



Outsourcing Tasks Online: Matching Supply and Demand on Peer-to-Peer Internet Platforms

Citation

Cullen, Zoë, and Chiara Farronato. "Outsourcing Tasks Online: Matching Supply and Demand on Peer-to-Peer Internet Platforms." *Management Science* 67, no. 7 (July 2021): 3985–4003.

Published version

<https://doi.org/10.1287/mnsc.2020.3730>

Link

<https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37379454>

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles (OAP), as set forth at

<https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFIOWI2M2EwYTg>

Accessibility

<https://accessibility.huit.harvard.edu/digital-accessibility-policy>

Share Your Story

The Harvard community has made this article openly available.

Please share how this access benefits you. [Submit a story](#)

Outsourcing Tasks Online: Matching Supply and Demand on Peer-to-Peer Internet Platforms*

Zoë Cullen[†], Chiara Farronato[‡]

June 2020

Abstract. We study the growth of online peer-to-peer markets. Using data from TaskRabbit, an expanding marketplace for domestic tasks at the time of our study, we show that growth varies considerably across cities. To disentangle the potential drivers of growth, we look separately at demand and supply imbalances, network effects, and geographic heterogeneity. First, we find that supply is highly elastic: in periods when demand doubles, sellers perform almost twice as many tasks, prices hardly increase, and the probability of requested tasks being matched falls only slightly. The first result implies that in markets where supply can accommodate demand fluctuations, growth relies on attracting buyers at a faster rate than sellers. Second and perhaps most surprisingly, we find no evidence of network effects in matching: doubling the number of buyers and sellers only doubles the number of matches. Third, we show that the cities where market fundamentals promote efficient matching of buyers and sellers are also those that are the fastest-growing. This heterogeneity in matching efficiency is related to two measures of market thickness: geographic density (buyers and sellers living close together) and level of task standardization (buyers requesting homogeneous tasks). Our results have two main implications for peer-to-peer markets in which network effects are limited by the local and time-sensitive nature of the services exchanged. First, marketplace growth largely depends on strategic geographic expansion. Second, a competitive rather than winner-take-all equilibrium may arise in the long run.

*We would like to thank Jonathan Levin, Susan Athey, and Liran Einav for their mentorship and guidance. We are also grateful to Juan Alcacer, Hemant Bhargava, Thomas Blake, Tim Bresnahan, Ramon Casadesus-Masanell, Michael Dinerstein, Pascaline Dupas, Andrey Fradkin, Shane Greenstein, Akshaya Jha, Brad Larsen, Michael Toffel, and many other seminar participants for their invaluable comments. Our appreciation goes to TaskRabbit for providing essential data access. Support for this research was provided by the B.F. Haley and E.S. Shaw Fellowship for Economics through the Stanford Institute for Economic Policy Research.

[†]Harvard Business School, zcullen@hbs.edu

[‡]Harvard Business School and NBER, cfarronato@hbs.edu

1 Introduction

The Internet has facilitated the diffusion of peer-to-peer markets for the exchange of goods and services. Users rent rooms on Airbnb, arrange rides on Uber, and find cleaning and moving help on TaskRabbit. These markets, which may compete with more traditional service providers, act as marketplaces for decentralized buyers and sellers to meet up and transact. Their growth and geographic reach have been highly heterogeneous. Uber, for example, is over ten times as big as Lyft, its closest competitor in the US; meanwhile, even though Chicago and New York have similar numbers of hotel rooms, Airbnb is much bigger in New York than in Chicago.¹ This paper studies the drivers of growth of peer-to-peer markets, and more specifically two key strategic decisions affecting their expansion. The first decision relates to a basic chicken-and-egg problem. Since peer-to-peer markets help match many fragmented buyers and sellers, their growth depends on balancing demand and supply. It is not enough to attract buyers if sellers' participation cannot accommodate their demand. At the same time, sellers may leave the marketplace if demand is too low. The second decision concerns entry across different geographies. When a peer-to-peer market is designed to match buyers and sellers of local services, cities can vary in propensity to conduct transactions on the platform. In addition, if network effects exist such that doubling the number of buyers and sellers in one city more than doubles platform matches, small initial differences across cities can compound into much larger differences in long-run platform success.

We use data from TaskRabbit, an online marketplace where buyers (formerly known as *posters*) can hire sellers (*rabbits*) to perform a wide range of domestic tasks and errands. We work with internal data from the company that allows visibility on all posted tasks, offers, and transactions. The setting allows us to focus on the efficiency and benefits of online marketplaces because successful matches must happen rapidly and locally. The geographic and time sensitive nature of services is not unique to TaskRabbit. Uber and Airbnb for example, share these characteristics. These time and location restrictions allow us to divide the activity on the marketplace into separate sub-markets by time and geography and use the large fluctuations in buyers and sellers to explore the matching of supply and demand.

The opportunity to observe multiple spot markets that match variable numbers of buyers and sellers with a single marketplace technology is empirically important. A common challenge faced by researchers in the study of online platforms has been the issue of defining and comparing separate markets. In examining eBay, for example, it is hard to divide buyers and sellers into geographically segregated markets given the prevalence of cross-state and cross-country transactions. A selection

¹<http://www.wired.com/2014/09/ubers-revenue-12-times-bigger-lyfts-new-study-says/> (accessed in June 2020), and Farronato and Fradkin (2018).

problem can also arise related to the fact that only platforms that have achieved a certain level of success are in use and can be studied. Given, however, that TaskRabbit operates essentially separate markets in different cities, and that we can observe these markets as they grow over time, we have useful variation for understanding demand and supply decisions and how scale economies might or might not arise.

In principle, there are several ways in which the market might function given fluctuations in buyers and sellers. One possibility is that with fewer sellers, buyers may not be able to have tasks performed, either because of higher prices that deter them, or because a smaller fraction of posted tasks receive offers. Another possibility is that seller labor supply expands. We find that in this context, labor supply is the key equilibrating factor. When demand is high relative to the number of sellers, the latter sharply expand their effort with very little price adjustment or reduction in the ability of buyers to trade. Elastic labor supply is a major contributor to the growth of the marketplace because it increases the number of matches created relative to a setting in which seller effort is fixed, and because it increases buyer retention. Because sellers adjust effort in response to demand increases, marketplace success depends on the level of task requests. This observation, combined with the finding that buyers' posting rate of tasks is fairly constant, implies that management efforts should be focused on attracting buyers at a higher rate than sellers.

Perhaps surprisingly, we find that matching does not display economies of scale. Doubling the number of requests and offers for tasks proportionally increases the number of successful matches. We confirm that this result is consistent with a possible mechanism driving scale economies. In particular, we find that the distance between the locations of buyers and sellers stays constant as the marketplace grows within a city. To partially address endogeneity concerns, we verify that our findings of an elastic labor supply and no returns to scale in matching are robust to a number of market definitions, and to using instrumental variables.

We find large differences across cities in the efficiency with which requests and offers for tasks are converted into successful matches. Cities where a higher share of requests and offers turn into matches also tend to attract and retain demand at higher rates. We relate matching efficiency to two measures of market thickness. The first is geographic density: the closer buyers and sellers live to each other, the higher the match rate between tasks and offers. The second measure is the level of task standardization: the match rate is higher when buyers request tasks in a few standardized categories, such as cleaning or delivery, for which sellers are more substitutable.

In what follows, we begin in Section 2 by presenting the relevant literature on multi-sided platforms, network effects, and platform growth. Section 3 describes TaskRabbit, and specifically how buyers post tasks and sellers submit offers to perform those tasks. We then introduce the key

motivating fact for our paper: growth is highly heterogeneous across the cities where TaskRabbit operates. We intentionally separate the growth of the peer-to-peer market in intensive and extensive margins. The extensive margins, the focus of Section 5, include platform adoption by new users and attrition, while the intensive margins, studied in Section 4, include the frequency of use by current users. Section 6 concludes, offering a reflection on the managerial implications of our findings for TaskRabbit and other peer-to-peer markets.

2 Literature Review

Our research contributes to a growing literature studying multi-sided platforms as a distinct business model (Evans (2003), Rysman (2009), Allon et al. (2012), Hagiu and Wright (2015)) that compete with more traditional service providers (Seamans and Zhu (2014), Zervas et al. (2017), and Farronato and Fradkin (2018)). Specifically, we touch on three distinct themes: platform market design, network effects, and platform adoption and growth.

From a market design perspective, theoretical work by Rochet and Tirole (2003), Parker and Van Alstyne (2005), Armstrong (2006), and later generalized by Weyl (2010) has focused on how to set platform fees as a function of users' participation decisions. Other non-price choices, such as information disclosure and search, have been explored by Boudreau and Hagiu (2009), Casadesus-Masanell and Halaburda (2014), and Economides and Katsamakas (2006) among others. Until recently, the empirical platform literature has focused on the interplay between the value created by a platform and that by complementary service providers, or sellers in our setting (Jiang et al. (2011), Huang et al. (2013), and Rietveld et al. (2019)).

A subset of the empirical platform literature has specifically focused on online peer-to-peer markets such as eBay, Airbnb, and Upwork, which match buyers and sellers of goods and services. Work in this area has studied the micro-structure of specific marketplaces, estimating search inefficiencies (Fradkin (2019)), heterogeneity in the matching process and problems of congestion (Horton (2019), Arnosti et al. (2019)), differences between distinct types of pricing mechanisms (Einav et al. (2018)), and the consequences of search frictions and platform design for price competition (Dinerstein et al. (2018) and Li et al. (2015)). There is also a large literature on trust and reputation systems (Luca and Zervas (2016), and Hui et al. (2016)), which dates back to early work by Resnick and Zeckhauser (2002), Dellarocas (2003), and Dellarocas and Wood (2008).

Our work complements this literature by empirically evaluating market equilibration in response to fluctuations in demand and supply. Furthermore, we study the effects of platform design and how market efficiency depends on city-level fundamentals. Our approach departs from the many forms of

individual heterogeneity and asymmetric information emphasized in other papers, as well as issues of strategic pricing and reputation. Instead, we offer a framework that enables us to shed light on the growth of peer-to-peer markets, commonly characterized by highly variable demand and supply (Hall and Krueger (2018) and Cachon et al. (2017)). Our analysis is applicable, in principle, to other peer-to-peer marketplaces that match buyers and sellers of local and time-sensitive services.

We also build on the network effects literature in our analysis of scale economies, and its implications for entry and competition. Users tend to limit themselves to a small number of platforms when network effects are strong, multi-homing is costly, and users have homogeneous needs (Katz and Shapiro (1985), Zhu and Iansiti (2012), and Gawer and Cusumano (2014)). At the extreme, a single platform may win the entire market. Indeed, since network effects compound small initial differences, a platform that by chance earns an early adoption advantage may end up as the leader (Arthur (1989) and Ellison and Fudenberg (2003)). One way to obtain an initial advantage is simply by moving first, as emphasized by Lieberman and Montgomery (1988). Empirical studies such as Bohlmann et al. (2002) and Cennamo and Santalo (2013) analyze the first-mover advantage and winner-take-all hypotheses in a multiplicity of industries, while Bresnahan and Greenstein (1999) focus on the evolution of platform competition in the computer industry.

In this specific context, where network effects can only arise locally and in a time-sensitive manner, we find that scale economies per se are not a major determinant of market efficiency, compared, for instance, to basic fixed features such as a city’s geography. The lack of scale economies is perhaps the most surprising result, and one that challenges some of the theoretical predictions in the existing literature on first-mover advantage and the winner-take-all equilibrium. Our findings lend support to alternative strategies for platform growth, such as platform envelopment (Eisenmann et al. (2006)) and product convergence (Greenstein and Khanna (1997)). Other more recent empirical tests of network effects include Li and Netessine (2020) and Farronato et al. (2020).

In examining platform growth through adoption and attrition decisions, this study also relates to a large literature on product and innovation diffusion pioneered by Griliches (1957) and Bass (1969), and how the speed with which new technologies grow can depend on information flows, technology improvements, and network effects (see for example, Tucker (2008) for an empirical analysis and Young (2009) for a theoretical contribution). The specific case of two-sided platforms requires balanced adoption by both buyers and sellers, a topic that has until now been studied mostly theoretically (Caillaud and Jullien (2003), and Hagiu (2009)). Here we empirically quantify demand and supply elasticities as a way of identifying a platform’s relative efforts to attract buyers and sellers.

3 Setting and Data

This section describes TaskRabbit. We first present how the marketplace operates, how tasks are posted, and how offers are made and accepted. We then show that matches are either made quickly and locally, or not at all. Finally, we provide an initial look at differences in market growth across cities, our study’s main focus, and some preliminary evidence on the drivers of growth.

3.1 The TaskRabbit Marketplace

TaskRabbit is an online peer-to-peer market that allows *posters* to outsource domestic tasks to *rabbis*.² Between 2009 and mid-2014, it operated in 18 major cities in the United States as well as London, UK.³ Posters post a description of the requested task. Rabbits can search through posted tasks on city-specific lists and respond with offers (Appendix Figure A1). We will refer to posters as *buyers* and rabbits as *sellers* of services.⁴

Buyers on TaskRabbit can post virtually any sort of domestic task or errand (e.g., pet sitting a goldfish), but the majority of tasks are relatively standard. The five largest categories of the 38 offered on the platform are delivery (20% of posted tasks), moving help (13%), cleaning (10%), minor home repairs (7%), and shopping (6%).⁵ These tasks do not typically require sellers with highly specialized skills. Services are generally provided locally and on relatively short notice, with 97% of the tasks that are matched (48.5% of all tasks) completed within one or two days.⁶ Almost all users (93.6%) participate in just one city.

The matching process works in two ways. A buyer can either post a task-specific price and then accept the first offer, or ask for bids and review the prices offered by sellers. Fixed-price tasks are

²In 2017, TaskRabbit was acquired by IKEA (<https://www.forbes.com/sites/blakemorgan/2017/09/30/3-lessons-learned-from-ikeas-acquisition-of-gig-economy-start-up-taskrabbit>, accessed in June 2020.)

³The active cities in the US are, by order of entry: Boston (2008), San Francisco (June 2010), Los Angeles (June 2011), New York (July 2011), Chicago (September 2011), Seattle (December 2011), Portland (January 2012), Austin (February 2012), San Antonio (August 2012), Philadelphia and Washington DC (July 2012), Atlanta, Dallas, and Houston (August 2013), Miami and San Diego (October 2013), and Phoenix and Denver (November 2013).

⁴Leah Busque first formulated her idea for TaskRabbit when she realized one evening that she had run out of dog food. Together with her husband, she started contemplating the idea of “a place online where we could say we needed dog food, name the price we’d be willing to pay, and see if there was someone in our neighborhood who would be willing to help us out” (<http://www.fatbit.com/fab/young-self-made-millionaires-women-entrepreneurs-making-difference-us-economy-part-1/>, accessed in July 2019).

⁵The 38 task categories include the following, from most to least popular: Delivery, Moving Help, Cleaning, Minor Home Repairs, Shopping, Research, Furniture Assembly, Yard Work & Removal, Event Staffing, Usability Testing, Computer Help, Office Administration, Packing & Shipping, Marketing, Organization, Carpentry & Construction, Writing & Editing, Painting, Data Entry, Pet Care, Electrician, Photography, Laundry, Selling Online, Automotive, Web Design & Development, Plumbing, Event Planning, Graphic Design, Cooking & Baking, Videography, Arts & Crafts, Sewing, Executive Assistant, Accounting, Child Care, Senior & Disabled Care. The remaining tasks are posted under “Other.” Table A13 displays the share of tasks posted in each category with a share higher than 1%.

⁶To add to the local and urgent nature of tasks, TaskRabbit’s ranking algorithm prioritizes newly posted tasks within each city, such that sellers see a list of local tasks, ranked according to their posting time (most recent at the top).

slightly more standard (65% are in the top five categories versus 48% of auctions), and prices are lower (\$49 versus \$63), but their share of the marketplace, at 41%, does not change considerably over time or across cities. Even if a task receives offers, a match can fail because the buyer finds a better alternative and does not select any of the bids received, or because the buyer and seller cannot agree on specific task details.

Users on TaskRabbit tend to be either buyers or sellers, but not both. Indeed, 80.3% of users have only ever posted task requests, and 16.3% have only ever submitted offers. Buyers are predominantly female (55%) and relatively affluent. Among users for which information is available, the modal buyer is a woman between the age of 35 and 44 with a household income between \$150,000 and \$175,000. The sellers are younger and, unsurprisingly, have lower incomes. The modal seller is 25-34 years old and has a household income between \$50,000 and \$75,000.

While buyers go through a basic verification process that checks their identity on social networks and their payment method, there is a more rigorous screening process for sellers. Until the spring of 2013, applicants were subject to a background check, were required to complete a digitized survey of their motivations, skills, and availability, and were interviewed by TaskRabbit employees to determine their fit. Acceptance rates of sellers' applications varied widely, ranging from 7% to 49% in different months, with a low average of just 13.6%. In the spring of 2013 TaskRabbit reduced the amount of screening in a successful attempt to add more sellers. As of May 2014, the end of our sample period, the process involved simpler background checks and social controls – Facebook or LinkedIn verification – paired with a system of user reviews.

3.2 Data

Our study uses internal data from TaskRabbit. We focus on the period from June 2010 to May 2014. During this period, TaskRabbit operated in 18 cities, although entry in these cities was staggered over time. We define the month of entry into a city as the first calendar month in which 20 or more local tasks were posted.⁷

The data include all posted tasks, offers, and matches that occurred on TaskRabbit during the study period. We exclude virtual tasks⁸ (10.4%) and tasks posted in cities that were not yet active (0.23%). We also drop the 10.3% of tasks that use assignment mechanisms other than auctions and fixed price tasks. We merge tasks with their corresponding offers, and drop extreme price outliers (the top and bottom 1% of bids or charged prices). In the occasional cases of fixed price tasks

⁷We have no record of the actual entry dates, but verified the accuracy of our definition through media coverage and by talking with TaskRabbit employees.

⁸A task is classified as virtual if the service does not require the seller to be at a specific location. Examples include writing and editing, or usability testing of mobile applications.

receiving multiple offers (6.04% of such posts), we keep only the matched offer in case of success (73% of cases), or select one of the received offers at random. This simplification restricts fixed price tasks to either one or no offers. Finally we aggregate activity at the city-month level, and drop city-months in which fewer than 50 buyers posted tasks or fewer than 20 sellers made offers.

Table 1 shows summary statistics for the data. In the first panel, an observation is a posted task. Of all the tasks posted, 78% receive offers, and those tasks receive 2.8 offers on average. Of the tasks receiving offers, 63% are successfully completed at an average price of \$57. TaskRabbit charges a 20% commission fee on successful tasks.⁹

In the second panel of Table 1, an observation is a city-month. We define a buyer as *active* in a city-month if she posts at least one task in that city-month. Analogously, a seller is active if he submits an offer to a task posted within the city-month. On average, there are 708 active buyers and 255 sellers in a city-month, although there is variation across cities and months. Each buyer posts 1.6 tasks, and each seller submits 6.4 offers. The task success rate is 46% and the average price paid is \$56. Of these four variables (tasks per buyer, offers per seller, task match rate, and prices), the number of offers per seller varies the most across city-months, with more limited variation in tasks per buyer, matches, and prices.

The platform’s success is highly heterogeneous across geographies. While TaskRabbit grew in all cities during the four-year period under study, this growth was much more rapid in some places compared to others. Figure 1 plots the number of successful matches for the ten oldest cities.¹⁰ Over the period considered, some cities grew from a few monthly matches to thousands of exchanges, like San Francisco and New York, while others grew at a slower pace, like Portland and Seattle.

The platform is also characterized by large fluctuations in demand and supply. The need for quick, local matches raises the question of what happens when demand is especially high or low relative to the number of sellers. Figure 2 shows the monthly variability of demand relative to supply in the 10 oldest cities. Specifically, it plots the number of active buyers per city-month divided by the number of active sellers. There are sizable fluctuations in this ratio, both within a city over time and across cities within a month. In San Francisco, for example, the monthly number of buyers per seller ranges from two to six. During a single month, some cities may have only one buyer per seller, while other cities have five. Month-to-month changes in the buyer to seller ratio are both positive and negative in no clear pattern, though there are some persistent differences across cities and across months. For instance, San Francisco consistently has more buyers per seller

⁹This commission fee can sometimes vary, for example in the case of coupons, referral bonuses, or other credits that reduce the price paid by buyers without affecting the price received by sellers.

¹⁰The 8 youngest cities show similar patterns to those in Figure 1.

than Los Angeles.

Table 2 provides initial evidence that a highly elastic labor supply equilibrates the market when buyer participation increases or decreases relative to the number of sellers. We divide the city-month observations into quartiles of the buyer to seller ratio. For each quartile we compute the average number of tasks per buyer and offers per seller as well as the task match rate and average transacted price. Table 2 shows that regardless of the number of buyers per seller, buyers always post around 1.6 tasks each. Meanwhile, sellers submit many more offers when the number of buyers per seller increases: that is, when sellers are scarce relative to demand. In the lowest quartile of the buyer to seller ratio (1.5 buyers per seller on average) sellers submit 4.4 offers on average, a number that more than doubles to 9.2 in the highest quartile (3.8 buyers per seller).

Table 2 also shows that the increased seller effort leads to a relatively constant probability – around 0.45 – that a task is successfully matched. Perhaps surprisingly, transacted prices move very little when sellers are scarce or abundant. The average transacted price is always between \$52 and \$60, even if the number of buyers per seller doubles and each seller chooses to work harder. Putting aside possible issues of task composition and seller heterogeneity, an apparent implication is that a small price increase can generate a large intensive margin response in labor supply. In the next section, we confirm these descriptive results and evaluate whether the market displays economies of scale.

4 Short-Run Market Equilibration

In this section, we study the market equilibrium as a function of the number of active buyers and sellers. We highlight two main results. First, we confirm that supply is highly elastic. When demand doubles, sellers work almost twice as hard and prices hardly increase. Second, we find no evidence of network effects in matching: doubling the number of buyers and sellers only doubles the number of matches.

We consider a static market, where the number of buyers and sellers is given. We assume for simplicity that buyers are all identical and choose to post the same number of tasks in equilibrium. Similarly, sellers are identical and choose to submit the same number of offers. We also treat tasks as homogeneous. Though this is certainly a major simplification, it does match our earlier observations that most tasks fall into a limited number of categories and do not require specialized skills. More importantly, it allows us to focus on the problem of supply and demand fluctuations and increasing market size without the complications of a heterogeneous matching framework. We take long-run participation decisions as given. Instead of focusing on users’ decisions to join or

leave the marketplace, we concentrate on how much participating users choose to engage with the market. Separately, Section 5 looks at adoption and attrition as additional drivers of growth.

We consider a setting in which each user is small relative to the market, such that a given participant cannot affect aggregate outcomes. In the market, buyers choose the number of tasks to post. Their choice is a function of the expected probability of them finding a match and the price they expect to pay. Buyers will find it easier to find a match and pay lower prices in a market with a small number of buyers and a large number of sellers. A higher match probability and lower prices, in turn, incentivize buyers to post more tasks, so we should expect that buyers' posting rate decreases with the number of buyers and increases with the number of sellers. Symmetrically, sellers choose the number of offers to submit, and we expect sellers' offer submission rate to increase with the number of buyers and decrease with the number of sellers.¹¹

These predictions hold as long as posting tasks and submitting offers is costly, or as long as there is an increasing cost for sellers to provide services on TaskRabbit and a decreasing value for buyers to find help. The elasticity of decisions to post tasks and submit offers with respect to the number of users in the market depends on the distribution of these costs. Buyers' elasticity is greater when more buyers are on the margin between posting and not posting tasks, and the same holds true for sellers. Measuring these elasticities is important for the platform's decision of which side – buyers or sellers – to attract at higher rates.

In addition to the number of buyers relative to sellers, users may also base their decisions on the overall market size. For example, if it becomes easier to find a match when the number of users increases, we would expect buyers to post more tasks and sellers to submit more offers. In a context where tasks are local and time sensitive, an increase in the number of buyers and sellers may lead to an increase in the availability of acceptable trading partners where and when they are most needed. Our data allows us to quantify whether doubling the number of users on the market more than doubles their propensity to post tasks and submit offers, and how these choices translate into matches.

First, we need to define the *market* within which buyers and sellers interact in the context of TaskRabbit. Given that 94% of users post or work in a single city, it is natural to treat cities as separate. The fact that 97% of successful tasks are matched to offers within 48 hours of posting suggests segmenting the data in time as well. As done in Section 3, we treat each city-month (e.g., San Francisco in October 2013) as a separate market. Within a city-month, we treat buyers and

¹¹Our conceptual framework is similar in approach to models of frictional search and matching in labor markets (Pissarides (2000)). In viewing price as being determined by the number of services requested and offered, we echo matching models in which wage is either a parameter in its own right or pinned down by other parameters. See, for example, Montgomery (1991) and Hall (2005).

sellers, as well as their tasks and offers, as homogeneous, which we discuss further in Appendix A2. This definition allows us to consider each participant as small relative to the size of the market, which is our modeling assumption, as well as to smooth smaller-scale time variation due to potential task heterogeneity.

4.1 Demand and Supply Elasticities

In this subsection we present our key finding of highly elastic supply. In order to quantify how buyers choose to post tasks and how sellers choose to submit offers as a function of the number of participants, we estimate OLS regressions of the following type:

$$\log(y_{ct}) = \alpha_1 \log(B_{ct}) + \alpha_2 \log(S_{ct}) + \eta_c + \eta_t + \nu_{ct}, \quad (1)$$

where c, t denote city c and year-month t , B_{ct} is the number of active buyers, and S_{ct} is the number of active sellers. The outcome is one of two relevant variables: tasks per buyer or offers per seller. The vector η_c controls for city-specific propensities to use TaskRabbit, which are time invariant. Similarly, η_t captures time-specific adjustments to usage intensities that are common across all cities. The identification of coefficients α_1 and α_2 is based on variation in the number of buyers and sellers within cities over time and within months across cities. We cluster standard errors at the city level.

The error term ν_{ct} represents a shock to the propensity to post tasks and submit offers. On the buyer side, it can be thought of as a city-month specific driver of demand for services among participating buyers. On the seller side, it can be interpreted as a city-month increase in time availability among participating sellers. For the OLS results to be interpreted as causal, we need to assume that these shocks are uncorrelated with users' decisions to join or stay on TaskRabbit. This is admittedly a strong assumption. For example, buyers and sellers might anticipate the future value of exchanges on the marketplace and base their decision to join, stay, or leave on these rational expectations. After presenting the OLS estimates, we demonstrate that our results do not change with instrumental variables and different market definitions, but also discuss the remaining threats to identification.

The first two columns of Table 3 show the OLS results for the number of tasks requested per buyer (Panel A) and for the number of offers submitted per seller (Panel B). The first column provides the regression results without fixed effects, and the second column with fixed effects. Adding fixed effects does not change the sign, size, or significance of the response of buyers or sellers to the number of active users, with the exception of buyers' response to the number of active

sellers. This supports the idea that TaskRabbit is used by buyers and sellers in a similar way over time and across cities.

The log-log specification means that the coefficients can be interpreted as elasticities. The estimates show that buyers' propensity to post tasks is not associated with the number of buyers or sellers active on the platforms. Yet, a doubling of the number of buyers is associated with a 54% increase in the number of offers submitted by each seller, and a doubling of the number of sellers is associated with a 29% decrease in the number of offers each seller submits.

The difference of the two coefficients $\alpha_1 - \alpha_2$ provides an estimate of users' response to the number of buyers *relative* to sellers, while the sum $\alpha_1 + \alpha_2$ is an estimate of users' response to market size.¹² The coefficient estimates in Panel A of table 3 imply that buyers' tasks are unrelated to changes in the number of buyers relative to sellers. Meanwhile, Panel B confirms that when buyers double relative to sellers, the number of offers submitted by each seller increases by 83%. In fact, we cannot reject the hypothesis that offers per seller double when the buyer to seller ratio doubles.¹³ As for market size, buyers' posting propensity is not significantly correlated with market size, while the number of submitted offers does seem to increase. We estimate the elasticity of offers per seller to market size to be 25%.

As anticipated, we cannot fully rule out the endogeneity of the number of buyers and sellers active on the platform. One possibility is that one side of the market might anticipate exogenous growth on the other side of the market. For example, if the weather forecast anticipates heavy snow in Chicago in January, more buyers will join the platform to request snow plowing. Since sellers also know about the weather forecast, they may join TaskRabbit in greater numbers in anticipation of higher demand.¹⁴ Platform policies can also facilitate seller response by advertising or loosening their screening process during these periods.

We can use instrumental variables to partially address these sources of endogeneity. To do so, we need to identify two instruments that shift the participation of buyers and sellers, respectively. The unemployment rate has been shown to affect the number of workers who join peer-to-peer platforms like TaskRabbit (Koustas (2019a) and Koustas (2019b)). Media coverage may also influence overall

¹²Given the log specification, we can transform the right-hand side of Equation 1 to be a function of the number of buyers relative to sellers and of overall market size. We can write $\log(y_{ct}) = \hat{\alpha}_1 \log\left(\frac{B_{ct}}{S_{ct}}\right) + \hat{\alpha}_2 \log(S_{ct}B_{ct}) + \eta_c + \eta_t + \nu_{ct}$, where $\hat{\alpha}_1 = \alpha_1 - \alpha_2$, $\hat{\alpha}_2 = \alpha_1 + \alpha_2$, $\frac{B_{ct}}{S_{ct}}$ is the buyer to seller ratio, and $S_{ct}B_{ct}$ is the product of the number of buyers and sellers. This transformation allows us to interpret the coefficient estimates in light of the particular nature of network externalities on TaskRabbit. A seller benefits from a market with relatively more buyers, where his services are highly demanded, but is hurt in a market with relatively more sellers, where his services face fierce competition. Yet when we hold the relative number of buyers and sellers constant, a seller may or may not benefit from a large market over a small one. A preference for larger markets can arise because of scale economies, while one for smaller markets may be related to congestion.

¹³The coefficient estimate of $\alpha_1 - \alpha_2$ cannot statistically be distinguished from 1.

¹⁴We thank an anonymous referee for providing this useful example of possible sources of endogeneity.

sign-ups. Indeed, after an initial marketing campaign at the time of entry, TaskRabbit advertising relied on articles mentioning the platform in newspapers and blogs. TaskRabbit was not directly involved in pitching some 40% of these articles, which independently referred to TaskRabbit while discussing the *sharing economy*.¹⁵ About 65% of these articles were in national media outlets, while the rest were published by local newspapers.¹⁶ We use a one-month lag for media articles and the concurrent unemployment rate, both interacted with city fixed effects for the largest cities, to instrument for the number of buyers and sellers. Media data come from Factiva and unemployment rates from the Bureau of Labor Statistics.¹⁷

We also use another set of instruments in the spirit of Bartik (1991). The idea behind this strategy is to use platform trends as instruments for local participation, removing local shocks that might affect both participation and usage intensity. Specifically, we instrument the number of buyers and sellers in city-month c, t with the average number of buyers and sellers in month t in cities other than c .

Instrumental variable estimates for Equation 1 are reported in the last two columns of Table 3. Column (3) uses Bartik instruments, while column (4) uses media articles and unemployment rates as instruments. The estimates are remarkably similar to the OLS results. We present the first stage estimates in Table 4. The coefficients confirm that both sets of instruments predict the number of buyers and sellers in the expected direction. The Bartik instrument for buyers (respectively sellers) positively affects the number of active buyers (respectively sellers). Unemployment rates positively affect seller participation in most cities and negatively affects buyer participation. Media coverage tends to increase participation of both user groups. The Kleibergen-Paap rk Wald F statistic, which tests whether instruments are weak for both our endogenous variables while adjusting for clustered standard errors, allows us to reject the null hypothesis of weak instruments (Stock and Yogo (2005)). In addition, we cannot reject the null that numbers of buyers and sellers are exogenous in the second-stage regressions.

The validity of the instruments hinges on the assumption that they shift users' decision of how many tasks or offers to post only by affecting the extensive margin of user participation. It is possible that an increase in unemployment affects not only the number of sellers active on the

¹⁵The sharing economy (or collaborative consumption) is a term often used to refer to online peer-to-peer marketplaces like Airbnb, Uber, or TaskRabbit. In the sharing economy, owners rent or share something they are not using (e.g., a car, house) or provide a service themselves to a stranger using peer-to-peer markets.

¹⁶These figures rely on TaskRabbit's tracking activity of its media presence in 2012 and on media articles mentioning TaskRabbit retrieved from Factiva. In particular, the share of articles that were not pitched by the company is based on media coverage that TaskRabbit tracked in 2012, while the share of articles in national media outlets comes from Factiva.

¹⁷Data on media coverage are available at <https://www.dowjones.com/products/factiva/>. Data on unemployment rates are publicly available at <https://beta.bls.gov/dataQuery/find>.

platform, but also how much time each seller can devote to selling services on TaskRabbit. Because we interact unemployment rate and media articles with city fixed effects, we have more instruments than endogenous variables, so we can test our overidentifying restrictions. The two Hansen J statistics, one for tasks and one for offers, fail to reject the null hypotheses that our instruments are valid and thus correctly excluded from the second stage. Appendix A2 presents a number of robustness checks using alternative market definitions and allowing for user heterogeneity.

We conclude this section by pointing out that despite the fact that our results are robust to a number of tests, we cannot fully rule out the possibility that shocks that are unobservable to the econometrician may be driving both extensive and intensive margins of user participation on the platform. The correlation patterns highlighted are, however, consistent with a very elastic labor supply, which has recently been replicated for Uber and Airbnb (Chen et al. (2019) and Farronato and Fradkin (2018)).

4.2 Returns to Scale in Matching

We now turn to an evaluation of the effect of buyer and seller behavior on the equilibrium number of matches and prices in order to quantify returns to scale in matching. To this end, we estimate regressions of the following type:

$$\log(y_{ct}) = \beta_1 \log b_{ct} + \beta_2 \log s_{ct} + \eta_c + \eta_t + \nu_{ct} , \quad (2)$$

where s_{ct} and b_{ct} are the total number of offers submitted and tasks requested in a city-month, and y_{ct} is either the number of transactions or the average transacted price. As in Section 4.1, we discuss OLS results, endogeneity issues, and IV estimates.

If we interpret the coefficients structurally, this specification is analogous to assuming that the total number of matches and average prices in the market are Cobb-Douglas functions of the number of tasks posted and offers submitted.¹⁸ The sum $\eta_c + \eta_t + \nu_{ct}$ is a market-level productivity shifter, where η_t is a month effect common across cities, η_c is a time-invariant city-specific parameter of match efficiency, and ν_{ct} is an idiosyncratic shock to matching.

In this specification, we expect that the number of matches will be increasing in both inputs, i.e. $\beta_1 \geq 0$ and $\beta_2 \geq 0$. The market exhibits increasing returns in matching if $\beta_1 + \beta_2 > 1$, such that doubling the number of tasks and offers more than doubles the number of matches. For pricing, we expect that more posted tasks will drive up prices and more offers will reduce them.

The first two columns of Panel A in Table 5 present OLS estimates of the number of matches.

¹⁸Petrongolo and Pissarides (2001) summarize the wide empirical support for a Cobb-Douglas matching function with constant returns to scale in labor markets. For its microfoundation, see Stevens (2007).

The results in column (2), once again show that including city fixed effects and month fixed effects does not affect our estimates of how tasks and offers contribute to matching. These estimates suggest that doubling the number of tasks while holding constant the number of offers increases the number of matches by 41%. Similarly, doubling the number of offers increases the number of matches by 52%. The estimates suggest that increases in either tasks or offers contribute about equally to the creation of successful matches.

The sum of the two elasticities provides an estimate of the returns to scale of the matching function. Work on two-sided platforms has emphasized the importance of increasing returns to scale for market structure (Parker and Van Alstyne (2005)). The hypothesis is that thick markets may lead to easier matching. Our estimates, however, show no evidence of increasing returns to scale. Returns are slightly (and significantly) less than constant when estimated by OLS.¹⁹

Our estimates of the matching parameters are robust to using instruments. We estimate Equation 2 using three separate sets of instruments, all of which are motivated by our analysis in Section 4.1. Since the total posted tasks and offers submitted are a direct function of the number of buyers and sellers, we can instrument for tasks and offers directly with the number of buyers and sellers. We can also employ the sets of instruments used in the previous subsection, namely Bartik-style instruments as well as media coverage and unemployment rate. Appendix Table A3 presents the first-stage regression results.²⁰ The results of the second stage are shown in the last three columns of Table 5 and confirm the OLS estimates. In addition, we cannot reject the null hypothesis that the number of tasks posted and offers submitted are exogenous. Table 5 does not report the city and time fixed effect estimates, which we will discuss in Section 5 when assessing the differences in the platform’s success across cities.²¹

Panel B of Table 5 reports estimates of the market pricing function. Including fixed effects changes both the magnitude and the significance of our price coefficients, suggesting that there is some degree of heterogeneity in the type of services provided across cities and months. We discuss this heterogeneity in Section 5. Instrumental variables do not affect the estimates. When fixed effects are included, prices move very little with the number of tasks and offers. Doubling the

¹⁹One might want to test directly whether increasing the number of buyers and sellers increases the number of matches more than linearly. This amounts to regressing the number of matches on B_{ct} and S_{ct} . Given that the propensity to use the platform is not greatly affected by market size, as shown in Section 4.1, the results do not change if we choose to include total buyers and sellers instead of total offers and tasks on the right-hand side. We confirm this in Appendix Table A17. Appendix Table A18 also confirms that entrants in large markets do not seem systematically different from entrants in small markets.

²⁰As before, the Cragg-Donald Wald F statistic and the Kleibergen-Paap rk Wald F statistic allow us to reject the null hypothesis of weak instruments. When using media coverage and unemployment rate as instruments, the two Hansen J statistics, one for matches and one for prices, do not allow us to reject the null that our instruments are valid.

²¹Estimates of city fixed effects are presented in Section 5, while those of time fixed effects are provided in Appendix A3.

number of tasks increases the average transacted price by 1.5%, while doubling offers decreases it by 1.3%. This may seem surprising from the standpoint of strategic pricing, especially for the auction tasks where buyers choose from competing offers, but it holds true even in a restricted sample of auctions (Appendix Table A6). In addition, the sum of the price elasticity to tasks and offers is virtually zero, suggesting that the composition of tasks does not change with market size.

As in the previous subsection, these results should be interpreted as associations given possibilities of endogeneity bias, although it is reassuring that our findings are robust to various market definitions and econometric approaches. Before concluding this section, we discuss the potential consequences of high supply elasticity, fixed prices, and constant returns in matching for the functioning of the platform.

The fact that prices do not seem to depend on fluctuations in buyers and sellers on the platform is consistent with the high level of competition in which TaskRabbit operates. Recent studies of Uber found similar results, namely that hourly wages are not dependent on platforms’ pricing decisions (Hall et al. (2019)). TaskRabbit was relatively small during the period of our study, and the tasks requested do not require high skills, so it is reasonable that the price for services exchanged on TaskRabbit be set by the external environment. Data from Crunchbase, a comprehensive database of startups and larger companies, suggests that TaskRabbit faces many competitors. A search for “Home Services” returns almost 250 companies, ranging from similar online platforms such as Thumbtack to publicly listed companies such as HomeServe.²² In addition, similar services could likely be exchanged through informal connections and word of mouth.

However, the fact that prices are fixed means that there are a couple of ways in which the market could equilibrate in response to demand increases. One option is adjustments in seller effort, the other is buyer rationing. If sellers did not expand effort, doubling the number of task requests would not double the number of matches unless the platform contemporaneously recruited new sellers. Our paper provides the first evidence of a high supply elasticity in online peer-to-peer markets. Subsequent work has confirmed that it is a common feature across many other platforms for local services, such as Uber (Chen et al. (2019)) and Airbnb (Farronato and Fradkin (2018)). In order to evaluate the benefits of this elastic labor supply for marketplace growth, we compare it to the alternative where supply is fixed.

We present an illustration of this intuition through a simple exercise. Consider a market with 1,000 posted tasks, and suppose that the number of offers submitted is 1,400. Using our estimates for the matching function, 488 matches would be created out of these two aggregate inputs. Now assume that demand doubles to 2,000 posted tasks. A perfectly elastic supply would lead to a

²²Data were accessed from <https://www.crunchbase.com/search/organization.companies> in December 2017.

doubling of the number of offers, and would create 930 total matches.²³ Analogously, if demand halved to 500 tasks and supply adjusted downward to 700 offers, the number of matches created would be 256. Regardless of the amount of demand, tasks would always match at the same rate.

In the alternative scenario, supply is held fixed at 1,400 offers and equilibration occurs through buyers’ rationing. It is easier for buyers to find a match when demand is low, and when demand is high it becomes harder to trade. In the low-demand market (500 tasks), the number of matches created would be 367 and the probability of matching a given task would be 73%. In the high-demand market (2,000 tasks), 649 tasks would be matched, implying a 32% match rate. Overall, if we compare the total number of matches between the two scenarios, the marketplace with an elastic labor supply is able to create 11% more matches.²⁴ This is evidently optimal from the platform’s perspective: since its revenues are a 20% commission on actual matches and prices barely move, in this simple example having an elastic supply raises its short-term revenue by 11%. Since it also raises retention, which we show in Section 5, it benefits the marketplace by accelerating growth.

The absence of increasing returns in the matching technology is perhaps unexpected. On a platform like TaskRabbit, in which tasks typically require a buyer and a seller to meet, efficiency can come from matching buyers and sellers who live close to each other. Interestingly, however, the distances between buyers and sellers do not appear to shrink as TaskRabbit grows in a city. Appendix Figure A2 provides more details and confirms that distances between buyers and sellers do not shrink as a market scales up.

The tendency of sellers to flexibly expand effort in response to increases in demand implies that peer-to-peer markets should focus on attracting buyers at higher rates than sellers, a strategy that we discuss further in the conclusion. We now turn to the extensive margins of user participation.

5 Long-Run Growth and City Heterogeneity

A noticeable feature of our data is that some cities exhibit striking growth in the number of matches while others exhibit more moderate growth (Figure 1). Given how flexibly sellers adjust their effort, we focus on buyers’ propensity to join and continue using TaskRabbit.

In principle, two types of theories can explain the differences in how cities attract and retain buyers. The first relies on scale economies and strategic complementarities between the adoption patterns of buyers and sellers. Under this hypothesis, we would expect cities that start off with a large user base to grow much faster than those with a smaller number of initial users, as growth

²³Note that doubling the number of tasks and offers does not double the number of matches because of the slightly decreasing returns to scale estimated for the matching function (shown in Table 5).

²⁴The result comes from the fact that the matching function is concave in both inputs, so that $M(b', s') + M(b, s) > M(b', s) + M(b, s')$ where $b' > b$ and $s' > s$.

leads to more growth. Yet as we find no evidence for scale economies in matching in Section 4, it would seem that initial differences in adoption cannot explain increasing heterogeneity over time.

A second set of hypotheses rely on city differences in facilitating interactions between buyers and sellers. To develop this idea, we show that user attrition is lower in cities where matches are created at higher rates, and that cities vary greatly in the efficiency with which tasks and offers are converted into matches. We conclude the section with evidence that relates match efficiency with measures of market thickness at the city level: geographic distance between buyers and sellers, and task standardization.

5.1 City Differences in Long-Run Growth

User growth is a combination of adoption and retention of existing users. Given the high elasticity of supply, growth depends on buyers' decisions to use the platform. The literature on innovation diffusion (Young (2009)) highlights three mechanisms that lead users to adopt new technologies: network effects, technology improvements, and information diffusion. The first two assume that different users adopt at varying points in time because of heterogeneous benefits: early adopters find significant intrinsic value from a new technology, while late adopters join because of scale economies or technical upgrades. We have argued, however, that marketplace efficiency does not increase with market scale. Furthermore, TaskRabbit did not implement major platform changes during the study period. Word of mouth and information diffusion, then, seem to be the most plausible alternatives in this context, and cities can differ both in the rate at which information spreads and their uptake rates once residents receive that information. For example, adoption might be fast in San Francisco because people there are eager to experiment with new technologies and because current users spread the information at a quicker rate. The latter factor may also be driven by positive experiences on TaskRabbit. We measure the aggregate effect by estimating the city-specific buyer adoption equation:²⁵

$$new_{ct} = \phi_c age_{ct} + \epsilon_{ct} , \quad (3)$$

where new_{ct} is the number of new buyers joining city c in calendar month t , and age_{ct} is the age of the marketplace in city c at time t . For example, $age_{ct} = 1$ if month t is the first since TaskRabbit

²⁵We assume a linear growth rate different across cities, given Appendix Figure A9. This can be rationalized within the Bass model of new product diffusion (Bass (1969)): $new_{ct} = \phi_c + new_{c,t-1}$, where new_{ct} is the number of new buyers joining city c in calendar month t . Our case differs from the standard specification in two ways. First, the total number of potential adopters is assumed to be large relative to the size of the marketplace, which is consistent with the population size of the metropolitan cities relative to the current TaskRabbit user base. Second, we assume that new adopters in the previous month are the only users spreading information.

became active in city c . A buyer is defined as new in a city-month if she posts her first task in that city during that particular month. Estimates are plotted in Figure 3, which we discuss after estimating retention parameters.²⁶

We now compare adoption rates with retention, which can be city-specific and, within each city, further depend on current outcomes, match rates and prices:

$$\log \left(\frac{stay_{ct}}{1 - stay_{ct}} \right) = \theta_0 X_{ct} + \theta_t + \theta_c + \epsilon_{ct} . \quad (4)$$

The variable $stay_{ct}$ is the share of buyers active in city-month c, t who were active again at least once in the following three months within the same city. X_{ct} includes realized buyer match rate and average transacted price.²⁷ In this regression, we are particularly concerned with city differences in user retention, θ_c .

Figure 3 plots the estimates of ϕ_c from Equation 3 and θ_c from Equation 4. A certain correlation exists between the rate at which a city is able to attract buyers and the rate at which it can retain them, although it is by no means perfect. San Francisco and New York are successful in both measures, while Houston, Atlanta and Phoenix lag behind on both. However, in San Diego, buyers join at a fast rate but are also likely to leave the marketplace, while Portland has a smaller number of new buyers, but they stay longer. Retention is arguably the decision that is most directly related to users' experience on the platform, and indeed in the next section we show that it is associated with how efficiently tasks and offers are converted into successful matches in each city.

5.2 City Differences in Match Efficiency and Market Thickness

For such mundane tasks as cleaning and delivery, it seems reasonable to expect that buyers would care about the ease with which they can find a seller willing to provide the service at the desired time and location. In this section, we explore how simple it is to find a match across different cities. To do so, we use earlier estimates of the matching function from Section 4.2.

Cities vary widely in the rate at which tasks and offers are converted into successful matches. Figure 4 plots estimates of η_c from Equation 2, ordering cities from most (San Francisco) to least efficient (Miami). San Francisco is 2.37 ($\frac{\eta_{SF}}{\eta_{Miami}}$) times as effective as Miami in creating matches: for every 100 matches created in Miami from a given number of tasks and offers, 237 matches are

²⁶Standard errors are reported in Appendix A4. In the appendix, we also verify that deviations from the linear adoption rate are not driven by contemporaneous market conditions, which supports our estimation strategy in Section 4.

²⁷We expect that a high match rate increases the odds that a buyer will be active again in the next three months, while a high price would drive away more buyers. Appendix A4 presents the estimates, which confirm our hypothesis. In the appendix, we also verify that retention is not correlated with future outcomes, which supports our estimation strategy in Section 4.

created in San Francisco. Other than San Francisco, the cities with the highest match productivity are Boston, Portland, Austin, and New York, with Miami, Denver, Phoenix, Philadelphia, and Atlanta at the opposite extreme.

The most efficient cities are those able to retain the most buyers.²⁸ In Figure 4 the size of the marker denoting the match efficiency parameter is proportional to the retention rate θ_c estimated from Equation 4. The cities with the highest match productivity η_c also have high retention rate θ_c .

The next step is to try to understand what explains efficiency differences across geographies. To this end, we examine two metrics related to market thickness. The first natural candidate is the proximity of buyers and sellers. Cities that more easily match tasks and offers might be those where buyers and sellers live closer together and can more easily meet and exchange services. The data provides strong support for this hypothesis.²⁹ Figure 5 plots the match efficiency parameter η_c and the median geographic distance between buyers and sellers of paired tasks and offers.³⁰ In cities like San Francisco, Boston, Portland, and New York, which have the highest rate of matches, the median distance between buyers and sellers of paired tasks and offers is around seven miles. At the other extreme, this distance in Philadelphia and San Diego is over 20 miles.

The second candidate measure of market thickness is related to task specificity. The hypothesis is that an idiosyncratic task, which might require specialized skills on the part of the seller, is harder to match than a standard task, for which a relatively unskilled seller’s location and availability are all that matter. To explore this idea, we look at the share of tasks posted in May 2014 within the top five categories, which include shopping, moving help, and cleaning. Figure 6 plots this share of more standard tasks against the match efficiency parameter η_c . Over 60% of the tasks posted in San Francisco, Boston, and New York fall within the top five categories, while such tasks are less than half of the total in Dallas and Atlanta.

All the correlations displayed in the plots are highly significant (Appendix Table A2). Even so, care should be taken in interpreting these correlations as causal given the small number of cities for which we have data. It is also possible that more experienced buyers learn to post tasks better, i.e., in homogeneous categories, which makes matching more efficient. Another possibility is that experienced buyers and sellers learn to transact with nearby users, but that the distance between buyers and sellers on TaskRabbit is unrelated to the distance between buyers and sellers of similar services outside of the platform. However, it is important to highlight that the differences

²⁸The most efficient cities also tend to be those where TaskRabbit entered earlier.

²⁹In Appendix A1 we run task-level regressions to verify that buyer-seller distance matters for match success.

³⁰The correlation persists with two other pairing definitions: median distance between buyers and sellers active around the same time window and median distance between successfully matched buyers and sellers.

between cities in terms of the distance between buyers and sellers emerges as soon as TaskRabbit becomes active in a city (Appendix Figure A2), reducing the likelihood of the learning hypothesis. It also seems that the growth of TaskRabbit in a city is correlated with the size of Craigslist, a competing platform.³¹ This correlation supports the possibility that the way buyers and sellers use TaskRabbit in a given city is common across similar online platforms.

6 Managerial Implications and Conclusions

In this paper, we study how supply elasticity, economies of scale, and city-specific characteristics related to market thickness contribute to the growth of online peer-to-peer markets. This is an important question for marketplaces such as Uber, Airbnb, and TaskRabbit that have recently gained in popularity. In our analysis, we separate users' decision of how intensely to use the marketplace in the short run from the decision to join or leave the marketplace, both of which affect growth. We explore short- and long-run participation decisions on TaskRabbit, a growing marketplace for domestic tasks by examining changes in market conditions across cities and over time.

Our analysis unveils three main findings. First, buyers post similar numbers of tasks regardless of the relative abundance of sellers, while sellers expand their effort when they are scarce and contract it when they are abundant. This adjustment on the seller side has virtually no impact on average prices nor on the probability that tasks are filled, suggesting a highly elastic labor supply. Second, the number of matches scales linearly with the number of market participants, suggesting constant rather than increasing returns to scale in matching. Third, match efficiency is positively correlated with the geographical proximity of buyers and sellers and with the level of task standardization, which helps to explain differences in platform success across cities.

Our findings have four important implications for managing the growth of peer-to-peer marketplaces. First, the managerial challenge unique to two-sided platforms is striking a balance between demand and supply. This is the classical chicken-and-egg problem: limited demand discourages sellers from joining, while limited supply discourages buyers (Caillaud and Jullien (2003) and Hagiu (2007)). Our results suggest that managerial efforts should focus on attracting the most inelastic side of the market, which constrains the volume of trades on the platform, and thus profits. On TaskRabbit, we find that demand is more inelastic, such that the platform should spend resources to attract buyers at a faster rate than sellers. Initially, TaskRabbit seems to have devoted more resources to recruiting and screening sellers than to attracting buyers, as evidenced by their tight

³¹The correlation plot is presented in Appendix Figure A3. We thank Robert Seamans and Feng Zhu for sharing Craigslist data from their research (Seamans and Zhu (2014)).

background checks. The fact that the platform changed their seller screening policies towards easier background checks in mid-2013 suggests a shift of focus away from sellers, in line with our findings. It is possible, however, that on other platforms the most inelastic side may be sellers rather than buyers. Our framework can be applied to other data to identify how responsive each side is to market conditions.

Second, our results have implications for entry and competition. The standard assumption is that two-sided platforms benefit from large network effects, implying that multiple platforms cannot coexist in equilibrium and a single winner will tend to dominate the market (Eisenmann et al. (2011)). The challenge then is to achieve scale quickly and deter competitive entry. When considering network effects in online platforms, it is important to distinguish between the scope of network effects (Zhu et al. (2019)), which depends on the network structure of exchanges among users, and the scale at which network effects manifest. In the context of peer-to-peer platforms for local and time-sensitive services, the scope of network effects is limited to a specific geography. This implies that entrants can successfully challenge incumbents at the local level without having to quickly expand geographically to reap the network benefits. For example, the ride-sharing apps Via and Juno have successfully entered and competed with leaders Uber and Lyft in a few local markets.

As to the scale at which network effects manifest, in this context we find that the latter are exhausted very quickly, and increasing the number of participants only proportionally increases the number of matches. This result implies that multiple such platforms can coexist in a given city without compromising the efficiency of matching. Thus, local entry and competition are much more likely here than in other markets, such as those for remote and specialized work, where network effects are not local in scope and may not be exhausted as quickly. Indeed, a simple count of active companies for “Home Services” on Crunchbase returns 249 firms, against just 67 “Freelance” companies.³² Combined with our finding of highly elastic labor supply, this result also implies that we should expect stronger price competition to attract buyers rather than sellers.

Third, our exploration of city heterogeneity can help guide platform expansion strategies. The results suggest that local characteristics are more important for growth than economies of scale. As a consequence, managerial efforts should be targeted at strategically picking cities whose characteristics promote market thickness. In the case of TaskRabbit, and of marketplaces for personal services more broadly, one relevant characteristic is the distance between the locations of buyers and sellers.

Fourth, our results shed some light on the potential benefits of economies of scale versus scope

³²On Crunchbase, platforms similar to TaskRabbit are classified as “Home Services,” while those akin to Upwork are categorized as “Freelance.” The database was accessed in December 2017.

in platform businesses. When scale does not increase the value that each user obtains from a platform, other strategies may help increase user surplus, especially those based on expanding scope. The strategic choice to diversify the service offering so as to attract and retain users is of primary importance and, based on our findings, may depend on the particular stage of platform development. Here we find that the level of service standardization is positively associated with matching efficiency because it guarantees a fairly undifferentiated supply of sellers. This ensures that task requests are more easily filled, which in turn increases user retention. This result suggests that managerial efforts should focus on choosing a few standard services to be exchanged on the marketplace, at least early on. Many other peer-to-peer marketplaces did in fact begin with narrow offerings – e.g., rides on Uber, accommodations on Airbnb – and later expanded to complementary services categories – Uber to food delivery and Airbnb to local experiences.³³ When and how to diversify a platform’s services is an important and open question for future empirical research.

A slightly less direct implication of our findings is that when services are relatively undifferentiated and buyers’ tastes are not highly heterogeneous, peer-to-peer markets can design the search and matching mechanism to reduce search costs. Given a platform design where sellers need to submit offers to each individual task request, it seems costly for sellers to search through posted tasks. In so doing, sellers potentially miss out on profitable tasks because of availability constraints or simply not finding such tasks in time. Indeed, 13% of tasks that did not receive any offers were canceled because the buyer indicated the task was already done, while 19% of tasks that did receive offers were canceled because of an inability to resolve scheduling conflicts. The benefits of reducing sellers’ search costs and improving match efficiency provide a rationale for re-designing the marketplace, which in fact occurred in the spring of 2014. Buyers can now select the category, location, and time for a given task request, and then either choose from among the sellers available to perform the task at the specified time and location, or have TaskRabbit automatically choose for them. However, the benefits of reducing search frictions and increasing match efficiency may come at a cost. Specifically, listing time availability may decrease the flexibility of seller schedules, which recent research shows to be valuable for service providers in peer-to-peer markets (Chen et al. (2019)). Automatic matching might also reduce sellers’ opportunity to trade more specialized skills at a premium. These examples highlight the fact that the two sides of a peer-to-peer market have conflicting interests, such that design choices benefiting one side might hurt the other.

Our paper primarily focuses on aggregate measures of network effects, supply elasticity, and city differences as contributors to the growth of peer-to-peer markets, leaving aside buyer and seller heterogeneity. A valuable direction for future research would be to study the dynamic participation

³³www.ubereats.com and www.airbnb.com/s/experiences, accessed in December 2017.

decisions of buyers and sellers in greater detail. A better understanding of the drivers of user adoption, including multi-homing, could help to explain the competition between peer-to-peer markets as well as that between the online marketplace model and more traditional service providers.

References

- Allon, Gad, Achal Bassamboo, and Eren B. Çil.** 2012. “Large-Scale Service Marketplaces: The Role of the Moderating Firm.” *Management Science*, 58(10): 1854–1872.
- Armstrong, Mark.** 2006. “Competition in two-sided markets.” *The RAND Journal of Economics*, 37(3): 668–691.
- Arnosti, Nick, Ramesh Johari, and Yash Kanoria.** 2019. “Managing Congestion in Matching Markets.” *Working Paper*.
- Arthur, W Brian.** 1989. “Competing Technologies, Increasing Returns, and Lock-In by Historical Events.” *The Economic Journal*, 99(394): 116–131.
- Bartik, Timothy J.** 1991. *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute for Employment Research.
- Bass, Frank M.** 1969. “A New Product Growth for Model Consumer Durables.” *Management Science*, 15(5): 215–227.
- Bohlmann, Jonathan D, Peter N Golder, and Debanjan Mitra.** 2002. “Deconstructing the pioneer’s advantage: Examining vintage effects and consumer valuations of quality and variety.” *Management Science*, 48(9): 1175–1195.
- Boudreau, Kevin J, and Andrei Hagiu.** 2009. “Platform rules: Multi-sided platforms as regulators.” *Platforms, markets and innovation*, 1: 163–191.
- Bresnahan, Timothy F, and Shane Greenstein.** 1999. “Technological competition and the structure of the computer industry.” *The Journal of Industrial Economics*, 47(1): 1–40.
- Cachon, Gerard P, Kaitlin M Daniels, and Ruben Lobel.** 2017. “The role of surge pricing on a service platform with self-scheduling capacity.” *Manufacturing & Service Operations Management*, 19(3): 368–384.
- Caillaud, Bernard, and Bruno Jullien.** 2003. “Chicken-and-egg: competition among intermediation service providers.” *The RAND Journal of Economics*, 34(2): 309–328.
- Casadesus-Masanell, Ramon, and Hanna Halaburda.** 2014. “When Does a Platform Create Value by Limiting Choice?” *Journal of Economics and Management Strategy*, 23(2): 259–293.
- Cennamo, Carmelo, and Juan Santalo.** 2013. “Platform competition: Strategic trade-offs in platform markets.” *Strategic Management Journal*, 34(11): 1331–1350.


- Chen, M. Keith, Judith A. Chevalier, Peter E. Rossi, and Emily Oehlsen.** 2019. “The Value of Flexible Work: Evidence from Uber Drivers.” *Journal of Political Economy*, 127(6): 2735–2794.
- Dellarocas, Chrysanthos.** 2003. “The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms.” *Management Science*, 49(10): 1407–1424.
- Dellarocas, Chrysanthos, and Charles A Wood.** 2008. “The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias.” *Management Science*, 54(3): 460–476.
- Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan.** 2018. “Consumer Price Search and Platform Design in Internet Commerce.” *American Economic Review*, 108(7): 1820–59.
- Economides, Nicholas, and Evangelos Katsamakas.** 2006. “Two-sided competition of proprietary vs. open source technology platforms and the implications for the software industry.” *Management Science*, 52(7): 1057–1071.
- Einav, Liran, Chiara Farronato, Jonathan Levin, and Neel Sundaresan.** 2018. “Auctions versus posted prices in online markets.” *Journal of Political Economy*, 126(1): 178–215.
- Eisenmann, Thomas, Geoffrey Parker, and Marshall W Van Alstyne.** 2006. “Strategies for two-sided markets.” *Harvard business review*, 84(10).
- Eisenmann, Thomas R, Geoffrey Parker, and Marshall W Van Alstyne.** 2011. “Platform Envelopment.” *Strategic Management Journal*, 32(12): 1270–1285.
- Ellison, Glenn, and Drew Fudenberg.** 2003. “Knife-edge or plateau: When do market models tip?” *The Quarterly Journal of Economics*, 118(4): 1249–1278.
- Evans, David S.** 2003. “The Antitrust Economics of Multi-Sided Platform Markets.” *Yale Journal on Regulation*, 20(2): 4.
- Farronato, Chiara, and Andrey Fradkin.** 2018. “The Welfare Effects of Peer Entry in the Accommodation Market: The Case of Airbnb.” *NBER Working Paper No. w24361*.
- Farronato, Chiara, Jessica Fong, and Andrey Fradkin.** 2020. “Dog Eat Dog: Measuring Returns to Scale Using a Digital Platform Merger.” *Working Paper*.

- Fradkin, Andrey.** 2019. “Search Frictions and the Design of Online Marketplaces.” *Working Paper*.
- Gawer, Annabelle, and Michael A. Cusumano.** 2014. “Industry Platforms and Ecosystem Innovation.” *Journal of Product Innovation Management*, 31(3): 417–433.
- Greenstein, Shane, and Tarun Khanna.** 1997. “What does industry convergence mean.” In *Competing in the age of digital convergence*. Harvard Business School Press.
- Griliches, Zvi.** 1957. “Hybrid Corn: An Exploration in the Economics of Technological Change.” *Econometrica*, 25(4): 501–522.
- Hagiu, Andrei.** 2007. “Merchant or Two-Sided Platform?” *Review of Network Economics*, 6(2): 115–133.
- Hagiu, Andrei.** 2009. “Two-Sided Platforms: Product Variety and Pricing Structures.” *Journal of Economics & Management Strategy*, 18(4): 1011–1043.
- Hagiu, Andrei, and Julian Wright.** 2015. “Marketplace or Reseller?” *Management Science*, 61(1): 184–203.
- Hall, Jonathan V., and Alan B. Krueger.** 2018. “An Analysis of the Labor Market for Uber’s Driver-Partners in the United States.” *ILR Review*, 71(3): 705–732.
- Hall, Jonathan V, John J Horton, and Daniel T Knoepfle.** 2019. “Labor Market Equilibrium: Evidence from Uber.” *Working Paper*.
- Hall, Robert E.** 2005. “Employment Fluctuations with Equilibrium Wage Stickiness.” *American Economic Review*, 95(1): 50–65.
- Horton, John J.** 2019. “Supply Constraints as a Market Friction: Evidence from an Online Labor Market.” *Working Paper*.
- Huang, Peng, Marco Ceccagnoli, Chris Forman, and D J Wu.** 2013. “Appropriability mechanisms and the platform partnership decision: Evidence from enterprise software.” *Management Science*, 59(1): 102–121.
- Hui, Xiang, Maryam Saeedi, Zeqian Shen, and Neel Sundaresan.** 2016. “Reputation and Regulations: Evidence from eBay.” *Management Science*, 62(12): 3604–3616.
- Jiang, Baojun, Kinshuk Jerath, and Kannan Srinivasan.** 2011. “Firm strategies in the “mid tail” of platform-based retailing.” *Marketing Science*, 30(5): 757–775.

- Katz, Michael L, and Carl Shapiro.** 1985. "Network externalities, competition, and compatibility." *The American Economic Review*, 75(3): 424–440.
- Koustas, Dmitri K.** 2019a. "Consumption Insurance and Multiple Jobs: Evidence from Rideshare Drivers." *Working Paper*.
- Koustas, Dmitri K.** 2019b. "What Do Big Data Tell Us about Why People Take Gig Economy Jobs?" *AEA Papers and Proceedings*, 109: 367–371.
- Lieberman, Marvin B, and David B Montgomery.** 1988. "First-mover advantages." *Strategic Management Journal*, 9(S1): 41–58.
- Li, Jun, and Serguei Netessine.** 2020. "Higher market thickness reduces matching rate in online platforms: Evidence from a quasiexperiment." *Management Science*, 66(1): 271–289.
- Li, Jun, Antonio Moreno, and Dennis J Zhang.** 2015. "Agent behavior in the sharing economy: Evidence from Airbnb." *Ross School of Business Working Paper Series*, 1298: 2015.
- Luca, Michael, and Georgios Zervas.** 2016. "Fake it till you make it: Reputation, competition, and Yelp review fraud." *Management Science*, 62(12): 3412–3427.
- Montgomery, James D.** 1991. "Equilibrium wage dispersion and interindustry wage differentials." *The Quarterly Journal of Economics*, 106(1): 163–179.
- Parker, Geoffrey G, and Marshall W Van Alstyne.** 2005. "Two-sided network effects: A theory of information product design." *Management Science*, 51(10): 1494–1504.
- Petrongolo, Barbara, and Christopher A Pissarides.** 2001. "Looking into the Black Box: A Survey of the Matching Function." *Journal of Economic Literature*, 39(2): 390–431.
- Pissarides, Christopher A.** 2000. *Equilibrium Unemployment Theory*. The MIT Press.
- Resnick, Paul, and Richard Zeckhauser.** 2002. "Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system." *The Economics of the Internet and E-commerce*, 127–157.
- Rietveld, Joost, Melissa A Scholling, and Cristiano Bellavitis.** 2019. "Reload and Relaunch: Strategic Governance of Platform Ecosystems." *Working Paper*.
- Rochet, Jean-Charles, and Jean Tirole.** 2003. "Platform competition in two-sided markets." *Journal of the European Economic Association*, 1(4): 990–1029.


- Rysman, Marc.** 2009. “The economics of two-sided markets.” *Journal of Economic Perspectives*, 23(3): 125–43.
- Seamans, Robert, and Feng Zhu.** 2014. “Responses to entry in multi-sided markets: The impact of Craigslist on local newspapers.” *Management Science*, 60(2): 476–493.
- Stevens, Margaret.** 2007. “New Microfoundations for the Aggregate Matching Function.” *International Economics Review*, 48(3): 847–868.
- Stock, James, and Motohiro Yogo.** 2005. “Testing for Weak Instruments in Linear IV Regression.” In *Identification and Inference for Econometric Models*. Cambridge University Press.
- Tucker, Catherine.** 2008. “Identifying formal and informal influence in technology adoption with network externalities.” *Management Science*, 54(12): 2024–2038.
- Vincenty, Thaddeus.** 1975. “Direct and inverse solutions of geodesics on the ellipsoid with application of nested equations.” *Survey review*, 23(176): 88–93.
- Weyl, E Glen.** 2010. “A price theory of multi-sided platforms.” *The American Economic Review*, 100(4): 1642–1672.
- Young, H Peyton.** 2009. “Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning.” *The American economic review*, 99(5): 1899–1924.
- Zervas, Georgios, Davide Proserpio, and John W. Byers.** 2017. “The Rise of the Sharing Economy: Estimating the Impact of Airbnb on the Hotel Industry.” *Journal of Marketing Research*, 54(5): 687–705.
- Zhu, Feng, and Marco Iansiti.** 2012. “Entry into platform-based markets.” *Strategic Management Journal*, 33(1): 88–106.
- Zhu, Feng, Xinxin Li, Ehsan Valavi, and Marco Iansiti.** 2019. “Network Interconnectivity and Entry into Platform Markets.” *Harvard Business School Working Paper*.

Figure 1: *Market Size.*



This figure plots the number of successful matches in selected cities over time. The y-axis, in log scale, is normalized by the value in San Francisco in May 2014. The picture shows the large heterogeneity in the growth of TaskRabbit across cities.

Figure 2: *Buyer to Seller Ratio.*




This figure plots the active number of buyers relative to sellers in selected cities over time. A buyer is active in a city-month if she posts at least one task within the city-month. A seller is active if he submits an offer to one of the tasks posted within that city-month. The figure shows the large fluctuations in the buyer to seller ratio across cities and over time.

Figure 3: *Buyer Adoption and Retention by City.*




This figure plots the estimates of ϕ_c (city-specific adoption rate) from Equation 3, and θ_c (retention rate) from Equation 4. For Denver and Miami, there are too few months of recorded activity to estimate θ_c . The picture shows that cities in which adoption occurs quickly do not always have high retention rates. Figures 5 and 6 relate retention to matching efficiency on the platform.

Figure 4: *City Heterogeneity in Match Efficiency.*




This figure shows η_c from the OLS regression of Equation 2 where the outcome variable is the log number of matches in a city-month. The marker size is proportional to the retention rate parameter θ_c from the estimation of Equation 4 (there are too few months for Denver and Miami to estimate θ_c). Vertical bars denote 95% confidence intervals. The picture shows the positive relationship between the retention rate (marker size) and matching efficiency (y-axis).

Figure 5: *City Heterogeneity in Geographic Distance Between Buyers and Sellers.*



This figure plots the median distance between buyers and sellers in a city against the estimated match productivity parameter η_c from Equation 2 for that city. The distance between the buyers and sellers is computed as the length of the shortest ellipsoidal curve between the zip codes of paired buyers and sellers (Vincenty (1975)). A buyer is paired with a seller if her task received an offer from that seller. Each pair is weighted by the number of their task-offer pairs within a city, and the median is computed among all such pairings at the city level. The size of each bubble is proportional to the retention rate parameter θ_c from the estimation of Equation 4. The correlation with overall marketplace growth by city is also positive. The full correlation matrix is displayed in Appendix Table A2.

Figure 6: *City Heterogeneity in Task Standardization.*



This figure is identical to Figure 5 except that on the x-axis we plot the share of tasks posted in May 2014 within the top five task categories. Alternative measures of standardization (other time frames, top three categories) display similar patterns. The correlation with the overall marketplace growth by city is also positive. The full correlation matrix is displayed in Appendix Table A2.

Table 1: *Summary Statistics.*

	N	Mean	Standard Deviation	25th Percentile	Median	75th Percentile
Share of Auction Tasks	459,879	0.59	0.49	0	1	1
Share Receiving Offers	459,879	0.78	0.41	1	1	1
Nr. Offers Received (if > 0)	358,557	2.82	5.1	1	1	3
Share Matched	459,879	0.49	0.50	0	0	1
Price of Successful Tasks (\$)	224,877	57	44.24	25	45	75
Commission Fee (%)	224,877	0.19	0.04	0.18	0.2	0.2

(a) *Task level summary statistics.*

	N	Mean	Standard Deviation	25th Percentile	Median	75th Percentile
Number of Active Buyers	336	708	1022	132	272	738
Number of Active Sellers	336	255	326	67	124	277
Buyer to Seller Ratio	336	2.52	0.96	1.87	2.36	3
Number of Tasks per Buyer	336	1.63	0.22	1.49	1.62	1.76
Number of Offers per Seller	336	6.45	3.04	4.22	5.62	7.59
Task to Offer Ratio	336	0.70	0.30	0.53	0.64	0.79
Task Match Rate	336	0.46	0.11	0.41	0.48	0.53
Average Transacted Price (\$)	336	56	8.69	52	57	61

(b) *City-month level summary statistics.*

Summary statistics. Data include fixed price and auction tasks, offers submitted to those tasks, and matches created in the 18 US cities between June 2010 and May 2014. In the first panel, an observation is a posted task. In the bottom panel, an observation is a city-month. We define a buyer to be active in a city-month if she posts at least one task in that city-month. Analogously, a seller is active if he submits an offer to a task posted within the city-month.

Table 2: *Equilibrium Outcomes Across Markets.*

Buyer to Seller Ratio	Tasks per Buyer	Offers per Seller	Task Match Rate	Average Transacted Price (\$)
1.49 (0.27)	1.63 (0.17)	4.43 (1.66)	0.45 (0.1)	55.93 (8.29)
2.13 (0.13)	1.6 (0.2)	5.66 (2.23)	0.5 (0.12)	52.12 (9.26)
2.64 (0.18)	1.63 (0.18)	6.56 (2.51)	0.46 (0.11)	56.2 (8.42)
3.84 (0.76)	1.68 (0.3)	9.15 (3.34)	0.44 (0.12)	59.67 (7.12)

This table shows average outcomes in four market groups. We divide the 336 city-months into four groups of 84 observations each, corresponding to the quartiles of the distribution of the buyer to seller ratio. The average number of buyers posting tasks in a city-month for every seller submitting offers is displayed in the first column. The other four columns compute, respectively, the average number of tasks per buyer, offers per seller, match rates, and average transacted prices. Standard deviations are shown in parenthesis.

Table 3: *Tasks per Buyer and Offers per Seller.*

Panel A: Tasks per Buyer (log)				
	(1)	(2)	(3)	(4)
Active Buyers (log)	-0.007 [0.024]	0.025 [0.030]	0.011 [0.045]	0.052 [0.067]
Active Sellers (log)	0.062 [0.031]*	0.025 [0.036]	0.048 [0.045]	0.002 [0.081]
Constant (SF Oct '13)	0.220 [0.041]***	0.205 [0.224]	0.150 [0.192]	0.152 [0.168]
R-squared	0.222	0.604	0.603	0.613
Panel B: Offers per Seller (log)				
Active Buyers (log)	0.561 [0.041]***	0.540 [0.043]***	0.599 [0.070]***	0.648 [0.095]***
Active Sellers (log)	-0.247 [0.031]***	-0.292 [0.049]***	-0.349 [0.061]***	-0.420 [0.113]***
Constant (SF Oct '13)	-0.277 [0.145]*	0.111 [0.247]	0.041 [0.341]	0.155 [0.181]
R-squared	0.869	0.939	0.939	0.938
City FE	No	Yes	Yes	Yes
Month FE	No	Yes	Yes	Yes
Instruments	No	No	Bartik	Media & Unempl.
Observations	336	336	336	319

*This table shows results from the OLS and IV regressions of Equation 1. The outcome variable is the log number of requests per buyer (Panel A) and the log number of offers per seller (Panel B). Column (1) shows estimates without city fixed effects or year-month fixed effects. Column (2) adds city fixed effects and year-month fixed effects. Column (3) uses Bartik instruments for the number of active buyers and sellers. Column (4) uses the one-month lag of the number of national and local media articles mentioning TaskRabbit and the concurrent unemployment rate as instruments. In column (4), the instruments are interacted with city-specific fixed effects for cities active for more than 25 months. Standard errors are clustered at the city level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robustness checks with alternative market definitions are shown in Appendix A2. The results of the first-stage regressions for columns (3) and (4) are presented in Table 4.*

Table 4: *Tasks per Buyer and Offers per Seller – First-Stage Results.*

	Buyers (log) (1)	Sellers (log) (2)	Buyers (log) (3)	Sellers (log) (4)
“Bartik” Buyers (log)	1.096 [0.434]**	-1.4 [0.859]		
“Bartik” Sellers (log)	-0.273 [0.478]	2.527 [1.008]**		
Media Articles (log)			0.031 [0.175]	-0.117 [0.174]
Media*Boston			0.229 [0.028]***	0.062 [0.035]**
Media*Austin			0.056 [0.037]	0.155 [0.043]***
Media*Chicago			0.016 [0.043]	0.136 [0.051]***
Media*Seattle			-0.01 [0.056]	0.007 [0.057]
Media*LA			0.021 [0.056]	0.122 [0.024]***
Media*NYC			0.095 [0.052]*	0.19 [0.05]***
Media*Other			-0.046 [0.066]	0.168 [0.093]**
Unemployment Rate * 100			-0.02 [0.093]	-0.051 [0.094]
Unemp*Boston			-0.502 [0.106]***	0.497 [0.084]***
Unemp*Austin			0.001 [0.023]	0.13 [0.027]***
Unemp*Chicago			-0.069 [0.046]	-0.086 [0.043]**
Unemp*Seattle			-0.01 [0.038]	0.107 [0.038]***
Unemp*LA			-0.06 [0.037]	-0.04 [0.023]*
Unemp*NYC			-0.201 [0.027]***	0.286 [0.024]***
Unemp*Other			-0.083 [0.067]	-0.032 [0.061]
Constant (SF Oct ‘13)	12.672 [0.651]***	12.759 [1.052]***	8.276 [0.857]***	7.918 [0.866]***
Observations	336	336	319	319
R-squared	0.97	0.961	0.987	0.977

First-stage results for IV regressions in Table 3 (the first stage is identical regardless of the outcome variable in the second stage). Columns (1) and (2) present the first stage of IV estimations that use Bartik-style instruments for the number of active buyers and sellers in a given city-month (column (3) in Table 3). Columns (3) and (4) present the first stage of IV estimations that use media articles and unemployment rate interacted with city fixed effects for the largest cities as instruments for the number of active buyers and sellers in a given city-month (column (4) in Table 3). Controls include city and year-month fixed effects.

Table 5: *Number of Matches and Transacted Prices.*

Panel A: Number of Matches (log)					
	(1)	(2)	(3)	(4)	(5)
Tasks (log)	0.405 [0.094]***	0.410 [0.054]***	0.615 [0.077]***	0.518 [0.071]***	0.416 [0.100]***
Offers (log)	0.583 [0.092]***	0.521 [0.043]***	0.344 [0.069]***	0.389 [0.061]***	0.529 [0.080]***
Constant (SF Oct '13)	-0.970 [0.173]***	-0.416 [0.168]**	-0.505 [0.113]***	-0.096 [0.224]	-0.552 [0.202]***
R-squared	0.973	0.996	0.996	0.996	0.997
Panel B: Transacted Prices (log)					
Tasks (log)	0.185 [0.057]***	0.015 [0.049]	0.016 [0.072]	0.020 [0.079]	0.108 [0.120]
Offers (log)	-0.137 [0.044]***	-0.013 [0.042]	-0.016 [0.063]	-0.032 [0.056]	-0.115 [0.106]
Constant (SF Oct '13)	3.771 [0.116]***	4.099 [0.152]***	4.129 [0.112]***	4.254 [0.274]***	4.277 [0.076]***
R-squared	0.108	0.727	0.727	0.726	0.718
City FE	No	Yes	Yes	Yes	Yes
Month FE	No	Yes	Yes	Yes	Yes
Instruments	No	No	Nr. Buyers & Sellers	Bartik	Media & Unempl.
Observations	336	336	336	336	319

*This table shows results from OLS and IV regressions of Equation 2. The outcome variable is the log number of matches, i.e., completed transactions (Panel A), and the log average of the transacted price (Panel B). Column (1) shows estimates without city fixed effects or year-month fixed effects. Column (2) adds city fixed effects and year-month fixed effects. Column (3) uses the number of active buyers and sellers to instrument for the number of tasks requested and offers submitted. Column (4) uses Bartik instruments for the number of active buyers and sellers. Column (5) uses the one-month lag of the number of national and local media articles mentioning TaskRabbit and the concurrent unemployment rate as instruments. In column (5), the instruments are interacted with city-specific fixed effects for cities active over more than 25 months. Standard errors are clustered at the city level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Robustness checks with alternative market definitions are shown in Appendix A2. The results of the first-stage regressions for columns (3), (4), and (5) are presented in Table A3.*