

The Role of Quality in Internet Service Markets*

Elena Krasnokutskaya[†]

Johns Hopkins University

Kyungchul Song[‡]

University of British Columbia

Xun Tang[§]

Rice University

June 28, 2016

Abstract

We use data from an on-line market for programming services to assess buyers' welfare gains associated with the globalization enabled by the Internet. Our study exploits the design of this market where projects are allocated through multi-attribute auctions, a mechanism that takes into account a seller's non-price characteristics as well as his bid. We focus on the increased variety and competitive effects associated with the presence of low cost foreign sellers as the main welfare-improving consequences of globalization. The paper proposes an empirical methodology to recover primitives of the model in the presence of unobserved seller heterogeneity. The methodology is designed to accommodate two important features of such markets: buyer-specific choice sets and the high turnover of sellers. We find that the Internet enables buyers to substantially improve on their outside (local) option, with large part of the gains arising from access to the international markets.

Keywords: quality, services, procurement, multi-attribute auctions, unobserved agent heterogeneity, Internet

JEL Classification: C14, C18, D22, D44, D82, L15, L86.

*This version: June 28, 2016. We would like to thank Susan Athey, Phil Haile, Ken Hendricks, Ariel Pakes, Dan Akerberg, Greg Lewis and Marc Rysman for helpful discussions. We are also grateful to seminar participants at the University of Wisconsin-Madison, Rice University, Harvard University, Columbia University, 2012 Winter Meeting of Econometrics Society, 2012 Cowles Foundation Conference for Applied Microeconomics, and 2012 Society for Economic Dynamics Annual Meeting.

[†]Email: ekrasno1@jhu.edu.

[‡]Email: kysong@mail.ubc.ca.

[§]Email: xun.tang@rice.edu.

1 Introduction

Until recently, markets for professional services¹ were local for all but a very few (large) buyers because the cost of searching for non-local providers, assessing their quality, and maintaining communication throughout the process was prohibitively high. The Internet facilitated the entry of intermediaries who were able to substantially mitigate such costs. The objective of this paper is to understand the sources and magnitudes of the gains to the buyers, many of whom were previously confined to small local markets, from the ability to access a large globalized pool of diverse sellers online. In addition to answering the main substantive question, this paper delivers a number of novel insights into the operation of these fast growing yet relatively little studied markets.

The design and the data generated by online markets present a set of challenges that are not yet addressed in the literature. Methodological contribution of this paper lies in developing a tractable framework as well as identification and estimation strategies that allow us to answer the substantive questions we study. Our analysis is based on data from a prominent online procurement market for programming services which is representative of many online markets that appeared during the last decade. We begin by summarizing the features of the market which are pertinent to our analysis and which differentiate this setting from other settings studied in the literature.

Transactions in this market are implemented in the form of auctions for individual projects where a buyer may either choose a seller from those who participated in the auction for his project or opt out in favor of offline opportunities.² The platform uses a multi-attribute auction format that allows each buyer to deviate from allocation based solely on price (as in standard auctions) and to choose instead a seller with a higher buyer-specific value. Such market design indicates potential for significant seller heterogeneity as well as the buyers' interest in having an option to differentiate among sellers on the basis of characteristics other than price. Under the multi-attribute format the weights for various seller attributes are not announced and may potentially differ from buyer to buyer. This means that the allocation rule is stochastic from sellers' point of view in contrast to standard auctions where the allocation rule is straightforward and deterministic.³ This feature affects sellers' pricing in an important way that to the best of

¹Services generate around 80% of the U.S. gross domestic product, a share that has increased by 20% over the last fifty years, with professional services accounting for half of this growth (according to Herrendorf, Rogerson, and Valentinyi (2009)).

²Recent analysis of eBay on-line auctions is represented by Bajari and Hortacsu (2003), Cabral and Hortacsu (2010), Akerberg, Hirano, and Shahriar (2006), Lewis (2011) and Backus and Lewis (2012), Decarolis, Goldmanis, and Penta (2014), Athey and Nekipelov (2012), Hendricks and Sorensen (2014), Akerberg, Hirano, and Shahriar (2014). For a survey of the older literature on on-line auctions, see Bajari and Hortacsu (2004).

³It is this unstructured nature of the auction format that distinguishes the service market we study from those studied in the previous auction literature, including the recent literature on "non-standard auction formats," which

our knowledge has not been previously studied in the literature.

In the data, buyers frequently chose sellers that charge prices above the lowest price submitted in the auction. Interestingly, a descriptive analysis that projects buyers' choices onto sellers' observable characteristics and prices reveals that buyers prefer sellers that charge higher prices, everything else equal. This suggests that some characteristics observed and valued by buyers are not recorded in the data available to the researcher. This is not surprising as the platform encourages and facilitates extensive buyer-seller communication related to sellers' qualifications and examples of their past work. Further, buyers may be specifically interested in certain seller characteristics such as the seller's country of origin which may indicate presence or absence of potential language barriers. However, they are likely to use most of the available seller characteristics to form an opinion about the seller's ability and, thus, to assess the quality of the product that he will deliver if chosen. The descriptive analysis mentioned above documents a positive relationship between the seller's price and the probability of winning. This indicates that the unobserved seller characteristics must be vertical, i.e., positively related to price. Thus, we refer to it as *unobserved quality* in this paper.

We formalize the features of this market in a model where each project attracts a set of sellers who submit bids for the buyer's consideration. The project is awarded to a seller who delivers the highest value over price only if it exceeds the value of the buyer's outside option. The buyer's valuation of a given seller is a function of the seller's characteristics which are weighted buyer-specifically. These buyer-specific weights are the buyer's private information and thus are not observed by sellers or the researcher. Our model assumes that the buyers are risk neutral and have full information on the sellers' characteristics. This assumption reflects features of many on-line settings such as the one we study. On-line platforms are often designed to minimize buyers' uncertainty about sellers' characteristics and to protect the participants from the ex post risks. In fact, the online platform we study maintains a database of performance-related measures, provides an arbitration service, and administers payments from an escrow account only after the buyer is satisfied with the delivered service. Hence informational concerns do not appear to be of first-order importance in this market: buyers often have access to sufficient information and it is unlikely that their risk aversion affects their decisions to a large degree. As we explain below, our estimation results appear to support this as well.

The structure of the data generated by online markets presents a set of novel challenges

assumes that the decision rule is known to the bidders, e.g., standard auctions with discrimination or preferential treatment, as in Marion (2007), Krasnokutskaya and Seim (2011), and Swinkels (2009), or scoring auctions where the award is based on a rule that aggregates several bid components, as in Athey and Levin (2001), Asker and Cantillon (2010), Asker and Cantillon (2008), and Bajari and Lewis (2011). While the multi-attribute auctions format is also prevalent in offline industry procurement, it is little studied, with the exception of Greenstein (1993, 1995).

despite being characterized by features typically arising in the discrete choice or auction settings. Specifically, while the overall number of sellers present in this market is very large, the number of sellers actively bidding for an individual project is relatively small. This means that a very large number of different choice sets are observed in the data and each choice set has only a negligibly small number of buyers choosing from this same set. As a result, “market shares” of individual sellers conditional on the choice set, which are conventionally used in analysis of discrete choice settings, cannot be precisely estimated.

The analysis is further complicated by the fact that similar to other service markets, online markets are characterized by high turnover. A large fraction of sellers leave the market after submitting only a small number of bids and winning only one or two auctions. We refer to such sellers as *transitory* sellers as opposed to *permanent* sellers who participate in many auctions.⁴ The presence of transitory sellers and their competitive pressure play an important role in online markets and in markets for services in general. For example, in our market every auction attracts several transitory sellers and projects are allocated to transitory sellers with high probability (38%), even in the presence of permanent sellers with comparable prices. Given the mechanism of buyer-seller communication it is entirely plausible that buyers are equally able to collect information about permanent and transitory sellers’ quality. Thus, the challenges created by the presence of “unobserved quality” apply to transitory sellers as well.

Our approach builds on the insights from the discrete choice and empirical auction literatures. Specifically, we treat the unobserved qualities of permanent sellers as parameters of the model (fixed effects) which may take a (relatively small) finite number of values. This naturally gives rise to a quality group structure where sellers within the same group are characterized by the same level of unobserved quality. In contrast, the transitory sellers’ qualities are modeled as random effects which could be correlated with their bids and observed characteristics. Since observable seller characteristics in our setting are discrete, both permanent and transitory sellers can be summarized by their membership in one of a finite number of groups.⁵ We derive pairwise inequality restrictions which link sellers’ relative performance to the order of their qualities. This property, in turn, allows us to recover the quality group structure of the population of permanent sellers through the testing procedure which does not require knowledge of the distributions of buyers’ weights or other model primitives. Empirical implementation of this step relies on the classification algorithm developed in Krasnokutskaya, Song, and Tang (2014).

The estimation methodology in this paper exploits the fact that once permanent sellers’ group memberships are identified, we can partially characterize the buyers’ choice sets in terms

⁴In the auction literature transitory sellers are sometimes referred to as “fringe” sellers.

⁵If an observable characteristic is continuous, our approach would require discretizing it. Such discretization is adopted in Chiappori and Salanie (2001) as well as in Ciliberto and Tamer (2009).

of the group composition of the participating permanent sellers. Our estimation procedure is based on the moments related to the permanent sellers' winning probabilities conditional on so characterized choice sets. We use this procedure to recover the quality levels associated with various quality groups as well as the distributions of buyers' weights and outside options. Finally, we rely on the structure of the sellers' pricing problem in the auction setting to recover the distribution of sellers' costs conditional on all characteristics (observable and unobservable).⁶

It is worth noting that this approach enables us to estimate the model's primitives jointly with the mixture components which capture the relationship between the transitory sellers' bids and their qualities. This offers a viable alternative to recovering this relationship from the model within the estimation routine which would be computationally infeasible in our environment, or to estimating some ad hoc parametric mixing distribution (reflecting such dependence) jointly with other elements of the model.⁷ Instead we exploit that the supports of the quality distributions in the subpopulations of permanent and transitory sellers may be linked so that in identifying the mixture components, we may use variations in the bids of permanent sellers, for whom endogeneity issue is resolved once we condition on their quality group memberships, to identify frequencies of different quality levels in the population of transitory sellers.

From our analysis of the data, we find that seller heterogeneity is important in the market for programming services. While buyers are willing to pay a substantial premium to sellers who are from certain countries or have a high level of performance measures, the premiums explained by the variations in observable characteristics are relatively small compared to the 50% of the project value premium that an average buyer is willing to pay for the increase in unobserved quality from the lowest possible to the highest possible level. Our estimates reveal heterogeneity in the unobserved seller quality within the same observed characteristics group, as well as significant differences in the distributions of the quality across different groups of country affiliation or performance measures.

We use the estimated parameters to evaluate the buyers' welfare gains from the availability of the online market. Since the buyer's outside option includes hiring from the offline local market, the difference in the net values from hiring in this market over that from the outside option provides a lower bound on the gains from market globalization. We estimate this lower bound on the average gain over a buyer's local option to be 73% of the project value. This number reflects the gain in utility from access to a more diverse set of sellers (both in terms of quality and in terms of costs). We further inquire into the source of the gains by examining the effect of having

⁶We accomplish this by relying on the inversion method first proposed by Guerre, Perrigne, and Vuong (2000) and later applied in various environments by Li, Perrigne, and Vuong (2000), Jofre-Bonet and Pesendorfer (2003), Li, Perrigne, and Vuong (2002), Krasnokutskaya (2011), Athey and Haile (2002) and others.

⁷This last estimation approach has been shown to perform poorly. See Heckman and Singer (1984) in the context of duration models.

access to international sellers facilitated by the Internet. In this analysis, we limit the diversity of choices available to buyers by replacing foreign sellers with US sellers of similar quality rank. Interestingly, this change impacts the market mainly by affecting the sellers' participation. US sellers, who are revealed in estimation to be weak competitors (they have higher costs and lower quality levels), participate at substantially lower rates relative to foreign sellers. As a result, under a counterfactual scenario a buyer faces a reduced set of alternatives which, given the higher realized prices, are less attractive to him. In the end, despite the number of potential bidders remaining the same, the reduction in the variety of potential bidders leads to the 32% decline in gains from the Internet market. We further establish that this loss of welfare arises in part due to the imperfect competition which leads to suboptimal participation by sellers. When the allocations in this market are implemented by the social planner the social welfare is reduced by less than 3% after the variety of sellers is reduced.

Further, the estimation results confirm our surmise that the buyers are informed about the transitory sellers' qualities.⁸ Indeed, we estimate a statistically significant relationship between the transitory sellers' bids and their unobserved qualities. Since such a relationship is not directly observed in the data, it cannot be explained unless the buyers observe transitory sellers' qualities. This feature also plays an important role in rationalizing allocative decisions. The model which imposes that transitory sellers are identical conditional on their observed characteristics substantially under-predicts the probability that a project will be allocated to a transitory seller (23% instead of 38% in the data), whereas our baseline model approximates this probability quite well (36%). Overall, our model fits the data well: it can explain the 75% of buyers' choices, in contrast to the 18% explained by the model without unobserved heterogeneity and the 52% by the model with unobserved heterogeneity allowed only for the permanent sellers.

Our estimation results provide a number of interesting insights into the operation of the on-line market for programming services and potentially on-line markets for other services as well. For example, we find that uncertainty about the buyers' allocation rule, which is inherent in multi-attribute auctions, induces gambling behavior from sellers with high cost realizations. This feature of the environment explains the high variance in the price distribution often observed in on-line markets by means of relatively tight cost distributions. We also find that while the distributions of costs generally appear to be stochastically increasing in unobserved quality (as well as in observed performance measures), a subset of sellers with the lowest quality levels appear to have costs comparable to the costs of high quality sellers. The latter speaks to the heterogeneity of the participants who are attracted to and are able to survive in on-line markets.

⁸Our results are consistent with findings in Cabral and Hortacsu (2010), who find that performance measures collected by the e-Bay platform were likely to serve an enforcing rather than an informative role and with Lewis (2011) who finds that the e-Bay auto market is able to deliver sufficient information about used product properties to buyers so as to overcome the "lemons problem."

The paper is organized as follows. Section 2 describes our market; the descriptive data statistics are reported in Section 3; and the basic model is summarized in Section 4. Sections 5 and 6 discuss our empirical methodology. Section 7 reports the results of the empirical analysis whereas Sections 8 and 9 describe analysis of counterfactual settings. Section 10 summarizes the findings and outlines directions for further research.

2 Market Description

We study a market mediated by an online platform that serves as a match-maker between the demand and supply for computer programming services. This company provides an environment that allows buyers (the demand side) to post job announcements. It also maintains the registry of potential bidders (the supply side). The registry provides limited information on verifiable “outside” credentials as well as history of seller’s performance on the platform. The latter includes instances of delays and disputes, as well as buyers’ feedback about working with a given seller in the form of numerical reputation scores or ratings. In the case of a dispute, the company provides professional arbitration services that ensure that a seller is paid if and only if the completed job satisfies industry standards.

The platform allocates jobs through multi-attribute auctions. Under the rules of such an auction, a buyer is allowed to take into account seller characteristics other than price. As a result, the winner is not necessarily the one who submits the lowest quote. An important feature of this mechanism is that the award rule is not announced and thus remains unknown to other market participants.

Sellers can communicate with buyers before posting price quotes. Such an exchange of messages is very common. On average, each seller submits three messages per auction in our data. A seller may attach an example of his work or a sketch of the proposed code. The number and the content of these communications are not observed by other sellers. Hence, while the buyer has an opportunity to form an opinion about each sellers’ quality, competing sellers have much less knowledge of their competitors’ characteristics. However, competitors might be able to infer a permanent seller’s quality from his long-run rate of winning.

A seller who decides to submit a bid has no information about other sellers who are also submitting bids for the same auction. His choice of price quote is thus based on his beliefs about potential competition.

3 Data Description

We have access to the data from the beginning and for the subsequent 6 years of the platform operation. The data include information on more than 600,000 projects that attract participation from close to 50,000 different sellers. For every project, we observe the description of work required, the size of the project as assessed by the platform, the deadline for the completion when it is imposed, and the location of the buyer. We also observe all bids submitted, characteristics of sellers who submitted bids, the identity of the winner, and measures of the winner’s subsequent performance.

The projects fall into several broad classes, such as system-based programming, databases, graphics programming and website design. The work is then further divided into finer categories within these classes. For example, one of the recurrent requirements is the specification that a particular programming language should be used. We focus on the projects requesting graphics-related programming which tend to be relatively homogeneous. In addition, the sellers participating in other segments of the market rarely submit bids for this type of projects due to the high degree of specialization required in this area of work. We further restrict our attention to US-based buyers who submit an overwhelming majority of projects in our dataset.

Outsourcing. The market we study is representative of the recent trend often referred to as ‘1099-economy’ where self-employed individuals contract for jobs through on-line platforms. Such markets, especially those associated with business services, are often international in their nature. Specifically, buyers participating in the market we study choose among sellers from a diverse set of countries. Table 1 lists seller countries with the largest presence in this market as well as reports seller presence by region. As the table indicates, the majority of bids are submitted by sellers from North America (16%), Eastern Europe (14%) and South and East Asia (48%). Respectively, a large fraction of the US projects are allocated to sellers from foreign countries: 32% to South and East Asia and 19% to Eastern European sellers. This market, therefore, provides the US buyers with an opportunity to acquire services of foreign sellers. In this analysis we investigate whether foreign sellers differ from those located in the US and, thus, whether buyers gain from being able to access these additional varieties of sellers through online market.

Project-Level Statistics. Table 2 provides some project-level statistics for the period covered by our data. Each row of the table summarizes a marginal distribution of the corresponding variable. Table 2 shows that a sizable number of the projects are very small (below \$150). On the other hand, some of the projects are quite big (above \$875). The projects are fairly short: the deadline for the majority of the projects is between one to three weeks. The average number of sellers submitting bids for a project is five while the median number of bidders is four. However, about 10% of projects receive more than 12 bids. The projects with a large number of bids tend

Table 1: Sellers' Composition by Country

Seller Country	Participation (share of bids)	Allocation (share of projects)
Major participants:		
US	0.132	0.205
Canada	0.025	0.044
UK	0.028	0.033
India	0.299	0.210
Pakistan	0.135	0.084
Bangladesh	0.024	0.017
Phillipines	0.024	0.013
Romania	0.096	0.121
Russia	0.018	0.025
Ukraine	0.022	0.029
By region:		
North America	0.157	0.249
South and East Asia	0.482	0.323
Eastern Europe	0.135	0.185
Number of Projects	24,116	
Number of Bids	128,580	

The entries in this table are based on a sample of projects with graphics-related programming posted by US buyers.

to be small.

The majority of buyers in our data are one-time participants. Less than 2% of buyers post multiple projects. In addition, returning buyers do not post the same type of projects. As a result, a buyer very rarely works with the same seller repeatedly.

Sellers Characteristics. Table 3 describes sellers present in the graphics segment of online market. As we emphasize in the introduction, our market attracts a large number of short-lived sellers. We define a seller's tenure as the length of time that elapses between his first and his last posting. In our data 65% of the seller population (not in the table) has a tenure under three weeks and 75% of the seller population has a tenure under eleven weeks (slightly less than three months). On the other hand, 10% of the seller population remains active for more than 144 weeks (or more than two years). The share of sellers with short tenure is larger in the beginning years but settles down, so that the distribution of tenure is almost constant over the last three years. In these years, 10% of the sellers stayed in the market for more than three years, whereas 85% of the sellers left the market in less than three months. Substantial seller turnover is an important feature of our market as well as of many other markets for services.

Table 2: Data Summary Statistics

	Mean	Std.Dev.	25%	50%	90%
Number of projects (per buyer)	1.05	0.02	1	1	1
Project Characteristics					
Size	\$525	\$221	\$150	\$500	\$875
Duration (days)	12	11	5	10	21
Number of bidders	5	3	3	4	12
Number of Projects	24,116				

The results in this table are based on a sample of projects with graphics-related programming posted by the US buyers. Duration of project is measured in days. Each row summarizes the inverse cumulative function of the corresponding variable. The number of potential bidders refers to the number of the sellers active on the platform in a given week and who specialize in graphics-related programming.

Table 3: Sellers' Characteristics

	Mean	Std.Dev.	25%	50%	75%	90%
Tenure (weeks)	23.66	54.42	1	1	11	144
Permanent Sellers						
Number of Projects Won	45.8	69	32	107	154	250
Average Score	9.8	0.061	9.7	9.87	10	10
Disputes	0.021	0.082	0	0	0	1
Delays	0.075	0.096	0	0	0	1
Fraction of Bids Resulting in Winning	0.11	0.09	0.045	0.082	0.143	0.219
Number of Bids Before First Success	11.3	6.7	5	9	17	42
Transitory Sellers						
Number of Projects Won	0.51	1.36	0	1	1	2
Average Score	9.58	0.09	9.5	9.78	10	10
Fraction of Bids Resulting in Winning*	0.12	0.13	0.035	0.076	0.151	0.25
Number of Bids Before First Success*	8.5	6.3	3	7	15	36
Number of Bids Before First Success	4.8	3.4	1	2	3	12
Number of Bids	128,580					
Number of Projects	24,116					

The results in this table are based on a sample of projects with graphics-related programming posted by US buyers. Each row summarizes the inverse cumulative function of the corresponding variable. Tenure is defined as the number of weeks between the last and the first posting of a given sellers. Disputes and delays variables reflect the number of disputes mediated by the platform and the number of missed deadlines reported to the platform respectively. For transitory sellers, 'Fraction of Bids Resulting in Winning*' and 'Number of Bids before First Success*' are computed conditional on winning at least one bid. 'Number of Bids before First Success' reports the overall number of bids a transitory sellers submitted before he wins his first project or exits the market.

In subsequent analysis we treat a bidder as permanent if he appears in our data for more

than 6 months. Otherwise, we label a bidder as transitory. In our sample, all permanent sellers complete more than twenty projects with the mean equal to 46 projects and the median to 47. The distribution of reputations scores for permanent sellers is quite tight with the mean score equal to 9.8 and the standard deviation of 0.06. Disputes and delays are rare, involving around 2% and 8% of permanent sellers respectively.

Roughly, half of transitory sellers ever registered with the platform leave the market before winning a single project but the other half completes at least one project. Overall, due to their significant presence in the market, transitory sellers win 38% of all projects (see Table 4). Further, transitory sellers' performance on the platform is comparable to that of the permanent sellers. Indeed, the distribution of average reputation scores in the population of transitory sellers who have completed at least one project is very similar to the distribution of average reputation scores in the population of permanent sellers. Additionally, the ratio of the number of bids to the number of projects won is very similar in these two populations. On the basis of these observations, it appears likely that transitory sellers are very similar to permanent sellers in their characteristics.

Further evidence supporting the conclusion above concerns the number of bids a seller submits before he wins his first project. As Table 3 indicates the distribution of this statistic in the population of transitory sellers who have completed at least one project is very similar to analogous distribution in the population of permanent sellers. At the same time the distribution of the number of bids before the first success in the full population of transitory bidders is shifted to the left relative to the distribution for the permanent sellers. This indicates that transitory sellers tend not to wait long enough for their first project. On the basis of these observations, it appears likely that transitory sellers are quite similar to permanent sellers in their characteristics but may have better outside opportunities.

Finally, very little information about transitory sellers is publicly available. Indeed, public information is released when a seller completes a project, and transitory sellers usually complete one or two projects and leave the market. It is plausible, therefore, that competing sellers are not informed about transitory sellers' qualities. The situation is different for permanent sellers since the market may infer their quality from the long-run rate of their successes.

Determinants of Buyers' Choices. Table 4 further documents that the multi-attribute feature of allocation mechanism is strongly supported in the data. Indeed, in our sample, 72% of the projects are awarded to a seller who quotes a price above the lowest price submitted in the auction. When such a seller is chosen his price on average exceeds the lowest price submitted in the auction by 56%. These results indicate that buyers consider seller characteristics other than price when choosing a winner. Thus, a model that takes sellers' heterogeneity into account is

Table 4: Auction Outcomes

	Mean	Std.Dev.
Buyers' Choice:		
Awarded to Lowest Bid	0.28	(0.20)
(Winning Bid-Lowest Bid)/Lowest Bid	0.56	(0.29)
Money-Left-on-the-Table	0.33	(0.26)
Share of Projects Won:		
Permanent sellers	0.59	
Transitory sellers	0.38	
Number of Projects	24,116	

The results in this table are based on a sample of projects with graphics-related programming posted by the US buyers. Standard deviations are reported in parenthesis. ‘Money-Left-on-the-Table’ variable is computed as the difference between the second lowest bid and the lowest bid divided by the lowest bid.

required to study this environment. The “Money-Left-on-the-Table” measure, which is computed as the average difference between the second lowest and the lowest bid submitted in the auction divided by the lowest bid, is equal to 0.33 in our sample. In the context of standard auction markets this statistic is often interpreted as indicative of the presence of private information about sellers’ costs. The interpretation is less straightforward in the multi-attribute setting. Indeed, the difference between the lowest and the second to the lowest bid may arise due to the difference in private costs but also due to the premium charged by sellers for the characteristics preferred by buyers.

We use a multinomial logit model to further explore how buyers’ choices are influenced by sellers’ characteristics. Here we assume that the award decision ($Y_{j,l} \in \{0, 1\}$) depends on the buyer’s net value from a specific alternative (seller), $Y_{j,l}^*$, which is modeled as a linear function of seller characteristics, $X_{p,j,l}$ (the number of ratings, delays, disputes and the average reputation score), seller location dummies, $\mu_{c(j)}$, and a seller’s bid, $B_{j,l}$:

$$Y_{j,l}^* = \sum_p X_{p,j,l} \beta_p + \gamma B_{j,l} + \mu_{c(j)} + \epsilon_{j,l}. \quad (1)$$

The project is awarded to bidder j , $Y_{j,l} = 1$, if and only if $Y_{j,l}^* \geq 0$ and $Y_{j,l}^* \geq Y_{i,l}^*$ for all $i \neq j$ who are present in the auction; $Y_{j,l}$ is equal to zero otherwise. The results of this analysis are reported in Tables 5.

We estimate the price coefficient to be positive and statistically significant. This result suggests an omitted variable bias since, in most markets, buyers prefer to pay less, other things equal. This means that some additional characteristic, not recorded in the data, affects buyers’ choice in conjunction with the price, location and performance measures. Such an omitted vari-

Table 5: Multinomial Logit

Variables	Estimates	Std.Errors
Constant	1.069	4.461
Number of Scores (NS):		
$1 \leq NS \leq 3$	1.331 ***	0.171
$3 < NS \leq 6$	1.496 ***	0.368
$6 < NS \leq 12$	1.511 ***	0.408
$NS \geq 12$	1.513 ***	0.517
Average Score:		
Average Score if $NS \leq 3$	0.016	0.041
Average Score if $3 < NS \leq 6$	0.048	0.075
Average Score if $6 < NS \leq 12$	0.058	0.094
Average Score if $NS \geq 12$	0.158	0.119
Disputes	-0.541 ***	0.064
Delays	0.232	0.164
Price	1.937 ***	0.073
Seller Country Dummies		
Number of Projects	Yes	
Number of Bids	24,116	
	128,580	

The results in this table are based on a sample of projects with graphics-related programming posted by US buyers. Country dummies for the seller countries with bid share exceeding 0.05% are included. Omitted category are US sellers who have not completed any projects yet. Above ‘***’ indicates significance at 1% significance level.

able should be positively aligned with the price and is, therefore, some vertical characteristic such as quality. Thus a model that rationalizes this pattern should allow for unobserved quality-like sellers’ attributes.

As for the other variables, the results indicate that having earned at least a few reputation scores increases the probability of winning in a statistically significant way. After that, however, the impact of the subsequent scores flattens out. Similarly, the impact of an extra point in the average reputation score does not appear to matter in a statistically significant way. Disputes (arbitrations) decrease future probabilities of winning whereas delay have no statistically significant effect.

To summarize, the preliminary analysis of our data indicates that (a) the buyers’ utility from hiring a given seller should non-trivially depend on sellers’ attributes; (b) it is important to account for the presence of a large number of transitory sellers; (c) the model should allow for the presence of an unobserved quality-like sellers’ attribute for permanent as well as for transitory sellers.

4 Model

Let \bar{N} denote the set of sellers who operate in an on-line programming market. Each seller $j \in \bar{N}$ is characterized by a vector of attributes $x_j \in \mathcal{X} \equiv \{\bar{x}_1, \bar{x}_2, \dots, \bar{x}_M\}$ and a quality index q_j which takes values from a discrete set $\mathcal{Q}(x_j) = \{\bar{q}_1(x_j), \dots, \bar{q}_{K(x_j)}(x_j)\}$. Notice that distribution of quality indices is x -specific in the sense that both the number of quality levels, $K(x_j)$, and the quality levels themselves may depend on x .

Seller j additionally belongs to one of two types $\rho_j \in \{p, t\}$ where ‘ p ’ denotes permanent and ‘ t ’ denotes transitory sellers. A permanent seller’s quality is known to all market participants; a transitory seller’s quality is his private information which is drawn from a distribution with the discrete support $\mathcal{Q}(x)$ for a seller characterized by a vector of attributes x . Each transitory seller reports his quality draw to the buyer when submitting a bid.

A buyer l seeks to procure services for a single indivisible project using a multi-attribute auction. Throughout the paper we use l to index a buyer or his project. Each project is associated with a set $N_l \subset \bar{N}$ of potential bidders who are available and interested in providing service.⁹ Each seller in N_l decides whether to participate in the auction for project l and what bid $(B_{j,l})$ to submit if he participates.

Buyers’ Choice. Let $A_l \subset N_l$ denote the set of sellers who participate in bidding for project l . The buyer either chooses a seller from A_l , or opts for an outside option that gives him a payoff $U_{0,l}$. The buyer is informed about the attributes x_j and the quality index q_j for each seller $j \in A_l$. The buyer’s payoff from choosing j as the service provider is

$$U_{j,l} = \alpha_l q_j + x_j \beta_l + \epsilon_{j,l} - B_{j,l},$$

where the random vector $\gamma_l = (\alpha_l, \beta_l, \epsilon_l)$ with $\epsilon_l \equiv \{\epsilon_{j,l} : j \in N_l\}$ denotes the buyer’s utility weights for seller characteristics. The vector γ_l and the outside option $U_{0,l}$ summarize the buyer’s preferences. The buyer chooses an option that maximizes his payoff. In keeping with the definition of a multi-attribute auction, sellers do not observe the utility weights or the outside option of a specific buyer, and consider these to be random draws from the corresponding distribution. Hence the quality of transitory sellers is assumed to be independent of buyers’ preference $(\alpha_l, \beta_l, \epsilon_l, U_{0,l})$.

Sellers’ Strategies. We use I_{N_l} to denote the composition of the set N_l with respect to sellers’ quality and attributes based on information that is common knowledge among sellers. We will formally define I_{N_l} in Section 5.2.2.

All potential bidders for project l know the composition of the set N_l , I_{N_l} . Each seller $j \in N_l$ draws a private entry cost $E_{j,l}$ from a distribution $F_{E|x,q}$ where $x_j = x$, $q_j = q$, and decides

⁹We explain how this feature is interpreted in the context of our data in Section 7.

whether to participate in bidding or not based on $E_{j,l}$ and I_{N_l} . If the seller decides to participate, he pays $E_{j,l}$, then draws a private cost for completing the project $C_{j,l}$ from a distribution $F_{C|x,q}$ and submits a bid $B_{j,l}$ to the buyer. We assume that private costs are independent across sellers. Sellers who do not participate obtain zero payoffs. Participation decisions of sellers are not observed by their competitors.

Equilibrium. Let \mathcal{I} denote the collection of all possible compositions of subsets of sellers in \bar{N} . For seller j , his *pure strategy* consists of two functions: an entry strategy $\tau_j : \mathbb{R}_+ \times \mathcal{I} \rightarrow \{0, 1\}$ which maps his realization of entry costs, and the composition of the set of potential bidders for auction l , $I_{N_l} \in \mathcal{I}$, into participation decision $\{0, 1\}$ with 1 denoting the decision to participate; and a bidding strategy $\sigma_j : \mathbb{R}_+ \times \mathcal{I} \rightarrow \mathbb{R}_+$ which maps his realization of project costs and the composition of the set of potential bidders into seller's bid. Let (τ_{-j}, σ_{-j}) denote a strategy profile of the sellers in N_l other than j ; and given composition $I_{N_l} \in \mathcal{I}$ of sellers in N_l , let $A_{j,l}(I_{N_l}; \tau_{-j})$ be the set of sellers *other than* j who decide to participate given their entry costs and strategies τ_{-j} . Note that by construction $A_{j,l}$ is a random set that depends on $E_{-j,l}$, which we suppress in notation for simplicity.

Given a composition $I_{N_l} \in \mathcal{I}$ of potential bidders in auction l , the ex ante payoff for a seller j who decides to participate, draws a project cost c and quotes b when other sellers use strategies (τ_{-j}, σ_{-j}) is:

$$\Pi_j(b, c; I_{N_l}; \tau_{-j}, \sigma_{-j}) = (b - c)P_j(b; I_{N_l}; \tau_{-j}, \sigma_{-j}) \quad (2)$$

where $P_j(b; I_{N_l}; \tau_{-j}, \sigma_{-j})$ denotes the conditional probability that bidder j wins the auction and is equal to

$$\Pr \left(\max \left(U_{0,l}, \max_{i \in A_{j,l}(I_{N_l}; \tau_{-j})} \{ \alpha_l q_i + x_i \beta_l + \epsilon_{i,l} - \sigma_i(C_{i,l}, I_{N_l}) \} \right) \leq \alpha_l q_j + x_j \beta_l + \epsilon_{j,l} - b \right). \quad (3)$$

The probability in (3) is with respect to the joint distribution of $(\alpha_l, \beta_l, \epsilon_l, U_{0,l})$, $C_{-j,l}$ and $E_{-j,l}$.

We follow the convention in the literature and focus on type-specific equilibria where sellers of the same type $\theta \equiv (\rho, x, q)$ with $\rho \in \{p, t\}$, use the same strategy. A *type-symmetric pure-strategy Bayesian Nash equilibrium (p.s.BNE)* is a profile of strategies (τ^*, σ^*) such that for any j with type $\theta_j = \theta$, $\tau_j^* = \tau_\theta$ and $\sigma_j^* = \sigma_\theta$, where for each c and e ,

$$\begin{aligned} \sigma_\theta(c, I_{N_l}) &= \arg \max_{b \geq 0} (b - c)P_j(b; I_{N_l}; \tau_{-j}^*, \sigma_{-j}^*) \text{ and} \\ \tau_\theta(e, I_{N_l}) &= 1 \{ e \leq \mathbf{E} [\Pi_j(\sigma_\theta(C_{j,l}, I_{N_l}), C_{j,l}; I_{N_l}; \tau_{-j}^*, \sigma_{-j}^*)] \}. \end{aligned}$$

Here $\mathbf{E} [\Pi_j(\sigma_\theta(C_{j,l}, I_{N_l}), C_{j,l}; I_{N_l}; \tau_{-j}^*, \sigma_{-j}^*)]$ summarizes the expected profit of seller j conditional on participation where expectation is further taken with respect to the distribution of own project costs, $C_{j,l}$.

Discussion. Notice that the multi-attribute environment substantially differs from a scoring auction, another mechanism which allows buyers to take into account seller’s attributes other than price at the allocation stage. First, in our case, a seller’s quality q_i and characteristics x_i are fixed as opposed to being part of the sellers’ strategic choices as in scoring auctions. Second, the allocation rule in scoring auctions is explicitly specified before bidding begins whereas buyer preferences for quality and seller attributes are not known to sellers in a multi-attribute auction. Thus in the multi-attribute setting the allocation rule is stochastic from a seller’s point of view. The latter property has an important implications for seller strategies, which to the best of our knowledge has not been studied in the empirical literature.

5 Identification

In this section we discuss how the primitives can be recovered from the data. Suppose that data contain information on many auctions that proceed as in the model described above. For each auction, the researcher observes the set of potential and actual bidders, submitted prices, and the buyer’s choice. For every seller the researcher observes the seller’s x -attributes but not his quality. Furthermore, the researcher observes whether each seller is permanent or transitory.

The model primitives to be identified include: (a) the quality of each permanent seller; (b) the distribution of a transitory sellers’ quality given observable characteristics; (c) the distribution of buyers’ utility weights $\alpha_l, \beta_l, \epsilon_l, U_{0,l}$; and (d) the distribution of sellers’ participation and project costs given sellers’ observable and unobservable characteristics.

Recovering the primitives of the model from the available data is quite challenging. To see this let us first consider an environment without transitory sellers. In this setting we only need to focus on recovering the quality levels (fixed effects) of permanent sellers. Recall that in a traditional discrete choice setting fixed effects associated with different alternatives are identified from the observed probabilities that a given alternative is chosen conditional on the choice set. In our setting choice sets are buyer-specific since sellers’ participation varies across auctions. Due to the large number of sellers, conditional choice probabilities cannot be precisely estimated. To get a sense of magnitudes consider that the number of permanent sellers present in the market for a given type of work is around 300 to 500 whereas only 2 or 3 permanent sellers participate in any given auction. This means that the number of possible choice sets is at least $C_{300}^3 = \frac{300!}{3!297!} = 8,910,200$ which exceeds the number of projects we have in our dataset. In fact, the *highest* number of projects sharing the same set of participating *permanent* sellers in our data is five. One way to deal with this issue would be to consider probabilities that aggregate

over buyers' choice sets, such as:

$$\begin{aligned}\Pr(j \text{ wins} | j \in A_l) &= \sum_{a: j \in a} \Pr(j \text{ wins} | A_l = a) \Pr(A_l = a, j \in A_l), \text{ and} \\ \Pr(j \text{ wins} | B_l = b, j \in A_l) &= \sum_{a: j \in a} \Pr(j \text{ wins} | B_l = b, A_l = a) \Pr(A_l = a, j \in A_l),\end{aligned}$$

where the sum above is over the choice sets a that contain j . While such aggregation is appealing, it is far from obvious that these moments could be used to identify seller-specific fixed effects and the distribution of buyers' tastes. Specifically, the invertibility argument underlying the standard approach (the most well-known exposition can be found in Berry, Levinsohn, and Pakes (1995)) does not apply to these moments because the probability of observing a given choice set, $\Pr(A_l = a | j \in A_l)$, depends in the model on the qualities of potential (and, in the consequence, actual) sellers. So further insight is necessary on how to achieve invertability in this context.¹⁰ It also has to be established that such moments allow us to exploit the exogenous variation present in the data to recover the distribution of buyers tastes. Further, even if this mechanism would work in theory, it is not at certain that it would perform well in practice given that the weighting probabilities used in aggregation above, $\Pr(A_l = a | j \in A_l)$, are very small. It might be preferable to consider an entirely different basis for aggregation which would maximize the performance of the estimator given the available data structure.

Let us now return to the realistic setting where transitory sellers differ in their qualities and these qualities are observable to buyers. Note that transitory sellers are an important part of this market since in the data every projects attracts several transitory sellers, and has a 38% chance of being allocated to a transitory seller. Buyers often chose transitory sellers over permanent ones even when the prices are comparable.

In contrast to permanent sellers we cannot use transitory sellers' identities as proxies for their quality. Thus, the information which underlies buyer's choice is not observed in the data. Instead, a researcher has to deal with a mixture problem where the probability distribution over the transitory sellers' qualities depends on these sellers' bids and observable attributes. More specifically, suppose the support of a transitory seller's quality $Q_{h,l}$ is $\{\bar{q}_1, \bar{q}_2\}$ and let x and b be the vectors of observable attributes and bids characterizing the entrants in the auction respectively. Then, the probability that the buyer chooses a permanent seller j while his choice set includes a single transitory seller h , $\Pr(j \text{ wins} | B_l = b, x)$, is a mixture of the following form:

$$\sum_{s=1,2} \Pr(j \text{ wins} | Q_{h,l} = \bar{q}_s, B_l = b, x) \Pr(Q_{h,l} = \bar{q}_s | B_l = b, x).$$

The mixing weights $\Pr(Q_{h,l} = \bar{q}_s | B_l = b, x)$ are unknown and correlated with the conditional

¹⁰For example, it is possible that the inversion could be made to work if we use empirical probabilities of observing different choice sets in the expression on the right-hand side.

choice probability through the bid vector b and attributes vector x . In practice, even writing the choice probability in this mixture form is not straightforward, because we do not observe the support of $Q_{h,l}$.

One might attempt to deal with this problem by solving for mixing probabilities from the model within the estimation routine. However, solving one such bidding and participation game would take a long time and solutions can be very fragile if parameter values are far from the truth. Further, a large number of possible seller types results in a very large number of possible choice sets for which the problem would have to be solved. These issues combined make such approach computationally infeasible. Alternatively, one may adopt an ad hoc functional form assumption for the mixing distributions and attempt to recover them jointly with other primitives of the model. It is doubtful that separate identification of these components can be established formally. In practice, such approach has been shown to perform poorly.^{11,12}

To overcome these difficulties we propose a methodology where we first classify permanent sellers into groups of equal quality. Next, we use this grouping to recover other primitives of the model. Such an approach facilitates the analysis in several ways. First, the buyers' choice sets may now be represented in terms of the participating sellers' group memberships rather than their identities. This offers a natural way for partial aggregation of buyers' choice sets and permits exploiting variation in buyers choices across choice sets which could be harnessed to identify the distribution of buyers' tastes. Further, this approach when combined with the second insight, that the support of the transitory sellers' quality distribution can be linked to the support of the permanent sellers' quality distribution, allows to identify payoffs associated with various bundles of (x, q) -attributes separately from the identification of the frequencies with which such bundles are observed in the population of transitory sellers. In the next section, we illustrate this identification strategy using a simple model.

5.1 Heuristics

Identifying the Group Structure. Consider sellers i and j with $x_i = x_j$ who participate in two separate but *ex ante* identical auctions (i.e., the project characteristics and the realized set of their competitors are the same, and both i and j belong to the set of potential bidders) and submit equal bids. Under such circumstances a seller with the higher value of q has a higher chance of winning. The ranking of winning probabilities is preserved in aggregation over possible sets of their competitors as long as the probability of encountering a given set of competitors is

¹¹See Heckman and Singer (1984) for details.

¹²A researcher may also consider an approach proposed by Kasahara and Shimotsu (2009) in the context of a dynamic discrete choice model. However, the model considered by these authors does not readily map into our environment so the applicability of this method, if possible, is far from obvious.

the same for both sellers. This condition holds if, for example, the pool from which competitors are drawn does not include either i or j .

More specifically, define

$$r_{i,j}(b) \equiv \Pr(i \text{ wins} \mid B_{i,l} = b, i \in A_l, j \notin A_l, i, j \in N_l) \quad (4)$$

where A_l denotes the set of entrants.¹³ Then

$$\begin{aligned} r_{i,j}(b) &> r_{j,i}(b) \text{ if and only if } q_i > q_j, \\ r_{i,j}(b) &< r_{j,i}(b) \text{ if and only if } q_i < q_j \text{ and} \\ r_{i,j}(b) &= r_{j,i}(b) \text{ if and only if } q_i = q_j. \end{aligned}$$

(A formal statement and conditions are provided in Proposition 1 in Section 5.2.) As long as the conditional winning probabilities are identified from data, we can use them to order sellers i and j nonparametrically with respect to their qualities. By implementing such comparison for every pair of permanent sellers within x -group, the quality ranking of the sellers within this group can be recovered.¹⁴ This identifies the quality group structure.

A Simple Example. Let us see how identifying the quality classification of the permanent sellers leads to the identification of the model primitives. Consider a simple setting with two groups of sellers defined by observable characteristics \bar{x}_1 and \bar{x}_2 ; each observable group is further partitioned into two unobservable subgroups based on quality levels $\bar{q}_1(\bar{x}_1), \bar{q}_2(\bar{x}_1)$ and $\bar{q}_1(\bar{x}_2), \bar{q}_2(\bar{x}_2)$ respectively. Some of the sellers are labeled as permanent and others as transitory. For simplicity, assume that all components of buyers' weights are mutually independent, which is relaxed in our formal identification results in Section 5.2 below. The remaining aspects of the model are just as described in Section 4.

Suppose further that a large number of sellers in each observable group is present. The total number of different choice sets (consisting of 3-4 sellers) defined in terms of the sellers' identities can be very large. This number is drastically reduced when the choice sets are defined in terms

¹³One could use an alternative index based on the winning probability conditional on $i \in A_l, j \in A_l, B_{i,l} = b$ and $B_{j,l} = b$ with the same b . In our data, for many pairs of bidders, there is only a small number of auctions where both i and j participate and submit similar bids. Hence precise estimation of this alternative index is much more problematic than the estimation of the index $r_{i,j}(b)$ which we propose here. In this paper we do not pursue this alternative strategy.

¹⁴Intuitively, if comparisons for all pairs of permanent sellers are available, we can always split a given x -group into two subgroups where the first subgroup consists of the sellers with the lowest quality among all the sellers in the x -group and the second subgroup consists of the remaining sellers. Then we split this second subgroup similarly into two further subgroups so that the first further subgroup consists of the lowest quality sellers within this second subgroup and the other further subgroup consists of the rest of the sellers. By continuing this process, we can identify the quality group structure.

of groups instead of identities.¹⁵ However, using such choice sets is possible only when we have a way to uncover the link between the sellers' identities and their unobserved quality groups, using, for example, a method such as one we propose in this paper.

Now suppose that the permanent sellers' group memberships are identified. Next, we identify the payoff structure (specifically the distributions of ϵ_l , α_l , β_l and quality levels associated with each quality group, $\bar{q}_k(x_m)$ where $k = 1, 2$ and $m = 1, 2$). After that, we identify the remaining model primitives (such as F_{U_0} and $F_{Q^t|B^t}$).

Identification of the Distribution of Payoffs. Identification of payoff components relies on the variation in the buyers' choices in response to the variation in the buyers' choice sets (defined in terms of the (x, q) -group memberships of permanent active bidders).

Let us demonstrate how to identify the distribution of a stochastic component ϵ_l as an example. For this, we focus on auctions which attract at least two permanent sellers with the same observed characteristics $x_i = x_j$ and the same unobserved quality $q_i = q_j$. We also allow for a transitory seller h to be present since an overwhelming majority of auctions attract transitory participants. The payoffs corresponding to different options available to this buyer are as follows:

$$\begin{aligned} U_{i,l} &= \alpha_l q_i + \beta_l x_i - B_{i,l} + \epsilon_{i,l}, \\ U_{j,l} &= \alpha_l q_j + \beta_l x_j - B_{j,l} + \epsilon_{j,l}, \\ U_{h,l} &= \alpha_l Q_{h,l} + \beta_l x_1 - B_{h,l} + \epsilon_{h,l} \text{ and } U_{0,l}. \end{aligned}$$

We cannot identify the scale of α_l and the quality levels at the same time. Hence we set $\mathbf{E}[\alpha_l] = 1$ as a scale normalization.

When $B_{i,l} = -t_2$ and $B_{j,l} = t_1 - t_2$ for some t_1, t_2 , the buyer chooses seller i with probability

$$\begin{aligned} & \Pr(i \text{ wins } l | B_{i,l} = -t_2, B_{j,l} = t_1 - t_2, B_{h,l} = b_{h,l}) \\ &= \Pr(\epsilon_{j,l} - \epsilon_{i,l} \leq t_1 \text{ and } Y_{i,l}(x_h) - \epsilon_{i,l} \leq t_2 | B_{i,l} = -t_2, B_{j,l} = t_1 - t_2, B_{h,l} = b_{h,l}) \\ &= \Pr(\epsilon_{j,l} - \epsilon_{i,l} \leq t_1 \text{ and } Y_{i,l}(x_h) - \epsilon_{i,l} \leq t_2 | B_{h,l} = b_{h,l}) \equiv F(t_1, t_2 | B_{h,l} = b_{h,l}), \end{aligned}$$

where $Y_{i,l}(x_h) = Y_{0,l}(x_h) - \alpha_l q_i(x_i) - \beta_l x_i$ and $Y_{0,l}(x_h) = \max\{\alpha Q_{h,l} + \beta_l x_h - B_{h,l} + \epsilon_{h,l}, U_{0,l}\}$. The last equality follows because the bids are independent of each other and of buyers' taste components, and $\epsilon_{j,l} - \epsilon_{i,l}$ and $Y_{i,l}(x_h) - \epsilon_{i,l}$ are independent of $(B_{i,l}, B_{j,l})$ due to our model assumptions.

The winning probability on the left hand side is identified from the data. Hence the joint CDF F on the right hand side is identified, and so are the marginal distribution of $\epsilon_{j,l} - \epsilon_{i,l}$ and that

¹⁵Specifically, if 100 different sellers are present, then the overall number of choice sets consisting of three sellers and defined in terms of their identities is $C_{100}^3 = 100!/97!3! = 100 * 99 * 98/6 = 161,700$. However, if we define choice sets in terms of groups we have only $4^3 = 64$ choice sets.

of $Y_{i,l}(x_h) - \epsilon_{i,l}$. Note that $\epsilon_{j,l}$, $\epsilon_{i,l}$ and $Y_{i,l}(x_h)$ are independent. Hence the marginal distributions of $\epsilon_{j,l}$, $\epsilon_{i,l}$ and $Y_{i,l}(x_h)$ are identified up to a location normalization if the support of $(B_{i,l}, B_{j,l})$ is large enough.^{16,17} (The latter support condition is not needed when the distributions of $(\epsilon_{i,l}, \epsilon_{j,l})$ and (α, β, U_0) are from a parametric family.)

Notice that this intuition continues to hold under the aggregation over many different choice sets that include two permanent sellers from the same (x, q) -group.

The identification of the quality levels $\bar{q}_1(\bar{x}_1)$, $\bar{q}_1(\bar{x}_2)$, $\bar{q}_2(\bar{x}_2)$ and the distribution of α_l and β_l can be done similarly. Specifically, we identify the distribution of $\alpha_l(\bar{q}_1(x) - \bar{q}_2(x))$ by applying similar arguments to the subset of auctions with choice sets consisting of permanent sellers i and j from the same observable group x but different quality groups, $\bar{q}_1(x)$ and $\bar{q}_2(x)$, a transitory seller h and the outside option. The mean of the distribution $\alpha_l(\bar{q}_1(x) - \bar{q}_2(x))$ identifies $\bar{q}_1(x) - \bar{q}_2(x)$ under the normalization $\mathbf{E}[\alpha_l] = 1$. Further, we consider the subset of auctions with choice sets consisting of permanent sellers i and j belonging to the lowest quality groups associated with different observable groups, \bar{x}_1 and \bar{x}_2 , a transitory seller h from \bar{x}_2 and the outside option. Then we can identify the distribution of β_l under additional normalization that $\bar{q}_2(\bar{x}_1) = \bar{q}_2(\bar{x}_2) = 0$, i.e., the lowest quality in each observable group is normalized to be zero. This restriction on the lowest quality levels can be relaxed if β_l is a fixed parameter (equal to β) rather than a random variable. In this case it is enough to normalize the lowest quality level for a single observable group, e.g., $\bar{q}_2(\bar{x}_1)$, whereas $\bar{q}_2(\bar{x}_2) - \bar{q}_2(\bar{x}_1)$ can be recovered.

Next, let us explain how we identify the distribution of the payoff from the outside option $U_{0,l}$, and the conditional quality distribution $\Pr(Q_{h,l} = \bar{q}_k(x_h) | B_{h,l} = b_{h,l}, X_h = x_h)$ for transitory sellers.

The Distribution of Outside Option and Transitory Sellers' Conditional Quality Distributions. For simplicity, let us focus on a set of auctions that attract one permanent bidder and one transitory bidder, and assume that $U_{0,l}$, β_l and α_l are independent.¹⁸ Note first that when $(\bar{q}_1(\bar{x}_1), \bar{q}_1(\bar{x}_2))$ and the distributions of $Y_{1,l}(x_h) = Y_{0,l}(x_h) - \alpha_l q_1(x_1) - \beta_l x_1$, α_l , and β_l are identified as above then the distribution of $Y_{0,l}(x_h)$, which is the payoff from the losing alternative (i.e., the maximum of the outside option and the utility from the transitory seller), is also identified. Let us now argue that the distributions of $Y_{0,l}(x_h)$, one for $x_h = \bar{x}_1$ and the other for $x_h = \bar{x}_2$, uniquely determine the distribution of the outside option and the conditional distribution of a transitory seller's quality.

¹⁶This is the consequence of Kotlarski Theorem. See Rao (1992). In short, it states that “the distributions of three mutually independent components X_1 , X_2 and X_3 are identified up to a location normalization if a joint distribution of $X_1 + X_3$, $X_2 + X_3$ is known.”

¹⁷The formal support requirements are stated in Section 5.2.2 and are further discussed in the Appendix.

¹⁸Our formal identification result does not require this independence assumption.

Specifically, since $U_{0,l}$, β_l and α_l are independent, the payoff to the outside option ($U_{0,l}$) and the payoff to the transitory seller ($U_{h,l}$) are independent. Hence, for each fixed number y_0 ,

$$\Pr(Y_{0,l}(x_h) \leq y_0 | B_{h,l}, X_h = x_h) = \Pr(U_{0,l} \leq y_0) \Pr(U_{h,l}(x_h) \leq y_0 | B_{h,l}, X_h = x_h). \quad (5)$$

From this, we obtain two equations by setting x_h equal to \bar{x}_1 first and then \bar{x}_2 in the above expression. After rearranging terms in the equations, we obtain the following functional equation:

$$\begin{aligned} g_1(y_0; B_{h,l}) \Pr(Q_{h,l} = q_1(\bar{x}_2) | B_{h,l}, X_h = \bar{x}_2) - g_2(y_0; B_{h,l}) \Pr(Q_{h,l} = q_1(\bar{x}_1) | B_{h,l}, X_h = \bar{x}_1) \\ = g_3(y_0; B_{h,l}), \end{aligned}$$

where for each $B_{h,l} = b_h$, $g_j(y_0; b_{h,l})$, $j = 1, 2, 3$, are known functions of y_0 , whereas $\Pr(Q_{h,l} = q_1(\bar{x}_j) | B_{h,l}, X_h = \bar{x}_j)$, $j = 1, 2$, are the unknown probabilities we seek to identify.¹⁹ These latter unknown probabilities are over-identified since we have infinitely many linear equations associated with different values of y_0 . Once these probabilities are identified, the distribution of the payoff from the outside option $U_{0,l}$ is identified from the equation (5).

5.2 Formal Results

This section summarizes formal details of our identification results.

5.2.1 Quality Classification of Permanent Sellers

The first step is to recover the classification of permanent sellers from the data. Let N_l^t denote the set of transitory potential bidders in auction l .

Assumption 1 (i) (a) For each $j \in N_l$, $E_{j,l}$ and $C_{j,l}$ are independent. (b) For each $j \in N_l$, $(\alpha_l, \beta_l, \epsilon_l, U_{0,l})$ and $(E_{j,l}, C_{j,l})$ are independent. (c) $C_{j,l}$'s are continuously distributed and i.i.d. across the sellers $j \in N_l$, with its distribution (and potentially its support) depending on (x_j, q_j) but not on ρ_j . (ii) If we let $\tilde{\alpha}_l \equiv (\alpha_l, U_{0,l})$, then $\tilde{\alpha}_l, \epsilon_l$, and β_l are mutually independent, each having a connected support, and $\epsilon_{j,l}$'s are i.i.d. across the sellers. (iii) The quality of transitory potential bidders in auction l , denoted by $Q_l^t \equiv (Q_{j,l} : j \in N_l^t)$, is independent of $(\alpha_l, \beta_l, \epsilon_l, U_{0,l})$ and $Q_{j,l}$ are independent across $j \in N_l^t$.

The independence between $(\alpha_l, U_{0,l})$ and β_l is not necessary for Propositions 1, 2 and the first part of Proposition 3 below. Let \mathcal{B}_i denote the support of the price quoted by a seller i in a pure-strategy Bayesian Nash equilibrium (BNE).

¹⁹Use $J_k(b_h, x)$ to denote $\Pr(U_h(x) \leq y_0 | b_h, x, Q_h = q_k(x))$ for $k = 1, 2$. Then,

$$\begin{aligned} g_1(y_0; b_h) &= \Pr(Y_0(\bar{x}_1) \leq y_0 | b_h, \bar{x}_1) [J_1(b_h, \bar{x}_2) - J_2(b_h, \bar{x}_2)] \\ g_2(y_0; b_h) &= \Pr(Y_0(\bar{x}_2) \leq y_0 | b_h, \bar{x}_2) [J_1(b_h, \bar{x}_1) - J_2(b_h, \bar{x}_1)] \\ g_3(y_0; b_h) &= \Pr(Y_0(\bar{x}_2) \leq y_0 | b_h, \bar{x}_2) J_2(b_h, \bar{x}_1) - \Pr(Y_0(\bar{x}_1) \leq y_0 | b_h, \bar{x}_1) J_2(b_h, \bar{x}_2). \end{aligned}$$

Assumption 2 For any permanent sellers i, j with $x_i = x_j$, $\mathcal{B}_i \cap \mathcal{B}_j$ contains an interval with a non-empty interior.

For each pair of permanent sellers i and j , let the index $r_{i,j}$ be defined as in (4):

$$r_{i,j}(b) \equiv \Pr(i \text{ wins} \mid B_{i,l} = b, i \in A_l, j \notin A_l, i, j \in N_l)$$

Note that A_l is a random set of entrants in the bidding stage, and in the expression for $r_{i,j}$ we are not conditioning on the identities of entrants other than i and j . In practice, we construct this index by pooling all the auctions under the conditioning event specified above.

Proposition 1 Suppose that Assumptions 1 - 2 hold. Then for each pair of permanent sellers i, j with $x_i = x_j$ and for each $b \in \mathcal{B}_i \cap \mathcal{B}_j$,

$$\text{sign}(r_{i,j}(b) - r_{j,i}(b)) = \text{sign}(q_i - q_j),$$

where $\text{sign}(z) \equiv 1\{z > 0\} - 1\{z < 0\}$, for $z \in \mathbb{R}$.

The proposition says that for each pair of permanent sellers i and j , we can determine their quality ordering by looking at the sign of $r_{i,j}(b) - r_{j,i}(b)$. Thus if such a sign is available for each pair of permanent sellers, we can identify the quality group structure among the permanent sellers. See a companion paper Krasnokutskaya, Song, and Tang (2014) for more details.

5.2.2 Quality Indices and the Distribution of Buyer Preferences

The next step is to identify the quality levels and the distribution of α_l and ϵ_l . Let \overline{N}^p be the total set of permanent sellers in the population, and \overline{N}^t that of transitory sellers in the population. For a generic set $a = a^p \cup a^t$ with $a^p \subset \overline{N}^p$ and $a^t \subset \overline{N}^t$, we denote (x, q) -specific subsets of a as follows:

$$\begin{aligned} a^p(x, q) &\equiv \{i \in a^p : x_i = x \text{ and } q_i = q\}, \text{ and} \\ a^t(x) &\equiv \{i \in a^t : x_i = x\}. \end{aligned}$$

Hence $a^p(x, q)$ denotes the (x, q) -subgroup of the permanent sellers in a , and $a^t(x)$ the x -subgroup of the transitory sellers in a . For each (x, q) , define $\lambda^p(a, x) \equiv (|a^p(x, q)| : q \in \mathcal{Q}(x))$ and $\lambda^t(a, x) \equiv |a^t(x)|$, i.e., the collections of cardinalities of subgroups. Then we define the *composition* of set a as

$$I_a \equiv (\lambda^p(a, x), \lambda^t(a, x))_{x \in \mathcal{X}}.$$

Hence the composition I_a depends on the set a of sellers only through the sizes of the subgroups, not through their identities.

For any generic set $a \subset \overline{N}$, let $b_a \equiv (b_s : s \in a)$. Next, for a given seller $j \in N_l$ and $a \subset N_l \setminus \{j\}$, define a random variable $Y_{j,l}(b_a, I_a)$ as the maximum of $U_{0,l} - \alpha_l q_j - x_j \beta_l$ and $(x_s - x_j) \beta_l + \alpha_l (Q_{s,l} - q_j) - b_s + \epsilon_s$ for $s \in a$. By construction, the distribution of $Y_{j,l}$ only depends on (b_a, I_a) but not on the realized identities in a . Further, for a pair of permanent sellers $\{i, j\}$, and a subset $a \subseteq N_l \setminus \{i, j\}$, define

$$V_{i,j,l}(b_a, I_a) \equiv (\alpha_l (q_j - q_i) + \Delta x_{j,i} \beta_l + \epsilon_{j,l} - \epsilon_{i,l}, Y_{i,l}(b_a, I_a) - \epsilon_{i,l}).$$

Assumption 3 *For each pair of groups $((x, q), (x', q'))$, and any pair of permanent sellers i, j such that i belongs to (x, q) -group and j belongs to (x', q') -group, there exists some composition I_a where $a \subset N_l \setminus \{i, j\}$ and a related bid vector $b_a \equiv (b_k)_{k \in a}$ such that (i) the support of $(B_{j,l} - B_{i,l}, -B_{i,l})$ contains the support of $V_{i,j,l}(b_a, I_a)$ given (b_a, I_a) , and (ii) $V_{i,j,l}(b_a, I_a)$ conditional on (b_a, I_a) has a non-vanishing characteristic function.*

To illustrate the support condition in Assumption 3, consider a pair of permanent sellers i, j such that $(q_i, x_i) = (q_j, x_j)$. Then this support condition holds with the set of other sellers being empty if the joint support of $\epsilon_{j,l} - \epsilon_{i,l}$ and $U_{0,l} - \alpha_l q_i - x_i \beta_l - \epsilon_{i,l}$ is a subset of the support of $(B_{j,l} - B_{i,l}, -B_{i,l})$. This requires the buyer-seller components to have a smaller support relative to that of the bids in equilibrium. In another example, again consider a pair of permanent sellers i, j such that $(q_i, x_i) = (q_j, x_j)$, and I_a consists of a single transitory seller k with $x_k = x_i$. Then this support condition holds if there exists an equilibrium bid $b_{k,l}$ such that the joint support of $\epsilon_{j,l} - \epsilon_{i,l}$ and $\max\{U_{0,l} - \alpha_l q_i - x_i \beta_l, \alpha_l (Q_{k,l} - q_i) - B_{k,l} + \epsilon_{k,l}\} - \epsilon_{i,l}$ is a subset of the support of $(B_{j,l} - B_{i,l}, -B_{i,l})$. In Appendix A3, we discuss how our model is capable of generating the price variation in Assumption 3.

Proposition 2 *Suppose that Assumptions 1-3 hold and $\epsilon_{j,l}$'s have non-vanishing characteristic functions. Then the distribution of $\epsilon_{j,l}$ is identified up to a location normalization (e.g., $\mathbf{E}(\epsilon_{j,l}) = 0$); and the distribution of α_l and the difference in quality are jointly identified up to a scale normalization (e.g., $\mathbf{E}(\alpha_l) = 1$).*

This proposition shows the quality difference between any two permanent sellers sharing the same characteristics x can be identified under appropriate conditions. The next proposition shows the identification of the distribution of the buyer's outside option $U_{0,l}$ and tastes for observed characteristics β_l .

Proposition 3 *Suppose that Assumptions 1-3 and Assumptions 4-5 (in Appendix A2) hold. Suppose further that the lowest quality level for a permanent seller is the same across groups with different observed characteristics x . Then the distribution of the value of outside option conditional on α_l , the distribution of the transitory seller's quality conditional on his bids, and the distribution of β_l are jointly identified.²⁰*

²⁰Since neither the quality indices nor the buyer's random tastes $(\alpha_l, \beta_l, \epsilon_l, U_{0,l})$ are recorded in the data, some minimal location normalization is required for full identification of the level of quality indices. Under the condition on lowest quality levels in Proposition 3, we identify the model by normalizing the lowest quality to zero.

To identify the distribution of β_l , we use the variation in bids submitted by permanent sellers. To do so, we need an extended version of the support condition in Assumption 3. We state this new condition as Assumption 5 in Appendix A2. In the special case where β_l is a constant vector, the identification of the model only requires a weaker location normalization that fixes the lowest quality level for permanent sellers in a single group with a fixed observed characteristic. We discuss this in more detail following the proof of this proposition in Appendix A2.

5.2.3 Entry and Project Costs

The identification of the project cost distribution follows from an argument similar to that in Guerre, Perrigne and Vuong (2000). Let $G_\theta(b) \equiv \Pr(j \text{ wins} \mid j \in A_l, B_{j,l} = b)$, when the type of seller j is summarized by $\theta \equiv (\rho, x, q)$.

Proposition 4 *Suppose that the conditions of Proposition 3 hold, and that G_θ is differentiable with non-zero derivatives over its domain for each θ . Then the distribution of C_j conditional on the seller type θ is identified.*

If the sellers' equilibrium bidding strategies are invertible and differentiable in their costs, Assumption 1 implies that the distribution of equilibrium bids is differentiable almost everywhere. Using a change-of-variable argument in Guerre, Perrigne and Vuong (2000), we identify the sellers' inverse bidding strategies and the distribution of project costs.

For identification of the entry cost distribution, we use exogenous variation in an observed variable, say, Z_l , which does not affect the entry cost distribution but in general enters the seller's ex ante payoff prior to entry decisions. For example, Z_l could be the number of sellers in N_l which affects sellers' ex ante payoff due to the competition effect. Alternatively, Z_l could be the average of observed characteristics x among all sellers that are irrelevant to entry costs. In this case, the average observed characteristics may affect ex ante payoff through its correlation with project costs.

Proposition 5 *Suppose that the conditions of Proposition 4 hold, and that there exists auction-level heterogeneity Z_l reported in the data such that Z_l and $E_l = \{E_{j,l}\}_{j \in N_l}$ are independent. Then for each z in the support of Z_l , the $p_\theta^*(z)$ -quantile of the entry cost distribution for each seller with type θ is identified, where $p_\theta^*(z)$ denotes the equilibrium entry probability for a seller of type θ when $Z_l = z$.*

In the empirical section, we parametrize the distribution of entry costs and use GMM for estimation, exploiting the variation in the observed characteristics of potential entrants as a source of exogenous cost shifters.

6 Estimation

We estimate the model in two steps. In the first step, we use a classification algorithm to uncover the (unobserved) group memberships for the permanent sellers, and in the second step, we perform GMM estimation to recover the rest of the model primitives.

Classification Algorithm. Our classification algorithm is based on Proposition 1. More specifically, for two permanent sellers i and j , we define $\hat{\delta}_{ij}(b) \equiv \hat{r}_{ij}(b) - \hat{r}_{ji}(b)$, where

$$\hat{r}_{ij}(b) \equiv \frac{\sum_{l=1}^L 1\{i \text{ wins}\} K_h(B_{i,l} - b) 1\{j \notin A_l\} 1\{i, j \in N_l\}}{\sum_{l=1}^L K_h(B_{i,l} - b) 1\{j \notin A_l\} 1\{i, j \in N_l\}},$$

where $K_h(v) = K(v/h)/h$ for a univariate kernel function K whereas N_l and A_l are the sets of potential and actual bidders in auction l respectively.²¹ Then we construct test statistics: $\hat{\tau}_{ij}^+ = \int \max\{\hat{\delta}_{ij}(b), 0\} db$, $\hat{\tau}_{ij}^- = \int \max\{-\hat{\delta}_{ij}(b), 0\} db$, and $\hat{\tau}_{ij}^0 = \int |\hat{\delta}_{ij}(b)| db$. For example, the test statistic $\hat{\tau}_{ij}^+$ is used to check whether $r_{ij}(b) > r_{ji}(b)$, i.e., whether $q_i > q_j$ (from Proposition 1.)

Next, we construct a pairwise bootstrap p -value for testing the quality ordering between i and j . For this we generate $\hat{r}_{ij}^*(b)$ and $\hat{\delta}_{ij}^*(b)$ using the bootstrap sample of data, and construct the re-centered bootstrap test statistics, $\hat{\tau}_{ij,s}^{*+}$, $\hat{\tau}_{ij,s}^{*-}$ and $\hat{\tau}_{ij,s}^{*0}$ in the same way as we constructed $\hat{\tau}_{ij}^+$, $\hat{\tau}_{ij}^-$, and $\hat{\tau}_{ij}^0$ except that we use recentered bootstrap quantities $\hat{\delta}_{ij,s}^*(b) - \hat{\delta}_{ij}(b)$ in place of $\hat{\delta}_{ij}(b)$.²² From these we find bootstrap p -values.

We use the algorithm proposed in Krasnokutskaya, Song, and Tang (2014) which is designed to recover the quality group structure so that transitivity of the quality ordering is retained in finite samples. Heuristics for the algorithm is as follows. For each seller i in an x -group, we first divide the remaining sellers into two groups, one with sellers likely to have higher quality than i and the other with sellers likely to have lower quality than i . We obtain this division by comparing the p -values from two pairwise bootstrap tests of the inequality restrictions $r_{i,j} \geq r_{j,i}$ and $r_{i,j} \leq r_{j,i}$. Next, we place seller i in one of the two groups depending on whether seller i is likely to have the same quality as the other sellers in the group. Thus we obtain one group structure for each seller i , and choose one of these structures (specifically, the one that has strongest empirical support in terms of average p -values). This gives the first division of the sellers into two subgroups.

We then sequentially select a subgroup with sellers most likely to have heterogeneous qualities, and divide the group similarly as before. To prevent overfitting (i.e., ending up with too many

²¹Test statistics is constructed using a triweight kernel function: $K(u) = 1\{|u| \leq 1\}(35/32)(1 - u^2)^3$. The bandwidth selection follows the usual Silverman's rule of thumb. Other parameters use in implementation are reported in the Web Supplement.

²²Note that Lee, Song, and Whang (2013) established asymptotic properties of these bootstrap test statistics in a more general set-up.

subgroups), we stop the division process when a goodness-of-fit measure defined in terms of average p -values is dominated by a penalty term. (See the Web Supplement for further details on the implementation of the algorithm.)

GMM Estimation: Moments. Given the estimated quality group structure in the first step, we proceed with GMM estimation of the model primitives.²³ For GMM, the moment conditions are primarily built around the permanent seller's winning probability given the seller's attributes, quality group membership and for a given configuration of the set of active permanent sellers.

To be specific, let B_l be the vector of submitted bids in auction l . For each bidder j in auction l , define

$$m_{j,l} = 1\{j \text{ wins } l\} - \Pr(j \text{ wins } | B_l, I_{A_l}),$$

where $\Pr(j \text{ wins } l | B_l, I_{A_l})$ is the conditional winning probability of seller j in auction l having a given composition of active sellers I_{A_l} when the vector of submitted bids is B_l . Then we construct a moment condition as follows:

$$\mathbf{E} \left[\sum_j g_j(B_l, I_{A_l}) m_{j,l} \right] = 0,$$

where $g_j(B_l, I_{A_l})$ is a function of B_l and I_{A_l} , and the summation is over j in the set of active permanent bidders at auction l .

For the implementation, it remains to choose functions $g_j(B_l, I_{A_l})$. The functions are chosen to exploit variations of the sellers' (x, q) -group memberships and variations in the compositions I_{A_l} of the active sellers. Specifically, we consider two types of moments. The moments of the first type are based on the subset of auctions such that the set A_l includes at least two permanent sellers from the same (x, q) -group. The moments of the second type are confined to the auctions where the set A_l contains permanent sellers belonging to two specific groups, (x, q) and (x', q') , for every possible pair of groups.

Formally, for each choice of (x, q) -, (x', q') -, and x_h -groups and each composition I_{A_l} , we form $g_j(B_l, I_{A_l})$'s as functions of $(B_{i,l}, B_{j,l}, B_{h,l})$, where $B_{j,l}$ is the bid of winning permanent seller j from group (x, q) , $B_{i,l}$ the bid of another permanent seller i from group (x', q') and $B_{h,l}$ the bid of a transitory seller from group x_h . As for the functions of $(B_{i,l}, B_{j,l}, B_{h,l})$, we consider: constant (equal to 1); $B_{j,l}$ and $B_{j,l}^2$; $B_{j,l} - B_{i,l}$ and $(B_{j,l} - B_{i,l})^2$; $(B_{j,l} - B_{i,l})B_{j,l}$; $B_{j,l}B_{h,l}$ and $B_{j,l}^2B_{h,l}$ for a transitory seller h ; $B_{j,l}x_h$ and $B_{j,l}^2x_h$. See the Appendix for details of our choice of the functions g_j 's.

²³The estimation error due to using the estimated quality groups does not affect the asymptotic distribution of the GMM estimator because it has an arbitrarily fast convergence rate due to the finite number of quality groups.

GMM Estimation: Accounting for Transitory Seller's Qualities. To use the moment conditions summarized above, we need to evaluate the conditional winning probability. However, note that a buyer observes (x, q) -group memberships of all sellers in his choice set, whereas the econometrician does not observe the transitory sellers' quality groups. This means that the conditional winning probability is the winning probability after integrating out the vector of participating transitory sellers' qualities. In other words,

$$\Pr(j \text{ wins } | B_l, I_l) = \sum_{\bar{q}_t} \Pr(j \text{ wins } l | B_l, I_{A_l}, Q_l^t = \bar{q}_t) \Pr(Q_l^t = \bar{q}_t | B_l, I_l), \quad (6)$$

where Q^t is the participating transitory sellers' quality vector in the auction and $I_l = \{I_{A_l}, I_{N_l}\}$ to denote the compositions of the sets of actual and potential sellers.

The conditional probability $\Pr(j \text{ wins } l | B_l, I_{A_l}, Q_l^t = \bar{q}_t)$ reflects buyers' decisions, and is determined by the distribution of buyers' weights and outside option. To obtain theoretical expression for this probability we parametrize the distributions of $\epsilon_{j,l}$ and $(\alpha_l, \beta_l, U_{0,l})$ in a standard way. However, it is not immediately obvious how to parametrize $\Pr(Q_l^t = q_t | B_l, I_l)$, because it involves the transitory sellers' behavior. To see how we address this, let us assume that the auction at hand contains only one transitory actual bidder, say, h . Then we express^{24,25} $\Pr(Q_{h,l} = q_h | B_l = b, x_h, I_l)$ as

$$\frac{\Pr(Q_{h,l} = q_h | x_h) \Pr(h \text{ is active} | x_h, Q_{h,l} = q_h, I_{N_l}) f_b(b_h | x_h, q_h, I_{N_l})}{\sum_{q'_h} \Pr(Q_{h,l} = q'_h | x_h) \Pr(h \text{ is active} | x_h, Q_{h,l} = q'_h, I_{N_l}) f_b(b_h | x_h, Q_{h,l} = q'_h, I_{N_l})}, \quad (7)$$

where the subscript h of a variable indicates that it belongs to the transitory seller h . In estimation we parameterize the bid density function $f_b(b_h | x_h, Q_{h,l} = q_h, I_{N_l})$ and the probability of entry $\Pr(h \text{ is active} | x_h, Q_{h,l} = q_h, I_{N_l})$.

Finally, since our identification strategy recovers only $\Pr(Q_{h,l} = q_h | B_l = b_h, x_h, I_l)$ we need to impose additional restrictions in estimation that would allow us to recover $\Pr(h \text{ is active} | x_h, Q_{h,l} = q_h, I_{N_l})$, $f_b(b_h | x_h, Q_{h,l} = q_h, I_{N_l})$, and $\Pr(Q_{h,l} = q_h | x_h)$ separately. Specifically, we impose that the conditional bid distributions and conditional participation probabilities aggregated to the level observed in the data correspond to their empirical counterparts. We also impose optimality of the transitory sellers' participation decisions. (See the Appendix for details.)

Estimation of Sellers' Costs. We estimate the distributions of the seller's costs conditional

²⁴The derivation for a general case is available in the Appendix.

²⁵Notice that information about x_h is contained in I_l . We condition on it separately here to make the expression more accessible to the reader.

on the seller's attributes by combining the bid distributions of permanent sellers with the corresponding inverse bid functions for a given composition of the set of potential bidders \bar{I}_N :

$$\hat{F}_C(c|x, q) = \hat{G}_{p,x,q}(\hat{\xi}_{p,x,q}^{-1}(c; \bar{I}_N) | I_{N_i} = \bar{I}_N),$$

where $\hat{\xi}_{p,x,q}^{-1}(c; \bar{I}_N)$ denotes the inverse of the estimated inverse bid function of a permanent seller in group (x, q) .²⁶ Recall that the distribution of costs depends only on sellers' attributes (both observable and unobservable), (x, q) , not on whether the seller is permanent or transitory. On the other hand, the sellers' bidding strategy and thus the distribution of bids depends on his full type $\theta = (\rho, x, q)$, $\rho \in \{p, t\}$ which is reflected in the expression above. Notice that we are using the distributions of bids and inverse bid functions associated with permanent sellers. The estimated inverse bid function, $\hat{\xi}_{p,x,q}(b; \bar{I}_N)$, is derived from the first order condition of the corresponding permanent seller's optimization problem: for a permanent seller i in group (x, q) ,

$$\hat{\xi}_{p,x,q}(b, \bar{I}_N) = b - \frac{\hat{P}(j \text{ wins} | B_{j,l} = b, I_{N_i} = \bar{I}_N)}{\frac{\partial}{\partial b} \hat{P}(j \text{ wins} | B_{j,l} = b, I_{N_i} = \bar{I}_N)},$$

where $\theta_j = (p, x, q)$.

We assess the magnitude of the costs of entering the auction using the conditions derived from the optimality of the permanent sellers' participation behavior. In the model, the participation probability satisfies the following equation

$$F_E(\Pi(\theta, \bar{I}_N) | x, q) = \Pr(j \text{ is active} | \theta, I_{N_i} = \bar{I}_N), \quad (8)$$

where $F_E(\cdot | x, q)$ is the distribution function of the entry costs and $\Pi(\theta, \bar{I}_N)$ the ex-ante expected profit of a seller with type θ .

We assume that the distribution of entry costs is given by the truncated normal distribution (truncated at zero) with seller-type-specific mean and standard deviation. We estimate the parameters of these distributions using a minimum distance estimation procedure based on the restrictions in equation (8) for various (x, q) -groups of permanent sellers and for time periods characterized by the high/low presence of potential bidders from different (x, q) -groups. The probability $\Pr(i \text{ is active} | \theta, I_{N_i} = \bar{I}_N)$ can be directly identified from the data. As for the expected profits conditional on entry, we compute them from the estimated distributions of bids and costs and the sellers' beliefs about their competitors' participation strategies approximated by the participation behavior observed in the data.

²⁶We implement kernel smoothing over \bar{I}_N when estimating $\hat{G}_{p,x,q}(\cdot | I_{N_i} = \bar{I}_N)$ or $\hat{P}(j \text{ wins} | B_{j,l} = b, I_{N_i} = \bar{I}_N)$, $\frac{\partial}{\partial b} \hat{P}(j \text{ wins} | B_{j,l} = b, I_{N_i} = \bar{I}_N)$ and $\Pr(i \text{ is active} | \theta, I_{N_i} = \bar{I}_N)$.

7 Empirical Results

This section summarizes the estimation results. We begin by describing the implementation details of the classification procedure and the estimated quality group structure. We then turn to the details and the results of the parametric estimation.

Implementation Details. We assume that the buyer’s utility from using a specific seller depends on the seller’s country affiliation, average reputation score, and his (residual) quality in addition to the seller’s bid. The seller’s country affiliation proxies for things such as work culture, convenience of working with a given seller related to the time difference, and the likelihood of language proficiency, whereas the reputation score may reflect public information about the seller’s quality. The distribution of residual quality may also plausibly depend on such factors so we allow the unobserved group structure, which captures this distribution, to depend on the sellers’ countries, and the long-run averages of the sellers’ reputation scores.²⁷

We are interested in the long-run differences among sellers. That is why, we focus our analysis on the last four years captured by our dataset (i.e., years three to six of the market operation). This ensures that the majority of permanent sellers (more than 95%) have been with the platform for more than a year by the starting date of our estimation sample.

We divide all the sellers into three cells according to the long-run average reputation score: average reputation score less than 9.7 (cell 1), average reputation score above 9.7 and below 9.9 (cell 2), average reputation score above 9.9 (cell 3). This results approximately in an allocation of 30%, 30%, and 40% of the sellers to the three cells.

We further group sellers into country groups by geographic proximity and similarity of language and economic conditions. We end up with seven country groups: North America (USA and Canada), Latin America, Western Europe, Eastern Europe, Middle East and Africa, South and East Asia, and Australia (grouped with New Zealand). In our data, North America, Eastern Europe and South or East Asia account for the majority of submitted bids.

We define the set of potential bidders as follows. We assume that the set of potential bidders for project l auctioned in week t consists of all the sellers who submit at least one message for projects of the same type of work as project l during this week.²⁸ This definition ensures that included sellers (a) might reasonably be expected to compete with each other in the auction for a given project; (b) are aware of each other’s presence in the market during the auction.

²⁷We have also verified the robustness of our results by repeating the analysis while including the number of arbitrations and delays as additional observable measures of quality. The results of this analysis are less precise since each cell contains a smaller number of observations but they are very similar to the results we report in the paper.

²⁸We have detailed data on the number of viewings by seller and by week so we can identify all the sellers who are searching in a given week. For those of them who do not submit any messages we use information on their history of past bidding to determine if they belong to the set of potential bidders.

7.1 Classification Results

The classification index is constructed for pairs of permanent sellers on the basis of projects where they both belong to the set of potential bidders. We follow the steps summarized in Section 6 and described in detail in the Web Supplement. That is, we start by estimating a group structure for different numbers of groups. We then apply a criterion function to select the structure with the number of groups most supported by the data. For this structure we then compute confidence sets. We demonstrate steps 1 and 2 for the group of Eastern European sellers with a medium level of average reputation score in a table included in the Web Supplement to the paper.

We have estimated the model for several cells of projects defined in terms of size and duration. The difference in the results across cells is not sufficiently large to warrant a separate discussion in the paper. The results presented below are for the projects owned by US buyers that are of medium size (between \$400 to \$600) and have the specified duration of two to three weeks. We use the results for all cells in our counterfactual analysis.

Table 6 reports the estimated group structures with corresponding confidence sets for cells of North American, Eastern European and East Asian sellers. We estimate multiple quality groups in each cell and the confidence sets associated with each group structure are quite tight. It is difficult to draw any substantive conclusions about the quality distribution on the basis of these results, since the classification into groups is ordinal and does not allow for the comparison of levels across countries or reputation scores. We note here that even the cells that correspond to a very narrow range of reputation scores (such as medium or high reputation scores) are classified into multiple quality groups. Also, allocation of mass between quality groups differs across cells.

We conduct extensive robustness analysis in order to verify robustness of our results to the assumptions of our model and various implementation details. Among other things we explore potential importance of unobserved auction heterogeneity and the possibility that the sellers' quality may vary across projects. We find that the results of classification analysis are quite robust and change very little across specifications that we consider. The results also indicate that unobserved auction heterogeneity if present plays limited role in our environment and that the sellers' qualities appear stable across projects. The results of this analysis are summarized in the Web Supplement to the paper.

7.2 The Results of GMM Estimation

In this section we present the results of the GMM estimation. We begin by summarizing our specification and then discuss the estimates of the objects of interest: the distribution of the buyers' utility weights, the quality distributions for a range of covariate values, the sellers' bidding strategies and the recovered cost distributions. The estimates of objects which are auxiliary to

Table 6: Estimated Quality Groups by Seller Covariates

Country Group	Average Score	Total Number of Sellers	$Q = 1$	$Q = 2$	$Q = 3$
North America	low	12	4 (6)	8 (10)	
North America	medium	13	4 (6)	9 (11)	
North America	high	17	12 (13)	5 (6)	
Eastern Europe	low	18	6 (8)	12 (14)	
Eastern Europe	medium	52	33 (37)	12 (14)	7 (9)
Eastern Europe	high	83	6 (7)	65 (69)	12 (15)
East Asia	low	91	62 (68)	18 (22)	11 (13)
East Asia	medium	66	6 (8)	53 (57)	7 (9)
East Asia	high	58	50 (53)	8 (11)	

This table shows the estimated group structure and a consistently selected number of groups for each cell determined by covariate values. Column 3 indicates the total number of the sellers in the cell. Columns 4-6 report the size of the estimated quality group. The number of the sellers in the corresponding confidence set with 90% coverage is reported in parenthesis. Note that the confidence set with the level $(1 - \alpha)$ for a given quality group is defined to be a random set whose probability of containing this quality group is ensured to be asymptotically bounded from below by $(1 - \alpha)$.

our analysis (such as the distribution of bids and participation probabilities) are reported and discussed in the Web Supplement to this paper.

Baseline Specification. We modify the specification of the buyer's utility function for the purpose of estimation. Specifically, we impose that $\beta_l = \beta$, i.e. β -coefficients are constant across buyers, as we find that such specification achieves the best fit to the data. Thus, the utility function used in estimation is given by

$$\alpha_l \tilde{q}_i - B_{i,l} + \epsilon_{i,l},$$

where $\tilde{q}_i = q_i + x_i \beta$.

Our baseline specification imposes that the distribution of the sellers' qualities conditional on the vector of observable covariates are the same in the populations of the permanent and transitory sellers. This restriction is more stringent than is necessary for our methodology which

requires only that the supports of the quality distributions should coincide. Imposing the equality of distributions enhances robustness and precision of our estimates. Such an assumption also appears to be consistent with the regularities documented in the data. Specifically, the descriptive results reported in Section 3 indicate that the population of permanent and transitory sellers are very similar in terms of the available performance indicators such as the distributions of average reputation score, the average number of bids per project won, etc. The Web Supplement reports the estimation results for two alternative specifications. We comment on these results later in this section.

We assume that buyer-seller-specific components, $\epsilon_{j,l}$, follow the Extreme Value Type I distribution with standard error σ_ϵ , while weight parameter α_l , and the buyer's outside option are assumed to be distributed according to the normal distribution $N(\mu_{\alpha,U_0}, \Sigma_{\alpha,U_0})$.²⁹ We impose the normalizations implied by our identification argument. That is, we normalize the expected value of $\epsilon_{j,l}$ to be equal to zero, the expected value of α_l to be equal to one, and one of the quality levels (quality level 1 of the low average score group, the South and East Asian country group) to be equal to zero. We, therefore, aim to estimate the vector of parameters $\theta = (\sigma_\epsilon, \sigma_\alpha, \rho_{\alpha,U_0}, \{\{\bar{q}_k(x)\}_{k=1,\dots,K(x)}\}_x)$ where $\{\bar{q}_k(x)\}_{k=1,\dots,K(x)}$ is the support of quality distributions that correspond to the covariate values x .

We assume that transitory and permanent sellers' bid distributions are well approximated by normal distributions $N(\mu_{B^t}, \sigma_{B^t}^2)$ and $N(\mu_{B^p}, \sigma_{B^p}^2)$,³⁰ respectively. The means of the bid distribution depend on the seller's quality, average reputation score and country group, as well as on the number of potential permanent competitors by group. We allow the bid distribution of transitory sellers to depend on the number of reputation scores and on both the current and the long-run average scores.³¹ Similarly, we approximate permanent and transitory bidders' respective probabilities of participation by normal distribution functions that depend on linear indices of the seller's quality, long-run average score and country group, the numbers of potential competitors by group (as well as the current number of reputation scores, and the current average of reputation scores in the case of transitory sellers).

The majority of transitory sellers complete only one or two projects. As a result their long-run average reputation scores are not observed in the data. We assume that buyers use public information to form beliefs about the probability that a beginning seller with a given number

²⁹Strictly speaking, the distribution of α_l should have been chosen to have a non-negative support. However, we estimate the standard error of this distribution to be quite small so that this assumption does not make any practical difference. The same comment applies to our assumption on the distribution of bids below.

³⁰See the comment for the distribution of α_l above.

³¹This is because the long-run average score is not observed in the data for transitory sellers. Therefore, the buyer has to base his expectation of the long-run average reputation score on contemporaneously available measures when awarding the project. This, in turn, implies that transitory bidders would incorporate their current average scores into their bids.

and sum of scores belongs to a particular long-run average score group. We recover these beliefs non-parametrically using the beginning of career and long-run data on permanent bidders.

Quality and Other Attributes as Determinants of Buyer’s Choice. Table 7 shows the estimated parameters of the distribution of the buyers’ utility weights and estimated quality levels across cells corresponding to different values of sellers’ x -attributes. In the estimation the prices are normalized by project size; therefore, the estimates for quality levels reflect the buyers’ willingness to pay in terms of the percentage of the project size.

The differences in the estimated quality levels are substantial in magnitude. Specifically, an average buyer is willing to pay an average premium of $(0.5 \times \text{the project size})$ in order to obtain services of a seller with the highest rather than the lowest quality. The quality levels have expected signs and are increasing according to group ranking. We observe that the quality levels are consistent across covariate cells. There appear to be roughly three quality levels present in this market, with the lowest normalized to be around zero, the medium quality level estimated to be somewhere in the range 0.1-0.3, and the highest quality level is between 0.45-0.68. The exact levels differ across country groups, with Eastern Europe characterized by the highest values for each quality level and North America characterized by the lowest “high” quality levels.

Having established that the quality levels are very similar across covariate groups, we can conclude, based on the results from the previous section (Table 6), that there exist important differences in the distribution of quality mass across covariate levels. In particular, North America is missing a middle quality level, whereas the lowest average score cell for Eastern Europe and the highest average score cell for South and East Asia are missing the lowest quality levels. Similarly, the medium score cell for Eastern Europe allocates the most mass to the lowest and medium quality levels, whereas the highest score cell allocates the most mass to the medium and high quality levels. We observe similar regularities in the case of South and East Asia. Hence, the distribution of qualities varies significantly with covariate values.

The country affiliation and the long-run average reputation score appear to have separate effects on the buyer’s utility. These effects, however, are rather small relative to the differences in quality levels. For example, using the quality levels reported in Table 7 we compute that a buyer with the average taste for quality ($\alpha = 1$) would be willing to pay almost 9% more of the project size, $(0.507 - 0.413 = 0.094)$, to obtain the service of a high-quality North American seller with a high reputation score rather than a high-quality North American seller with a low reputation score. Similarly, a buyer with the average taste for quality would be willing to pay 12% more of the project size, $(0.668 - 0.544) = 0.124$, to hire a medium score, high-quality supplier from Eastern Europe rather than a medium score, high-quality supplier from South or East Asia.

Notice that buyers are quite heterogeneous in their willingness to pay for quality. For example,

Table 7: Buyers' Tastes and Quality Levels

	Score Group	Quality Group	Estimated Parameters	Standard Errors
Estimated Quality Levels				
North America	Low	1	-0.016 ***	(0.007)
North America	Low	2	0.413 ***	(0.009)
North America	Medium	1	-0.016 ***	(0.008)
North America	Medium	2	0.433 ***	(0.008)
North America	High	1	-0.016 ***	(0.003)
North America	High	2	0.507 ***	(0.004)
Eastern Europe	Low	1	0.263 ***	(0.003)
Eastern Europe	Low	2	0.625 ***	(0.005)
Eastern Europe	Medium	1	-0.103 ***	(0.005)
Eastern Europe	Medium	2	0.255 ***	(0.003)
Eastern Europe	High	1	-0.107 ***	(0.006)
Eastern Europe	High	2	0.263 ***	(0.005)
Eastern Europe	High	3	0.668 ***	(0.004)
South and East Asia	Low	1	normalized to 0	
South and East Asia	Low	2	0.089 ***	(0.008)
South and East Asia	Low	3	0.449 ***	(0.008)
South and East Asia	Medium	1	-0.019 ***	(0.003)
South and East Asia	Medium	2	0.105 ***	(0.007)
South and East Asia	Medium	3	0.544 ***	(0.006)
South and East Asia	High	1	0.105 ***	(0.004)
South and East Asia	High	2	0.556 ***	(0.007)
Other Parameters				
			$\log(\sigma_\epsilon)$	-0.915 *** (0.041)
			$\log(\sigma_\alpha)$	-1.118 *** (0.012)
			μ_{U_0}	-1.841 *** (0.035)
			$\log(\sigma_{U_0})$	-0.329 *** (0.046)
			σ_{α, U_0}	0.242 *** (0.063)

The quality level for South and East Asia, low score, $Q = 1$, is normalized to be equal to zero. The columns in the table show the estimated coefficients and corresponding standard errors for our baseline specification which corresponds to the case when the distributions of qualities for transitory and permanent sellers are restricted to be equal. The stars, ***, indicate that a coefficient is significant at the 95% significance level.

whereas an average buyer would be willing to pay a 0.507 premium for a high-score, high-quality North American seller, about 5% of the buyers would pay less than 0.20 ($0.507 \times (1 - 1.96 \times \sigma_\alpha) = 0.202$) premium and another 5% would pay more than a 0.81 ($0.507 \times (1 + 1.96 \times \sigma_\alpha) = 0.812$) premium.

The estimated mean of the distribution of the outside option, μ_{U_0} , measured relative to the quality level of a South or East Asian, low-score, low-quality seller is somewhat lower than the average value from this inside option.³² The variance of the distribution of the outside option is larger than the variance of the stochastic match component (ϵ). In our sample the outside option is positively correlated with price sensitivity, i.e., buyers with the high outside option also tend to be more price sensitive.

Fit to the Data. The estimates reported in Table 7 allow us to predict empirical market shares of different seller groups with a precision of one to two percentage points. In addition, using Efron definition of pseudo R^2 , we determine that our model explains auction outcomes at the level of individual bidder for 75% of our observations.³³ This is a big improvement relative to the multinomial logit estimates reported in Section 3 that explain auction outcomes only for 18% of observations.

It is also worth noting that the estimated distribution of transitory sellers' bids and their participation probabilities (reported in the Web Supplement to the paper) indicate a statistically significant dependence of these objects on the transitory sellers' quality levels. Our estimates, therefore, support the assumptions of our model as well as validate our identification strategy.

To summarize, our estimates indicate significant difference in quality levels across sellers. In addition, accounting for unobserved quality substantially improves the fit of the model to the data which indicates that quality plays an important role in this environment.

Alternative Specifications. We have estimated two alternative specifications which differ from the baseline specification in their treatment of transitory sellers. The results for these specifications are reported in the Web Supplement.

The first specification aims to demonstrate the ability of our methodology to handle more general specifications. In particular, it allows the distributions of the transitory and permanent sellers' qualities potentially to be different. Under this specification the frequencies of different quality groups in the population of transitory sellers are estimated from the data. As the results in the Web Supplement show the methodology performs quite well and the parameter estimates obtained in the context of this more general specification are broadly consistent with baseline estimates.

The second specification maintains the 'no unobserved heterogeneity' assumption for transitory sellers. The estimates obtained under this specification are less plausible. This specification also performs poorly in terms of fit to the data. Indeed, both the baseline specification and the

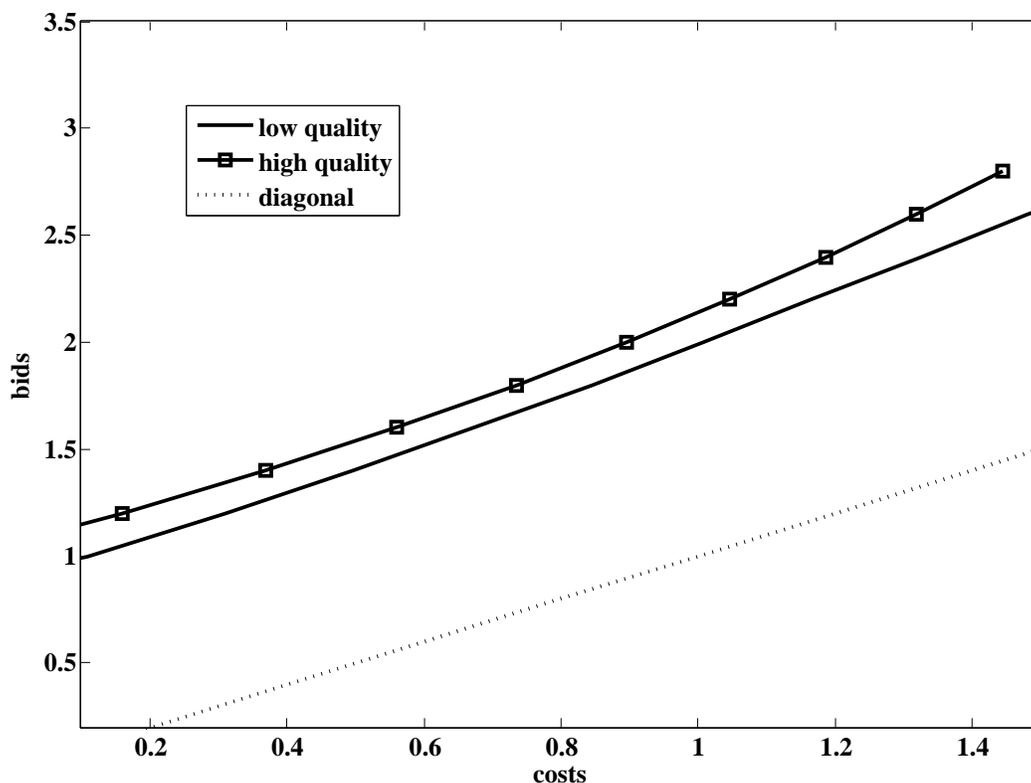
³²The relevant average price scaled by the project size is 1.34.

³³Efron pseudo R^2 is defined as $R^2 = 1 - \sum_{l,j} (y_{l,j} - \pi_{l,j})^2 / \sum_{l,j} (y_{l,j} - \bar{y})^2$ where $y_{l,j}$ is an indicator variable which is equal to one if bidder j wins auction l ; $\pi_{l,j}$ is the predicted probability that bidder j wins auction l ; and \bar{y} is the predicted probability which ignores bidder heterogeneity.

first specification (which permits unobserved heterogeneity for transitory sellers) predict quite precisely the probability that a project is allocated to a transitory seller. This probability is equal to 0.38 in the data (see Table 4) whereas the predicted probabilities computed from the baseline estimates and the estimates obtained under specification one are 0.36 and 0.41 respectively. However, the probability computed from the specification without unobserved heterogeneity of transitory sellers (0.23) substantially underpredicts the probability observed in the data. On the basis of these results we conclude that the assumption that buyers are not informed about the qualities of transitory sellers does not appear to be consistent with the data.

Pricing Strategies and Project Cost Distribution. We recover the sellers' bidding strategies and the distributions of project costs following the approach summarized in Section 6. Figure 1 demonstrates the properties we observe in the estimated bid functions across countries and score levels. The full set of estimated bid functions is reported in the Web Supplement.

Figure 1: Bid Functions



The figure shows the equilibrium bidding strategies of North American permanent sellers with medium score levels. Bidding strategies are recovered from the first order conditions of bidders' optimization program. The convexity at the upper end of the costs' support arises due to presence of stochastic component in buyers' tastes.

We find that the estimated bid functions are increasing in costs, which is consistent with the

theoretical predictions for the environment with private values. The low quality group always follows the most aggressive bidding strategy. The mark-ups over the sellers' costs change very slowly with the cost level and, in fact, for some groups increase as costs reach the upper end of the support. This feature arises because the buyer's choice is based in part on a purely stochastic (from the seller's point of view) component, $\epsilon_{j,l}$. At the high cost realizations where the seller's ability to compete on price is low, his probability of winning increasingly depends on the realization of the $\epsilon_{j,l}$ component, which in turn induces him to choose less aggressive bids. This effect essentially reflects the "gambling" behavior of bidders in the presence of uncertainty about the allocation rule used by buyers. In general, stochasticity plays an important role in our environment: sellers are uncertain about buyers' utility weights as well as their actual competition. This accounts for the relatively large mark-ups we document in our environment.

The estimated project cost distributions (means and standard deviations are reported in Table 8) are typically "increasing" in sellers' quality. More specifically, the cost distribution of the high-quality group is always shifted to the right relative to the distribution of the medium-quality group. However, the low-quality group often has costs that are comparable to the costs of the high-quality group. Further, the distribution of costs differ across countries. In general, high quality US sellers tend to have higher costs relative to foreign sellers of comparable rank.

Many researchers have commented on the fact that the distributions of bids in the online setting tend to have high variance relative to other environments. We document similar regularity: the estimated standard deviation of the bid distribution (reported in the Web Supplement) is equal to 0.24. Notice that the estimated project cost distributions have substantially lower standard deviations relative to the standard deviations of the bid distributions. Thus, our model is capable of rationalizing the highly variable pricing environment through reasonably tight cost distributions. The "gambling" property of the bid functions described above explains this effect. Indeed, convexity or increasing mark-up near the end of the support induces a high variance in sellers' prices and also explains the presence of high bids in this environment.

The Distribution of Entry Costs. Table 9 reports the estimated means and standard deviations of the distributions of entry costs for North American, Eastern European and Asian sellers. The results indicate that North American sellers tend to have lower entry costs relative to sellers from other country groups. At the same time the entry costs are very similar for various average reputation score levels and residual quality groups within country. We use these estimates to assess that the participation costs incurred by entrants in this market constitute around 8% to 12% of the mean project costs. This number is slightly higher than that documented in other markets.³⁴ The relatively large entry costs estimated in this market may reflect large opportu-

³⁴Studies of the US highway procurement market have estimated entry costs to be around 2 – 5% of the engineer's estimate.

Table 8: Project Costs Distributions (Permanent Sellers)

Country Group	Score Group	Quality Group	Mean		Standard Deviation	
			Estimate	Std.Error	Estimate	Std.Error
North America	Low	1	1.239	(0.013)	0.053	(0.006)
North America	Low	2	1.275	(0.012)	0.087	(0.005)
North America	Medium	1	1.627	(0.012)	0.039	(0.005)
North America	Medium	2	1.551	(0.010)	0.073	(0.004)
North America	High	1	1.165	(0.011)	0.053	(0.006)
North America	High	2	1.535	(0.013)	0.099	(0.004)
Eastern Europe	Low	1	1.169	(0.011)	0.037	(0.005)
Eastern Europe	Low	2	1.328	(0.010)	0.079	(0.004)
Eastern Europe	Medium	1	1.576	(0.012)	0.119	(0.005)
Eastern Europe	Medium	2	1.202	(0.009)	0.062	(0.002)
Eastern Europe	Medium	3	1.331	(0.009)	0.101	(0.003)
Eastern Europe	High	1	1.575	(0.011)	0.119	(0.004)
Eastern Europe	High	2	0.981	(0.008)	0.048	(0.002)
Eastern Europe	High	3	1.354	(0.009)	0.089	(0.002)
South and East Asia	Low	1	1.621	(0.010)	0.056	(0.005)
South and East Asia	Low	2	1.119	(0.011)	0.051	(0.002)
South and East Asia	Low	3	1.317	(0.011)	0.169	(0.004)
South and East Asia	Medium	1	1.609	(0.008)	0.107	(0.002)
South and East Asia	Medium	2	1.289	(0.009)	0.121	(0.003)
South and East Asia	Medium	3	1.255	(0.009)	0.124	(0.003)
South and East Asia	High	1	1.074	(0.008)	0.036	(0.002)
South and East Asia	High	2	1.235	(0.008)	0.047	(0.002)

This table summarizes the means and standard errors of the estimated distributions of permanent sellers' project costs.

nity costs and the fact that active bidding for a project involves substantial interaction with the buyer and possibly the preparation of supplementary materials.

Last, we would like to comment on the limitations of the analysis presented in this section. In this analysis, we take the seller's reputation score as given and ignore the possible dynamic considerations associated with reputation building. To mitigate this concern, we base our estimation of the distribution of the sellers' costs on the optimization problem of a permanent seller. While permanent sellers may still take reputation-related concerns into account, the incentives associated with these concerns are likely to be quite weak. A single score does not make a large impact on the average reputation score once a seller has completed three or more projects. Indeed, in the data a bad score does not make a statistically significant impact on the probability of winning or on the bid of an established seller.

Table 9: Entry Cost Distribution (Permanent Sellers)

Country Group	Score Group	Quality Group	Mean		Standard Deviation	
			Estimate	Std.Error	Estimate	Std.Error
North America	Low	1	0.081	(0.035)	0.054	(0.022)
North America	Low	2	0.074	(0.041)	0.055	(0.024)
North America	Medium	1	0.072	(0.045)	0.050	(0.021)
North America	Medium	2	0.073	(0.033)	0.057	(0.029)
North America	High	1	0.083	(0.029)	0.058	(0.029)
North America	High	2	0.082	(0.025)	0.059	(0.021)
Eastern Europe	Low	1	0.112	(0.033)	0.072	(0.034)
Eastern Europe	Low	2	0.113	(0.038)	0.067	(0.029)
Eastern Europe	Medium	1	0.103	(0.034)	0.068	(0.031)
Eastern Europe	Medium	2	0.104	(0.021)	0.065	(0.029)
Eastern Europe	Medium	3	0.101	(0.027)	0.064	(0.023)
Eastern Europe	High	1	0.104	(0.021)	0.075	(0.034)
Eastern Europe	High	2	0.103	(0.013)	0.073	(0.032)
Eastern Europe	High	3	0.105	(0.023)	0.076	(0.034)
South and East Asia	Low	1	0.107	(0.043)	0.067	(0.033)
South and East Asia	Low	2	0.118	(0.026)	0.065	(0.036)
South and East Asia	Low	3	0.117	(0.034)	0.063	(0.033)
South and East Asia	Medium	1	0.091	(0.025)	0.059	(0.023)
South and East Asia	Medium	2	0.096	(0.023)	0.061	(0.029)
South and East Asia	Medium	3	0.097	(0.017)	0.060	(0.029)
South and East Asia	High	1	0.104	(0.013)	0.064	(0.033)
South and East Asia	High	2	0.112	(0.021)	0.062	(0.031)

This table summarizes the means and standard errors of the estimated distributions of permanent sellers' entry costs.

8 Buyers' Gains from Market Globalization

We use the estimated parameters to evaluate the average gain in value over the outside option collected by buyers in our market using the following measure:

$$\mathbf{E} \left[\max_{i \in A_l \cup \emptyset} U_{i,l} - U_{0,l} \right] = \frac{1}{L} \sum_l \mathbf{E}_{\alpha_l, \epsilon_l, U_{0,l}, Q_l^t | b_l^t} \left[\max_{i \in A_l \cup \emptyset} U_{i,l} - U_{0,l} \right]$$

Here \emptyset denotes an outside option and $\max_{i \in A_l \cup \emptyset} U_{i,l}$ represents the buyer's utility collected from participating in the online market.

Recall that bids are scaled by the size of the project; thus, the welfare gain is measured as a fraction of the project size. We find that the buyers who had access to this market on average are able to improve their welfare relative to the outside option by 73% of the project value. The outside option in our setting represents a traditional procurement process, which implies hiring

somebody locally or not hiring anyone at all. In this case our measurement captures the value of the Internet as an alternative marketplace.

This assessment has a number of caveats. First, we are working with a selected set of buyers who perhaps are best able to extract value from the online market. It is possible that the general buyer population still perceives an Internet transaction as high-cost (perhaps in terms of psychic cost) and prefers to use traditional markets. So, perhaps, our finding mostly applies to the “sophisticated” segment of the demand. Second, the offline markets are likely to respond to the emergence of the online market by adjusting prices or product selection. In such a case our measurement would provide a lower bound on the gains to buyers from the Internet since we assess buyers’ gains relative to such improved outside options. Finally, the outside option may potentially include using an alternative online platform or re-auctioning the project on our platform but to a different set of sellers. Even if these concerns are valid it would only indicate that our measurement may underestimate the value of this market.

Welfare Gains from International Trade. We expect that welfare gains created by the Internet arise in part because it provides opportunity for international trade. This potentially creates multiple benefits. First, the number of sellers participating in the market might increase and this would intensify competition and potentially result in lower prices. Further, our estimates indicate that international sellers attracted to this market tend to have lower costs conditional on quality and higher quality levels within the quality rank. Thus, through the Internet, buyers may be able to access these low-cost and high-quality options that are not otherwise available to them. In addition, presence of foreign (more competitive) sellers may put downward pressure on the prices in this online market.

In this section we assess the importance of the last two effects. Specifically, we compare buyers’ welfare under the market conditions in the data to the welfare which obtains when country affiliations of all foreign potential bidders are changed to US country affiliation. In practical terms, when implementing this analysis we keep characteristics of the US potential bidders unchanged (as they are observed in the data). For the foreign sellers, we keep their long-run average score group, and their quality rank (high, medium or low) fixed but replace their quality levels and their distributions of private costs (both participation and project costs) with those which are estimated for the US sellers of the same average score group and the same quality rank. Thus, this experiment is implemented in such a way that the number of potential bidders remains unchanged but the set of potential bidders is homogenized in terms of quality levels and the distributions of costs.³⁵ In this analysis we additionally replace all the transitory sellers

³⁵This assessment reflects only short-run benefits since entry/exit of US buyers and sellers in the absence of international trade may result in the change of numbers as well as of composition of participating sellers. We leave the analysis of this later effect for the future research since it requires conceptually different modeling framework

with permanent sellers (both in the baseline and counterfactual scenario) to simplify the task of solving the model.³⁶ We believe that this does not impact our assessment of the variety effect to an important degree since transitory sellers are substantively very similar to the permanent sellers in our setting, although their presence causes a number of methodological difficulties in the empirical analysis. We then solve for the equilibrium participation and bidding strategies for the baseline and counterfactual environments. These strategies are used to simulate average auction outcomes which are reported in column one and two of Table 10 respectively.

A small implementation nuance concerns the medium quality level. Since medium quality is not present among the US sellers, we replace the medium-quality level by the low- and high-quality levels while maintaining the relative proportions of the high and low quality sellers constant. We verify the impact of this assumption by re-computing equilibrium outcomes for the intermediate step where the set of the quality levels is reduced to “high” and “low” for all countries and reputation score groups while the original country affiliation is preserved for all potential bidders. The results of this analysis are reported in the last column of Table 10. The results indicate that this step does not have any important implications for our subsequent conclusions. It also provides a useful alternative benchmark for our analysis.

The algorithm we use to solve for bidding and participation strategies combines a numeric method, which relies on the local approximation of seller objective function by means of Taylor expansion (first proposed by Marshall, Meurer, Richard, and Stromquist (1994)), with the projection method which allows us to solve for the whole vector of Taylor coefficients at a given grid point simultaneously. The use of the projection method is necessary because the first order condition in our case contains the competitors’ inverse bid functions evaluated at several different points on the support. This is in contrast to the standard case where a single bid level is present.³⁷ As a result, the iterative procedure developed in Marshall, Meurer, Richard, and Stromquist (1994) is not feasible in our setting. On the other hand, in our algorithm, the projection method is applied locally, which is why, it retains the precision and robustness properties that are so attractive in the algorithms based on local approximation.³⁸ Participation strategies, which are type-specific, are summarized by the equilibrium probabilities of participation. We use

as well as data on sellers opportunities in the offline markets.

³⁶We implement this by drawing for each transitory seller his quality level using the appropriate distribution of quality.

³⁷In our case multiple bid values can be pivotal for winning since in addition to sellers’ own bid the allocation depends on his and competitors’ other characteristics.

³⁸Because of these features, the criterion functions typically used to judge convergence in either type of the algorithms (either those based on local approximation or projection ones) can not be used. Instead, we have to design an alternative criterion function which reflects the fact that (a) approximation is local and thus the grid points at which the inverse bid function is evaluated (and their number) may change during the computation and that (b) the performance cannot be evaluated on the basis of the fit of the boundary condition only (as is done in local approximation methods).

Table 10: Welfare Gain from International Internet Trade

	All Groups	US sellers Only	Low and High Quality Only
Participation (average number of bidders):			
Total	3.813	2.151	4.104
Low Quality	0.380	0.781	0.501
Medium Quality	1.151	-	-
High Quality	2.281	1.37	3.603
Allocation (share of projects):			
Low Quality	0.067	0.283	0.081
Medium Quality	0.279	-	-
High Quality	0.610	0.569	0.884
Outside Option	0.043	0.148	0.040
Prices:			
Average	1.563	1.626	1.631
Low Quality	1.647	1.450	1.616
Medium Quality	1.412	-	-
High Quality	1.607	1.726	1.651
Welfare Measures:			
Average Buyers Surplus	0.746	0.565	0.771
Average Seller Profit before Entry Costs	0.266	0.180	0.274
Average Profit (Low Quality)	0.067	0.098	0.065
Average Profit (Medium Quality)	0.147	-	-
Average Profit (High Quality)	0.315	0.220	0.350
Fixed Participation:			
Average Price	-	1.643	-
Price (Low Quality)	-	1.513	-
Price (High Quality)	-	1.680	-
Average Buyers Surplus	-	0.744	-

This table reports the results of a counterfactual analysis investigating the welfare gains to US buyers from access to the international market. The first column presents the results for the benchmark setting when all quality groups are present. The second column presents the outcomes from the setting where foreign potential bidders are replaced by US potential bidders while preserving sellers' quality ranks (medium-quality is replaced by high- and low-quality while preserving original shares of these quality levels in population). The last column is included for comparison purposes. It reports the results for the intermediate step where medium-quality potential bidders are replaced by high- and medium-quality potential bidders without changing the seller's country of origin. Average prices are computed as a share-weighted average of submitted bids. Buyer surplus is measured relative to the expected value of an outside option.

a system of equations similar to the one in 8 to solve for participation probabilities imposing that the seller entry threshold is given by expected profit conditional on participation which depends on the competitors' participation strategies.³⁹ The full details are reported in the Web

³⁹It is well known that multiple equilibria in participation strategies may arise in the settings such as the

Supplement.

The change in the country affiliation of potential bidders has a substantial negative impact on the market. Specifically, the overall participation under the counterfactual scenario is reduced and the set of active sellers has a higher proportion of low quality sellers relative to the overall participation and the composition of sellers under the baseline and intermediate market structures. This regularity arises because the high quality US sellers have higher costs on average and slightly lower quality levels in comparison to the high quality foreign sellers. Thus, the post-entry profit that high quality sellers are able to obtain in the counterfactual scenario is lower than the post entry profit achieved by high quality foreign sellers in the baseline and intermediate scenarios. This effect is somewhat mitigated by the higher participation of low quality sellers under counterfactual scenario relative to the baseline and intermediate scenarios (since the low quality US sellers have lower costs relative to the low quality foreign sellers) and by the fact that the US sellers of all quality levels have lower entry costs relative to the foreign sellers.

As a result of these participation outcomes, the prices charged by high quality sellers are higher on average under the counterfactual scenario. Thus, buyers in this setting are presented with fewer choices and the choices are less attractive (lower quality at higher price) than in the baseline scenario or under the intermediate scenario.⁴⁰ As a result buyers' surplus conditional on purchase declines and the fraction of buyers choosing the outside option increases from 3% to 14.5%. The overall buyer surplus declines by 32% relative to the baseline scenario and by 35% relative to the intermediate step.

The analysis above highlights the importance of the sellers' participation in this market. In fact, the main channel through which elimination of 'foreign' variety of sellers impacts the market is through the reduction in participation which is driven not by the reduction in the *number* of potential bidders but by the change in the *characteristics* of potential bidders. We illustrate this point by re-computing the equilibrium while holding participation frequency fixed at the levels arising in the intermediate step. Under this restriction the reduction in the buyers' surplus relative to the level under the intermediate step is only 3.63%. This result arises because under fixed participation the difference in prices is substantially reduced (the price of high quality is 1.680 comparing to 1.726 under the counterfactual scenario with full adjustment in participation) and the differences in the size and the composition of the buyers' choice sets are eliminated. The reduction in surplus occurs because the overall price level and the prices charged by high quality

one we study. This issue does not impact our estimates since we are not solving the model in estimation. This problem might potentially affect our counterfactual analysis. To address this issue, we re-solve the model using 500 different starting points. We do not find any indication of multiple equilibria.

⁴⁰Prices charged by low quality sellers are lower than in baseline case but in general the choices available to buyers are less attractive than before: choice sets contain higher fraction of low quality sellers who even with low prices deliver lower surplus in comparison to high quality sellers.

sellers remain somewhat higher relative to those in the intermediate step (the price of high quality is 1.680 under the counterfactual scenario with fixed participation and 1.651 under the intermediate scenario) since they are based on less advantageous cost distributions and quality levels of the US sellers.

9 Participation and Impact of International Trade

In this section we investigate the social cost associated with restrictions on the international trade while emphasizing adjustment in the sellers' participation choices. Specifically, we contrast the market outcomes arising in the auction environment with those that would realize if allocation decisions were implemented by a social planner.⁴¹ We consider two scenarios: (a) the unrestricted match of a buyer to a seller where a social planner has full information about potential bidders' costs and qualities as well as about the realization of the buyer's utility coefficients; (b) an allocation mechanism where a social planner has to pay entry fee in order to learn the sellers' costs but he is fully informed about the sellers' qualities and the realization of the buyer's utility coefficients.

Let us begin by describing the details of these experiments. In experiment (a) a social planner chooses the best match for a given buyer among the set of potential bidders associated with his project.⁴² In experiment (b) the social planner observes the entry costs of the potential bidders at the time when he decides for which sellers the entry costs should be paid in order to uncover their project costs; after that he observes the project costs of the chosen "entrants".

As in the auction analysis we assume that all potential bidders are permanent and hold the total number of potential bidders fixed at the level observed in the data. Again, the variety of the sellers available to the buyers is restricted by replacing the quality levels and the cost distributions of the foreign sellers with those of the US sellers conditional on the average score group and quality rank. The medium quality sellers are relabeled as either low or high quality so that the original ratio between the sizes of these groups remains constant.

Since the social planner observes the buyer's preferences (α_l and $\epsilon_{j,l}$) and because the composition of the set of potential bidders varies across projects, all seller groups receive a non-zero share of projects. In the experiment (a), we measure social welfare delivered by matching seller

⁴¹We are grateful to one of the referees for suggesting this exercise.

⁴²We do not restrict the set of potential sellers in any way except to allow it to be sufficiently large. Such approach is without a loss of generality. Indeed, in the data the number of sellers who are allocated to the projects in any given week is only a small subset of the sellers who submit messages to the buyers (potential bidders). Therefore, a surplus of potential bidders exists which, theoretically, is available for the social planner to draw on in order to decide on the allocation for a given project. Further, in the analysis with the entry costs only the number of potential bidders selected to draw project costs but not their fraction is important.

Table 11: Further Analysis on the Impact of International Trade

	All Groups	US Sellers Only	Low and High Quality Only
Social planner: Free match			
Total Surplus	1.190	1.179	1.186
Share of outside option	0.003	0.003	0.003
Number of sellers	20	20	20
Social planner: Costly Participation			
Total Surplus before Entry Costs	1.063	1.048	1.036
Social Welfare net of Entry Costs	0.859	0.872	0.858
Share of outside option	0.005	0.006	0.006
Number of sellers	8.88	8	9.2
Composition of sellers	high quality (0.88,5,3)	high quality (8,0,0)	high quality (0,4,4.2)
Auction			
Buyer Surplus	0.746	0.565	0.771
Social Welfare net of Entry Costs	0.777	0.605	0.798
Share of outside option	0.043	0.148	0.040
Number of sellers	3.813	2.151	4.104

This table reports the results of a counterfactual analysis investigating the allocation decisions in this market from social planner's point of view. The analysis emphasizes the importance of participation effects which here are interpreted as the number and composition of the sellers that the social planner chooses to consider. The results for the auction allocations are reported for comparison in the bottom panel of the table. The first column presents the results for the baseline setting where all quality groups are present. The second column presents the outcomes from the setting where foreign potential bidders are replaced by the US potential bidders while preserving the sellers' quality ranks (medium-quality is replaced by high- and low-quality preserving original shares of these quality levels in population). The last column is included for comparison purposes. It reports the results for the intermediate step where medium-quality potential bidders are replaced by high- and medium-quality potential bidders without changing the seller's country of origin. The composition of the sellers shown in the brackets refers to the average number of the sellers by the country of origin listing the US first, then Eastern Europe and South and East Asia last. Buyer and social surpluses are measured relative to the expected value of the outside option.

j to buyer l as

$$w(q_j, C_{j,l}; \alpha_l, \epsilon_l) = \tilde{u}_{j,l} + (B_{j,l} - C_{j,l}),$$

where $\tilde{u}_{j,l} = -\mu_{U_0} + \alpha_l q_j - B_{j,l} + \epsilon_{j,l}$. Here $\tilde{u}_{j,l}$ reflects buyer l utility from seller j relative to the mean of the outside option. The social planner chooses a seller j_0 from the set of potential bidders associated with project l , N_l , so that social welfare for a given project is maximized:

$j_0 = \operatorname{argmax}_{j \in N_l} w(q_j, C_{j,l}; \alpha_l, \epsilon_l)$. Then, we define

$$W(N_l, \bar{C}_l; \alpha_l, \epsilon_l) = \tilde{u}_{j_0,l} + (B_{j_0,l} - C_{j_0,l}),$$

where \bar{C}_l denoted the vector of project costs of the sellers in N_l . The average per project welfare in this market is given by

$$\bar{W} = \sum_{N_l} \int_{\bar{C}} \int_{\alpha, \epsilon} W(N_l, \bar{C}_l; \alpha_l, \epsilon_l) dF_{\alpha, \epsilon}(\alpha_l, \epsilon_l) dF_C(\bar{C}_l) \Pr(N_l).$$

In the experiment (b), the social welfare for project l is modified as follows

$$W(A_l, \bar{C}_{A_l}, \bar{E}_{N_l}; \alpha_l, \epsilon_l) = \tilde{u}_{j_0,l} + (B_{j_0,l} - C_{j_0,l}) - \sum_{j \in A_l} E_{j,l},$$

where A_l is a subset of potential bidders for whom the social planner chooses to pay the entry fee in order to learn their project cost realizations; $\{E_{j,l}\}_{j \in A_l}$ is the vector of entry costs for this subset of potential bidders; j_0 is defined as before ($j_0 = \operatorname{argmax}_{j \in N_l} w(q_j, C_{j,l}; \alpha_l, \epsilon_l)$); and A_l is chosen to maximize the expected welfare for this project with the expectation taken with respect to the distribution of the project costs for the sellers in A_l :

$$A_l = \operatorname{argmax}_{A \subset N_l} \int_{\bar{C}_A} W(A, \bar{C}_A; \alpha_l, \epsilon_l) dF_C(\bar{C}_A).$$

Notice that in both experiments there is no need to compute prices because they drop out from the welfare expression.

Table 11 summarizes the results of this analysis. As can be easily seen, the unrestricted social planner achieves the highest social surplus before the entry costs. The social planner subject to the entry costs achieves the second highest level of surplus whereas the auction setting delivers the lowest welfare. In the first case the social welfare is high because the planner has an opportunity to choose among a large number of sellers. His decision thus reflects the optimum based on a large number of random draws. In case (b) and in the auction environment, the set of alternatives available to the buyer is substantially reduced since adding an alternative to the choice set is costly. The social planner in case (b) performs better than the auction mechanism because he is able to fully internalize the benefit to the buyer's welfare from a given cost draw. In the auction environment this social benefit is ignored since sellers base their participation decisions only on private profitability. Additionally, prices (and profitability) which endogenously decline in the number of participants further limit entry in the auction environment. Thus, the auction setting is characterized by insufficient entry from the social point of view.

Internalizing the benefits from additional participation helps the social planner to limit the losses associated with the reduction in the variety of seller types. Indeed, under scenario (a), the loss is minimal (1.1%) since the ability to optimize over the cost draws within quality group

compensates for the less desirable group-level characteristics of the available seller groups. Under counterfactual scenario (b), the social planner is actually able to achieve higher social surplus after entry costs under reduced variety relative to the baseline case (the surplus before the entry cost is higher in the baseline case). This outcome arises because the US sellers have lower entry costs than the sellers from other countries. As a result under scenario (b) the welfare is improved by 1.3% in contrast to 32% reduction which is realized in the auction market.

10 Conclusion

This paper makes a two-fold contribution to the literature. First, it exploits the structure of an online service market to provide one of the few available assessments of the welfare gains associated with the globalization of trade in services facilitated by the Internet. Second, it develops a tractable framework that enables analysis of such markets.

We find that the gains to the buyers are quite substantial, at 73% of the project value. The paper emphasizes two channels through which globalization impacts the buyers' welfare: the increase in the variety of available seller types and the competitive effect of the presence of low-cost providers. The analysis of these effects is enabled by our methodology, which allows us to account for the sellers' quality differences that are not observable in the data and to obtain unbiased estimates for the distribution of the buyers' weights, outside options, and the distribution of sellers' costs conditional on sellers' characteristics (observable and unobservable) in the presence of potential endogeneity of sellers' observable attributes and prices.

We obtain a number of important insights into the operation of online procurement markets. To the best of our knowledge, our paper is the first one to inquire into the competitive implications of the market's organization in the form of multi-attribute auctions, which is becoming prevalent in the online (as well as offline) procurement markets. Specifically, we document "gambling"-motivated pricing at high cost realizations that arises due to uncertainty about the buyer's allocation rule. This regularity works well in rationalizing the high variability of prices in our data and is likely to explain similar price variability that has been observed in other online markets.

The methodological part of the paper contains several innovative steps. First, we deviate importantly from the traditional discrete choice approach by structuring our estimation in two steps such that the unobserved group structure of permanent sellers is recovered in the first step and then is subsequently used in the second step to facilitate the identification of the distribution of the buyers' weights and outside options as well as to relieve the computational burden associated with accounting for endogeneity of transitory sellers' observable characteristics and prices. An important insight underlying this procedure is that the unobserved group structure

could be recovered separately from the estimation of the buyers' components.

Second, our estimation procedure does not rely on the moments that condition on the buyers' choice set, as is typical in discrete choice estimation. Instead, we exploit moment conditions that aggregate over the choice sets that have certain common properties. This is necessitated by the presence of transitory sellers and buyer-specific choice sets that are prevalent in our setting.

Third, our estimation approach leverages a large amount of data typically available to a researcher in the Internet markets. Specifically, we are able to uncover permanent sellers (unobserved) groupings nonparametrically in the presence of selection into participation (especially by transitory sellers whose selection is difficult to control for in estimation) by conditioning on buyers' and sellers' observable characteristics in estimation.

Some of the features of our setting, such as buyer-specific choice sets and self-selection into participation by agents, have been previously addressed in the context of college choice, online marriage markets, enrollment in residential medical training and other environments with a matching component. Our environment differs from these settings in that in addition to unobserved heterogeneity of the supply-side agents we also have private information on both sides of the market, and that the pricing, which is based on private information, has an important impact on equilibrium outcomes. Further, the large number and anonymity of players on the demand side in our setting facilitates identification of unobserved player heterogeneity despite the endogeneity of the sellers' participation and pricing decisions. In this we have advantage over other studies where researchers have to impose stronger assumptions in order to uncover primitives in the presence of unobserved heterogeneity.

To the best of our knowledge, this paper marks the first effort to estimate a tractable model of the online procurement market. Consequently, we focus on the factors we believe are of the first-order importance - unobserved heterogeneity of the sellers, private information of the sellers about their costs and private information of the buyers about the weights they use and their outside options - while making simplifying assumptions about the issues that are likely to be less important. We expect that the basic insights of our methodology will carry over to richer settings that elaborate on these issues in future research.

In summary, we believe that our results shed light on the operation of online markets for services and on the gains accrued to the buyers participating in these markets. The methodology developed in this paper can be applied in other settings characterized by unobserved agent heterogeneity. Among other things it could be used to further study various aspects of (online) service markets: from optimal pricing and optimal procurement to the analysis of reputation-building in this environment.

References

- ACKERBERG, D., K. HIRANO, AND Q. SHAHRIAR (2006): “The Buy-it-Now Option, Risk Aversion and Impatience in Empirical Model of eBay Bidding,” Working paper, University of Michigan.
- (2014): “Identification of Time and Risk Preference in Buy Price Auctions,” Working paper, University of Michigan.
- ASKER, J., AND E. CANTILLON (2008): “Properties of Scoring Auctions,” *RAND Journal of Economics*, 39.
- (2010): “Procurement when Both Price and Quality Matter,” *RAND Journal of Economics*, 41.
- ATHEY, S., AND P. A. HAILE (2002): “Identification of Standard Auction Models,” *Econometrica*, 70(6), 2107–2140.
- ATHEY, S., AND J. LEVIN (2001): “Information and Competition in US Forest Service Timber Auctions,” *Journal of Political Economy*, 109(2).
- ATHEY, S., AND D. NEKIPELOV (2012): “A Structural Model of Sponsored Search Advertising Auctions,” Working paper, Stanford University.
- BACKUS, M., AND G. LEWIS (2012): “A Demand System for a Dynamic Auction Market with Directed Search,” Working paper, Harvard University.
- BAJARI, P., AND A. HORTACSU (2003): “The Winner’s Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions,” *RAND Journal of Economics*, 34(2), 329–355.
- (2004): “Economic Insights from Internet Auctions: A Survey,” *Journal of Economic Literature*, 42(June), 457–486.
- BAJARI, P., AND G. LEWIS (2011): “Procurement with Time Incentives,” *Quarterly Journal of Economics*, 126.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile prices in market equilibrium,” *Econometrica*, 63, 841–890.
- CABRAL, L., AND A. HORTACSU (2010): “Dynamics of Seller Reputation: Theory and Evidence from eBay,” *Journal of Industrial Economics*, 58(1), 54–78.
- CHIAPPORI, P.-A., AND B. SALANIE (2001): “Testing for Asymmetric Information in Insurance Markets,” *Journal of Political Economy*, 108, 56–78.
- CILIBERTO, F., AND E. TAMER (2009): “Market Structure and Multiple Equilibria in the Airline

- Industry,” *Econometrica*, 77, 1791–1828.
- DECAROLIS, F., M. GOLDMANIS, AND A. PENTA (2014): “Common Agency and Coordinated Bids in Sponsored Search Auctions,” Working paper, Boston University.
- GREENSTEIN, S. (1993): “Did Installed Base Give an Incumbent any (Measurable) Advantage in Federal Computer Procurement,” *Rand Journal of Economics*, 24(1), 19–39.
- (1995): “Sole-Sourcing versus Competitive Bidding: U.S. Government Agencies’ Procedural Choices for Mainframe Computer Procurement,” *Journal of Industrial Economics*, XLIII(2), 125–140.
- GUERRE, E., I. PERRIGNE, AND Q. VUONG (2000): “Optimal Nonparametric Estimation of First-Price Auctions,” *Econometrica*, 68(3), 525–574.
- HECKMAN, J., AND B. SINGER (1984): “A Method of Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data,” *Econometrica*, 52, 271–320.
- HENDRICKS, K., AND A. SORENSEN (2014): “Evaluating the Role of Intermediary in On-line Auctions,” Working paper, University of Wisconsin-Madison.
- HERRENDORF, B., R. ROGERSON, AND A. VALENTINYI (2009): “Two Perspectives on Preferences and Structural Transformation,” Working paper 15416, NBER.
- JOFRE-BONET, M., AND M. PESENDORFER (2003): “Estimation of a Dynamic Auction Game,” *Econometrica*, 71(5), 1443–1489.
- KASAHARA, H., AND K. SHIMOTSU (2009): “Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices,” *Econometrica*, 77(1), 135–175.
- KRASNOKUTSKAYA, E. (2011): “Identification and Estimation in Highway Procurement Auctions under Unobserved Auction Heterogeneity,” *Review of Economic Studies*, 28, 293–323.
- KRASNOKUTSKAYA, E., AND K. SEIM (2011): “Bid Preference Programs and Participation in Highway Procurement,” *American Economic Review*, 101.
- KRASNOKUTSKAYA, E., K. SONG, AND X. TANG (2014): “Uncovering Unobserved Agent Heterogeneity through Pairwise Comparisons,” Working paper, University of Pennsylvania.
- LEE, S., K. SONG, AND Y.-J. WHANG (2013): “Testing Functional Inequalities,” *Journal of Econometrics*, 172, 14–32.
- LEWIS, G. (2011): “Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors,” *American Economic Review*, 101(4), 1535–1546.
- LI, T., I. PERRIGNE, AND Q. VUONG (2000): “Conditionally independent private information in OCS wildcat auctions,” *Journal of Econometrics*, 98(1), 129–161.

- (2002): “Structural Estimation of the Affiliated Private Value Auction Model,” *RAND Journal of Economics*, 33(2), 171–193.
- MARION, J. (2007): “Are bid preferences benign? The effect of small business subsidies in highway procurement auctions,” *Journal of Public Economics*, 91(7-8), 1591–1624.
- MARSHALL, R. C., M. J. MEURER, J.-F. RICHARD, AND W. STROMQUIST (1994): “Numerical Analysis of Asymmetric First Price Auctions,” *Games and Economic Behavior*, 7(2), 193–220.
- RAO, B. (1992): *Identifiability in Stochastic Models: Characterization of Probability Distributions*. Academic Press, New York.
- SWINKELS, J. (2009): “First and Second Price Mechanisms in Procurement and Other Asymmetric Auctions,” *Working Paper*.

Appendix A: Identification Proofs

Throughout the proof, we suppress the auction subscript l to simplify notation, and write $(C_{i,l}, E_{i,l}, \epsilon_{i,l}, B_{i,l}, U_{0,l}, \alpha_l, \beta_l, N_l, A_{i,l}, I_{N_l}, I_{A_l})$ as $(C_i, E_i, \epsilon_i, B_i, U_0, \alpha, \beta, N, A_i, I_N, I_A)$, which are considered i.i.d. draws across auctions reported in the data. In comparison, sellers’ characteristics x_i, q_i are fixed for each seller throughout all auctions in the data and therefore do not need any auction subscript.

A1. Proof of Proposition 1

Throughout Appendix A we suppress the dependence of strategies and winning probabilities on the set of potential bidders N and its composition I_N . Partition the set of entrants A according to whether a seller is weakly preferred to the outside option or not. That is $A = A^{(1)} \cup A^{(0)}$, where

$$A^{(1)} \equiv \{k \in A : U_k \geq U_0\} \text{ and } A^{(0)} \equiv A \setminus A^{(1)}.$$

In general, both $A^{(1)}$ and $A^{(0)}$ contain permanent and transitory sellers with various quality and observed characteristics.

For any pair of *permanent* potential sellers i, j , let $\mathcal{A}_{i,j}$ denote the support of $(A^{(1)}, A^{(0)})$ after excluding i and j . That is, $\mathcal{A}_{i,j} \equiv \{(a, a') : a, a' \subseteq N \setminus \{i, j\} \text{ with } a \cap a' = \emptyset\}$. For any $(a, a') \in \mathcal{A}_{i,j}$, define

$$\mathcal{P}_{i,j}(b; a, a') \equiv \Pr \left\{ i \text{ wins } \left| \begin{array}{l} i \in A, j \notin A, B_i = b, i, j \in N \\ A^{(1)} \setminus \{i\} = a, A^{(0)} \setminus \{i\} = a' \end{array} \right. \right\}$$

for any $b \in \mathcal{B}_i$.

Lemma A1 *Suppose Assumptions 1 and 2 hold. Fix a pair of permanent potential bidders (i, j) with $x_i = x_j$.*

(a) *For any $(a, a') \in \mathcal{A}_{i,j}$ and $b \in \mathcal{B}_{i,j}$,*

$$q_i \begin{cases} > \\ = \\ < \end{cases} q_j \Rightarrow \mathcal{P}_{i,j}(b; a, a') \begin{cases} \geq \\ = \\ \leq \end{cases} \mathcal{P}_{j,i}(b; a, a'). \quad (9)$$

(b) *For some $(a, a') \in \mathcal{A}_{i,j}$, the two weak inequalities implied in (9) hold strictly with positive probability for all b in a subset of $\mathcal{B}_{i,j}$ that has positive Lebesgue measure.*

Proof: Part (a). Recall in our model each seller $i \in N$ draws a private entry cost E_i independently from $F_{E|x_i, q_i}$ and makes an entry decision based on E_i and the composition of N . As a result, sellers' entry decisions are independent. Besides, Assumption 1 implies that in equilibrium sellers' bidding strategies are functions of private costs and do not depend on the buyer's preference taste $(\alpha, \beta, \epsilon, U_0)$. Given any $(a, a') \in \mathcal{A}_{i,j}$, let $\mathcal{E}(a, a')$ be a shorthand for the event " $U_s \geq U_0 \forall s \in a$ and $U_{s'} < U_0 \forall s' \in a'$ ". Let $\Delta\epsilon_{s,i} \equiv \epsilon_s - \epsilon_i$, $\Delta x_{i,s} \equiv x_i - x_s$ and $\Delta q_{i,s} \equiv q_i - q_s$. Then

$$\begin{aligned} \mathcal{P}_{i,j}(b; a, a') &= \int \Pr \left\{ \begin{array}{l} -\epsilon_i \leq x_i \beta + \alpha q_i - b - u_0 \text{ and} \\ \Delta\epsilon_{s,i} - B_s \leq \Delta x_{i,s} \beta + \alpha \Delta q_{i,s} - b \forall s \in a \end{array} \middle| \alpha, \beta, u_0, \mathcal{E}(a, a') \right\} \\ &\quad \times dF(\alpha, \beta, u_0 | \mathcal{E}(a, a')). \end{aligned} \quad (10)$$

The independence between sellers' private entry costs, project costs and the buyer's tastes α, β and outside option U_0 implies that the event that $i, j \in N$ satisfy " $i \in A, j \notin A, B_i = b$ " can be excluded from the events conditioned on the conditional distribution of (α, β, U_0) on the right-hand side of (10). Furthermore, that the entry decisions are independent from (α, β, U_0) , ϵ and C , that the project costs are independent across the sellers, and that these project costs are independent from (α, β, U_0) and ϵ imply the event that $i, j \in N$ satisfy " $i \in A, j \notin A, B_i = b$ " can be excluded from the conditioning set in the integrand once we condition on α, β, u_0 and $\mathcal{E}(a, a')$. By similar arguments, $\mathcal{P}_{j,i}(b; a, a')$ takes a form that is identical to $\mathcal{P}_{i,j}$ in (10), only with i replaced by j .

By Assumption 1, the joint distribution of ϵ_i and $(\epsilon_s, B_s)_{s \in a}$ is identical to that of ϵ_j and $(\epsilon_s, B_s)_{s \in a}$ once conditioned on α, β, u_0 and $\mathcal{E}(a, a')$. Besides, the conditional distribution $F(\alpha, \beta, u_0 | \mathcal{E}(a, a'))$ does not depend on the identity of sellers i and j . It then follows that the claim in part (a) of the lemma holds for all pairs of permanent sellers i, j with $x_i = x_j$.

Part (b). Consider a simple case where $a = \{s\}, a' = \{t\}$ where $q_s = q_i$ and $x_s = x_i$ but q_t is unrestricted. This case happens with positive probability under our maintained assumptions

about the entry stage. Let U_i, U_s, U_t denote the payoff for the buyer from sellers i, s, t respectively. By definition,

$$\mathcal{P}_{i,j}(b; a, a') = \frac{\Pr\{U_i > U_s, U_s \geq U_0 > U_t \mid B_i = b, i \in A, j \notin A, i, j \in N\}}{\Pr\{U_s \geq U_0 > U_t \mid B_i = b, i \in A, j \notin A, i, j \in N\}}. \quad (11)$$

Under Assumptions 1, the denominator on the right-hand side of (11) does not depend on events in the conditioning set, and the denominator is $\Pr\{U_s \geq U_0 > U_t\}$ and does not depend on the identities i or j .

On the other hand, the numerator, by the Law of Total Probability and the maintained independence assumptions, is

$$\begin{aligned} & \int \Pr\{U_i > U_s, U_s \geq U_0 > U_t \mid \alpha, \beta, u_0, B_i = b\} dF(\alpha, \beta, u_0) \\ = & \int \Pr\left\{ \begin{array}{l} B_s - \epsilon_s \leq x_s \beta + \alpha q_s - u_0, \quad B_t - \epsilon_t > x_t \beta + \alpha q_t - u_0 \\ \epsilon_i + B_s - \epsilon_s > (x_s - x_i) \beta + \alpha(q_s - q_i) + b \end{array} \mid \alpha, \beta, u_0 \right\} dF(\alpha, \beta, u_0) \\ = & \int \left(\Pr\left\{ \begin{array}{l} B_s - \epsilon_s \leq x_s \beta + \alpha q_s - u_0, \\ B_s + \epsilon_i - \epsilon_s > (x_s - x_i) \beta + \alpha(q_s - q_i) + b \end{array} \right\} \times \Pr\{B_t - \epsilon_t > x_t \beta + \alpha q_t - u_0\} \right) dF(\alpha, \beta, u_0) \end{aligned} \quad (12)$$

where the second equality follows from the independence between C, ϵ and (α, β, U_0) . A similar expression holds for $\mathcal{P}_{j,i}$, only with the subscripts i in (12) replaced by j .

Without loss of generality, suppose $q_i > q_j$. This implies $q_s - q_i = 0 < q_s - q_j$ (because $q_s = q_i$ by our supposition). Also, by construction $x_i - x_s = x_j - x_s = 0$. Under Assumptions 1, the joint distribution of $(\epsilon_i, B_s - \epsilon_s)$ is identical to that of $(\epsilon_j, B_s - \epsilon_s)$. Furthermore, B_s is independent from B_i and $(\epsilon_i, \epsilon_s, \alpha, \beta, U_0)$; and both $(\epsilon_i - \epsilon_s)$ and $(\epsilon_j - \epsilon_s)$ are continuously distributed with positive density over the support $[\underline{\epsilon} - \bar{\epsilon}, \bar{\epsilon} - \underline{\epsilon}]$.

It then follows from Assumption 2 that there exists a subset of $\mathcal{B}_{i,j}$ with positive Lebesgue measure so that for all b in this subset, there is positive probability that $(B_s - b)$ is close enough to 0. Also recall the distribution of α is independent from the sellers' bids. Hence there is positive probability that α is sufficiently small so that for all b in this afore-mentioned subset of $\mathcal{B}_{i,j}$, we have $\mathcal{P}_{i,j}(b; a, a') > (<) \mathcal{P}_{j,i}(b; a, a')$ whenever $q_i > (<) q_j$. ■

Proof of Proposition 1: By definition and the Law of Total Probability, we can write $r_{i,j}(b)$ as:

$$\sum_{(a,a')} \Pr\{i \text{ wins} \mid A^{(1)} \setminus \{i\} = a, A^{(0)} \setminus \{i\} = a', B_i = b, i \in A, j \notin A, i, j \in N\} \times \Pr\{A^{(1)} \setminus \{i\} = a, A^{(0)} \setminus \{i\} = a' \mid B_i = b, i \in A, j \notin A, i, j \in N\}$$

Under Assumption 1, the first conditional probability in the summand above is $\mathcal{P}_{i,j}(b; a, a')$. Besides, the second conditional probability in the summand does not depend on the event " $B_i =$

b_i ” and

$$\begin{aligned} & \Pr\{A^{(1)} \setminus \{i\} = a, A^{(0)} \setminus \{i\} = a' \mid i \in A, j \notin A, i, j \in N\} \\ = & \Pr\{A^{(1)} \setminus \{j\} = a, A^{(0)} \setminus \{j\} = a' \mid j \in A, i \notin A, i, j \in N\} \end{aligned}$$

under the same assumptions. It then follows from Lemma A1 that the claim in the proposition is true. ■

A2. Proof of Propositions 2-5

Proof of Proposition 2: Consider a pair of permanent sellers i, j with $(x_i, q_i) = (x_j, q_j)$ and a composition of *other* entrants I_a with $a \subset N \setminus \{i, j\}$ which satisfy Assumption 3 for some price vector b_a . Let $b \equiv (b_i, b_j, b_a)$. Let $\varphi_{i,j}(a, b, I_a)$ denote the probability that i wins conditional on “ $\{i, j\} \in A, B_i = b_i, B_j = b_j$ and the composition of the other entrants in A is I_a and quotes b_a ”. Under Assumptions 1,

$$\varphi_{i,j}(a, b, I_a) = \Pr\{\epsilon_j - \epsilon_i \leq b_j - b_i, Y_i(b_a, I_a) - \epsilon_i \leq -b_i \mid I_a\}.$$

Note the event conditioned on the right-hand side is not “ $\{i, j\} \in A, B_i = b_i, B_j = b_j$ and the composition of the other entrants in A is I_a and quotes b_a ” any more. Instead, the right-hand side is only conditioning on the composition of the other entrants than i, j . This is because sellers’ private entry and project costs and the quality of transitory sellers are independent from $(U_0, \alpha, \beta, \epsilon)$. In fact $\varphi_{i,j}$ is a function of b and I_a only. Also, under these maintained assumptions, ϵ_i, ϵ_j and $Y_i(b_a, I_a)$ are mutually independent once conditional on b_a and I_a . With B_j and B_i being independent from $(U_0, \alpha, \beta, \epsilon)$, evaluating $\varphi_{i,j}(a, b, I_a)$ at different quoted prices b_i, b_j amounts to evaluating a fixed joint distribution of $\epsilon_j - \epsilon_i$ and $Y_i(b_a, I_a) - \epsilon_i$ conditional on I_a at different points on the support. By Assumption 3, the joint distribution of $(\epsilon_j - \epsilon_i, Y_i(b_a, I_a) - \epsilon_i)$ conditional on I_a is identified from $\varphi_{i,j}$ over its full support. By the mutual independence between ϵ_i, ϵ_j and $Y_i(b_a, I_a)$ given I_a and the non-vanishing characteristic functions of ϵ_i ’s, the proposition follows from the Kotlarski’s Theorem (or Theorem 2.1.1 in Rao (1992)).

Next, without loss of generality, consider a pair of permanent sellers i, j such that $x_i = x_j$ and $q_i > q_j$. Fix another composition I_a for other entrants $a \subset N \setminus \{i, j\}$ and a bid vector b_a that satisfy Assumption 3 for such a pair of permanent sellers i, j . By the same argument as that used above, the joint distribution of $(\alpha \Delta q_{j,i} + \Delta \epsilon_{j,i}, Y_i(b_a, I_a) - \epsilon_i)$ conditional on I_a is fully identified over its support for the fixed (b_a, I_a) under Assumption 3. Because the marginal distribution of match components is identified as above, so is the distribution of the difference $\Delta \epsilon_{j,i}$. With $\alpha \Delta q_{j,i}$ independent from $\Delta \epsilon_{j,i}$, the distribution of $\alpha \Delta q_{j,i}$ is identified from the marginal distribution of $\alpha \Delta q_{j,i} + \Delta \epsilon_{j,i}$ (which does not depend on I_a under independence assumptions maintained), using

the non-vanishingness of the characteristic functions of $\epsilon_j - \epsilon_i$. It then follows that the constant parameter $\Delta q_{j,i}$ (recall i, j are permanent sellers) and the distribution of α are jointly identified up to a scale normalization. ■

Proposition 2 shows the quality difference between all pairs of permanent sellers sharing the same observed characteristics is identified. Under the condition that the lowest quality level for a permanent seller is the same across groups with different observed characteristics x_i , we can recover the level of quality for sellers with various observed characteristics up to a location normalization that sets the (constant) lowest quality to zero.

We now introduce the additional rank condition used in Proposition 3. Suppose Assumption 3 holds for a pair of permanent sellers $\{i, j\}$ with $q_i \neq q_j$ and $x_i = x_j$, and let I_a, b_a denote a composition for a set of other entrants than i, j and a bid vector associated with this set, which satisfy Assumption 3 for such i, j . Let N^p (and N^t) denote the set of permanent (and transitory) potential bidders in an auction. Partition a set $a \subset N \setminus \{i, j\}$ into a set that consists of transitory sellers only $a^t \equiv a \cap N^t$ and one that consists of permanent sellers only $a^p \equiv a \cap N^p$. Let I_{a^t} denote the composition of a^t , and let $Q_{a^t} \equiv (Q_s : s \in a^t)$. Let $b_{a^t} \equiv (b_s : s \in a^t)$ denote the vector of bids submitted by entrants who are transitory sellers; and likewise define b_{a^p} . By construction, $b_a \equiv (b_{a^p}, b_{a^t})$ and

$$\begin{aligned} \psi_i(y, \alpha, b_a, I_a) &\equiv \Pr\{Y_i(b_a, I_a) + \alpha q_i \leq y \mid \alpha, I_a\} \\ &= \Pr\{U_0 \leq y \mid \alpha\} \sum_{q_{a^t}} \lambda(q_{a^t}; b_{a^t}, I_{a^t}) \Lambda_i(y, q_a; \alpha, b_a, I_a) \end{aligned} \quad (13)$$

where $\lambda(q_{a^t}; b_{a^t}, I_{a^t}) \equiv \Pr\{Q_{a^t} = q_{a^t} \mid \alpha, b_{a^t}, I_{a^t}\}$ (which does not depend on α under independence assumptions maintained) and

$$\Lambda_i(y, q_a; \alpha, b_a, I_a) \equiv \Pr \left\{ \begin{array}{l} (x_s - x_i)\beta + \alpha q_s - b_s + \epsilon_s \leq y \quad \forall s \in a^p \\ (x_{s'} - x_i)\beta + \alpha q_{s'} - b_{s'} + \epsilon_{s'} \leq y \quad \forall s' \in a^t \end{array} \middle| \alpha, I_a \right\}; \quad (14)$$

and the summation $\sum_{q_{a^t}}$ is over all possible quality profile for transitory sellers in a^t .⁴³

In the derivations below, we fix I_a and b_{a^p} and suppress them in the notation for ψ_i, λ and Λ_i . Likewise, we can construct a similar equation for the same y, α with $\Pr\{U_0 \leq y \mid \alpha\} > 0$ but different vector $\tilde{b}_{a^t} \neq b_{a^t}$ such that $(b_{a^p}, \tilde{b}_{a^t})$ also satisfy Assumption 3. Taking the ratio of these two equations associated with b_{a^t} and \tilde{b}_{a^t} , and re-arranging terms, we get

$$\psi_i(y, \alpha, \tilde{b}_{a^t}) \sum_{q_{a^t}} \lambda(q_{a^t}; b_{a^t}) \Lambda_i(y, q_{a^t}; \alpha, b_{a^t}) = \psi_i(y, \alpha, b_{a^t}) \sum_{q_{a^t}} \lambda(q_{a^t}; \tilde{b}_{a^t}) \Lambda_i(y, q_{a^t}; \alpha, \tilde{b}_{a^t}) \quad (15)$$

⁴³The last equality in (13) follows from the law of total probability, the independence between (α, U_0) and ϵ , and the fact that for all $s' \in a^t$, $Q_{s'}$ is an independent draw from a distribution conditional on $x_{s'}$. That the distribution of U_0 given α is independent from I_a is due to the independence between sellers' entry costs and (U_0, α) ; and that $\lambda(q_{a^t}; b_{a^t}, I_{a^t})$ is not a function of α is because the price quoted are functions of project costs and privately observed quality q_{a^t} , which independent from α once I_{a^t} is controlled for.

which is an equations in $2K^*$ unknown weights, or probability masses,

$$\{\lambda(b_{a^t}), \lambda(\tilde{b}_{a^t})\} \equiv \{\lambda(q_{a^t}; b_{a^t}), \lambda(q_{a^t}; \tilde{b}_{a^t})\}_{q_{a^t}},$$

with K^* being the cardinality of the support of Q_{a^t} given x_{a^t} . In addition each group of the K^* weights need to sum to one:

$$\sum_{q_{a^t}} \lambda(q_{a^t}; b_{a^t}) = \sum_{q_{a^t}} \lambda(q_{a^t}; \tilde{b}_{a^t}) = 1. \quad (16)$$

We need the follow condition on the support of (U_0, α) .

Assumption 4 *There exists a set $\mathcal{U} \equiv \{(y^{(r)}, \alpha^{(r)}) : r = 1, 2, \dots, R\}$ from the joint support of (U_0, α) with $R \geq 2(K^* - 1)$ such that the matrix of coefficients in the linear system of $(R + 2)$ equations in $\{\lambda(b_{a^t}), \lambda(\tilde{b}_{a^t})\}$, which is constructed by stacking (16) with R equations of (15) evaluated at the pairs in \mathcal{U} , has full rank.*

With (U_0, α) continuously distributed, this full rank condition can be expected to hold for a set \mathcal{U} with positive measure, provided there is sufficient variation in ψ_i, Λ_i over the support of (U_0, α) . We also need the following extended version of the support condition in Assumption 3.

Assumption 5 *There exists a composition I of permanent sellers such that for some composition of the other sellers I' and a vector of their bids b' , the support of the random vector $((\alpha \Delta q_{s,i} + \Delta x_{s,i} \beta + \Delta \epsilon_{s,i})_{s \in n \setminus \{i\}}, Y_i(b', I') - \epsilon_i)$ conditional on (b', I') is a subset of the support of $(B_j - B_i, -B_i)$ for an $i \in n$ whenever n is a set of permanent sellers that has composition I with $|n| \geq \dim(\beta) + 1$ and $(x_s - x_i)_{s \in n \setminus \{i\}}$ having full rank. The distribution of the random vector conditional on (b', I') has a non-vanishing characteristic function.*

Assumption 5 requires that there exists a composition I for a set of permanent sellers such that for some composition of the other sellers I' and a vector of their bids b' , the support condition similar to part (i) in Assumption 3 holds for some $i \in a$ whenever the set of permanent sellers a has composition I , with $|a| \geq \dim(\beta) + 1$ and $(x_s - x_i)_{s \in a \setminus \{i\}}$ being non-singular.

Proof of Proposition 3. Without loss of generality, consider a composition I for some permanent sellers $\{i, j\}$ and a composition I_a for other entrants $a \subset N \setminus \{i, j\}$ and an associated bid vector b_a that satisfy Assumption 5. Then by the same argument in Proposition 2, the distribution of $(\alpha \Delta q_{j,i} + \epsilon_j - \epsilon_i, Y_i(b_a, I_a) - \epsilon_i)$ for the given I_a, b_a is identified over the full support from the impact of independent variation in B_i and B_j on the probability that i wins when $A = a \cup \{i, j\}$, $B_i = b_i, B_j = b_j$ and the other entrants quote b_a .

Next, recover the distribution of $(\alpha \Delta q_{j,i}, Y_i(b_a, I_a))$ given I_a, b_a using the distribution of $(\alpha \Delta q_{j,i} + \epsilon_j - \epsilon_i, Y_i(b_a, I_a) - \epsilon_i)$ given I_a, b_a and the distribution of $(\epsilon_j - \epsilon_i, \epsilon_i)$, with the latter joint distribution already identified due to Proposition 2. To see why, note $(\epsilon_j - \epsilon_i, \epsilon_i)$ is independent from $(\alpha \Delta q_{j,i}, Y_i(b_a, I_a))$ for any fixed I_a and b_a due to Assumptions 1 and the characteristic function of (ϵ_i, ϵ_j) is non-vanishing. With the quality levels q_i and q_j already identified

due to Proposition 2, we can recover the distribution of $Y_i(b_a, I_a) + \alpha q_i$ conditional on α and I_a, b_a using the joint distribution of $(\alpha \Delta q_{j,i}, Y_i(b_a, I_a))$ conditional on I_a, b_a . This means $\psi_i(y, \alpha, b_a, I_a)$ is identified for all y .

Evaluate the linear system mentioned in Assumption 4 at the set of other entrants a with $x_s = x_i$ for all $s \in a$. Again, fix I_a, b_{a^p} and a^p (and therefore the quality and characteristics of sellers in a^p), and suppress them in the notation for ψ_i, λ and Λ_i . Then the right-hand side of (14) is simplified to

$$\Pr \{ \epsilon_s \leq y + b_s - \alpha q_s \ \forall s \in a \mid \alpha \} \quad (17)$$

which is identified because the distribution of match components is already recovered as above. Thus, by evaluating (15) at the same $(I_a, b_{a^t}, \tilde{b}_{a^t})$ but different values of (α, y) with $\Pr\{U_0 \leq y \mid \alpha\} > 0$ provides us with the linear system of $R + 2$ equations in $2K^*$ unknown weights $\{\lambda(b_{a^t}), \lambda(\tilde{b}_{a^t})\}$. Under the full rank condition in Assumption 4, $\{\lambda(b_{a^t}), \lambda(\tilde{b}_{a^t})\}$ is identified. Finally, with the $2K^*$ probability masses $\{\lambda(b_{a^t}), \lambda(\tilde{b}_{a^t})\}$ recovered, we can identify the conditional distribution $F_{U_0|\alpha}$ using the equality in (13).

Next, without loss of generality, suppose Assumption 5 holds for $i \in n$ (where n is a set of permanent sellers with the composition I) and a set of transitory sellers a with the composition $I' = I_a$ and associated bid vectors $b' \equiv b_a$. Let b_n denote the vector of bids from n . Under our assumptions,

$$\begin{aligned} & \Pr \{ i \text{ wins} \mid A = n \cup a, b_n \} \\ = & \Pr \left\{ \begin{array}{l} \Delta \epsilon_{s,i} + \alpha \Delta q_{s,i} + \Delta x_{s,i} \beta \leq b_s - b_i \ \forall s \in n \setminus \{i\} \\ Y_i(b_a, I_a) - \epsilon_i \leq -b_i \end{array} \right\} \end{aligned}$$

where as before $Y_i(b_a, I_a)$ denotes the maximum of $U_0 - \alpha q_i - x_i \beta$ and $\Delta x_{k,i} \beta + \alpha(Q_k - q_i) - b_k + \epsilon_k$ for $k \in a \subset N_t$. The support condition of the proposition implies the marginal distribution of $(\Delta \epsilon_{s,i} + \alpha \Delta q_{s,i} + \Delta x_{s,i} \beta)_{s \in n \setminus \{i\}}$ is identified. With the distribution of ϵ, α and the quality differences already identified, it follows from the independence between β and (α, U_0) that the joint distribution of $(\Delta x_{s,i} \beta)_{s \in n \setminus \{i\}}$ is identified over its full support. With $(x_s - x_i)_{s \in n \setminus \{i\}}$ having full-rank, the distribution of β is identified from that of $(\Delta x_{s,i} \beta)_{s \in n \setminus \{i\}}$ using Jacobian transformation (by multiplying the joint density of $(\Delta x_{s,i} \beta)_{s \in n \setminus \{i\}}$ with the absolute value of the determinant of $(x_s - x_i)_{s \in n \setminus \{i\}}$). \square

If the distribution of β is degenerate at a constant vector, the model is identified under weaker restrictions. To see this, recall that the first step in the proof of the second claim in Proposition 3 is to identify the joint distribution of $(\Delta \epsilon_{s,i} + \alpha \Delta q_{s,i} + \Delta x_{s,i} \beta)_{s \in n \setminus \{i\}}$. With ϵ independent from (α, β) and with the distribution of $(\Delta \epsilon_{s,i})_{s \in n \setminus \{i\}}$ identified, this implies means the joint distribution of $(\alpha \Delta q_{s,i} + \Delta x_{s,i} \beta)_{s \in n \setminus \{i\}}$ is identified. If β is a constant vector, then there are only

$\dim(\beta) + |n| - 1$ parameters to recover from the continuous distribution of such a $(|n| - 1)$ -vector $(\alpha\Delta q_{s,i} + \Delta x_{s,i}\beta)_{s \in n \setminus \{i\}}$. Recall that $\{x_i : i \in n\}$ is a constant vector of observed characteristics that do not vary throughout the data, and that the distribution of α is identified in Proposition 2. Hence we can use the second moments of $(\alpha\Delta q_{s,i} + \Delta x_{s,i}\beta)_{s \in n \setminus \{i\}}$ to identify $(\Delta q_{s,i})_{s \in n \setminus \{i\}}$, which in turn can be used to recover β under the rank condition maintained in the proposition. Note that this implies that, in order to recover the quality index of each permanent seller when β is constant, we only need to normalize the location of the lowest quality level for sellers in any one (as opposed to all) of the groups defined by observed characteristics.

Proof of Proposition 4: For a given seller i with type $\theta = (p, x, q)$, define

$$G_\theta(b) \equiv \Pr\{i \text{ wins} \mid i \in A, B_i = b\} = \Pr\{\varpi_i \leq -b\}$$

where ϖ_i denotes the maximum of $\max_{j \in A \setminus \{i\}} [\alpha(Q_j - q_i) + \Delta x_{j,i}\beta + \Delta \epsilon_{j,i} - B_j]$ and $U_0 - \alpha q_i - x_i\beta - \epsilon_i$. Note the randomness in ϖ_i comes from $\alpha, \beta, U_0, \epsilon$ and A as well as the bids $B_j, j \in A \setminus \{i\}$. (Note the definition of ϖ_i differs from that of $Y_i(b_a, I_a)$ because the former does not conditional on any vector of bids b_a or the composition of competing entrants I_a .) Provided sellers' equilibrium bidding strategies are differentiable, the smoothness conditions maintained in Assumptions 1 imply that G_θ is differentiable almost everywhere. Hence the first-order condition for bidder i choosing price b_i in equilibrium is:

$$(b_i - c_i) G'_\theta(b_i) + G_\theta(b_i) = 0. \quad (18)$$

Note that the conditional winning probability as a function of b_i , $G_\theta(b_i)$ is identified. Hence so is its derivative $G'_\theta(b_i)$. This then implies that the inverse bidding strategy (and consequently the distribution of private project cost C_i which depends on x_i, q_i but not ρ_i under Assumption 1) is identified for sellers with type $\theta = (p, x, q)$ as long as $G'_\theta(b_i) \neq 0$. ■

Proof of Proposition 5: The equilibrium entry threshold for a seller with type θ (which is a cutoff in the support entry signals below which the seller decides to enter), in the presence of exogenous project cost shifters $Z = z$, is characterized by:

$$t_\theta(z) = \mathbf{E} [\Pi_i(\sigma_\theta^*(C_i, Z), C_i, Z; p^*, \sigma^*) \mid Z = z] \quad (19)$$

where p^* is the vector of type-symmetric equilibrium entry probabilities, which are directly identifiable from the data. (Note there is a slight abuse of notation in that we now write Π_i as a function of p^* as opposed to τ^* . Also recall that, as stated at the beginning of Appendix A we have been suppressing the dependence on the set of potential bidders N as well as its composition I_N throughout the proofs.) The inverse equilibrium bidding strategies are recovered

as Proposition 4 and the distribution of bidders are identified; and the distribution of buyer tastes $(\alpha, \beta, U_0, \epsilon)$ is also identified. Hence the right-hand side of (19) is identified for all z . The left-hand side by definition is the $p_\theta^*(z)$ quantile of the distribution of entry costs for a seller with type θ .⁴⁴ ■

A3. Discussion of Support Conditions

We discuss how our model is capable of generating the price variation in the support condition used in our identification argument. We do so in the context of a stylized model that abstracts away from observed characteristics x_i , seller endogenous participation and presence of transitory sellers, but that has an allocation rule involving unobserved quality indices and stochastic match components. (Adding endogenous entry and transitory sellers complicates the algebra but does not involve any additional insight.)

Suppose the set of entrants in the stylized model is $A = \{i, j\}$, which is known to both participants. In this case, the support condition for identifying the distribution of match components is reduced to: “there exist i, j with $q_i = q_j$ such that the support of $(B_j - B_i, -B_i)$ includes that of $(\epsilon_j - \epsilon_i, \tilde{U}_0 - \epsilon_i)$ ”, where we let \tilde{U}_0 denote the difference of outside option U_0 and the weighted quality index αq_i . Note that, in a type-symmetric equilibrium which we consider here, the support of bids from i, j are identical, and denoted by $[b_i, \bar{b}_i]$.

A set of sufficient conditions for the support restrictions is: “there exist i, j with $q_i = q_j$ s.t. $\bar{b} > \bar{\epsilon} - \underline{u}_0$ and $\underline{b} < \underline{\epsilon} - \bar{u}_0$,” which can be satisfied if both (a) $\bar{b} > \bar{\epsilon} - \underline{u}_0$ and (b) $\bar{b} - \underline{b} > (\bar{u}_0 - \underline{u}_0) + (\bar{\epsilon} - \underline{\epsilon})$ hold. Condition (a) holds provided $\bar{b} \geq \bar{c}_i > \bar{\epsilon} - \underline{u}_0$, where \bar{c}_i is the supreme of the support of private costs for i and j . Condition (b) essentially requires the support of bids to be large relative to that of outside utility and match components. Intuitively, (b) also holds when the support of sellers’ private costs is sufficiently large. We now provide an argument for this intuition.

The idea is to show that the bidding strategy is continuous in the length of the support of match component $\bar{\epsilon} - \underline{\epsilon}$ and the support of outside option. Under type-symmetric pure-strategy BNE the bidders’ strategies solve the maximization problem:

$$\sigma_i(c) \equiv \arg \max_b (b - c) \Pr(\tilde{U}_0 - \epsilon_i \leq -b) \Pr(-B_j + \epsilon_j - \epsilon_i \leq -b) \quad (20)$$

The second probability in (20) represents i ’s belief, which is formed from i ’s knowledge of the distribution of private costs, and the distribution of quality indices in the population of sellers.

Suppose ϵ_i are i.i.d. uniform over $[\underline{\epsilon}, \bar{\epsilon}]$. Applying the Law of Total Probability, we can write

⁴⁴Note that this proof allows for the dependence between Z and I_N . For example, Z could be the number of potential bidders in an auction, i.e., the cardinality of N .

the objective function for seller i with costs c as

$$(b - c) \int_{\underline{\varepsilon}}^{\bar{\varepsilon}} \left[\left(\int_{\underline{\varepsilon}}^{\bar{\varepsilon}} \frac{1 - F_{B_j}(b - \Delta\varepsilon_{i,j})}{\bar{\varepsilon} - \underline{\varepsilon}} d\varepsilon_j \right) \frac{F_{\tilde{U}_0}(-b + \varepsilon_i)}{\bar{\varepsilon} - \underline{\varepsilon}} \right] d\varepsilon_i. \quad (21)$$

Changing variables between ε_r and $\tau_r \equiv \frac{\varepsilon_r - \underline{\varepsilon}}{\bar{\varepsilon} - \underline{\varepsilon}}$ for $r = i, j$, we can write (21) as

$$(b - c) \int_0^1 \left(\int_0^1 F_{\tilde{U}_0}(\underline{\varepsilon} - b + \delta\tau_i) [1 - F_{B_j}(b - \delta\Delta\tau_{i,j})] d\tau_j \right) d\tau_i \quad (22)$$

where $\delta \equiv \bar{\varepsilon} - \underline{\varepsilon}$ is the length of support of match components. Note (22) is continuous in both δ and the length of support for \tilde{U}_0 . It then follows from an application of the Theorem of Maximum that the support of bids is continuous in the size of the support of match component and outside options. Provided private costs vary sufficiently, the support of bids in a standard auction model with no match components (i.e. ϵ is degenerate at 0) and no outside option is an interval with non-degenerate interior. It follows from the implication of the Theorem of Maximum that condition (b) holds whenever the support of ϵ and \tilde{U}_0 is small enough. By the same token, we can provide similar structural justifications for the support conditions for recovering quality levels using i, j with $q_i \neq q_j$, as long as variation in private costs and quality differences are sufficiently large relative to that of buyers' tastes in $(\alpha, \epsilon, \tilde{U}_0)$. In the current example, adding transitory sellers would not require any qualitative change in the argument for the support conditions.

Appendix B. Representation Result

In this section, we derive a generalized version of the representation in (7) in the main text which is used in the estimation step. For each project l , recall that $N_l = N_l^p \cup N_l^t$ denotes the set of potential bidders for project l , where N_l^p denotes the set of potential permanent bidders and N_l^t the set of potential transitory bidders. Similarly, $A_l = A_l^p \cup A_l^t$ denotes the set of active bidders for projects l , where A_l^p and A_l^t denote the sets of permanent and transitory active bidders respectively. (From here on, we suppress the auction index l from notation for simplicity.)

Recall that the sets \bar{N}^p and \bar{N}^t denote the total set of permanent sellers and transitory sellers. As in the main text, we write A, A^p, A^t , etc. to denote the random set of entrants, and a, a^p, a^t , etc., to denote their particular realizations. For each set $a = a^p \cup a^t$ with $a^p \subset \bar{N}^p$ and $a^t \subset \bar{N}^t$, we recall the definitions of the compositions, I_a in Section 5.2.2.

When collecting individual observations to the project level we find it convenient to arrange observations in a certain order. More specifically, the observations for permanent and transitory sellers are allocated into separate vectors. Within each vector we enumerate observations for actual entrants first then for non-entrants. Then within the set of actual entrants, we order the

observations for permanent sellers according to (x, q) -characteristics in the order of

$$(\bar{x}_1, \bar{q}_1(\bar{x}_1)), \dots, (\bar{x}_1, \bar{q}_{K(\bar{x}_1)}(\bar{x}_1)), \dots, (\bar{x}_m, \bar{q}_{K(\bar{x}_m)}(\bar{x}_m))$$

and then for transitory sellers according to x -characteristics in the order of $\bar{x}, \dots, \bar{x}_m$. We order observations similarly for the set of non-entrants as well. Thus we maintain the order of individual bids as described above so that it is easy to match individual bids to seller characteristics using only information contained in $I = (I_A, I_N)$, i.e., the information on the sizes of (x, q) -groups of the potential and active permanent sellers and the sizes of x -groups of the potential and active transitory sellers in auction l . Indeed, since the bids are submitted by active sellers, the relevant part of I is I_{A^p} and I_{A^t} . Let us denote a vector of bids by $b = [b^{p'}, b^{t'}]'$, where b^p and b^t are shorthand for b_{a^p} and b_{a^t} , i.e., vectors of bids submitted by active permanent and transitory sellers respectively with their j -th entries denoted by b_j^t and b_j^p . If I_{A^p} reports that A^p contains two bidders from group $(\bar{x}_1, \bar{q}_1(\bar{x}_1))$ and three bidders from group $(\bar{x}_2, \bar{q}_1(\bar{x}_2))$ then we know that bids b_1^p and b_2^p are submitted by bidders from group $(\bar{x}_1, \bar{q}_1(\bar{x}_1))$ and bids b_3^p, b_4^p, b_5^p are submitted by sellers from group $(\bar{x}_2, \bar{q}_1(\bar{x}_2))$.

Recall that Q_i denotes the random quality of transitory seller i , and write

$$Q_A = (Q_i)_{i \in A} \text{ and } Q_a = (Q_i)_{i \in a}.$$

Similarly, for a nonstochastic vector $q = (q_i)$ of numbers q_i , we write $q_A = (q_i)_{i \in A}$ and $q_a = (q_i)_{i \in a}$. To simplify notation, we let

$$\omega(q_a; I_{a^t}) \equiv \prod_{i=1}^{|a^t|} \Pr(Q_i = q_i | x_i),$$

where a^t denotes the realized set of active transitory bidders, and $\Pr(Q_i = q_i | x_i)$ denotes the probability of transitory bidder i 's quality being q_i when his characteristic is x_i . This latter probability is a model primitive. We also let

$$g(b^t; q_a, I_{a^t}) \equiv \omega(q_a; I_{a^t}) \prod_{i=1}^{|a^t|} f(b_i | Q_i = q_i, I) \Pr(i \in A^t | Q_i = q_i, I).$$

Here $\Pr(i \in A^t | Q_i = q_i, I)$ is the conditional probability of a transitory bidder entering the auction when his quality is q_i . Note that $g(b^t; q_a, I_{a^t})$ depends on I_N though we suppress it from notation for simplicity. Let $\Pr(Q_A = q_A | b, I)$ denote the conditional probability of $Q_A = q_A$ given the bid vector b and composition I . (Note that this probability depends on A only through I .) Our representation result characterizes this conditional probability in terms of $g(b^t; q_a, I_{a^t})$'s.

Proposition 2. *Suppose that Assumptions 1-3 hold. Then,*

$$\Pr(Q_A = q_A | b, I) = \frac{g(b^t; q_A, I_{A^t})}{\sum_{q'_A \in \mathbb{Q}_A} g(b^t; q'_A, I_{A^t})}.$$

Proof: First, note that by the Bayes Rule, we can write

$$\begin{aligned} \Pr(Q_A = q_A | b, I) &= \frac{f_b(b | Q_A = q_A, I) \Pr(Q_A = q_A | I)}{f_b(b | I)} \\ &= \frac{f_b(b^t | Q_A = q_A, I) \Pr(Q_A = q_A | I)}{f_b(b^t | I)}. \end{aligned} \quad (23)$$

The second equality holds because the bids of permanent sellers are independent of bids and qualities of transitory sellers, i.e., $f_b(b^p | Q_A = q_A, I) = f_b(b^p | I)$.

Notice that b_j^t 's are independent across j 's conditional on $Q_A = q_A$ and I , and that out of I_{A^t} only x_j^t matters for the distribution of bids that could be submitted by a transitory seller j . Therefore,

$$f_b(b^t | Q_A = q_A, I) = \prod_{j=1}^{|A^t|} f_b(b_j^t | Q_A = q_A, I) = \prod_{j=1}^{|A^t|} f_b(b_j^t | Q_j = q_j, I). \quad (24)$$

The last equality holds because the transitory seller knows his quality but not the qualities of his transitory competitors.

On the other hand, applying the law of total probability we obtain

$$\begin{aligned} f_b(b^t | I) &= \sum_{q_A \in \mathbb{Q}_A} f_b(b^t | Q_A = q_A, I) \Pr(Q_A = q_A | I) \\ &= \sum_{q_A \in \mathbb{Q}_A} \left(\prod_{j=1}^{|A^t|} f_b(b_j^t | Q_j = q_j, I) \right) \Pr(Q_A = q_A | I). \end{aligned} \quad (25)$$

We will use this expression later, after we deal with $\Pr(Q_A = q_A | I)$.

First, we write $I_A = I_{A^p} \cup I_{A^t}$ and let \bar{I}_A be a realized set of I_A . Then notice that

$$\begin{aligned} \Pr(Q_A = q_A | I_N, I_A = \bar{I}_A) &= \frac{\Pr(I_{A^p} = \bar{I}_{A^p}, I_{A^t} = \bar{I}_{A^t} | Q_A = q_A, I_N) \Pr(Q_A = q_A | I_N)}{\Pr(I_{A^p} = \bar{I}_{A^p}, I_{A^t} = \bar{I}_{A^t} | I_N)} \\ &= \frac{\Pr(I_{A^t} = \bar{I}_{A^t}, Q_A = q_A | I_N)}{\sum_{\tilde{q}_A \in \mathbb{Q}_A} \Pr(I_{A^t} = \bar{I}_{A^t}, Q_A = \tilde{q}_A | I_N)}. \end{aligned} \quad (26)$$

The second equality holds because the events $I_{A^p} = \bar{I}_{A^p}$ and $I_{A^t} = \bar{I}_{A^t}$ are independent conditional on $Q_A = q_A, I_N$, and the event $I_{A^p} = \bar{I}_{A^p}$ is independent of $Q_A = q_A$ conditional on I_N .

We write $\Pr(I_{A^t} = \bar{I}_{A^t}, Q_A = q_A | I_N)$ as

$$\begin{aligned}
& \sum_{q_{N \setminus A} \in \mathbb{Q}_{N \setminus A}} \Pr(I_{A^t} = \bar{I}_{A^t}, Q_A = q_A, \text{ and } Q_{N \setminus A} = q_{N \setminus A} | I_N) \\
&= \sum_{q_{N \setminus A} \in \mathbb{Q}_{N \setminus A}} \Pr(I_{A^t} = \bar{I}_{A^t}, Q_N = [q_A; q_{N \setminus A}] | I_N) \\
&= \sum_{q_{N \setminus A} \in \mathbb{Q}_{N \setminus A}} \Pr(I_{A^t} = \bar{I}_{A^t} | Q_N = [q_A; q_{N \setminus A}], I_N) \Pr(Q_N = [q_A; q_{N \setminus A}] | I_N).
\end{aligned} \tag{27}$$

As for the summands in the last expression, we define

$$\begin{aligned}
p_i^A(q_i) &= \Pr(i \in A^t | Q_i = q_i, I_N) \text{ and} \\
p_i^{N \setminus A}(q_i) &= \Pr(i \in N^t \setminus A^t | Q_i = q_i, I_N),
\end{aligned}$$

where q_i denotes the i -th entry of q_N . Thus $p_i^A(q_i) > 0$ only when $i \in A^t$ and similarly $p_i^{N \setminus A}(q_i) > 0$ only when $i \in N^t \setminus A^t$. Rewrite

$$\Pr(I_{A^t} = \bar{I}_{A^t} | Q_N = q, I_N) = \left(\prod_{i=1}^{|\alpha^t|} p_i^A(q_i) \right) \left(\prod_{i=|\alpha^t|+1}^{|N^t|} p_i^{N \setminus A}(q_i) \right).$$

This is because participation decisions of transitory sellers do not depend on the qualities of their transitory competitors since the qualities of transitory sellers are only privately known. On the other hand, we write

$$\Pr(Q_N = q | I_N) = \prod_{i=1}^{|N^t|} \Pr(Q_i = q_i | I_N) = \prod_{i=1}^{|N^t|} \Pr(Q_i = q_i | x_i^t),$$

because potential bidders are randomly drawn from the population characterized by a given vector x and from the whole structure I_N only the information about x_i^t is relevant in this probability. Combining this, we write $\Pr(I_{A^t} = \bar{I}_{A^t}, Q_A = q_A | I_N)$ as

$$\begin{aligned}
& \prod_{i=1}^{|N^t|} \Pr(Q_i = q_i | x_i^t) \left(\prod_{i=1}^{|\alpha^t|} p_i^A(q_i) \right) \left(\prod_{i=|\alpha^t|+1}^{|N^t|} p_i^{N \setminus A}(q_i) \right) \\
&= \left(\prod_{i=1}^{|\alpha^t|} p_i^A(q_i) \Pr(Q_i = q_i | x_i^t) \right) \left(\prod_{i=|\alpha^t|+1}^{|N^t|} p_i^{N \setminus A}(q_i) \Pr(Q_i = q_i | x_i^t) \right).
\end{aligned} \tag{28}$$

Applying this result to the denominator of (26), we find that the numerator and the denominator have a common factor:

$$\prod_{i=|\alpha^t|+1}^{|N^t|} p_i^{N \setminus A}(q_i) \Pr(Q_i = q_i | x_i^t).$$

After canceling out this factor, we conclude that

$$\Pr(Q_A = q_A | I) = \frac{\prod_{i=1}^{|A^t|} p_i^A(q_i) \Pr(Q_i = q_{A,i} | x_i^t)}{\sum_{\tilde{q}_A} \prod_{i=1}^{|A^t|} p_i^A(\tilde{q}_i) \Pr(Q_i = \tilde{q}_i | x_i^t)} = \frac{\left(\prod_{i=1}^{|A^t|} p_i^A(q_i) \right) \omega(q_A; I_{A^t})}{\sum_{\tilde{q}_A} \left(\prod_{i=1}^{|A^t|} p_i^A(\tilde{q}_i) \right) \omega(\tilde{q}_A; I_{A^t})}, \quad (29)$$

recalling the definition of $\omega(q_a; I_{A^t})$. Now, plugging this expression into (25), we can write

$$f_b(b^t | I) = \sum_{q_A \in \mathbb{Q}_A} \frac{\omega(q_A; I_{A^t}) \prod_{j=1}^{|A^t|} f_b(b_j^t | Q_j = q_j, I) p_i^A(q_i)}{\sum_{\tilde{q}_A} \omega(\tilde{q}_A; I_{A^t}) \prod_{j=1}^{|A^t|} p_i^A(\tilde{q}_i)}. \quad (30)$$

Thus we plug (24), (30) and (29) into (23) to obtain the desired result. ■

Appendix C. Choice of Moment Conditions

10.1 Basic Moment Restrictions

For the choice of g_j 's, we use functions that vary with the active sellers' (x, q) -group memberships and their composition I_A 's. For each auction l such that $I_{A_l} = \bar{I}_A$ for a given composition \bar{I}_A and have a winning seller j in (x, q) -group, we take $g_j(B_l, \bar{I}_A)$ as follows. (Here, we let $A_{l,x,q}$ be the set of active sellers in auction l in (x, q) -group and $A_{l,x}$ that in x -group. Then for each (x, q) and (x', q') and for a given composition \bar{I}_A , we consider three types of g_j 's (except for constant 1)) as follows.

(1) Moments involving a single winning bid:

$$g_j(B_l, I_{A_l}) = B_{j,l} 1\{I_{A_l} = \bar{I}_A\} 1\{j \in A_{l,x,q}\}.$$

(2) Moments involving a winning bid and another bid:

$$g_j(B_l, I_{A_l}) = \sum_{i \in A_{l,x',q'}} (B_{j,l} - B_{i,l}) 1\{I_{A_l} = \bar{I}_A\} 1\{j \in A_{l,x,q}\}.$$

(3) Moments involving two permanent sellers' bids and one transitory seller's bid:

$$g_j(B_l, I_{A_l}) = \sum_{i \in A_{l,x,q}} \sum_{h \in A_{l,x}} (B_{j,l} - B_{i,l}) B_{h,l} 1\{I_{A_l} = \bar{I}_A\}.$$

We also use moments in (i) with $B_{j,l}$ replaced by $B_{j,l}^2$, and in (2) with $(B_{j,1} - B_{i,l})$ replaced by $(B_{j,1} - B_{i,l})^2$ and by $(B_{j,1} - B_{i,l}) B_{j,l}$, and in (3) with $(B_{j,l} - B_{i,l}) B_{h,l}$ replaced by $B_j B_{h,l}$, $B_j^2 B_{h,l}$, $B_j x_h$ and $B_j^2 x_h$. In the actual implementation, we re-weight all the moments by the frequency

with which relevant observations appear in the data.

10.2 Additional Restrictions

We also use additional restrictions that stem from restrictions on transitory sellers' bid distributions, transitory sellers' participation probabilities, and restrictions that come from expected profit conditions.

(a) The restriction associated with transitory sellers' bid distribution:

$$f(b_{a^t}|I) = \frac{1}{J} \sum_{q'_a \in \mathbb{Q}_a} \omega(q_a; I_{a^t}) \prod_{i=1}^{|a^t|} f(b_i|Q_i = q_i, I) \Pr(i \in A^t|Q_i = q_i, I), \quad (31)$$

where

$$J = \sum_{q'_a \in \mathbb{Q}_a} \omega(q_a; I_{a^t}) \prod_{i=1}^{|a^t|} \Pr(i \in A^t|Q_i = q_i, I).$$

Moment conditions associated with this restriction relate the empirical moments of $f(b_{a^t}|I)$ to the theoretical moments based on the restriction (31).

(b) The restriction associated with the transitory sellers' participation probability:

$$\begin{aligned} & \Pr(I_{A^t} = \bar{I}_{A^t}|I_N) \\ &= \sum_{q_A \in \mathbb{Q}_A} \left(\prod_{i=1}^{|a^t|} p_i^A(q_i) \Pr(Q_i = q_i|x_i^t) \right) \left(\prod_{i=|a^t|+1}^{|N^t|} p_i^{N \setminus A}(q_i) \Pr(Q_i = q_i|x_i^t) \right). \end{aligned} \quad (32)$$

This restriction is based on (28) in the proof of Proposition 2. Moment conditions associated with this restriction relate the transitory sellers' empirical probability of participation and expected x -characteristics of entrants conditional on I to their theoretical counterpart using (32).

(c) The restriction related to the expected profit condition. This restriction summarizes optimal participation behavior. It is summarized by the threshold strategy where potential bidders with entry cost draws below the ex-ante expected profit participate in the auctions and those with higher draws stay out. This implies that in equilibrium: for each $q^t \in \mathbb{Q}_x$ and $\theta_j = (t, x_j, q^t)$,

$$\Pr(j \in A^t|\theta = \theta_j, I_N) = F_E(\mathbf{E}[\Pi(\theta_j, I_N)]),$$

where $F_E(\cdot)$ is the distribution function of entry costs E .