PEER-TO-PEER MARKETS

Liran Einav
Chiara Farronato
Jonathan Levin

Peer-to-Peer Markets
Liran Einav, Chiara Farronato, and Jonathan Levin
NBER Working Paper No. 21496
August 2015
JEL No. D02,D47,L86

ABSTRACT

Peer-to-peer markets such as eBay, Uber, and Airbnb allow small suppliers to compete with traditional providers of goods or services. We view the primary function of these markets as making it easy for buyers to find sellers and engage in convenient, trustworthy transactions. We discuss elements of market design that make this possible, including search and matching algorithms, pricing, and reputation systems. We then develop a simple model of how these markets enable entry by small or flexible suppliers, and the resulting impact on existing firms. Finally, we consider the regulation of peer-to-peer markets, and the economic arguments for different approaches to licensing and certification, data, and employment regulation.

Liran Einav
Stanford University
Department of Economics
579 Serra Mall
Stanford, CA 94305-6072
and NBER
leinav@stanford.edu

Chiara Farronato
Stanford University
Department of Economics
579 Serra Mall
Stanford, CA 94305-6072
chiarafarronato@gmail.com

Jonathan Levin
Stanford University
Department of Economics
Stanford, CA 94305-6072
and NBER
jdlevin@stanford.edu

1 Introduction

In 1995, Pierre Omidyar started the consumer auction website that became eBay. An often cited story is that he recognized its potential when he put a broken laser pointer up for sale, and it sold for $14.83. The buyer turned out to be a collector of broken laser pointers. As the story suggests, the internet is a powerful tool to help buyers and sellers find each other. It has enabled the creation of marketplaces for local goods and services (e.g. Craiglist), computer programming (Freelancer, oDesk), consumer loans (Prosper, Lending Club), crafts (Etsy), start-up financing (Kickstarter), accommodation (Airbnb), baby-sitting (Care.com) and currency exchange (Transferwise, CurrencyFair). These days, hundreds of companies are trying to create markets for "on demand" services such as rides (Uber, Lyft, Blabla Car), deliveries (Instacart, Postmates), and household tasks (TaskRabbit, Handy).[1]

While these businesses each have specific features, they share common and innovative elements. They lower entry costs for sellers, allowing individuals and small businesses to compete with traditional firms. They rely on spot transactions, often eschewing long-term contracts or employment relationships. They take advantage of technology to improve the matching of buyers and sellers, or to implement flexible or auction-based pricing. They frequently do little up-front screening or certification, and instead try to maintain quality by using reputation and feedback mechanisms. And at least in some cases, they have made inroads by skirting regulatory barriers.

For economists, the rise of marketplace businesses has provided a fascinating set of examples of innovative market design. Companies such as eBay, Etsy and Airbnb allow thousands of sellers to experiment with prices, selling mechanisms, and advertising strategies. Finance platforms such as Prosper or Kickstarter use a variety of public good mechanisms to enable individuals to collectively fund loan or project investments. Labor markets such as oDesk

---

[1] As of June 2015, there were 583 peer-to-peer start-ups listed on Angel List, a website that tracks early-stage investment opportunities. Trying to figure out which of these are viable turns out to be an enjoyable form of procrastination. We were delighted to learn about Drizly, a Uber-esque service that delivers beer and alchohol at the press of a button, only to learn that it faced a dozen or more competitors. We again were enthused when we learned that Dufl had raised 2.5 million dollars to develop a service that moves your travel bag from city and city, washing your clothes between trips, although its price tag of $99 plus a monthly subscription fee led one caustic online commentator to remark: "it seems like an expensive solution to a problem that doesn't really exist."

or TaskRabbit allow buyers to run small-scale procurement auctions for specialized tasks. Businesses such as Instacart and Uber use centralized mechanisms to assign workers to jobs, but these mechanisms also rely on market forces. When a rider submits a desired route, Uber advertises the job to nearby drivers. The allocation of the job is invisible to riders, but Uber tries to balance demand and supply and limit wait times by adjusting prices to current market conditions.

These new businesses also collect large amounts of data, which has helped to advance a growing body of academic research on their market mechanisms, and on the impact of peer-to-peer businesses on traditional industries. In this paper, we take stock of this work, and the economic issues around peer-to-peer markets. We divide our review into three parts, first looking at the design of internnmet marketplaces, then the economics of peer production, and finallly at some regulatory issues that are emerging as peer-to-peer platforms become more prominent.

We identify and discuss some key issues in internet market design in Section 2. Businesses that hope to create successful marketplaces or platforms for matching buyers and sellers have to solve several problems. They need to help buyers and sellers find each other, either by developing a centralized assignment mechanism or by allowing for effective search. They need to set prices that balance demand and supply, or alternatively ensure that prices are set competitively in a decentralized fashion. And importantly, they have to maintain an adequate level of trust in the market, by developing mechanisms to guard against low quality, misbehavior and outright fraud.

In solving these problems, peer-to-peer businesses usually have to trade off between two important objectives: designing market mechanisms that efficiently elicit and incorporate dispersed information, and minimizing search and deliberation to keep the user experience convenient. To see an example, consider the cases of Airbnb and Uber. In accommodation, heterogeneity is a central problem. Given a choice of Paris apartments, not everyone will agree on a common ranking and sellers may have widely varying costs. The dispersed nature of information calls for a market design that prioritizes choice. In ride sharing, matches need to be made in real time. Most people heading home from a bar on a Friday night want an immediate and safe ride, and care less about choosing between a nicer car versus

a more experienced driver. It therefore makes sense for Airbnb to create a decentralized marketplace, while Uber uses centralized assignment mechanisms. We argue in Section 2 that many aspects of peer-to-peer market design can be analyzed similarly.

The way that peer-to-peer markets deal with quality assurance and trust is perhaps more surprising. One hardly needs a PhD in economics to see the potential incentive problems in paying upfront to order something online from an semi-anonymous seller. Yet we know plenty of professional game theorists who are perfectly content to spend the night in a random person's apartment that they found on the internet. What is especially striking about how this degree of trust has developed is that peer-to-peer businesses often do not perform much ex ante screening or certification. London drivers historically spent years studying to obtain a black cab license, and becoming licensed to run a bed and breakfast could take months. The application process to become an Uber driver or Airbnb host takes days or hours. In place of heavy up-front requirements, these businesses rely on user feedback to provide ongoing monitoring. We discuss the evidence on these trust mechanisms in Section 2.3, including subtleties such as one-sided versus two-sided reviews, what information to collect, whether to limit the set of reviewers, and how to provide incentives to leave meaningful and truthful feedback.

In Section 3, we turn from market design to the broader effects of peer-to-peer businesses on traditional industries. To do this, we set out a simple model in which goods and services can be produced by dedicated professionals or flexible peers. Certain features of markets make them more amenable to peer production. These include variability or diversity in demand, the absence of scale economies in production, and of course the existence of well-functioning spot markets to match buyers and sellers effectively. We think of the entry of peer-to-peer platforms as potentially performing two tasks: improving the efficiency of spot transactions and lowering the cost required for sellers to advertise themselves to buyers. These innovations especially benefits flexible sellers, who might not otherwise do enough business to justify large investments in advertising, reputation or customer relationships. We also discuss the extent to which platforms can generate revenue from users, and the longer-term prospects for peer production.

The final section of the paper considers the regulation of peer-to-peer marketplaces.

One highly publicized issue is how peer-to-peer businesses entering into local taxi and hotel markets should be incorporated into the existing regulatory structure. Another issue that will grow in importance if flexible work becomes more prevalent is at what point contract workers on peer-to-peer platforms should be viewed as employees of the platform. A further set of questions surrounds the collection of data on workers and customers, and the ways in which people's histories can be used or shared. Each of these areas is relatively novel, so after laying them out, we focus on posing questions for future research.

2 Market Design: Search, Pricing, and Trust

The goal of peer-to-peer businesses is to create trade between large numbers of fragmented buyers and sellers. Doing this efficiently requires solving several core market design problems. One is to match buyers and sellers effectively, while keeping search frictions low. A related problem is to establish prices, or to organize the market so that prices will be set competitively. Finally, an important and potentially difficult problem is to ensure that transactions are safe and reliable for buyers and sellers. Recent research sheds light on each of these problems, which we consider in turn.

2.1 Matching Buyers and Sellers

Peer-to-peer markets often are characterized by a high degree of heterogeneity. Buyers may be interested in very specific products or services, and sellers may be quite differentiated. This creates a substantive problem of matching buyers and sellers, and figuring out appropriate, and perhaps personalized, prices. Both problems have an important informational component. Information is dispersed about who should be matched and at what prices, so an effective market must aggregate information successfully. At the same time, there is a practical problem: platforms need to minimize mundane transaction costs, such as the time it takes to sort through options or communicate information. Many aspects of internet market design can be viewed as trading off between these two priorities: keeping transaction costs low and using information efficiently.

One solution to matching buyers and sellers is to centralize the process. This is the

4

strategy followed by on-demand services such as Uber. When Uber customers look for a ride, they specify the type of car they want (e.g. a black car, or an SUV), but not the exact driver. Drivers see a request and can choose whether to respond, but they aren't shown where the rider wants to go.[2] Although some riders might value the option to pick their driver, most are apparently happy to delegate the choice provided they get to their destination quickly and safely. It then becomes Uber's problem to ensure it has a sufficient supply of drivers at any given time, and to weed out problematic or unsafe one. Provided that this happens, centralization keeps transaction costs for riders and drivers low, certainly compared to hailing a taxi, but also compared to showing buyers a list of drivers to try to contact.[3]

In contrast, decentralized markets facilitate individual choice. In peer-to-peer markets where sellers are diverse and offer a wide array of products and services, a main challenge is to create a streamlined and informative search process. Platforms use different strategies to facilitate search or make recommendations, but the process often begins with buyers specifying what they want and being presented with search results. This can be straightforward in some cases and less so in others. A buyer who searches for a used textbook on Amazon probably cares about getting it quickly, reliably, cheaply and in good condition. Not surprisingly, Amazon displays all of this information prominently, placing the cheapest option at the top of its search results and highlighting products whose delivery is handled through the Amazon Prime program. In contrast, buyers looking for a weekend apartment in Barcelona may prefer different neighborhoods and different types of apartments. Consequently, Airbnb presents buyers with an initial set of options but prominently offers a variety of ways for buyers to refine their search, for instance by narrowing down the location or type of apartment or price range.

---

[2]Sometimes this mechanism does not work well, as one of us discovered when he ordered an Uber car in London after midnight, and the driver refused to drive him back to his house an hour away, on the grounds that it was too far a drive.

[3]Although the parallel may not be immediately obvious, the advertising markets run by Google and Facebook are centralized in a similar way. When there is an opportunity to show an advertisement, these companies run spot auctions to allocate the opportunity. Users get some say about what they want to see (someone who searches for auto repair on Google will not see ads for shoe stores), as do advertisers (who can submit bids targeted to certain keywords or demographics or browsing histories). However, Google and Facebook then adjust the bids to reflect their best guesses of what the user wants to see, and in that way try to leverage their superior data.

A consistent empirical finding is that the presentation of search results matters a great deal, even in settings where it seems as if buyers should be able to browse easily through multiple listings. In internet search advertising, buyers are about twice as likely to click a listing in the top position as they would be if it were moved one position down (Goldman and Rao, 2014). A recent theoretical literature has taken this as motivation and asked whether intermediaries have an incentive to present search results in ways that create maximal benefits for users. These papers point out that platform incentives may not align fully with consumers, especially if platforms obtain higher revenue if a buyer chooses a specific seller (Armstrong and Zhou, 2011; Eliaz and Spiegler, 2011; Hagiu and Jullien, 2011). This can happen if certain sellers pay higher fees, or are vertically integrated with the platform (de Corniere and Taylor, 2014).

A few recent studies have used data from internet markets to try to quantify search frictions. Dinerstein et al. (2014) compare shopping behavior and price competition on eBay under alternative search designs: rankings based on a "relevance score," and a multi-stage search process where the buyer first chooses the exact product, and then sees sellers ranked by price. They show that guiding buyers toward a price ranking can lead to higher surplus, but only when the relevant product is clearly defined with few variants. Fradkin (2015) looks at search frictions in Airbnb's more complex apartment rental market. He shows that on Airbnb, even after buyers identify apartments of interest, many transactions fall through. Transactions can fail because the seller rejects the buyer, or because multiple buyers contact the seller at the same time. Horton (2014) shows that the latter congestion problem is also common in the oDesk online labor market. He argues that it results in part from showing buyers similar seller rankings in a setting where sellers have limited capacity (see also Arnosti, Johari and Kanoria, 2015, for a theoretical analysis).[4]

---

[4]In these studies, buyers know what they prefer once they see different options. In practice, one of the ways in which internet marketplaces are able to create value is by creating recommender systems that steer buyers toward potentially desirable products. These systems generally rely on historial purchase patterns (in the case of Amazon's "people who bought X also bought Y") or other user feedback such as reviews (in the case of Netflix's recommendation system).

2.2 Pricing Mechanisms

The internet enables sellers and marketplaces to use a wide array of different pricing mechanisms. In the early days of electronic commerce, one of eBay's innovations was to introduce the use of proxy bidding, which enabled dynamic auctions to run over a period of days without buyers being attentive at every minute. Prosper, an early entrant in peer-to-peer lending, introduced an auction model in which borrowers posted a maximum interest rate, and lenders were able to make offers at lower rates. Labor markets such as oDesk allow buyers to post jobs and invite bids from potential suppliers. The internet advertising markets run by Google, Facebook and other firms also rely on spot auctions.

Auctions are appealing because they allow prices to respond to market conditions. However, contingent pricing does not necessarily require an auction. Marketplaces such as Airbnb, Etsy and Amazon make it easy and inexpensive for sellers to adjust prices in real time, which some sellers do using automated algorithms. Lending Club uses a proprietary algorithm to asses the riskiness of each potential borrower and sets interet rates based on this score, adjusting for market conditions and required risk premia. Uber similarly uses its surge pricing algorithm to vary the per-mile price of a taxi ride as supply and demand conditions change. In the latter cases, information collected and processed by the platform effectively substitutes for an auction mechanism.

We argued above that platforms often have to trade off between keeping transaction costs low and eliciting and using dispersed information. It is useful to think of pricing mechanisms from this perspective. In economic theory, the canonical pricing problem involves a seller with several potential buyers, each of whom privately knows his or her willingness to pay. Auctions tend to be the optimal mechanism to ensure an efficient allocation or maximize revenue, while maintaining proper incentives for buyers. But in practice, it can be cumbersome to identify potential buyers and sellers, and elicit information from them. As a result, simpler pricing mechanisms can make sense if information is available from other sources, or there is not too much uncertainty about what is the right price.

In Einav et al. (2015), we studied the trade-offs between auctions and posted prices using data on eBay sellers. As one would expect, sellers tend to use auctions for used goods,

or when they have less selling experience. More surprising is the fact that auctions have been in steady decline for more than a decade. One might hypothesize that this is due to compositional changes — that is, by trade shifting toward more commoditized products or more professional sellers. As it turns out, this is not the case. Instead, we show that for a given seller, offering a given item, the returns to using an auction were relatively high fifteen years ago and much lower today, partly because competition has become more intense and lowered seller margins, and partly because buyer's simply seem less interested in participating and bidding in auctions, even though they can expect on average to obtain a better price.

One of the points we make in Einav et al. (2015) is that the decline of auctions in e-commerce extends to other peer-to-peer markets. For instance, Prosper replaced its auction mechanism for funding loans in 2010, with a system where interest rates are set centrally based on the borrower's credit score. Gomez Lemmen Meyer (2015) studied this evolution and found that the centralized algorithm managed to price risk with just about the same effectiveness as the auction format, while simplifying the funding process. Another example comes from Cullen and Farronato's (2015) study of TaskRabbit's peer-to-peer labor market. TaskRabbit initially allowed buyers to either post a price for their job or to request bids and then pick their preferred offer. Because the number of active buyers and sellers turns out to be quite volatile, one might expect an auction mechanism to be particularly useful. But Cullen and Farronato found that auction prices do not vary much with market conditions (mainly because worker labor supply is quite elastic), suggesting that a simpler mechanism might be preferable. Indeed, TaskRabbit subsequently has moved toward a mechanism where workers post an hourly wage and schedule, giving buyers a simple and convenient way to hire.

The TaskRabbit example also illustrates another interesting point, that the pricing unit can matter a great deal. On TaskRabbit, most jobs are relatively standard, involving for instance house-cleaning or delivery. But they might take very different amounts of time. So setting a fixed price for cleaning jobs would be problematic, but setting an hourly rate makes a lot of sense. A similar insight was crucial in search advertising. Advertisers initially paid for impressions or pageviews. But the value of an impression can vary depending on where the ad is placed on a webpage and who is looking at the page. By instead pricing clicks

or conversions, advertisers have a better sense of what they are paying for, which makes it easier for them to bid for advertising opportunities (Milgrom, 2008; Varian, 2010). Getting the right transaction unit also can reduce monitoring costs. For instance, apartment rentals on Airbnb generally charge by the night, not the number of people staying in the apartment, which is much harder to monitor.

2.3   Trust and Reputation

Market transactions require trust, and this is especially true in markets that seek to facilitate spot trades between large numbers of dispersed buyers and sellers. In fact, one might be surprised these markets can work at all. When eBay started in 1995, it was not a slam dunk that people would send money across the country to nearly anonymous sellers, and these sellers would reciprocate by sending back items as advertised. Similarly one might have doubted that people would repay peer-to-peer loans, hire baby-sitters on the strength of a few online reviews, or rent rooms in their house to lightly vetted strangers. Yet all of these transactions seem to be workable.[5] Apart from general goodwill, what are the mechanisms that make this possible?

Trust can derive from upfront inspection, from reputation, and from external enforcement. Internet markets rely on all three, but often in different degrees from traditional markets. Inspection is more difficult when buyers and sellers meet online. This creates opportunities for misrepresentation, which Jin and Kato (2007) attempted to test in an early study where they compared the quality of graded baseball cards purchased online and offline. They found that online sellers were more likely to overstate the quality of their cards. Related evidence comes from Lewis (2011), who studied auto sales on eBay and found that buyers tend to be skeptical when sellers post few pictures of the car they are selling.

An alternative way to generate trust is for market makers to impose external regulations: limiting entry, certifying quality or insuring bad transactions. Amazon's sellers and Uber's drivers must adhere to minimum quality standards. Airbnb offers apartment owners the

---

[5]This is not to say that there are zero problems, as a recent New York Times story about a truly horrendous experience on Airbnb illustrated (Leiber, 2015). We have our own limited first-hand knowledge of bad experiences, for instance when one of us (Farronato) attempted to rent her car to a complete stranger using RelayRides, who promptly crashed it.

opportunity to have certified photos taken of their apartments, providing a signal to buyers that the apartment is accurately represented. Platforms may also offer to compensate users for bad experiences. In 2010, eBay introduced a buyer guarantee, and now compensates buyers if they purchase a product and the seller does not deliver as advertised (Hui et al., 2014). These interventions are costly but direct ways to ensure quality.

The most pervasive approach used by peer-to-peer platforms, however, is to employ reputation or feedback mechanisms. These mechanisms provide a substitute to inspection or upfront screening, are easy to set up online, and are used in some form in virtually all of our examples. They appear to have bite even though rsearchers have pointed out many of their flaws and shortcomings. For instance, eBay developed one of the first feedback mechanisms, allowing buyers and sellers to trade under pseudonyms rather than their real-world names. Disappointed buyers often do not leave feedback (Nosko and Tadelis, 2014), buyers can be deterred from truthful reporting by the threat of retaliatory feedback (Bolton et al., 2013), and since 98% of positive/negative feedback is positive, average feedback scores appear to have relatively little information content. (Horton and Golden (2015) report on a similar review inflation on oDesk.) All the same, eBay's reputation system seems to have worked well enough to screen out most of the really bad actors and deter highly fraudulent behavior (Resnick et al., 2002; Dellarocas, 2003; Cabral and Hortacsu, 2010).

In markets where the stakes to individual transactions are higher, or where personal safety is a concern, reputation mechanisms have become increasingly sophisticated. Prosper, for instance, collects and posts credit bureau information about potential borrowers, and Airbnb verifies the true identify of both buyers and sellers. Two-sided reviews also play an important role. For instance, Uber uses customer reviews to screen out problematic drivers, and it shows drivers the ratings of potential riders, so that riders who behave badly may have a harder time finding a ride in the future. Of course, heavy reliance on feedback scores raises the concern that users will seek to manipulate these scores. Mayzlin et al. (2015) argue that on review platforms such as Trip Advisor or Yelp, where anyone can post a review, manipulation is pervasive. One might expect manipulation to be more limited when reviews can only be written after a confirmed transaction.[6] But since reviews are in some sense a public good,

---

[6]Dowd (2015), however, amusingly describes her discovery that she could raise her low Uber rider rating

people may still under-report information that would be helpful to future customers. Fradkin et al. (2015) report on an experiment they conducted at Airbnb, showing how additional incentives for reviewing can improve information aggregation.

One interesting question is the extent to which a well-functioning feedback system removes the need for upfront screening or quality certification. If low-quality sellers or service providers can be quickly identified, the need to screen them out is arguably less important. Of course, if buyers insist on trading only with sellers who have strong feedback, the reputation system can create an entry barrier to new sellers, who might in fact benefit from a certification process. These trade-offs also arise and have been discussed in traditional markets, but it is interesting to consider how technology has affected the underlying trade-offs by making continuous monitoring so much easier. We will return to this theme below in discussing the general economics of peer-to-peer markets and approaches to regulation.

3   Peer Production and Traditional Industries

We now turn to the more general economics of peer production. There is less existing work to draw on, so we develop a simple, stylized model. We consider a market where flexible peer producers can offer services in competition with professionals who make upfront investments in dedicated capacity. For instance, Hilton building a hotel and an apartment owner renting a spare room are alternative ways to provide short-term accommodation. The introduction of peer-to-peer marketplaces makes it easier for small flexible suppliers to reach consumers, lowering the barrier for them to enter and compete. This can lead to changes in market structure, allow for trade in new services, and generate lower consumer prices. We use the model to identify conditions that are favorable for peer production, and to provide some context for our earlier discussion of market design, and subsequent discussion of regulation. In order to keep the presentation light, we defer some details of the model to the appendix.

---

by promising her driver "five for five" (a perfect rating in exchange for a perfect rating).

3.1 Peers versus Professional Sellers

We consider a market for product or service that can be produced by two types of sellers. A dedicated or professional seller incurs an upfront cost $k(q)$, to create q units of capacity — think of the construction of a hotel, the purchase of inventory or the hiring of full-time employees — and subsequently has marginal cost c_0 for each unit. In contrast, a flexible or peer seller has no upfront cost, and a marginal cost $c_0 + c$. The cost c is drawn from a distribution G, with support $[0, \infty)$. Sellers offer their services to a pool of buyers whose demand is variable. We write demand as $D_s(p)$, where s is the demand state drawn from a distribution H, and p is the market price. We assume that high values of s are associated with high demand, so $D_s(p)$ is increasing in s and decreasing in p.

Both professional and peer sellers must advertise their services to buyers in order to be recognized and make sales. We assume this advertising requirement takes the form of a fixed cost f that each seller must incur to become visible to buyers. The larger the cost of visibility f, the larger the advertising or reputational barrier to entry, and the fewer sellers who will be active in market equilibrium. Of course, professional sellers have an advantage when it comes to advertising because they can spread the fixed advertising cost over a larger number of sales. Later, we will think of peer-to-peer markets as providing a cheaper alternative to spending f, making it easier for peer sellers to become visible.

To streamline the model, we abstract from some realistic features of most peer-to-peer markets such as search frictions, product differentiation and seller market power. Instead, we assume that selling is competive, and that in each state the market price adjusts to equate demand and supply. In this spirit, we assume that despite the fixed entry cost f, scale economies for dedicated sellers are sufficiently limited to justify a competitive analysis. In particular, $(f + k(q))/q$ has a unique minimum q^*, which implicitly is small relative to market demand.

The timing of the model is as follows. First, potential sellers decide whether to enter the market. We will let Q_k denote the amount of dedicated capacity, and Q_c the amount of flexible capacity. Second, the demand state s is realized and peer sellers realize their marginal costs. Third, the market clears at a price p that equates demand and supply.

To solve, the model, we start by identifying the market clearing price that results in each demand state. Because all sellers have marginal cost of at least c_0, it is convenient to write the market price as $p = c_0 + \pi$, where $\pi \geq 0$ is the "price premium."[7] With this notation, supply is $Q_k + G(\pi) Q_c$, and demand in state s is $D_s (c_0 + \pi)$. Let

$$\pi^* (s) = \{\pi | D_s (c_0 + \pi) = Q_k + G(\pi) Q_c\}$$

denote the price premium that clears the market when $D_s (c_0) \geq Q_k$.

The market-clearing premium $\pi(s)$ is given by

$$\pi(s) = \begin{cases} \pi = 0 & \text{if } D_s (c_0) < Q_k \\ \pi^* (s) & \text{if } D_s (c_0) \geq Q_k \end{cases}. \tag{1}$$

Figure 1 illustrates how the market clears when there are only dedicated sellers, and when there is a mix of dedicated and flexible sellers. If demand is low, $\pi = 0$ and dedicated sellers are paid their marginal cost. If demand is high there is a positive price premium. With only dedicated capacity, capacity is fixed at Q_k and there may be a great deal of price variability, as in Prescott's (1975) well-known "hotel" model. With flexible sellers, short-run supply will be more elastic. In this case, demand variability is partially accommodated by a supply increase so prices fluctuate less. For instance, Cullen and Farronato (2015) studied the labor supply of peer workers on TaskRabbit and found it to be very elastic. They estimated that a 10 percent increase in the wage rate leads workers to apply for around 30 percent more jobs (see also Sheldon, 2015).

We next derive seller profits. A dedicated seller expects to make a profit $\pi(s)$ for each unit of capacity in state s. So her expected profit, per unit of capacity, is

$$U_k = \int_s \pi(s) \, dH(s) - \frac{f + k(q)}{q}. \tag{2}$$

It is apparent from this expression that each dedicated seller will maximize per-unit profits

---

[7]We have assumed that dedicated sellers always have lower marginal costs, so they will serve the market in low demand states, but it is straightforward to allow flexible sellers to have marginal costs that are sometimes below c_0.

by choosing the capacity level q^* that minimizes $(f + k(q))/q$, and in equilibrium this is exactly what will happen.

A flexible seller expects to sell only when $\pi(s)$ exceeds her marginal cost c. So her expected profit is

$$U_c = \int_s \left[\int_0^{\pi(s)} (\pi(s) - c) \, dG(c) \right] dH(s) - f. \tag{3}$$

Any increase in seller capacity reduces both spot market prices and seller profits. Increases in dedicated and flexible capacity are not identical, however. In the Appendix (Result 1), we note that an extra unit of dedicated capacity has a bigger effect on prices than a unit of flexible capacity, because of the lower associated marginal costs. Similarly, dedicated sellers in the market are effected more negatively by any increase in capacity — because they hope to sell in every state, their expected profits fall more sharply as spot market prices are reduced.

3.2   When is Peer Production Efficient?

Under our assumption that the industry is competitive, capacity will adjust to drive seller profits to zero. An equilibrium is therefore a pair of capacity levels (Q_c^*, Q_k^*) such that there is no opportunity for further entry, and no active seller wants to exit. That is, $U_c(Q_c^*, Q_k^*) \leq 0$ and $U_k(Q_c^*, Q_k^*) \leq 0$, meaning there is no incentive for profitable entry, and also, if a given type of seller $\tau \in \{c, k\}$ is active, $U_\tau(Q_c^*, Q_k^*) = 0$. We show in the Appendix (Results 2 and 3) that an equilibrium always exists, is unique, and is efficient in the sense that it maximizes consumer and total surplus.

Figure 2 illustrates the two zero-profit lines defined by $U_k = 0$, and $U_c = 0$. The number of flexible sellers, Q_c is on the x-axis and Q_k is on the y-axis. Each zero-profit line is downward sloping, and divides the space into regions where entry by the respective type of seller is profitable or unprofitable. The equilibrium, which involves a mixed market structure, involves each type of seller making zero profit. Other market structures are also possible. For instance, if the zero-profit line for dedicated sellers were everywhere above the line for flexible sellers, the equilibrium involves only dedicated sellers. And if fixed costs are sufficiently high, no sellers will enter at all.

14

What are the conditions that favor peer production? One factor is relative costs. If the upfront costs of capacity are low, or flexible sellers tend to have high marginal costs, peer entry will be difficult. To add a bit more detail, suppose we let $K^* = k(q^*)/q^*$ denote the per-unit capacity cost for a professional at efficient scale, and assume there is only a single demand state. In an equilibrium with only professional sellers, the price premium π must just compensate for upfront costs, so $\pi = f + K^*$. A flexible seller who enters expects to sell with probability $G(\pi)$, and have an average marginal cost $c(\pi) = \mathbb{E}[c|c \leq \pi]$. Its break-even condition is therefore $\pi - c(\pi) \geq f/G(\pi)$. Higher capacity costs, which mean a larger price premium, and more frequent low-end marginal costs, both help peer sellers.

Advertising or visibility costs are also a critical factor for peer entry. The higher is f, the fewer sellers in general but also the more the market structure will favor dedicated professionals. An interesting point here is that even in the absence of a specialized peer-to-peer marketplace, one might think of the internet in general as lowering advertising or visibility costs. In the Appendix, we work out the comparative statics associated with a fall in f. The result is more capacity in equilibrium and lower prices, but it is flexible sellers who benefit the most from the lower fixed costs. As a result, dedicated capacity actually can be crowded out by the entry of new flexible sellers.

A further condition favoring peer production is variability in demand. The logic is as follows (we provide a simple analysis in the Appendix). When demand is variable, the efficient form of production is to have at least some capacity that operates only part of the time. Dedicated investments are a costly way to provide capacity that will be used only some of the time, in constrast to flexible sellers who provide more elastic short-run supply. This observation seems germane for businesses such as Airbnb. In the hotel industry, demand can vary widely between peak and off-peak periods. Building hotel rooms is an efficient way to serve a fixed number of consumers, but variation in demand means either periods of very high prices or alternatively empty rooms. A similar story can be told about the taxi industry, where instead of prices adjusting to clear the market, drivers or riders tend to end up queueing or waiting around. In these settings, there is an economic logic to having an elastic supply of flexible drivers or accommodations.

3.3 The Rise of Peer-to-Peer Markets

We can now provide a stylized account of how peer-to-peer marketplaces enable entry and changes in market structure. We think of a new marketplace as introducing a technology — at some investment cost F — that allows sellers to become visible to buyers. For concreteness, imagine that marketplace fees are set competitively, so that when a platform enters it sets fees to just recover its investment cost F. An entering seller then has a choice between paying the direct advertising cost f, or instead joining the platform and paying its fees. Assuming buyers can purchase readily on or off the platform, sellers can expect the same market price π either way they offer their services.

The structure of fees will be important. If the platform charges a fixed fee to all sellers, then depending on whether its fee is above or below the direct advertising cost f, it will attract either all or none of the sellers. In the former case, it will still change the market structure toward peer production because of the lower advertising costs. However, if we think of the platform entering an existing market where dedicated sellers, such as hotel chains, already have built reputations, the platform may want to appeal specifically to peer producers. This makes it interesting to consider the case, which is indeed the common one we observe, where platform fees are per transaction. This structure will be particularly attractive to flexible sellers because they only have to pay the fee when they trade.

To see what might happen, suppose that there are Q_k dedicated sellers already in the market. If a peer platform enters and enables the entry of Q_c flexible sellers, the market price will be determined as described above. To break even, the platform will need to collect F/Q_c on average from each flexible seller. So these sellers expect a profit

$$U_c = \int_{\underline{s}}^{\overline{s}} \left[\int_0^{\pi(s)} (\pi(s) - c)\, dG(c) \right] dH(s) - \frac{F}{Q_c}. \tag{4}$$

Whether the platform can succeed depends on the general conditions for peer production, and on how effectively it can bring down visibility costs. If peers were previously deterred by high upfront costs (large f), then F/Q_c will need to be substantially smaller than f for the platform to be viable. Because the platform costs per seller fall with scale, a new feature of the model is that flexible sellers may in fact benefit from having other flexible sellers to

16

support a viable marketplace. Figure 3 illustrates this point, showing a situation where, with Q_k fixed, there are multiple equilibrium levels of Q_c.

Of course, this provides only a partial analysis. If a peer marketplace enters successfully, it will lower market prices. If dedicated sellers need to renew their capacity investments over time, the longer term effect is that peer entry will crowd out some professional sellers (Zervas, 2014; Farronato and Fradkin, 2015). Dedicated sellers also may want to use the platform alongside flexible ones. Indeed, eBay started as a consumer auction platform but became a sales channel for many larger retailers. Similarly, labor markets such as oDesk and Freelance have organized firms that bid for jobs, and some peer-to-peer financial service platforms have tried with varying degrees of success to attract professional lenders.

The proceeding discussion focuses on the way the peer-to-peer marketplaces can lower entry barriers. In practice, as we discussed in Section 2, they also can play an important role in creating more efficient spot markets for transactions. Our model envisions an initially competitive setting, but in many industries where peer-to-peer marketplaces are getting traction, there are pervasive market frictions. For instance, consumer lending can be a confusing market for borrowers, and is often associated with hidden or high fees. In traditional taxi markets, licensing restricts the entry of new drivers and prices are often completely independent of current demand conditions. In the case of deliveries or household tasks, a buyer looking for a one-off transaction historically would have had high search costs, little idea of what to pay, and perhaps little assurance that the person they hired would be reliable.

We conclude this section by noting a useful parellel between our model of peer production, and the so-called "long tail" of internet commerce which refers to the idea that local book or shoe stores might stock a few thousand items, while Amazon or Zappos stock hundreds of thousands (see Quan and Williams (2014) for recent work). The economics of the long tail are similar to the way falling fixed costs enable flexible entry in our model. Stocking a book involves an upfront fixed cost, so if books $1, 2, ...$ are ordered by the probability that there will be potential buyer in any given geographic market, a local retailer will want to stock some set of books $1, ..., n^*$, with $n > n^*$ unavailable. Because the internet enables a seller to reach buyers in many markets, it reduces the (per-market) fixed cost of stocking a book. As a result, it becomes economical to stock a larger product line. The "entry" of the book

17

that sells only with small probability is parallel to the entry of peer seller in our model, and creates a similar benefit for consumers.

4   Regulation of Peer-to-Peer Markets

The entry of peer-to-peer businesses has raised a number of a new regulatory issues that in some cases have become quite contentious. At this point, there has not been a lot of economic research focused on these issues, so we try in this section to lay out some of the questions that economists might hope to answer in future work.

4.1   Entry and Licensing Standards

Peer-to-peer businesses such as Airbnb and Uber have attracted a great deal of press attention because of their struggles with local regulators and incumbent businesses. A main cause of these conflicts is their perceived advantages when it comes to local regulation. Most cities limit the number of hotel rooms, the use of residential property for short-term rental, the number of taxis or for-hire cars, and the way in which taxi companies and hotel operators can set prices. Taxi and hotel businesses often are subject to specific licensing requirements, taxes and fees, and health and safety rules and checks. Peer-to-peer platforms, at least initially, managed to skirt many of these requirements.[8]

The resulting controversy has led to some hasty regulatory responses. Seattle's city council responded to the introduction of peer-to-peer ride services by limiting the number of cars and the number of driving hours per car (300 cars, across all ride-sharing platforms, and a maximum of 16 hours per week for each car). It also imposed regulations on cars and drivers, requiring vehicle inspections, driver background checks, and commercial insurance. The city council explained these regulations as necessary to protect both public safety and the interests of existing taxi drivers (Grossman, 2015). In Paris, regulators were prevented by a court ruling from banning Uber, but have interpreted the law as saying that drivers of paying customers must have special licenses, in effect making most Uber rides illegal

---

[8]Some of the recent debates eacho an earlier controversy about the tax advantage enjoyed by online retailers in the United States. Even today, many small online sellers are exempt from collecting state sales taxes as brick-and-mortar retailers must do (Goolsbee, 2000; Einav et al., 2014).

(Fourquet and Scott, 2015). Other cities, including New York, Berlin, and San Francisco have moved to regulate peer-to-peer apartment rental businesses.

What is the appropriate economic framework for thinking about these decisions? There are competing economic viewpoints. One account sees licensing and operating regulations as a response to market failures. For example, unregulated taxi drivers might take advantage of tourists, operate unsafe cars, refuse to serve certain neighborhoods, or create traffic congestion. From this perspective, regulations exist largely to protect consumers, especially vulnerable ones, from unscrupulous operators or adverse market forces. A contrary view, however, is that entry and licensing restrictions primarily serve the interests of incumbents by limiting competition (Stigler, 1971). From this perspective, restricting taxi medallions or hotel rooms raises prices and creates rents for incumbent operators, and peer-to-peer entry stands to increase availability, enhance comeptition, and perhaps raise service quality.

Apart from regulatory goals, one might also inquire about the form that local regulations take. Many requirements for taxi and hotel businesses, and for other service occupations, take the form of lengthy licensing or certification processes. Once a business is approved, there is relatively infrequent monitoring by regulators. As we noted above, marketplace businesses tend to have adopted a very different approach to ensure quality standards, with far more reliance on user feedback to provide continuous monitoring and far less upfront screening. It is natural to ask whether regulators could learn something from these markets, and try to incorporate new technology in the design of regulatory mechanisms.

Many economists are inherently sympathetic to the idea that new entry and competition, and technological advances, generally benefit consumers. In a recent survey of economists by the Chicago Booth IGM Forum, respondents uniformly agreed that "letting car services such as Uber or Lyft compete with taxi firms on equal footing regarding genuine safety and insurance requirements, but without restrictions on prices or routes, raises consumer welfare." (Disclosure: two of the authors of this paper are part of the IGM panel.) That said, the effect of new platforms for ride-sharing, short-term accommodation or other services on prices and quality, and their consequences for incumbent businesses, are really empirical questions. Providing convincing analyses seems like a useful topic for future work.

4.2 Flexible Workers and Employment Regulation

The use of contractors rather than employees is a common feature of recent on-demand service businesses. Because businesses such as Instacart, Postmates, Uber, or TaskRabbit run spot markets where workers are hired by the job, they are not subject to employment regulations, nor are they required to provide benefits such as health or disability insurance. While these businesses are a tiny fraction of the overall labor market, the idea that contract work and piece rates might encroach on long-term employment relationships is potentially very controversial. Indeed, the issue was highlighted recently when the California Labor Commissioner ruled that an Uber driver was technically an employee of Uber.

Our analysis in the previous section identified some of the possible efficiency benefits of part-time work. Some workers may value having flexible hours (Hall and Krueger, 2015). In other cases, variability in demand favors flexible work arrangements. Moreover, in a competitive labor market, a business that offers contract work must compete with firms offering stable employment, so piece rates must compensate for any lack of associated benefits. At the same time, however, employment regulations often are justified on the grounds that market forces do not sufficiently protect workers, something that did not enter into our model of peer production.

The question of whether technology could make spot labor markets more common, relative to long-term employment contracts, is potentially very consequential. Even absent regulation, however, it is possible that some peer-to-peer platforms, having used contract workers to rapidly scale up their businesses, might turn to stable employees to provide more consistent or highly-trained service. In this sense, there is room for both theoretical work to lay out the trade-offs that firms face in relying on contract workers, and empirical work to estimate the consequences, and evaluate the arguments for and against imposing employment regulations on peer-to-peer businesses.

4.3 Data and Privacy Regulation

A central feature of marketplace businesses is their reliance on user data and algorithms to match buyers and sellers, set prices, and monitor behavior. The detailed data collection, and

the use of personalized algorithms, raises some potentially challenging regulatory issues. For instance, what rights should consumers have to limit the way platforms use data? Should a platform be able to share or sell individual feedback ratings or purchase histories? If a worker performs poorly, or a consumer gets bad feedback when using a particular service, should they be able to expect a fresh start with another one?

A similarly thorny set of issues pertains to the use of personalized algorithms. Consumer finance, housing, and many service businesses are governed by strict regulations guarding against discrimination. Lenders, for instance, cannot use an applicant's race, gender, religion, age, or marital status to reject a loan or set interest rates. Presumably algorithms used by internet platforms will avoid using proscribed variables directly, but what if an algorithm that is trained to effectively use all permissible data ends up placing weight on a set of close proxies? For instance, if some platform users consistently give lower feedback to other users on the basis of gender or race, the latter group may get worse opportunities even if the platform does not explicitly discriminate. In response, platforms might argue that they are simply intermediaries and cannot be held accountable for the behavior of their users.[9]

These issues are only partly economic. A worker whose poor feedback follows her to her next job may suffer economic harm, but even if she does not, she may suffer an inherent — and difficult to quantify — cost from the loss of privacy. Advocates for data regulation often focus on the latter. Nonetheless, despite the potential measurement difficulties, there is room for economic analysis that looks at the arguments for regulating both the sharing and use of individual user data.

4.4   When to Regulate?

We finish this section by highlighting a point that has not received much attention in the current debate about regulating peer-to-peer businesses. It concerns the timing of when regulators should act. The businesses we have discussed, like many internet platforms, have the feature that once they can grow extremely fast if they start to succeed. And as they grow, their businesses can change rapidly. Sellers may become more professionalized, market

---

[9]This argument also comes up when a transaction goes wrong, and liability must be assigned. Platforms sometimes offer compensation but typically present it as being voluntary rather than required.

mechanisms are modified, technology can alter the way people use a service. One difficulty in trying to impose regulations on rapidly growing and evolving businesses is that regulations cannot easily be changed or withdrawn, so rules that look sensible at the time they were imposed may appear outdated or misguided. This logic suggests taking a relatively lenient early-stage approach to regulation.

The counter-argument is that in platform businesses, there can be a great deal of path dependence. Twenty years ago there were many marketplaces for used goods, and many competitors in search advertising, but eBay and Google became dominant by offering thicker markets and better services (Brown and Morgan, 2009; Levin, 2013). Today, regulators in Europe seem inclined to pursue antitrust action against Google, which would be a massive and lengthy undertaking given Google's size and importance. While one can argue about the likelihood of one of the newer platforms we have discussed achieving anything like Google's success, it is certainly the case that all of these businesses are characterized by at least some degree of network effects and scale economies. As a result, regulators may worry that by delaying action, they may miss the opportunity. As with the above issues we have raised, the question of when and how regulators should intervene in rapidly evolving markets deserves more thought.

5   Conclusion

Economic and business opportunities inevitably depend on the available technology. In the last two decades, information technology has facilitated the creation of larger, faster, and more geographically diverse marketplaces. In doing so, it has allowed consumers access to a wider and more personalized range of products and services. Internet marketplaces also have managed to deal fairly successfully with the incentive problems that arise in long-distance and semi-anonymous trade, and in doing so have enabled the entry and participation of small suppliers and flexible workers into many markets. In this review, we have tried to provide some insight into the market mechanisms used by internet marketplaces, and the economics that underlie peer production.

A significant risk in writing this review is that the businesses we have discussed are not

static. As marketplace businesses such as eBay became successful the sellers became more professionalized and larger retailers adopted the platform. Today, only a minority of eBay's business could be described as peer-to-peer. As we noted above, it may well turn out that providing rides, deliveries and other services with flexible contract workers does not prove to be a winning model, whether for economic, legal, or regulatory reasons. Nonetheless, peer-to-peer markets are likely to be an exciting area for at least some time, and hopefully this paper will encourage even more interest among economists.

References

Arnosti, Nick, Ramesh Johari and Yash Kanoria (2015). "Managing Congestion in Dynamic Matching Markets," Stanford University, unpublished.

Armstrong, Mark, and Jidong Zhou (2011). "Paying for Prominence," *Economic Journal*, 121, F368-F395.

Bolton, Gary, Ben Greiner, and Axel Ockenfels (2013). "Engineering Trust: Reciprocity in the Production of Reputation Information," *Management Science*, 59(2), 265-285.

Brown, Jennifer and John Morgan (2009). "How Much is a Dollar Worth? Tipping versus Equilibrium Coexistence on Competing Auction Sites," *Journal of Political Economy*, 117(4), 668-700.

Cabral, Luis, and Ali Hortacsu (2010). "The Dynamics of Seller Reputation: Evidence from eBay," *Journal of Industrial Economics,* 58(1), 54-78.

Cullen, Zoe, and Chiara Farronato (2015). "Outsourcing Tasks Online: Matching Demand and Supply on Peer-to-Peer Internet Platforms," Stanford University, unpublished.

De Corniere, Alexandre, and Greg Taylor (2014). "Integration and Search Engine Bias," *RAND Journal of Economics*, 45(3), 576-597.

Dellacrocas, Chris (2003). "The Digitization of Word-of-Mouth: Promise and Challenges of Online Feedback Mechanisms," *Management Science*, 49(10), 1407-1424.

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan (2014). "Consumer Price Search and Platform Design in Internet Commerce," Stanford University, unpublished.

Dowd, Maureen (2015). "Driving Uber Mad," *New York Times*, May 23, 2015.

Einav, Liran, Chiara Farronato, Jonathan Levin, and Neel Sundaresan (2015). "Auctions versus Posted Prices in Online Markets," Stanford University, unpublished.

Einav, Liran, Dan Knoepfle, Jonathan Levin, and Neel Sundaresan (2014). "Sales Taxes and Internet Commerce," *American Economic Review*, 104(1), 1-26.

Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan (2015). "Assessing Sales Strategies in Online Markets Using Matched Listings," *American Economic Journal: Microeconomics*, 7(2), 215-247.

Eliaz, Kfir, and Ran Spiegler (2011). "A Simple Model of Search Engine Pricing," *Economic Journal*, 121, F329-F339.

Farronato, Chiara, and Andrey Fradkin (2015). "Market Structure with the Entry of Peer-to-Peer Platforms: the Case of Hotels and Airbnb," Stanford University, in progress.

Forquet, Laure, and Mark Scott (2015). "Uber Drivers Face Fines in Paris" *New York Times*, February 22, 2015.

Fradkin, Andrey (2015). "Search Frictions and the Design of Online Marketplaces," MIT, unpublished.

Fradkin, Andrey, Elena Grewal, David Holtz, and Matthew Pearson (2015). "Bias and Reciprocity in Online Reviews: Evidence from Field Experiments on Airbnb," MIT, unpublished.

Goldman, Matthew, and Justin Rao (2015). "Experiments as Instruments: Heterogeneous Position Effects in Sponsored Search Auctions," unpublished.

Gomez Lemmen Meyer, Ana (2015). "Pricing Mechanisms in Peer-to-Peer Online Credit Markets," PhD Dissertation, Stanford University.

Goolsbee, Austan (2000). "In a World without Borders: The Impact of Taxes on Internet Commerce," *Quarterly Journal of Economics*, 115(2), 561-576.

Grossman, Nick (2015). "Regulation, the Internet Way," Harvard University Kennedy School, unpublished.

Hagiu, Andre, and Bruno Jullien (2011). "Why Do Intermediaries Divert Search?" *RAND Journal of Economics*, 42, 337-362.

Hall, Jonathan and Alan Krueger (2015). "An Analysis of the Labor Market for Uber's Driver-Partners in the United States," Princeton University, unpublished.

Horton, John J. (2014). "Misdirected Search Effort in a Matching Market: Causes, Consequences and a Partial Solution," *Proceedings of the fifteenth ACM Conference on Economics and Computation*, 357-357.

Horton, John J., and Joseph M. Golden (2015). "Reputation Inflation: Evidence from an Online Labor Market." NYU, unpublished.

Hui, X. A., M. Saeedi, Neel Sundaresan, and Zeqian Shen (2014). "From Lemon Markets to Managed Markets: the Evolution of eBay's Reputation System," unpublished.

Jin, Ginger Zhe, and Andrew Kato (2007). "Dividing Online and Offline: A Case Study" *Review of Economic Studies*, 74(3), 981-1004.

Leiber, Ron (2015). "Airbnb Horror Story Points to Need for Precautions," *New York Times*, August 14, 2015.

Levin, Jonathan (2013). "The Economics of Internet Markets," in D. Acemoglu, M. Arellano, and E. Dekel eds., *Advances in Economics and Econometrics*, Cambridge University Press, 2013.

Lewis, Greg (2011). "Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors," *American Economic Review*, 101(4), 1535-1546.

Luca, Michael (2014). "Reviews, Reputation and Revenue: The Case of yelp.com," Harvard University, unpublished.

Mayzlin, Dina, Yaniv Dover, and Judy Chevalier (2014). "Promotional Reviews: An Empirical Investigation of Online Review Manipulation," *American Economic Review*, 104(8), 2421-2455.

Milgrom, Paul (2008). "Simplified Mechanisms with an Application to Sponsored-Search Auctions," *Games and Economic Behavior*, doi: 10/1016/j.geb.2008.12.003.

Milgrom, Paul and John Roberts (1994). "Comparing Equilibria," *American Economic Review*, 84(3): 441-459.

Nosko, Chris, and Steven Tadelis (2015). "The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment," UC Berkeley, unpublished.

Prescott, Edward C. (1975). "Efficiency of the Natural Rate," *Journal of Poilitical Economy*, 83(6), 1229-1236.

Quan, Thomas W., and Kevin R. Williams (2014). "Product Variety, Across Market Demand Heterogeneity, and the Value of Online Retail," Yale University, unpublished.

Resnick, Paul, Richard Zeckhauser, Eric Friedman, and Ko Kuwabara (2002). "Reputation Systems: Facilitating Trust in Internet Interactions," *Communication of the ACM*, 43(12), 45-48.

Sheldon, Michael (2015). "Income Targeting and the Ride-Sharing Market," University of Chicago, unpublished.

Stigler, George (1971). "The Theory of Economic Regulations," *Bell Journal of Economics and Management Science*, 2(1), 3-21.

Varian, Hal (2010). "Computer-Mediated Transactions," *American Economic Review*, 100(2), 1-10.

Zervas, Georgias, Davide Proserpio, and John W. Byers (2014). "The Rise of the Sharing Economy: Estimating the Impact of Airbnb on the Hotel Industry," Boston University, unpublished.

Appendix: More Detailed Results from The Model

This section presents additional details on the model in Section 3. Recall that in the model, flexible peer sellers have unit capacity, pay no upfront capactiy cost and have marginal cost $c_0 + c$. The marginal cost c is drawn from a distribution G on support $[0, \infty)$. We assume it has density g. Dedicated professional sellers incur an upfront cost $k(q)$ to create q units of capacity, and have marginal cost c_0 for each unit. Each type of seller must incur a fixed visibility cost $f > 0$ upon entry. We assume that there is some q^* that minimizes per-unit fixed costs for dedicated sellers, that is $q^* = \arg\min_q (f + k(q))/q$. All dedicated sellers will enter at this scale. Demand is uncertain the time when sellers enter, and we write it as $D_s(p)$, where D is increasing in the demand state s, and strictly decreasing in the market price p.

Once entry occurs and demand is realized, the product trades on a competitive spot market. With dedicated and flexible capacity Q_k and Q_c, the market clearing price in state s equates demand $D_s(c_0 + \pi)$ with supply $Q_k + G(\pi)Q_c$. The quantity $\pi = p - c_0$ is the "price premium." As shown in Figure 1, if $D_s(c_0) \leq Q_k$, market clearing involves $\pi(s) = 0$. If $D_s(c_0) > Q_k$, then $\pi(s)$ solves $D_s(c_0 + \pi) = Q_k + G(\pi)Q_c$. It is convenient to define $s^*(Q_k)$ as the maximum value of s for which $\pi(s) = 0$. An entry equilibrium in the model is a pair of capacity levels Q_c^*, Q_k^* such that there is no further entry opportunity, $U_c \leq 0$ and $U_k \leq 0$, and no seller wants to exit. So $U_c = 0$ if $Q_c > 0$ and $U_k = 0$ if $Q_k > 0$. Expressions for the expected seller profits U_c and U_k are shown in the text.

Our first result shows the effect of capacity on spot prices and seller profits.

Result 1. *The spot market premium $\pi(s)$, and expected seller profits U_k and U_c are decreasing in seller capacity Q_c, Q_k. Furthermore,*

$$\frac{\partial U_k}{\partial Q_k} \leq \frac{\partial U_k}{\partial Q_c} = \frac{\partial U_c}{\partial Q_k} \leq \frac{\partial U_c}{\partial Q_c} \leq 0,$$

and the inequalities are strict provided that $s > s^*(Q_k)$ with positive probability.

Proof. First consider the spot market price. If $s > s^*(Q_k)$, then

$$\frac{\partial \pi(s)}{\partial Q_c} = G(\pi(s)) \frac{\partial \pi(s)}{\partial Q_k} = \frac{G(\pi(s))}{D'_s(c_0 + \pi(s)) + g(\pi(s))Q_c},$$

which is negative. And of course $\pi(s) = 0$ if $s \leq s^*(Q_k)$.

Now consider profits. For $\tau = \{C, K\}$, we have

$$\frac{\partial U_k}{\partial Q_\tau} = \int_{\underline{s}}^{\bar{s}} \frac{\partial \pi(s)}{\partial Q_\tau} dH(s)$$

and

$$\frac{\partial U_c}{\partial Q_\tau} = \int_{\underline{s}}^{\bar{s}} G(\pi(s)) \frac{\partial \pi(s)}{\partial Q_\tau} dH(s).$$

Combining these equations yields the claimed comparative static result. Q.E.D.

Our second result shows that equilibrium is unique and socially efficient.

Result 2. *There is a unique entry equilibrium Q_c^*, Q_k^*.*

Proof. Figure 2 in the main text shows the iso-profit curves defined by $U_c(Q_c, Q_k) = 0$ and $U_k(Q_c, Q_k) = 0$. Equilibrium is given by their intersection, and what ensures uniqueness is that the zero-profit line for flexible sellers is more steeply downward sloped than the corresponding line for dedicated sellers. We first show that at any point (Q_c, Q_k), the iso-profit line for flexible sellers will be more steeply sloped than the corresponding iso-profit line at that point for dedicated sellers. To see why, imagine a small amount of entry by flexible sellers and a small amount of exit by dedicated sellers that just offset to leave U_k unchanged. Such a change holds the average spot market price (which determines U_k) unchanged. But by replacing dedicated capacity with flexible capacity, the market price becomes higher in low demand states and lower in high demand states. This hurts flexible sellers, who are more likely to be selling in the high demand states. So to preserve flexible profits, Q_k must fall farther. Formally, at any point (Q_c, Q_k),

$$\frac{\partial U_k/\partial Q_c}{\partial U_k/\partial Q_k} = \frac{\int_{\underline{s}}^{\overline{s}} G(\pi(s)) \frac{\partial \pi(s)}{\partial Q_k} dH(s)}{\int_{\underline{s}}^{\overline{s}} \frac{\partial \pi(s)}{\partial Q_k} dH(s)} \leq \frac{\int_{\underline{s}}^{\overline{s}} G(\pi(s))^2 \frac{\partial \pi(s)}{\partial Q_k} dH(s)}{\int_{\underline{s}}^{\overline{s}} G(\pi(s)) \frac{\partial \pi(s)}{\partial Q_k} dH(s)} = \frac{\partial U_c/\partial Q_c}{\partial U_c/\partial Q_k}.$$

The inequality follows from the fact that G is increasing in s.

Now to show that equilibrium is unique, note that if U_c and U_k are negative for all (Q_c, Q_k), the unique equilibrium has no entry. If U_c is always negative but U_k can be positive, the equilibrium is given by the unique Q_k^* that solves $U_k(0, Q_k) = 0$. Similarly if U_k is always negative, the equilibrium is given by the unique Q_c^* that solves $U_c(Q_c, 0) = 0$. This leaves the case where profits for both types of sellers can be positive, as depicted in Figure 2. In this case, both zero-profit curves are downward sloping (Result 1) and the flexible seller zero-profit curve is more steeply sloped (from above). So the zero profit curves can cross at most once. If they do cross, the crossing point is the unique equilibrium. If they do not cross, one curve lies always above the other. The unique equilibrium has only the type of seller with the higher zero-profit curve, and the equilibrium capacity for this type of seller equates profit to zero. $\hfill$ Q.E.D.

Result 3. *The entry equilibrium Q_c^*, Q_k^* capacities maximize consumer plus producer surplus given the efficient spot market.*

We show that the marginal return to an entering seller is just equal to the marginal social surplus. Given capacity Q_c, Q_k in the spot market, the sum of producer and consumer surplus, ignoring fixed entry costs, is (note that to compute surplus, we integrate along the price axis in Figure 1, rather than the quantity axis):

$$W(s) = \int_{\pi(s)}^{\infty} D(c_0 + \widetilde{\pi}) d\widetilde{\pi} + \int_0^{\pi(s)} (Q_k + Q_c G(\widetilde{\pi})) d\widetilde{\pi}.$$

27

If $s \leq s^*\left(Q_k\right)$, then an extra unit of capacity adds nothing to welfare. If $s > s^*\left(Q_k\right)$, the welfare contribution of an additional unit of capacity is

$$\frac{\partial W\left(s\right)}{\partial Q_k} = \pi\left(s\right) \qquad \text{and} \qquad \frac{\partial W\left(s\right)}{\partial Q_k} = \int_0^{\pi(s)} G\left(\tilde{\pi}\right) d\tilde{\pi} = \int_0^{\pi(s)} \left(\pi\left(s\right) - c\right) dG\left(c\right).$$

This coincides exactly with the private return to entry, so entry will occur whenever it is socially efficient, but will not occur if it is socially inefficient. The equilibrium capacities will therefore maximize consumer plus producer surplus, taking fixed costs of entry into account. Q.E.D.

Our next results show the comparative statics results asserted in the text:

Result 4a. (Capacity Cost) *A decrease in $k\left(q\right)$ increases Q_k^* and decreases Q_c^*.*

Proof. For any $\left(Q_c, Q_k\right)$, a decrease in $k\left(q\right)$ has no effect on U_c but it increases U_k. (Note that the efficient scale may change if $k\left(q\right)$ changes but the per-unit profit for a dedicated seller choosing optimal scale will weakly increase with any decrease in $k\left(q\right)$.). The first conclusion then follow from Milgrom and Roberts (1994).

Result 4b. (Cost Structure) *A first order stochastic dominance decrease in c makes it more likely that flexible sellers can exist profitably alongside dedicated sellers.*

Proof. Suppose only dedicated sellers are present in the market and the spot price premium is $\pi\left(s\right)$.

$$U_c = \int_s \int_c \max\{0, \pi\left(s\right) - c\} dG\left(c\right) dH\left(s\right) - f$$

A decrease G in the sense of FOSD increases U_c, and hence makes it more likely that the equilibrium involves flexible sellers. Q.E.D.

Result 4c. (Visibility Costs) *A decrease in f increases flexible capacity Q_c^*, and total capacity $Q_c^* + Q_k^*$. It decreases expected spot market prices. It also decreases the efficient scale of dedicated sellers q^*, and reduces dedicated capacity Q_k^*, provided that flexible sellers were present in the initial equilibrium.*

Proof. We start with the case at some initial f, the equilibrium involves $Q_c^*, Q_k^* > 0$, and $U_c\left(Q_c^*, Q_k^*, f\right) = U_k\left(Q_c^*, Q_k^*, f\right) = 0$. Now consider a drop from f to f'. Each dollar drop in f increases the profit of both types of sellers by a dollar, so at Q_c^*, Q_k^*, we have $U_c = U_k = f - f' > 0$. To restore equilibrium, there must be entry by at least one type of seller. If both Q_c, Q_k increase, however, it will have a bigger effect on the profit of flexible sellers, so we will have $U_c > U_k$, which is inconsistent with having an equilibrium with both types of sellers present. The same will be true of Q_k increases and Q_c falls, which means that to restore equilibrium Q_c must increase and Q_k must fall.

The new equilibrium has a lower average spot prices. To see why, note that at the new equilibrium $U_k = 0$, so the average price premium now equals f' rather than f. Moreover, because entry by a flexible seller has a smaller effect on average prices than the exit of a

dedicated seller, it follows that for prices to fall, there must be net positive entry. Of course, prices need not be lower in every state. The exit of dedicated sellers means that there could be some demand states where previously we would have had $\pi = 0$, and now have $\pi > 0$ to induce flexible sellers into the market.

The same argument also applies if the initial equilibrium involved only flexible sellers. That is, Q_c must increase from its initial level Q_c^*, and dedicated sellers will not enter. However, if the initial equilibrium had no flexible sellers, and $0 \geq U_k > U_c$ at equilibrium, the situation is slightly different. By definition the new equilibrium must have at least as many flexible sellers, but it also can have more dedicated sellers. A simple example is where flexible sellers have such unfavorable marginal costs that they never enter, and then a fall in f clearly will increase the equilibrium number of dedicated sellers. *Q.E.D.*

Result 4d. (Demand Uncertainty) *Suppose demand is zero with probability α and $\frac{1}{1-\alpha}D(p)$ with probability α. Then greater demand variability (an increase in α) makes it more likely flexible sellers can enter against dedicated sellers.*

Proof. To enter, dedicated sellers require a price premium $\pi_\alpha = (f + k^*)/(1 - \alpha)$ in the event of positive demand. A flexible seller can expect to make $U_f = (1 - \alpha) \int_0^{\pi_\alpha} (\pi_\alpha - c) \, dG(c) - f$. It follows that

$$\frac{d}{d\alpha} U_f = \frac{d}{d\alpha} (1 - \alpha) \int_0^{\pi_\alpha} (\pi_\alpha - c) \, dG(c) = \int_0^{\pi_\alpha} c \, dG(c).$$

Flexible sellers are relatively more advantaged with variable demand. because they can avoid incurring their (higher) marginal costs in the event of zero demand, whereas a dedicated seller must incur its (higher) up-front cost in all demand states.[10]. *Q.E.D.*

Our final observation relates to the entry of a peer-to-peer market. With a peer-to-peer market that charges a break-even transaction fee, the cost of the platform will be split across participants. If only flexible sellers use the platform, the platform free per seller will be F/Q_c, and the fee decreases with the entry of additional flexible sellers. At the same time, the entry of additional flexible sellers reduces the spot market price. So in general flexible sellers may have their profits increased or decreased by the entry of additional flexible sellers, that is U_c may be increasing or decreasing in Q_c.

---

[10]There is an envelope theorem logic. In considering the effect of an increase in α, we can ignore the change in π_α which is set efficiently. So the first-order effect is that there is a slightly higher chance of zero demand, which results in dedicated sellers avoiding their marginal cost of c_0 and a flexible seller avoiding its expected cost of $c_0 + \int_0^{\pi_\alpha} c \, dG(c)$.

FIGURE 1: Spot Market with Dedicated and Flexible Sellers

Price

Supply with dedicated and flexible sellers

Supply with only dedicated sellers

$D_{High}(p)$

$D_{Low}(p)$

c_0

Q_k

Q_k+Q_c

FIGURE 2: Equilibrium Market Structure

Q_k

Q_c

$U_c(Q_c, Q_k) = 0$

$U_k(Q_c, Q_k) = 0$

(Q_c^*, Q_k^*)

Both Seller Types
Unprofitable

Both Seller Types
Profitable

FIGURE 3: Equilibria with a Peer-to-Peer Platform