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

Zoë Cullen and Chiara Farronato†

November 2014

Job Market Paper

Latest version at http://web.stanford.edu/~chiaraf/matching_p2p_latest.pdf

Abstract. We study a central economic problem for peer-to-peer online marketplaces: how to create successful matches when demand and supply are highly variable. To do this, we develop a simple static model of a frictional matching market for services, which lets us derive the elasticity of labor demand and supply, the split of surplus between buyers and sellers, and the efficiency with which requests and offers for services are successfully matched. We estimate the model using data from TaskRabbit, a rapidly expanding platform for domestic tasks, and report three main findings. First, supply is highly elastic: seller effort adjusts to short-term fluctuations in demand to equilibrate the market. Prices do not increase much when sellers are scarce, nor does the probability of tasks being matched fall. Second, we estimate average gains from each trade to be $37. Given the modest value generated from a match, efficient matching is important for the market to function well. The elastic labor supply facilitates this efficient matching by creating 15 percent more matches than alternative equilibrating mechanisms. Third, we estimate how market efficiency changes as more users join, and across different cities. We find a limited degree of scale economies, but efficiency varies greatly across cities. The cities which grow fast in the number of users are also those where market thickness, as measured by geographic density and task standardization, promotes efficient matching.

*We would like to thank Jon Levin, Susan Athey, and Liran Einav for their mentorship and guidance. We are also grateful to Tim Bresnahan, Pascaline Dupas, Brad Larsen, Andrey Fradkin, Michael Dinerstein, Akshaya Jha, and many other seminar participants at Stanford University for their invaluable comments. We appreciate TaskRabbit for providing essential data access. Support for this research was provided through the B.F. Haley and E.S. Shaw Fellowship for Economics through the Stanford Institute for Economic Policy Research.

†Department of Economics, Stanford University. Email: zcullen@stanford.edu, chiaraf@stanford.edu.

1 Introduction

The Internet has facilitated the growth of peer-to-peer marketplaces for the exchange of underutilized goods and services. Users rent rooms on Airbnb, arrange rides on Uber, and find cleaning and moving help on TaskRabbit. These platforms, which may compete with more traditional service providers, act as marketplaces for decentralized buyers and sellers to meet up and transact. This paper studies a basic economic problem for peer-to-peer marketplaces: how to equilibrate highly variable demand and supply when matches often need to be made locally and rapidly.

One obvious answer is that a decentralized market should equilibrate on price. Prices should rise when sellers are in short supply, causing buyers to pull back demand and sellers to expand supply. Of course, this leaves open the question of which side adjusts more: are peer-to-peer markets characterized by elastic demand or elastic labor supply? Moreover, recent research emphasizes that peer-to-peer markets are inherently frictional (Fradkin, 2014, and Horton, 2014). Perhaps when sellers are scarce, prices do not adjust and buyers simply fail to find matches, as in theories of frictional labor markets (Diamond, 1982, Mortensen, 1982, and Pissarides, 1985, among many).

In this paper, we use an analytical framework to analyze the possible mechanisms contributing to market equilibration when demand and supply fluctuate. We develop a simple but fully specified static model of a frictional matching market for services. Conditional on parameter values, the model lets us derive the elasticity of labor demand and supply, the split of surplus between buyers and sellers, and the efficiency with which requests and offers for services are successfully matched. All these results combined allow us to measure the aggregate value created by the peer-to-peer platform where these exchanges take place. Our framework can easily be applied to a variety of online peer-to-peer platforms for local services.

We apply our model to data from TaskRabbit, an online marketplace where buyers (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 into all posted tasks, offers, and transactions. The setting allows us to think about the efficiency and benefits of online marketplaces because successful matches must happen rapidly and locally, so we can divide the activity on the platform into separate sub-markets by time and geography, and use the large and plausibly exogenous fluctuations in buyers and sellers to estimate the demand and supply parameters of our model.[1]

[1]The opportunity to observe multiple spot markets where variable numbers of buyers and sellers match while using the same platform technology is empirically relevant. A common challenge that economists face in studies of online platforms is that it can be difficult to define and compare separate markets. In studying eBay, for example, it is hard to divide buyers and sellers into geographically segregated markets given the prevalence of cross-state and

We establish that seller effort, or 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 and little reduction in the ability of buyers to consummate trades. Our estimates imply that average gains from trade are $37 for each successful match. Because of matching frictions and search costs needed to find potential matches, the ex-ante gains are more modest. But so long as there are a lot of requests for tasks, the market can still be successful if matching frictions can be reduced and labor supply is elastic, as in this case. Specifically, the elastic supply leads to a 15 percent increase in the value of matches created relative to a setting in which seller effort does not adjust to equilibrate the market. Both the platform and the buyers equally benefit from this elastic supply. We discuss our results in light of the recent platform change on TaskRabbit, which reduced both matching and search frictions, while efficiently making use of sellers' slack capacity. We also explain why on TaskRabbit some cities are more successful than others using our model estimates. Successful cities attract and retain demand at higher rates, and are more efficient in converting requests and offers for tasks into productive matches. We relate matching efficiency to two measures of market thickness: geographic density (buyers and sellers leaving close together), and level of task standardization (buyers requesting homogeneous tasks). We conclude by combining our results to estimate the aggregate value created by peer-to-peer platforms for domestic tasks.

We start in Section 2 by describing TaskRabbit, in particular how buyers post tasks such as cleaning or grocery shopping, and how sellers submit offers to perform those tasks. We then introduce the key economic problem faced by the platform, which is to balance demand and supply to create valuable matches when these matches need to be found quickly and locally. Section 3 provides preliminary evidence that when demand is high relative to supply, sellers adjust the number of offers they submit, without significant adjustments on the buyer side. We also show that price does not adjust very much, and that the fraction of posted tasks that are completed stays relatively constant.

We then propose a simple economic model that captures the labor demand and supply decisions of buyers and sellers. In the model of Section 4, buyers choose how many tasks to post, sellers choose how intensively to search for possible jobs, and these choices determine the number of matches and the price at which trade occurs. In the model, matching frictions prevent some offers

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

from being accepted and some tasks from being filled. We allow for scale effects, both in the matching technology and in the seller cost of search. We use comparative statics to discuss the response of buyers and sellers' individual choices to changes in the aggregate number of buyers and sellers, and discuss how the size of this response depends on the model parameters.

In Section 5, we take our model to the data. The model offers a set of moment conditions: matching and pricing technologies, individual choices to post tasks and submit offers. Estimation is carried out in two steps by method of moments. First, price and number of matches are estimated as a function of the total number of offers and tasks. We allow for city heterogeneity in the efficiency with which each city is able to match tasks and offers. Second, the utility parameters of buyers and sellers are estimated from their choices to post tasks and submit offers as functions of expected match rates and prices.

Our estimates rely on variation across cities and over time in the number of buyers, both in absolute levels and relative to the number of sellers. Our key assumption is that the decision to join or leave the platform is not affected by the anticipation of unobserved match effectiveness or price shocks, nor by the expectation on others' adoption or attrition decisions. Roughly, we assume that prospective buyers and sellers have access to the same historical data on the market that we have. In Section 5 we discuss this assumption and provide supporting evidence.

Section 6 answers our first two questions. First, we quantify the high labor supply elasticity that allows the market to equilibrate in response to supply and demand shocks. We find that seller search costs are low in general ($7 for each accepted offer in the median market) and that sellers, paid $48.5 for each completed task, work at close to their opportunity cost, $33. The highly elastic supply curve, together with an average cost of listing tasks equal to 50c, implies that buyers do not need to adjust their rate of posting tasks when they are abundant relative to sellers. Buyers are neither rationed - i.e. they still match with similar probability - nor are they charged considerably higher prices. Second, we estimate the gains from trade and the surplus created by TaskRabbit. Gains from each successful match, excluding seller search costs, are equal to $37, and are shared similarly among the buyer, the seller, and the platform. However, matching frictions and seller search costs considerably affect the aggregate surplus generated by the platform. We also find that the surplus does not increase considerably with market scale: the matching technology does not display increasing returns to scale, although a larger market size moderately lowers sellers' search costs. The existence of an elastic supply curve, however, allows the market to efficiently accommodate variable demand and to create 15 percent higher value from aggregate matches. In the conclusion, we come back to this feature and our estimates of search and matching frictions to support the recent platform change.

Section 7 focuses on our third question, related to city heterogeneity. The biggest reason why some cities are more successful than others on TaskRabbit seems to be that in those cities demand is higher and the matching of buyers and sellers is more efficient. We find that buyers are somewhat sensitive to recent outcomes: the higher the probability that their task is completed today, the higher their probability to post again in the future. The general matching efficiency of the city's platform thus affects buyer retention, which together with adoption is the crucial ingredient to growth given the elasticity of existing sellers' labor supply. The geography of the city seems to be an important determinant of matching efficiency, as well as the level of task standardization. There is little evidence for economies of scale leading to increasingly larger heterogeneity across cities after small initial differences in platform success, but it is possible that more experienced buyers learn to post tasks better, i.e. in homogeneous categories, and that makes matching more efficient. We conclude our work by discussing some implications of our findings for platform design, and for other peer-to-peer markets in Section 8.

Our research contributes to a growing literature studying the economics of online marketplaces and especially peer-to-peer platforms. Recent work in this area has focused on the micro-structure of specific marketplaces, estimating search inefficiencies (Fradkin, 2014), heterogeneity in the matching process and problems of congestion (Horton, 2014), the consequences of search frictions and platform design for price competition (Dinerstein et al., 2014), the differences between distinct types of pricing mechanisms (Einav et al., 2014). There is also a large literature on trust and reputation systems (e.g. Nosko and Tadelis, 2014, and Pallais, 2014), which dates back to early work by Resnick and Zeckhauser (2002) and Bajari and Hortaçsu (2003).

Our work is complementary to this literature in that we abstract from many forms of heterogeneity and asymmetric information that are emphasized in other papers, and from issues of strategic pricing and reputation. Instead, we offer a framework that enables us to examine in detail the particular problem, which we view as both important and common, of balancing highly variable demand and supply, and the process of market equilibration. The model we propose is in principle applicable to other peer-to-peer matketplaces that match buyers and sellers of local and time-sensitive services.

In studying the balancing of demand and supply on TaskRabbit, our modeling approach draws on the literature on frictional search and matching in labor markets (Petrongolo and Pissarides, 2001). In particular, Pissarides (2000, ch.5) most closely resembles our model. Workers submit applications to posted vacancies, at a cost, and employers select among applications received. In equilibrium, the number of applications submitted reflects the expectation of matching. Most theoretical work assumes constant returns to scale in the functional form of the matching technology, supported in part by results in the empirical literature - for example, in Anderson and Burgess (2000). The are two differences in our setup. First, ours is a market for services in the spirit of Michaillat and Saez (2013), where each buyer (resp. seller) can be matched to multiple sellers (buyers). Second, we allow for economies of scale in the costly search process.

Our analysis of scale economies and platform growth and success touches on two additional areas of research. First, papers such as Ellison and Fudenberg (2003) have emphasized that an important issue when marketplaces compete for buyers and sellers is whether increasing scale makes a marketplace more efficient. We find that scale economies per se are not a major determinant of market efficiency, for instance compared to basic fixed features such as the geography of a given city. In modeling platform growth, through adoption and attrition decisions, we also connect to a large literature on innovation diffusion and how the speed of growth of new technologies can depend on information flows, technology improvements, and network effects (Young, 2009). We mention other related papers in Section 7.

2 Setting and Data

This section describes the TaskRabbit platform, the data, and some salient facts that are important for our analysis. We first describe how the platform 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. As a result, a central problem for the platform is to balance supply and demand, which are highly variable, on a local high-frequency basis. In the next section we provide some initial evidence that the market equilibrates mainly through variation in seller effort, with only minimal price responsiveness. We tailor our model in Section 4 to capture this feature. Finally, we provide a first look at differences in market success across cities, an issue we return to in Section 7.

2.1 The TaskRabbit Platform

TaskRabbit is an online platform that allows "posters" to outsource domestic tasks to "rabbits". It currently operates in 18 cities in the United States.[2] Posters post a description of the requested task in a flexible manner. Rabbits can search through posted tasks on city-specific lists and respond with offers (Fig. 1). We will refer to posters as *buyers* and rabbits as *sellers* of services.[3]

Buyers on TaskRabbit can post virtually any sort of domestic tasks or errand (e.g. babysitting for a goldfish), but the majority of tasks are relatively standard and generic. The five largest categories are shopping and delivery (24%), moving help (12%), cleaning (9%), home repairs (6%), and furniture assembly (4%). These tasks typically do not require sellers with highly specialized skills. The nature of the tasks implies that services generally are provided locally and on relatively short notice. Almost all users (93.6 percent of them) participate in just one city. At the same time, of the 48.5 percent of tasks that are matched, 97 percent are filled within one or two days.[4]

The matching process can work in two ways. A buyer can 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 slightly more standardized (65 percent of them are in the top 5 categories versus 48 percent of auctions), and prices are lower ($49 versus $63), but their share on the platform, at 41 percent, has not changed considerably over time or across cities. About 77 percent of tasks receive an offer, and of them 63 percent result in a match. Matches 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 coordinate on specific task details.

Platform users tend to be either buyers or sellers, but not both. Indeed, 80.3 percent of users have only ever posted task requests, and 16.3 percent have only ever submitted offers. The buyers on the site are predominantly female (55 percent of buyers) and relatively affluent. The modal buyer is a woman between the age of 35 and 44 with a household income between $150,000 and

[2]The active cities in the US are, in 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 2010) Phoenix and Denver (November 2011).

[3]Leah Busque first formulated her idea for TaskRabbit when one evening she realized she had ran out of dog food. 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/).

[4]To add to the local and urgent nature of tasks, the platform's ranking algorithm prioritizes newly posted tasks within each city. Indeed, to every seller searching through posted tasks, the platform shows a list of local tasks, ranked according to their posting time (most recent at the top).

$175,000. The sellers are younger and not surprisingly have lower income. The modal seller is 25-34 years old and has a household income between $50,000 and $75,000.

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 March 2013, applicants received a background check, 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. They ranged between 7 and 49 percent in different months, and on average they were very low - only 13.6 percent. In the spring of 2013 TaskRabbit reduced the amount of screening in a successful attempt to add more sellers. The current process involves simpler background checks and social controls - Facebook or Linkedin verification - paired with a system of users' reviews.

2.2 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. Since we have no record of the actual entry date, we define the month of entry into a city as the first calendar month in which 20 or more local tasks were posted.[5]

The data include all posted tasks, offers, and matches that occurred on the platform during the study period. We exclude virtual tasks[6] (10.4%) and tasks posted in not yet active cities (0.23%). We also drop 10.3 percent of tasks that use other assignment mechanisms and keep only auction and posted price tasks. We merge the tasks with the corresponding offers, and we drop extreme price outliers (top and bottom 1 percent in bids or charged prices). To deal with the fact that posted price tasks occasionally receive multiple offers (6.04 percent of them did), we only keep the matched offer in case of success, or select one of the received offers at random. This simplification restricts posted price tasks to receive either one or no offers. Finally, for much of the paper we will aggregate activity at the city-month level, and drop city-months with less than 50 buyers posting tasks or less than 20 sellers making offers.

Table 1 shows summary statistics for the data. In the first panel, an observation is a posted task. Out of all posted tasks 78 percent receive offers, and those tasks receive 2.8 offers on average.

[5] We verified the accuracy of our definition through media coverage of the platform and by talking with TaskRabbit employees.

[6] A 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.

Of the tasks receiving offers, 63 percent are successfully completed at an average price of $57. The platform charges a 20% commission fee on successful tasks.[7]

In the second panel of Table 1, 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. On average, there are 708 active buyers and 255 sellers in a city-month, but there is large variation across cities and months. Each buyer posts 1.6 tasks, and each seller submits 6.4 offers. The task success rate is 46 percent 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 sellers varies the most across city-month observations, with limited variation in tasks per buyer, matches, and prices.

During the 4-year period we study, the platform was growing in all cities, and quite rapidly in some. Figure 2 plots the number of successful matches for the 10 oldest cities.[8] Over the period considered, some cities grew from a few monthly matches to thousands of exchanges, like San Francisco and New York, while some others grew at a reduced pace, like Portland and Seattle. We will use the cross-city and over time variation in market size in our empirical section to study the effect of scale. We will also examine the dynamic forces underlying the platform growth in Section 7.

3 Descriptive Evidence

A key feature of the platform is that there are large fluctuations in demand and supply. Since matches must be made quickly and locally, this raises the question of what happens when demand is especially high or low relative to the number of sellers. Here we provide some initial evidence. In the next section we develop a theoretical model of market equilibration which allows us to analyze labor demand, labor supply, and market clearing in more detail.

Figure 3 shows the variability of demand relative to supply in the 10 oldest cities at a monthly level. Specifically the figure plots the number of active buyers in the city-month divided by the number of active sellers. As before, activity is defined as posting at least one task (for buyers) or submitting at least one offer (for sellers). There are sizable fluctuations in the ratio of active buyers to active sellers, both within a city over time, and across cities within a month. In San Francisco,

[7]The commission fee can sometimes depart from 20 percent, 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.

[8]Similar patterns to those in Figure 2 are found in the 8 youngest cities.

for example, certain months have two buyers per seller, while other months have six buyers. During the same calendar month, some cities may have only one buyer per seller, while other cities have five. The variability is not due to a single time trend. Month-to-month changes in the buyer to seller ratio are both positive and negative in no particular order. Finally, we emphasize some persistent heterogeneity across cities and across months. For instance, San Francisco has many more buyers per seller than Los Angeles.

In principle, there are several ways in which the market might function given this variability. One possibility is that with fewer sellers, buyers may not be able to have tasks performed, either because of higher prices which deter them, or because a smaller fraction of posted tasks receive offers. Another possibility is that seller labor supply expands. We show that the latter occurs, and that labor supply is sufficiently elastic that the level of price increase needed to generate a supply response is small.

Figure 4 first shows that the number of posts per buyer does not adjust when sellers are in short supply. Here, we divide the 336 city-months into four groups, corresponding to the four quartiles of the distribution of the buyer to seller ratio. For each group we compute the average number of tasks per buyer and offers per seller. The figure shows that regardless of the number of buyers per seller, buyers always post 1.6 tasks each.

In contrast, Figure 4 shows that sellers submit many more offers when they are scarce relative to demand. For the city-months in the lowest quartile of the buyer to seller ratio (1.5 buyers per seller on average) sellers submit 4.4 offers on average. For the city-months at the other extreme (3.8 buyers per seller) sellers each submit twice as many offers, 9.1. Offers do not fully double as buyers double relative to sellers, so the match rate of tasks slightly declines (Figure 5). However, the sellers' intensive margin response, together with buyers' constant rate of task posting, translate into a large expansion in the number of trades as the number of buyers per seller increases.

Perhaps surprisingly, transacted prices move very little when sellers are scarce or abundant. Figure 5 shows the average price of completed tasks for the city-months sorted by the buyer to seller quartiles. Average transacted price is always between \$52 and \$59, 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 not much price increase is needed to generate a large intensive margin increase in labor supply.

So far we have not ruled out the possibility that there might be something special about certain cities or certain months that leads to sellers making more offers when they are scarce for reasons that are not causal. For example, San Francisco tends to have a higher number of buyers per seller than other cities, so San Francisco city-months are disproportionately represented in the upper quartiles of the buyer to seller ratio distribution in Figure 4. If sellers in San Francisco submit more offers for reasons unrelated to the number of buyers per seller - for example because they can find tasks that are closer to them - from Figure 4 we might wrongly conclude that a higher number of buyers per seller leads sellers to work harder.

A first step towards establishing causality is to consider a simple difference in differences specification that includes city and time fixed effects, and therefore controls for factors leading certain cities or months to have more buyers to focus instead on idiosyncratic time variation in demand conditions within cities and within months. Specifically, we estimate OLS regressions of the following type:

$$\log(y_{tc}) = \theta_1 \log\left(\frac{B_{tc}}{S_{tc}}\right) + \theta_2 \log\left(\sqrt{S_{tc}B_{tc}}\right) + \eta_c + \eta_t + \nu_{tc} \ , \tag{1}$$

where c, t denote city c and calendar month t, $\frac{B_{tc}}{S_{tc}}$ is the buyer to seller ratio, $\sqrt{S_{tc}B_{tc}}$ is the geometric average of buyers and sellers, and y_{tc} is one of the four relevant variables: users' choices (tasks per buyer, offers per seller), and outcomes (task match rate, prices). η_c controls for city-specific propensities to use TaskRabbit which are time invariant. Similarly, η_m captures time-specific adjustments to usage intensities that are common across all active cities. Standard errors are clustered at the city level.[9]

The results are shown in Table 2. The top panel shows the regression results without fixed effects, the bottom panel shows those with fixed effects. We first call attention to the comparison of the coefficients between the two panels: adding fixed effects does not change the response of sellers to fluctuations in the buyer to seller ratio, not in sign, size, or significance. The same can be said for the task match rate. While the coefficients on price and tasks per buyers are one or two

[9] Given the log-specification, we can transform the right-hand side to be a function of the number of buyers and sellers: $\log(y_{tc}) = \hat{\theta}_1 \log B_{tc} + \hat{\theta}_2 \log S_{tc} + \eta_c + \eta_m + \nu_{tc}$, where $\hat{\theta}_1 = \theta_1 + 0.5\theta_2$ and $\hat{\theta}_2 = -\theta_1 + 0.5\theta_2$. The results in Table 2 imply that buyers post the same number of tasks, regardless of how many users are active. Each seller submits more offers when there are more buyers, holding constant the number of sellers, but submits fewer offers when there are more sellers. The task match rate goes down as more buyers post tasks, but goes up when there are more sellers. Finally the price stays relatively constant as a function of buyers and sellers. The specification of the regression in terms of number of buyers and sellers helps interpret the effects of the number and composition of users in terms of network externalities: the utility a user derives from participating in a city-month depends on the number and type of other active users. We do still prefer the specification from equation 1 because of the particular nature of network externalities on TaskRabbit: users benefit from the platform insofar as it allows them to trade services, and users' participation affects the terms of trade. 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. At the same time, holding the relative number of buyers and sellers constant, a seller can like a large market more or less than a small market. A preference for larger markets can arise because of scale economies, while one for smaller markets may be due to congestion.

orders of magnitude smaller when controlling for city and time characteristics, they are in both cases quantitatively small and statistically insignificant. This provides some confidence that the platform is used by buyers and sellers in a similar way both over time and across cities.

The size of the coefficients confirm what was shown in the plots. We discuss those from the bottom panel of Table 2, obtained controlling for potential city and month differences. An increase in the number of buyers per seller has virtually no effect on how many tasks each buyer posts. On the other hand, doubling the number of buyers per seller of the median city-month, where the median is selected according to the distribution in the buyer to seller ratio and holding everything else constant, increases the number of offers submitted by each seller from 5.6 to 7.5. The effect on the task match rate is negative, but smaller in percentage terms: doubling the median number of buyers per seller decreases the match rate of tasks from 65.6 to 56.2 percent. Finally, buyers pay just a few cents more for completed tasks when they are twice as prevalent relative to sellers.

The regressions also estimate the effect of market size, previewing possible mechanisms for economies of scale. A city-month with more active participants significantly increases the number of offers submitted by sellers, holding constant the relative number of active buyers and sellers. In particular, holding the ratio of buyers to sellers fixed and doubling the number of participants of the median sized city-month increases the number of offers submitted from 5.5 to 6.5. More active participants also seem to raise buyers' rate of task posting and the task match rate, but the effect is not statistically significant when including city and month fixed effects. Finally, the price appears invariant to the number of participants, consistently across the two panels. We will capture each of these features in our model.

4 Model of a Market for Services

We now propose a model of how the TaskRabbit marketplace matches tasks and offers, and how buyers and sellers make decisions about whether to post tasks and how much effort to put into making offers. We then use the model to explain what happens when there is variation in the number of active buyers and sellers, and explain how this variation can be used to identify the elasticity of labor demand and supply, the division of surplus in the market, the effects of increased market size, and the efficiency of matching in different cities and different market conditions.

We assume for simplicity that buyers are all identical and in equilibrium choose the same number of tasks to post. Similarly, sellers are identical and choose the same intensity with which to search and submit offers. We also treat tasks as homogeneous. Obviously, this is a large simplification, but it does correspond to our earlier observations that most tasks on the platform are relatively standard and generic, and that they do not require specialized skills. More importantly, it allows us to focus on the problem of widely fluctuating supply and demand, without being bogged down by a complicated heterogeneous matching framework.

Market Technology. There is a measure B of identical buyers and a measure S of identical sellers. Each buyer will choose a number of tasks, β, to post. Each seller chooses a number of offers, σ, to make. The total number of services requested in a market is $b \equiv B\beta$, while the total number of offers submitted is $s \equiv S\sigma$.

The number of trades between buyers and sellers is given by the matching function:

$$m = M(s, b). \tag{2}$$

$M(s, b)$ is continuous and differentiable, and increasing in both its arguments. Each request is matched with probability $q^b = \frac{m}{b}$ and each offer is successful with probability $q^s = \frac{m}{s}$. We assume that $M(s, b) \leq b$ and $M(s, b) \leq s$ to guarantee that matches are never larger than total requests or offers. In each match, the buyer pays price $p = P(s, b)$, the seller receives $(1 - \tau)p$, and the platform keeps τp as commission fee. In particular, price is determined as a function of services requested and offered, and is assumed to be a continuous and differentiable function, increasing in b and decreasing in s.[10] Later we will estimate the matching and price functions from the data, but we will not provide a more micro-level model of price determination - e.g. by modeling the posted pricing decision or the bidding game between sellers.

Buyers and sellers choose how many requests to post and how many offers to submit with full knowledge of the matching and price determination processes, but without the possibility to affect either of those with their individual choices because each participant is small relative to the market.

Buyer's Choice to Post Requests for Services. Each buyer randomly receives a number of potential needs to outsource. We assume that the number of service needs is a random draw from a Poisson distribution, iid across buyers, with mean arrival μ.[11] Each service is worth $v - p$ to

[10] In fact, several matching models of the labor market assume that the wage is either a parameter altogether or pinned down by other parameters. See, for example, Montgomery (1991), Hall (2005), Blanchard and Gali (2010), Michaillat (2012), and Michaillat and Saez (2013).

[11] For the conditions under which a continuum of independent and identically distributed random variables sum to

its buyer, where v is the fixed value of having the task completed and p is the price paid. There is a cost ξ of posting each task, drawn from an exponential distribution, iid across needs and buyers: $\xi \sim exp(\eta)$. The average cost of posting a task is therefore equal to $\frac{1}{\eta}$.

The buyer's problem is to choose whether to post each needed service. The decision is separable across service needs. If a buyer makes a request for need t, she pays cost ξ_t and expects payoff $q^b(v-p)$. She optimally chooses to submit a request whenever the listing cost is small enough: $\xi_t \leq q^b(v-p)$. The expected number of requests posted by a representative buyer is:

$$\beta = \mu Pr\left(\xi \leq q^b(v-p)\right) = \mu\left(1 - e^{-\eta q^b(v-p)}\right). \tag{3}$$

Seller's Choice to Submit Offers for Services. Each seller chooses a level of effort σ spent searching through buyers' requests. An effort level σ corresponds to a discovery process of profitable requests, to which the seller submits offers. Higher effort σ makes it more likely to find a higher number of profitable submissions.[12] Specifically, we assume that the number of suitable tasks identified and offers submitted is a random draw from a Poisson distribution $Poi(\sigma)$, with mean equal to the chosen effort level and independent across sellers. Given this assumption, we will interchangeably refer to σ as the level of search effort or the expected number of offers submitted by a representative seller. Search effort is costly, and its cost rises at an increasing rate. In particular, we assume that the cost of search effort is equal to $\frac{1}{2\gamma(b)}\sigma^2$, with $\gamma(b)$ being a continuous and increasing function of the total number of tasks posted.[13] Conditional on matching a submitted offer, the seller's profit is $(1-\tau)p - c$, where τ is the platform commission fee and c is the fixed cost of completing the task.

The problem of a representative seller is to choose the optimal level of search intensity subject to expectations on matching and prices:

$$\underset{\hat{\sigma}}{Max}\ \hat{\sigma} q^s \left[(1-\tau)p - c\right] - \frac{1}{2\gamma(b)}\hat{\sigma}^2,$$

a nonrandom quantity in large economies, see Judd (1985), and Duffie and Sun (2012).

[12] Specifically, we assume that the distribution of application arrivals for a given σ first order stochastically dominates the distribution for any $\sigma' \leq \sigma$.

[13] We model sellers' search costs as increasing in the intensity of search at an increasing rate. This assumption can be better understood in terms of time needed before finding a new task to which a seller chooses to make an offer. Conditional on a level of effort, it is likely that the first profitable task is easier to find than the second, the second is easier than the third, and so on. If a seller wanted to double the number of profitable tasks found, his level of effort would then be more than twice as costly. In addition, search costs are decreasing in the number of total tasks posted. In a market with many posted tasks, a seller is likely to spend less time finding the same number of profitable applications as in a smaller market. If a seller wanted to send the same expected number of offers in a large market his level of effort would then be less costly.

The optimal level of search effort satisfies:[14]

$$\sigma = \gamma(b)q^s\left[(1-\tau)p - c\right]. \tag{4}$$

Equilibrium. Equilibrium in the market is defined as a state in which buyers and sellers maximize their objective functions subject to the matching and pricing technologies and correct expectations of other agents' behavior. The equilibrium requires consistency of individual optimal choices (β and σ) with expectations on average behavior in the market ($\overline{\beta}$ and $\overline{\sigma}$). Given the size of buyers B and sellers S present in the market, we define the competitive equilibrium as a vector (β, σ, p, m) such that:

- The transacted price is determined according to $p = P(s, b)$, and the number of matches is determined according to $m = M(s, b)$, where $b = B\overline{\beta}$ and $s = S\overline{\sigma}$.
- Taking $q^b = \frac{m}{b}$ and p as given, buyers list the number of service requests to maximize utility. The number of requests β of the representative buyer is given by equation (3).
- Taking $q^s = \frac{m}{s}$ and p as given, sellers choose the level of search intensity to maximize utility. The level of search intensity σ (i.e. of offers submitted) of the representative seller is given by equation (4).
- The actual average number of requests posted is $\beta = \overline{\beta}$ and offers submitted is $\sigma = \overline{\sigma}$.

In equilibrium, all buyers choose the same strategy in terms of the decision to post tasks, which in turn is consistent with the expected posting rate. The model explains differences in the actual number of requests across buyers as arising from the Poisson arrival rate of needs and from different draws of listing costs. Analogously, in equilibrium, all sellers choose the same level of search intensity, which in turn is consistent with the market average intensity. Differences in the rate of offer submission across sellers arise from the Poisson process with which they discover profitable requests.

[14] Buyers' choice to post tasks and sellers' choice of search effort are not symmetric. On the buy side, there is an exogenous arrival of tasks, and a decision to post each of them separately conditional on arrival. A buyer in need of moving help selects whether to post it or find an alternative solution - another service provider or informal help - as a function of the expected value from each option. On the sell side, the setup is truly a choice of platform usage intensity. A seller selects his optimal level of search effort, and if he finds profitable tasks he submits offers for sure. In this case, a seller chooses his time allocation between leisure and searching for services to sell as a function of the expected benefits from the two activities.

Assumptions on Matching and Price. The matching technology displays constant returns to scale if a doubling of the number of tasks b and offers s doubles the number of matches m. Analogously, the price function is invariant to scale if doubling tasks and offers does not affect the price p. If the matching technology displays constant returns to scale, the total number of matches (equation 2) can be rewritten as $m = M\left(1, \frac{b}{s}\right) s$, where $\frac{b}{s}$ is the task to offer ratio, the offer match rate is equal to $q^s = M\left(1, \frac{b}{s}\right)$ and the task match rate is $q^b = \frac{M\left(1, \frac{b}{s}\right)}{\frac{b}{s}}$. If the price function is invariant to scale, it can also be rewritten just in terms of the task to offer ratio: $p = P\left(1, \frac{b}{s}\right)$. With a slight abuse of notation, we let $m = M\left(\frac{b}{s}\right) s$ and $p = P\left(\frac{b}{s}\right)$. This reformulation implies that the match probabilities of tasks and offers, as well as the price, are just a function of demand relative to supply, and not on the overall level of demand or supply. If, in addition to these two conditions, seller search costs $\gamma(b)$ do not decrease much with market scale, the equilibrium is unique. In Section 6 we test that these conditions hold on TaskRabbit. Anticipating this, we maintain them for the rest of our discussion.

Optimal Choices. Figure 6 illustrates the optimal individual choice of a seller. We discuss the seller side, noting that for buyers the reasoning is analogous. The figure plots the individual level marginal benefit (solid red line) and marginal cost curves (solid blue line) as a function of search effort σ. The marginal cost curve $\frac{\sigma}{\gamma(b)}$ is increasing in effort, while the marginal benefit curve $q^s\left[(1-\tau)p - c\right]$ is independent on the single seller's choice. This is because in a large market the offer match rate and price depend on the market average effort level $\overline{\sigma}$ and posting rate $\overline{\beta}$, which cannot be affected by any single participant alone. The shaded area in the picture between the marginal benefit and marginal cost curve is the seller surplus at equilibrium, and can be inferred knowing the matching and pricing functions, as well as seller costs. The flatter the marginal cost curve, in the case of low search costs, the smaller is the seller surplus.

In equilibrium, a seller optimal choice of σ must be consistent with the market-average search effort: $\overline{\sigma} = \sigma$. We can therefore also draw the market level marginal benefit curve, where every seller in the market chooses effort level σ, and the buyers' posting rate is held constant at $\overline{\beta}$ (red dotted line in Figure 6). In this case, both the offer match rate and the price will be affected by σ: if every seller in the market were to increase his search intensity, each offer would be less likely to be matched, and sellers would receive a lower price for every trade. Thus the market level marginal benefit is decreasing in σ. Consistency of σ with $\overline{\sigma}$ requires that the market level and the individual level marginal benefit curves must cross the marginal cost curve at the same point ($Eq1$ in Figure

6): each individual seller takes as given the flat marginal benefit curve generated by every other seller in the market choosing the same effort level as his own.

Changes in the match probabilities and prices affect the best response functions. Holding price constant, an increase in the offer match rate q^s increases search effort. The increase in q^s corresponds to an upward shift of the marginal benefit curve (upper red dotted line in Figure 7). The size of the increase in search - i.e. the horizontal difference between $Eq1$ and $Eq2$ - depends on the function $\gamma(b)$ and on the cost of completing the task c. The lower the cost of search, corresponding to higher $\gamma(b)$, the flatter is the marginal cost curve in the figure. With an almost flat marginal cost curve even a small change in q^s can lead to a large change in search effort. At the same time, the smaller the seller profit $(1-\tau)p-c$ for each task, the smaller is the upward shift of the marginal benefit curve, thus the smaller the increase in effort. Holding the offer match rate constant, an increase in price will raise a seller's search effort since each task will pay more. As before, the flatter the marginal cost curve (high $\gamma(b)$) the larger the change in σ. Moreover, in percentage terms, the smaller the seller profit $(1-\tau)p-c$, the higher the change in search effort.

Comparative Statics. The model generates predictions about the effects of both market size and market composition. We capture these by thinking about changes in B holding $\frac{B}{S}$ fixed, and changes in $\frac{B}{S}$ holding B fixed. First consider an increase in B holding $\frac{B}{S}$ fixed. This lowers the cost to sellers of finding tasks, and raises their search effort. This in turn makes it more attractive to post tasks, raising buyers' posting rates. The result is an increase in β and σ, with sellers responding more. Proposition 1 states it formally (proofs are in Appendix A).

Proposition 1. *(Effect of scale) An increase in B, holding $\frac{B}{S}$ fixed, leads to an increase in β and σ, and a decrease in $\frac{\beta}{\sigma}$. This in turn implies lower prices and seller match rates, and higher buyer match rates.*

Now consider an increase in $\frac{B}{S}$, holding market size B fixed. The direct effect is to make the market more attractive to sellers, thus raising σ, and less attractive to buyers, lowering β. There is an indirect effect because once buyers lower β the market contracts, which raises sellers' costs. We show that at equilibrium this effect cannot dominate, so that the new equilibrium involves more seller effort and less buyer posting. However, the endogenous response of buyers and sellers does not fully compensate for seller scarcity, and the task to offer ratio is still higher at the new equilibrium.

Proposition 2. *(Effect of user composition) An increase in* $\frac{B}{S}$*, holding B fixed, leads to a decrease in* β*, an increase in* σ*, a decrease in* $\frac{\beta}{\sigma}$ *and an increase in* $\frac{b}{s}$*. This in turn implies higher prices and seller match rates, and lower buyer match rates.*

These comparative statics predictions are consistent with our empirical results from the previous section, and will form the basis for our identification strategy. To see how this works, consider the supply side where the unknown parameters will be (c, δ, γ). The demand side is similar. If the search cost parameters (γ, δ) were known and given that p and q^s are observable, the choice of σ in a single market would directly pin down the marginal cost c of performing each task, regardless of the number and composition of participating users. A low number of offers σ would imply a low c, holding everything else constant. Next consider adding variation in market size (left-hand side panel of Figure 7). We discuss the intuition holding constant buyers' β, although by increasing their posting rate, buyers' response actually amplifies the direct effect. An increase in B pivots the marginal cost curve downward. If we knew the other parameters affecting seller utility (c, γ), the magnitude of this shift would only depend on δ, the extent of scale economies in search effort: when δ is large, search costs decrease considerably with market scale, and this directly translates into a large increase in offer submission. So roughly speaking the response of seller offers σ to changes in B identifies δ. Finally, we add variation in the number of buyers relative to sellers $\frac{B}{S}$ (right-hand side panel of Figure 7). Again we discuss the intuition holding constant buyers' β. An increase in $\frac{B}{S}$ shifts and pivots the marginal benefit curve upward by increasing the match rate and the price of each additional offer. If c and δ were known, the magnitude of this shift would only depend on γ, the slope of the marginal cost curve: when γ is large, the marginal cost curve is flat and an increase in the relative number of buyers translates into a large increase in offer submission. So effectively changes in $\frac{B}{S}$ and the response of σ allows us to identify γ.

5 Econometric Model

In this section, we describe how we move from the theoretical model to an econometric model, and how we apply the model of the TaskRabbit data.

We then describe our estimation strategy and discuss our identification assumptions. Our main assumption is that the number of buyers and sellers who consider using the platform in any city-month does not depend on contemporaneous unobserved characteristics that affect the efficiency of matching or price determination, the cost of search or posting, or the value of trade. We also present the functional forms of the matching and pricing functions, and make assumptions about how buyers and sellers form beliefs about them when making search intensity and posting decisions. Results are presented in the next section.

5.1 Market Definition

Our model envisions a single static market, while trades occur continuously in the data. To create an empirical analogue to the model, we define distinct markets in the data. Given that 94 percent of users post or work in a single city, it is natural to treat cities as separate. The fact that 97 percent of successful tasks are matched to offers within 48 hours of posting suggests segmenting the data in time as well. One option is to treat each city-month (e.g. San Francisco in October 2013) as a separate market. Within a city-month, we treat buyers and sellers, as well as their tasks and offers, as homogeneous, following the model, and discuss this further in Appendix B. This definition allows us to consider each participant as small relative to the size of the market, which is our modeling assumption, and also lets us smooth shorter time variation due to potential task heterogeneity. Other market definitions do not change our qualitative results, as shown in the Appendix.

Our market definition is motivated by several additional considerations. First, we do not separate markets along the various task categories - cleaning, furniture assembly, and so on - because sellers do not specialize: of the sellers who submitted 10 offers or more 63.6 percent did so in more than 10 categories, and of the sellers who were successfully matched to more than 10 tasks, 43 percent did so for tasks in more than 10 categories. Second, we follow TaskRabbit business practice and do not separate markets into geographic partitions smaller than the metropolitan boundaries.[15] Third, we choose the calendar month as the relevant time window as a way to balance the short time period over which tasks receive offers with the need to have enough offers and tasks in each market to estimate match probabilities, average prices, and search and posting intensities.

There is one further data issue we must address in moving between the model and the data. In the model, all buyers choose the same posting threshold, but the distribution of posted tasks across buyers is Poisson. This means that some participating buyers post zero tasks. Similarly the distribution of seller offers is Poisson and some participating sellers make zero offers. However, in the data we cannot distinguish between buyers and sellers who were considering posting tasks and

[15] We do not observe any sort of clear neighborhood partitioning in the data, although the platform's setup does not preclude it. We provide additional details in the Appendix.

submitting offers but did not, and those who were completely disengaged from TaskRabbit.

Our solution is to rely on the Poisson assumption in the model and use it to impute the number of buyers posting zero tasks and sellers submitting zero offers. Specifically, if the number of buyers posting at least one task is $\tilde{B}$ and the average number of posts among these buyers is $\tilde{\beta}$, then under the Poisson assumption the average number of posts β among all buyers solves $\frac{\beta}{1-e^{\beta}} = \tilde{\beta}$. Under the same assumption, the total number of participating buyers B, which is the sum of buyers posting zero tasks and those posting one task or more, is equal to $B = \frac{\tilde{B}}{1-e^{\beta}}$. We perform a parallel exercise to impute the total number of participating sellers making zero offers, and to appropriately construct S and σ.[16]

5.2 Econometric Model

We now describe the econometric model that we take to the data. It has three components: the aggregate pricing and matching functions that map offers and tasks to market outcomes; the expectations formed by buyers and sellers about their probability of matching and the market price (assumed to be rational); and their optimal search and posting decisions.

Throughout, $n = (c, t)$ identifies a market, our unit of observation: c denotes the city, and t denotes the calendar month. We let B_n and S_n denote the number of participating buyers and sellers, β_n and σ_n denote the participants' average posting and offer intensities, and b_n and s_n denote the total number of posts and offers in a market. Finally, we let M_n denote the number of matches, and P_n the average transacted price.

Matching and Pricing Functions. We assume that the total number of matches and average prices in a market are Cobb-Douglas functions of the number of tasks posted and offers made.[17] Specifically,

$$M(s_n, b_n) = A_n s_n^{\alpha_1} b_n^{\alpha_2} \tag{5}$$

$$P(s_n, b_n) = K_n s_n^{\rho_1} b_n^{\rho_2} \ , \tag{6}$$

The variables A_n and K_n are market level productivity and pricing shifters. We assume that

[16] This solution is somewhat imperfect as it relies on a relatively strong distributional assumption. However, we examine the validity of the Poisson assumption, and check the robustness of our results to alternative imputation strategies in the Appendix, finding that our empirical results are not very sensitive to alternative approaches.

[17] Petrongolo and Pissarides (2001) summarize the wide empirical support for a Cobb-Douglas matching function with constant returns to scale. For its micro-foundation, see Stevens (2002).

each has an error component structure. That is,

$$A_n = A_t A_c \epsilon_n^a$$

where A_t is a month effect, A_c is city-specific match efficiency, and ϵ_n^a is an idiosyncratic shock to matching, which has expected value 1 and is not anticipated by buyers or sellers. Similarly, we assume that $K_n = K_t K_c \epsilon_n^k$, and again assume that ϵ_n^k is not anticipated by market participants.

In this specification, we expect that the number of matches will be increasing in both inputs, s and b - i.e. $\alpha_1 \geq 0$ and $\alpha_2 \geq 0$. The market exhibits increasing returns in matching if $\alpha_1 + \alpha_2 > 1$, and constant returns if $\alpha_1 + \alpha_2 = 1$. Under increasing returns, doubling the number of tasks and offers more than doubles the number of matches. For pricing, we expect more posted tasks will drive up prices, and more offers will reduce them, so that $\rho_1 \leq 0 \leq \rho_2$. If $\rho_1 + \rho_2 = 0$ the pricing is not affected by scale: doubling both the number of offers and tasks has no price effect.

Participants' Expectations. Therefore, the expected matching probabilities are $q_n^b = \frac{Q_n}{b_n}$ and $q_n^s = \frac{Q_n}{s_n}$, where

$$Q_n = E\left[A_n s_n^{\alpha_1} b_n^{\alpha_2} | s_n, b_n\right] = A_t A_c s_n^{\alpha_1} b_n^{\alpha_2} \ , \tag{7}$$

and the expected price is

$$p_n = E\left[K_n s_n^{\rho_1} b_n^{\rho_2} | s_n, b_n\right] = K_t K_c s_n^{\rho_1} b_n^{\rho_2} \ . \tag{8}$$

Buyers expect to pay p_n, and sellers to receive $(1-\tau)p_n$. Given the constant 20 percent platform commission fee, we fix $\tau = 0.2$.

Optimal Decisions. Finally we can write the buyer and seller optimal decisions. From Section 4, buyers choose a posting intensity β_n equal to

$$\beta_n = \epsilon_n^b \mu \left[1 - e^{-\eta q_n^b (v_n - p_n)}\right] \ . \tag{9}$$

ϵ_n^b is a task arrival shock, known to participating buyers and sellers but independent of their (prior) decisions to participate in the market. It can be thought of as a city-month specific driver of demand for services among participating buyers. An example might be an increase in requests for shopping deliveries due to December snowstorms in Chicago, which does not drive buyers' or sellers' decision to stay or join TaskRabbit in that particular market, but does increase the posting intensity of participating buyers.[18]

On the seller side, we specify search costs as $\frac{1}{\epsilon_n^s \gamma b_n^\delta}\sigma_n^2$ - i.e. decreasing in the total number of posted tasks at rate δ. Each seller chooses a search intensity σ_n equal to

$$\sigma_n = \epsilon_n^s \gamma b_n^\delta q_n^s (0.8 p_n - c_n) \ . \tag{10}$$

As for buyers, ϵ_n^s is a supply shock that reduces the cost of search. It can be interpreted as a city-month specific increase in time availability among participating sellers.

Recall from equation 8 that price has a time-specific component K_t. In order to homogenize values and costs over time, we assume that the same time parameter multiplicatively changes buyer values and seller costs of performing tasks. Specifically, we assume that $v_n = K_t v$ and $c_n = K_t c$.

5.3 Identification

We make two key identification assumptions: i) participating buyers and sellers do not anticipate the idiosyncratic pricing and matching shocks in making their posting and offer decisions, and ii) the number of participating buyers and sellers in a given market does not depend on these shocks, nor on the unobserved components of buyer and seller utility. Lastly, we make the additional and convenient assumption that the unobservable shocks are iid across markets. Formally, we write our assumptions as follows:

Assumptions: *We assume that:*

1. *Pricing and matching shocks are not anticipated: the vector* $(\epsilon_n^a, \epsilon_n^k)$ *is independent of* (σ_n, β_n).
2. *Limited predictability for prospective users: the vector* $(\epsilon_n^b, \epsilon_n^s, \epsilon_n^a, \epsilon_n^k)$ *is independent of* (S_n, B_n).
3. *IID: the vector* $(\epsilon_n^b, \epsilon_n^s, \epsilon_n^a, \epsilon_n^k)$ *is independently and identically distributed across markets, with mean* $(1, 1, 1, 1)$ *and variance covariance matrix* Σ.[19]

An issue of splitting cities into multiple markets over time is that buyers and sellers might anticipate the future value of exchanges on the platform and base their decision to stay or leave on these rational expectations. Two empirical features of TaskRabbit lead us to think that forward-looking behavior is not prevailing: the low level of retention and its response to future outcomes. Only a small share of buyers active in a given market (31 percent on average across markets) post

[18] Note that it does increase the number of buyers who actively post at least one tasks.

[19] Given Assumption 1, which implies that $(\epsilon_n^a, \epsilon_n^k)$ is independent of $(\epsilon_n^b, \epsilon_n^s)$, Σ is block diagonal.

again at least once in the subsequent three months. For sellers, this share is 66 percent. Moreover, in Section 7, we will consider the decision of current buyers and sellers to stay on the platform and find that there is very little empirical support for forward-looking anticipation of platform outcomes, although there is evidence that sellers respond to past outcomes.

We are not overly concerned with the possibility that marketing and advertising could affect both the number of buyers and sellers and their posting and search decisions. This is because, during the period of our study, the platform did not spend heavily to attract buyers and sellers. Advertising relied on articles mentioning TaskRabbit in newpapers and blogs. 40 percent of these articles were not pitched by the platform, but rather made reference to TaskRabbit while discussing the sharing economy.[20] In addition, more than 70 percent of them were on national media, as opposed to local newspapers. Finally, presence on the media was fairly uniform across months.[21] This media coverage is unlikely to be specifically tied to market conditions affecting posting or search effort. Marketing targeted at the city level occurred only for a few weeks around the time of entry into that city: TaskRabbit would start by acquiring some sellers before opening the platform to buyers, and would train them to perform services by assigning them to a small number of marketing tasks - e.g. flyer distribution. By only keeping markets with more than 50 active buyers and 20 active sellers, we are fairly confident that the TaskRabbit's marketing efforts are not the driving activity within a market.

This leaves us to consider the participation decisions of new buyers and sellers. Our basic premise is that prospective new buyers and sellers do not have much information about the specific idiosyncratic conditions on the platform. On the buyer side, there is relatively little cost of joining the platform, and we believe that during our study period adoption may have been driven significantly by people simply becoming aware that the platform existed. Our indication is that buyer sign-ups tended to increase notably after media mention, which we do not believe were tied to specific market conditions.

On the seller side, a significant source of month to month variation in new participation was driven by changes in the screening process. For a period, sellers were rigorously screened and interviewed by TaskRabbit employees. Acceptance rates of received applications depended on employees' time to conduct interviews, were usually very low (13.6 percent) and varied greatly month

[20]The sharing economy (or collaborative consumption) is the 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 platforms.

[21]The numbers rely on TaskRabbit's tracking activity of its media presence in 2012.

to month. Further, they introduced a certain delay between the sign-up decision and the actual participation on the platform.We have no evidence that they varied by city-months in response to expectations of higher demand or of lower time availability of each seller. In the spring of 2013 the platform decided to ease sellers' screening, and started to require simpler background checks and social controls (Linkedin, Facebook verification). This resulted in an acceleration of sellers' acquisitions. Together, the varying screening policies and acceptance rates led to fluctuations in the relative number of buyers and sellers for reasons arguably unrelated to individuals' activity within each city-month.

We have also assumed that pricing and matching shocks $(\epsilon_n^a, \epsilon_n^k)$ are not anticipated when buyers and sellers make their posting and search decisions. These shocks can result from unexpected concentration of offers among a small number of tasks, for example due to variation in the time when users access the platform: if all sellers in a market find themselves looking for tasks at the same time within a month, offers will tend to be sent to the same tasks, more so than if sellers search for tasks at different times. This could decrease the rate at which offers and tasks are converted into matches and the price at which matches trade. Our assumption essentially requires that buyers and sellers cannot anticipate these coordination problems.

5.4 Estimation

The estimation consists of two steps. First, we estimate the pricing and matching functions by ordinary least squares. To do that, we transform equations 5 and 6 by taking logs to obtain

$$\log Q_n = \log A_t + \log A_c + \alpha_1 \log s_n + \alpha_2 \log b_n + \log \epsilon_n^a \tag{11}$$

and

$$\log P_n = \log K_t + \log K_c + \rho_1 \log s_n + \rho_2 \log b_n + \log \epsilon_n^k. \tag{12}$$

Second, we estimate the utility parameters by method of moments using our assumption of equilibrium behavior and the orthogonality of S and B from contemporaneous effort shocks. The moment conditions are

$$E\left[x_n\left(\frac{\beta_n}{\mu\left(1-e^{-\eta q_n^b (v_n - p_n)}\right)} - 1\right)\right] = 0$$

and

$$E\left[x_n\left(\frac{\sigma_n}{\gamma b_n^\delta q_n^s (0.8 p_n - c_n)} - 1\right)\right] = 0,$$

where $x_n = (1,\ B_n/S_n,\ B_n)'$ is the three-element column vector of instruments. The model is exactly identified.

6 Results

This section presents our empirical estimates of the pricing and matching functions, and the demand and supply parameters. We use these to derive estimates of the gains from trade, and the role of labor supply elasticity in promoting efficient matching.

6.1 Matches and Prices

We start by discussing our results on the aggregate pricing and matching functions.

Table 3 presents results from ordinary least square regressions of equations 11 and 12 above. The coefficients can be interpreted as elasticities, because of the log-log specification. The first column shows that doubling the number of tasks, holding constant the number of offers, increases the number of matches by 41 percent (α_1). Similarly, doubling the number of offers, holding fixed the number of tasks, increases the number of matches by 52 percent (α_2). The estimates suggest that scaling up either tasks or offers contributes about equally to the creation of successful matches.

The sum of the two elasticities provides an estimate of the returns to scale in the matching technology. Work on two-sided platforms has emphasized the importance of increasing returns to scale for market structure (Ellison and Fudenberg (2003)). The hypothesis is that active and thick markets may lead to easier matching. In a platform like TaskRabbit where tasks typically require a buyer and a seller to meet, efficiency can come from matching buyers and sellers who live close to each other. Our estimates, however, show no evidence of increasing returns to scale. Returns are slightly (and significantly) less than constant ($\alpha_1 + \alpha_2 < 1$) when estimated by ordinary least squares, and slightly over 1 when the number of buyers and sellers are used as instruments for tasks and offers.

The absence of increasing returns was perhaps unexpected given the specific nature of the market. But interestingly it does not seem to be the case that distances between matched buyers and sellers shrink as a market grows. Figure 13 plots the median distance within a city-month, where the distance is measured between the zip codes reported by buyers and sellers.[22] The blue

[22] We compute the geodetic distance, i.e. the length of the shortest curve between the two zip codes, where the input coordinates are assumed to be based on the WGS 1984 datum. The distances are ellipsoidal distances computed using equations from Vincenty (1975).

line takes a seller who made an offer at a specific time, and pairs him to every buyer who posted tasks in the preceding 48 hours. The median distance is computed among all such pairs within a city-month. The orange line is just the pairing of tasks and their corresponding offers, and the distance is computed between the zip codes of buyers who posted those tasks and sellers who submitted those offers. The grey line is the pairing of buyers and sellers from successful matches. The figure plots the median distance for the six largest cities over time.[23] None of these measures of buyer-seller distance shrinks as a market scales up.

Table 3 reports estimates of the market pricing function, again estimated by ordinary least squares (second column). Price moves very little with the number of tasks and offers. Doubling the number of tasks, holding constant the number of offers, increases the average transacted price by 1.5 percent, while doubling offers decreases it by 1.3 percent. This is perhaps a little 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. More details are in the Appendix.

The results further confirm that the average price is invariant to market scale: the sum of the price elasticity to tasks (1.5) and to offers (-1.3) is virtually zero. The two results of constant returns to scale in the matching technology and scale invariance of price empirically confirm our earlier assumptions from Section 4.

Table 3 does not report the city and time fixed effect estimates. We will return to these estimates in Section 7 where we discuss the differences in platform success across cities.

We conclude this section with two observations. First, our identification assumptions provide an over-identifying restriction that we can test because we have assumed that both (S, B) and (σ, β) are independent of the pricing and matching errors (ϵ^k, ϵ^a). We can test the latter assumption by running a Hausman specification test where we re-estimate equations 11 and 12 using (S, B) as instruments and compare the set of OLS and IV estimates, testing for their equality. With a $\chi^2(65)$-distributed test statistic of 10.39 for the matching function and 0.15 for the pricing function, both tests fail to reject that (σ, β) are independent of (ϵ^k, ϵ^a).[24]

Our estimation of the matching function leaves very little unexplained. The R-squared of the regression (first column of Table 3) is 0.996. By construction $E(\epsilon^a_n) = 1$, and its standard deviation is only 0.08. About 50 percent of the differences between actual and predicted matches across

[23] Other cities display similar time trends.

[24] We run the original Hausman test and compare the full set of estimates, rather than just the estimates of the elasticities to tasks and offers.

markets is less than 9. The number of matches formed between tasks and offers is thus very accurately predicted by the Cobb-Douglas matching function. The amount of residual variation in the pricing function is a little higher, given a R-squared of 0.73 (second column of Table 3). However it corresponds to a discrepancy of \$3 or less in most markets.

6.2 Gains from Trade

We now turn to our estimates of the utility parameters, presented in Table 4. To discuss them, we consider a market with the median number of buyers ($B = 447$), the median relative number of buyers ($\frac{B}{S} = 3.72$), and the price and matching parameters from San Francisco in October 2013.

We consider sellers first. The estimates imply that search costs are relatively low, and, consistent with our earlier evidence, labor supply is highly elastic. Estimates of $\hat{\gamma} = 0.41$ and $\hat{\delta} = 0.25$ imply that the search costs in the median market are 25 cents for one application, \$1 for two, and, per our assumption, continue to increase at twice the rate of offers. The predicted number of offers per seller in that market is 9, corresponding to \$21 in individual search costs, and to \$4.60 in marginal search costs. We also find that search costs decrease with market size, but only slightly. Holding constant the buyer to seller ratio, an increase in the absolute number of buyers from the 25^{th} ($B = 220$) to the 75^{th} ($B = 1,085$) percentile of the distribution of market size increases the number of offers per seller from just above 8 to 10, a 22 percent increase.

Search effort is much more responsive to price and the expected match rate than it is to market size (Figure 8). The elasticity of search effort to the offer match rate q^s is equal to one by assumption: a doubling in the offer match rate doubles search intensity. Effectively, this implies that the elasticity of tasks supplied is twice as large. Doubling the offer match rate doubles the number of offers submitted, and, because now each offer is twice as likely to be accepted, each seller works fours times as hard. The elasticity of search effort to the price is equal to the inverse of the seller markup. We estimate the seller cost of performing a task to be $\hat{c} = \$33$.[25] In the median market, sellers are paid \$48.50, so the supply elasticity to price is equal to 3.21. Even a \$5 price increase can raise the number of offers per seller from 9 to 12.

Lastly, we find that sellers receive relatively little surplus from completed tasks. The per-task profit is \$48.50 - 33 = \$15.50 in the median market. Given that an offer is expected to match with

[25] In principle, it can happen that the seller opportunity cost c_n in a market be smaller than the realized average price paid to sellers. However, c_n should always be lower than the expected price reduced by the commission fee $0.8p_n$ in order to rationalize a positive number of submitted offers. That is the case in our estimates. Analogously on the buyer side, p_n should always be (and in fact always is) lower than the buyer value for the task v_n.

30 percent probability and each seller submits 9 offers, the search cost per completed task is \$7, or half the task profit. Even if seller search costs are low in absolute dollars, they represent a large share of the per-task profit, and a seller ex-ante expected surplus from the median market is \$21.[26]

Next we turn to the buyer side.[27] We find that the mean arrival rate of tasks is $\hat{\mu} = 1.23$. Tasks are cheap to post, with an average cost equal to $1/\hat{\eta} = 50$ cents, and the buyer value from task completion is $\hat{v} = \$70$. Given the small cost of posting tasks, in the median market, all needs are posted, so $\beta = 1.23$. In this market, tasks are successfully matched with 61 percent probability, and buyers pay \$61 for each completed task.

Task demand is inelastic. At equilibrium in the median market, the elasticity of buyers' posting rate to the task match rate is basically zero (0.0002), and the elasticity to price is similarly small (-0.0015). Given the low cost of posting, Figure 9 shows that buyers are going to post all their needs for most equilibrium values of match rate and prices. In a standard setting, low elasticity of demand would imply that buyers receive a large share of the surplus. However, a seller willingness to pay for the average task, at \$70, is not much higher than the price she actually pays, and a buyer ex-ante expected surplus from the median market is \$6.

Combining our estimates, we find that the gains from trade for each task are \$34, without factoring in seller search costs (\$7 per completed task) and buyer posting costs (less than \$1). This seems quite plausible given the nature of most tasks, and combined with an elastic supply it means that maintaining a relatively efficient matching process is as crucial for the platform success as attracting a large number of buyers. We discuss the efficiency of the matching process in the next subsection, and buyers' growth in Section 7.

We conclude this section by examining the fit of our model. Figure 10 compares the actual aggregate number of posted tasks and submitted offers with those predicted by our model for the city of San Francisco. Other cities are in Appendix C. Overall the model does a good job at tracking task and offer activity over time for each city, albeit in some cities supply is consistently underestimated (as in Boston) or overestimated (as in New York). The discrepancies are typically

[26] This is possibly consistent with a young peer-to-peer platform, where the average seller only performs a few tasks to fill in his schedule. Because we focus on market averages, we do not consider the more professional sellers, those who perform tasks on a more regular basis.

[27] On the buyer side, the estimation is complicated by the fact that variation in the posting rate is very limited. This implies that the parameter estimates for the mean task arrival rate μ, the cost distribution parameter η, and the value for each task v have to be such that the response of β to changes in the task match rate and price have to be small. Using equation 3, this requires that both $\frac{\partial \beta}{\partial q^b} = \mu f\left(q^b(v-p)\right)(v-p)$ and $\frac{\partial \beta}{\partial p} = \mu f\left(q^b(v-p)\right)(-q^b)$ be small. The second partial derivative is close to zero only if the density $f\left(q^b(v-p)\right)$ is small, given that both the task arrival rate and the task match rate are not small. The buyer net benefit $v-p$ (divided by q^b) is equal to the ratio of the two partial derivatives, which are both close to zero. This limits our ability to easily estimate v.

less than 25 percent relative to the actual value, with the only notable exception being the last fews months in New York.

6.3 Benefits of an Elastic Labor Supply

When the cost of posting tasks is small, fluctuations in the number of buyers translate in proportional fluctuations in the number of posted tasks, where the coefficient of proportionality is the mean arrival rate of needs. In addition, when buyers have low willingness to pay for tasks, and gains from trade are relatively modest, the range of price adjustment is very limited. Together, these two results imply that the market can clear in only one of two ways: through an elastic labor supply, or through buyers' rationing. We found the first to be the dominant equilibrating mechanism on TaskRabbit, and in order to evaluate its benefits, we compare it to the second alternative.

We start with a simple exercise to illustrate the intuition. 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 doubling of the number of offers, and would create 930 total matches.[28] 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 size 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: when demand is low, it is easier for buyers to find a match, 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 each task would match with a 73 percent probability. In the high-demand market (2,000 tasks), 649 tasks would be matched, implying a 32 percent match rate.

Overall, if we compare the total number of matches between the two scenarios, the platform with an elastic labor supply is able to create 11 percent more matches.[29] This is evidently optimal from the platform 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 percent. Since it also raises retention, it benefits the platform by accelerating its growth. Equilibration through seller effort is also optimal from the buyers' perspective. Given that the cost

[28] Note 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 3).

[29] The 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 $s' > s$.

of listing a task is low enough that rationing does not prevent them from posting them in the short run, effectively increasing the number of matches by 11 percent raises buyers' surplus by the same percentage. An elastic supply creates more matches but it is also more costly, given that sellers' costs increase in their search effort at an increasing rate. Again we can use the cost estimates from our model in the simple example above, assuming that offers come from 200 active sellers (so that in the market with 1,000 tasks and 1,400 offers each seller submits 7 offers). The total search costs in the first scenario, where supply fully adjusts to accommodate demand, are 54 percent higher than in the second scenario, where supply is fixed at 7 offers per seller.

We now apply this intuition to our context. To do so, we consider all 336 markets and simulate interaction among buyers and sellers under two scenarios: the first considers the labor supply elasticity directly estimated from the model, which implies that the market predominantly equilibrates through seller effort; the second fixes individual supply at 7.58, the average number of offers per seller across our markets, which implies that the market will equilibrate through buyer rationing.

We start by measuring the aggregate value of matches created under the two scenarios, which is given by the following formula: $\sum_{n=1}^{336} M(B_n\hat{\beta}_n, S_n\hat{\sigma}_n)(\hat{v}_n - \hat{c}_n)$. A flexible supply allows a 15 percent increase in the value of matches created. The increase in buyer and platform surplus is analogous for two reasons: prices do not adjust much, and buyers' posting costs are small. So measuring the platform's aggregate revenue as $\sum_{n=1}^{336} 0.2\hat{p}_n M(B_n\hat{\beta}_n, S_n\hat{\sigma}_n)$ results in a 15.5 percent increase in revenue relative to an inelastic supply. Buyers' aggregate surplus, or $\sum_{n=1}^{336} M(B_n\hat{\beta}_n, S_n\hat{\sigma}_n)(\hat{v}_n - \hat{p}_n) - E\left(\eta | \eta \leq \hat{q}^b(\hat{v}_n - \hat{p}_n)\right) B_n\hat{\beta}_n$, also increases by the same percentage. Sellers' aggregate surplus is however reduced by their elasticity, due to the increasing search costs. Relative to providing a constant level of effort, sellers' aggregate surplus, defined as $\sum_{n=1}^{336} M(B_n\hat{\beta}_n, S_n\hat{\sigma}_n)(0.8\hat{p}_n - \hat{c}_n) - \frac{1}{2\hat{\gamma}(B_n\hat{\beta}_n)^{\hat{\delta}}} S_n\hat{\sigma}_n^2$, is reduced by 6 percent.

If the platform were able to reduce, or even eliminate, seller search costs, the benefits of an elastic supply would be large and positive for sellers as well. We discuss how the platform can achieve this in the conclusions. We now turn to discuss how market efficiency and growth differ by city.

7 Platform Growth and City Heterogeneity

A notable feature of the data is that some cities exhibit striking growth in participation, and others exhibit more moderate growth (Figure 2).

In principle, two types of theories can explain differences in how cities attract and retain a large number of users. The first type relies on scale economies and strategic complementarities between the adoption patterns of buyers and sellers. If market frictions were reduced by market scale, we would expect that cities which started off with a large user base grew much faster than cities of modest size, exactly because growth led to more growth. However in Section 6 we estimated only moderate scale economies. Therefore 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 more efficient markets and that markets vary greatly in their matching efficiency, summarized by the fixed effects of equation 11. Combining these results, we see a strong relationship between the rate at which tasks and offers are converted into successful matches and city growth rates. 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 specificity.

7.1 City Differences in Growth

Platform growth is a combination of adoption and retention of existing users. Given that supply is so elastic that buyers do not have a considerably harder time finding matches when abundant, and given that active buyers seem to post on average the same number of tasks every month, growth depends on buyers' participating decisions.

Figure 11 plots buyer adoption and retention separately for the 10 largest cities. The left-hand side panel shows the number of new buyers, in log scale, over time. A buyer is defined as new in a city-month if she posts her first task in that city during that particular month. Buyers adopt the platform at a linear rate, different in all cities. At visual inspection, this rate seems to be correlated with the city-specific retention rate. For every city, the right-hand side of Figure 11 plots the share of active users in a month who posted again at least once task in the following three months. San Francisco is successful at both attracting new buyers and retaining current ones, while Philadelphia has both lower adoption and retention rates.

The literature on innovation diffusion (Young, 2009) has focused on three types of mechanisms leading users to adopt new technologies: network effects, technology improvements, and information diffusion. The first two assume that different users adopt at different points in time because of heterogeneous benefits: early adopters have a large intrinsic value from a new technology, while late adopters join because of scale economies or technical upgrades. We have argued that platform efficiency does not increase with market scale, and for the period under consideration TaskRabbit did not implement major platform changes. Word of mouth and information diffusion, then, seem to be the most plausible alternative in this context, and cities can differ both in the rate at which information spreads and the rate of take-up conditional on receiving that information. For example, in San Francisco adoption might be fast because people there are eager to experiment with new technologies and because current users spread the information at a faster rate, with the second factor possibly driven by a positive experience on the platform. We measure the aggregate effect by estimating the city-specific growth rate:[30]

$$new_{tc} = \phi_c age_{tc} + \nu_{tc} \ , \tag{13}$$

where new_{tc} is the number of new buyers joining city c in calendar month t, and age_{tc} is the age of the platform in city c at time t. For example, $age_{tc} = 1$ if month t is the first since TaskRabbit became active in city c. In Appendix D, Table A1 shows the results, and we verify that deviations from the linear adoption rate are not driven by contemporaneous market conditions, in support of our earlier identification assumptions.

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

$$\log\left(\frac{stay_{tc}}{1 - stay_{tc}}\right) = \theta_0 X_{tc} + \eta_t + \eta_c + \nu'_{tc} \ . \tag{14}$$

$stay_{tc}$ is the share of users active in city-month t, c who were active again at least once in the following three months within the same city. X_{tc} is two-element vector of relevant outcomes in city-month t, c: realized buyer match rate and average transacted price. We expect that a high match rate would increase the odds that a buyer will be active again in the next three months, while a high price would drive away more buyers. As with equation 13, Table A3 in Appendix E shows the results, which confirm our hypothesis, and we verify that retention is not driven by

[30]We assume a linear growth rate different across cities, given Figure 11. It can be rationalized within the Bass model of new product diffusion (Bass, 1969): $new_{tc} = \phi_c + new_{t-1,c}$, where new_{tc} is the number of new buyers joining city c in calendar month t. Two things differ from the standard specification. First, the total number of potential adopters is assumed to be large relative to the platform size, which is consistent with the population size of the metropolitan cities relative to the current users on TaskRabbit. Second, we assume that new adopters in the previous month are the only users spreading information, and not adopters of previous months. Each new adopter diffuses information so that exactly one extra adopter joins the platform in the following month.

expectations on future outcomes, in support of our earlier identification assumption.

Figure 14 plots the estimates of ϕ_c (city-specific adoption rate) and η_c (retention rate) from equations 13 and 14. 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 on both measures, while Houston, Atlanta and Phoenix lag behind on both. However, in San Diego buyers adopt at a fast rate but are also likely to leave the platform, while in Portland new buyers are just a few but they stay longer. Retention is arguably the decision that is mostly related to the experience on the platform, and indeed in the next section we show that it is associated with how efficiently the platform matches buyers and sellers in each city.

7.2 City Differences in Match Efficiency and Market Thickness

For tasks like cleaning and delivery, it is obvious to expect that buyers would care about how easy it is to find a seller willing to provide the service at the desired time and location, and the price to pay for the service. In this section we explore how differently cities perform in this respect. To do so, we take advantage of earlier estimates of the matching and pricing functions (equations 11 and 12).

Cities vary widely in the rate at which tasks and offers are converted into successful matches. Figure 12 plots estimates of A_c from equation 11, ordering cities from the most efficient (San Francisco) to the least efficient (Miami). San Francisco is 2.37 ($\frac{A_{SF}}{A_{Miami}}$) times as effective as Miami in creating matches: out of the same number of tasks and offers, 100 matches are created in Miami while 237 matches are created in San Francisco. Other than San Francisco, among the cities with the highest match productivity are Boston, Portland, Austin, and New York City. The other extreme includes Miami, Denver, Phoenix, Philadelphia, and Atlanta.

There is limited heterogeneity in prices across cities, as Figure A3 shows. Here cities are ordered according to their ranking in the match efficiency parameter from Figure 12, and the plot displays estimates of K_c from equation 12.[31] Most prices range between \$54 and \$65, with Denver (\$42) and Miami (\$68) as outliers.[32]

[31] $K_{Oct2013}$ is normalized to 1.

[32] The two equations 11 and 12 also include time effects, which we briefly discuss here. Over time, there is a limited decline in match efficiency, but not sizable nor statistically significant. Figure A2 plots estimates of A_t from the matching equation. The line is fairly flat between Spring 2011 and Summer 2013, with higher variability before and a small downward jump afterwards. This variation coincides with the staggered entry of TaskRabbit into the various cities. Indeed, prior to Spring 2011, only two cities were active, San Francisco and Boston. Around the Summer of 2013, TaskRabbit became active in the nine youngest cities, which are also those with lower match productivity estimates. Instead, there is a substantial increase in transacted prices over time. As Figure A4 shows, estimates of K_t

The most efficient cities are those able to retain the most buyers.[33] In Figure 12 the size of the marker for the match efficiency parameter is proportional to the retention rate η_c estimated from equation 14. The cities with high match productivity A_c also have high retention rate η_c.

The next step is to try to understand what can explain efficiency differences across geography. To this purpose we look at 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. This hypothesis is strongly supported by the data. Figure 15 plots the match efficiency parameter A_c and the median geographic distance between buyers and sellers of paired tasks and offers.[34] In cities like San Francisco, Boston, Portland, and New York, which convert tasks and offers into matches at the highest rate, the median distance between buyers and sellers of paired tasks and offers is around 7 miles. At the other extreme, the distance in Philadelphia and Miami is over 20 miles.

The second candidate measure of market thickness is related to task specificity. The idea is that an idiosyncratic task, which possibly requires specialized skills on the part of the seller, is harder to match than a standardized cleaning task, for which all that matters is location and time availability of one seller out of many good alternatives. To explore this idea, we look at the share of tasks posted in May 2014 within the top five categories (Shopping and delivery, moving help, cleaning, minor home repairs, and furniture assembly). Figure 16 plots this share of more "standard" tasks against the match efficiency parameter A_c. In San Francisco, Boston, and New York over 60 percent of the tasks are posted within the top five categories, while Dallas, Miami, Atlanta, and Denver all have shares smaller than 50 percent.

8 Conclusions

In this paper, we have studied the problem of balancing highly variable demand and supply. This is a basic problem of the recently popular online peer-to-peer marketplaces for local and time-sensitive services, such as Uber, Airbnb, and TaskRabbit. We have presented a static model of a

monotonically increase, from $30 in June 2010 to just above $60 in May 2014. The increase in price seems to closely track the task diversification on the platform towards more expensive tasks. Figure A5 presents the share of posted tasks in the 10 largest categories over time, combining all other categories in an eleventh group. It demonstrates how the period of fastest diversification, where the cumulative share of the top categories fell considerably, occurred in the Spring of 2011, exactly when the average price experienced the highest increase.

[33] The most efficient cities also tend to correspond to those where TaskRabbit entered earlier on, and this still holds under different specifications of the matching function, as shown in the Appendix.

[34] The correlation is maintained with the 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.

frictional market where buyers and sellers post requests and offers for services. Our model specifies the possible mechanisms through which the market clears in terms of the elasticity of demand and supply to price, as in standard product markets, and to rationing, as in frictional labor markets.

We have applied the theoretical model to study TaskRabbit, a growing platform for domestic tasks. We estimated utility, matching, and pricing parameters using variation in the number of buyers and sellers across cities and over time. The empirical application has allowed us to measure the gains from trade facilitated by the platform, the particular mechanism that equilibrates the market (highly elastic supply), and how market efficiency varies with location and market size.

The natural level of market efficiency is modest, although certain cities, such as San Francisco and New York, are largely more efficient than others. The efficient cities are also those which grow the fastest, by attracting new users and retaining existing ones at a higher rate. City differences in the efficiency with which tasks and offers are converted into successful matches seem to be related to at least two sources of frictions: geographic distance between buyers and sellers, and task specificity.

The market is able to efficiently accommodate fluctuations in buyers and sellers thanks to a highly elastic labor supply. When demand is capped by an exogenous arrival process of needs for cleaning or furniture assembly, and sellers flexibly adjust their effort in response to changes in relative demand, buyers find it profitable to post all their tasks, and can match at the same rate and price regardless of the number of sellers present in the market.

On TaskRabbit, the elasticity of supply is likely due to the fact that sellers might be available to perform tasks only within a defined time window in any given month (say, every Saturday) or within a few miles from their house. A doubling (resp. halving) of demand would thus imply a doubling (halving) of the profitable opportunities for each individual seller, thereby affecting their offer submission. However, it is costly for sellers to search through posted tasks, and with some probability profitable tasks are not found by any seller and are left unmatched, or are found by sellers who turned out to be unavailable to perform them. Indeed, 13 percent of the tasks that did not receive any offers were canceled because the buyer reported having the task done sooner, and 19 percent of the tasks that did receive offers were canceled because one of the parties reported inability to resolve scheduling conflicts. So, what would happen if instead of having sellers search for profitable submissions, the platform let sellers directly list their availability on a calendar and within a specific geographic area? This would eliminate search costs of supply, or at least make them independent on the size of demand, and would raise seller welfare while benefiting buyers and the platform by raising task match rates.

The benefits of reducing sellers' search costs and improving match efficiency provide a rationale for a platform re-design, which actually occurred in the Spring of 2014. Figure 17 is a screenshot of the current platform. Buyers can now select the category, location, and time for a given task request, and then either choose among the sellers available to perform that task type at the specified time and location (similar to the auction mechanism prior to the change), or have the platform automatically choose for them (similar to the posted price mechanism).

We can actually also use our model estimates on TaskRabbit in a broader perspective, and roughly calculate the value generated by online peer-to-peer markets for domestic tasks. This includes TaskRabbit, as well as other platforms, such as Craigslist, ThumbTack, or Handy. We estimated the value created in each match to be $37, and a monthly average number of tasks per buyer equal to 1.23. Taking the number of households present in the US cities where TaskRabbit is currently active, and assuming that 20 percent of them will be outsourcing domestic tasks online in the long run, an industry able to match 80 percent of those tasks would generate $920 million in value, and this in 18 cities alone.

The advantage of a highly flexible supply is appealing for other peer-to-peer platforms. Consider Uber for example, the fastest growing ride-sharing marketplace. Matches between people wanting a ride and drivers are even more local and time-sensitive that on TaskRabbit. A person at the airport is likely to use alternative transportation if it takes long to find an Uber car (rationing) or the price is too high. Uber relies on having enough cars on the road to ensure reliability during normal and busy times, and achieves this by adjusting its price, balancing buyers' response to use substitute services and sellers' willingness to provide more rides. An elastic supply would require a lower price increase to adjust its effort and accommodate demand, thus limiting the buyers' response to request fewer rides.

The same delicate balancing of demand and supply for short-term accommodation occurs on Airbnb. In this case, in order to satisfy demand for accommodation the platform relies on a wide inventory of available listings to adapt to normal and high traveling seasons.

Our paper has primarily focused on the short-term balancing of demand and supply. A valuable avenue for future research, facilitated by the large availability of data, would be to study dynamic participation decisions of buyers and sellers in more detail than done here. Better understanding the drivers of user adoption and retention can help explain platform competition, both between multiple peer-to-peer marketplaces and between the online marketplace model and the more traditional service providers.

References

[1] Patricia M Anderson and Simon M Burgess. Empirical matching functions: estimation and interpretation using state-level data. *Review of Economics and Statistics*, 82(1):93–102, 2000.

[2] Patrick Bajari and Ali Hortacsu. Economic insights from internet auctions. *Journal of Economic Literature*, 42(2):457–486, 2004.

[3] Frank M. Bass. A new product growth for model consumer durables. *Management Science*, 15(5):215–227, 1969.

[4] Olivier Blanchard and Jordi Galí. Labor markets and monetary policy: a new keynesian model with unemployment. *American economic journal: macroeconomics*, 2(2):1–30, 2010.

[5] Peter A Diamond. Aggregate demand management in search equilibrium. *The Journal of Political Economy*, pages 881–894, 1982.

[6] Michael Dinerstein, Liran Einav, Jonathan Levin, and Neel Sundaresan. Consumer price search and platform design in internet commerce, 2014.

[7] Darrell Duffie and Yeneng Sun. The exact law of large numbers for independent random matching. *Journal of Economic Theory*, 147(3):1105–1139, 2012.

[8] Liran Einav, Chiara Farronato, Jonathan Levin, and Neel Sundaresan. Sales mechanisms in online markets: What happened to internet auctions? *NBER Working Paper*, (w19021), 2013.

[9] Glenn Ellison and Drew Fudenberg. Knife-edge or plateau: When do market models tip? *The Quarterly Journal of Economics*, 118(4):1249–1278, 2003.

[10] Andrey Fradkin. Search frictions and the design of online marketplaces. *Working Paper*, 2014.

[11] Robert E Hall. Employment fluctuations with equilibrium wage stickiness. *American economic review*, pages 50–65, 2005.

[12] John Joseph Horton. Misdirected search effort in a matching market: causes, consequences and a partial solution. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 357–357. ACM, 2014.

[13] Kenneth L Judd. The law of large numbers with a continuum of iid random variables. *Journal of Economic theory*, 35(1):19–25, 1985.

[14] Pascal Michaillat. Do matching frictions explain unemployment? not in bad times. *The American Economic Review*, 102(4):1721–1750, 2012.

[15] Pascal Michaillat and Emmanuel Saez. A model of aggregate demand and unemployment. Technical report, National Bureau of Economic Research, 2013.

[16] Paul Milgrom and John Roberts. Comparing equilibria. *The American Economic Review*, pages 441–459, 1994.

[17] James D Montgomery. Equilibrium wage dispersion and interindustry wage differentials. *The Quarterly Journal of Economics*, pages 163–179, 1991.

[18] Dale T Mortensen. The matching process as a noncooperative bargaining game. In *The economics of information and uncertainty*, pages 233–258. University of Chicago Press, 1982.

[19] Chris Nosko and Steven Tadelis. The limits of reputation in platform markets: An empirical analysis and field experiment. *Working Paper*, 2014.

[20] Amanda Pallais. Inefficient hiring in entry-level labor markets. *American Economic Review*, Forthcoming.

[21] Geoffrey G Parker and Marshall W Van Alstyne. Two-sided network effects: A theory of information product design. *Management Science*, 51(10):1494–1504, 2005.

[22] Barbara Petrongolo and Christopher A Pissarides. Looking into the black box: A survey of the matching function. *Journal of Economic literature*, 39(2):390–431, 2001.

[23] Christopher A Pissarides. Short-run equilibrium dynamics of unemployment, vacancies, and real wages. *The American Economic Review*, pages 676–690, 1985.

[24] Christopher A Pissarides. Equilibrium unemployment theory. 2000.

[25] Paul Resnick and Richard Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. *Advances in applied microeconomics*, 11:127–157, 2002.

[26] Margaret Stevens. New microfoundations for the aggregate matching function, with empirical and theoretical implications. 2002.

[27] Thaddeus Vincenty. Direct and inverse solutions of geodesics on the ellipsoid with application of nested equations. *Survey review*, 23(176):88–93, 1975.

[28] H Peyton Young. Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *The American economic review*, pages 1899–1924, 2009.

Figure 1: *The TaskRabbit Platform between January 2009 and May 2014*

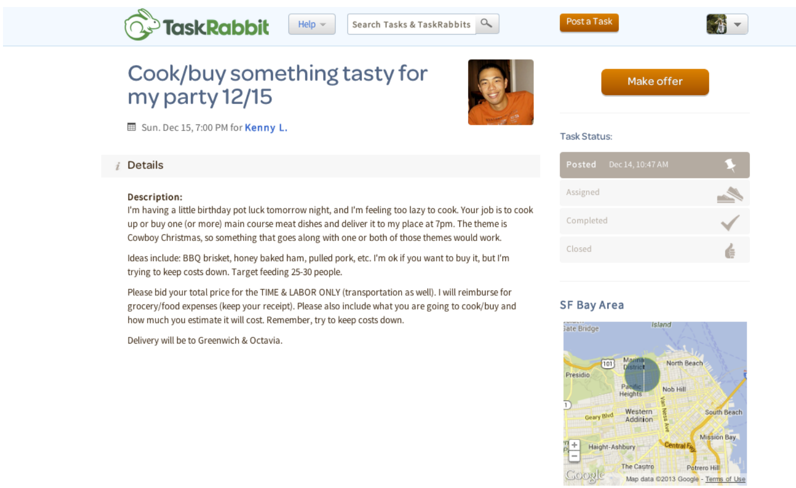

(a) *Example of a posted task.*

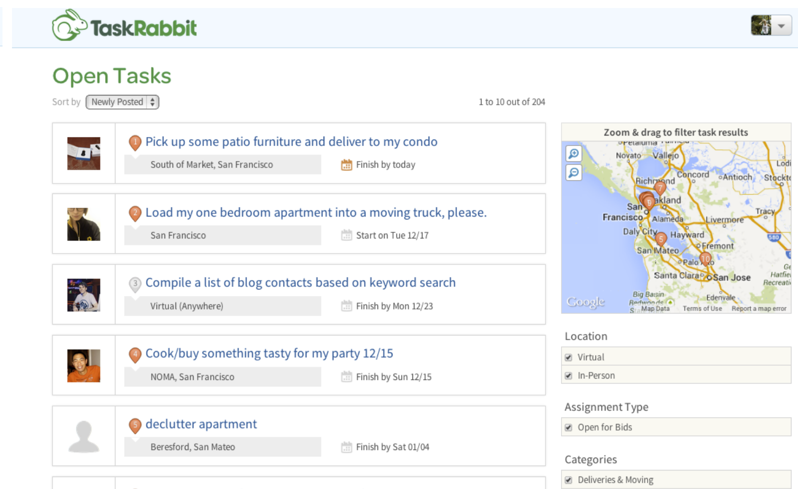

(b) *City list of posted tasks for sellers to search.*

Screenshots from TaskRabbit (accessed on December, 14th, 2013). The screenshots display information available to sellers about posted tasks.

Figure 2: *Market Size*

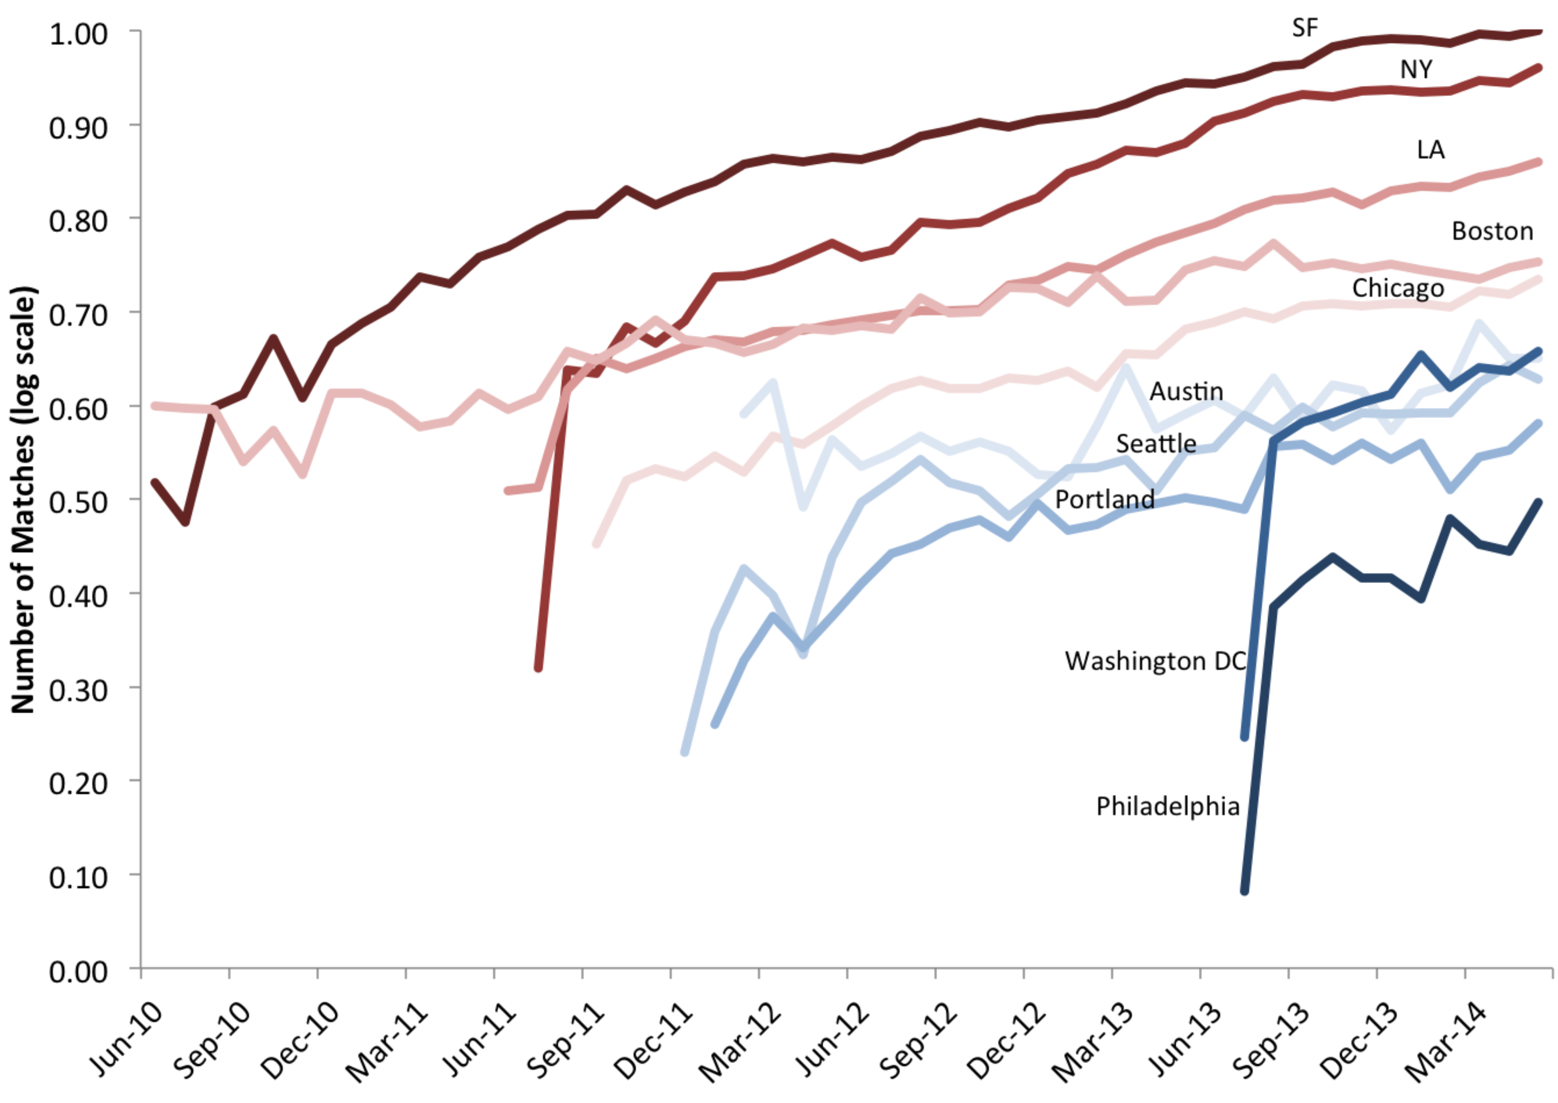

Number of successful matches in selected cities over time. The y-axis, in log scale, is normalized by the value in San Francisco in May 2014.

Figure 3: *Buyer to Seller Ratio*

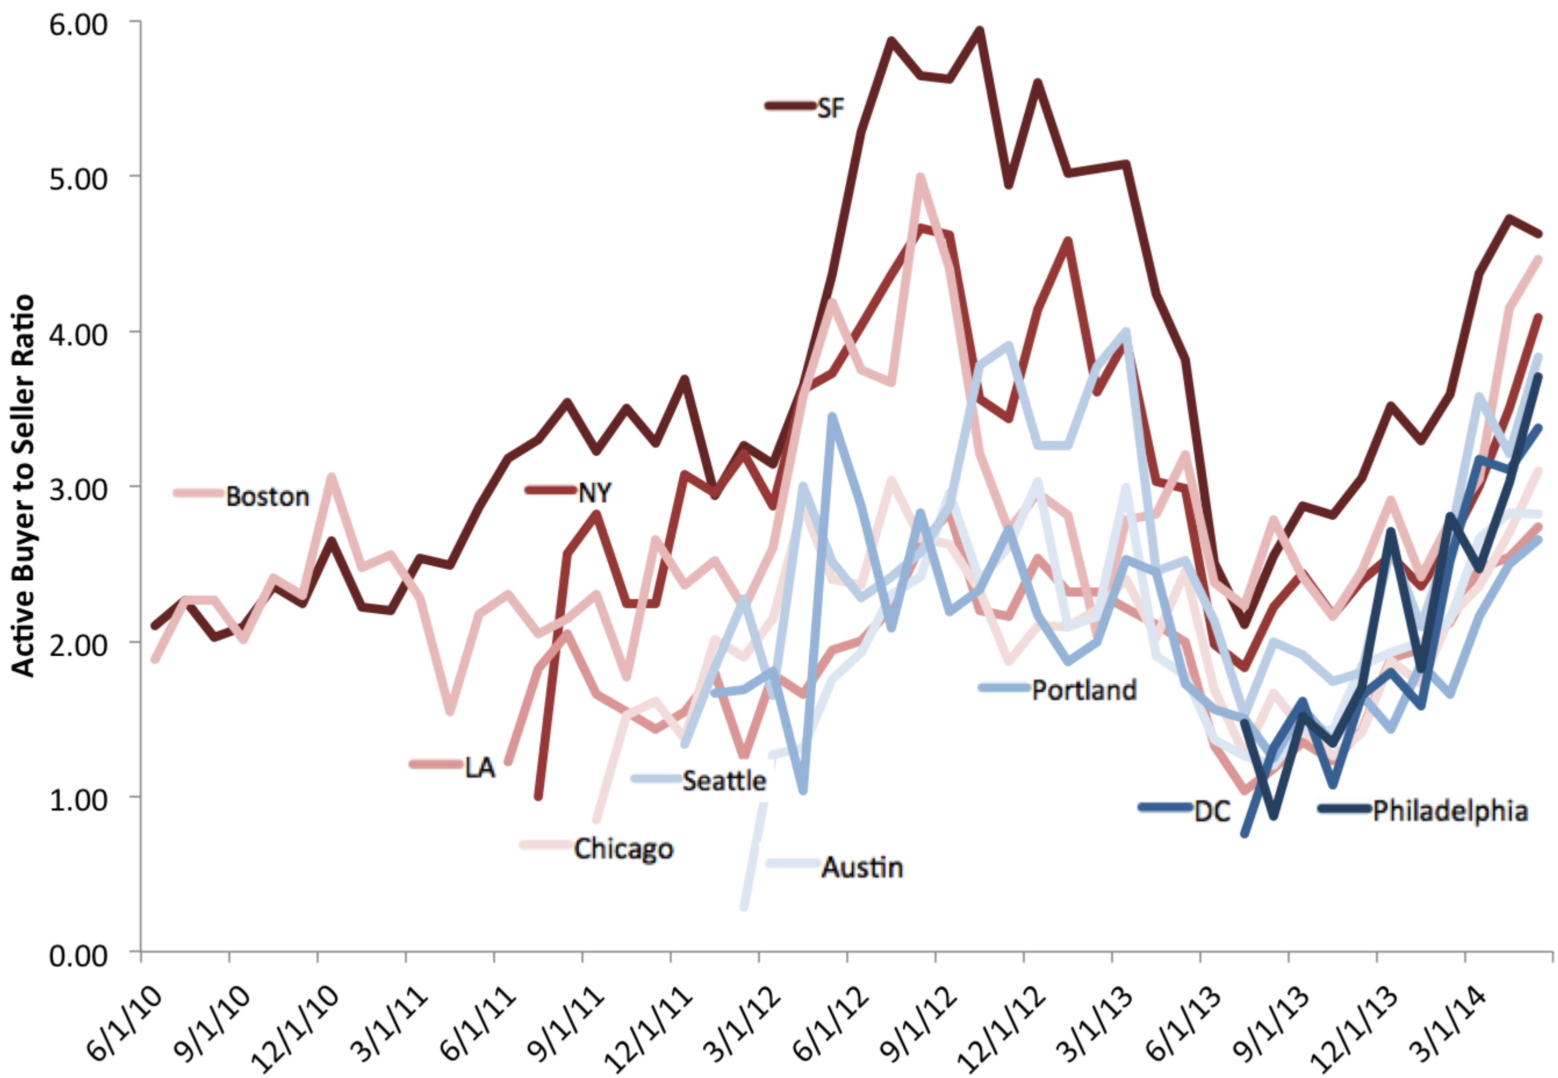

Active buyer to seller ratio in selected cities over time. A buyer is active in a city-month if he posts at least one task within the city-month. A seller is active if he submits an offer to one of tasks posted within that city-month.

Figure 4: *Tasks per Buyer and Offers per Seller*

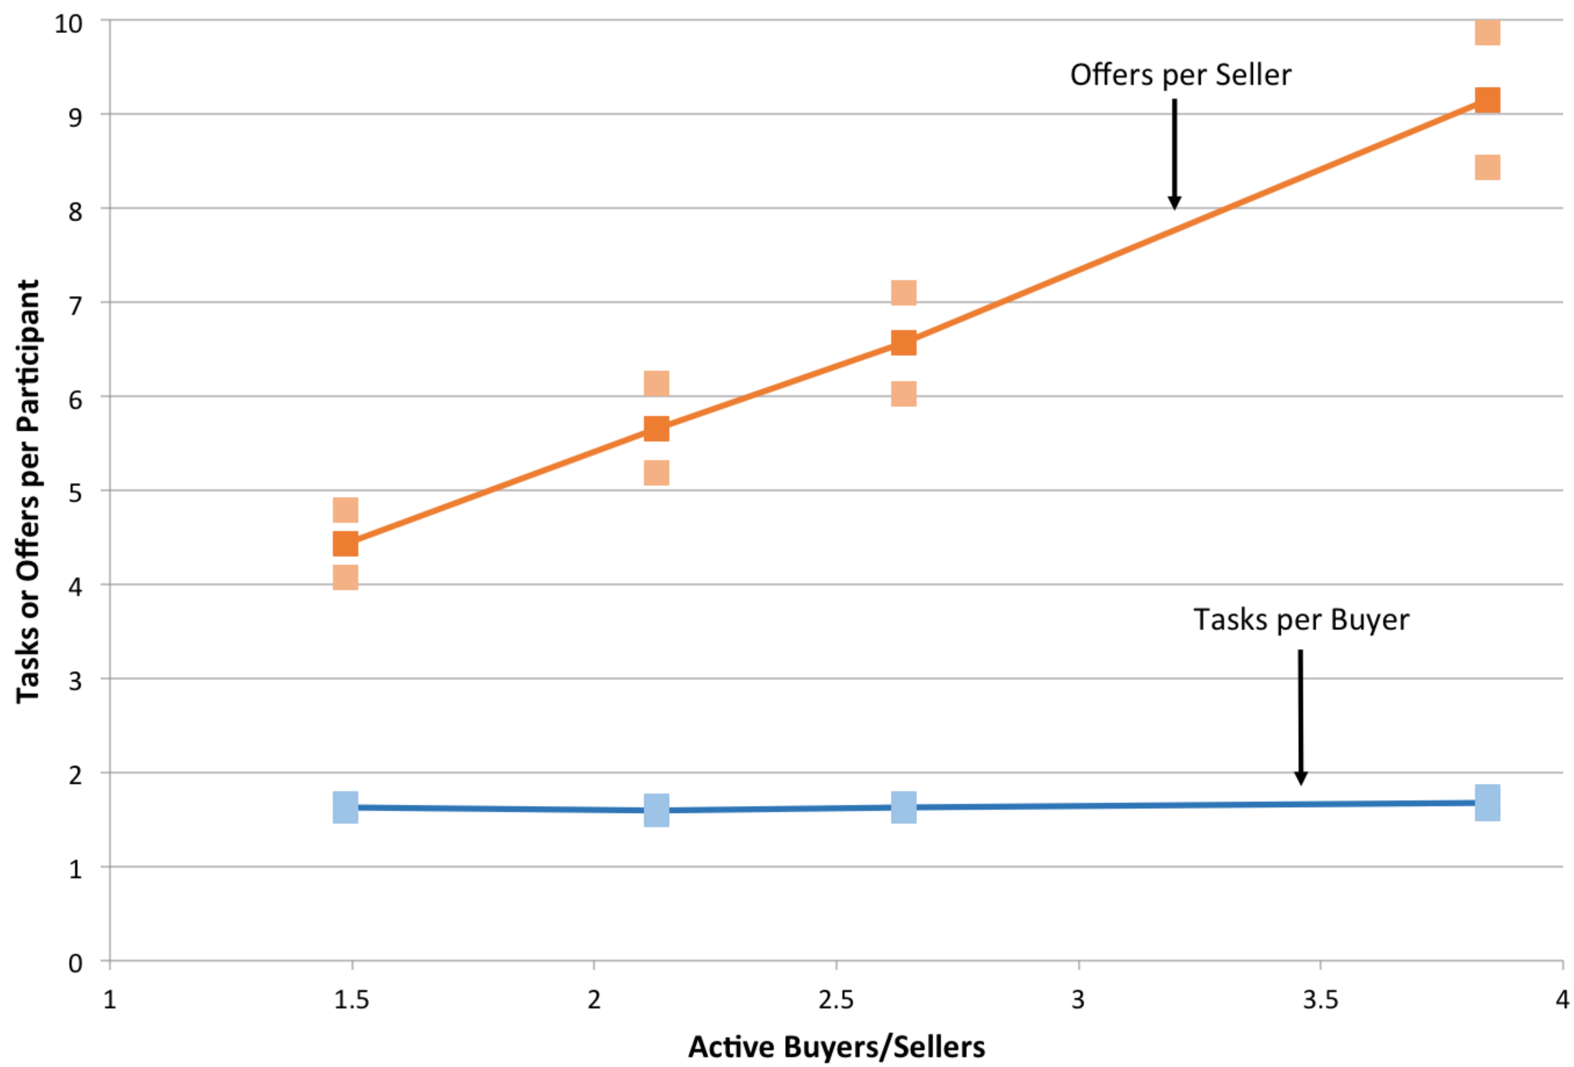

City-month average of tasks per buyer and offers per seller. We divide the 336 city-months into four groups, corresponding to the four quartiles of the distribution of the buyer to seller ratio. For each group we compute the average number of tasks per buyer (blue) and offers per seller (orange). Confidence intervals are displayed in lighter colors.

Figure 5: *Task Match Rates and Prices*

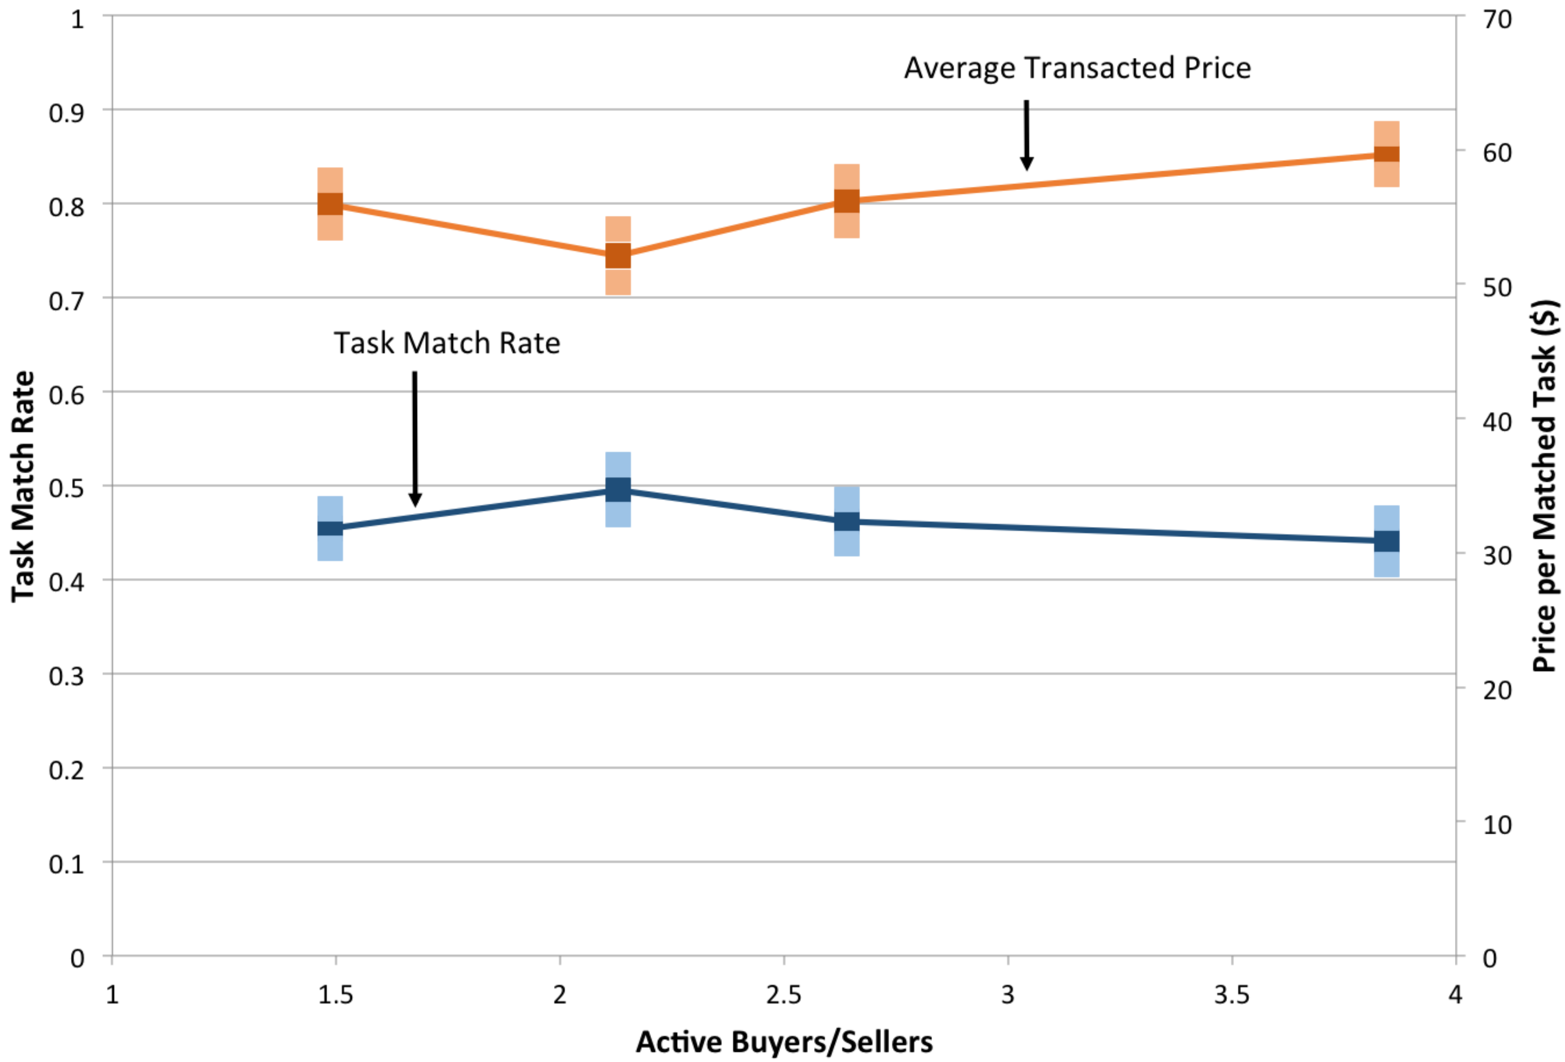

City-month average of task match rate and transacted price. We divide the 336 city-months into four groups, corresponding to the four quartiles of the distribution of the buyer to seller ratio. For each group we compute the average share of completed tasks (blue) and price paid (orange). Confidence intervals are displayed in lighter colors.

Figure 6: *Seller's Optimal Choice.*

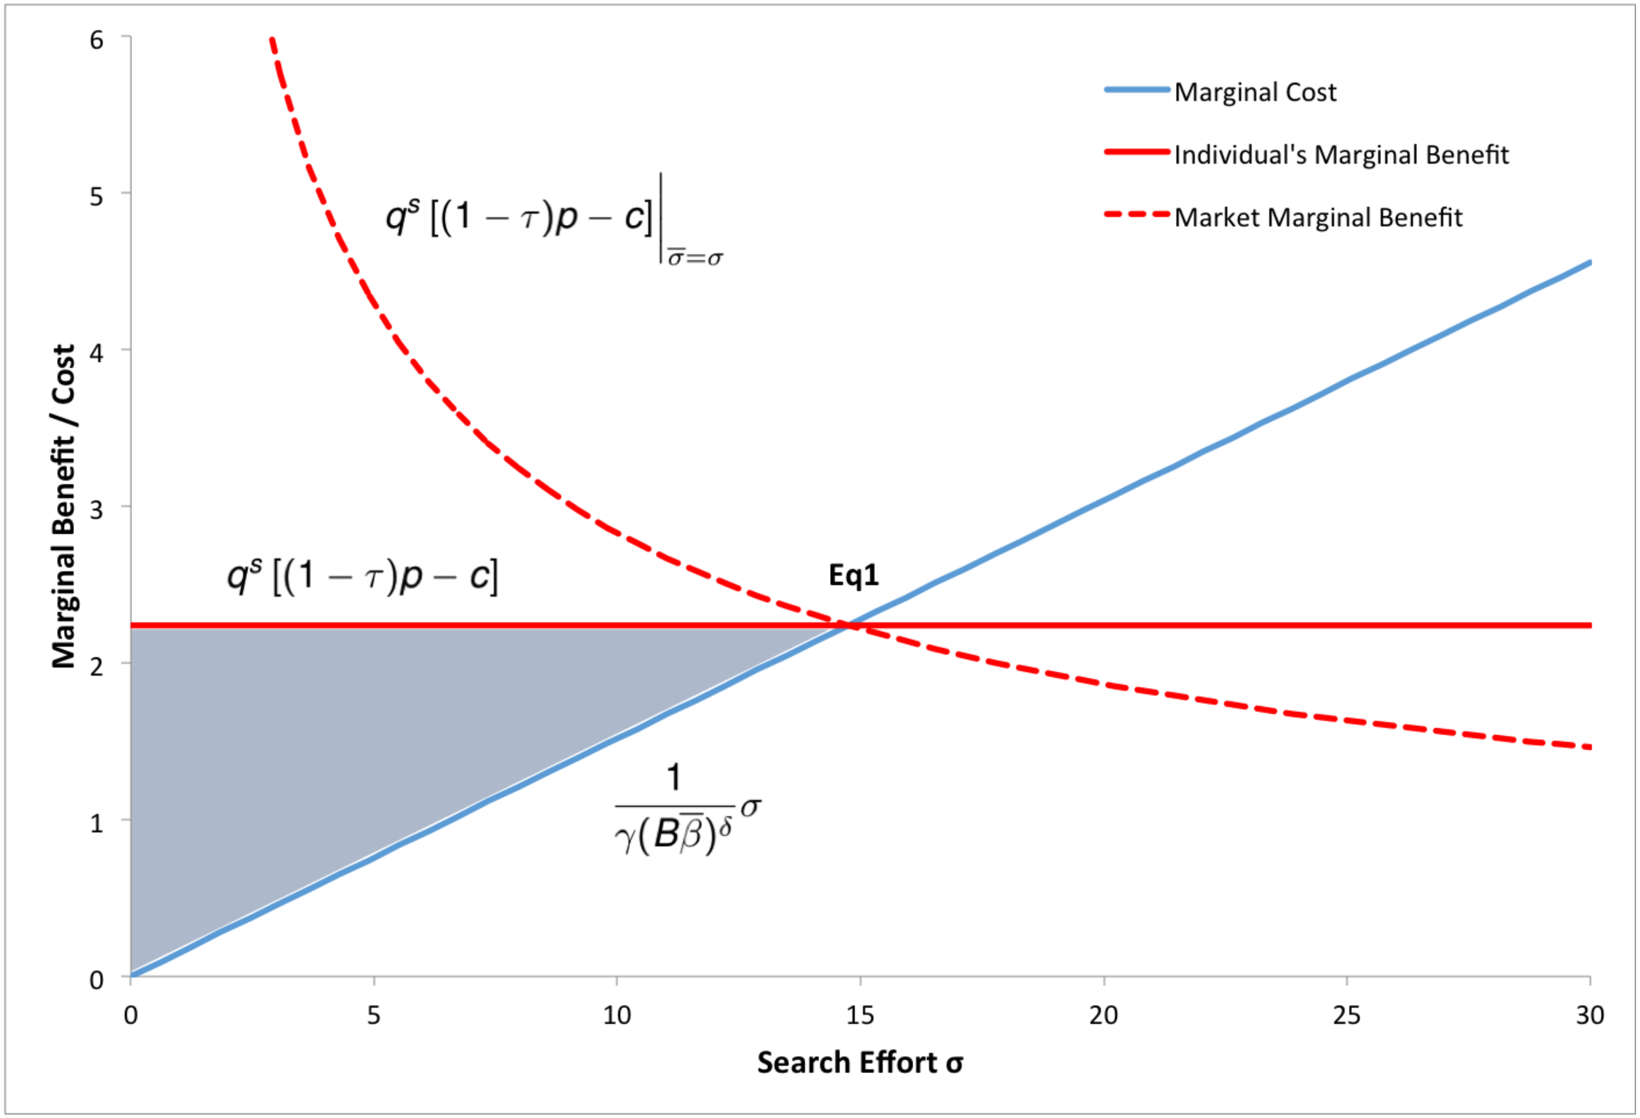

The figure plots a seller's optimal choice of search effort, which is equivalent to the expected number of offers submitted. The blue line corresponds to his marginal cost curve $\frac{\sigma}{\gamma(b)}$. The solid red line is his individual level marginal benefit curve $q^s\left[(1-\tau)p - c\right]$, where q^s and p are determined by the average search intensity $\overline{\sigma}$ and posting rate $\overline{\beta}$ of all sellers and buyers in the market (and cannot be affected by a single seller). The dotted red line corresponds to the market level marginal benefit curve: here, the offer match rate and prices are determined by the average posting rate $\overline{\beta}$, and an average search intensity equal to σ (x-axis). Together, all sellers can affect the probability that each offer is matched and the price they receive: an increase in search intensity by all sellers decreases both the probability that any one offer is accepted and the price received. Thus the market-level marginal benefit curve is decreasing in σ. Consistency requires that $\sigma = \overline{\sigma}$, i.e. that the market and individual level marginal benefit curves cross the marginal cost curve at the same point (Eq1). The area shaded in grey is the seller surplus.

Figure 7: *Comparative Statics.*

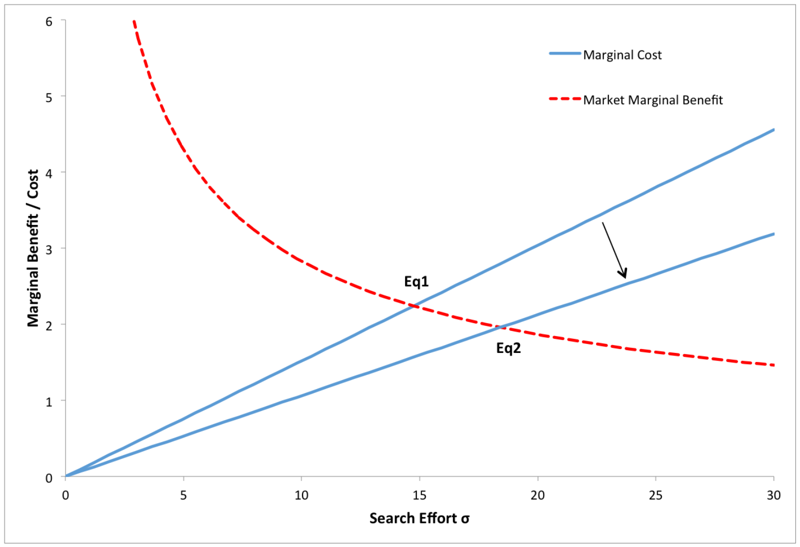

(a) *Increase in B holding constant $\frac{B}{S}$.*

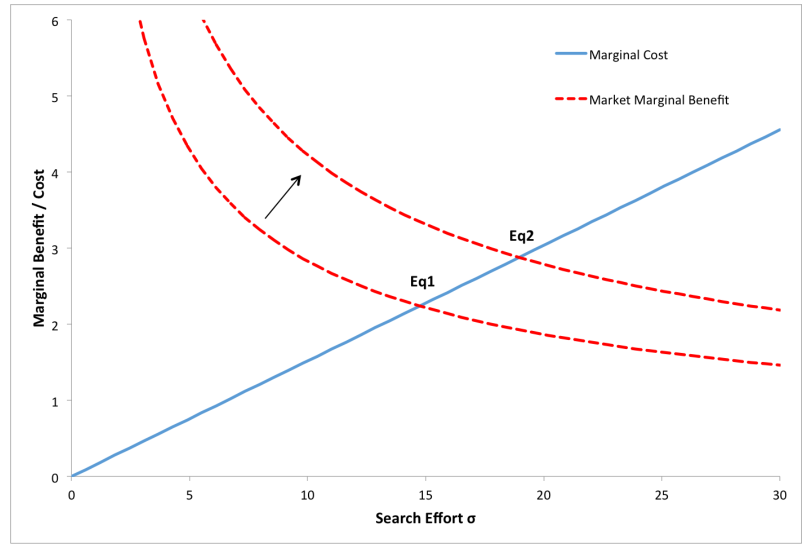

(b) *Increase in $\frac{B}{S}$ holding constant B.*

The figures plot the direct shifts (i.e. holding β constant) in the market-level marginal benefit and marginal cost curves of sellers, following changes in market scale and user composition. An increase in the number of buyers B, holding $\frac{B}{S}$ constant, shifts the marginal cost curve downward (left panel). The reduction in search costs induces more seller effort. An increase in the relative number of buyers $\frac{B}{S}$, holding B constant, shifts the marginal benefit curve upward (right panel) because each offer is now more likely to be accepted and pays a higher price. The increase in search benefits induces more seller effort.

Figure 8: *Individual Level Offer Supply Curves.*

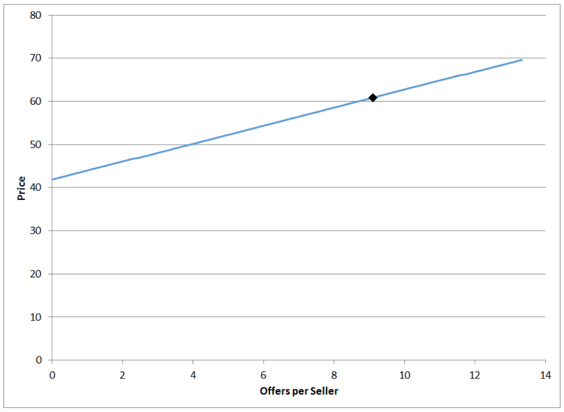

(a) *Seller effort as a function of price, holding constant the offer match rate at the equilibrium level in the median market.*

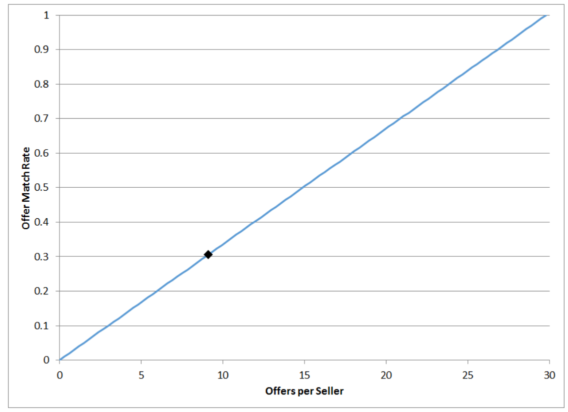

(b) *Seller effort as a function of the offer match rate, holding constant the price at the equilibrium level in the median market.*

The figures plot the individual level supply curves in the median market. Buyers' behavior is held constant at their equilibrium levels.

Figure 9: *Individual Level Task Demand Curves.*

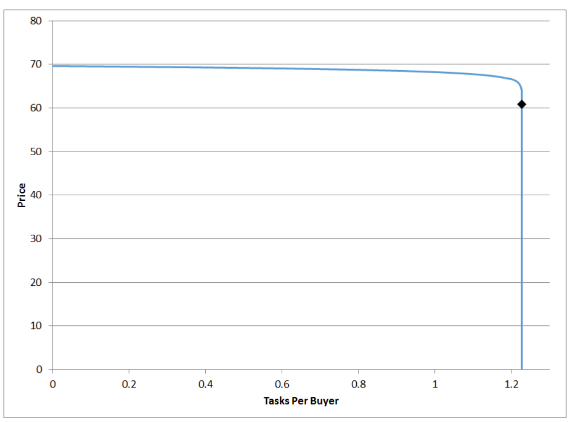

(a) *Tasks posted as a function of price, holding constant the task match rate at the equilibrium level in the median market.*

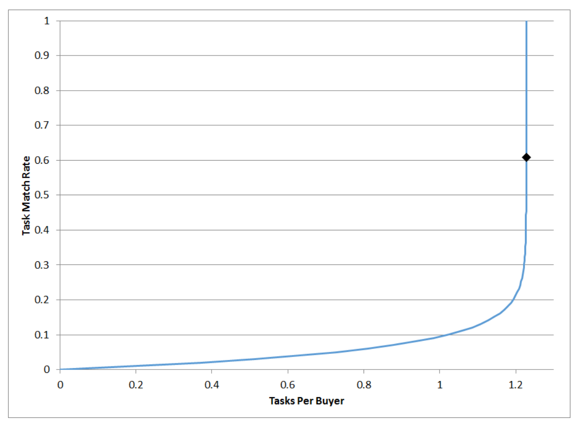

(b) *Tasks posted as a function of the task match rate, holding constant the price at the equilibrium level in the median market.*

The figures plot the individual level demand curves in the median market. Sellers' behavior is held constant at their equilibrium levels.

Figure 10: *Actual and Predicted Offers and Tasks.*

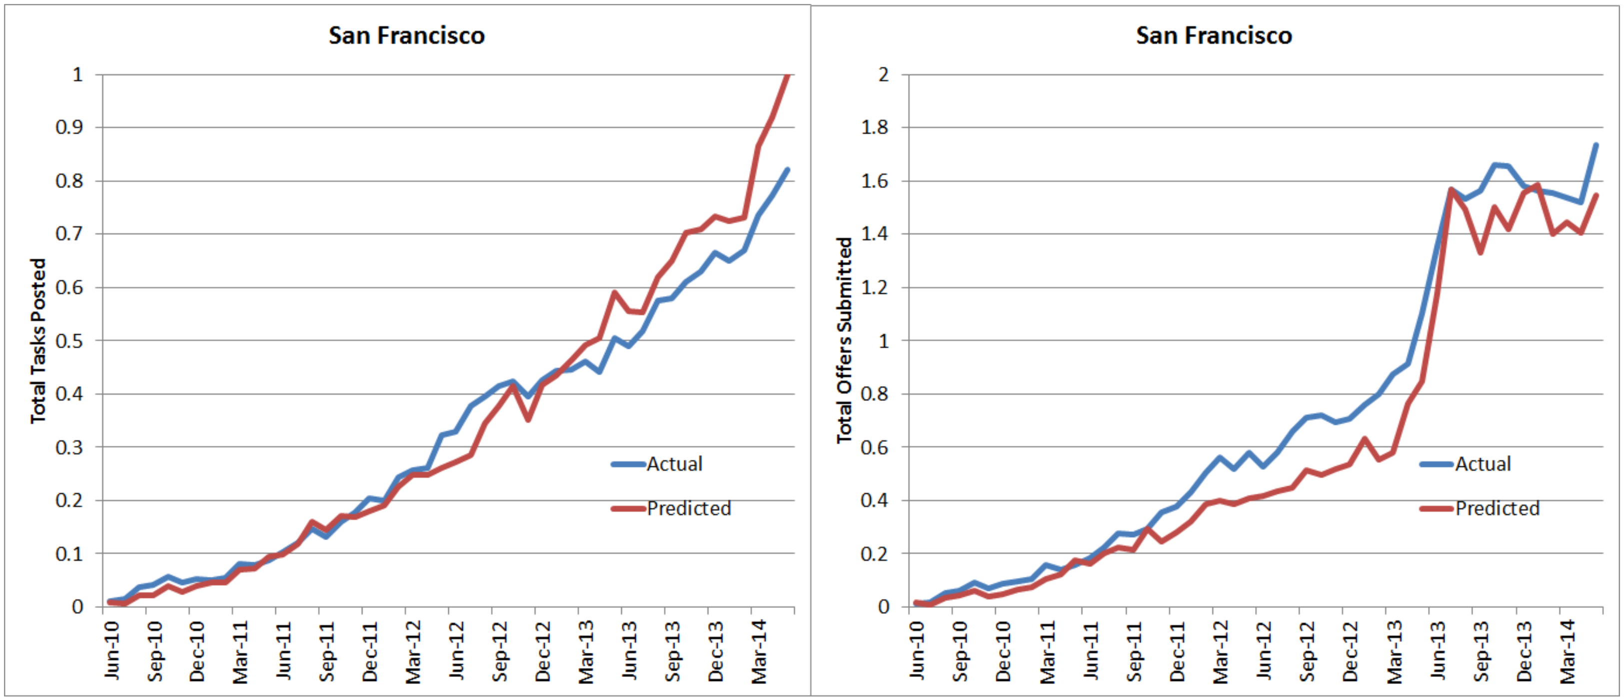

The figures plot the aggregate number of tasks (left) and offers (right) in San Francisco, comparing the actual and the predicted values. Every value is divided by the maximum number of tasks posted in a given month to protect company's privacy. Other cities are in the Appendix

Figure 11: *Buyer Adoption and Retention.*

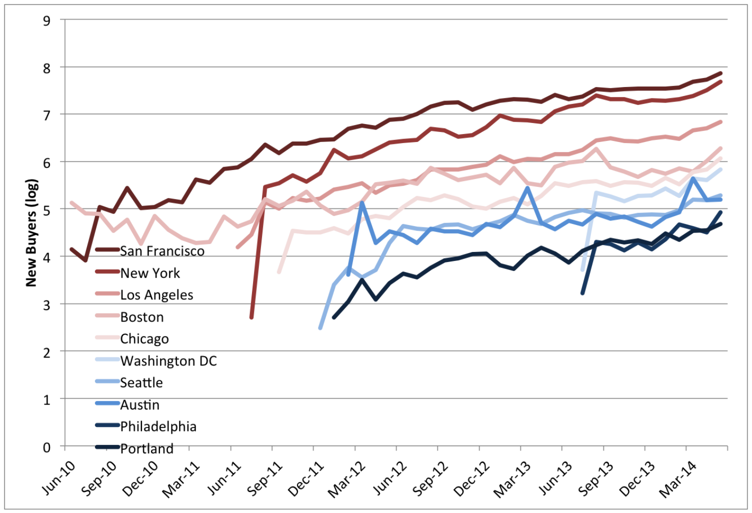

(a) *Number of new buyers by city over time. A buyer is defined as new in a city-month if she has never posted tasks in that city prior to that month.*

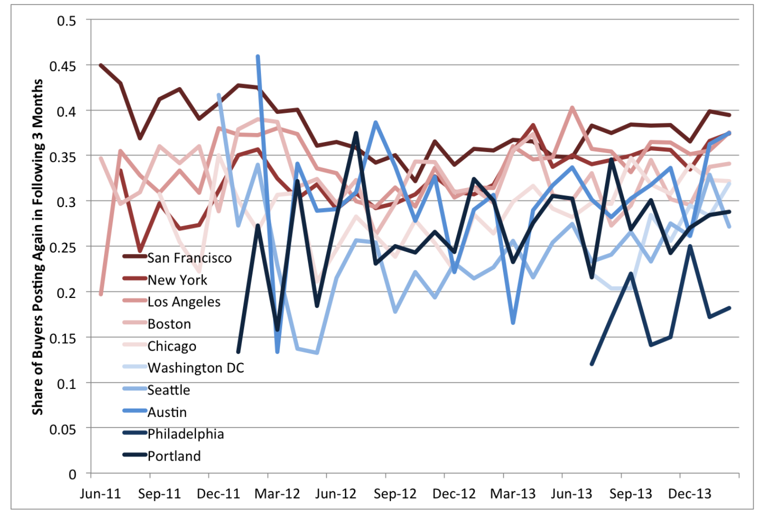

(b) *Buyer retention by city over time. For every city-month, the figure plots the share of current buyers posting again in the same city in the following three months.*

The figures plot adoption and retention patterns of buyers across cities over time, for selected cities.

Figure 12: *City Heterogeneity in Match Efficiency.*

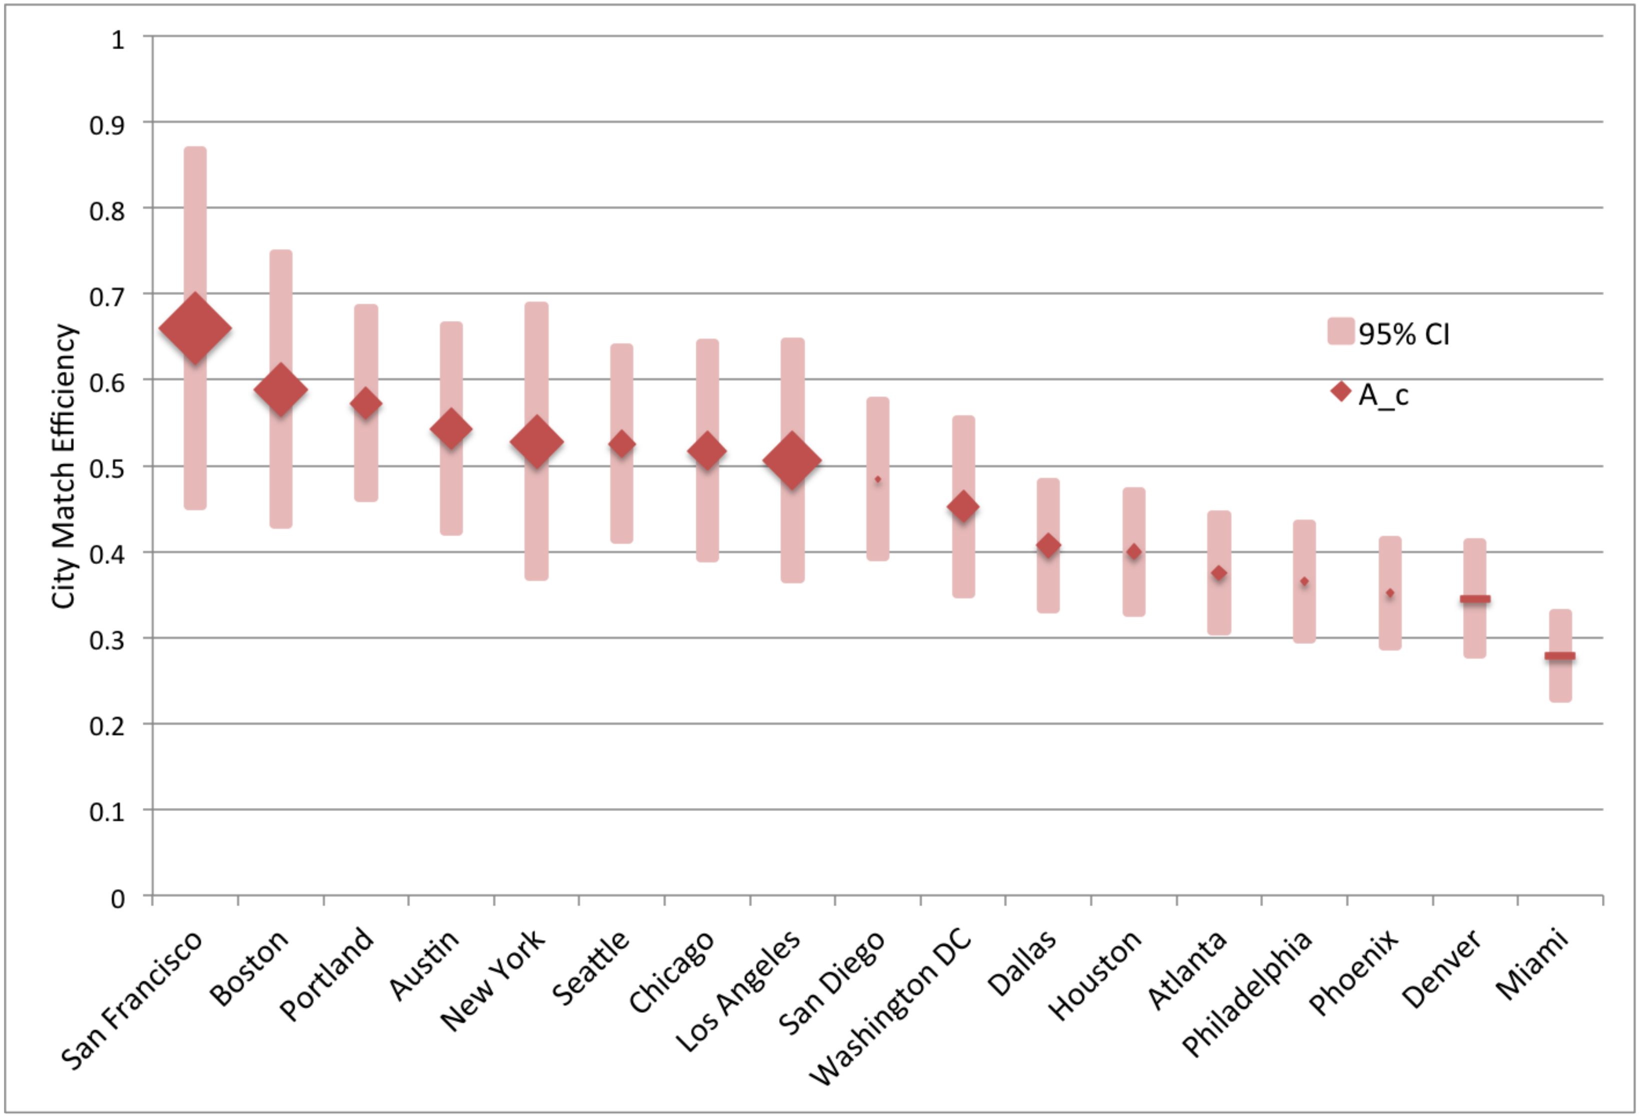

The figure shows A_c from the OLS regression $\log M_n = \log A_t + \log A_c + \alpha_1 \log b_n + \alpha_2 \log s_n + \log \epsilon_n^a$ of the (log-transformed) matching function from equation 7. Each city-month market is identified by $n = (t, c)$. The marker size is proportional to the retention rate parameter from the estimation of equation 14 (for Denver and Miami too few months of activity are recorded to estimate it).

Figure 13: *Geographic Distance Between Buyers and Sellers and Platform Growth*

The figure plots the median distance within a city-month between three different pairings of buyers and sellers, for the six largest cities over time. The distance is measured as the length of the shortest ellipsoidal curve between the buyer zip code and the seller zip code (Vincenty, 1975, integrated in Stata). Buyers and sellers are paired in three different ways. The first pairing (blue line) takes a seller who made an offer at a specific time, and pairs him to every buyer who posted tasks in the preceding 48 hours. The median distance is computed among all such pairs within a city-month. In the second pairing (orange line) 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-month, and the median is computed among all such pairings. In the third pairing (grey line) a buyer is paired with a seller if they actually exchanged services. As before, each buyer-seller pair is weighted by the number of their successful matches within a city-month.

Figure 14: *Buyer Adoption and Retention.*

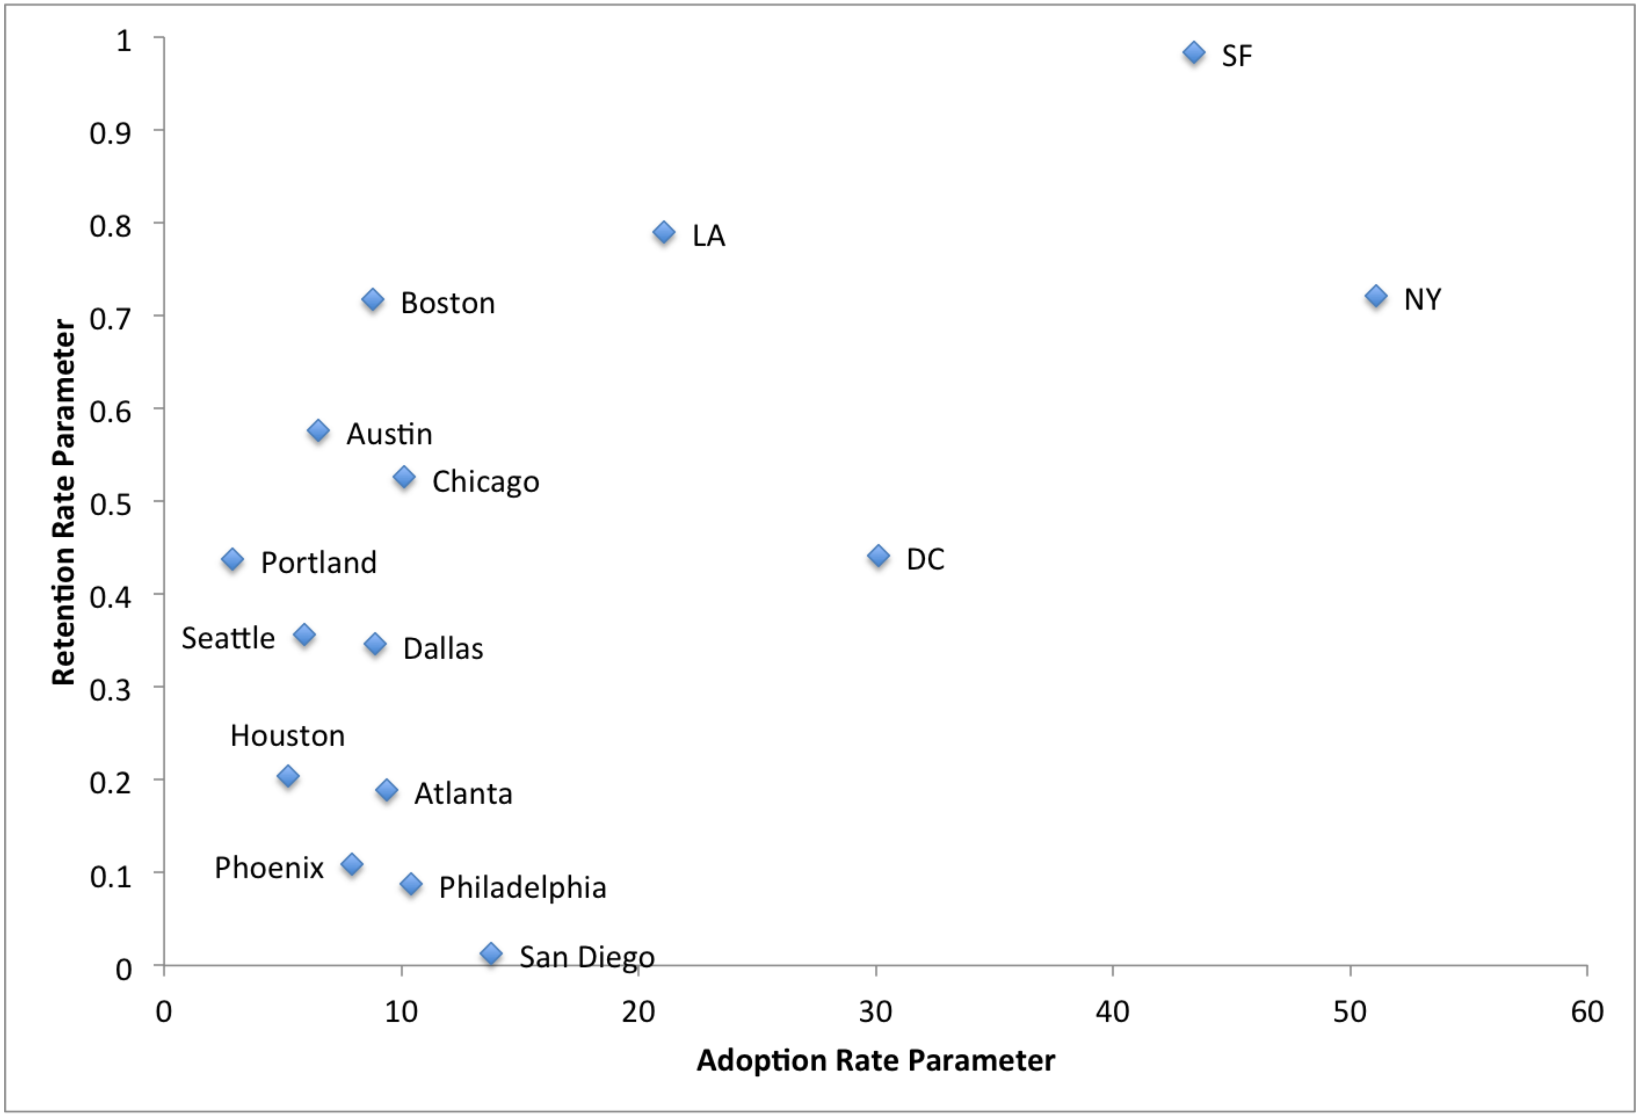

The figure plots the estimates of ϕ_c (city-specific adoption rate) and η_c (retention rate) from equations 13 and 14. η_c cannot be estimated for Denver and Miami, since we only have data for the last three months of activity.

Figure 15: *Geographic Distance Between Buyers and Sellers and City Heterogeneity*

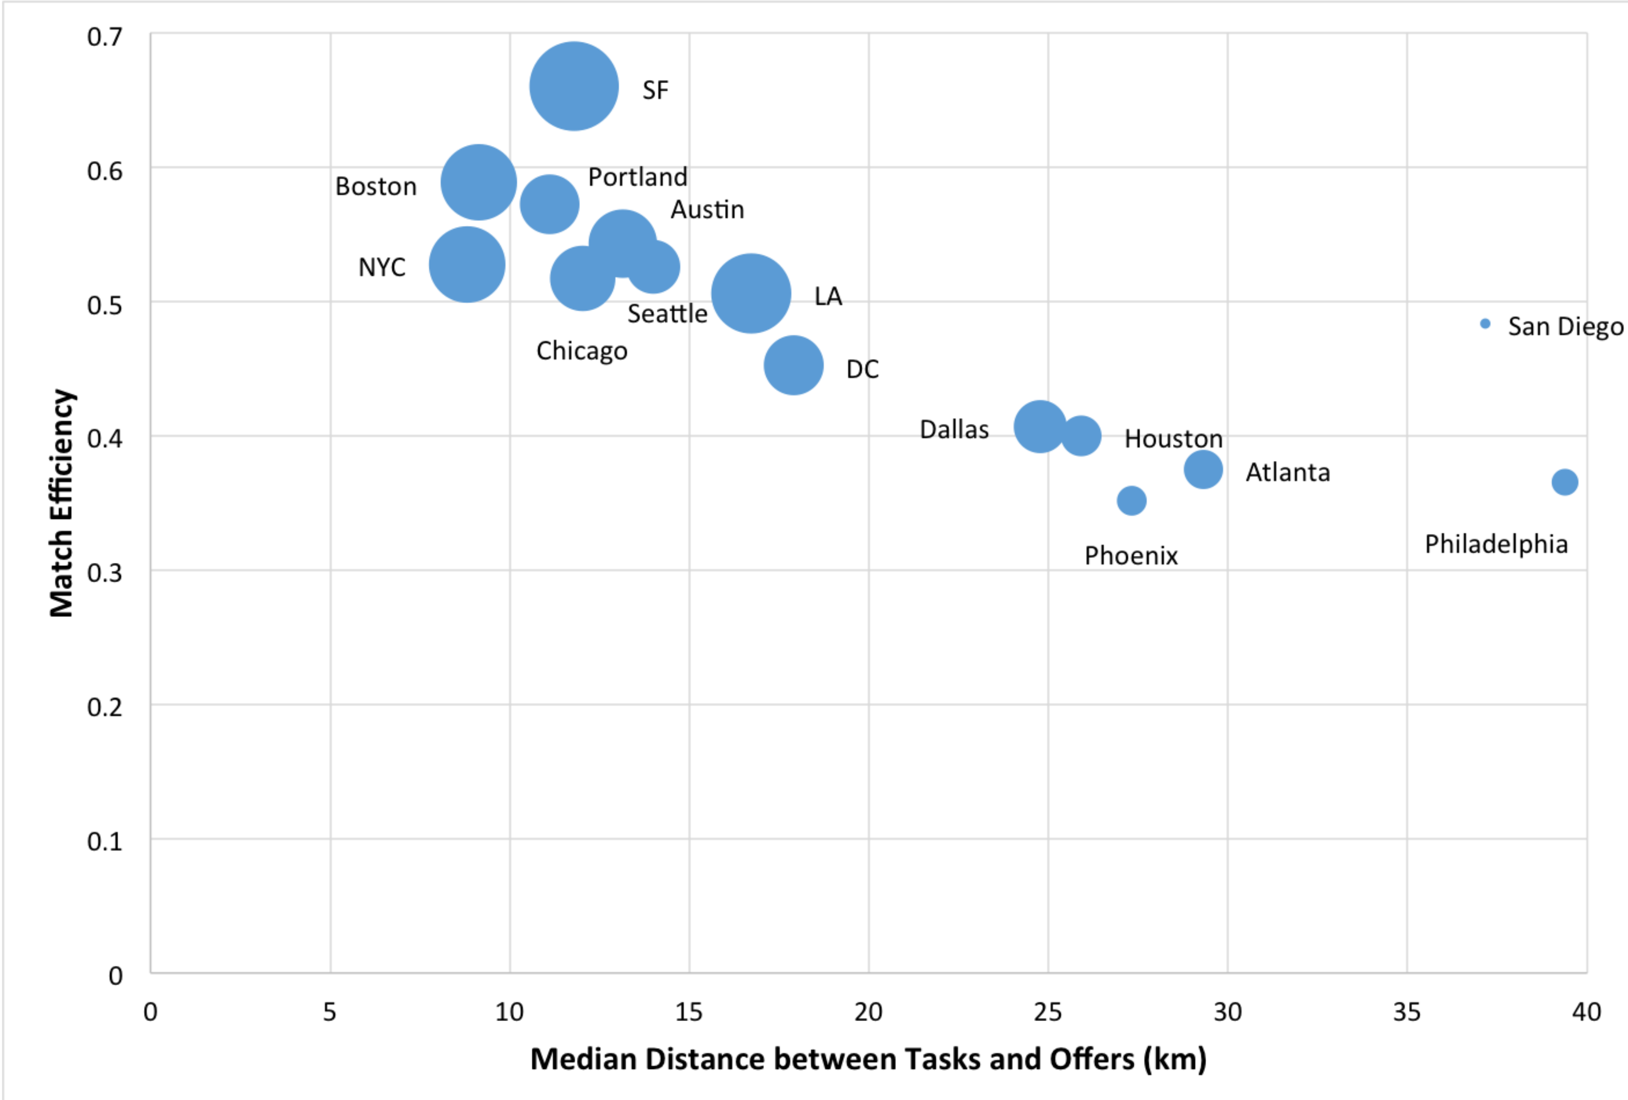

For each active city on TaskRabbit, the figure plots the median distance between buyers and sellers against the city estimate of the match productivity parameter A_c from equation 7. The distance between buyers and sellers is computed according to the second definition from Figure 13: a buyer is paired 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 from the estimation of equation 14 (for Denver and Miami too few months of activity are recorded to estimate it). The correlation with the overall platform growth by city is also positive and shown in Appendix G.

Figure 16: *Task Standardization and City Heterogeneity*

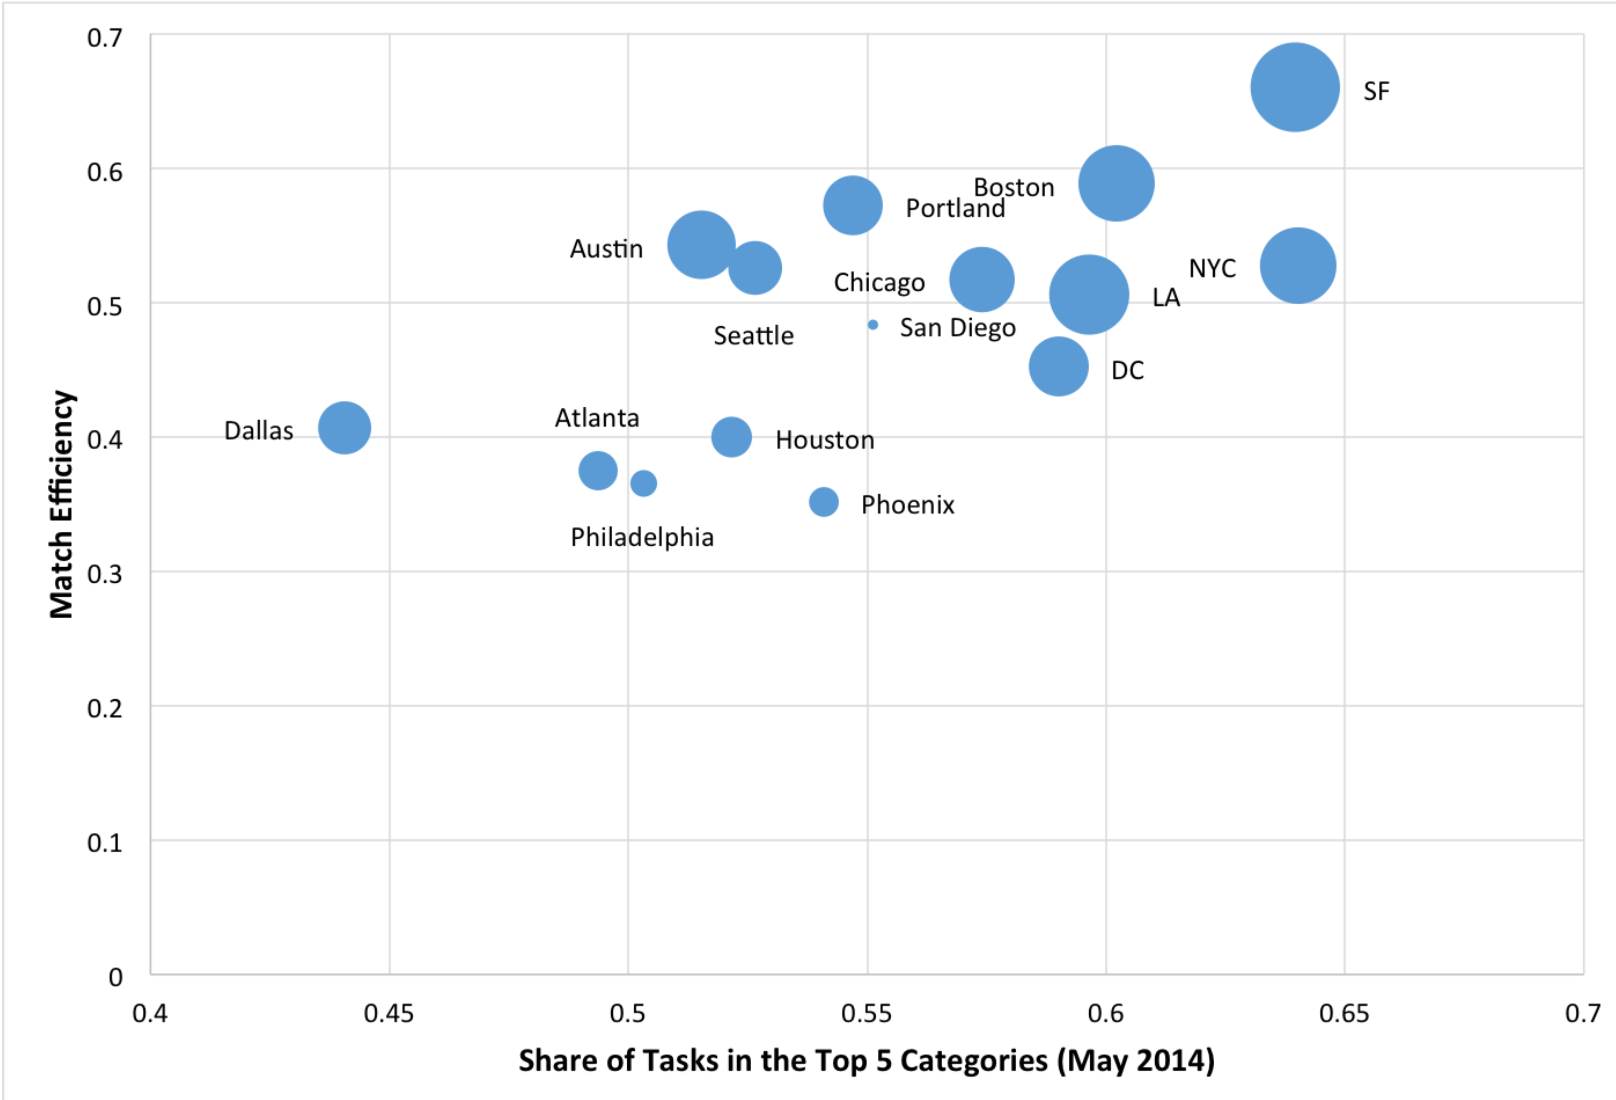

For each active city on TaskRabbit, the figure plots a measure of task standardization against the city estimate of the match productivity parameter A_c from equation 7. The measure of task standardization is the share of tasks posted in May 2014 within the top 5 categories - Shopping and Delivery, Moving Help, Cleaning, Minor Home Repairs, and Furniture Assembly. Alternative measure of standardization (other time frames, top 3 categories) display similar patterns. The size of each bubble is proportional to the retention rate parameter from the estimation of equation 14 (for Denver and Miami too few months of activity are recorded to estimate it). The correlation with the overall platform growth by city is also positive and shown in Appendix G.

Figure 17: *New TaskRabbit.*

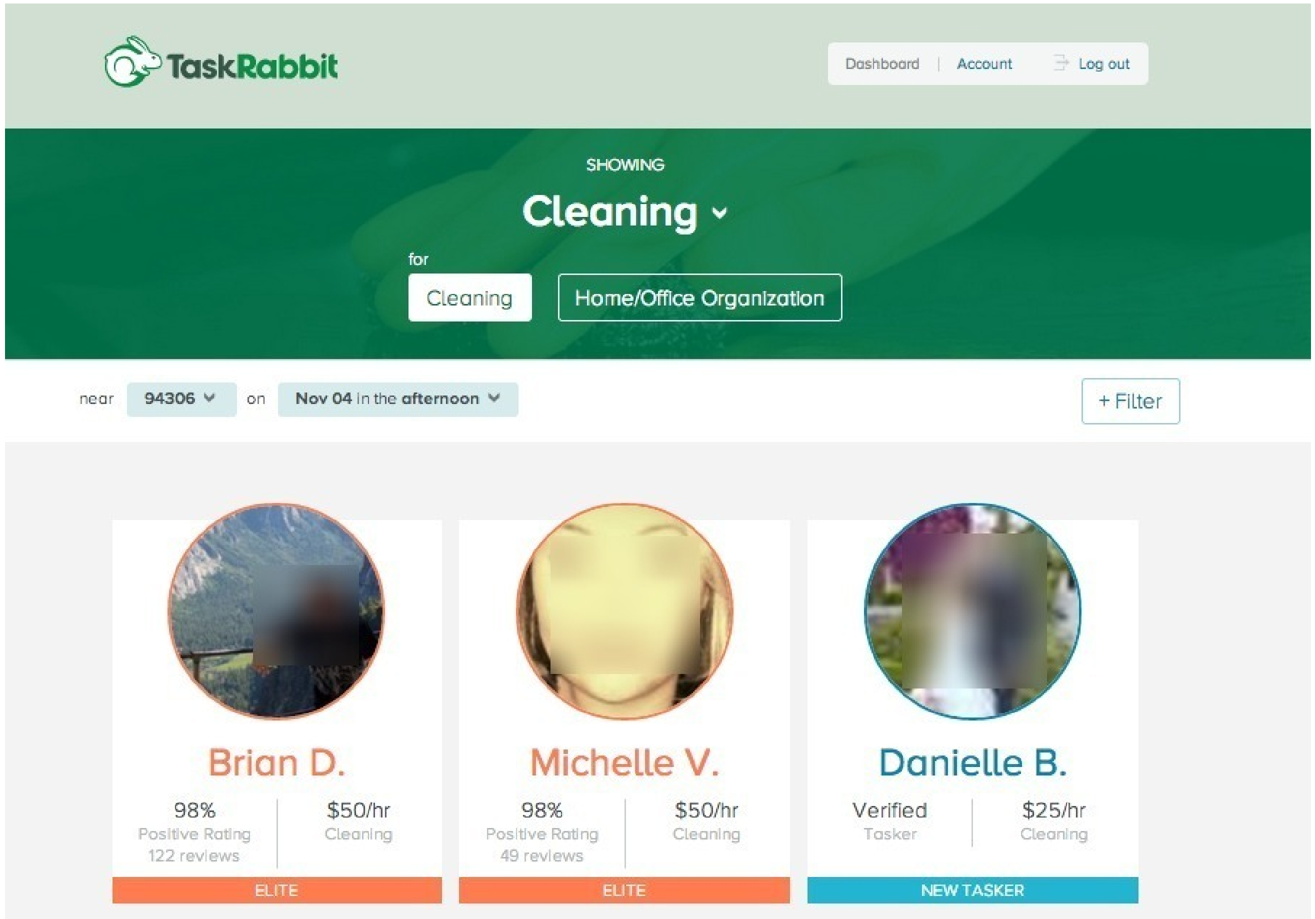

Task posting on the new TaskRabbit. Buyers can now select the category, location, and time for a given task request, and then either choose among the sellers available to perform that task type at the specified time and location (similar to the auction mechanism prior to the change), or have the platform automatically choose for them (similar to the posted price mechanism). The screenshot was taken on October 30th, 2014.

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 (≥ 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 Price Charged($)	336	56	8.69	52	57	61

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

Summary statistics. Data include posted price and auction tasks, offers submitted to those tasks, and matches created in the 18 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: *Tasks per Buyer, Offers per Seller, Match Rates and Prices*

	Tasks per Buyer	Offers per Seller	Task Match Rate	Prices
Buyer to Seller Ratio	-0.035	0.404	-0.154	0.053
	[0.027]	[0.034]***	[0.062]**	[0.032]
Nr. of Participants	0.054	0.314	0.113	0.022
	[0.009]***	[0.025]***	[0.040]**	[0.018]
Constant (SF Oct '13)	0.22	-0.277	-1.282	3.849
	[0.041]***	[0.145]*	[0.217]***	[0.116]***
City FE	No	No	No	No
Month FE	No	No	No	No
N Markets	336	336	336	336
R-squared	0.222	0.869	0.155	0.047

(a) *Regressions without fixed effects.*

	Tasks per Buyer	Offers per Seller	Task Match Rate	Prices
Buyer to Seller Ratio	-0.0002	0.416	-0.222	0.009
	[0.030]	[0.043]***	[0.050]***	[0.047]
Nr. of Participants	0.051	0.248	0.025	-0.003
	[0.029]	[0.034]***	[0.022]	[0.015]
Constant (SF Oct '13)	0.205	0.111	-0.546	4.131
	[0.224]	[0.247]	[0.168]***	[0.126]***
City FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
N Markets	336	336	336	336
R-squared	0.604	0.939	0.851	0.727

(b) *Regressions with city fixed effects and calendar month fixed effects.*

The tables show results from OLS regressions of the following type: $\log(y_{tc}) = \theta_1 \log\left(\frac{B_{tc}}{S_{tc}}\right) + \theta_2 \log\left(\sqrt{S_{tc}B_{tc}}\right) + \eta_c + \eta_t + \nu_{tc}$*, where* c, t *denote city* c *and calendar month* t*.* $\frac{B_{tc}}{S_{tc}}$ *is the active buyer to seller ratio,* $\sqrt{S_{tc}B_{tc}}$ *is the geometric average of the number of active buyers and sellers, and* y_{tc} *is one of the four relevant variables: users' choices (tasks per buyer, offers per seller), and outcomes (task match rate, prices). The top panel shows estimates without city fixed effects or month fixed effects. The bottom panel shows estimates with those fixed effects. Standard errors are clustered at the city level. **** $p < 0.01$*, *** $p < 0.05$*, ** $p < 0.1$*. Robustness checks are shown in the Appendix.*

Table 3: *Pricing and Matching Function Parameters*

	Number of Matches (log)	Average Price (log)
Tasks (log)	0.41	0.015
	[0.033]***	[0.035]
Offers (log)	0.521	-0.013
	[0.028]***	[0.029]
City FE	Yes	Yes
Month FE	Yes	Yes
N Markets	336	336
R-squared	0.996	0.727

The table shows results from OLS regressions of the (log-transformed) price and matching functions from equations 7 and 8. City and time heterogeneity parameters are displayed in figures 12 through A4.

Table 4: *Utility Parameters.*

Demand Parameters	
Task arrival μ	1.23
Posting cost η	2.01
Task value v	70

Supply Parameters	
Search cost γ	0.41
Scale effect in search δ	0.25
Task cost c	33

The table presents the estimates of the buyer and seller utility parameters. They are estimated by method of moments using our assumption of equilibrium behavior, the orthogonality of S and B from contemporaneous demand and effort shocks, and with the first-stage estimates of the pricing and matching functions.

A Appendix: Proofs

Proposition 1. *(Effect of scale) An increase in B, holding $\frac{B}{S}$ fixed, leads to an increase in β and σ, and a decrease in $\frac{B\beta}{S\sigma}$. This in turn implies lower prices and seller match rates, and higher buyer match rates.*

Proof. Given homogeneity of degree one in tasks and offers for the matching function, and homogeneity of degree zero for the pricing function, the conditions for market equilibrium are the following:

$$\beta = \mu\left(1 - e^{-\eta q^b (v-p)}\right) \tag{15a}$$

$$\sigma = \gamma(B\beta)q^s\left[(1-\tau)p - c\right] \tag{15b}$$

$$p = P\left(\frac{B\beta}{S\sigma}\right) \tag{15c}$$

$$q^b = \frac{M\left(\frac{B\beta}{S\sigma}\right)}{\frac{B\beta}{S\sigma}} \tag{15d}$$

$$q^s = M\left(\frac{B\beta}{S\sigma}\right) \tag{15e}$$

where p and q^s are increasing in $\frac{B\beta}{S\sigma}$ (hence increasing in $\frac{B}{S}$ and β, and decreasing in σ) and q^b is decreasing in $\frac{B\beta}{S\sigma}$. Keeping this in mind, and given that $\frac{B}{S}$ is held constant, we can rewrite:

$$\beta = f_1(\beta, \sigma)$$
$$\sigma = f_2(\sigma, \beta, B).$$

Given our assumptions, f_1 is continuous and decreasing in β and f_2 is continuous and decreasing in σ. Moreover, f_1 is increasing in σ, while f_2 is increasing in β. Intuitively, β and σ are strategic complements: an increase in the number of offers, by reducing price and increasing the task match rate, increases the number of posted tasks. Analogously, an increase in the number of tasks, by increasing price and increasing the offer match rate, increases the number of submitted offers. Finally, f_2 is increasing in B. Therefore, by theorem 4 of Milgrom and Roberts (1994), an increase in B leads to an increase in both β and σ.

In order to prove that $\frac{B\beta}{S\sigma}$ decreases in B, assume, by contradiction, that $\frac{\beta}{\sigma}$ increases. An increase in the task to offer ratio implies that p increases and q^b decreases, in turn leading to a decrease in β. This implication contradicts the earlier result that β increases with B. Therefore it must be that $\frac{B\beta}{S\sigma}$ decreases in B.

Given the reduction in $\frac{B\beta}{S\sigma}$, the effects of B on $q^{b,s}$, and p follow directly from the assumptions in the paper, i.e. a matching function which is increasing in both inputs and displays constant returns to scale, in addition to a pricing function which is increasing in tasks, decreasing in offers, and invariant to scale. $\square$

Proposition 2. *(Effect of user composition) An increase in $\frac{B}{S}$, holding B fixed, leads to an increase in β and a decrease in $\frac{B\beta}{S\sigma}$. This in turn implies higher prices and seller match rates, and lower buyer match rates.*

Proof. The equilibrium conditions are as in the proof of Proposition 1. We can take the ratio of equation 15a and 15b $\frac{\beta}{\sigma} = \frac{\mu\left(1-e^{-\eta q^b(v-p)}\right)}{\gamma(B\beta)q^s[(1-\tau)p-c]}\frac{B}{S} \rightarrow \frac{B\beta}{S\sigma}\frac{q^s[(1-\tau)p-c]}{\mu\left(1-e^{-\eta q^b(v-p)}\right)} = \frac{1}{\gamma(B\beta)}\frac{B}{S}$. The left hand side is an increasing function of $\frac{B\beta}{S\sigma}$. We can restate all conditions in terms of $\frac{B}{S}$, B, and the endogenous variables β and $\lambda = \frac{B\beta}{S\sigma}$ (and q^b, q^s, p):

$$\beta = \mu\left(1 - e^{-\eta q^b(v-p)}\right) \tag{16a}$$

$$\lambda = \frac{\mu\left(1 - e^{-\eta q^b(v-p)}\right)}{q^s\left[(1-\tau)p - c\right]}\frac{1}{\gamma(B\beta)}\frac{B}{S} \tag{16b}$$

$$p = P(\lambda) \tag{16c}$$

$$q^b = \frac{M(\lambda)}{\lambda} \tag{16d}$$

$$q^s = M(\lambda) \tag{16e}$$

We can rewrite equations 16a and 16b as:

$$\beta = g_1(\lambda)$$
$$\sigma = g_2\left(\lambda, \beta, \frac{B}{S}\right),$$

where g_1 is decreasing in λ, g_2 is decreasing in λ and β and increasing in $\frac{B}{S}$. Since λ and $-\beta$ are strategic complements, by Milgrom and Roberts (1994) an increase in $\frac{B}{S}$ leads to an increase in λ and a decrease in β. Because the pricing and matching functions are assumed to be increasing in

λ, and the latter displays diminishing returns, the increase in λ in turn implies higher q^s and p, and lower q^b. $\square$

B Appendix: Homogeneity of Tasks, Offers, Buyers, and Sellers

Homogeneity of tasks seems reasonable because so many of the tasks are relatively standard, not requiring specialized skills, and most sellers send offers across multiple task categories. However, we can gain a more nuanced view of search frictions by thinking explicitly about task heterogeneity, as we show in Section 7.

Homogeneity of buyers and sellers implies that all buyers choose the same task posting strategy, and all sellers choose the same level of search intensity. Homogeneity of buyers is less of a concern given that they tend to post few tasks and repeated platform use over time is limited. Sellers, on the other hand, can build experience and reputation on the platform. In the Appendix we provide evidence that our main results do not change when we account for some degree of seller experience.

This assumption allows us to simplify the space of choices available to buyers and sellers to just posting tasks and submitting offers. We do not explicitly model the selection of tasks to which a seller makes offers, nor a buyer's choice to assign the task to a specific bidder. We assume that the number of matches formed between buyers and sellers of services, as well as the price at which they trade, are determined by a matching and a pricing function. Underlying frictions due to actual task heterogeneity and information asymmetries are not made explicit, but summarized in a match productivity parameter. In this sense, the paper complements work by Fradkin (2014) on search inefficiencies on Airbnb, and Horton (2014) on congestion on Odesk. These papers study their respective platforms and quantify the efficiency losses due to specific types of frictions that an improved platform design can help alleviate. Our approach takes the platform "natural" level of frictions as given, and shows that it does not change considerably as more or less buyers, both in absolute and relative terms, are present in the market.

Our framework also assumes that services are independent of each other, both within a market and across markets. This implies that there are no externalities to other services from completing one task with a specific partner. Most buyers only post one task in a a city-month, providing some justification for this assumption. Moreover, matches in one market have no externalities on future matches. This effectively assumes that the benefit of trading one service with a specific partner does not carry over to future services. This is because a buyer might receive moving help today from a specific seller, but for cleaning tomorrow the same seller is unavailable, or does not have the right supplies. In practice this is true on TaskRabbit: the share of repeated buyer-seller pairs is only 6 percent of buyer-seller pairs that were ever matched.

C Appendix: Model Fit for Other Cities

Figure A1: *Actual and Predicted Offers and Tasks.*

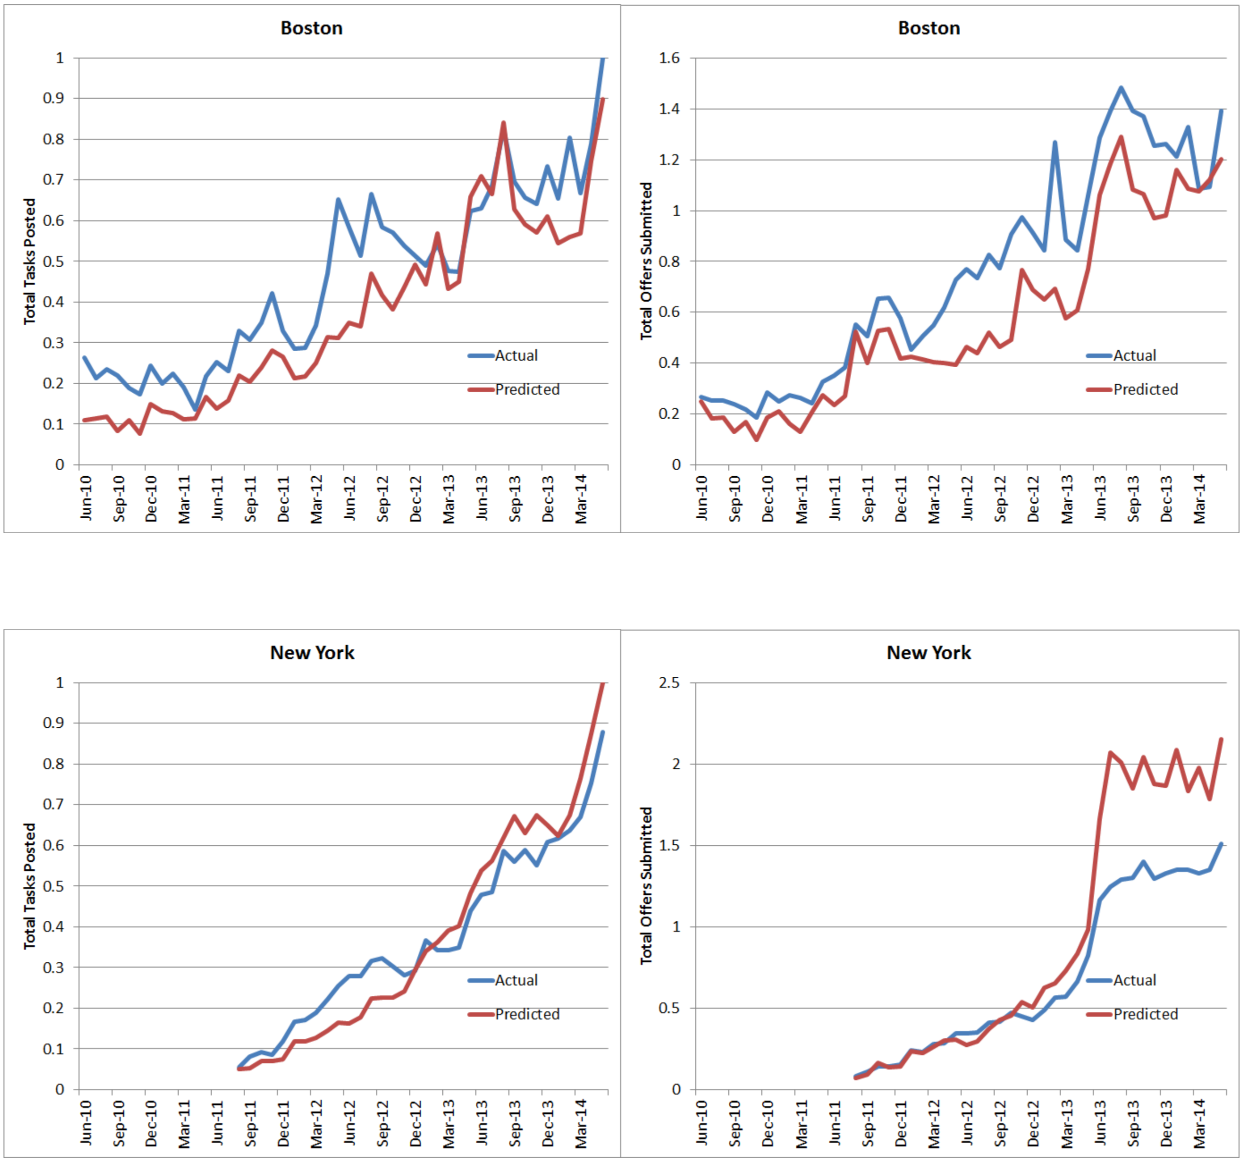

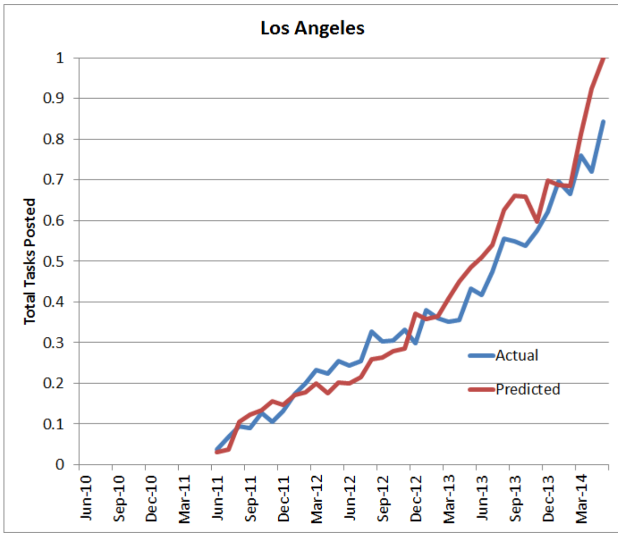
Los Angeles
Total Tasks Posted
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Jun-10
Sep-10
Dec-10
Mar-11
Jun-11
Sep-11
Dec-11
Mar-12
Jun-12
Sep-12
Dec-12
Mar-13
Jun-13
Sep-13
Dec-13
Mar-14
Actual
Predicted

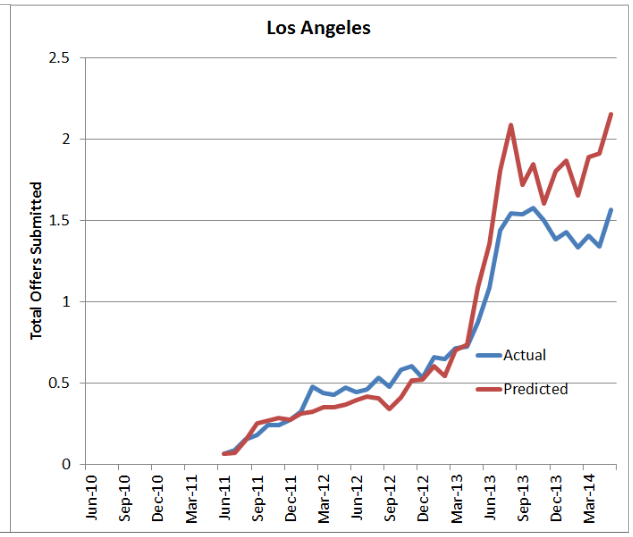
Los Angeles
Total Offers Submitted
2.5
2
1.5
1
0.5
0
Jun-10
Sep-10
Dec-10
Mar-11
Jun-11
Sep-11
Dec-11
Mar-12
Jun-12
Sep-12
Dec-12
Mar-13
Jun-13
Sep-13
Dec-13
Mar-14
Actual
Predicted

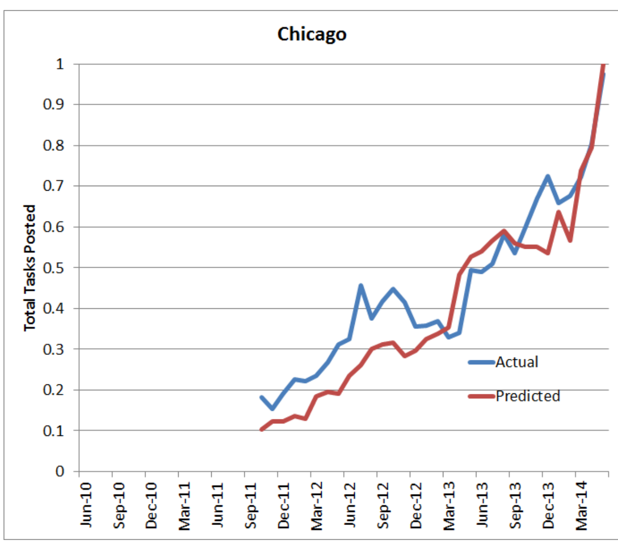
Chicago
Total Tasks Posted
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Jun-10
Sep-10
Dec-10
Mar-11
Jun-11
Sep-11
Dec-11
Mar-12
Jun-12
Sep-12
Dec-12
Mar-13
Jun-13
Sep-13
Dec-13
Mar-14
Actual
Predicted

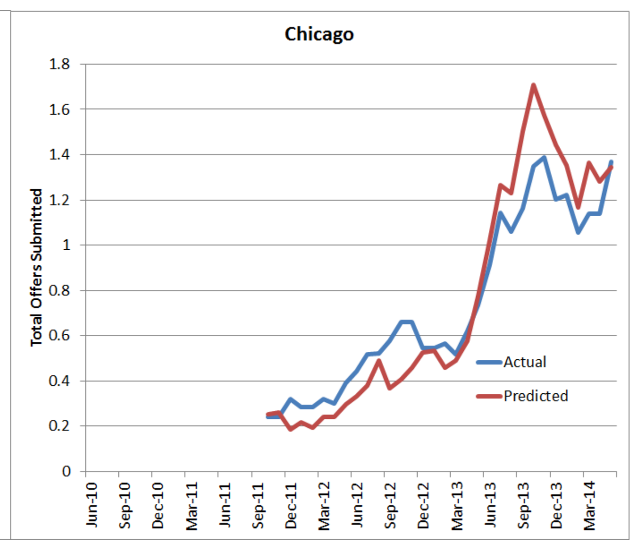
Chicago
Total Offers Submitted
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
Jun-10
Sep-10
Dec-10
Mar-11
Jun-11
Sep-11
Dec-11
Mar-12
Jun-12
Sep-12
Dec-12
Mar-13
Jun-13
Sep-13
Dec-13
Mar-14
Actual
Predicted

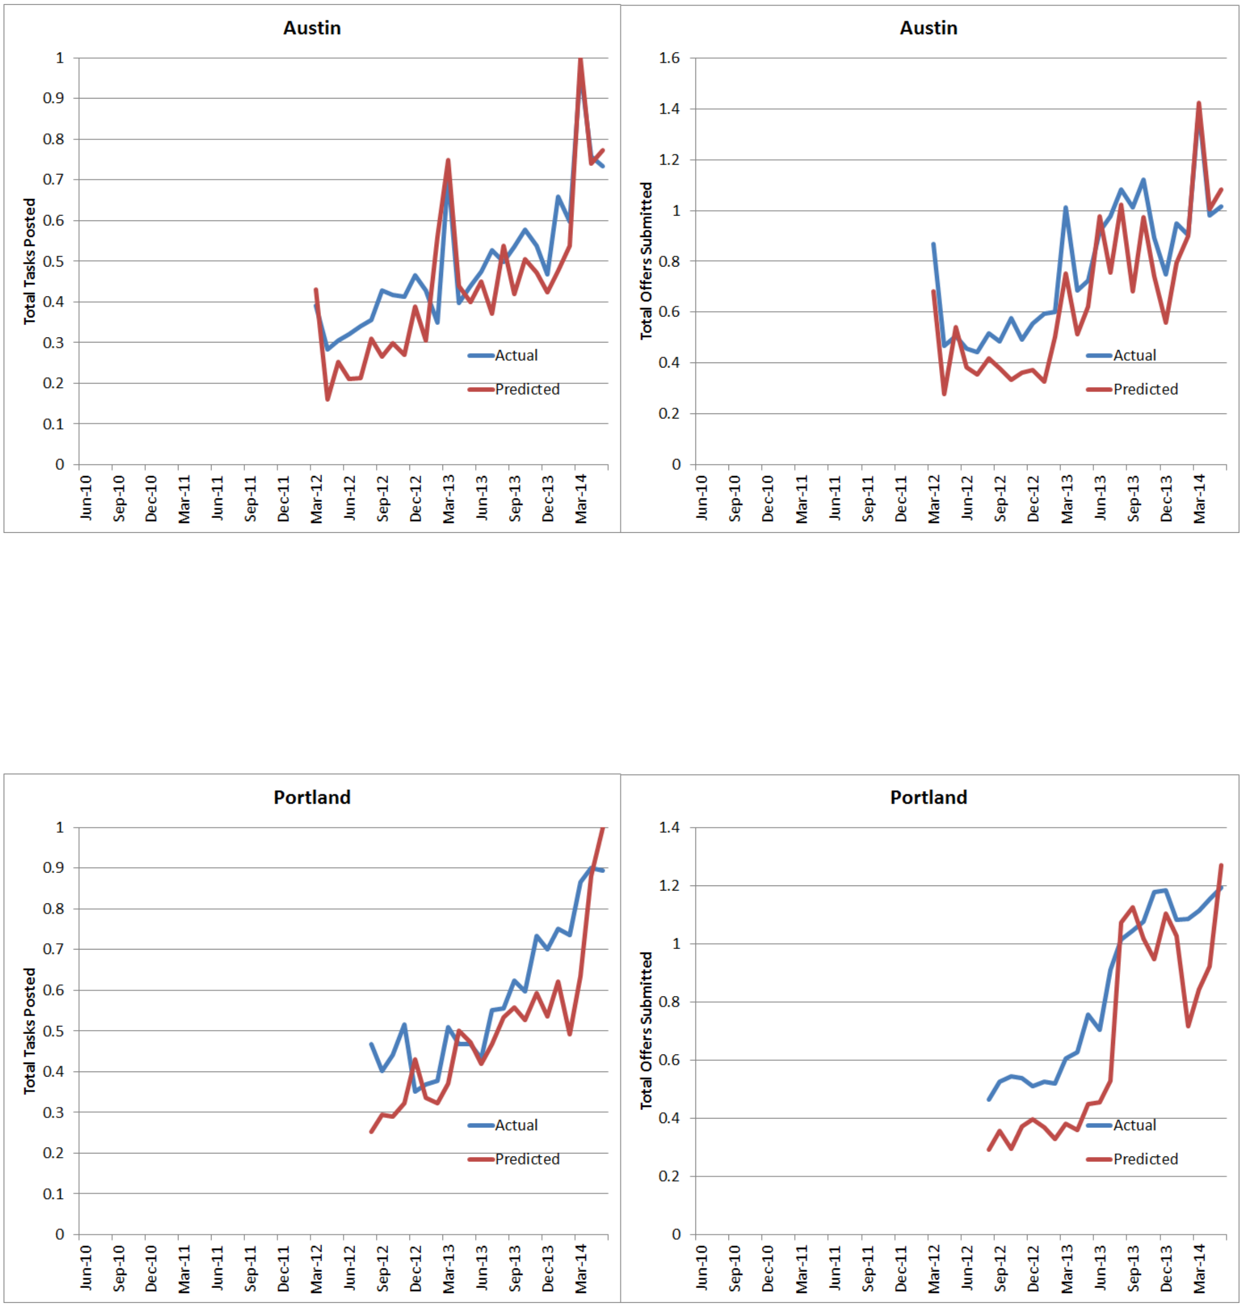

The figures plot the aggregate number of tasks (left) and offers (right) in cities other than San Francisco, comparing the actual and the predicted values. For every city, each value is divided by the maximum number of tasks posted in a given month to protect company's privacy.

D Appendix: Adoption

Table A1 shows the coefficients from the estimation of equation 13.

Table A1: *Buyer Adoption.*

	Coefficient	Standard Error
Boston	8.782	[0.871]***
SF	43.41	[0.871]***
Austin	6.518	[1.482]***
Chicago	10.141	[1.261]***
Seattle	5.9	[1.385]***
Portland	2.885	[1.432]*
LA	21.091	[1.157]***
NY	51.078	[1.190]***
Atlanta	9.364	[4.023]**
Dallas	8.876	[4.023]**
Houston	5.224	[4.023]
Miami	7.336	[4.970]
Philadelphia	10.41	[3.673]**
Phoenix	7.935	[5.632]
San Diego	13.777	[4.970]**
Washington DC	30.096	[3.673]***
Denver	8.421	[5.632]
Constant	12.43	[28.162]
N Markets	391	
R-squared	0.976	

*The table shows the OLS estimates of equation 13. Standard errors are clustered at the city level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.*

We also verify that deviations from the linear trend in buyer adoption are not driven by contemporaneous market conditions. To do this, we take the residuals from equation 13 and run the following regression:

$$\hat{\nu}_{tc} = \theta X_{tc} + \eta_c + \eta_t + \epsilon_{tc}.$$

X_{tc} is two-element vector of relevant outcomes in city-month t, c: realized buyer match rate and average transacted price. As table shows, the estimates of θ are not statistically significant, and not even of the expected sign.

Table A2: *Buyer Adoption and Contemporaneous Market Conditions.*

	Predicted Residual	
Task Match Rate*100	-0.879	-1.42
	[0.768]	[0.817]
Average Transacted Price	0.541	1.126
	[0.910]	[0.769]
City FE	No	Yes
Month FE	No	Yes
N Markets	336	336
R-squared	0.028	0.298

The table shows the OLS estimates of $\hat{\nu}_{tc} = \theta X_{tc} + \eta_c + \eta_t + \epsilon_{tc}$*, where* $\hat{\nu}_{tc}$ *are the predicted residuals from equation 13, and* X_{tc} *is two-element vector of relevant outcomes in city-month* t, c*: realized buyer match rate and average transacted price. Standard errors are clustered at the city level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.*

E Appendix: Retention

To support our main identification assumption, we verify that, conditional on the outcomes (matches and prices) in a current market, expectations on future outcomes do not affect the propensity to stay or leave the platform. To do this we run OLS regressions similar to equation 14:

$$\log\left(\frac{stay_{tc}}{1-stay_{tc}}\right) = \theta_0 X_{tc} + \theta_1 X_{t+1,c} + \theta_2 X_{t+2,c} + \theta_3 X_{t+3,c} + \eta_t + \eta_c + \nu_{tc} \ ,$$

X_{tc} is defined as in equation 14. The regression is run separately for buyers and sellers, so the match rate for buyers is the task success probability, while the match rate for sellers is the offer acceptance rate. The 6-element vector $(X_{t+1,c}, X_{t+2,c}, X_{t+3,c})$ contains the realized match rates and prices in the following three months within the same city. If users did not base their decision to stay or leave the platform on expectations of future outcomes we would expect the 6-element coefficient vector $(\theta_1, \theta_2, \theta_3)$ to be non significant, both for buyers and for sellers.

Results are presented in Tables A3 (for buyers) and A4 (for sellers). Each table has four columns, corresponding to four different specifications. The first specification estimates θ_0 without including forward variables. Each other specification sequentially adds the outcomes of the following month, two-months ahead, and three-months ahead. The final column is the full specification. With the exception of just one coefficient in the buyers' regression (the coefficient on the three-month-ahead transacted price), all other coefficients from $(\theta_1, \theta_2, \theta_3)$ are not statistically different from zero, and in some cases even have the opposite sign. Despite the small number of observations, another important observation from the two tables is that users appear to base their decisions to stay on the platform particularly on the contemporaneous match rate. Adding forward variables does not change the effect of match rate much nor helps improve the goodness of fit of the estimation.

Table A3: *Buyers' Retention*

	Buyer 3-month Retention			
Price Paid	-0.296	-0.293	-0.286	-0.321
	[0.146]*	[0.166]	[0.185]	[0.202]
Task Match Rate	0.341	0.32	0.267	0.215
	[0.117]**	[0.137]**	[0.126]*	[0.135]
Price Paid (t+1)		0.162	0.22	0.273
		[0.147]	[0.159]	[0.155]
Task Match Rate (t+1)		0.058	0.041	0.058
		[0.086]	[0.098]	[0.111]
Price Paid (t+2)			-0.219	-0.119
			[0.147]	[0.155]
Task Match Rate (t+2)			0.053	0.122
			[0.114]	[0.116]
Price Paid (t+3)				-0.503
				[0.193]**
Task Match Rate (t+3)				-0.215
				[0.133]
Constant	0.984	-0.332	0.2	1.565
	[0.566]	[0.729]	[0.775]	[0.754]*
City FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
N Markets	282	258	246	238
R-squared	0.758	0.75	0.746	0.757

The table shows results from OLS regressions of the following type: $\log(stay_{tc}) - \log(1 - stay_{tc}) = \theta_0 X_{tc} + \theta_1 X_{t+1,c} + \theta_2 X_{t+2,c} + \theta_3 X_{t+3,c} + \eta_t + \eta_c + \nu_{tc}$*, where* $stay_{tc}$ *is the share of buyers active in city-month* t, c *who were active again at least once in the following three months within the same city.* X_{tc} *is a two-element vector of relevant outcomes in city-month* t, c*: realized match rate and average transacted price (log scale). The match rate for buyers is the task success probability. Activity is defined as posting a task for buyers. Standard errors are clustered at the city level. **** $p < 0.01$*, *** $p < 0.05$*, ** $p < 0.1$*.*

Table A4: *Sellers' Retention*

	Seller 3-month Retention			
Fee Received	-0.079	-0.108	0.028	0.024
	[0.215]	[0.311]	[0.297]	[0.302]
Offer Match Rate	0.522	0.588	0.453	0.416
	[0.191]**	[0.231]**	[0.255]*	[0.282]
Fee Received (t+1)		-0.059	0.037	0.058
		[0.305]	[0.329]	[0.339]
Offer Match Rate (t+1)		-0.143	-0.215	-0.215
		[0.213]	[0.237]	[0.254]
Fee Received (t+2)			-0.067	-0.1
			[0.393]	[0.379]
Offer Match Rate (t+2)			0.279	0.33
			[0.276]	[0.320]
Fee Received (t+3)				0.141
				[0.234]
Offer Match Rate (t+3)				-0.18
				[0.167]
Constant	1.429	1.649	1.175	0.58
	[1.012]	[1.018]	[1.464]	[1.907]
City FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
N Markets	282	258	246	238
R-squared	0.725	0.696	0.674	0.664

The table shows results from OLS regressions of the following type: $\log(stay_{tc}) - \log(1 - stay_{tc}) = \theta_0 X_{tc} + \theta_1 X_{t+1,c} + \theta_2 X_{t+2,c} + \theta_3 X_{t+3,c} + \eta_t + \eta_c + \nu_{tc}$, *where* $stay_{tc}$ *is the share of sellers active in city-month* t, c *who were active again at least once in the following three months within the same city.* X_{tc} *is a two-element vector of relevant outcomes in city-month* t, c*: realized match rate and average transacted price (log scale). The match rate for sellers is the offer acceptance rate. Activity is defined as sending an offer for sellers. Standard errors are clustered at the city level.* *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

F Appendix: Estimates of the Pricing and Matching Functions

Figure A2: *Time Heterogeneity in Match Efficiency*

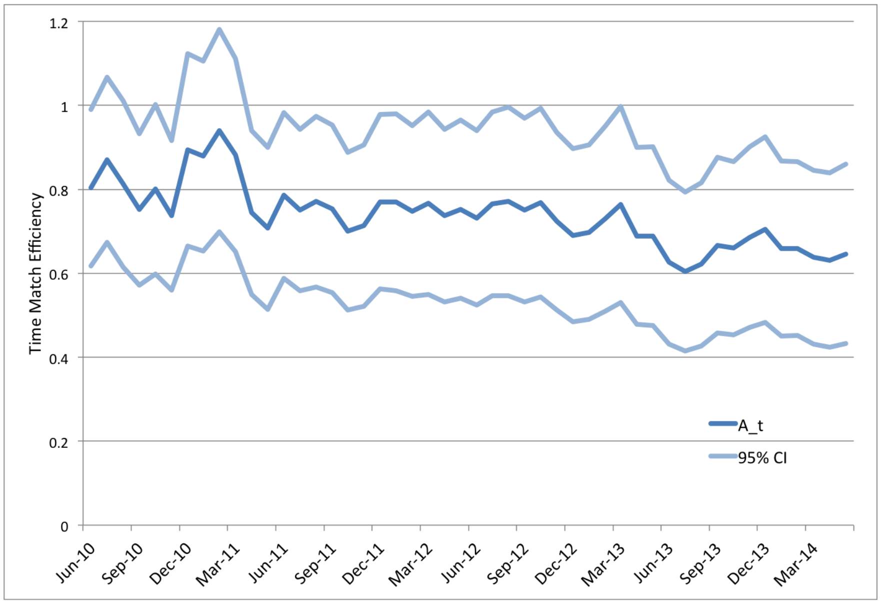

The figure shows A_t *from the OLS regression* $\log M_n = \log A_t + \log A_c + \alpha_1 \log b_n + \alpha_2 \log s_n + \log \epsilon_n^a$ *of the (log-transformed) matching function from equation 7. Each city-month market is identified by* $m = (t, c)$*.*

Figure A3: *City Heterogeneity in Prices*

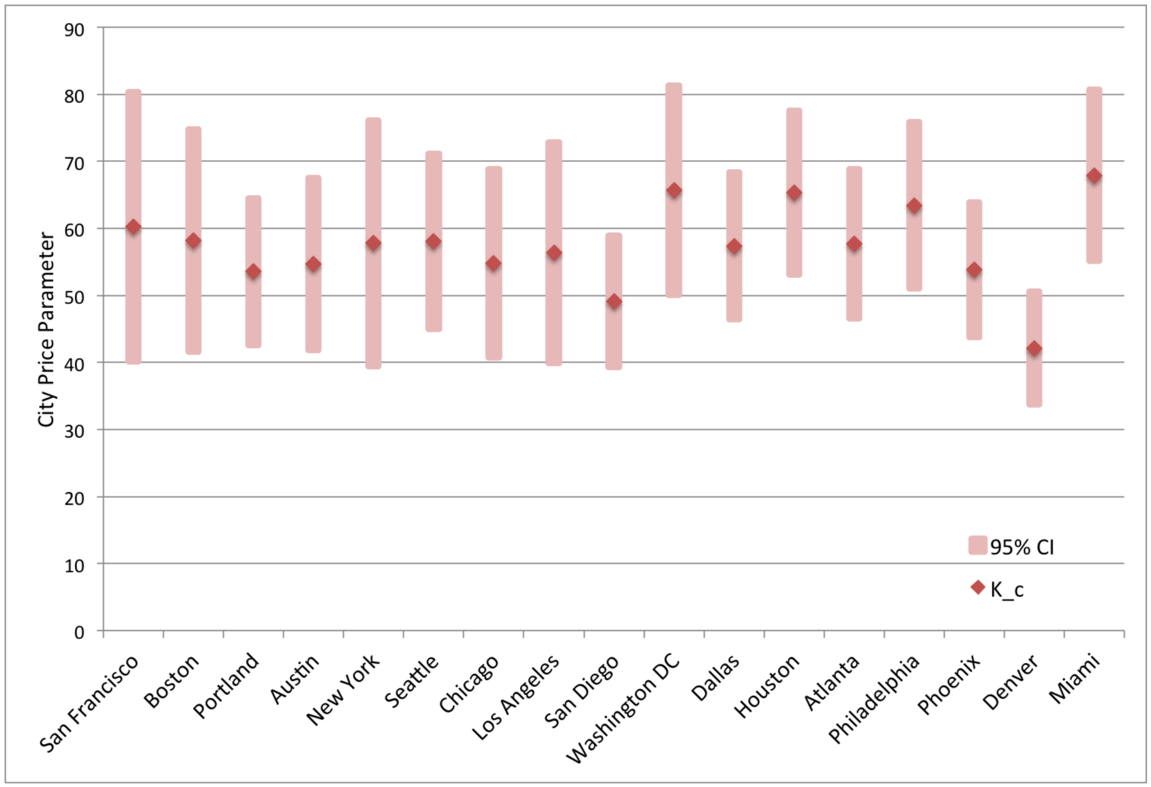

The figure shows K_c from the OLS regression $\log P_n = \log K_t + \log K_c + \rho_1 \log b_n + \rho_2 \log s_n + \log \epsilon_n^k$ of the (log-transformed) pricing function from equation 8. Each city-month market is identified by $m = (t, c)$.

Figure A4: *Time Heterogeneity in Prices*

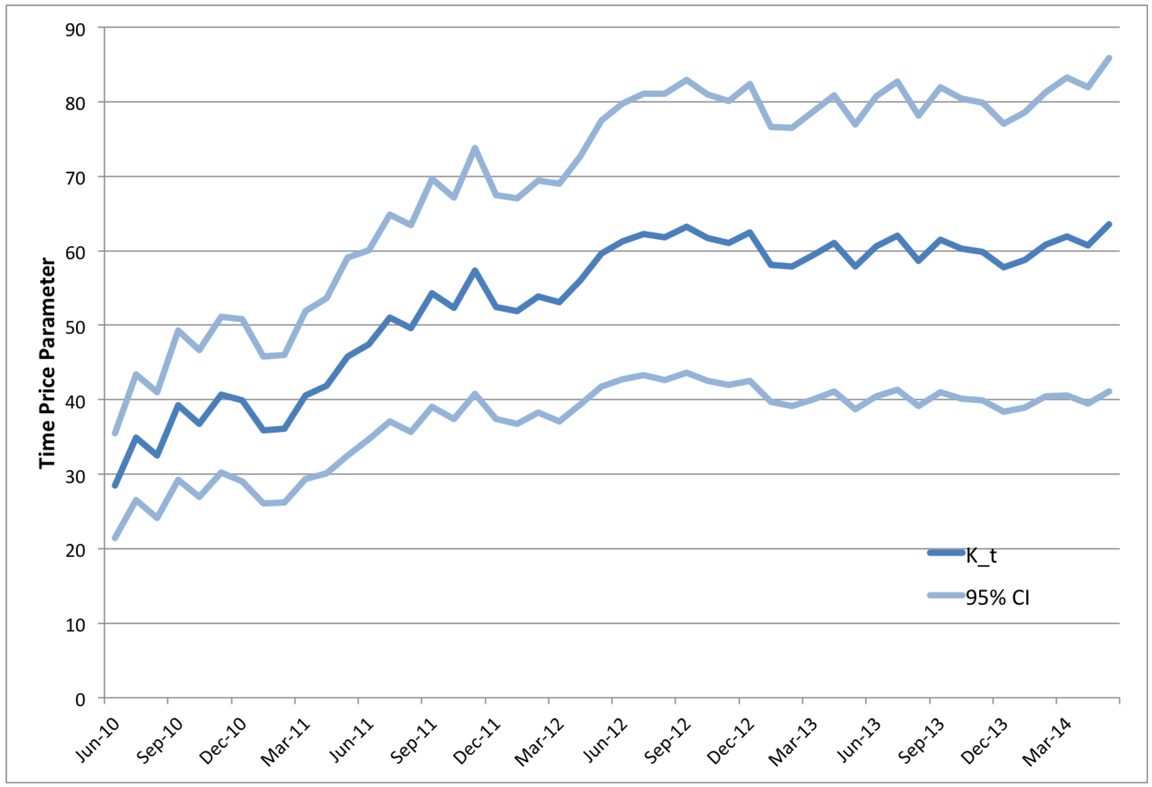

The figure shows K_t from the OLS regression $\log P_n = \log K_t + \log K_c + \rho_1 \log b_n + \rho_2 \log s_n + \log \epsilon_n^k$ of the (log-transformed) pricing function from equation 8. Each city-month market is identified by $m = (t, c)$.

Figure A5: *Task Categories over Time*

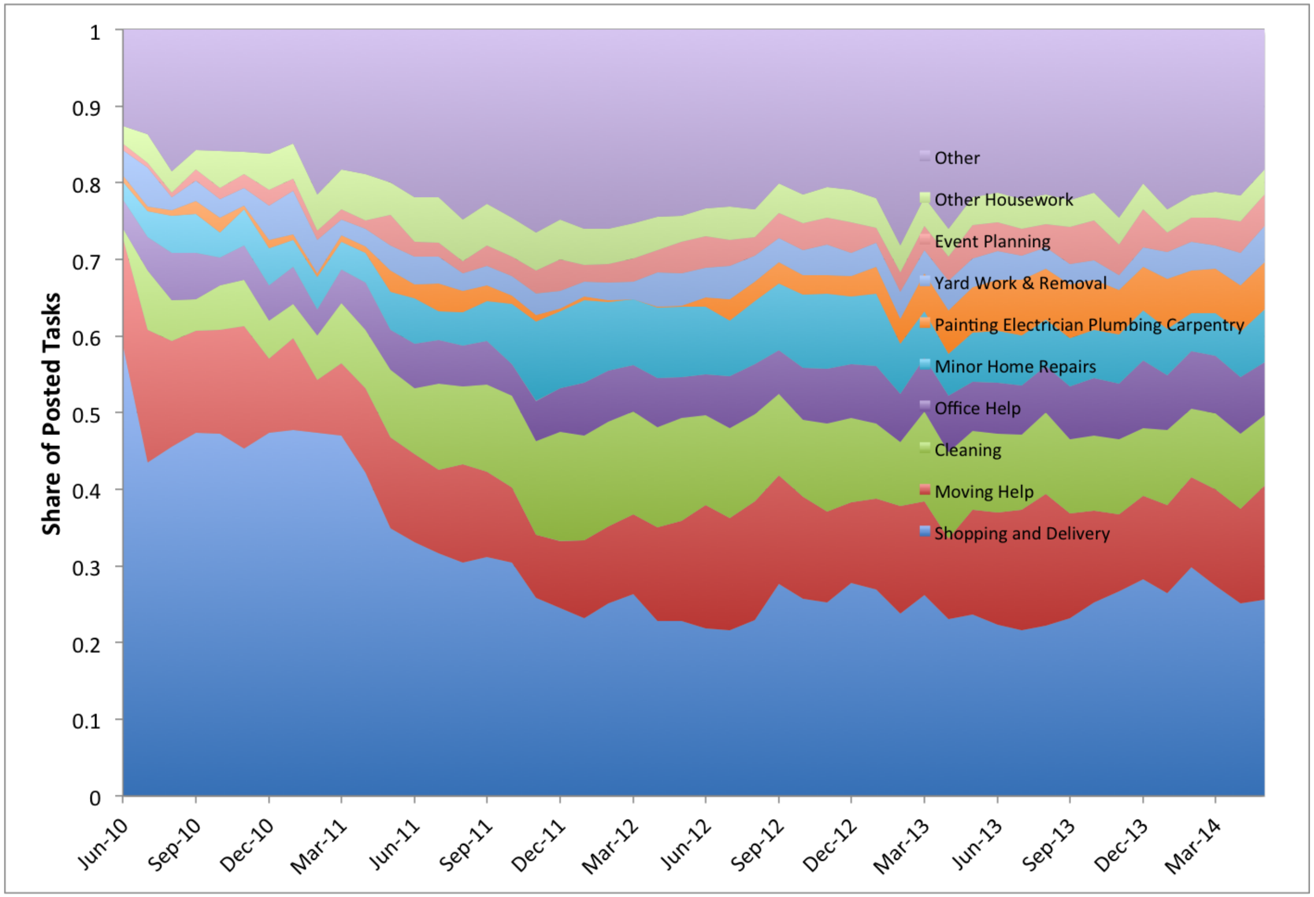

For each month on TaskRabbit, the figure plots the relative share of posted tasks falling in each of the main categories. We classified the 38 categories on the platform into 10 larger groups: Shopping and Delivery, Moving Help, Cleaning, Minor Home Repairs, Yard Work and Removal coincide with TaskRabbit's classification. "Painting Electrician Plumbing Carpentry" includes the four corresponding categories on TaskRabbit. "Other Housework" includes Laundry, Cooking and Baking, Sewing, Pet Care, Child Care, and Senior and Disabled Care. "Office Help" includes Packing and Shipping, Organization, Office Administration, Executive Assistant, Accounting. "Event Planning" includes Event Planning, and Event Staffing.

G Appendix: City Heterogeneity and Market Thickness

Figure A6: *City Heterogeneity and Market Thickness.*

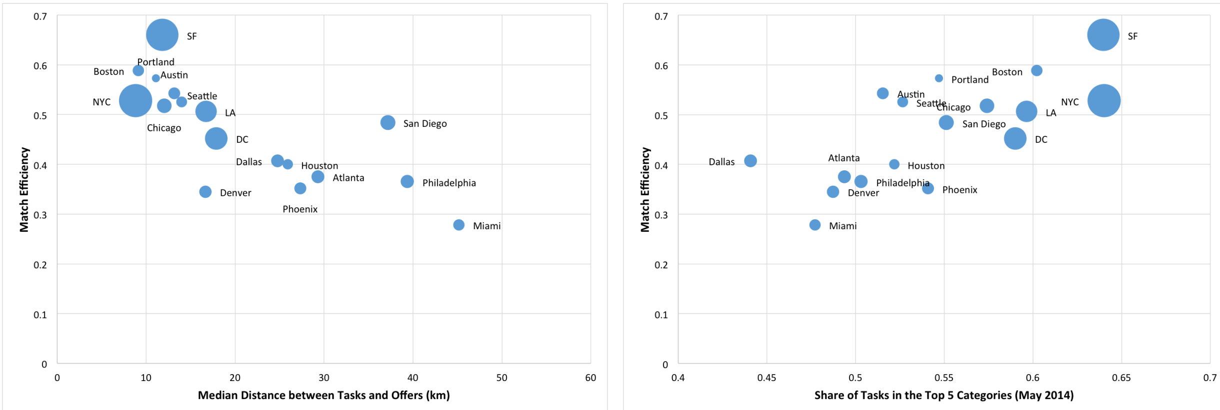

The figures plot the same market thickness and match efficiency metrics as Figure 15 (left) and Figure 16 (right). This time, however, the size of each bubble is proportional to the overall platform growth rate at the city level, which is the combination of buyer adoption and retention decisions.