ECON-C5100 - Digital Markets

Lecture notes 2022 – Draft

Iivo Vehviläinen

February 9, 2022

1 Introduction

The ingresses for the sections will provide some clues on what I consider important.

Economics has been at the heart of much of the digital market revolution, from Google to Amazon and beyond. While success of these companies builds on new technologies, economic theory has found some surprising uses from auctions to market design and the understanding of how modern online platforms should price their products. The course "Digital Markets" aims to increase the understanding of the economics of firm decisions and strategies, and competition in online and digital markets and how the special features of such markets affect decisions at the firms and by the regulators.

I am more than conscious about the fact that the course has no text book to rely on, but so far there is no good alternative. As a partial remedy, this draft is intended to be helpful for the course 2022 participants together with the lecture slides. This may or may not be helpful in addition to your own notes. The idea with these notes is to include most of the relevant economics from the slides in an easy to read-through format. There are some additions to clarify a few points where the slides are not self-sustained, and much repetition from the slides. However, not everything is replicated, in particular tables, figures, and maths. I recommend to keep the lecture slides close by when reading through these notes. Any suggestions on how to improve these notes greatly appreciated.

2 Prerequisites and markets

The course material builds on at least principles level understanding of economics theory, which is not repeated here nor will this material be questioned as such in the exam. The reading list below should provide a sufficient material to follow the other parts of the course. If these seem too hard, you may want to familiarize yourself with the other CORE units as well.

Reading materials

- Preferences (CORE 3.2-3.5)
- Basics of game theory (CORE 4.1-4.3)
- Pareto efficiency (CORE 5.2)
- Institutions, supply and demand (CORE 8.1, 8.2)
- Competitive equilibrium (CORE 8.5)
- Perfect competition (CORE 8.8)

3 Auctions

Auction theory, for the purposes of this course, is a mandatory step in order to understand how much of the money online is made, i.e. through advertisement auctions. Hopefully understanding how auctions work is useful otherwise as well. Some key points from the standard introduction to auction theory are briefly repeated here for completeness.

Why auctions online

In the standard analysis, it is common to start with a situation where a seller has a single item for sale and there a number of potential buyers. What is the price that the seller should set? It will depend on how much buyers are willing to pay; but they are not going to tell you their true value. Auction is a mechanism for price discovery. At the same time, auctions also force buyers to compete with each other in order to win. Indeed, if the cost of arranging an auction can be kept low, then auctioning an item can shown to be superior compared to other ways of selling an item, i.e. posting a single list price or bargaining over the price.

Arranging an auction is a trade-off between the benefit of price discovery and competition vs. the cost of arranging an auction. Both the advantage and disadvantage change online: many more participants are possible online vs. physically and the costs of arranging auctions are lower. Though the early enthusiasm for selling stuff via auctions online (e.g. eBay) has reduced, auctions are the backbone of much of the revenue made in online advertisements.

The analysis of auctions builds on the understanding of how bidders behave within a given auction mechanism. Game theory is extensively used to analyze the strategic interaction that takes place through bidding strategies. The behavior of bidders will be different in different auction mechanisms, so we are somewhat forced to go through the mechanisms to gain some insight on what drives the bidding behavior and why some particular auction form is preferred over the others. The common auction forms are 1) *ascending price auctions*, 2) *descending price auctions*, 3) *first price auctions*, and 4) *second price auctions* or *Vickrey auctions*.

Additionally, the context of the auction matters. Things are different if the auctioned item has the same value for all (*common value* or *known value*) or if everyone values the item differently (*private value* or *unknown value*). The latter is more prevalent in online world applications and is the focus of most of the discussion during the course as well. In particular, online auctions that sell slots for advertisements are an example of private value auctions: each advertiser will have a private value for showing an advertisement that is different from the values of the others.

Example of truthful bidding behavior

Easy access to online marketplaces has enabled auction mechanisms where the optimal bidding strategies are kept simple. In particular ascending price auctions have been popular in online marketplaces like eBay. In this auction seller asks for bids from potential buyers. Several rounds of bidding are possible. Price starts low and is increased until only one bidder is left and the remaining bidder pays her/his bid. This mechanism is often also called an English auction.

In the English auction it is optimal for you to bid until the price is higher than your *private value*: If you continue bidding above your value and win, you need to pay a price that is higher than your private value. If you stop bidding below your value, you lose the potential gain from buying the item below your value. Bidder with the highest valuation will win and pay the second highest value. The same outcome can also be achieved with a second price auction: The highest valuation will win and the winner will pay the second highest valuation. This is the crucial difference to a first price auction where your optimal bidding strategy will deviate from truthful bidding.

Efficiency and expected revenue

Auction outcome is efficient if the high value bidder wins. Given the differences in bidding strategies, it may be somewhat surprising that the Nash equilibrium outcome is the same in several efficient auction mechanisms: The high value bidder wins and the expected auction price equals the expected value of the second highest bidder. This is the essence of the *Revenue Equivalence Theorem*: Every auction that allocates the goods efficiently has the same expected profits for every bidder valuation and the same expected revenue for the seller.

4 Market design

Like auction theory, mechanism design has a rich set of theoretical results of which we will try to find few key insights for the online applications. In the real world, the detail of designing a market will be dependent on the particular case. We have spent quite a bit of time to study the Google search ad auction as an example that is important on its own and as a basis of understanding how some of the theory work can be applied.

Mechanism design

Game theory is helpful in describing the rational strategic interactions between agents, such as the bidding strategies in auctions. Mechanism design turns the question around: Assuming that the agents will play a strategic interaction game, what should the rules for the game be?

In the study of existing markets, mechanism design can be used to identify the "rules of the game", the incentives for the participants, and how they behave. This understanding is required to try to comprehend why the market functions well, or not so well. In the design of new markets, mechanism design can be used to identify the economic problem to be solved, the players, and their incentives and information. The task is then try to understand what sort of market rules would lead to desired outcomes.

Impact of design to auction outcomes

Auction theory tells us what to expect in a Nash equilibrium: In an efficient allocation, the high value bidder wins and expected revenue is the same in many auction methods. But the outcomes of the auction will still depend on the level of competition in the market and the implementation detail. The expected revenue for the seller can be increased, for example, by facilitating competition and the introduction of reserve prices. The effect of competition is clear. If there are more buyers then the chances are better that some of them have higher valuations for the item that is being auctioned. This can make it worthwhile for the seller to even subsidize participation to auctions, but at least make the access as easy as possible. A reserve price works almost as a phantom bidder because the other bidders need to bid higher than the reserve price to win. When setting the reserve price, there is a balance though: a higher reserve price will reduce the probability of sales but increase the expected sales price.

Ad auctions

Advertisement auctions are the way many internet companies make their money. Google was the first to successfully switch to auctions as the mechanism for selling advertisement slots in their search platform. They operate a *generalized second price auction*, which tries to capture the essence of the second prices auction: Because the the payment of the winner is not dependent on her/his own bid, it shouldn't be too risky to bid your own true value for the ad slot. The method is also relatively easy to understand. Both of these factors should make it easier (=less costly) for advertiser to participate to the auction.

The generalized second price auction works as follows. An advertiser chooses a set of keywords that are related to the product it wishes to sell. Each advertiser states a bid for each keyword that can be interpreted as the amount that it is willing to pay if a user clicks on its ad. When a user's search query matches a keyword, a set of ads is displayed. These ads are ranked by bids (or a function of bids) and the ad with the highest bid receives the best position; i.e., the position that is mostly likely to be clicked on by the user. If the user clicks on an ad, the advertiser is charged an amount that depends on the bid of the advertiser below it in the ranking (from Varian 2007). Though in theory the bidders could gain by deviating from their private values, the mechanism has been widely adopted and is used in millions of transactions daily, and strategic bidding is not a widely reported problem. Facebook auctions are so-called Vickrey-Clarke-Groves (VCG) auctions which induce truthful bidding, but are more complex.

Other considerations for market design

Market design is always a balance between what tasks are left to the competition/game in the marketplace to decide and what is decided by the rules. Choice of implementation detail will depend on the desired outcomes. The following list, by Alvin Roth (2008), gives an idea of the different aspects that can be relevant in the design of successful markets.

1. Provide *thickness*. Attract a sufficient proportion of potential market participants to come together ready to transact with one another.
2. Overcome the *congestion* that thickness can bring. Ensure that market participants can consider enough alternative possible transactions to arrive at satisfactory ones.
3. Make it *safe* to participate in the market as simply as possible. As opposed to trading outside of the marketplace or engaging in strategic behavior.
4. Some markets can be *repugnant*: they should not exist. The design goal here is to try to avoid such outcome.
5. *Experimentation* to diagnose and understand market failures and successes, and to communicate results (at a firm level or as an outside observers for the regulation of the firms).

There needs also to be sufficient enforcement mechanisms in place to ensure that the actual outcomes meet desired outcomes (see section on Regulation).

5 Online markets

A lot of the efficiency gains from online and digital markets result from the reduction of frictions. Online marketplaces are a prime example of these gains. They also show how data can and is used to extract value from consumers.

Reducing frictions and managing competition

An online marketplace needs to serve two competing groups: Buyers want to find the products quickly and with low prices, sellers want to maintain high margins. The layout of an online marketplace affects consumer search and seller incentives at the same time. First, the search can be used to predict consumers' demand, and to guide them toward their most desired product. Second, the search can help consumers to find a retailer who offers an attractive price for the product. This increases the effective price elasticity faced by sellers, and increases competition.

The amount of competition among the sellers is determined on the basis of how many items the buyer can (effectively) see when making the search. This will affect which choices the buyer can make. Also, the order of the items or sellers displayed by the marketplace will affect which seller gets most sales. This ordering may be based on other criteria than price, giving control to the marketplace. The available customer choices will start to affect the pricing (and other) decisions by the sellers. Different marketplaces have adopted different strategies on how easy they make for the buyer to find the lowest prices, and how they treat the sellers, for example how much consumer data is shared with the sellers. Notice that these seemingly trivial details can greatly affect the success of a given platform as network effects will amplify the impact of design choices.

Use of data by the marketplace operators

Companies can collect and use data to their advantage in the short-run at least in three ways: *dynamic pricing*, *price discrimination*, and *product steering*. In the long-run data can be affect decisions on product development, investments, logistics and so forth. These long-term impacts may be more important than the short-term impacts, but the study of them will require even more machinery from the economics toolbox (such as more on competitive market structures and their impact on innovations, long-term development of the economy), and these will be mostly left to other courses. Focus here is more in the tangible short-term phenomena.

With dynamic pricing, prices for an item are adjusted depending on demand and/or competition. Online information on the bids of competitors makes it easy to adjust prices dynamically to match prices or have prices higher than the competitors. For example, if and when consumers are inattentive, Amazon has the advantage of deciding when it wants to sell the product and at what price. This also works the other way around: Other companies can start to match Amazon prices. Empirical evidence suggest that this price convergence is a widespread phenomena also in the offline world, and may even start to influence macroeconomic indicators, such as inflation.

Product steering and price discrimination are ways by which the firm tries to extract the maximum amount the buyers are willing to pay. If the firm has market power and is only able to offer *uniform pricing*, i.e. the same price for everyone, then it will withhold output to maximize its profit but that results in an efficiency loss as some trade will not happen. In such cases, the possibility for price discrimination can in fact increase the total welfare in the short-term allocation (but see section on Regulation for concerns). Based on the available empirical evidence, direct price discrimination doesn't seem to be that common in the digital realm. Instead firms seem to use product steering to push consumers towards higher margin products, and also they try to make money through the sale of additional services or high fees that are obfuscated.

6 Networks

It is easier to form networks online, for example, you can have more access to a larger group of friends through social media than offline. Understanding that the value of a network good depends on the number of other users is one of the key points in order to understand the economics of platforms.

Network effects

Network effects or *network externalities* are essential in understanding how interactions between users benefit larger firms that dominate the online world. We need to start with the analysis of traditional networks where network externalities arise because one particular subscriber can reach more subscribers in a larger network. The number of potential transactions increases with network size, see Figure 1. If every new transaction has positive value, a larger network gives higher value to a subscriber. The key reason for the appearance of network effects is the complementarity between network users: each user adds value to all other users. This is an example of *direct network effect*.

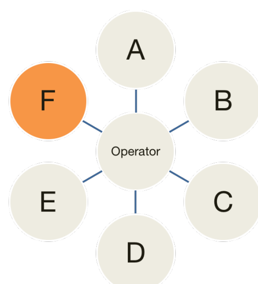

Figure 1: Example of direct network effects: Enabling connection F–A through the operator enables also connections F–B, . . . , F–E.

The network effects can also be *indirect*, typical examples are found in platforms that serve several different sides of a market. Users on side A of a platform do not need to benefit directly from other users on the same side, but they can create a benefit to the other side of the platform, which in turn benefits all users on side A. A credit card is an example of indirect network effects: User does not directly gain if one more person has the same credit card, but the additional person will encourage more merchants to accept a credit card. From a cardholder's perspective there is more choice and variety of merchants accepting this card.

Demand for a network good

The existence of positive network effects implies that, as sales expand, people are willing to pay more for the last unit. Figure 2 shows an example of how the demand for a network good may look like. Take 100 people who are in a market for a network good. Index the people from $v = 1, \ldots, 100$. v can here be thought to be the person's private type: the ones with a high type value the network more, the ones with a low type less. Now assume that the value (the reservation price or willingness-to-pay) of the network good to person v is vn, where n is the number of other subscribers to the network. If the price is set at some low enough p, then some individual is indifferent between buying the good or not buying it. Let's denote the type of this person by $\hat{v}$. So the indifference price is $p = \hat{v}n$ for some number of network users n.

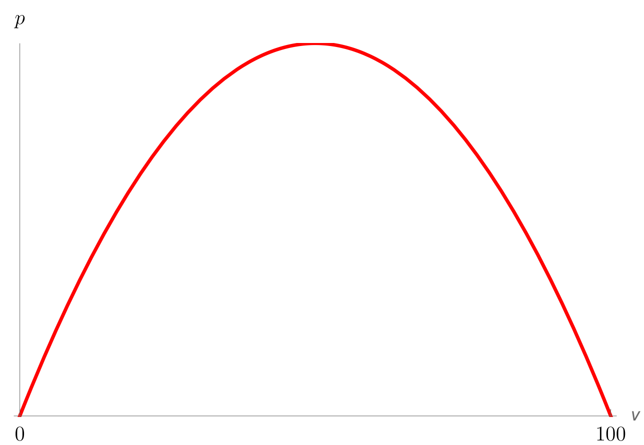

Figure 2: Equilibrium prices as a function of the network size.

Now how many users will join the network? If person $\hat{v}$ was indifferent, then everyone with a higher type will value the network more, so everyone whose type is higher than $\hat{v}$ will join the network and this gives the number of users. Remember, the total number of people in the market is 100. The number of people whose valuation is greater than index value $\hat{v}$ is $n = 100 - \hat{v}$. Solving for $\hat{v}$ and substituting to the price equation above gives the following equilibrium prices:

$$p = (100 - n)n.$$

Though the shape of the demand curve may look strange, it follows from common sense: If you consider joining a network, say a messenger app, its value to will depend on how many others you can reach through the network. Overall, some people will value communicating with others more and some less. The people who value the network most will join first, and the value increases rapidly. But as new persons join the network they will value it less and less, and at some point (the top of the curve) the value for those considering to join the network starts to decline despite the fact that the network is still growing.

Efficiency with network goods

In perfect competition for a traditional good, prices will equal marginal cost of production, and the allocation will be socially efficient. But if there are network effects, then marginal cost pricing is no longer efficient. When making a decision to join the network, individuals count only their own benefits. With a positive network externality, the marginal benefit to the society is greater than the marginal benefit to any one individual. In some cases, this could be an argument for government subsidies or regulation.

If a network operator is a monopoly, it will also not in general choose socially optimal size. Incentive of the monopolist to restrict output can outweigh its incentive to increase consumer demand. If the network firm is allowed to internalize the external benefits (e.g. via price discrimination), then the service level can be closer to socially optimal than with strict enforcement of competition. This price discrimination can be achieved also indirectly. As an example, a free-to-play game with optional purchases within the game allows a large number of players to join the game, increasing welfare[1]. The players who have a higher willingness-to-pay for the game can do so via the within game paid add-ons.

7 Platforms

Most successful online companies have organized themselves around platforms so it is given that the emphasis carries over to the course as well. Platforms are complex entities. The course tries to aid in understanding 1) why these platform form (frictions they solve, sides they serve), 2) how do the network effects between and within the different sides affect the dynamics and equilibria within a platform, and 3) how pricing works in principle.

[1]Not considering if playing can be consider a social gain.

Platforms over networks

The popular belief in 1990s was that the introduction of internet would lead to winner-take-all competition in many areas because of the network effects between internet users. With network goods, a service is more valuable if more customers are using it because customers want to interact with each other. If a firm moves fast and gets some customers, those customers make the network more valuable to next group of customers, which would attract even more people to join. In a simplistic model of competition, this could lead to explosive growth and result in a single firm owning the market forever.

The network effects story turned out to be only part of the competition narrative. First, as Rochet & Tirole (2003) realized: "[M]any, if not most markets with network externalities are characterized by the presence of two distinct sides whose ultimate benefit stems from interacting through a common platform." Network effects in many circumstances are indirect between different kinds of customers rather than direct between the same kind of customers. Second, network effects result from getting the right customers, and not just more customers. Finally, the same dynamics that can strengthen a network rapidly can work in reverse leading to rapid decline of networks.

Externalities within the platforms

There can be many sides within a platform, depending on the market and application. The following taxonomy (Evans & Schmalensee, 2007) has gained some traction:

- Ads: Find right ads to viewers in traditional media or users online.
- Exchanges: Match buyers with sellers.
- Software: Match developers, users, and hardware.
- Transaction: Match consumers with merchants.

The taxonomy arises from the differences between the externalities on the platform, both between the sides (inter-group) and within the sides (intra-group). For example, the inter-group externalities on an ad platform – advertisers like users, but users dislike advertisers – is different from an exchange where buyers and sellers both benefit from more participants on the other side. As an example of the differences in within the side or intra-group externalities, advertisers may dislike competition between other advertisers, but users in a social media platform may enjoy the presence of other users. See Evans & Schmalensee (2007) for more discussion.

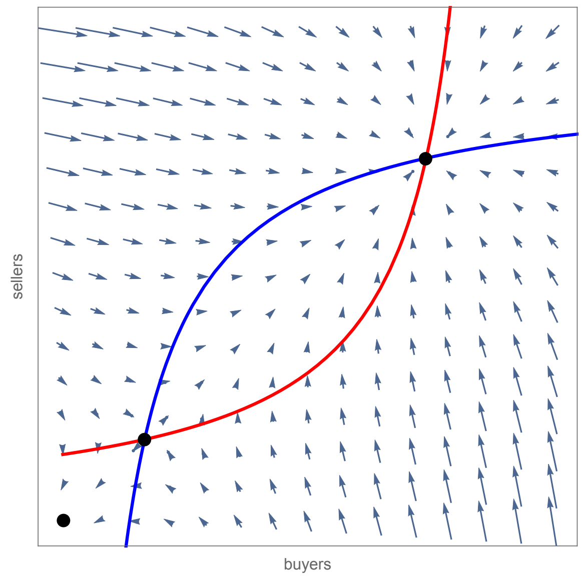

Figure 3: Indifference curves for buyers (red line) and sellers (blue line) with fixed prices. Pricing on each side affects demand on that side, but also on the other sides through externalities.

Figure 3 present a simple case of an exchange platform with buyers and sellers. The users on one side like to have more users on the other side, but do not have a preference on the number of participants on their own side. There is a fixed price for the buyers and a fixed price for the sellers. Red line is the indifference curve for the buyers, blue line for the sellers. For buyers to join, there needs to be a minimum number of sellers to make the platform worth the price, and vice versa. The prices determine where the indifference curves cross the horizontal and vertical axes, i.e. the number of sellers or buyers that would be needed for any participants from the other side to join to make the platform worth the price. The indifference curves intersect at two points: 1) The intersection with higher number of buyers and sellers is an equilibrium. In an equilibrium both the sellers and buyers who are in the platform are happy to stay with the given prices, and there are no additional sellers or buyers who would prefer to join the platform. 2) The other intersection point is a tipping point or critical mass; above that point the platform can start to grow towards the higher equilibrium, below it will shrink towards zero. Reaching a sufficient number of users to go over the tipping point is crucial for the success of the platform. Notice also, that zero is always one equilibrium point (no platform). See also Unit 21.5 of the Economy by CORE for discussion and a nice illustration of the dynamics.

Key strategic choice 1: Openness

Sometimes a firm can decide if it is a platform or not and how many sides it opens. For example, Apple launched iOS as a one-sided business with in-house apps, the possibility for third parties to sell through App Store came later. Sometimes the "platformness" is part of the business, like for many exchanges, software, and transaction systems.

Links on a network are potentially complementary, but it is compatibility that makes complementarity actual; a real world example is compatibility of nuts and bolts (if a bolt isn't the right size for a nut, then it is of little use). Some network goods are immediately combinable because of their inherent properties, like a road network. However, for many complex products, actual complementarity can be achieved only through the adherence to specific technical compatibility standard, like the internet. Within a platform, the firm chooses whether to provide all services itself or allow others to do some.

Producers' profits can be enhanced when they coordinate on a standard that permits the production of compatible components as there is less competition. A network good has higher value than a traditional because of the network effects. Different firms conforming to the same technical standard can create a larger network effect while still competing with each other in other dimensions, such as quality and price. Android is a good example. Main code is developed by Google and shared as open source for mobile phone makers to customize.

Providers of platforms can prefer incompatibility because it locks in current customers and locks out competitors. Network effects that are associated with the installed base generate *switching costs*, which are the costs of switching from one brand to another incompatible brand. Oftentimes in digital world compatible services are also hard to envision: e.g. Google search, Facebook or Amazon. Also, if users "pay" with e.g. the data that they generate, it may be harder to split profits compared to users paying money for compatible products to different firms (nuts and bolts).

Key strategic choice 2: Pricing

Platform pricing needs to take into account the externalities that a change in demand on one side causes in the other sides. Choice of price on one side affects how many users will use the platform. Number of users on that side will affect the interest on other sides through externalities. The changes in prices affect market equilibria. Returning to the example in Figure 4: A change in prices on e.g. the buyer side will change the demand on that side, but also affect the demand on the other side through the externalities.

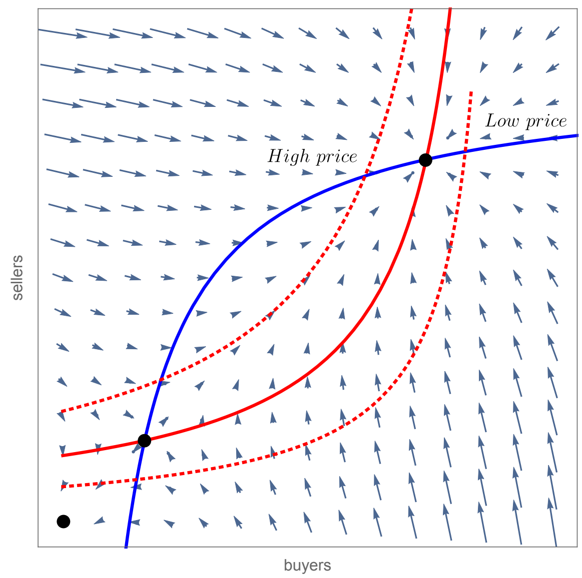

Figure 4: Indifference curves for buyers (red line) and sellers (blue line) with fixed prices. Pricing on each side affects demand on that side, but also on the other sides through externalities.

How the prices are chosen will depend on the level of competition. In general, prices are lower on the market side where users are more price responsive and which exerts a stronger externality on the other side. Many internet firms operate multi-sided platforms where content for users is free and the money comes from advertisements. By keeping the content free, the firms can attract more users than if they would require a payment. The large user base is attractive for the advertisers who are willing to pay for access to the users. This price subsidy paid by the platform to one side in order to attract more demand on the other sides is true in many cases, for example marketplaces charge transaction costs from the seller side to attract more buyers, credit cards offer rewards to the card holders etc.

A social planner in charge of platform pricing is interested in maximizing efficiency, which in this case means the surplus of all the users on all sides. In order to achieve this, the social planner would set the prices based on marginal cost of providing the service to each side plus the externality to the average user on the other sides of the platform.

Monopoly prices would be different from social optimum, because the monopoly is interested in maximizing its profits. In general this means that the sum of payments from all sides of the platform is higher than in social optimum and the platform is not used as much, in a similar way as a monopoly of traditional goods withholds output. A second distortion away from socially efficient prices results because the monopoly is only interested in the externality to the marginal user on the other side of the platform. Additional distortions away from socially optimal prices result as the number of participants in a monopoly controlled platform will be different from the socially optimal levels: this affects the quantity and quality of the possible interactions within the platform and the value of the platform to the participants.

Platform competition

In general, the effects of competition between platforms are complex and outcomes market specific. For example, in the video game industry (like Xbox vs. Playstation) competition can improve the availability of games, but in newspapers competition may further distort the amount of advertisement.

Multiple platforms are more likely to coexist, if platforms are sufficiently different (Android vs. Apple iOS), multihoming is feasible (travel search engines) or it is hard to differentiate within a platform. The companies can and do affect the competitive situation by their strategic choices. For example, if two marketplace platforms have a different pricing strategy (Taobao vs. Amazon) it may help them to differentiate on one or many sides of the platform. Because of the externalities, these differences may affect demand on the other sides of the platforms as well, leading to overall differences between the platforms. As a result, the platforms do not need to compete as fiercely on participants as they would in a symmetrical situation.

8 Sharing economy

The efficiency improving drivers of sharing economy are good to recognize on their own, and in the regulation we return to the trade-offs that such efficiency improvement can bring. Peer-to-peer platforms are also useful to discuss trust and reputation mechanisms that affect many other platforms as well.

Peer-to-peer market platforms

Sharing economy or *peer-to-peer markets* increase efficiency of the use of durable goods or labor. In traditional rental markets, owners hold assets to rent them out in a professional capacity. Sharing economy has introduced a new kind of rental market, in which owners sometimes use their assets for personal consumption and sometimes rent them out. Such markets are referred to as peer-to-peer or sharing economy markets. Sharing economy platforms within these markets solve matching, search, and transaction frictions between the parties.

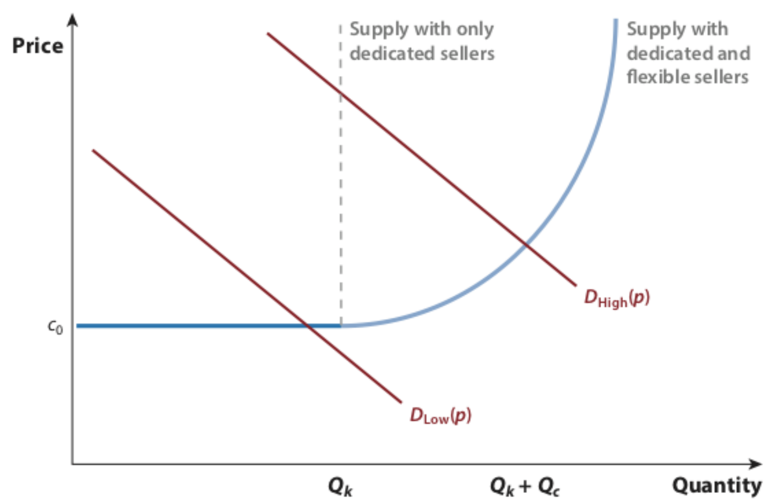

Figure 5: Sharing economy can increase capacity in traditional businesses (Einav et al. 2016).

Figure 5 gives an simple illustration of the economics. For example, Q_k could be the number of available rooms in Helsinki. In times of low demand (e.g. winter time), the capacity is sufficient to meet demand and prices are low. But in peak times (summer), there would be more demand and the market becomes capacity constrained. The total number of transactions is limited to the capacity and the prices paid by visitors increase. Now if there are homeowners who are willing to rent their own apartments for visitors, there is additional flexible supply. In the case of high demand, the supply from these homeowners allows more demand to be matched. As a result, the allocation becomes more efficient, because more transactions take place, and also the prices for visitors are lower.

The losers in this picture are the hotel owners, who cannot get the high rents anymore[2].

In addition to the potential efficiency gains, the sharing platforms have become more popular because they are often a cheaper option. Critics charge that the primary competitive advantage of P2P platforms is their ability to duck costly regulations that protect third parties. Rules and regulation in place for traditional businesses to protect consumers and limit externalities, for example: environmental regulation, consumer protection law, copyright law, health and safety laws, labor laws etc. Conflicts when non-traditional business models make rights and compliance requirements of the platforms participants unclear.

Establishing trust

Information asymmetry

In a peer-to-peer platform, the lack of responsible professional seller means that buyers face higher risks, but also the sellers may not have the capacity that the professional operators have to handle problems. The heterogeneity on both sides creates also information asymmetries and the problem for both parties to signal their quality.

In order to work, peer-to-peer markets have to maintain an adequate level of trust by developing mechanisms to guard against low quality, misbehavior, and outright fraud. The success of the markets as a whole is built on trust. These problems with establishing trust are evident also in "regular" online marketplaces.

Reviews

In addition to information asymmetries, the offering on sharing economy platforms is often not standardized (taxi vs. Uber), and can be complex (e.g. coding services). Reviews are a natural response to each of these problems. They describe the performance of a transaction counterpart, a user can alert others to what went right and what went wrong. Simultaneously reviews improve future matches and penalize bad behavior.

Review process however suffers from several complications. First, the reviews are subject to reviewers behaving strategically. Especially if both sides can make reviews, then the first one to give a review has to take into account that the review she is going to write may affect the review she is going to get. This leads to positive bias in reviews. Other problems include 1) that at the start, there are no reviews available, 2) user reviewed unfavorably

[2] Note that this analysis ignores externalities and potential longer-term implications: the supply from hotels may go down or they can response with pricing, etc.

might decide to start over and 3) users who submit reviews might differ systematically from normal users, by selection or by collusion as the targets of the reviews have an incentive to inflate their ratings to get more sales on the platforms.

To tackle the problem with reviews, platforms have added other information to reduce the need for reviews. Collecting and publishing such other information may however have unwanted discriminatory consequences (see section on Regulation).

More recently, AI algorithms have been developed to automatize the monitoring and enforcement of representative feedback. Examples include the assessment of the quality of sellers by analyzing the communication between the buyers and the sellers, and the creation of a market for feedback by automatic ranking the informativeness of user feedback.

9 Digital goods

The characteristics of digital goods determine why digital markets are different from traditional markets and how that affects the monetization of digital goods. Selling eyeballs is closely linked to ad auctions and the data questions very much relate to privacy concerns and the regulatory questions. Streaming services are an example of how the digital nature of good affects pricing decisions, beyond what has been covered in platform pricing.

Characteristics of digital goods

Digital goods are non-rival. Consumption of a digital good does not take away the possibility from someone else to consume that same good. Almost everyone could be watching the same movie from Netflix, but a movie theater would be constrained by the seat capacity.

The replication costs of digital goods is (nearly) zero. It costs nearly nothing to create and deliver a new copy of a digital good for consumption. A simple competition argument suggest that with sufficient competition, the price should go down to marginal cost. The question for the firms is how can digital goods be priced profitably?

Selling eyeballs

Selling advertisement space has long history in traditional media: newspapers, radio, TV. The main principle remains the same online, the main difference is the ability to track consumers more accurately and use the collected data to make the sale of eyeballs very precise, down to an individual search query or visit to a web page.

Individualized sales would not be possible without data crunching capability that scales to automatize the process, and the right mechanism design. Online ad auctions are a way to automatically extract value from advertisers: they force advertisers to reveal their willingness-to-pay for an ad slot and induce competition.

For the ad sponsored business model to be viable, the firms need to have content that provides utility to users. This value takes many forms, such as email, search, maps, marketplace, social media, messaging, and entertainment. Companies that can provide such utility can attract users to their pages. This allows for repeated opportunities to sell ad space.

Free access to a web page maximizes the number of users. Charging even a small fee would create transaction costs, both monetary and in terms of hassle. Also, because consumers seem to strongly prefer zero prices, there might be a large response to even a very small price charge for the service. If users are very price responsive, it follows from the platform pricing theory that their price should be subsidized. It would also seem plausible that consumers dislike advertisement less than advertiser like consumers, supporting the argument of lower prices in comparison to the advertisers. Finally, from the auction theory we know that increasing competition in the ad market side is bound to increase revenues more than what can be achieved by other means.

Selling data

Advertising revenues support new services on the Internet. Firms use tracking to collect data that is used to target ads, develop services, and engage customers. The data can also be sold forward, and such data can be collected only for the purposes of sale.

Targeted advertising gives consumers useful information. User see ads that should more precisely reflect their needs and less of the ads that are irrelevant to them. This improves the efficiency of advertisement. The users however face a trade-off: Sharing some information may be beneficial to them, but so can withholding some other information. The possibility that data is sold further can reduce the incentives to disclose data. The concept of data as a commodity has created new business for data intermediaries, whose significance is still somewhat unclear (these companies do not have incentives to be very visible).

Beyond the privacy concerns (see Regulation), the sale of data is done by the firms that collect the data, not the individuals themselves. They cannot yet efficiently buy back their data, or offer their data for sale. There is an argument that by assigning the property rights of data clearly to the individuals that would solve the data issues. However, it is clear that should change would require changes to (global) legislation and some means of enforcing the rules. Also, in some cases it would be hard to disentangle the data of the user from the supplementary data created by the services used.

Selling services

If marginal costs are near zero, why are not all digital services prices near zero? In many services network effects play a role and give the network operators pricing power. However, it is good to remember that the profits are due to innovation and investments. For example, the search algorithm by Google was more efficient in scaling up as internet grew compared to the search at Yahoo, the previously dominant search platform.

The non-rival nature of digital goods makes it possible to bundle products. Bundling several non-rival goods to heterogeneous consumers will extract higher rents from consumers than pricing the products individually. Bundling can also increase the overall efficiency as it reduces the dead-weight losses from unused goods compared to the sale of individual items at fixed prices.

Competition between bundle providers with similar content can result in price competition that the firms would prefer to avoid. One possibility is that competition will lead to the creation of more exclusive content with which the firms can attract/retain consumers, in particular those who are willing to purchase only one bundle. For example, increased competition to the incumbent, say Netflix, may induce it to introduce more exclusive content that it could have done already earlier but chose not to. This will enable the incumbent to protect its price against competition. The choice of the dominant player to not compete with prices allows also competitors to charge higher prices. In the end, there may be several competing bundles that each try to have exclusive content to attract customers with.

10 Regulation and policy

Regulation is done to protect consumers, to limit market power, or to control some externality that the markets create. It is always the case that there the regulator needs to balance the cost of action with the benefits that the regulation can achieve; the rapid pace of development makes this more true online.

Regulatory concerns

Typical motivations for regulation are the limitation of market power, consumer protection and externalities. When considering which action to take, the regulator needs to balance the potential harm with how regulation can be implemented and what the cost of regulation to consumers will be. Benefits of online markets and other platforms can be large enough to compensate for the costs and risks involved. Also, rapid development of new innovations may be unnecessarily hampered with too stringent or early regulation. Some of the current real or perceived problems may be solvable by the markets. Others won't.

Privacy

Nobel-winner Joseph Stiglitz has stated (2018) that "Big data and privacy... represents one of the biggest challenges to our society, and to competition law and consumer protection." Individuals often don't know that their data is being taken and what is really being done with it, and what the value of the data is. Privacy is consider to be human right on its own. The recurrent breaches of privacy where sensitive or economically valuable personal data is misused or ends up in the hands of bad actors can create other concerns, such as safety, frauds, and political interference. Big data and algorithms that process data are relatively new phenomena, that could warrant caution to limit unforeseen misuses.

Market power and efficiency

In some cases, there may be social value in the collection and use of data. In many other cases, data is being used to extract more consumer surplus out of buyers. This is possible if the firms collecting data can exercise market power, either because of the data or because of their market position otherwise.

The direct potential way to use data in an online market context is that it provides new means for price discrimination and product steering. These may be beneficial strategies for the companies, but the empirical evidence suggests that they are not that widely used. This may be a result of industry self-regulation, getting caught with price discrimination may alienate customers and lead to a regulatory backslash.

Longer-term, data can give a competitive advantage to those with more data, not necessarily the most efficient firms. Data can be used to improve e.g. logistics and marketing, offer ancillary services like financing, or develop new products and services.

Seemingly harmless use of data, e.g. through algorithmic pricing, may lead to (tacit) collusion. If the algorithms can immediately respond to price discounts by some seller that reduces the potential gain from setting the discounts in the first place. Sophisticated price discrimination may lead to narrower relevant product markets, i.e. if a company can target a group of consumers very precisely with a certain offering, the other companies' algorithms may decide that that group of consumers is not worth the effort, reducing competition.

Finally, a dominant platform may use market power in a new way by controlling what a user of that platform is shown. Few prominent examples are Google and Amazon. The algorithm that determines search results in Google is made by Google, giving it at least implicit control over what the users see. Likewise, the search within Amazon gives it similar power, but in addition Amazon can choose whether to sell a product themselves or let the third parties to compete over the product.

Discrimination

Use of data by algorithms is highly selective and may lead to discrimination of certain user groups. To increase trust, many platforms encourage users to provide personal profiles and even to post pictures of themselves. These features may facilitate discrimination based on sellers' race, gender, age, or other aspects of appearance. Discrimination has been documented in various types of services. Building trust without discrimination can be challenging. The inherent informational asymmetries need to be solved, but it can be hard to find the proper signal.

Other externalities

The development of technology drives the costs of collecting, storing, and processing big data down. Data will continue to be processed as long as it has value. This can increase the electricity consumption from data centers to many times over the current consumption. However, this does not necessarily mean that computing power should be regulated: If electricity generation is polluting, tax the pollution there, which will increase the price and affect the consumption of electricity in all sectors.

Gig economy platforms (e.g. Uber in the U.S., Wolt in Finland) typically try not to have employees, but contract their workers as self-employed. Workers are outside the normal labor law protections, and may find it hard to organize: there is no office and workers have no knowledge on their colleagues. Bargaining power in employment terms is more in the hands of the platforms. Algorithms that control the worker hours and pay bring efficiency to the platform, but are insensitive to workers. As a result, the rights of the workers for these sharing economy platforms have been challenges in the courts of many jurisdictions.

The possibility to rent out assets in a sharing economy platform may lead to increase in consumer purchase of such assets. For example with Airbnb, increased apartment prices, together with preference of short-term over long-term rentals, leads to more congested housing markets. Change from long-term rentals to short-term stays causes externalities to the neighbors. While the evidence is still being gathers, in some studies, the increase of ride-sharing has been associated with adverse environmental impacts. Overall the novelty of these phenomena and heterogeneity of effects across markets mean that the regulatory wisdom of how to best handle these externalities is still limited.

Who's to regulate

Regulatory concerns can be seen to require government intervention. The questions on should the regulation be on local, national or global level is much trickier as the interest of different regions differ. For example, EU has a higher interest in protecting its citizens against the interests of U.S. based companies than the U.S. Government: the monetary benefits end up in the U.S. but the potential harm is spread worldwide.

At the moment not clear how to achieve socially optimal market conditions. Public regulation may reduce the efficiency of the markets. For example, it may hinder innovation, reduce the entry of new competition, and there can be actual costs of organizing effective regulation. Information asymmetries support letting the firms decide on the best course of action as they are superior both with data and other resources to the regulators. The firms themselves have also an interest in self-regulation to remain viable. However, whether or not such private regulation is sufficient will still require public oversight.

Supplementary materials

As a final note, there is plenty of additional material if you are interested in some particular topics. See the separate note on supplementary readings.