i . We assume actions are given by

$$a_i = \beta_{i0} - \beta_{i1}P_{im} + \gamma_i q_{jm} + \xi_{im} + \varepsilon.$$

Thus, we have dropped β_2 (the direct network effect) and replaced it with the value of activity on the other side of the market. The function captures agents that are indifferent to how many agents on their side of the market use the product but care about how many agents on the other side do so. That is, we study across-side network effects rather than within-side network effects, consistent with the approach in sections 1-6. It is unusual to study models with both direct and indirect network effects, although it could happen. For instance, consumers might value Windows operating system both because there is a variety of software available (an indirect network effect) and because it is easier to share files with neighborhood agents (a direct network effect).

From the perspective of the reflection problem, the correlated effect from the direct network effect case ζ_m is replaced by $\gamma_i q_{jm} + \xi_{im}$. That is, $\gamma_i q_{jm}$ represents a correlated effect, a factor that affects all agents in a market. Whereas in the network effect setting, correlated effects were something to be controlled for in order to estimate the endogenous effect, in the case of indirect network effects, the correlated effect becomes the primary interest. However, in addition to the network effect $\gamma_i q_{jm}$, there may be other market level shocks or selection effects captured by ξ_{im} . In this case, the realization of q_{im} does not create the same sort of problem as before as it does not show up on both sides of the equation:

$$q_{im} = \beta_{i0} - \beta_{i1} P_{im} + \gamma_i q_{jm} + \xi_{im} \quad (31)$$


There is still a simultaneity problem between q_{im} and q_{jm} because the equation holds for both sides of the market. Realizations of ξ_{im} (or in the case of individual-level data, ε) on side i affect q_{im} which affects q_{jm} , leading to correlation between the error term and a regressor.

However, this case provides natural instruments, which are variables that affect one side but not the other. In Equation 31, P_{im} affects q_{im} but does not otherwise affect q_{jm} . Similarly, P_{jm} affects q_{jm} but does not otherwise enter into Equation 31 and thus can play the role of an instrumental variable in estimating that equation. In practice, we often think of prices as endogenous so then we require some other variable that affects one side of the market but not the other, similar to the way P_{im} is written here. We turn to specific examples below, but the key point is that this appears quite feasible.

In the direct network effects case, we could not distinguish between shifts along a demand curve and shifts of a demand curve resulting from a change in price. In the indirect network effects case, we have separate prices for each side of the market. Price changes on side 2 shift out the demand curve on side 1 and do not otherwise lead to shifts along side 1's demand curve, leading to identification.

For intuition, consider Figure 7.3. In this figure, we distinguish between the demand curve and the willingness-to-pay curve. The figure shows the willingness-to-pay at price P'_i . Willingness-to-pay is provided by $D(P_i, q'_j)$ where $q'_j = q_j(P'_i)$. That is, we graph willingness-to-pay by varying price P_i but holding quantity on the other side fixed at the level induced by P'_i . The demand curve is labeled as $D(P_i, q_j(P_i))$. Writing the argument as $q_j(P_i)$ rather than q_j emphasizes the dependence of q_j on P_i through q_i . This is demand as observed by the firm; it maps out how q_i changes as the firm adjusts P_i .

Demand is less than willingness-to-pay above P'_i . A consumer on side i may be willing to pay more than P'_i when price is at P'_i but if price actually rose to that level, the reduction in q_i reduces q_j and so the consumer is no longer willing to pay that amount. The slope of willingness-to-pay is β_1 , whereas the slope of demand also involves γ_i , which causes the slope to be less steep. For instance, the graph shows that when price drops from P'_i to P''_i , the willingness-to-pay shifts upward. As econometricians, we can never actually observe a change in price P_i with quantity on the other side q_j held constant. But instrumental variables approaches allow us to shift q_j while holding P_i constant and thus separately identify these different parameters. Relative to the identification problem in Equation 28,



the following assumptions on that equation provide identification:

$$\beta_2 = 0 \quad \zeta_m = \gamma_i q_{jm} + \xi_{im} \quad E[\xi_{im} | P_{im}, P_{jm}] = 0.$$

The instrumental variables idea is not restricted to indirect network effects. If we can find instruments that affect some agents and not others, we can use instrumental variables in the case of direct network effects. Moffitt (2001) presents a social-spillover model of two agents that highlights this point, and we could see the use of pre-existing heterogeneity in Equation 29 along these lines. However, in practice, this approach is much more prevalent in the case of indirect network effects and the actual instruments used tend to be quite different. We take that up now.

8 Estimating indirect network effects

This section discusses approaches to estimating the strength of indirect network effects from the perspective of implementation. We focus particularly on the use of exclusion restrictions in simultaneous equations representations of platform markets following the approach laid out in Section 7.3. We briefly discuss the use of dynamics as well, similar to Section 7.2.

8.1 Dynamics

Lee (2013) studies consumers buying video game consoles and video games as well as game developers releasing consoles. His model captures that, as more video games become available for a given console, consumers value that console more. Consumers are forward-looking both in their choices over games and consoles. He estimates his model in first-differences which allows him to use lagged variables as instruments. This is, in part, reliant on assumptions of how unobserved terms are not correlated. Lee (2013) argues that because of the long lag between when a game developer commits to producing a video game and when the game is released, these assumptions are likely to hold. Zhou (2017) also studies video games in a fully dynamic game between consumers and firms and relies on similar timing

assumptions in a Bayesian econometric framework. Dube, Hitsch, and Chintagunta (2010) emphasize tipping in a related dynamic model of video games.

8.2 Exclusions in two-sided markets

A number of papers have used exclusion restrictions by assuming that some variables affect outcomes on one side of the market but do not otherwise affect outcomes on the other side. Early papers in this vein are Gandal, Kende, and Rob (2000) and Rysman (2004).¹²⁸ Gandal, Kende, and Rob (2000) study consumer demand for compact disk players and compact disks. One might worry about simultaneity in the determination of these outcomes and the paper uses the evolving cost of building a plant to produce disks as a shifter for the availability of titles. Dranove and Gandal (2003) analyze sales of DVD players during its competition with the DIVX standard. They allow sales of DVD players to depend on the number of titles available. Rather than use installed base of DVD players as an instrument, they use installed base of related technologies such as CDs and camcorders, which should better satisfy exogeneity requirements. An related early paper is Nair, Chintagunta, and Dube (2004), which studies the dynamics of the market for personal digital assistants and software.

In many industries, we see a competitive bottleneck, with single-homing on one side of the market and multi-homing on the other. For instance, consumers may read a single newspaper but advertisers appear in multiple newspapers, or consumers may use a single payment card but merchants accept many cards. A typical empirical model that captures this treats the single-homing side, such as consumers, as making a discrete choice. Let side $i = 1$ be the consumer side and let consumers that select platform $k = 1, \dots, K$ in market m obtain utility:

$$u_{1m}^k = \beta_{10} - \beta_{11}P_{1m}^k + x_{1m}^k\beta_{12} + \gamma_{1m}^kq_{2m}^k + \xi_{1m}^k + \varepsilon_{1m}^k$$

where ε_{1m}^k is distributed according to the extreme value distribution, a popular choice in empirical discrete choice models (see Train, 2009). Here, x_{1m}^k are some exogenous platform-market characteristics and ξ_{1m}^k is a quality index unobserved by the researcher.. We assume each consumer picks platform k such that $u_{1m}^k \geq u_{1m}^z, \forall z$. Assuming there is also an outside option $k = 0$ with utility normalized to $u_{1m}^0 = \varepsilon_{1m}^0$, we have the common logit form for market share s_{1m}^k :

$$s_{1m}^k = \frac{\exp(\beta_{10} - \beta_{11}P_{1m}^k + x_{1m}^k\beta_{12} + \gamma_{1m}^kq_{2m}^k + \xi_{1m}^k)}{1 + \sum_{z=1}^K \exp(\beta_{10} - \beta_{11}P_{1m}^z + x_{1m}^z\beta_{12} + \gamma_{1m}^zq_{2m}^z + \xi_{1m}^z)}$$

Taking the log of the ratio s_{1m}^k/s_{1m}^0 (see Berry, 1994) we obtain that:

$$\ln(s_{1m}^k) - \ln(s_{1m}^0) = \beta_{10} - \beta_{11}P_{1m}^k + x_{1m}^k\beta_{12} + \gamma_{1m}^kq_{2m}^k + \xi_{1m}^k. \quad (32)$$

Let $i = 2$ be the seller side. We assume the demand for the (multi-homing) seller side can be expressed as a linear (or log-linear) demand function:

$$q_{2m}^k = \beta_{20} - \beta_{21}P_{2m}^k + x_{2m}^k\beta_{22} + \gamma_{2m}^kq_{1m}^k + \xi_{2m}^k. \quad (33)$$

¹²⁸Both of these papers are far pre-dated by Rosse (1970), which studies the interaction of readers and advertisers at newspapers although the focus is on the cost function rather than this feedback loop.

If the set of consumers is of size M , then $q_{1m}^k = Ms_{1m}^k$, which links together Equations 32 and 33. For estimation purposes, we assume $\{\xi_{1m}^k, \xi_{2m}^k\}$ are mean independent of exogenous variables.

In the seller demand function (33), we have left off prices of rival products in determining seller demand for platform k . This could be justified if consumers single-home and sellers have constant returns to scale. In that case, selling on one platform does not make selling on another platform more or less attractive and sellers make decisions independently across platforms, which justifies excluding rival prices from demand on the seller side. See Rysman (2004) for a formal derivation of a Cobb-Douglas demand function in a platform setting (in which case, many of these explanatory variables would be expressed in logs). Of course, we could add rival prices to Equation 33 if we would like (and then perhaps side 1 quantities as well).

This set-up clarifies the simultaneous equations nature of the estimation problem and makes clear our identification requirements. The variables q_{1m}^k and q_{2m}^k are endogenous. If we are taking prices as exogenous, then prices serve as instrumental variables because they enter in only one equation. In many cases, we will treat prices as endogenous. In those cases, we need variables in x_{im}^k excluded from x_{jm}^k to identify parameters in equation j . If we estimate additional equations, such as platform first-order conditions, there may be additional excluded variables available from those equations.

There are many examples of this approach. Rysman (2004) studies Yellow Pages directories and models consumer demand for Yellow Pages advertising and advertiser demand for readership as a simultaneous equations problem that leads to excluded variables. For instance, the number of people covered by a directory affects advertiser demand but does not enter into the consumer choice, and some consumer demographics are assumed to affect whether consumers use a directory but not affect the value that advertisers place on those consumers conditional on usage. This general approach has proven popular for media studies, such as Fan (2013) in newspapers and Song (2021) in magazines. Similarly, Clements and Ohashi (2005) study video games and use software age as an instrument for software demand.

A more recent application of these ideas appears in Caoui (2020), who studies the adoption of digital projection technology by movie theaters and the production of movies in the digital format. The paper focuses on France and uses the digital production of movies in the United States as an exogenous shifter of digital movie availability in France, as many US movies are released in France. A related idea is in Gowrisankaran, Park, and Rysman (2014) which studies the adoption of DVD players in response to DVD titles and uses box office outcomes as a shifter of the attractiveness of titles. Like Lee (2013) and Caoui (2020), Gowrisankaran, Park, and Rysman (2014) use a dynamic model of consumer demand. Sokullu (2016) studies the interaction of advertising and readership in newspapers relying on characteristics from each side of the market as excluded variables to identify the indirect network effect. In particular, she allows the network effect to be non-monotonic and utilizes nonparametric instrumental variables in this context.

These ideas are being applied in a wide variety of industries. For instance, Li, Tong, Xing, and Zhou (2017) and Li (2019) study the diffusion of electric vehicles and charging stations. Li, Tong, Xing, and Zhou (2017) assume that gas prices affect only consumer demand and identify the effect of consumers on charging stations. Li (2019) uses, among other

instruments, federal and state subsidies as instruments. As subsidies exist on both sides of the market (subsidies to vehicles and subsidies to charging stations), these can identify the network effect in both directions. Rysman (2007) studies correlation between consumer usage of payment card networks and merchant acceptance, although does not attempt to establish causality. Akerberg and Gowrisankaran (2006) study ACH banking adoption as a two-sided market, analyzing both consumer adoption and bank adoption. They use adoption behavior in other markets as a proxy for costs, which provides identification. They quantify the importance of network effects and the extent of multiple equilibria.

Farronato, Fong, and Fradkin (2020) study digital platforms that match pet owners to pet sitters. From the perspective of a pet owner, a merger between two platforms creates a larger set of potential pet sitters to interact with on the remaining platform than either platform provided before the merger. As platform market shares differ across markets, this provides an opportunity to evaluate the presence of network effects. In a sense, this is a version of random assignment discussed in Section 7.2.1 adapted to a two-sided setting. Pet owners are randomly assigned different increases in the size of the network of pet sitters with which they may interact, with assignment depending on what city the pet owner lives in. Farronato, Fong, and Fradkin (2020) find that in markets where the acquiring platform was dominant, there is little change after the merger. Network effects are apparent in markets where the two platforms were similar in size, although pet sitters of the acquired platform often did not participate after the merger. This evidence of switching costs provides a potential downside to network effects realized through merger.

Similarly, Reshef (2020) exploits a partnership between Yelp and Grubhub. Because market shares differ for these firms across cities, their partnership generates plausibly exogenous variation across cities in the size of the network expansion. Obviously, this version of random assignment is different from the social economics literature in which random assignment typically takes the form of literal randomness imposed at an administrative level, such as random assignment of college roommates. In the platform economics setting, establishing randomness (or more realistically, exogeneity) requires specialized arguments, but conditional on that, the thought experiment is very similar.

9 Empirical work on pricing in platform studies

With our discussion of the identification and estimation of network effects completed, we can now turn to approaches to study pricing in these markets and their results, which more closely connects with the earlier sections in the chapter. We distinguish between two approaches, which loosely correspond to treating price as endogenous or exogenous. In the first approach, the researcher observes many platforms setting prices and studies how observed prices respond to factors such as market structure and determines whether platform behavior is consistent with theory. In this approach, we find papers that are both structural and reduced-form in their econometric techniques. Next, we discuss papers in which the researcher observes one platform or set of platforms with relatively few sets of prices, such as a single pricing schedule used in many geographic markets. The researcher typically models how agents interact on this platform, and studies pricing by analyzing how agents would behave under some counterfactual pricing regime. Modeling agent interactions in ways that make the counterfactual compelling can often lead to very involved models

and this approach tends to be structural. Naturally, the two approaches answer different questions and are complementary for understanding platform markets.

9.1 Price as endogenous

Early papers in the first approach focus on media. Rysman (2004), discussed earlier in the context of identification of network effects, also fits here. The paper observes many markets, each populated by several Yellow Pages directories, which are the platforms, and models quantity (and by extension, price) as the choice variable. It uses a model similar to the one represented in Equations 32 and 33, along with a pricing first-order condition. Rysman (2004) simulates outcomes as the number of competing directories change. More directories reduce market power but also dissipate network effects. The paper finds that the former effect dominates, and so market efficiency is enhanced as the number of directory increases, at least within the range of the data.

Kaiser and Wright (2006) study German magazines, again studying the interaction of readers and advertisers. They estimate a structural model of both consumer and advertiser demand, in this case, assuming that both sides single-home. They use several instrumental variables, in particular, prices of publishers in related markets, which can be interpreted as related to cost-shifters in the spirit of Hausman (1996). They find that profits are collected largely from the advertiser side rather than the consumer side. Higher demand on the advertiser side actually reduces consumer prices as magazines seek to attract these advertisers with large readership, similar to the seesaw effect described in Section 4.3.2. Another paper on media is Argentesi and Filistrucchi (2007), which provides a structural model of Italian newspapers. They extend the approach of Kaiser and Wright (2006) to include a publisher first-order condition, as in Rysman (2004), which allows them to study market power. Chandra and Collard-Wexler (2009), mentioned in Section 4.3.2, study mergers between newspapers and show theoretically and empirically that the logic of two-sided markets implies that prices do not necessarily increase. Song (2021) analyzes pricing, market power, and merger in German magazines accounting for two-sidedness. Fan (2013) emphasizes how a merger leads newspapers to change not only the prices but also quality characteristics such as news content and how these features interact. Jeziorski (2014) embeds a related model of radio stations in a dynamic model of mergers. He finds that a merger results in market power that allows radio stations to lower the quantity and increase the price of advertisements. However, because consumers dislike advertising, this outcome enhances consumer welfare.

Reduced-form approaches often provide compelling tests of whether theoretical models make accurate predictions. Empirical papers confirm the existence of the seesaw effect of Section 4.3.2 in a variety of industries. That is, a shock that tends to raise price on one side often reduces prices on the other side. In media, in addition to Kaiser and Wright (2006) mentioned above, Seamans and Zhu (2014) study how newspaper prices and outcomes respond to the entry of Craigslist, which cannibalizes classified advertisement revenue, an important source of newspaper revenue up until that point. The response is complex. For instance, consumer prices increase and advertiser prices decrease. Their interpretation is that the reduction in demand for classified advertisements reduces the value of consumers leading to an increase in consumer price, and the resulting reduction in readers lowers value,

and thus price, to advertisers.

In telecom, Boik (2016) studies cable systems, which are platforms between content providers, such as local television stations, and consumers. Boik confirms the two-sided intuition that local television stations in markets with high advertisement rates set lower fees to cable systems as the advertising rates incentivize the stations to seek more consumers. Similarly, Genakos and Valletti (2011) find that pricing reductions imposed by regulators in how mobile carriers price in-coming calls lead to increases in price to consumers as the consumers become less valuable. Genakos and Valletti (2011) show how this effect differs across different types of consumers.

In payments, Kay, Manuszak, and Vojtech (2018) study the debit card market in the US. Under the Financial Reform Act of 2010, the Federal Reserve Bank capped the interchange fee on debit cards for banks above a size threshold, which effectively reduced the amount that these banks charge to merchants. Kay, Manuszak, and Vojtech (2018) compare banks around the threshold to document that the regulation led to more fees to consumers. Manuszak and Wozniak (2017) take a similar approach to analyze the nature of fees added to checking accounts and similar products. Carbo-Valverde, Chakravorti, and Rodriguez-Fernandez (2015) study the Spanish payment card market during a period of declining interchange fees, which reduced effective payments from merchants to consumers. Although this likely led to higher consumer fees, the paper shows that it also led to increased merchant acceptance. Overall, Carbo-Valverde, Chakravorti, and Rodriguez-Fernandez (2015) find that consumer adoption increased.

It would seem unintuitive that competition between firms would lead to worse prices for some agents but that is natural in a two-sided context. As described above (see Section 4.3), competition between platforms can lead to better terms for one side but worse terms for the other. The study of local television fees of Boik (2016) finds that in markets with more substantial competition between television delivered by cable companies and telephone companies, local television stations charge lower rather than higher fees. An explanation consistent with a number of theoretical models is that local television stations multi-home and the benefits of competition between platforms (cable and telephone) go to the single-homers, the consumers.¹²⁹

Similarly, Jin and Rysman (2015) analyze sports card conventions, which bring together buyers and sellers of sports memorabilia. Taking the convention organizer as the platform, they observe prices charged to each side at thousands of conventions. They find that pricing responds to competition between platforms much more on the single-homing side. We discussed prices that face a lower bound at zero in Section 2.5.3, and about 50% of conventions offer free admission to consumers. Jin and Rysman (2015) find that when price on the single-homing side hits zero, pricing on the other side shows a stronger response to competition, as predicted by Armstrong and Wright (2007).

¹²⁹Boik (2016) finds that fees to local television stations increase when there is more significant competition from satellite, but he points out that it is likely driven by the advertisement market. Cable competes with local television for local advertisements whereas satellite does not, giving the local television stations incentives to direct consumers to satellite service.

9.2 Price as exogenous

Studies in the previous section use covariation in prices with other variables. The second approach studies a platform with fixed prices (or fixed commission rates) and studies how agents interact on the platform. Typically using structural estimation, these papers study counterfactual outcomes under alternative pricing structures. Several papers study markets with many competing intermediaries, such as brokers. An example is Gavazza (2016) who studies intermediaries in the market for trading used aircraft and looks at outcomes without intermediaries. Barwick and Pathak (2015) study the effects of the commission rate for agents in the Boston real estate market, which is typically fixed across brokers. They find that allowing brokers to compete on commission rates leads to a reallocation of market share to more efficient brokers and substantially more efficient outcomes. Similarly, Hendel, Nevo, and Ortalo-Magné (2009) show in a reduced-form comparison between outcomes on an internet platform for selling homes and the traditional broker market that high commissions drive sellers to the on-line platform, which creates important selection effects in which sellers participate on each platform. Robles-Garcia (2020) models the relationship between mortgage lenders and borrowers with a focus on mortgage brokers that connect them. The paper finds that caps on commissions to brokers increase welfare in part by aligning incentives with consumers rather than lenders. A common theme across these papers appears to be that if the market-wide commission structure is not efficient, competition between brokers does not necessarily correct for that.

Some recent papers apply this method to online markets. Marra (2020) studies an online platform for trading wine. Echoing the selection effects of Hendel, Nevo, and Ortalo-Magné (2009), she shows how fees on one side of the platform affect not only the number of participants but also the quality of those participants, which has implications for agents on the other side. Castillo (2020) provides a detailed model of participants in the ride-sharing platform Uber. In ride-sharing, the value to riders from participating in the platform depends on the number of drivers participating, and vice versa. Castillo considers the effect of eliminating surge pricing, the increase in fees from riders to drivers during high demand periods. He finds that riders in particular benefit from surge pricing, contrary to some of the popular discussion, as they benefit from increased service levels. Rosaia (2021) studies a model of drivers choosing not only whether to drive but also choosing between Uber and Lyft. He studies pricing power of the two companies and finds that merger leads to efficiencies due to the management of traffic congestion. These papers highlight the complex interactions between agents in a platform and the sometimes surprising outcomes that arise from policy interventions that would be more straightforward in a one-sided market.

9.3 Matching

Above, Section 6 highlights the importance of matching and platform design. As pointed out there, the line between indirect network effects and matching can be blurred. For instance, empiricists may model the value of a platform as a simple function of the number of sellers even if the platform plays a substantial role in matching consumers to individual sellers or groups of sellers. Some papers however engage with this matching process and the consumer's individual choice of which seller to utilize. For example, Fox and Bajari (2013) study FCC spectrum auctions as a matching problem between bidders and spectrum

with the government as the “platform” connecting them. The econometrics of matching models is discussed in Fox (2018), who estimates a model of parts buyers and sellers in the automobile market. Crawford and Yurukoglu (2012) and Crawford, Lee, Whinston, and Yurukoglu (2018) study the market for cable television and model the choices of cable systems over what channels to include and consumers’ prices, as well as consumers’ choices over which channels to view. Similarly, Ho (2009), Gowrisankaran, Nevo, and Town (2015) and Ho and Lee (2017) are part of a literature that studies how health insurance companies form bundles of providers (i.e., hospitals) to offer consumers. They also model consumer choice of a provider upon a health event. These papers emphasize the role of bargaining between the intermediary and the sellers and the complex ways in which market power and product differentiation at both the seller and intermediary level affect consumer pricing and choices.

Salz (2020) studies intermediaries in commercial waste collection in New York City and observes contracts reached by buyers and sellers including price. He finds that, because buyers with high search costs gravitate to the intermediaries, the bilateral bargaining segment of the market is particularly competitive. Thus, matching by intermediaries and matching to intermediaries together causes low prices in the bilateral market. Biglaiser, Li, Murry, and Zhou (2020) also highlight selection effects when dealers compete with bilateral bargaining in the context of used cars. They use pricing patterns to infer the importance of addressing asymmetric information in explaining the prevalence of dealers. Buchholz (2021) and Fréchette, Lizzeri, and Salz (2019) study dynamic matching models of taxi cabs in New York City and view the regulator as a platform that can implement policies that can reduce search costs and affect both the entry choice and locations of taxicabs. Both papers find that appropriate policies can substantially increase welfare. Buchholz, Doval, Kastl, Matejka, and Salz (2020) study a Czech ride-hailing platform that allows drivers to bid on riders so riders choose between rides with different prices and waiting times. This study emphasizes price-setting between the buyer and seller which plays a role in the models in Section 6. There is a large literature in finance on the value of centralized markets relative to individual traders. A recent example is Allen and Wittwer (2021) who study the treasury bond market in Canada and observe individual trades and prices.

10 Conclusion

This chapter focuses on pricing in two-sided markets, the most well-studied phenomenon in the growing literature on platform economics. Platforms set prices on each side of the market accounting for the complex interactions between the various platform’s users. Pricing strategies respond to the existence of network effects but also contribute to the realization of such network effects.

We start by reviewing the theoretical literature on monopoly platforms, comparing profit-maximizing to welfare-maximizing solutions, and analyzing different ways in which network effects shape prices and market outcomes.

Next, we turn to competition between platforms, which has been the subject of recent research. We emphasize the (somewhat fuzzy) difference between “competition for the market” and “competition on the market,” and review the different approaches to modeling platform competition. We point to the importance of accounting for the various ways that

agents may interact with the platforms, such as by single-homing or multi-homing. In addition, we discuss the nascent literature on the closely-related topic of matching design, opening the “black box” of network effects and allowing for individualized matching and within-side price discrimination.

Finally, we discuss empirical approaches to the study of these markets. We discuss identification issues in the estimation of network effects. We argue that connecting this literature to the well-developed literature on the estimation of social spillovers provides valuable insights and opens the door to promising new directions for future research. We review a number of empirical papers that rely on a variety of strategies to identify network effects. We review the study of pricing in these markets and show support for several of the theoretical results.

Despite the fact that many issues related to platforms are dynamic in nature, the literature on platforms is largely static; very few works study ignition and dynamic platform competition. More progress is needed to understand the evolution of platform markets. Moreover, most of the literature on competing platforms focuses on symmetric market configurations. More work on the relation between platform size and dominance is needed for policy analysis.

Similarly, more work on the effects of multi-homing when agents have different values for interacting with the other sides is necessary for a better understanding of pricing in such markets.

Discriminatory practices have also become prevalent in many platforms and the literature still lacks a flexible but tractable model accounting for such practices. The literature is also evolving towards a deeper understanding of the interactions between pricing and other aspects of platform design such as feedback and recommendation systems, bundling, integration, ranking systems, marketing, information design, and entry rules.

Lastly, specialized industry studies in areas such as media, finance, transportation, housing, and health insurance, which have been covered only tangentially in this chapter, represent important directions for current and future research.

References

- Acemoglu, D. and M. K. Jensen (2013). Aggregate comparative statics. *Games and Economic Behavior* 81, 27–49.
- Akerberg, D. A. and G. Gowrisankaran (2006). Quantifying equilibrium network externalities in the ACH banking industry. *RAND Journal of Economics* 37, 738–761.
- Acquisti, A., C. Taylor, and L. Wagman (2016). The economics of privacy. *Journal of Economics Literature* 54, 442–92.
- Allen, J. and M. Wittwer (2021). Centralizing over-the-counter markets? Unpublished Manuscript, Stanford University.
- Ambrus, A. and R. Argenziano (2009). Asymmetric networks in two-sided markets. *American Economic Journal: Microeconomics* 1(1), 17–52.
- Ambrus, A., E. Calvano, and M. Reisinger (2016). Either or both competition: A “two-sided” theory of advertising with overlapping viewerships. *American Economic Journal: Microeconomics* 8(3), 189–222.
- Ambrus, A. and M. Reisinger (2006). Exclusive vs overlapping viewers in media markets. Technical report, SFB/TR 15 Discussion Paper.
- Amelio, A., L. Giardino-Karlinger, and T. Valletti (2020). Exclusionary pricing in two-sided markets. *International Journal of Industrial Organization* 73, 102592.
- Amelio, A. and B. Jullien (2012). Tying and freebies in two-sided markets. *International Journal of Industrial Organization* 30(5), 436–446.
- Anderson, S., A. de Palma, and Y. Nesterov (1995). Oligopolistic competition and the optimal provision of products. *Econometrica* 63(6), 1281–1301.
- Anderson, S. P. and S. Coate (2005). Market provision of broadcasting: A welfare analysis. *The Review of Economic Studies* 72(4), 947–972.
- Anderson, S. P., N. Erkal, and D. Piccinin (2020). Aggregative games and oligopoly theory: Short-run and long-run analysis. *The RAND Journal of Economics* 51(2), 470–495.
- Anderson, S. P., Ø. Foros, and H. J. Kind (2018). Competition for advertisers and for viewers in media markets. *The Economic Journal* 128(608), 34–54.
- Anderson, S. P., Ø. Foros, and H. J. Kind (2019). The importance of consumer multihoming (joint purchases) for market performance: Mergers and entry in media markets. *Journal of Economics & Management Strategy* 28(1), 125–137.
- Anderson, S. P. and B. Jullien (2015). The advertising-financed business model in two-sided media markets. In *Handbook of Media Economics*, Volume 1, pp. 41–90. Elsevier.
- Anderson, S. P. and M. Peitz (2020a). Ad clutter, time use, and media diversity. *CEPR Discussion Paper No. DP15130*.
- Anderson, S. P. and M. Peitz (2020b). Media see-saws: Winners and losers in platform markets. *Journal of Economic Theory* 186, 104990.

- Angrist, J. and K. Lang (2004). Does school integration generate peer effects? Evidence from Boston's Metco program. *American Economic Review* 94, 1613–1634.
- Angrist, J. D. (2014). The perils of peer effects. *Labour Economics* 30, 98–108.
- Argentesi, E. and L. Filistrucchi (2007). Estimating market power in a two-sided market: The case of newspapers. *Journal of Applied Econometrics* 22, 1247–1266.
- Armstrong, M. (2006). Competition in two-sided markets. *RAND Journal of Economics* 37, 668–691.
- Armstrong, M. and J. Wright (2007). Two-sided markets, competitive bottlenecks and exclusive contracts. *Economic Theory* 32, 353–380.
- Arthur, W. B. (1989). Competing technologies, increasing returns and lock-in by historical events. *Economic Journal* 99, 106–131.
- Athey, S., E. Calvano, and J. S. Gans (2018). The impact of consumer multi-homing on advertising markets and media competition. *Management Science* 64 (4), 1574–1590.
- Augereau, A., S. Greenstein, and M. Rysman (2006). Coordination vs. differentiation in a standards war: 56k modems. *RAND Journal of Economics* 37, 887–909.
- Bakos, Y. and H. Halaburda (2020). Platform competition with multihoming on both sides: Subsidize or not? *Management Science* 66(12), 5599–5607.
- Baranes, E., T. Cortade, and A. Cosnita-Langlais (2019). Horizontal mergers on platform markets: Cost savings v. cross-group network effects? Unpublished manuscript, University Montpellier.
- Barwick, P. J. and P. A. Pathak (2015). The costs of free entry: An empirical study of real estate agents in Greater Boston. *RAND Journal of Economics* 46(1), 103–145.
- Baxter, W. P. (1983). Bank interchange of transactional paper: Legal perspectives. *Journal of Law and Economics* 26, 541–588.
- Bedre-Defolie, Ö. and E. Calvano (2013). Pricing payment cards. *American Economic Journal. Microeconomics* 5(3), 206.
- Belleflamme, P. and M. Peitz (2019a). Platform competition: Who benefits from multihoming? *International Journal of Industrial Organization* 64, 1–26.
- Belleflamme, P. and M. Peitz (2019b). Price disclosure by two-sided platforms. *International Journal of Industrial Organization* 67, 102529.
- Belleflamme, P. and M. Peitz (2020). Network goods, price discrimination, and two-sided platforms. Technical report, CORE.
- Belleflamme, P. and M. Peitz (2021). The economics of platforms: Concepts and strategy. Forthcoming, Cambridge University Press.
- Belleflamme, P., M. Peitz, and E. Toulemonde (2020). The tension between market shares and profit under platform competition. *CEPR Discussion Paper No. DP15204*.
- Belleflamme, P. and E. Toulemonde (2009). Negative intra-group externalities in two-sided markets. *International Economic Review* 50(1), 245–272.
- Benhabib, J., A. Bisin, and M. O. Jackson (Eds.) (2011). *Handbook of Social Economics*. Elsevier.

- Berry, S. (1994). Estimating discrete choice models of product differentiation. *RAND Journal of Economics* 25, 242–262.
- Biglaiser, G. and J. Crémer (2020). The value of incumbency when platforms face heterogeneous customers. *American Economic Journal: Microeconomics* 12(4), 229–69.
- Biglaiser, G., F. Li, C. Murry, and Y. Zhou (2020). Intermediaries and product quality in used car markets. *RAND Journal of Economics* 51, 905–933.
- Björkegren, D. (2019). The adoption of network goods: Evidence from the spread of mobile phones in Rwanda. *The Review of Economic Studies* 86, 1033–1060.
- Bloch, F. and N. Quérou (2013). Pricing in social networks. *Games and Economic Behavior* 80, 243–261.
- Blume, L. E., W. A. Brock, S. N. Durlauf, and Y. M. Ioannides (2011). Identification of social interactions. In J. Benhabib, A. Bisin, and M. O. Jackson (Eds.), *Handbook of Social Economics*, pp. 853–964. Elsevier.
- Blume, L. E., W. A. Brock, S. N. Durlauf, and R. Jayaraman (2015). Linear social interactions models. *Journal of Political Economy* 2, 444–496.
- Boik, A. (2016). Intermediaries in two-sided markets: An empirical analysis of the U.S. cable television industry. *American Economic Journal: Microeconomics* 8, 256–282.
- Bramoullé, Y., H. Djebbari, and B. Fortin (2009). Identification of peer effects through social networks. *Journal of Econometrics* 150, 41–55.
- Bramoullé, Y., H. Djebbari, and B. Fortin (2020). Peer effects in networks: A survey. *Annual Review of Economics* 12, 603–629.
- Brock, W. A. and S. N. Durlauf (2001). Discrete choice with social interactions. *Review of Economics Studies* 68, 235–260.
- Brock, W. A. and S. N. Durlauf (2007). Identification of binary choice models with social interactions. *Journal of Econometrics* 140, 52–75.
- Buchholz, N. (2021). Spatial equilibrium, search frictions and dynamic efficiency in the taxi industry. In Press, *Review of Economic Studies*.
- Buchholz, N., L. Doval, J. Kastl, F. Matejka, and T. Salz (2020). The value of time: Evidence from auctioned cab rides. Unpublished Manuscript, MIT.
- Cabral, L. (2011). Dynamic price competition with network effects. *The Review of Economic Studies* 78(1), 83–111.
- Cabral, L. (2019). Towards a theory of platform dynamics. *Journal of Economics & Management Strategy* 28(1), 60–72.
- Cabral, L. M., D. J. Salant, and G. A. Woroch (1999). Monopoly pricing with network externalities. *International Journal of Industrial Organization* 17(2), 199–214.
- Caillaud, B. and B. Jullien (2001). Competing cybermediaries. *European Economic Review* 45, 797–808.

- Caillaud, B. and B. Jullien (2003). Chicken & egg: Competition among intermediation service providers. *RAND Journal of Economics* 34, 309–328.
- Calvano, E. and M. Polo (2020). Strategic differentiation by business models: Free-to-air and pay-TV. *The Economic Journal* 130(625), 50–64.
- Candogan, O., K. Bimpikis, and A. Ozdaglar (2012). Optimal pricing in networks with externalities. *Operations Research* 60(4), 883–905.
- Caoui, E. H. (2020). Estimating the costs of standardization: Evidence from the movie industry. Unpublished manuscript, University of Toronto.
- Carbo-Valverde, S., S. Chakravorti, and F. Rodriguez-Fernandez (2015). The role of interchange fees in two-sided markets: An empirical investigation on payment cards. Forthcoming, *Review of Economics and Statistics*.
- Carroni, E., L. Madio, and S. Shekhar (2020). Superstars in two-sided markets: Exclusives or not? Unpublished manuscript, Toulouse School of Economics (TSE).
- Casadesus-Masanell, R. and H. Halaburda (2014). When does a platform create value by limiting choice? *Journal of Economics & Management Strategy* 23(2), 259–293.
- Castillo, J. C. (2020). Who benefits from surge pricing. Unpublished manuscript, University of Pennsylvania.
- Chandra, A. and A. Collard-Wexler (2009). Mergers in two-sided markets: An application to the Canadian newspaper industry. *Journal of Economics and Management Science* 18, 1045–1070.
- Chen, J., U. Doraszelski, and J. E. Harrington, Jr (2009). Avoiding market dominance: Product compatibility in markets with network effects. *The RAND Journal of Economics* 40(3), 455–485.
- Chen, Y. and M. H. Riordan (2008). Price-increasing competition. *The RAND Journal of Economics* 39(4), 1042–1058.
- Chen, Y.-J., Y. Zenou, and J. Zhou (2018). Competitive pricing strategies in social networks. *The RAND Journal of Economics* 49(3), 672–705.
- Choi, J. P. (2010). Tying in two-sided markets with multi-homing. *The Journal of Industrial Economics* 58(3), 607–626.
- Choi, J. P. and D.-S. Jeon (2020). Two-sided platforms and biases in technology adoption.
- Choi, J. P. and D.-S. Jeon (2021). A leverage theory of tying in two-sided markets with nonnegative price constraints. *American Economic Journal: Microeconomics* 13(1), 283–337.
- Choi, J. P., B. Jullien, and Y. Lefouili (2017). Tying in two-sided markets with multi-homing: Corrigendum and comment. *The Journal of Industrial Economics* 65(4), 872–886.
- Chou, C.-f. and O. Shy (1990). Network effects without network externalities. *International Journal of Industrial Organization* 8, 259–270.
- Chowdhury, S. M. and S. Martin (2017). Exclusivity and exclusion on platform markets. *Journal of Economics* 120(2), 95–118.
- Church, J. and N. Gandal (1992). Network effects, software provision, and standardization. *Journal of Industrial Economics* 40, 85–103.

- Clements, M. and H. Ohashi (2005). Indirect network effects and the product cycle: Video games in the U.S., 1994-2002. *Journal of Industrial Economics* 53, 515–542.
- Correia-da Silva, J., B. Jullien, Y. Lefouili, and J. Pinho (2019). Horizontal mergers between multisided platforms: Insights from Cournot competition. *Journal of Economics & Management Strategy* 28(1), 109–124.
- Cournot, A.-A. (1838). *Recherches sur les principes mathématiques de la théorie des richesses par Augustin Cournot*. chez L. Hachette.
- Crampes, C., C. Haritchabalet, and B. Jullien (2009). Advertising, competition and entry in media industries. *The Journal of Industrial Economics* 57(1), 7–31.
- Crawford, G. S., R. S. Lee, M. D. Whinston, and A. Yurukoglu (2018). The welfare effects of vertical integration in multichannel television markets. *Econometrica* 86, 891–954.
- Crawford, G. S. and A. Yurukoglu (2012). The welfare effects of bundling in multi-channel television markets. *American Economic Review* 102, 643–685.
- Cusumano, M. A., A. Gawer, and D. B. Yoffie (2019). *The Business of Platforms*. HarperCollins.
- Damiano, E. and H. Li (2007). Price discrimination and efficient matching. *Economic Theory* 30(2), 243–263.
- Damiano, E. and H. Li (2008). Competing matchmaking. *Journal of the European Economic Association* 6(4), 789–818.
- De Cornière, A. and G. Taylor (2014). Integration and search engine bias. *The RAND Journal of Economics* 45(3), 576–597.
- De Cornière, A. and G. Taylor (2019). Upstream bundling and leverage of market power. Unpublished manuscript, Toulouse School of Economics.
- Dewenter, R., J. Haucap, and T. Wenzel (2011). Semi-collusion in media markets. *International Review of Law and Economics* 31(2), 92–98.
- Doganoglu, T. (2003). Dynamic price competition with consumption externalities. *Netnomics* 5(1), 43–69.
- Doganoglu, T. and J. Wright (2006). Multihoming and compatibility. *International Journal of Industrial Organization* 24(1), 45–67.
- Doganoglu, T. and J. Wright (2010). Exclusive dealing with network effects. *International Journal of Industrial Organization* 28(2), 145 – 154.
- Dranove, D. and N. Gandal (2003). The DVD vs. DIVX standards war: Network effects and empirical evidence of preannouncement effects. *Journal of Economics and Management Strategy* 12, 363–386.
- Driskill, R. (2007). Monopoly and oligopoly supply of a good with dynamic network externalities. Unpublished manuscript, Vanderbilt University.
- Dube, J.-P. H., G. J. Hitsch, and P. K. Chintagunta (2010). Tipping and concentration in markets with indirect network effects. *Marketing Science* 29, 216–249.

- Durlauf, S. N. and Y. M. Ioannides (2010). Social interactions. *Annual Review of Economics* 2, 1–78.
- D’Annunzio, A. (2017). Vertical integration in the TV market: Exclusive provision and program quality. *International Journal of Industrial Organization* 53, 114–144.
- Einav, L., C. Farronato, and J. Levin (2016). Peer-to-peer markets. *Annual Review of Economics* 8, 615–635.
- Ellison, G. and D. Fudenberg (2003). Knife-edge or plateau: When do market models tip? *Quarterly Journal of Economics* 118, 1249–1278.
- Ellison, G., D. Fudenberg, and M. Möbius (2004). Competing auctions. *Journal of the European Economic Association* 2(1), 30–66.
- Evans, D. S. and R. Schmalensee (2010). Failure to launch: Critical mass in platform businesses. Unpublished manuscript, University College London.
- Evans, D. S. and R. L. Schmalensee (2015). The antitrust analysis of multi-sided platform businesses. In D. D. Sokol and R. D. Blair (Eds.), *Oxford Handbook of International Antitrust Economics*, pp. 404–447. Oxford University Press.
- Fainmesser, I. P. and A. Galeotti (2016). Pricing network effects. *The Review of Economic Studies* 83(1), 165–198.
- Fan, Y. (2013). Ownership consolidation and product characteristics: A study of the U.S. daily newspaper market. *American Economic Review* 103, 1598–1628.
- Farrell, J. and P. Klemperer (2007). Coordination and lock-in: Competition with switching costs and network effects. In M. Armstrong and R. Porter (Eds.), *Handbook of Industrial Organization, Vol. 3*. Elsevier.
- Farrell, J. and G. Saloner (1986). Installed base and compatibility: Innovation, product preannouncements, and predation. *American Economic Review* 76, 940–955.
- Farronato, C., J. Fong, and A. Fradkin (2020). Dog eat dog: Measuring network effects using a digital platform merger. Unpublished manuscript, Boston University.
- Fershtman, D. and A. Pavan (2017). Pandora’s auctions: Dynamic matching with unknown preferences. *American Economic Review* 107(5), 186–90.
- Fershtman, D. and A. Pavan (2021). Matching auctions.
- Foros, Ø., H. J. Kind, and L. Sjørgard (2015). Merger policy and regulation in media industries. In *Handbook of Media Economics*, Volume 1, pp. 225–264. Elsevier.
- Fox, J. (2018). Estimating matching games with transfers. *Quantitative Economics* 9, 1–38.
- Fox, J. T. and P. Bajari (2013). Measuring the efficiency of an FCC spectrum auction. *American Economic Journal: Microeconomics* 5, 100–146.
- Fréchet, G. R., A. Lizzeri, and T. Salz (2019). Frictions in a competitive, regulated market: Evidence from taxis. *American Economic Review* 109, 2954–2992.
- Fudenberg, D. and J. Tirole (2000). Pricing a network good to deter entry. *The Journal of Industrial Economics* 48(4), 373–390.

- Fumagalli, C., M. Motta, and C. Calcagno (2018). *Exclusionary Practices: The Economics of Monopolisation and Abuse of Dominance*. Cambridge University Press.
- Gabszewicz, J. J. and X. Y. Wauthy (2014). Vertical product differentiation and two-sided markets. *Economics Letters* 123(1), 58–61.
- Galichon, A. (2018). *Optimal Transport Methods in Economics*. Princeton University Press.
- Gandal, N. (1994). Hedonic price indexes for spreadsheets and an empirical test for network externalities. *RAND Journal of Economics* 25, 160–170.
- Gandal, N., M. Kende, and R. Rob (2000). The dynamics of technological adoption in hardware/software systems: The case of compact disc players. *RAND Journal of Economics* 31, 43–61.
- Gans, J. S. (2019). The specialness of zero. Technical report, National Bureau of Economic Research.
- Gavazza, A. (2016). An empirical equilibrium model of a decentralized asset market. *Econometrica* 84, 1755–1798.
- Genakos, C. and T. Valletti (2011). Testing the “waterbed” effect in mobile telecommunications. *Journal of the European Economic Association* 9, 1114–1142.
- Goldfarb, A. and C. Tucker (2019). Digital economics. *Journal of Economic Literature* 57, 3–43.
- Gomes, R. and A. Pavan (2016). Many-to-many matching and price discrimination. *Theoretical Economics* 11(3), 1005–1052.
- Gomes, R. and A. Pavan (2019). Price customization and targeting in matching markets. Unpublished manuscript, Northwestern University and Toulouse School of Economics.
- Gomes, R. and A. Pavan (2021). Two-sided markets and matching design. In F. Echenique, N. Immorlica, and V. Vazirani (Eds.), *Online Matching Theory and Market Design*. Cambridge University Press.
- Goolsbee, A. and P. J. Klenow (2002). Evidence on learning and network externalities in the diffusion of home computers. *Journal of Law and Economics* 45, 317–343.
- Gowrisankaran, G., A. Nevo, and R. Town (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review* 105, 172 – 203.
- Gowrisankaran, G., M. Park, and M. Rysman (2014). Measuring network effects in a dynamic environment. Unpublished Manuscript, Boston University.
- Gowrisankaran, G. and J. Stavins (2004). Network externalities and technology adoption: Lessons from electronic payments. *RAND Journal of Economics* 35, 260–276.
- Graham, B. S. (2008). Identifying social interactions through conditional variance restrictions. *Econometrica* 76, 643–660.
- Graham, B. S. (2018). Identifying and estimating neighborhood effects. *Journal of Economic Literature* 56, 450–500.
- Griva, K. and N. Vettas (2011). Price competition in a differentiated products duopoly under network effects. *Information Economics and Policy* 23(1), 85–97.

- Haeringer, G. (2018). *Market Design: Auctions and Matching*. MIT Press.
- Hagiu, A. (2006). Pricing and commitment by two-sided platforms. *The RAND Journal of Economics* 37(3), 720–737.
- Hagiu, A. and H. Halaburda (2014). Information and two-sided platform profits. *International Journal of Industrial Organization* 34, 25–35.
- Hagiu, A. and D. Spulber (2013). First-party content and coordination in two-sided markets. *Management Science* 59(4), 933–949.
- Hagiu, A. and J. Wright (2015a). Marketplace or reseller? *Management Science* 61(1), 184–203.
- Hagiu, A. and J. Wright (2015b). Multi-sided platforms. *International Journal of Industrial Organization* 43, 162–174.
- Halaburda, H., B. Jullien, and Y. Yehezkel (2020). Dynamic competition with network externalities: how history matters. *The RAND Journal of Economics* 51(1), 3–31.
- Halaburda, H., M. J. Piskorski, and P. Yildirim (2018). Competing by restricting choice: The case of matching platforms. *Management Science* 64(8), 3574–3594.
- Halaburda, H. and Y. Yehezkel (2013). Platform competition under asymmetric information. *American Economic Journal: Microeconomics* 5(3), 22–68.
- Halaburda, H. and Y. Yehezkel (2016). The role of coordination bias in platform competition. *Journal of Economics & Management Strategy* 25(2), 274–312.
- Hausman, J. A. (1996). Valuation of new goods under perfect and imperfect competition. In T. F. Bresnahan and R. J. Gordon (Eds.), *The Economics of New Goods*, Chapter 5, pp. 209–248. University of Chicago Press.
- Hendel, I., A. Nevo, and F. Ortalo-Magné (2009). The relative performance of real estate marketing platforms: MLS versus FSBOMadison.com. *American Economic Review* 99, 1878–1898.
- Ho, K. (2009). Insurer-provider networks in the medical care market. *American Economic Review* 99, 393–430.
- Ho, K. and R. S. Lee (2017). Insurer competition in health care markets. *Econometrica* 85, 379–417.
- Hurkens, S. and A. L. López (2014). Mobile termination, network externalities and consumer expectations. *The Economic Journal* 124(579), 1005–1039.
- Innes, R. and R. J. Sexton (1994). Strategic buyers and exclusionary contracts. *American Economic Review* 84(3), 566–584.
- Ishihara, A. and R. Oki (2020). Exclusive content in two-sided markets. Unpublished manuscript, University of Tokyo.
- Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press.
- Jeitschko, T. D. and M. J. Tremblay (2020). Platform competition with endogenous homing. *International Economic Review* 61(3), 1281–1305.
- Jeon, D.-S., B.-C. Kim, and D. Menicucci (2021). Second-degree price discrimination by a two-sided monopoly platform. *American Economic Journal: Microeconomics* (forthcoming).

- Jeziorski, P. (2014). Effects of mergers in two-sided markets: Examination of the U.S. radio industry. *American Economic Journal: Microeconomics* 6, 35–73.
- Jin, G. and M. Rysman (2015). Platform pricing at sports cards conventions. *Journal of Industrial Economics* 63, 704–735.
- Johnson, T. R. (2013). Matching through position auctions. *Journal of Economic Theory* 148(4), 1700–1713.
- Jullien, B. (2011). Competition in multi-sided markets: Divide and conquer. *American Economic Journal: Microeconomics* 3(4), 186–220.
- Jullien, B. (2012). Two-sided B to B platforms. In *The Oxford Handbook of the Digital Economy*. Oxford University Press, New York.
- Jullien, B. and A. Pavan (2019). Information management and pricing in platform markets. *The Review of Economic Studies* 86(4), 1666–1703.
- Jullien, B. and W. Sand-Zantman (2020). The economics of platforms: A theory guide for competition policy. *Information Economics and Policy*.
- Kaiser, U. and J. Wright (2006). Price structure in two-sided markets: Evidence from the magazine industry. *International Journal of Industrial Organization* 24, 1–28.
- Karle, H., M. Peitz, and M. Reisinger (2020). Segmentation versus agglomeration: Competition between platforms with competitive sellers. *Journal of Political Economy* 128(6), 2329–2374.
- Karlinger, L. and M. Motta (2012). Exclusionary pricing when scale matters. *The Journal of Industrial Economics* 60(1), 75–103.
- Katz, M. and C. Shapiro (1992). Product introduction with network externalities. *Journal of Industrial Economics* 40(1), 55–83.
- Katz, M. L. and C. Shapiro (1985). Network externalities, competition and compatibility. *American Economic Review* 75, 424–440.
- Katz, M. L. and C. Shapiro (1986). Technology adoption in the presence of network externalities. *Journal of Political Economy* 94, 822–841.
- Katz, M. L. and C. Shapiro (1994). Systems competition and network effects. *Journal of Economic Perspectives* 8, 93–115.
- Kay, B. S., M. D. Manuszak, and C. M. Vojtech (2018). Competition and complementarities in retail banking: Evidence from debit card interchange regulation. *Journal of Financial Intermediation* 34, 91–108.
- Kim, H. and K. Serfes (2006). A location model with preference for variety. *The Journal of Industrial Economics* 54(4), 569–595.
- Kim, J.-H., P. Newberry, L. Wagman, and R. Wolff (2021). Local network effects in the adoption of a digital platform. Forthcoming, *Journal of Industrial Economics*.
- Kind, H. J., T. Nilssen, and L. Sjørgard (2007). Competition for viewers and advertisers in a TV oligopoly. *Journal of Media Economics* 20(3), 211–233.

- Kojima, F. (2017). Recent developments in matching theory and their practical applications. *Advances in Economics and Econometrics 1*, 138.
- Kourandi, F., J. Krämer, and T. Valletti (2015). Net neutrality, exclusivity contracts, and internet fragmentation. *Information Systems Research 26*(2), 320–338.
- Laussel, D., M. de Montmarin, and N. Van Long (2004). Dynamic duopoly with congestion effects. *International Journal of Industrial Organization 22*(5), 655–677.
- Laussel, D. and J. Resende (2014). Dynamic price competition in aftermarkets with network effects. *Journal of Mathematical Economics 50*, 106–118.
- Lee, R. (2013). Vertical integration and exclusivity in platform and two-sided markets. *American Economic Review 103*, 2960–3000.
- Lefouili, Y. and J. Pinho (2020). Collusion between two-sided platforms. *International Journal of Industrial Organization 72*, 102656.
- Leonello, A. (2010). Horizontal mergers in two-sided markets. Unpublished manuscript, European University Institute.
- Lerner, J. and J. Tirole (2006). A model of forum shopping. *American Economic Review 96*, 1091–1113.
- Li, J. (2019). Compatibility and investment in the U.S. electric vehicle market. Unpublished manuscript, MIT.
- Li, S., L. Tong, J. Xing, and Y. Zhou (2017). The market for electric vehicles: Indirect network effects and policy design. *Journal of the Association of Environmental and Resource Economists 4*, 89–133.
- Liu, C., T.-H. Teh, J. Wright, and J. Zhou (2020). Multihoming and oligoplistic platform competition. Working paper.
- Manski, C. (1993). Identification of endogenous social effects: The reflection problem. *Review of Economic Studies 60*, 531–542.
- Manski, C. (1995). *Identification Problems in the Social Sciences*. Harvard University Press.
- Mansour, Y., S. Muthukrishnan, and N. Nisan (2012). Doubleclick ad exchange auction. Unpublished manuscript, Tel Aviv University.
- Manuszak, M. D. and K. Wozniak (2017). The impact of price controls in two-sided markets: Evidence from us debit card interchange fee regulation. Finance and Economics Discussion Series 2017-074, Board of Governors of the Federal Reserve System.
- Markovich, S. (2008). Snowball: A dynamic oligopoly model with indirect network effects. *Journal of Economic Dynamics and Control 32*(3), 909–938.
- Markovich, S. and J. Moenius (2009). Winning while losing: Competition dynamics in the presence of indirect network effects. *International Journal of Industrial Organization 27*(3), 346–357.
- Markovich, S. and Y. Yehezkel (2021). Group hug: Platform competition with user-groups. *American Economic Journal: Microeconomics (forthcoming)*.

- Marra, M. (2020). Estimating an auction platform game with two-sided entry. Unpublished manuscript, Sciences Po.
- Martens, B. (2016). An economic policy perspective on online platforms. Unpublished manuscript, Institute for Prospective Technological Studies.
- Martimort, D. and L. Stole (2009). Market participation in delegated and intrinsic common-agency games. *The RAND Journal of Economics* 40(1), 78–102.
- Miao, C.-H. (2009). Limiting compatibility in two-sided markets. *Review of Network Economics* 8(4), 1–19.
- Mitchell, M. F. and A. Skrzypacz (2006). Network externalities and long-run market share. *Economic Theory* 29.
- Moffitt, R. A. (2001). Policy interventions, low-level equilibria, and social interactions. In S. N. Durlauf and H. P. Young (Eds.), *Social Dynamics*, pp. 45–82.
- Moghaddam, M. and S. Y. Nof (2017). *Best Matching Theory & Applications*. Springer.
- Mussa, M. and S. Rosen (1978). Monopoly and product quality. *Journal of Economic theory* 18(2), 301–317.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research* 6(1), 58–73.
- Myerson, R. B. and M. A. Satterthwaite (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29(2), 265–281.
- Nair, H., P. Chintagunta, and J.-P. Dube (2004). Empirical analysis of indirect network effects in the market for personal digital assistants. *Quarterly Marketing and Economics* 2, 23–58.
- Nocke, V., M. Peitz, and K. Stahl (2007). Platform ownership. *Journal of the European Economic Association* 5(6), 1130–1160.
- Ohashi, H. (2003). The role of network effects in the U.S. VCR market, 1978-86. *Journal of Economics and Management Strategy* 12, 447–494.
- Park, S. (2004). Quantitative analysis of network externalities in competing technologies: The VCR case. *Review of Economics and Statistics* 86, 937–945.
- Parker, G. and M. Van Alstyne (2005). Two-sided network effects: A theory of information product design. *Management Science* 51, 1494–1504.
- Pathak, P. A. (2017). What really matters in designing school choice mechanisms. *Advances in Economics and Econometrics* 1, 176–214.
- Peitz, M., S. Rady, and P. Trepper (2017). Experimentation in two-sided markets. *Journal of the European Economic Association* 15(1), 128–172.
- Peitz, M. and M. Reisinger (2015). The economics of internet media. In S. P. Anderson, J. Waldfogel, and D. Strömberg (Eds.), *Handbook of Media Economics*, pp. 445–530. Elsevier.
- Peitz, M. and T. M. Valletti (2008). Content and advertising in the media: Pay-TV versus free-to-air. *International Journal of Industrial Organization* 26(4), 949–965.
- Perry, M. K. (1989). Vertical integration: Determinants and effects. In *Handbook of Industrial Organization*, Volume 1, pp. 183–255. Elsevier.

- Pinker, E. J., A. Seidmann, and Y. Vakrat (2003). Managing online auctions: Current business and research issues. *Management Science* 49(11), 1457–1484.
- Rasmusen, E. B., J. M. Ramseyer, and J. S. Wiley Jr (1991). Naked exclusion. *The American Economic Review*, 1137–1145.
- Reisinger, M. (2014). Two-part tariff competition between two-sided platforms. *European Economic Review* 68, 168–180.
- Reshef, O. (2020). Smaller slices of a growing pie: The effects of entry in platform markets. Unpublished manuscript, University of Washington, St. Louis.
- Rey, P. and J. Tirole (2007). A primer on foreclosure. In *Handbook of Industrial Organization*, Volume 3, pp. 2145–2220. Elsevier.
- Robles-Garcia, C. (2020). Competition and incentives in mortgage markets: The role of brokers. Unpublished manuscript, Stanford University.
- Rochet, J.-C. and J. Tirole (2002). Cooperation among competitors: Some economics of payment card associations. *RAND Journal of Economics* 33, 549–570.
- Rochet, J.-C. and J. Tirole (2003). Platform competition in two-sided markets. *Journal of the European Economic Association* 1, 990–1029.
- Rochet, J.-C. and J. Tirole (2006). Two-sided markets: A progress report. *RAND Journal of Economics* 37, 645–667.
- Rohlfis, J. (1974). A theory of interdependent demand for a communications service. *The Bell Journal of Economics and Management Science* 5(1), pp. 16–37.
- Rosaia, N. (2021). Competing platforms and transport equilibrium: Evidence from New York City. Unpublished Manuscript, Harvard University.
- Rosse, J. N. (1970). Estimating cost function parameters without using cost data: Illustrated methodology. *Econometrica* 38, 256–275.
- Roth, A. E. (2008). Deferred acceptance algorithms: History, theory, practice, and open questions. *International Journal of Game Theory* 36(3), 537–569.
- Roth, A. E. (2018). Marketplaces, markets, and market design. *American Economic Review* 108, 1609–1658.
- Roth, A. E. and M. Sotomayor (1992). Two-sided matching. In *Handbook of Game Theory with Economic Applications*, Volume 1, pp. 485–541. Elsevier.
- Ryan, S. P. and C. Tucker (2012). Heterogeneity and the dynamics of technology adoption. *Quantitative Marketing and Economics* 10, 63–109.
- Rysman, M. (2004). Competition between networks: A study of the market for Yellow Pages. *Review of Economic Studies* 71, 483–512.
- Rysman, M. (2007). Empirical analysis of payment card usage. *Journal of Industrial Economics* 60, 1–36.
- Rysman, M. (2009). The economics of two-sided markets. *Journal of Economic Perspectives* 23, 125–144.

- Rysman, M. (2019). The reflection problem in network effect estimation. *Journal of Economics and Management Strategy* 28, 153–158.
- Rysman, M. and J. Wright (2014). The economics of payment cards. *Review of Network Economics* 13, 303–353.
- Sacerdote, B. (2001). Peer effects with random assignment: Results for Dartmouth roommates. *Quarterly Journal of Economics* 116, 681–704.
- Sacerdote, B. (2014). Experimental and quasi-experimental analysis of peer effects: Two steps forward? *Annual Review of Economics* 6, 53–72.
- Salz, T. (2020). Intermediation and competition in search markets: An empirical case study. In Press, *Journal of Political Economy*.
- Schiff, A. (2003). Open and closed systems of two-sided networks. *Information Economics and Policy* 15(4), 425–442.
- Seamans, R. and F. Zhu (2014). Responses to entry in multi-sided markets: The impact of Craigslist on local newspapers. *Management Science* 60, 476–493.
- Segal, I. R. and M. D. Whinston (2000). Naked exclusion: Comment. *American Economic Review* 90(1), 296–309.
- Seim, K. (2006). An empirical model of firm entry with endogenous product-type choices. *RAND Journal of Economics* 37, 619–640.
- Selten, R. (1973). A simple model of imperfect competition, where 4 are few and 6 are many. *International Journal of Game Theory* 2(1), 141–201.
- Shaked, A. and J. Sutton (1983). Natural oligopolies. *Econometrica*, 1469–1483.
- Sokullu, S. (2016). A semi-parametric analysis of two-sided markets: An application to the local daily newspapers in the USA. *Journal of Applied Econometrics* 31, 843–864.
- Song, M. (2021). Estimating platform market power in two-sided markets with an application to magazine advertising. *American Economic Journal: Microeconomics* 13, 35–67.
- Spence, A. (1981). The learning curve and competition. *The Bell Journal of Economics* 12, 49–70.
- Stennek, J. (2014). Exclusive quality—why exclusive distribution may benefit the TV-viewers. *Information Economics and Policy* 26, 42–57.
- Tadelis, S. (2016). Reputation and feedback systems in online platform markets. *Annual Review of Economics* 8, 321–340.
- Tan, G. and J. Zhou (2019). Price competition in multi-sided markets. Available at SSRN 3029134.
- Tan, G. and J. Zhou (2020). The effects of competition and entry in multi-sided markets. *The Review of Economic Studies*.
- Tan, H. and J. Wright (2018). A price theory of multi-sided platforms: Comment. *American Economic Review* 108(9), 2758–60.
- Teh, T.-H. (2021). Platform governance. *American Economic Journal: Microeconomics* (forthcoming).

- Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton University Press.
- Train, K. (2009). *Discrete Choice Methods with Simulation* (2nd ed.). Cambridge University Press.
- Tucker, C. (2008). Identifying formal and informal influence in technology adoption with network externalities. *Management Science* 12, 2024–2039.
- Valenzuela-Stookey, Q. (2020). Platform-mediated competition. Unpublished manuscript, Northwestern University.
- Vasconcelos, H. (2015). Is exclusionary pricing anticompetitive in two-sided markets? *International Journal of Industrial Organization* 40, 1–10.
- Veiga, A., E. G. Weyl, and A. White (2017). Multidimensional platform design. *American Economic Review* 107(5), 191–95.
- Weiergraeber, S. (2019). Network effects and switching costs in the US wireless industry. Unpublished manuscript, Indiana University.
- Wen, W. and F. Zhu (2019). Threat of platform-owner entry and complementor responses: Evidence from the mobile app market. *Strategic Management Journal* 40, 1336–1367.
- Weyl, E. G. (2010). A price theory of multi-sided platforms. *American Economic Review* 100, 1642–1672.
- White, A. and E. G. Weyl (2016). Insulated platform competition. Unpublished manuscript, Tsinghua University.
- Zhou, Y. (2017). Bayesian estimation of a dynamic model of two-sided markets: Application to the U.S. video game industry. *Management Science* 63, 3874–3894.
- Zhu, F. and M. Iansiti (2012). Entry into platform-based markets. *Strategic Management Journal* 33(1), 88–106.
- Zhu, F. and Q. Liu (2016). Competing with complementors: An empirical look at Amazon.com. Harvard Business School Working Paper 15-044.