Getting Started

Motivation

This book analyzes firms that manage the interaction between their users. More specifically, the book considers markets in which users enjoy benefits that depend on the decisions of other users (meaning that users are subject to network effects) and a firm, operating the "platform," takes decisions that partly determine how large those benefits are and who will obtain which benefit.

Such platforms pervade our everyday lives and contribute to an increasing share of economic activity; they also raise important societal issues through their innovative activities and practices. Of particular interest are "two-sided platforms" in which a firm caters to heterogeneous users with different needs or interests (think, for instance, of an intermediary that runs an e-commerce platform on which buyers and sellers interact). Intermediaries operating a two-sided platform determine which products to show to which consumers and which product information to release; they also guide information gathering of consumers and often control or regulate information release by sellers; they may further determine the terms and conditions under which a transaction is to be carried out by fixing the price sellers can charge to buyers or by fixing nonprice elements of the contract between buyer and seller (e.g., cancellation terms, dispute settlement on the platform).

The operations of platforms put into question the way many academic economists have thought about the functioning of markets. First, the way markets operate (and which allocation will result) depends on the incentives and means of a platform operator that monitors and controls the interaction between its users. A for-profit platform provides services and taxes trade and interaction (by charging a price to at least some of its users or by bundling its service with other services that generate revenues elsewhere); in many ways, the platform plays the role that the state or some nonprofit actors could play.[1] Second, processes like the exchange of information and even, possibly, the formation of prices are no longer decentralized but controlled by the platform. An example of a platform with centralized price-setting is Uber, which sets the prices at which drivers and travelers can interact. This stands in sharp

[1] For instance, Booking and Airbnb offer services to hotels and landlords on one side and guests on the other side for short-term stays in a hotel or apartment, and charge for their services. Similarly, some nonprofit local tourist boards offer services to local landlords, which include a portal for hotel and apartment bookings, and charge a fee for its services.

contrast with standard microeconomic teaching, which imagines markets operated by some "invisible hand" that brings supply and demand together. Instead, what we describe here are situations in which trade is carried out under the *visible* hand of the intermediary managing the platform.[2]

Intermediaries operating platforms have created possibilities for interactions between users where nonorganized marketplaces often have failed to properly solve coordination problems and asymmetric information problems. In particular, intermediaries can provide a safe trading environment, attract key participants with special offers, and make infrastructure investments that facilitate trade or interaction. They may thus fully or partially replace other forms of markets or institutions; they may also create markets that previously did not exist. Many of those intermediaries are driven by profit motives (at least from some stage of their existence), but there remain other intermediaries that keep on operating on a nonprofit basis (e.g., Wikipedia). In the past, many such central places for trade or interaction were (and still are) subject to specific rules set (or directly controlled) by the government.[3]

A Little Bit of History

Marketplaces as meeting points of traders and consumers have existed for thousands of years. For example, traders and consumers in ancient Greece could find a meeting place in the *agora*, while in the Roman Empire, the *forum* and the *macellum* served as, respectively, open-air and indoor marketplaces.[4] These marketplaces operated under local rules and were not privately run. Meeting places (such as the Roman forum) also served for the exchange of news, which again features network effects because, with more participation, more diverse and more credible news could be obtained.

Starting with the late Middle Ages, exchanges allowing trade in obligations and government securities were first set up in France, in some city states of Northern Italy and later in what is now Belgium. In the seventeenth century, with the emergence of publicly traded companies, the trading of stocks started to flourish on the Amsterdam bourse; the Dutch East India Company and the Dutch West India Company were the two major companies whose stocks were traded. Also in the seventeenth century, auction houses for arts and valuables appeared in London. These platforms benefited from network effects by providing liquidity.[5] A seller was confident

[2] We are not claiming that it is a novel approach of the literature on two-sided platforms to have a party implementing a pricing system that balances demand and supply. The literature on market intermediation has analyzed a variety of settings in which intermediaries run marketplaces that are preferred by at least some users and thus replaces decentralized trade because they fix prices on both sides of the market; for an overview, see Spulber (1999). Furthermore, double auctions are used to find market-clearing prices (see Wilson, 1985; Satterthwaite and Williams, 1989; McAfee, 1992); such allocation mechanisms have been applied in a number of markets, e.g., wholesale electricity markets.

[3] Languages and currencies also allow for interactions to take place, feature coordination problems, and were often, at least partly, government-controlled.

[4] In the macellum, vendors sold food products. They were constrained in their pricing and their produce was subject to some quality control by a public authority (see de Ruyt, 1983, pp. 356–358).

[5] For a discussion of network effects on stock exchanges, see Pagano and Padilla (2005). In particular, risk-averse participants benefit from a more liquid market; see Pagano (1989).

about finding interested buyers, and a buyer was confident about finding appropriate supplies.

Jumping straight into the twenty-first century, we observe the important role of digital platforms in creating economic value, the breathtaking rise of stock market valuations of some of them, and the growing policy concern about the economic and political power of GAFAM (Google, Amazon, Facebook, Apple, and Microsoft) in the US-dominated part of the world. The policy reports that were released recently in several places around the world are evidence of this growing concern.[6] We also see the power of other non-US-based platforms in the China-dominated part of the world (the acronym there is BATX, which stands for Baidu, Alibaba, Tencent, Xiaomi), with important overlaps in some countries and homegrown behemoths in others.

Preliminaries

Network effects and scale economies lie at the core of the success of platforms. When we say that a platform *enjoys* network effects or scale economies, this is often misleading, as it often takes ingenuity and effort to make network effects (and scale economies) happen. Sometimes luck replaces ingenuity (e.g., a good platform design may have been picked by accident), and often effort involves a lot of experimentation (e.g., by testing out modifications on a small group of users).

Network effects can be considered as scale economies on the demand side. In a way, they are the mirror image of scale economies of production: In the case of network effects, an expansion of the firm's customer base increases the benefit that goes to each consumer while, in the case of scale economies, this expansion reduces the average cost of production and thus generates a benefit that goes to the firm (for a given price). Given the proximity – but also the differences – between the two concepts, it is useful to take a closer look at them.

For simplicity, consider a population of identical users with unit demand. A user's net utility is computed as the difference between their gross utility, U, and the price they must pay to access the platform, A. With network effects, we have $U(N) - A$, where U is increasing in the number of users N under positive network effects. With constant marginal costs of production, c, and price equal to marginal cost ($A = c$), the net utility is $U(N) - c$. With scale economies in production, the marginal cost depends on N and is decreasing in N. With price equal to marginal cost, the price is $c(N)$ and the net utility of a user is $U - c(N)$.[7]

[6] In particular, we think of the reports from the EC (Crémer et al., 2019), the UK (Furman et al., 2019), and the USA (Scott Morton et al., 2019). For a comparison of the policy recommendations made in these three reports, see Ennis and Fletcher (2020). Australia, Benelux, BRICS, France, Germany, Italy, Japan, the Netherlands, Portugal, and UNCTAD also released their own reports in 2019.

[7] In the real world, network effects and scale economies may be present at the same time. For instance, when Amazon (as a retailer) ships more products, it gains from scale economies of the logistics network (therefore $c(N)$ is decreasing) and, at the same time, consumers may benefit from it through quicker delivery (as demand becomes more predictable, Amazon's stocking decisions better reflect consumer tastes). Thus, if Amazon were to price at marginal cost, users' net benefit would be $U(N) - c(N)$.

In both cases – network effects and scale economies – the net utility is increasing in the number of users N. What is the socially optimal size of the user base? The optimality condition is that the marginal social benefit is equal to the marginal cost. With network effects, the total gross benefit is $N \times U(N)$ and, thus, the marginal benefit is $U(N) + N \times U'(N)$. Hence, the optimal N satisfies $U(N) + N \times U'(N) = c$. Here, marginal-cost pricing does not implement the socially optimal outcome but is feasible for a private firm in the sense that a platform forced to set price equal to marginal cost does not make a loss and, thus, does not exit (since we assume that there are constant marginal cost and no fixed cost). Nonoptimality follows from the fact that participation will increase as long as $U(N) - A \geq 0$. To obtain socially optimal participation levels, we would need $A = c - N \times U'(N) < c$. That is, the firm would need to implement below marginal-cost pricing. In the presence of scale economies, the total cost is $C(N)$ with $C''(N) = c'(N) < 0$. The socially optimal N satisfies $U = c'(N)$. Marginal-cost pricing implements the socially optimal outcome but is not feasible in the sense that a private firm would make losses under marginal-cost pricing, $N \times c(N) < C(N)$.

The relationship between network effects and scale economies is less straightforward if a platform caters to two different user groups that are possibly connected through cross-group network effects or that feature decreasing unit costs.[8] Suppose that there are two groups, X and Y, that group X exerts a positive cross-group effect on group Y, and that marginal costs are constant and equal to c; that is, a user of group Y obtains a gross benefit of $U_Y(N_X) - c$ with $U_Y(N_X)$ increasing in N_X (and with N_K, $K \in \{X, Y\}$, denoting the number of users of group K that are present on the platform). Alternatively, there may be scale economies in the provision to group Y from more participation of group X; that is, $c(N_X)$ is a decreasing function of N_X. Let us compare the net surplus for group Y users under marginal-cost pricing in the two cases. With network effects, we have $U_Y(N_X) - c$ and with scale economies, we have $U - c(N_X)$. In both cases, the net surplus is increasing in N_X. While this shows that the similarity between network effects and scale economies extends to two-sided platforms, the property that the unit costs of serving a user from group Y decreases with the participation of group-X users appears to be less relevant in practice.[9] Moreover, a platform may be restricted in its ability to set prices that take into account the opportunity cost of losing a user of one group. For instance, the platform may face pricing restrictions such that users of one group cannot receive negative prices.[10]

[8] Cross-group network effects exist when users in one group care about the participation of users of another group; we define them formally in Chapter 1.

[9] One case in point could be the following. If users of group Y are distributed in space and a platform compensates them for travel time, then an increase in N_X reduces the expected cost to reach a user of group X since average travel time goes down. Takeout meals from restaurants whose ordering and delivery is facilitated by platforms are a possible application: Suppose that the number of restaurants on a platform is given; group Y consists of available riders to deliver meals, while group X contains the consumers ordering meals.

[10] We address the pricing strategy of a platform subject to network effects and such pricing constraints in Chapter 5.

In the formal analysis of the book, we focus on network effects and mostly assume that there are no scale economies of production. We explain the mechanisms by which different types of network effects arise and how a platform's strategy affects the strength of those network effects. We acknowledge that it may depend on the exact types of contract between one group of users and the platform as to whether scale economies are better classified as being on the supply side (because they reduce the platform's unit costs) or on the demand side (because they increase the user's willingness to pay, leading then to network effects). To see this, take the example of a food-delivery platform. If the platform pays riders per hour they are available, then a better prediction of demand (because of increased volume) reduces in expectation the number of riders (relative to the number of expected deliveries) that have to be hired in any given hour to make sure that a certain delivery time is typically met; as a result, the platform experiences decreasing unit costs – an instance of supply-side economies of scale. By contrast, if riders are paid per drop, improved demand predictions make the platform more attractive to riders because they obtain more reliable information as to whether they may actually be needed; this is then an instance of demand-side economies of scale.[11]

Approach

In this book, we mostly follow the theoretical literature on multi-sided platforms that emerged in the early 2000s.[12] This literature first focused on card payment systems such as Visa and Mastercard, which Rochet and Tirole (2002) analyzed as two-sided platforms with merchants on one side and consumers on the other side.[13] This

[11] Both types of contract can be observed in the industry. For example, in the UK Just Eat (part of Just Eat Takeaway) pays per hour including waiting time (according to its website for riders, www.takeaway.com/drivers/uk/, last accessed November 24, 2020). Another food-delivery platform is Deliveroo. According to the UK tabloid *Daily Mirror*, "there are two pay structures with Deliveroo. With the original model, riders are paid an hourly wage that is then topped up with a small additional payment for each delivery. There is now a second model where riders are exclusively paid based on deliveries, but with a larger fee." (John Fitzsimons, "Deliveroo: How to become a rider and what you can earn," *Daily Mirror*, July 12, 2018, updated May 28, 2020).

[12] An early literature studied markets with direct and indirect network effects that were started with communication networks in the 1970s (Artle and Averous, 1973; Rohlfs, 1974) and continued in the 1980s with a particular focus on compatibility (Farrell and Saloner, 1985, 1986; Katz and Shapiro, 1985, 1986) – we will have a closer look at this issue in Chapter 3. There is also an abundant literature on platforms in management science; for recent and comprehensive books grounded in this literature, see, e.g., Parker, Van Alstyne, and Choudary (2016) and Cusumano, Gawer, and Yoffie (2019).

[13] The peculiar feature of Visa and Mastercard (in contrast, e.g., to American Express) is that they do not contract with consumers and merchants directly but let banks do so: On the one hand, banks are issuers of cards to consumers and, on the other, they provide transaction processing services to merchants and play the role of "acquirer." Typically the issuing bank is different from the acquiring bank. When the merchant pays a fee to the acquiring bank, whereas consumers are not charged by the issuing bank (or even receive a subsidy), the acquiring bank has to make a payment to the issuing bank – this is the "interchange fee." Gans and King (2003) show that the interchange fee is neutral and, thus, does not affect the equilibrium outcome when there is payment separation between cash and credit transactions. Payment separation is violated if the no-surcharge rule applies to sellers, which prevents them from charging different prices depending on the type of transaction.

concept of two- or multi-sided platforms was then applied more generally – see the seminal research articles by Caillaud and Jullien (2003), Rochet and Tirole (2003, 2006), Anderson and Coate (2005), Parker and van Alstyne (2005), and Armstrong (2006). Since then, an impressive body of research has developed and examined an ever-increasing range of issues relating to markets with multi-sided platforms.

Economic theory helps us to build rigorous arguments, pointing to the implications of particular assumptions about the market environment faced by a platform and the strategic choices available to it. Throughout the book, we try to keep theoretical models as simple as possible; when suitable, models are exposed in a general way; otherwise, they are formulated as numerical examples or just explained in words. Important takeaways from our investigation are formulated as "lessons." At many points, we provide real-world facts and insights formulated as "cases," which complement our theory-based exposition. We also point to empirical work that provides further guidance.

Our theoretical analyses are hopefully accessible to readers without advanced training in calculus and, therefore, serve the purpose of making arguments precise without creating barriers to understanding because of the mathematics that are needed. Somewhat more complicated arguments are summarized in words, which allows for a quicker reading by skipping the associated formal exposition.

We provide relevant references in this book, but do not claim to provide an exhaustive literature guide. Our aim is to provide some basic insights based on platforms being seen as managers or organizers of the interaction between users. The key feature on which we base our exposition is the presence of network effects. We rely on advances in the literature on (two-sided) platforms that have been made over the last twenty years. Our focus is on economic insights that can be shown in simple, stylized settings.

Outline

This book is organized in six chapters. In **Chapter 1**, we distinguish different types of network effects and define platforms as facilitators of interaction and trade and, more specifically, as managers of network effects. For-profit platforms face two issues: (i) how to create value for participants and (ii) how to manage interaction or trade on the platform. Many digital platforms initially focus on the first issue and do not necessarily have a concrete monetization model in mind. We provide a classification of different types of platforms, which we illustrate with a number of real-world examples.

In **Chapter 2**, we take a closer look at multiple sources of network effects and how platform design choices and user behavior relate to those network effects. In particular, we elaborate on rating and recommender systems. The former aggregate user

Rochet and Tirole (2002) and Wright (2004) compare private to social incentives of setting the interchange fee and its impact on the prices on the two sides. Wright (2004) provides conditions under which private and social incentives coincide and under which they do not.

experiences and, thus, address asymmetric information problems. We show under which conditions rating systems lead to positive network effects. Also, recommender systems rely on information about user behavior and possibly enable users to make better informed decisions or reduce the cost of decision-making (in particular, users have to search less to find good matches). Again, we show under which conditions recommender systems lead to network effects. In addition, we discuss their effects on the distribution of sales between "mass-market" and "niche" products. Both rating and recommender systems rely on the processing of choice data. More generally, platforms may collect or have access to data about user choices and other variables that are relevant to predict these choices. They may use this information not only to inform and steer some group of users but also for monetization purposes on other sides of the market. In particular, this information may affect, in a buyer-seller relationship, which gains from trade are realized and how these gains are shared. This raises the question about whether and how the collection and processing of more data relate to network effects. Here, we establish conditions under which more and better data lead to network effects and discuss the types of network effects that may arise.

In **Chapter 3**, we consider platforms catering to one group of users who are connected through network effects. We first formalize the impact of network effects on user demand. Users have to form expectations about the participation decision of fellow users who have not yet chosen whether to adopt a network good. As we will see, there are some economic environments in which, for some set of prices, different allocations can be rationalized through self-fulfilling expectations. In other environments, this issue does not arise and only a single allocation can be supported with self-fulfilling expectations. We then consider platform pricing for this network good and uncover to what extent a platform internalizes network effects in its pricing decisions. Furthermore, we investigate how network effects shape a platform's compatibility choice when there are two network goods. In particular, we formalize the compatibility decision of a platform introducing a new network good that can either use the same standard as a preexisting network good or be incompatible with it.

In **Chapter 4**, we ask whether and how a firm wants to establish a platform and how it can grow. We explore a firm's economic trade-offs between choosing a (two-sided) platform model and alternative modes of organization. Within the two-sided platform model, we expose the difficulties that a firm will inevitably encounter when trying to bring two groups of agents together. Potential users need to be convinced that they will find other users on the platform with whom they can interact. The key question is thus how to convince them. In this context, we formalize the "chicken-and-egg-problem" and discuss firm strategies that may solve it. We also discuss the strategies that platforms can implement to increase the level of trust among users, thereby securing their participation and, possibly, intensifying the network effects. Finally, we discuss how a platform can use its strong position for some intermediation service to succeed when offering other intermediation services.

In **Chapter 5**, we take a closer look at pricing decisions by a two-sided platform. The presence of various user groups opens the possibility for differential pricing. As we will see, differential pricing is desirable to tackle the interdependence between

the users' decisions. In the presence of heterogeneous users, the profit-maximizing price structure will be jointly determined by the price elasticities and the network effects. We describe the different types of prices that a platform might choose. We will see that the profit-maximizing pricing structures often have the feature that different groups of users face different price-cost margins; we also address the question of whether a profit-maximizing platform charges users only for access or also for the transactions they conduct on the platform. A platform may be restricted in the available price instruments; also participation may be sequential or some users may not be able to observe all prices. We analyze how a platform responds in its choice of strategy in such circumstances.

In **Chapter 6**, we consider economic environments that are richer in the ways the two user groups interact and in the ways a platform can manage this interaction. First, we extend the analysis to two-sided e-commerce platforms on which sellers compete with each other and discuss in which way a platform manages competition on its platform. In particular, if a platform can only charge sellers, it makes profits from participating in the sellers' gross profits and may therefore be inclined to safeguard high industry profits. However, since high industry profits stem from high prices, this discourages user participation. We address how this trade-off affects platform pricing and product variety. Second, we consider a platform that allows buyers to obtain information about more products as is the case with price comparison engines. We show how the platform's price strategies affect the market outcome and, in particular, the degree of price dispersion that arises naturally with differential information among users (some knowing only their local sellers and others obtaining information by accessing the platform). Third, we revisit the issue of product variety, which a platform can also manage through its design of rating, reviews and recommender systems; we assess the extent to which the incentives of a profit-maximizing platform when designing these systems are aligned with those of the users. We also examine the extent to which an intermediary wants to increase price transparency on the platform. Fourth, we consider platform design regarding the information and price instruments that are made available to sellers on the platform. For instance, the platform may provide sellers with buyers' personal data and, thus, facilitate differential pricing.

We close the book with an Epilogue in which we point towards a number of issues that we leave for our next book.

How to Use the Book

We hope to have written a book that turns out to be useful for students of economics and business, researchers, practitioners, and policy makers. We aimed at making technical material accessible so that it can be digested by advanced undergraduate and master students in economics. For those readers who are less inclined to follow mathematical expressions, we also endeavor to guide them smoothly through the book.

The book aims to introduce key concepts and, even if it does not claim to be all-encompassing, it will hopefully serve as a reference for economists in academia

and be of relevance to competition economists working at agencies and consultancies. Introductory notions of industrial organization, microeconomics, and game theory will certainly facilitate the reading, but they are not necessary. Our intention is indeed to present all the main concepts in a simple and clear way, to illustrate them with short case studies, to develop simple models that establish formally several insights, and to summarize the main lessons that can be drawn from more complex developments. As a result, most of the contents of the book should also be accessible to students, practitioners, and professionals with an interest in the topic, who have been exposed to some economics thinking (even if economics is not their main background); we think, in particular, of lawyers, managers, and IT professionals.

The book can be used, in economics and business programs, for an elective course on network effects and platforms. Until our second book on competition and regulation is published, to complement this book, teachers may want to add some basic models of competing platforms; to this end, they may rely on the exposition in surveys such as Belleflamme and Peitz (2018a) or our textbook treatment in Belleflamme and Peitz (2015, chapter 22); they may also use the theory guide to competition policy contributions with respect to platforms that is provided by Jullien and Sand-Zantman (2021).

More broadly, our book can be used as a source for a course on the digital economy or on peer-to-peer networks; it can further serve as supplementary material for a course on industrial organization that takes a particular interest in platforms. Finally, we also hope that our book will prove useful to competition practitioners when considering some of the conceptual issues regarding platforms that arise when changing the law, when formulating guidelines, and in competition cases.

1 Platforms: Definitions and Typology

According to historians, European trade took off at the end of the twelfth century in what is now the North of France, in the county of Champagne.[1] It is in this period that this county started to host regular trade fairs, which lasted for six weeks and rotated among six cities. Merchants came from all over Europe because they were confident that they would meet each other at these fairs. This confidence was instilled by the count of Champagne through his authoritative and clever running of the fairs. Everything was done to provide merchants with a safe and efficient business environment. The count of Champagne actively selected the participants, especially by keeping away dubious businessmen. Once admitted, all participants were on a level playing field, as the count carefully avoided granting any privilege to anyone. The fair locations were fortified, and impartial institutions were put in place to enforce contracts and resolve disputes. The count also guaranteed loans and the replacement of cash by notary bills to settle transactions. In exchange for all these services, the count took a small share of each transaction and quickly amassed a fortune.

What the count of Champagne started around 1180 is known today as a "platform." A platform can be roughly seen as *an entity that enables interactions between users so as to generate value from these interactions* (we provide a more precise definition in Section 1.3). Implicit behind this description is the observation that although users benefit from interacting, they have a hard time organizing the interaction by their own means. This is because the interaction generates external effects: One user's decisions as to whether and how much to interact affect the well-being of other users. In general, the more users there are, the more valuable the interaction becomes for each user. External effects of this specific sort are called (positive) "network effects" (for reasons that will be clarified in this chapter). The problem is that network effects make the interaction hard to organize, as users fail to take them into account when making their decisions: As long as users are not compensated for the benefits they bring to others, they are insufficiently keen to start the interaction on their own.

As a result, a valuable interaction may fail to occur, unless some third-party intermediary finds ways to "internalize" the network effects. In troubled medieval Europe,

[1] See Fisman and Sullivan (2016).