EMPIRICAL ESSAYS ON BIDDER BEHAVIOR IN AUCTIONS

by

Lucia Tiererova

A dissertation submitted to The Johns Hopkins University in conformity with the requirements for the degree of Doctor of Philosophy.

Baltimore, Maryland

October, 2014

Abstract

This dissertation empirically studies bidder behavior in procurement auction environments. I apply auction theory to develop models explaining market outcomes in two distinct settings. The empirical analysis uses detailed bid-level data and takes advantage of the particular features of each market to estimate the models. I then use the estimated parameters to conduct counterfactual analyses and identify specific drivers of bidder behavior.

In the first setting, I study government procurement auctions for road repair and construction projects in California. This market is characterized by significant turnover of participants and there is a number of government programs in place designed to prolong firm tenure. The first aspect of this market which I investigate in Chapter 2 is learning-by-doing among inexperienced firms that have recently entered this market. I extend previous methods for estimation of dynamic auction games by developing a dynamic model of experience accumulation that allows for intertemporally-linked costs and endogenous auction participation. I document the presence of learning-by-doing in the data, and use the theoretical model to show that bidders with a higher level of experience have lower average costs than inexperienced bidders. Consistent with dynamic incentives, my empirical results indicate that inexperienced firms charge lower markups and bid more aggressively in order to leverage learning-by-doing and achieve cost reductions in the future.

Next, in Chapter 3, I address the subcontracting side of the highway procurement auction market in California. I empirically investigate whether experience in the subcontracting market can improve future prospects of new entrants. The evidence on learning-by-doing from Chapter 2 suggests that experience obtained in the primary market plays an important role in the future success of such firms. However, many recent entrants participate as subcontractors before actively submitting bids as primary contractors. I conclude that despite the fact that experience gained as a subcontractor might be more limited in scope and task-specific, it can be a valuable stepping stone for firms that have only recently entered this market.

Finally, in Chapter 4, which is joint work with Dr. Elena Krasnokutskaya and Dr. Christian Terwiesch, I turn my attention to a very different auction market. In this case, we analyze data from an online market for programming services to study bilateral trading patterns that arise between buyers and sellers from distinct world regions. Our methodology allows us to control for both the country-specific differences in seller quality and the country-specific differences in buyers' preference for particular bidder attributes relative to price, as well as for differences in their outside options. We decompose the clustering in trading into effects generated by various components

ABSTRACT

of the model. We find that differences in cost distributions across seller countries drive most of the clustering pattern.

Primary Reader: Professor Elena Krasnokutskaya

Secondary Reader: Professor Richard Spady

Acknowledgments

I am indebted to my main advisor, Professor Elena Krasnoskutskaya whose guidance and comments shaped this dissertation. She helped me learn how to apply auction theory to data, and how to transform a messy dataset into economic wisdom. However, she has taught me more than just economics. She has been the most generous advisor and an inspiring mentor, and I thank her wholeheartedly for her advice and encouragement.

I would also like to thank Professors Richard Spady, Andrew Sweeting, and Jorge Balat, as well as all the participants of the JHU Applied Micro Seminar and IO Lunch Seminar for comments and helpful discussions. My training as a teacher and instructor owes a great deal to Professor Bruce Hamilton, who served as a great source of teaching advice.

Many of my graduate student colleagues made the past six years memorable but I absolutely could not have done it without the support of Sangah Lee and Ruli Xiao. I also appreciate the encouragement of my job-market colleagues Jiae Yoo, Natsuki Arai and Yoichi Goto.

ACKNOWLEDGMENTS

I am greatly appreciative of the support and encouragement of my parents, Eva and Jozef, as well as my sister Katka. While the dissertation is dedicated to my parents, Andrew Stokes, my partner of seven years, is the one who inspired me daily and helped me persevere even in the darkest moments. He is the one who shared the pain and the joy of producing this dissertation, and I am thankful for his unconditional support.

Dedication

For my loving parents, Eva and Jozef,

who believed in me throughout this journey...

Contents

List of Tables

List of Figures

Chapter 1

Introduction

This dissertation empirically studies bidder behavior in procurement auction environments. I apply auction theory to develop models explaining market outcomes in two distinct settings. The empirical analyses use detailed bid-level data and take advantage of the particular features of each market to estimate the models. I then use the estimated parameters to conduct counterfactual analyses and identify the main drivers of bidder behavior. Specifically, in the first setting I study the effects of prior contracting and subcontracting experience on bidding and participation strategies among recent entrants in highway procurement auctions; in the second setting, I analyze bilateral trading patters between buyers and sellers in an online market for programming services.

I begin by considering government procurement auctions for road repair and construction projects in California. This market is characterized by significant turnover of participants, and there are a number of government programs in place designed to prolong firm tenure. In this dissertation I study two related aspects of this market. First, in Chapter 2 I investigate learning-by-doing among inexperienced firms that have only recently entered this market. In Chapter 3 I consider an even earlier period in firms' histories, which is their subcontracting experience prior to joining the market as primary contractors.

In Chapter 2 I develop a dynamic model of experience accumulation that allows for intertemporally-linked costs and endogenous auction participation. I extend previously developed methods for estimation of dynamic auction games and document learning-by-doing in highway procurement auctions in California. I find that the average cost for bidders with the highest level of experience is 10% below the average cost for the least experienced bidders. Consistent with dynamic incentives, my empirical results indicate that inexperienced firms charge lower markups and bid more aggressively in order to leverage learning-by-doing and achieve cost reductions in the future. Using counterfactual analysis I conclude that if all new bidders became experienced the stronger competition would benefit the auctioneer through a 7% decrease in the average procurement cost. This evidence provides a novel insight for affirmative action in procurement auctions: in presence of learning-by-doing, bid preference programs help bidders more quickly reap the benefits of experience and thus have important long-term considerations extending beyond the period in which the discount is given.

Chapter 3 examines firms' level of experience as subcontractors prior to joining the main market. Evidence on learning-by-doing presented in Chapter 2 suggests that experience likely plays an important role in the future success of recent entrants. However, many firms in this market participate as subcontractors before actively submitting bids as primary contractors. Serving as a potential subcontractor to one or more firms increases the probability of carrying out at least a portion of the project workload. While experience gained as a subcontractor is more limited in scope and task-specific, it can be a valuable stepping stone for firms that have only recently entered the market. This chapter therefore focuses on subcontractor experience accumulated before entering the procurement auction market as a primary contractor.

In Chapter 4, which is joint work with Dr. Elena Krasnokutskaya and Dr. Christian Terwiesch, we study bilateral trading patterns between different countries using data from an online market for programming services. Our methodology allows us to control for the country-specific differences in the quality of participating sellers as well as the country-specific differences in buyers' price sensitivity and their outside options. We decompose the clustering in trading into the effects generated by different components of the model and assess gains from trade in this market.

Chapter 2

Dynamic Model of Bidder Learning in Procurement Auctions

Procurement auctions account for a large proportion of expenditures at all levels of government. According to the U.S. Economic Census in 2005, the U.S. federal government spent \$378 billion using procurement auctions, with roughly 20% of these funds going to highway, street and bridge construction.[1] Government procurement auctions often include controversial bid preference programs that provide some bidders favorable treatment.[2] Frequent targets of such policies are small businesses, women, minority or veteran-owned firms, or firms hiring local subcontractors. The

[1]California receives a significant proportion of these expenditures, with \$10.6 billion of spending on construction projects in 2007.

[2]For example, a favored bidder with a 5% bid preference would win an auction despite not submitting the lowest bid, if her bid was within 5% of the lowest bid. However, bid preference is only applied for bid comparison, so the bidder would receive her full bid amount for carrying out the work.

controversy arises because the motivation for these programs is often political rather than economic and the discussion centers around fairness of redistribution of government resources rather than the underlying economic rationale. I contribute to the limited evidence on the long-run effects of bid preference programs by providing a novel economic insight regarding the benefits arising from the prolonged presence of firms in market environments with learning.

In this chapter, I study learning-by-doing among inexperienced bidders in repeated procurement auctions for highway construction projects. When bidders' costs improve with additional experience the benefits of winning a contract extend beyond the current period by helping the winning firm face more advantageous cost distribution in the future. Dynamic incentives then drive firms to bid more aggressively at lower levels of experience and forgo current profits in order to gain more experience. The learning effect depicts the improving ability of firms to master the complex construction process that often consists of a large number of tasks and requires a high-degree of coordination and government supervision.

To capture bidder learning, I develop and estimate a dynamic auction model of experience accumulation that allows for intertemporal linking of costs for new firms and endogenous auction participation. In the first stage, eligible firms simultaneously decide whether or not to participate in the auction. In the second stage, participating firms submit bids and the winner is determined in a first-price sealed-bid auction. To allow for learning-by-doing firms that only recently entered the procurement market are modeled to have flexible cost distributions that can evolve over time. In particular, the average cost is allowed to depend on the level of experience, as measured by the cumulative number of all contracts awarded to a firm in the past. Winning a project then has implications extending into the future by improving the average cost in all following auction rounds. This incentive leads inexperienced firms to become more competitive and lower current markups to acquire additional experience. The participation and bidding decisions therefore depend not only on the set of competitors but also on their respective levels of experience.

The model is estimated using data from procurement auctions administered by the California Department of Transportation (Caltrans) during the period from July 2003 to June 2009. The focus of this paper is on inexperienced firms, which are not observed sufficiently often to estimate firm-specific behavior. However, I expect firms with similar levels of experience to undergo comparable learning-by-doing, which allows me to pool firms together by experience level. I extend and apply the estimation methods for dynamic auction games previously developed by Jofre-Bonet and Pesendorfer (2003).[1] The two-step estimation starts by using the observed data to compute the entry probabilities, state transitions and bidding strategies conditional on state variables. In the second step, I employ the theoretical predictions of optimal behavior to recover the parameters of the cost distributions.

The empirical results document learning-by-doing in this market. I find that for recent entrants accumulating experience results in driving down future average cost and helping them become more competitive. Comparing bidders across experience levels, the average cost of the most experienced bidders is about 10% below the cost of the least experienced ones. Consistent with dynamic incentives, the empirical results show that bidders charge lower markups at lower levels of experience. As expected, bidders who consider the dynamic effects of learning-by-doing will forgo current profits in order to achieve cost reductions in the future.

The counterfactual analysis further explores the economic significance of bidder learning. I find that bidder experience has important implications for firms' costs and markups as well as for the auction outcomes from the perspective of the auctioneer. As a result of leveraging learning-by-doing at low levels of experience, more experienced firms are able to bid more competitively while charging higher markups. As new firms gain experience their costs and markups approach the costs and markups of regular firms. The implications of bidder learning are substantial, especially on large projects, where exogenous improvement of a bidder's experience from the minimum to the maximum level would result in a 11% drop in the average cost. In addition, I examine changes in the auction outcome when bidders are forced to face stronger competition. Since bidders respond strategically to their opponents, changing the experience level of all new bidders to the highest possible level would intensify competition and result in a 7% decrease in the average procurement cost, benefiting the auctioneer.

These results confirm the importance of dynamic incentives in environments with learning-by-doing and suggest that policies affecting the speed of experience acquisition have important long-term benefits that have not yet been considered. This evidence provides a novel insight for affirmative action in procurement auctions, and indicates that the effects of bid preference programs can extend well beyond the period in which the discount has originally been provided.

This paper relates to the literature on the estimation of dynamic auctions and builds on the model developed by Jofre-Bonet and Pesendorfer (2003)[1] to examine the effects of capacity constraints. Jofre-Bonet and Pesendorfer consider dynamic behavior in response to the level of backlog and studied the interaction among the largest and most frequent bidders. Balat (2013)[2] extended this setup by accounting for unobserved auction heterogeneity and endogenous participation. In contrast, the focus of this paper is on firms which can benefit from learning-by-doing *while* establishing themselves in the market. Unlike backlog, which measures how stretched out a firm's resources are, I use the number of projects won in the past as a measure of experience that captures the cumulative knowledge and ability to carry out construction projects. Importantly, backlog and experience drive costs in the opposite directions: high levels of backlog tend to increase both future costs and the level of experience; but with bidder learning, high levels of experience tend to decrease costs.

Groeger (2013)[3] considers dynamics in the bid preparation rather than construction costs, and studies the relationship between participation in the most recent auction and future entry costs. This type of learning can be exploited by re-ordering of contracts. Other changes to auction rules have been analyzed in a variety of static auctions with endogenous participation by Athey et al.(2010),[4] Levin and Smith (1994),[5] Roberts and Sweeting (2010).[6] Saini(2012)[7] studies the optimal auction schedule with capacity-constrained bidders in a theoretical model.

Secondly, this paper relates to the large existing literature on learning-by-doing. Learning-by-doing has been explored in a variety of settings, from aircraft manufacturing (Benkard (2000)[8]) to service organizations (Darr et al.(1995)[9]). Despite abundance of research on knowledge transfers within and across organizations, learning-by-doing among bidders in repeated procurement auctions has not yet been examined.

Finally, the results presented in this paper have implications for the literature on affirmative action in procurement auctions. Theoretically, bid preference programs could increase or decrease the cost of procurement depending on the participation and competitive response of favored and non-favored firms. Empirically, Marion (2007)[10] and Krasnokutskaya and Seim (2011)[11] have estimated the cost of such programs to be relatively small. However, these estimates are based on static models and therefore capture only the immediate effects of bid preference programs without accounting for the long-term benefits that would arise from bidder learning.

2.1 Caltrans Data and Empirical Motivation

This section describes the data used in the estimation, provides summary statistics of relevant variables and presents preliminary evidence showing that participation and bidding behavior of new firms depend on their experience. I find that more experienced firms participate more often and submit lower bids than less experienced firms. While these reduced-form results suggest that firms become stronger with experience, without a full dynamic model it is not possible to determine whether this is due to changing costs or changing markups.

The data consist of 3,415 contracts awarded by Caltrans over the time period from July 2003 through June 2009. Data from the initial six months are used to construct history of firms, leaving 3,292 projects that are used in the estimation. Projects are awarded to the lowest bidder[3] in repeated first-price sealed-bid auctions and vary in scale and type, ranging from simple lane markings and signage, through medium-size projects like lane resurfacing to large ones involving bridge or highway construction. The data contain information on project characteristics such as the type of work and location of the project, estimated number of working days, engineer's estimate of project cost, source of funding (federal/state), date at which the project was first advertised, letting date, as well as the winning bidder and the winning bid.

[3]Caltrans runs a pre-qualification process which ensures required quality of work.

Furthermore, the data include information on bidders and planholders such as the name and address of the business, decision to submit a bid as well as the bid amount.

Overall, there are 2,482 unique firms that request planholder materials but only 1,057 of these firms submit at least one bid, and only 461 of these firms win at least one auction. I categorize firms according to their level of experience into regular, new, and fringe firms. There are 18 regular firms, which are the largest, most frequent participants and have won more than 30 auctions over the period. On the other hand, there are 500 small firms that are potentially subject to learning-by-doing and are the focus on this paper. The remaining category consists of 1,964 fringe firms, which never won and submitted fewer than 6 bids in their lifetime. Only 541 out of these fringe firms ever submitted any bid.

Table 2.1 presents summary statistics for the size and duration of projects, as well as the number of planholders and bidders per project. The average size of a project is $4.5 million, with the average duration of about 5 months. However, there is substantial variability among projects, with projects ranging in size from the smallest of $74,000 to the largest of $1.5 billion. Similarly, projects range in duration from 5 days to almost 7 years.

Another view of the data is provided by looking at summary statistics by bidder type, presented in Table 2.2. As expected, regular bidders are the most active and the most frequent participants both in the role of planholders as well as bidders. On average, regular bidders ordered planholder materials for 377 projects, resulting

Table 2.1: Summary Statistics of Contract Characteristics

	Mean	SD	Min	Max	Total
Engineer's Estimate ($million)	4.5	29.1	0.1	1,480	14,883
Working days	148	213	5	2,490	-
Number of planholders	13.6	8.6	2	73	44,856
Regular	2.1	1.7	0	9	6,790
New	8.7	5.3	0	38	28,601
Fringe	2.9	4.2	0	53	9,465
Number of bidders	5.4	3.3	1	32	17,861
Regular	1.2	1.2	0	7	4,102
New	3.7	2.6	0	21	12,253
Fringe	0.5	1.1	0	13	1,506

Note: Based on the sample of 3,292 projects used in the estimation.

in 228 submitted bids of which 57 were winning bids. In contrast, new bidders ordered planholder materials for 61 projects and submitted 26 bids leading to 5 wins. There is substantial variability even within bidder types, with the most successful firm winning 260 projects based on 1,083 bids, after ordering planholder materials for 1,412 projects, which is almost half of all auctioned projects. There rest of the market consists of fringe firms, which are likely subcontractors ordering planholder materials with the intention of becoming familiar with the project for negotiating purposes rather than with the intention to submit a bid. Overall, more than 68% of projects are awarded to new firms with the rest awarded to regular firms, indicating the economic significance of small firms in this market.

Table 2.2: Summary Statistics - Behavior of Firms by Type

	Mean	SD	Min	Max	Total
Number of plans ordered					44,856
Regular bidders	377.2	300.7	128	1,412	6,790
New bidders	60.6	73.4	1	598	30,297
Number of bids submitted					17,861
Regular bidders	227.9	230.2	90	1,083	4,102
New bidders	25.6	32.5	0	264	12,782
Number of contracts awarded					3,292
Regular bidders	57.0	53.2	30	260	1,026
New bidders	4.5	6.0	0	30	2,266

Note: Based on the sample of 3,292 projects used in the estimation.

2.1.1 Reduced-Form Results

Using the Caltrans data described in the preceding section, I present reduced-form evidence showing the relationship between bidder's history of successful bidding and future participation and bidding behavior. The results indicate that additional past experience, even after controlling for backlog, increases the probability of participation in the future and allows bidders to submit more competitive bids. While this preliminary evidence suggests the presence of learning, only a full dynamic model can provide insight into the effects of experience on the underlying costs.

I consider the effect of experience on the probability of entry, followed by results exploring the relationship between experience and the level of the bid. Table 2.3 shows the results of a probit regression of the decision to participate in an auction regressed on the level of past experience, estimated on the set of new planholders. In order to separate the effects of backlog and experience, I partition the measure of

Table 2.3: Probability of Participation Regressed on the Number of Past Awarded Contracts)

	(1)	(2)
log(number of contracts within past 6 months)	0.032***	0.036***
	(0.007)	(0.009)
log(number of contracts 6 to 12 months in the past)	0.050***	0.064***
	(0.007)	(0.009)
log(number of contracts 12 to 18 months in the past)	0.011	0.015
	(0.008)	(0.011)
log(number of contracts more than 18 months in the past)	0.022***	0.000
	(0.002)	(0.004)
log(Engineer's estimate)	−0.124***	−0.015
	(0.006)	(0.009)
Funding Source (Federal=1)	−0.049***	−0.014
	(0.017)	(0.028)
Number of Planholders - Regular	−0.042***	−0.103***
	(0.005)	(0.008)
Number of Planholders - New	−0.008***	−0.031***
	(0.001)	(0.002)
Number of Planholders - Fringe	0.000	0.011***
	(0.002)	(0.003)
Standardized backlog	−0.090***	−0.117***
	(0.012)	(0.018)
Year, Month and District Indicators	N	Y
N	30.297	30,297

Note: Standard errors provided in parentheses. All specifications also include a constant. Statistical significance: 1%(***), 5%(**), 10%(*).

experience into four non-overlapping intervals. The regression presented in the second column controls for seasonality and regional differences by including year, month and district dummies. The coefficients on the past experience are positive and significant, suggesting that more experienced bidders participate in auctions more frequently, even after controlling for their level of backlog. As expected, new bidders are less likely to participate when they face a larger number of competitors, whether those competitors are regular or new firms.

Table 2.4: Log(Bid) Regressed on the Number of Past Awarded Contracts

	(1)	(2)
log(number of contracts within past 6 months)	−0.001	−0.004**
	(0.002)	(0.002)
log(number of contracts 6 to 12 months in the past)	−0.011***	−0.013***
	(0.002)	(0.002)
log(number of contracts 12 to 18 months in the past)	0.000	−0.003
	(0.003)	(0.002)
log(number of contracts more than 18 months in the past)	−0.009***	0.001*
	(0.001)	(0.001)
log(Engineer's estimate)	0.987***	0.994***
	(0.004)	(0.004)
Funding Source (Federal=1)	0.016	−0.024**
	(0.011)	(0.011)
Number of Planholders - Regular	−0.032***	−0.016***
	(0.004)	(0.004)
Number of Planholders - New	−0.018***	−0.007***
	(0.002)	(0.002)
Number of Planholders - Fringe	−0.037***	−0.024***
	(0.007)	(0.005)
Standardized backlog	−0.003	0.008**
	(0.004)	(0.004)
Year, Month and District Indicators	N	Y
N	12.782	12.782

Note: Standard errors provided in parentheses. All specifications also include a constant. Statistical significance: 1%(***), 5%(**), 10%(*).

Table 2.4 presents the results of an OLS regression of the logarithm of the bid amount on the past experience. Again, past experience is partitioned into non-overlapping intervals. Even after controlling for backlog, past experience is negatively correlated with bid amount, indicating that more experienced firms are able to submit lower bids. Note that the effect of backlog is positive, as expected, since higher levels of backlog constrain firm's resources, tend to increase marginal costs, and make winning less desirable. On the other hand, if learning-by-doing is present in this market, experience will have the opposite effect on costs, and can be consistent with a decrease in bid levels.

Number of contracts won in the past is one among many possible measures of experience. To confirm that other measures would yield similar results, Appendix 2.9 shows corresponding regression results based on the dollar value of past awarded projects as the independent variable, rather than the number of awarded projects. The results are qualitatively similar, confirming that the number of contracts won in the past is an appropriate measure of the overall level of experience.

Preliminary reduced-form evidence suggests the presence of bidder learning: firms new to the market are less active (submit fewer bids) than their competitors, and conditional on participation the bid amount decreases with past experience, as measured by either the number of submitted bids or the value of awarded contracts. New firms are able to lower their bids perhaps due to facing more advantageous cost distributions, helping them to be better able to compete against their stronger competitors; this holds even when controlling for backlog. This is consistent with long-run effects of bidder learning, when costs are lowered for firms that are able to remain in the market for longer periods of time and become more competitive. However, without a full dynamic model it is not possible to determine whether this decline in bids is due to cost synergies or simply reduced markups.

2.2 Model

I start by constructing a theoretical model of repeated procurement auctions with flexible cost functions, which can capture learning-by-doing experienced by new firms as they enter the market for highway procurement, and firm-level participation decision. After explaining the setup, timing, and assumptions used in the model, I proceed by defining the state variable and the state transition. I conclude by formulating the components of the participation and bidding decisions, and outlining the steps involved in computing the value function and expressing it in terms of observables.

The model is a dynamic bidding game in which new firms' costs are inter-temporally linked through their past experience; as new firms gain experience through project construction, it translates into lower future construction costs. Furthermore, this model takes advantage of the availability of data on the set of planholders and explicitly includes firm-level participation decision.

The framework builds on the model originally developed by Jofre-Bonet and Pesendorfer (2003)[1] to study capacity constraints among large firms. In contrast, I focus on smaller firms and their bidding experience as they establish themselves in the market.

2.2.1 Setup

There are three types of risk-neutral firms: regular firms, which have extensive experience and therefore their costs are unlikely to change due to winning additional auctions; new firms, which are the focus of this paper, and which have costs intertemporally linked through their level of experience, and therefore subject to learning-by-doing and most likely to benefit from additional experience; and, finally, short-lived fringe firms, which are subcontracting firms that request project materials but rarely proceed in submitting a bid. Denote the type of firm i by $k(i) \in \{reg, new, fri\}$.

In each of the infinite number of discrete time periods $t \in 1, 2, \ldots$ a single contract is offered in a first-price sealed-bid auction.[4] Contract in period t has observable characteristics z_t, like type of work, size, or duration, which are assumed to be i.i.d. draws from distribution function of project characteristics F_z. Firms discount future with a common discount factor $\beta \in (0, 1)$, which is fixed over time.

[4] Following previous literature (Jofre-Bonet and Pesendorfer (2003),[1] Balat (2013),[2] and Groeger (2013)[3]), I treat each auction as independent even though multiple projects are sometimes auctioned off on the same day.

2.2.2 Stage Game

The stage game can be split into two stages: in the first stage, an exogenous subset of firms currently in the market that have requested the project documents (so called *planholders*) decide whether or not to participate in the auction; this step gives rise to the set of *bidders*. In the second stage, bidders decide on the dollar value of their bids.

Participation stage

1. In period t, an auction for a single contract with observable characteristics $z_t \sim F_z(.)$, such as type of work, size, or duration of the project is announced. State vector s_t, measuring the stock of past experience of all new bidders, is also publicly observable and contains the level of experience of all new firms.

2. An exogenous set of planholders ($\mathcal{N}_t$) receive private draws of participation costs from a common distribution, $\phi_{it} \sim F_\phi(.|z_t; \lambda)$, and decide whether or not to participate in the auction

Bidding stage

1. The outcome of the participation stage, the set of bidders ($N_t \subseteq \mathcal{N}_t$), is realized and observed by all firms.

2. Bidders receive private draws from type-specific distributions of construction costs: $c_{it} \sim F_c^{new}(.|z_t, N_t, s_{it}; \theta_c^{new})$ for new firms; and $c_{it} \sim F_c^{reg}(.|z_t; \theta_c^{reg})$ for experienced firms. Note that unlike for new firms, the cost distribution for regular firms does not depend on the state variable and hence on their past experience.

3. All bidders simultaneously submit bids (b_{it}) that maximize their expected discounted future profits. The submitted bids are ranked and the contract is awarded to the lowest bidder.

The long-run benefits of learning for new firms occur when winning contracts and accumulating experience improves their future costs. This process is captured by new firms' costs which are inter-temporally linked through the stock of their past experience, unlike the costs of regular firms.

Following previous literature, I assume there is no binding reserve price. For auctions with a single bidder, I simulate a competitor by taking a draw from the distribution of construction costs of regular firms.

2.2.3 State Variables and State Transitions

The state variable is publicly observable and defined as the stock of experience of new firms, s_t, and is measured by the number of contracts previously awarded to each new bidder. Therefore, a new bidder's experience increases by one after winning a project and therefore evolves deterministically conditional on the state variable in the previous period and the identity of the current winner. The component of the transition function for a new firm i is given by

$$s_{it} = \omega(s_{i(t-1)}, j_{win}) = \begin{cases} s_{i(t-1)} + 1 & \text{if } j_{win} = i \\ s_{i(t-1)} & \text{otherwise,} \end{cases}$$

where $s_{i(t-1)}$ is bidder i's stock of experience at the end of period $t-1$ and j_{win} is the identity of the winning bidder in the current period. While the empirical section of this paper uses the number of contracts won in the past, there are other variables that could be used to capture the level of construction experience in this market and provide additional variation, such as the number of submitted bids or the value of projects for which bids were submitted. An important issue, though not addressed in this paper, is horizontal subcontracting of firms. In any auction, firms can choose to participate in the role of a primary bidder or as a potential subcontractor for one or more of its competitors. This affects the bidding strategy as firms consider the tradeoff between higher likelihood of performing at least part of the work in the role of subcontractor versus the lower likelihood of becoming the main contractor. Since the experience accumulated as a subcontractor is very specific and limited in scope, I assume that only winning as the primary contractor can improve future costs. I investigate the effects of subcontractor experience in Chapter 3.

2.2.4 Equilibrium

Assuming conditional independence of contract characteristics and private participation and construction cost realizations allows me to focus on anonymous type-symmetric strategies, so that conditional on observables only current state and private shocks determine optimal actions; hence, I drop the time subscript t from the notation. The strategy of bidder i of type $k(i)$ consists of a participation strategy $\sigma^d_{k(i)}(\phi_i, z, \mathcal{N}, s)$ mapping private participation cost draws to the decision to submit a bid, $\sigma^d_{k(i)} : (\phi_i, z, \mathcal{N}, s) \to d_i \in \{0, 1\}$; and a bidding strategy, mapping private construction cost draws to bid amounts, $\sigma^b_{k(i)} : (c_i, z, N, s) \to b_i \in [0, \infty)$. In order to ensure invertibility, I assume that the bidding function is monotonically increasing in the construction cost. The full strategy profile of bidder i can be expressed as $\sigma_i = (\sigma^d_{k(i)}, \sigma^b_{k(i)})$, and the strategy profile of its rivals by σ_{-i}.

A Markov Perfect Equilibrium (MPE) of this game is the set of strategies σ^* such that the prescribed strategy for any bidder i given by σ^*_i is the best response to σ^*_j for all bidders j, and the equilibrium beliefs on the probability of entry are consistent with the strategies in σ^*. Furthermore, I assume that the data is generated by a single MPE strategy profile.

2.2.5 Participation and Bidding Decisions

When deciding on entry into an auction, planholders compare the net expected profits from entry against the expected profits from staying out of the auction, taking into account the impact of their decisions on the expected future profits through state transitions. In particular, firms anticipate that acquiring additional experience will lead to better costs in the future, while by not participating in an auction they certainly forgo the possibility of any additional experience. Conditional on entry, bidders maximize their expected discounted future profits, incorporating their expectations about future distributions of costs, project characteristics, the set of planholders, and states.

Assuming that other bidders follow their optimal strategy as prescribed by σ^*, the continuation value for a planholder i of type $k(i)$ on a project with characteristics z, set of rival planholders $\mathcal{N}$, in state s with a draw of participation costs ϕ_i can be expressed as

$$W_{k(i)}(z, \mathcal{N}, s, \phi_i; \sigma^*) = \max_{d \in \{0,1\}} \left\{ d(W^1_{k(i)}(s, \mathcal{N}, z; \sigma^*) - \phi_i) + (1-d) W^0_{k(i)}(s, \mathcal{N}, z; \sigma^*) \right\}, \tag{2.1}$$

where $W^1_{k(i)}$ and $W^0_{k(i)}$ are the expected value of entering and not entering the auction, respectively.

In particular, conditional on the decision to enter, the expected payoff from entry before accounting for participation costs consists of the sum of the expected payoff in the current period and the expected discounted future payoffs,

$$\begin{aligned} W^1_{k(i)}(z,\mathcal{N},s;\sigma^*) = E_{N_i}\Bigg[\int \max_b \Bigg\{ &(b-c_i)\Pr(i \text{ wins}|b,z,N,s;\sigma^*) \\ &+ \beta \sum_{j\in N} \Pr(j \text{ wins}|b,z,N,s;\sigma^*) V_i(\omega(s,j)) \Bigg\} dF_c^{k(i)}(c_i|z,N,s)\Bigg] \end{aligned} \tag{2.2}$$

where the expectation is taken over the set of possible bidder configurations which include firm i (i.e. all possible subsets of planholders that contain i). The first line in Equation (2.2) is the expected payoff in the current period and depends on the markup and the probability that bidder i wins the current auction, conditional on his bid. The second line represents the expected discounted future profit, and depends on bidder i's continuation value in the future period as well as on the transition into the future state, which in turn is a function of the probability of winning of different bidders. $W^1_{k(i)}(z,\mathcal{N},s;\sigma^*)$ is evaluated before firm i receives its draw of the construction cost, and hence must integrate this cost out over the distribution of possible costs.

On the other hand, the current period expected payoff from not entering an auction is zero, so the expected value of staying out of an auction depends only on the probability that *other* participating firms win and the resulting continuation value,

$$W^0_{k(i)}(z, \mathcal{N}, s; \sigma^*) = E_{N_{\neg i}}\left[0 + \beta \sum_{j \in N} \Pr(j \text{ wins}|z, N, s; \sigma^*) V_i(\omega(s, j))\right]. \qquad (2.3)$$

In this case, the expectation is taken over all subsets of planholders that exclude firm i, denoted as $N_{\neg i}$.

In equilibrium, bidders will decide to enter the auction if the expected value of entry net of entry costs exceeds the expected value of non-entry,

$$W^1_{k(i)}(z, \mathcal{N}, s; \sigma^*) - \phi_i \geq W^0_{k(i)}(z, \mathcal{N}, s; \sigma^*).$$

The ex-ante value function before the project characteristics, the set of planholders, and the entry cost are realized can be obtained by integrating out $z, \mathcal{N}$, and ϕ_i,

$$\begin{aligned} V_i(s; \sigma^*) = E_{z\mathcal{N}} & \left[p_{k(i)}(z, \mathcal{N}, s; \sigma^*) \left(W^1_i(z, \mathcal{N}, s; \sigma^*) - E[\phi_i | \phi_i \leq \zeta_{k(i)}(z, \mathcal{N}, s; \sigma^*)]\right) \right. \\ & \left. + \left(1 - p_{k(i)}(z, \mathcal{N}, s)\right) W^0_i(z, \mathcal{N}, s; \sigma^*)\right], \end{aligned} \qquad (2.4)$$

where $\zeta_{k(i)}(z, \mathcal{N}, s; \sigma^*)$ denotes the incremental value of entry over non-entry, and corresponds to the entry cost that makes a bidder just indifferent between entry and non-entry,

$$\zeta_{k(i)}(z, \mathcal{N}, s; \sigma^*) = W^1_{k(i)}(z, \mathcal{N}, s; \sigma^*) - W^0_{k(i)}(z, \mathcal{N}, s; \sigma^*),$$

while $p_{k(i)}(z, \mathcal{N}, s)$ is the equilibrium conditional choice probability of entry for a bidder of type $k(i)$, defined by

$$p_{k(i)}(z, \mathcal{N}, s; \sigma^*) = \int 1\{\zeta_{k(i)}(z, \mathcal{N}, s; \sigma^*) \geq \phi\} dF_\phi(\phi). \tag{2.5}$$

2.2.6 Value Function

To proceed, I first express the expected participation cost conditional on entry as a function of the equilibrium probability of participation. To simplify notation, I omit the explicit dependence of the probability of participation and the incremental value of entry on the project characteristics, set of planholders and the state variable, and let $p_{k(i)} \equiv p_{k(i)}(z, \mathcal{N}, s; \sigma^*)$ and $\zeta_{k(i)} \equiv \zeta_{k(i)}(z, \mathcal{N}, s; \sigma^*)$. Assuming that participation costs follow the exponential distribution, so $\phi_i \sim F_\phi(\phi; \lambda)$, where $F_\phi(\phi; \lambda) = 1 - e^{\lambda\phi}$, I can express the conditional expectation as a function of the incremental value of participation over non-participation and the equilibrium participation probability as[5]

$$E\left[\phi_i | \phi_i \leq \zeta_{k(i)}\right] = -\frac{\zeta_{k(i)}}{p_{k(i)}} \left[\frac{p_{k(i)}}{\log(1 - p_{k(i)})} + (1 - p_{k(i)})\right]. \tag{2.6}$$

[5]Details of the derivation are provided in Appendix 2.8.

Substituting this conditional expectation of the participation cost given in Equation (2.6) back into the value function in Equation (2.4) and using the definition of the incremental value of participation over non-participation, the value function can be simplified as

$$
\begin{aligned}
V_i(s;\sigma^*) &= E_{z\mathcal{N}}\left\{p_{k(i)}W^1_{k(i)} + (1-p_{k(i)})W^0_{k(i)} + \zeta_{k(i)}\left[\frac{p_{k(i)}}{\log(1-p_{k(i)})} + (1-p_{k(i)})\right]\right\} \\
&= E_{z\mathcal{N}}\left\{p_{k(i)}W^1_{k(i)} + (1-p_{k(i)})W^0_{k(i)} + (W^1_{k(i)} - W^0_{k(i)})\left[\frac{p_{k(i)}}{\log(1-p_{k(i)})} + (1-p_{k(i)})\right]\right\} \\
&= E_{z\mathcal{N}}\left\{p_{k(i)}W^1_{k(i)} + \left[\frac{p_{k(i)}}{\log(1-p_{k(i)})} + (1-p_{k(i)})\right]W^1_{k(i)} + (1-p_{k(i)})W^0_{k(i)}\right. \\
&\qquad \left. -W^0_{k(i)}\left[\frac{p_{k(i)}}{\log(1-p_{k(i)})} + (1-p_{k(i)})\right]\right\} \\
&= E_{z\mathcal{N}}\left\{\left[1 + \frac{p_{k(i)}}{\log(1-p_{k(i)})}\right]W^1_{k(i)} - \frac{p_{k(i)}}{\log(1-p_{k(i)})}W^0_{k(i)}\right\}. \qquad (2.7)
\end{aligned}
$$

Furthermore, substituting for the value of entry $W^1_{k(i)}$ and non-entry $W^0_{k(i)}$ from Equations (2.3) and (2.2), respectively, the value function can be expressed in recursive notation as

$$
\begin{aligned}
V_i(s;\sigma^*) = E_{z\mathcal{N}}\Bigg\{&\left[1 + \frac{p_{k(i)}}{\log(1-p_{k(i)})}\right]E_{N_i}\Bigg[\int \max_b \Bigg\{(b-c_i)\Pr(i \text{ wins}|b,z,N,s;\sigma^*) \\
&+ \beta\sum_{j\in N}\Pr(j \text{ wins}|b,z,N,s;\sigma^*)V_i(\omega(s,j);\sigma^*)\Bigg\}dF^{k(i)}_c(c_i|z,N,s)\Bigg] \\
&- \frac{p_{k(i)}}{\log(1-p_{k(i)})}E_{N_{\neg i}}\left[\beta\sum_{j\in N}\Pr(j \text{ wins}|z,N,s;\sigma^*)V_i(\omega(s,j);\sigma^*)\right]\Bigg\}. \qquad (2.8)
\end{aligned}
$$

Since the value function specifies behavior in equilibrium assuming that all other bidders behave optimally as prescribed by σ^*, the probability of a bidder winning an auction can be expressed as a function of the equilibrium bid distributions, which can be estimated from the available data. Let $G^{new}(.|z, N, s)$ denote the distribution function of equilibrium bids of a new bidder is state (z, N, s) with associated density $g^{new}(.|s, N, z)$. Similarly, let $G^{reg}(.|z, N)$ be the distribution function of equilibrium bids for a regular, experienced bidder, with associated density $g^{reg}(.|z, N)$.

Three different types of probabilities need to be specified. First is the probability that bidder i wins when he enters the auction and submits bid b, which only happens if all other participating bidders submit bids above b,

$$\Pr(i \text{ wins}|b, z, N, s; \sigma^*) = \Pr(b \leq b_j, \forall j \in \mathcal{N}|z, N, s; \sigma^*) = \prod_{j \in N, j \neq i} \left[1 - G^{k(j)}(b|z, N, s)\right]$$

Second is the probability that bidder j wins given that bidder i submits bid b, which occurs when bidder j submits a bid below b, while all remaining bidders bid above bidder j's bid,

$$\Pr(j \text{ wins}|b, z, N, s; \sigma^*) = \int_{\underline{b}}^{b} g^{k(j)}(x|z, N, s) \prod_{\substack{l \in N \\ l \neq i,j}} \left[1 - G^{k(l)}(x|z, N, s)\right] dx. \tag{2.9}$$

Last is the probability that bidder j wins when bidder i does not enter the auction, which happens when all other bidders bid above bidder j's bid,

$$\Pr(j \text{ wins}|z, N, s; \sigma^*) = \int_{\underline{b}}^{\overline{b}} g^{k(j)}(x|z, N, s) \prod_{\substack{l \in N \\ l \neq j}} \left[1 - G^{k(l)}(x|z, N, s)\right] dx. \tag{2.10}$$

New firms face construction costs that depend on their state, and therefore can improve their future cost by acquiring additional experience. Regular and fringe firms are assumed to face a stationary construction cost distribution that does not depend on their experience, and hence maximize their current period payoff since their expected future payoffs are going to be the same regardless of their current action.

In order to compute the expectation over the set of bidders, I need to specify the probability of observing a particular subset of planholders entering the auction: $Pr(N|\mathcal{N})$. This probability depends on the probability of entry of each type of planholder,

$$\Pr(N|z, \mathcal{N}, s) = \prod_{k \in \{new, reg, fri\}} C_{N^k}^{\mathcal{N}^k} p_k(z, \mathcal{N}, s)^{N^k} (1 - p_k(z, \mathcal{N}, s))^{\mathcal{N}^k - N^k}, \tag{2.11}$$

where $\mathcal{N}^k$ is the number of planholders of type k, while N^k is the number of bidders of type k in a particular configuration N. To avoid computing this combinatorial problem, I simulate different entry configurations based on estimated probabilities of entry and average across these simulations.

2.3 Identification

In this section, I provide the intuition for the identification of construction costs and participation costs. The identification arguments follow previous literature on the structural estimation of dynamic auctions and dynamic discrete choice models. The market is assumed to be in steady state and all auctions are treated as independent, which allows me to identify the conditional probability of participation, state transitions and bid distributions. The distribution of the construction costs is traced from the empirical bid distribution using the inverse bid function that characterizes optimal behavior, and relies on the monotonicity assumption. On the other hand, the parameters of the participation cost distribution are over-identified and depend on the specified functional form.

2.3.1 Construction Costs

The first order condition (FOC) for optimal bidding behavior of new firms implies that the privately known project construction cost c_i for a bidder who submitted bid b is given by[6]

[6]The derivation of this equation is provided in the appendix in Jofre-Bonet and Pesendorfer (2003).[1]

$$c_i = b - \frac{1}{\sum_{j \neq i} h^{k(j)}(b|z, N, s_j, s_{-j})} + \beta \sum_{j \neq i} \frac{h^{k(j)}(b|z, N, s_j, s_{-j})}{\sum_{l \neq i} h^{k(l)}(b|z, N, s_l, s_{-l})} \Big[V_i(\omega(s,i)) - V_i(\omega(s,j)) \Big], \tag{2.12}$$

where $h^{k(i)}(.|z, N, s_i, s_{-i}) = \frac{g^{k(i)}(.|z,N,s_i,s_{-i})}{1-G^{k(i)}(.|z,N,s_i,s_{-i})}$ is the hazard function of bids submitted by bidder i of type $k(i)$. The FOC relates the construction cost c_i to the bid, the distribution and density functions of equilibrium bids, and the value function. The second term on the right hand side corresponds to the markup accounting for the competition in the current period. However, unlike in a static auction, there is now an additional term accounting for the incremental effect on future discounted profits that arise when firm i wins the contract instead of another firm.

In order to recover the private construction costs, all objects on the right-hand side of Equation (2.12) must be obtained. These include the bid hazard function, the discount factor, the state transition function and the value function. The hazard function depends on the distribution and density functions of equilibrium bids and can be directly estimated from data on bids. The discount factor is not identified from the data and is fixed.[7] The state transition function is a deterministic function described in the model setup.[8] The remaining object is the value function, which is given in Equation (2.8). However, this expression cannot be directly taken to data as

[7] In estimation, I fix β=0.995. If a firm obtains planholder materials for one project per month, this would correspond to an annual discount factor of 0.94.

[8] However, in practice I aggregate levels of experience into a smaller number of bins. Transition between different levels of experience is then estimated by a frequency estimator.

it includes integration over the unobserved distribution of costs.

Following Jofre-Bonet and Pesendorfer (2003)[1] and Balat (2013)[2] and substituting the FOC from Equation (2.12) into the value function in Equation (2.8), the value function can be expressed as a function of the equilibrium bids only,

$$
\begin{aligned}
V_i(s) = E_{z\mathcal{N}} \Bigg\{ & \left[1 + \frac{p_{k(i)}}{\log(1-p_{k(i)})}\right] E_{N_i} \Bigg[\int \frac{1}{\sum_{j\in N, j\neq i} h^{k(j)}(b|z,N,s_j,s_{-j})} dG^{(i)}(b|z,N,s) \\
& + \beta \sum_{j\in N, j\neq i} \left(\Pr(j \text{ wins}|z,N,s) + \int \frac{h^{k(i)}(b|z,N,s_i,s_{-i})}{\sum_{l\in N} h^{k(l)}(b|z,N,s_l,s_{-l})} dG^{(j)}(b|z,N,s) \right) V_i(\omega(s,j)) \Bigg] \\
& - \frac{p_{k(i)}}{\log(1-p_{k(i)})} E_{N_{\neg i}} \left[\beta \sum_{j\in N} \Pr(j \text{ wins}|z,N,s) V_i(\omega(s,j)) \right] \Bigg\}, \qquad (2.13)
\end{aligned}
$$

where $G^{(j)}(b|z,N,s)$ corresponds to the ex-ante probability that bidder j wins with a bid of b or less. The derivation is based on two observations: First, the probability of winning can be expressed as a function of the distribution of bids by other bidders, ignoring the dependance of other bidders' bids on their own private cost draws, thus reducing the dynamic game to a single agent dynamic decision problem, in which the bidder i takes as given the equilibrium bid distribution of the other bidders. The second issue arises from the presence of the private cost, which is not observed. This can be solved by substituting the FOC into the value function. The derivation is explained in detail in Appendix 2.8.

For a new bidder i at any state s from the set of states $(s^1, s^2, \ldots, s^m)$, let the components of the value function be defined as

$$
\begin{aligned}
A_i(s) &= E_{z\mathcal{N}}\left\{\left(1+\frac{p_{k(i)}}{\log(1-p_{k(i)})}\right)E_{N_i}\left[\int\frac{1}{\sum_{j\in N,j\neq i}h^{k(j)}}dG^{(i)}(b|z,N,s)\right]\right\},\\
B_i(s) &= E_{z\mathcal{N}}\left\{\left(1+\frac{p_{k(i)}}{\log(1-p_{k(i)})}\right)E_{N_i}\left[\sum_{j\in N,j\neq i}\int\left(1+\frac{h^{k(i)}(b|z,N)}{\sum_{l\in N}h^{k(l)}(b|z,N,s)}\right)dG^{(j)}(b|z,N,s)\right.\right.\\
&\qquad\left.\left.\times\left(1\{\omega(s,j)=s^1\},\ldots,1\{\omega(s,j)=s^m\}\right)\right]\right\},\\
D_i(s) &= E_{z\mathcal{N}}\left\{-\frac{p_{k(i)}}{\log(1-p_{k(i)})}E_{N_{\neg i}}\left[\sum_{j\in N,j\neq i}\int dG^{(j)}(b|z,N,s)\right.\right.\\
&\qquad\left.\left.\times\left(1\{\omega(s,j)=s^1\},\ldots,1\{\omega(s,j)=s^m\}\right)\right]\right\},
\end{aligned}
$$

where A_i is the expected current period profit net of entry costs and accounting for the probability of entry, while B_i and D_i are the $1\times m$ vectors of transition probabilities for entry and non-entry, respectively. These components are specified from the perspective of bidder i before making his own entry decision. For example, the kth element of $B_i(s)$ contains the probability of reaching state s^k from the state s when bidder i decides to enter, incorporating his probability of entry.

Aggregating over all states allows me to write the value function in matrix notation. Let V_i denote the vector $(V_i(s^1),\ldots,V_i(s^m))'$, $A_i = A_i(s^1),\ldots,A_i(s^m))'$, $B_i = B_i(s^1),\ldots,B_i(s^m))'$ and $D_i = D_i(s^1),\ldots,D_i(s^m))'$. Then

$$
\begin{aligned}
V_i &= A_i + \beta B_i V_i + \beta D_i V_i\\
&= A_i + \beta\left(B_i + D_i\right)V_i,
\end{aligned}
$$

with the solution given by

$$V_i = \left[\sum_{k=0}^{\infty} \beta^k (B_i + D_i)^k \right] A_i = \left[I - \beta(B_i + D_i) \right]^{-1} A_i. \tag{2.14}$$

With value function estimated using Equation (2.14), all objects on the right-hand side of the FOC in Equation (2.12) can be estimated from data, and the parameters of the construction cost distribution can be recovered. The FOC together with the monotonicity assumption imply that the inverse bid function conditional on state variables is defined as $c = (\sigma^b)^{-1}(b|z, N, s)$. Furthermore, because of the relationship between the distributions of costs and the associated bids,

$$F_c^k(c|z, N, s) = G^{k(i)}(\sigma_k^b(c, z, N, s)|z, N, s).$$

Once the value function is estimated, $F_c^k(.|z, N, s)$ can be recovered by integrating the estimated bid distribution functions $\hat{G}^k(.|z, N, s)$ and using the inverse bid function derived from the FOC,

$$F_c^k(c|z, N, s) = \int_{\{b|(\sigma^b)^{-1}(b|z,N,s) \leq c\}} d\hat{G}^k(.|z, N, s) \tag{2.15}$$

2.3.2 Participation Costs

The parameter of the participation cost distribution can be identified by taking advantage of the functional form assumption on the participation cost distribution and employing the fact that in equilibrium beliefs are going to be correct.

Entry costs are assumed to be distributed exponentially with parameter λ, with distribution $F_\phi(\phi_{it}|z_t) = 1 - \mathrm{e}^{-\lambda(z_t)\phi}$ and with the associated probability density function $f_\phi(\phi_{it}|z_t) = \lambda(z_t)\mathrm{e}^{-\lambda(z_t)\phi}$, where $\lambda(z_t) = \lambda_0 + \lambda_1 z_t$. The conditional equilibrium probability of entry can be expressed as a function of the incremental value of entry over non-entry, and corresponds to the participation cost that makes a firm just indifferent between entry and non-entry,

$$p_k(z, \mathcal{N}, s; \lambda) = F_\phi(\zeta(z, \mathcal{N}, s); \lambda) = F_\phi(W_k^1(z, \mathcal{N}, s) - W_k^0(z, \mathcal{N}, s); \lambda). \tag{2.16}$$

2.4 Estimation

I use the data from Caltrans described in Section 2.1 to estimate the parameters of construction and participation cost distributions. Construction costs are recovered using the inverse bid function that characterizes the optimal bidding behavior. This requires estimates of the bid distribution and value function, which in turn depends on the estimates of the probability of entry and state transitions. Bid distribution together with the conditional entry probabilities and state transitions are obtained in the first stage of the estimation, which then allow me to calculate the value function at all states. Using the inverse bid function I then trace out the underlying distribution of construction costs and estimate its moments. In the final step, I use a moment estimator to recover the parameters of the participation cost distribution.

To make estimation feasible, the level of experience for new bidders is aggregated into Q non-overlapping categories. In practice, I split experience into 5 categories: category 1 (no contracts awarded in the past), category 2 (one contract awarded in the past), category 3 (2-3 contracts awarded in the past), category 4 (4-10 contracts won in the past), and category 5 (more than 10 contracts won in the past).

I use the engineer's estimate and a measure of overall backlog as observable project characteristics. Engineer's estimate is categorized into five quintiles, while the backlog measure is an indicator based on the proportion of regular bidders with above-average backlog. This variable measures how constrained overall are the firms in the market.

The goal of the first step of the estimation is to flexibly estimate the probability of entry, state transition, and conditional bid functions.

2.4.1 Conditional Probability of Entry and State Transitions

The conditional probability of entry ($\hat{p}_k(z, \mathcal{N}, s)$) is estimated using the following probit specification separately for each bidder type,

$$Pr(d_{it} = 1|z, \mathcal{N}, s) = \Phi\Big(\beta_0 + \beta_1 \text{EngEst} + \beta_2 \text{Backlog} + \beta_3 \#\text{Regular} + \sum_{q=1}^{Q} \beta_{3+q} \#\text{New(q)} + \beta_9 \#\text{Fringe} + \beta_{10} \text{OwnExp} * \mathbf{1}\{k(i) = new\}\Big)$$

On the other hand, state transitions are estimated using a frequency estimator. While the transition process is deterministic in the sense that the experience of a firm increases by one every time this firms wins an auction, in practice the level of experience is aggregated into five levels. Transition between different levels then depends on the probability of a firm having the highest level of experience within a level, and thus progressing to the next level. The probability that a bidder with the current level of experience $\bar{s}$ will move to the next level of experience $\bar{\bar{s}}$ is estimated using a frequency estimator,

$$Pr(\hat{\omega}(\bar{s}, \text{winner} = i) = \bar{\bar{s}}, \tilde{z}, \tilde{\mathcal{N}}) = \frac{\sum_i \sum_t \mathbf{1}\{s_{it} = \max(\bar{s}), z_t = \tilde{z}, \mathcal{N}_t = \tilde{\mathcal{N}}\}}{\sum_i \sum_t \mathbf{1}\{s_{it} = \bar{s}, z_t = \tilde{z}, \mathcal{N}_t = \tilde{\mathcal{N}}\}},$$

where

2.4.2 Conditional Bid Distribution

Following Jofre-Bonet and Pesendorfer (2003)[1] and Balat (2013),[2] I estimate the parameters of the conditional bid distribution ($\hat{\theta}$) using maximum likelihood. I assume that the scaled bids follow Weibull distribution, so the corresponding density function for a bidder of type $k \in \{new, reg\}$ is

$$g^k(b|z, s, N, \theta^k) = \left[\frac{\theta_1^k}{\theta_2^k} \left(\frac{b}{\theta_2^k} \right)^{\theta_1^k - 1} e^{-(\frac{b}{\theta_2^k})^{\theta_1^k}} \right],$$

and the parameters of the distribution are defined as

$$\begin{aligned}
\log(\theta_1^k) &= \gamma_{1,0}^k \\
\log(\theta_2^k) &= \gamma_{2,0}^k + \gamma_{2,1}^k \text{EngEst} + \gamma_{2,3}^k \text{Backlog} + \gamma_{2,4}^k \#\text{Regular} \\
&\quad + \sum_{q=1}^{Q} \gamma_{2,4+q}^k \#\text{New(q)} + \gamma_{2,10}^k \#\text{Fringe} + \gamma_{2,11}^k \text{OwnExp} * \mathbf{1}\{k(i) = new\}.
\end{aligned}$$

Parameters are estimated using maximum likelihood, with the likelihood specified as

$$L(\theta) = \prod_t \prod_{k \in \{new, reg\}} \prod_{i \in N^k} g^{k(i)}(b_{it}|z, N, s; \theta^k)$$

2.4.3 Value Function

With the estimates of the bid distribution, probability of entry and state transitions in hand, I use Equation (2.14) to compute the value function on the vector of states. Given the parameters of the bid distribution $\hat{\theta}$, I can compute the distribution and density functions of the equilibrium bids, $\hat{G}^k(b|z,s,N,\hat{\theta})$ and $\hat{g}^k(b|z,s,N,\hat{\theta})$, and hence also the hazard function $\hat{h}^k(b|z,s,N,\hat{\theta})$ for any particular combination of $(\tilde{z},\tilde{s},\tilde{N})$. Together with the estimated probability of entry $\hat{p}_{\tilde{k}}(\tilde{z},\tilde{s},\tilde{\mathcal{N}})$ this allows me to calculate the vector of current period expected profits $A_i(s)$ and the matrices of transition probabilities $B_i(s)$ and $D_i(s)$ conditional on entry and non-entry, respectively.

The expectation over the project characteristics is approximated by sampling from the observed set of projects. Similarly, to compute the expectation over the set of bidders, I consider a given set of planholders and simulate entry of individual planholders based on the estimated conditional probability of entry.

2.4.4 Construction and Participation Costs

After obtaining the value function, I can directly apply the inverse bid function defined by the FOC in Equation (2.12) to recover the cost for any bid amount, given the combination of $(\tilde{z},\tilde{s},\tilde{\mathcal{N}})$. To recover the full distribution of construction costs, I use Equation (2.15) to integrate over the estimated bid distribution.

On the other land, λ can be recovered by employing the fact that in equilibrium beliefs are going to be correct. I use a moment estimator that minimizes the square of the distance between the probability of entry implied by the model $p(\lambda)$ and its empirical counterpart $\hat{p}$ obtained from data,

$$\hat{\lambda} = \arg\min_{\lambda}(\hat{p} - p(\lambda))'W(\hat{p} - p(\lambda))$$

2.5 Results

The estimation results from the structural model are organized in the following way. I start with a discussion of the first-stage estimates of entry probabilities and bid distribution parameters, and proceed by presenting the recovered parameters of the construction and entry costs. The estimated construction costs for new bidders demonstrate the dependence on past experience, and converge to the distribution of costs of regular bidders for the maximum level of experience. Results indicate the presence of bidder learning in this market and show that recent entrants benefit from acquiring additional experience. The estimated average cost of the most experienced among new bidders are substantially lower than the average cost of the least experienced firms bidding on the same project, and this difference is especially pronounced for large projects.

2.5.1 Probability of Entry and Bid Distribution

The parameters of the conditional probability of entry for each type of firm are presented in Table 2.5. Larger projects tend to discourage entry, especially among small and fringe bidders. Average backlog, which measures how capacity-constrained are the regular planholders, does not seem to affect entry into auctions in a significant way. It is likely that capacity-constrained firms, which would be most affected, are not in the set of planholders and and choose to stay out of the auctions due to their high backlog and inability to carry out construction work. On the other hand, the competition effects are important. The presence of a larger number of regular firms in an auction has a significant negative effect on the probability of entry for all types of bidders. For new firms, additional competitors of any type decrease probability of entry. Finally, the effect of experience (measured as the number of awarded projects in the past) has the expected positive sign and is significant, meaning that more experienced bidders participate in auctions relatively more often.

Table 2.6 presents the estimated parameters of the bid density function for each type of bidder obtained using maximum likelihood and assuming Weibull distribution. Since the dependent variable is the bid amount scaled by the engineer's estimate, the negative effect of engineer's estimate indicates that bidders compete more aggressively for larger projects. Similarly, a larger number of regular, fringe, or relatively experienced new bidders intensifies competition. Furthermore, the effect of own experience for new bidders is negative, implying that higher level of experience allows

Table 2.5: Probability of Entry

	New	Regular	Fringe
Constant	0.2655***	0.6711***	−0.4188***
Engineer's Estimate	−0.1248***	−0.0475***	−0.2254***
Backlog (Above/Below Average)	0.0005	0.0056	0.0027
# Regular Competitors	−0.0396***	−0.0669***	−0.0413***
# Fringe Competitors	−0.0029*	−0.0294***	0.0209***
# New Competitors (Exp1)	−0.0151***	−0.0221**	−0.0279***
# New Competitors (Exp2)	−0.0102**	−0.0068	0.0063
# New Competitors (Exp3)	−0.0070*	−0.0136	0.0192**
# New Competitors (Exp4)	−0.0135***	−0.0182***	−0.0219***
# New Competitors (Exp5)	−0.0033	0.0601***	−0.0167**
# Past wins	0.0241***		

Note: Significance level at 1%(***), 5%(**), 10%(*)

these bidders to submit lower bids. However, without recovering the underlying cost structure it is not possible to know whether bids decline because of lower markups or changes in the construction costs.

The estimated bid distribution and density functions are presented in Figure 4.1. The bid distribution functions for less experienced bidders generally stochastically dominate the ones corresponding to more experienced new bidders and regular bidders.

2.5.2 Construction and Participation Costs

Using the estimated bid distribution and applying the first-order condition, I can recover the distribution of the corresponding construction costs. Since experience of new bidders is likely going to be of special importance in auctions for the largest projects, I focus my attention to projects in the highest quintile of the engineer's

Figure 2.1: Estimated Bid Distribution and Density Functions

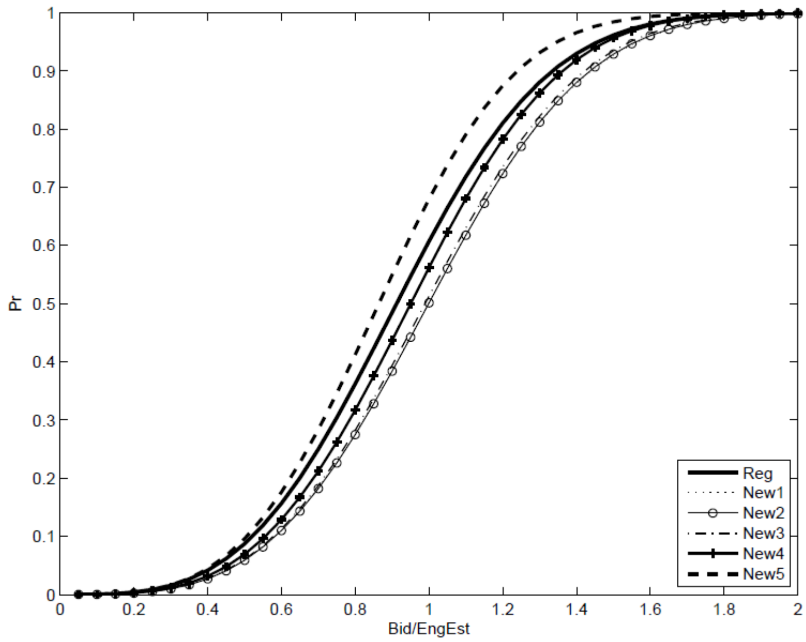

(a) Bid Distribution

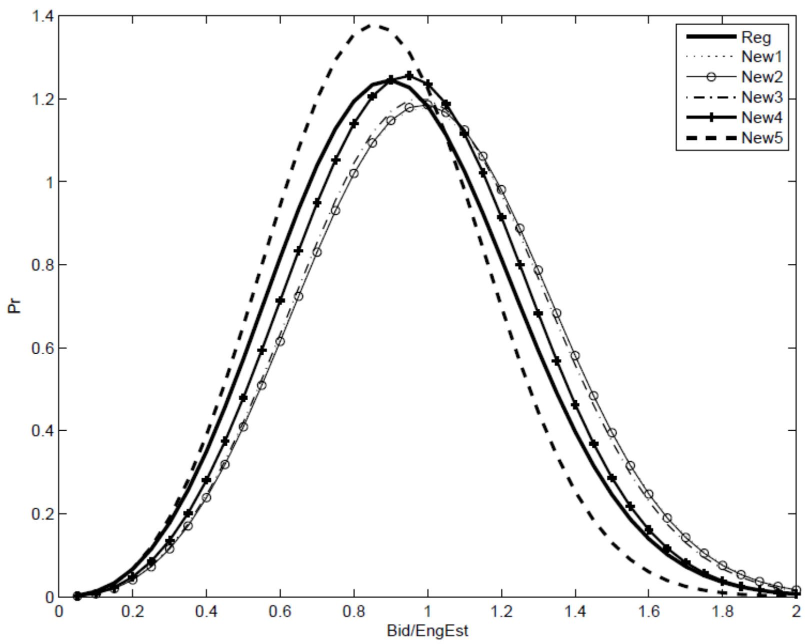

(b) Bid Density

Table 2.6: Bid Density

	New	Regular
Constant (shape)	1.285	1.245
Constant (scale)	0.315	0.240
Engineer's Estimate	−0.022	−0.017
Backlog (Above/Below Average)	0.011	−0.008
# Regular Competitors	−0.038	−0.047
# Fringe Competitors	−0.032	−0.044
# New Competitors (Exp1)	0.009	0.012
# New Competitors (Exp2)	−0.019	−0.009
# New Competitors (Exp3)	0.000	0.013
# New Competitors (Exp4)	−0.026	−0.019
# New Competitors (Exp5)	−0.034	−0.057
# Past wins	−0.007	

Note: Estimated using maximum likelihood with Weibull distribution. Dependent variable is submitted bid scaled by the corresponding engineer's estimate.

estimate and with constrained bidders.

The estimated distribution and density functions of the recovered construction costs for a large project with constrained bidders are presented in Figure 4.2. Again, the cumulative distribution functions for less experienced bidders generally stochastically dominate the ones for experienced and regular bidders. With more experience, the density of costs for new bidders shifts to the left, meaning that new bidders with higher levels of experience face more advantageous cost distribution.

In order to quantify bidder learning that occurs as firms gain experience, Table 2.7 reports the estimated average cost by project size and backlog status. Overall, the average cost declines for larger projects, which confirms the results detected from estimated bid distributions that bidders are able to compete more aggressively for

Figure 2.2: Estimated Cost (Large Project with Constrained Bidders)

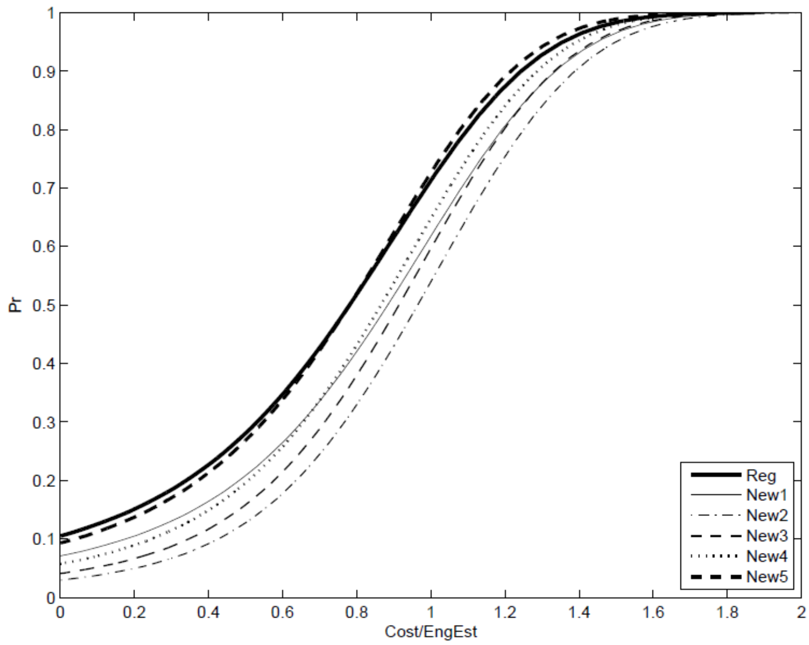

(a) Cost Distribution

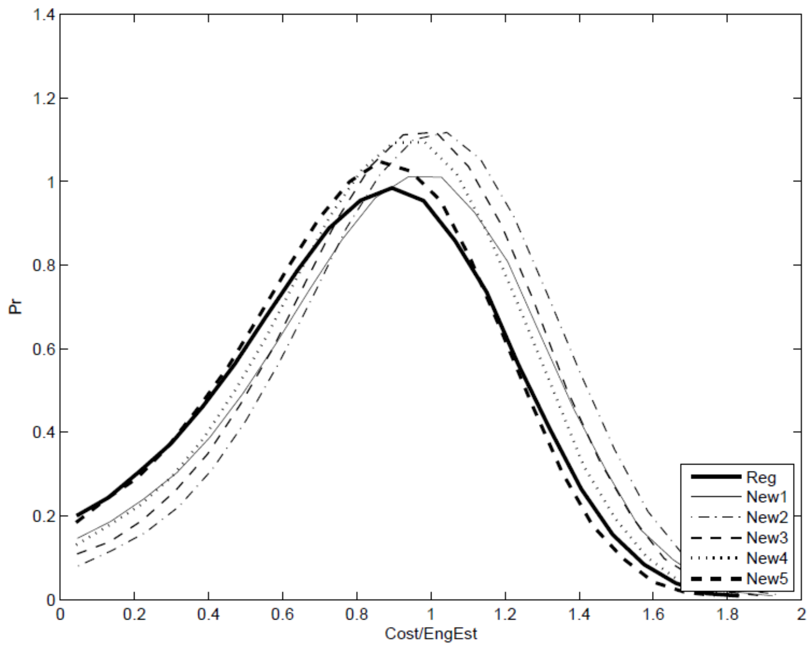

(b) Cost Density

larger projects. For any particular project size and backlog status, the average cost for new bidders generally declines with experience, with the exception of going from the first to the second level of experience. Level 1 corresponds to no experience, so consists of bidders before winning their first auction.

Comparing average cost across bidders with different levels of experience, I find that the most experienced bidders have average cost about 10% below the average cost of the least experienced bidders. For small project, the effect of experience is more important for projects with less constrained bidders, while for large projects the effect is greater for projects with constrained bidders.

Table 2.7: Average Cost by Project Size and Backlog Status

Eng Est (quintile)	Backlog (0/1)	Reg	New1	New2	New3	New4	New5
1	0	0.763	0.967	1.063	1.024	0.957	0.811
	1	0.828	0.952	0.971	0.986	0.915	0.848
2	0	0.752	0.931	1.014	0.944	0.890	0.785
	1	0.816	0.939	1.017	0.964	0.900	0.808
3	0	0.770	0.620	0.989	0.906	0.813	0.754
	1	0.754	0.923	0.975	0.932	0.850	0.776
4	0	0.730	0.792	0.931	0.751	0.821	0.737
	1	0.743	0.849	0.979	0.915	0.833	0.760
5	0	0.719	0.817	0.910	0.882	0.802	0.725
	1	0.732	0.834	0.932	0.876	0.822	0.731

Note: Projects are categorized according to the corresponding quintile of the engineer's estimate and a backlog indicator. Backlog indicator measure the proportion of regular planholders with above-average backlogs.

Furthermore I decompose bids into the underlying costs and markups. Figure 2.3 shows the inverse bid function by bidder type for a large project with constrained

Figure 2.3: Inverse Bid Function for a Large Project with Constrained Bidders

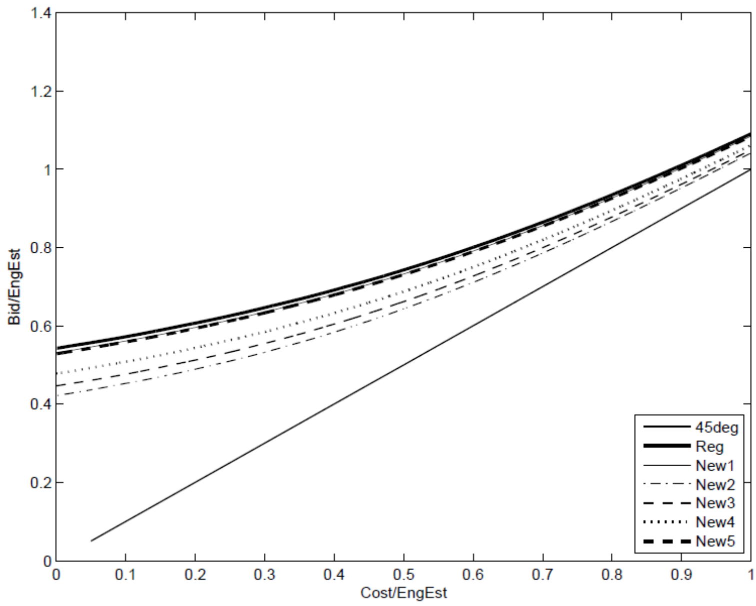

bidders. It shows that markups decline monotonically with cost. Less experienced bidders have lower markups (with the exception of new bidders without any experience), suggesting that firms are willing to forgo some profits today in order to increase their probability of winning and thus acquiring additional experience that would allow them to face more advantageous costs in the future. As new firms gain experience, their markups, and hence their inverse bid function, approach the inverse bid function of regular bidders.

Finally, the parameter of the exponential participation distribution is estimated to be 0.754.

2.6 Simulation Results

I use the parameter estimates reported in the previous section in order to perform a number of counterfactual experiments evaluating the importance of bidder learning in the procurement auction market. In the first case, I quantify the effect of increasing experience of a new bidder from the minimum to the maximum possible level and I decompose this effect into the corresponding change in markup and the change in cost. While both bids and costs decrease with more experience, the change in average cost is relatively greater, allowing more experienced bidders to charge higher markups. In the second case, I consider the effect of an exogenous change in the experience of all new bidders. In particular, I look at the implications of making all new bidders have the highest level of experience on the auction outcome. I find that average cost of procurement declines as bidders become more competitive and markups drop. Both of these scenarios confirm the importance of bidder learning in the long run.

Table 2.8: Effect of Increasing Level of Experience

	1	2	3	4	5
Freq of entry	0.294	0.294	0.315	0.348	0.438
Freq of winning	0.201	0.141	0.186	0.216	0.252
Bid	0.949	0.988	0.966	0.925	0.860
Cost	0.825	0.898	0.857	0.808	0.735
Markup	0.124	0.090	0.109	0.116	0.125
Procurement cost (lowest bid)	0.661	0.662	0.663	0.656	0.639
Winners cost	0.323	0.333	0.338	0.337	0.336
Winners markup	0.338	0.329	0.324	0.319	0.303
Inefficiency	0.012	0.012	0.012	0.011	0.009

Note: Compares outcomes when a new bidder's experience is increased from the minimum level to the maximum.

2.6.1 Effect of Bidder Learning on Cost and Markup

Bidder learning affects not only the frequency of entry and winning of new firms but also the procurement cost and bidders' profitability. Learning experienced by new firms as they enter the market is exhibited through two channels that determine the competitiveness of their bids in an auction. First, learning affects the cost distribution, with more experienced firms having more advantageous cost distributions. Second, the level of bid is determined by the firm's markup, with less experienced firms charging lower markups in order to acquire additional experience, and more experienced, stronger firms being able to charge higher markups.

The importance of the effect of learning can be evaluated by comparing the outcomes from a simulated set of auctions in which the level of experience for a bidder is exogenously increased. Table 2.8 presents the results of a simulation in which the level of experience for a new bidder is increased from the lowest to the highest level. With the exception of the initial level of no experience, the average bid declines with additional experience by 13%. Note, however, that this change is a result of a 18% drop in the average cost, and an almost 40% increase in the average markup. With more experience, bidders get stronger and better able to compete against their competitors. After initially forgoing profits by lowering markups and gaining experience, the resulting lower cost allows them to increase the markup. Also note that the increased competition benefits the auctioneer by a 3.5% decrease in the average procurement cost.

2.6.2 Effect of a Change in Experience of All Bidders

Table 2.9 compares the outcome of a simulation in which the experience of all new bidders is exogenously increased to the maximum level with the outcome under the actual observed distribution of experience among new bidders. Since improving competitiveness of any bidder affects not only the bidder himself but also the response, in term of both participation and bidding strategy, of all his competitors, it is not surprising to see that the average procurement cost declines by 7%. Since bidders are now competing against stronger opponents, the winner's average markup decreases as well.

Table 2.9: Effect of Overall Increase in Experience

	Actual	Maximum Experience
Procurement cost (lowest bid)	0.702	0.651
Winners cost	0.337	0.344
Winners markup	0.365	0.307
Inefficiency	0.015	0.087

Note: Compares outcomes under actual observed distribution of new bidders and a hypothetical in which all new bidders have maximum possible level of experience.

2.7 Conclusion

This paper considers the long-run effects of bidder learning by examining the changes in the cost distribution for bidders with different levels of experience. Using data from highway procurement auctions in California I demonstrate that recent entrants in the market experience substantial cost improvements that arise with winning additional contracts. This learning effect has important implications for the evaluation of bid preference programs from the perspective of economic efficiency. The results presented in this paper show that in presence of learning-by-doing it is also in the auctioneer's interest to encourage bidders to promptly acquire experience and become more competitive. This novel insight for the economic justification of bid preference programs also implies that any evaluation of the effect of bid preference programs on the cost of procurement in a static setting leaves out the important long-run consequences of learning and underestimtes the benefits of such programs.

2.8 Appendix: Derivations and Proofs

2.8.1 Conditional Expectation of Entry Costs

Let X be a truncated exponential distribution with parameter $\lambda > 0$, $0 < x \leq a$. Given the pdf and cdf of exponential distribution $f(x) = \lambda \mathrm{e}^{-\lambda x}$ and $F(x) = 1 - \mathrm{e}^{-\lambda x}$, respectively, we can derive the conditional expectation as follows:

$$
\begin{aligned}
E[X|0 < X \leq a] &= \int_0^a x f(x|0 < x \leq a) dx = \int_0^a x \frac{f(x)}{F(a) - F(0)} dx = \int_0^a x \frac{\lambda \mathrm{e}^{-\lambda x}}{1 - \mathrm{e}^{-\lambda a}} dx \\
&= \frac{\lambda}{1 - \mathrm{e}^{-\lambda a}} \int_0^a x \mathrm{e}^{-\lambda x} dx = \frac{\lambda}{1 - \mathrm{e}^{-\lambda a}} \left[\frac{-a}{\lambda} \mathrm{e}^{-\lambda a} + \frac{1}{\lambda} \frac{1}{-\lambda} \mathrm{e}^{-\lambda x} \Big|_0^a \right] \\
&= \frac{1}{1 - \mathrm{e}^{-\lambda a}} \left[-a \mathrm{e}^{-\lambda a} + \frac{1}{\lambda} (1 - \mathrm{e}^{-\lambda a}) \right] = \frac{1}{\lambda} - \frac{a \mathrm{e}^{-\lambda a}}{1 - \mathrm{e}^{-\lambda a}}
\end{aligned}
$$

Now letting $a = \frac{1}{\lambda} k$ for some $k > 0$, we can further rewrite this expression as

$$
E[X|0 < X \leq a] = \frac{a}{k} - \frac{a \mathrm{e}^{-k}}{1 - \mathrm{e}^{-k}} = a \left[\frac{1}{k} - \frac{\mathrm{e}^{-k}}{1 - \mathrm{e}^{-k}} \right]
$$

Next, let $p = F(a) = 1 - \mathrm{e}^{-\lambda a} = 1 - \mathrm{e}^{-k}$, which implies that $k = -\log(1 - p)$. Substituting for k in the last equation and manipulating the result,

$$
E[X|0 < X \leq a] = a \left[\frac{-1}{\log(1 - p)} - \frac{1 - p}{p} \right] = -a \left[\frac{1}{\log(1 - p)} + \frac{1 - p}{p} \right].
$$

Finally, multiplying by p, we obtain the desired result,

$$pE[X|0 < X \leq a] = -a\left[\frac{p}{\log(1-p)} + (1-p)\right]. \tag{2.17}$$

2.8.2 Value Function as a Function of Equilibrium Bids

Following is the derivations of Equation (2.13). To simplify, let $\mathcal{P} = 1 + \frac{p_{k(i)}}{\log(1-p_{k(i)})}$ and $1 - \mathcal{P} = -\frac{p_{k(i)}}{\log(1-p_{k(i)})}$. Then,

$$V_i(s) = E_{z\mathcal{N}}\left\{\mathcal{P}E_{N_i}\left[\int \max_b\left\{(b-c_i)\Pr(i \text{ wins}|b,z,N,s;\sigma^*) + \beta\sum_{j\in N}\Pr(j \text{ wins}|b,z,N,s;\sigma^*)V_i(\omega(s,j))\right\}dF_c^{k(i)}(c_i|z,N,s)\right]\right.$$
$$\left. + (1-\mathcal{P})E_{N_{\neg i}}\left[\beta\sum_{j\in N}\Pr(j \text{ wins}|z,N,s;\sigma^*)V_i(\omega(s,j))\right]\right\}.$$

Use the relationship between cost and bid from the FOC to substitute for $(b - c_i)$:

$$V_i(s) = E_{z\mathcal{N}}\left\{\mathcal{P}E_{N_i}\left[\int\left\{\frac{1-\beta\sum_{j\neq i}h^{k(j)}(b|z,N,s)[V_i(\omega(s,i)) - V_i(\omega(s,j))]}{\sum_{j\neq i}h^{k(j)}(b|z,N,s)} \times \Pr(i \text{ wins}|b,z,N,s;\sigma^*)\right.\right.\right.$$
$$\left.\left.\left. + \ \beta\sum_{j\in N}\Pr(j \text{ wins}|b,z,N,s;\sigma^*)V_i(\omega(s,j))\right\}dF_c^{k(i)}(c_i|z,N,s)\right] + (1-\mathcal{P})E_{N_{\neg i}}\left[\beta\sum_{j\in N}\Pr(j \text{ wins}|z,N,s;\sigma^*)V_i(\omega(s,j))\right]\right\}.$$

Next, note that the term for bidder i from the first sum on the second line, $\beta\Pr(i \text{ wins}|b,z,N,s;\sigma^*)V_i(\omega(s,i))$, cancels out with the same term on the first line, resulting in

$$
\begin{aligned}
V_i(s) = E_{z\mathcal{N}} \Bigg\{ \mathcal{P} E_{N_i} \Bigg[\int \Bigg\{ \frac{1+\beta \sum_{j\neq i} h^{k(j)}(b|z,N,s) V_i(\omega(s,j))}{\sum_{j\neq i} h^{k(j)}(b|z,N,s)} \times \Pr(i \text{ wins}|b,z,N,s;\sigma^*) \Bigg\} dF_c^{k(i)}(c_i|z,N,s) \\
+ \beta \sum_{j\in N, j\neq i} \Pr(j \text{ wins}|z,N,s;\sigma^*) V_i(\omega(s,j)) \Bigg] + (1-\mathcal{P}) E_{N_{\neg i}} \Bigg[\beta \sum_{j \in N} \Pr(j \text{ wins}|z,N,s;\sigma^*) V_i(\omega(s,j)) \Bigg] \Bigg\} \\
= E_{z\mathcal{N}} \Bigg\{ \mathcal{P} E_{N_i} \Bigg[\int \Bigg\{ \frac{\Pr(i \text{ wins}|b,z,N,s;\sigma^*)}{\sum_{j\neq i} h^{k(j)}(b|z,N,s)} + \beta \frac{\sum_{j\neq i} h^{k(j)}(b|z,N,s) V_i(\omega(s,j))}{\sum_{j\neq i} h^{k(j)}(b|z,N,s)} \times \Pr(i \text{ wins}|b,z,N,s;\sigma^*) \Bigg\} dF_c^{k(i)}(c_i|z,N,s) \\
+ \beta \sum_{j\in N, j\neq i} \Pr(j \text{ wins}|z,N,s;\sigma^*) V_i(\omega(s,j)) \Bigg] + (1-\mathcal{P}) E_{N_{\neg i}} \Bigg[\beta \sum_{j \in N} \Pr(j \text{ wins}|z,N,s;\sigma^*) V_i(\omega(s,j)) \Bigg] \Bigg\}
\end{aligned}
$$

Changing the variable of integration from costs to bids using the facts that $db = \frac{\partial \sigma^b(c)}{\partial c} dc$ and $f((\sigma^b)^{-1}(b)|s) \cdot \frac{\partial (\sigma^b)^{-1}(b)}{\partial b} = g(b|s)$, where $\frac{\partial (\sigma^b)^{-1}(b)}{\partial b} = \frac{1}{\frac{\partial \sigma(c)}{\partial c}}$, and expressing the probability of winning in terms of equilibrium bid distributions,

$$
\begin{aligned}
V_i(s) = E_{z\mathcal{N}} \Bigg\{ \mathcal{P} E_{N_i} \Bigg[\int \frac{\prod_{j\neq i} [1 - G^{k(j)}(b|z,N,s)]}{\sum_{j\neq i} h^{k(j)}(b|z,N,s)} g^{k(i)}(b|z,N,s) db \\
+ \beta \int \Bigg[\sum_{j\neq i} \frac{h^{k(j)}(b|z,N,s) V_i(\omega(s,j))}{\sum_{l\neq i} h^{k(l)}(b|z,N,s)} \Bigg] \prod_{j\neq i} [1 - G^{k(j)}(b|z,N,s)] g^{k(i)}(b|z,N,s) db \\
+ \beta \sum_{j\in N, j\neq i} \Pr(j \text{ wins}|z,N,s;\sigma^*) V_i(\omega(s,j)) \Bigg] + (1-\mathcal{P}) E_{N_{\neg i}} \Bigg[\beta \sum_{j\in N} \Pr(j \text{ wins}|z,N,s;\sigma^*) V_i(\omega(s,j)) \Bigg] \Bigg\}
\end{aligned}
$$

The expression on the second line can be modified further, by taking the summation outside of the integral and multiplying the term by $\frac{1-G^{k(i)}(b|z,N,s)}{1-G^{k(i)}(b|z,N,s)}$ and $\frac{g^{k(j)}(b|z,N,s)}{g^{k(j)}(b|z,N,s)}$,

$$
\begin{aligned}
&\beta \sum_{j\neq i} \Bigg\{ \int \frac{h^{k(j)}(b|z,N,s) V_i(\omega(s,j))}{\sum_{l\neq i} h^{k(l)}(b|z,N,s)} \prod_{l\neq i} [1 - G^{k(l)}(b|z,N,s)] g^{k(i)}(b|z,N,s) db \frac{1 - G^{k(i)}(b|z,N,s)}{1 - G^{k(i)}(b|z,N,s)} \frac{g^{k(j)}(b|z,N,s)}{g^{k(j)}(b|z,N,s)} \Bigg\} \\
=&\beta \sum_{j\neq i} \Bigg\{ \int \frac{\frac{g^{k(j)}(b|z,N,s)}{1-G^{k(j)}(b|z,N,s)} \frac{1}{g^{k(j)}(b|z,N,s)}}{\sum_{l\neq i} h^{k(l)}(b|z,N,s)} \prod_{l} [1 - G^{k(l)}(b|z,N,s)] \frac{g^{k(i)}(b|z,N,s)}{1 - G^{k(i)}(b|z,N,s)} g^{k(j)}(b|z,N,s) db \Bigg\} V_i(\omega(s,j)) \\
=&\beta \sum_{j\neq i} \Bigg\{ \int \frac{h^{k(i)}(b|z,N,s)}{\sum_{l\neq i} h^{k(l)}(b|z,N,s)} \prod_{l\neq j} [1 - G^{k(l)}(b|z,N,s)] g^{k(j)}(b|z,N,s) db \Bigg\} V_i(\omega(s,j)) \\
=&\beta \sum_{j\neq i} \Bigg\{ \int \frac{h^{k(i)}(b|z,N,s)}{\sum_{l\neq i} h^{k(l)}(b|z,N,s)} dG^{k(j)}(b|z,N,s) \Bigg\} V_i(\omega(s,j)),
\end{aligned}
$$

where $dG^{k(j)}(b|z,N,s) = \prod_{l\neq j}[1 - G^{k(l)}(b|z,N,s)]g^{k(j)}(b|z,N,s)$ is the derivative of the probability that a bid of bidder j is the lowest bid. Substituting this expression back into the equation for the value function, I obtain the final equation

$$V_i(s) = E_{z\mathcal{N}}\Bigg\{\mathcal{P}E_{N_i}\Bigg[\int \frac{1}{\sum_{j\neq i} h^{k(j)}(b|z,N,s)} dG^{k(i)}(b|z,N,s) \\ + \beta\sum_{j\neq i}\Bigg(\Pr(j \text{ wins}|z,N,s;\sigma^*) + \int \frac{h^{k(i)}(b|z,N,s)}{\sum_{l\neq i} h^{k(l)}(b|z,N,s)} dG^{k(j)}(b|z,N,s)\Bigg) V_i(\omega(s,j))\Bigg] \\ + (1-\mathcal{P})E_{N_{\neg i}}\Bigg[\beta\sum_{j\in N}\Pr(j \text{ wins}|z,N,s;\sigma^*)V_i(\omega(s,j))\Bigg]\Bigg\}$$

Furthermore, since I am considering type-symmetric strategies, the expression above can be simplified as follows,

$$V_i(s) = E_{z\mathcal{N}}\Bigg\{\mathcal{P}E_{N_i}\Bigg[\int \frac{1}{(N_{new}-1)h^{new}(b|z,N,s) + N_{reg}h^{reg}(b|z,N,s)} dG^{k(i)}(b|z,N,s) \\ + \beta\sum_{j\neq i}\Bigg(\Pr(j \text{ wins}|z,N,s;\sigma^*) + \int \frac{h^{k(i)}(b|z,N,s)}{(N_{new}-1)h^{new}(b|z,N,s) + N_{reg}h^{reg}(b|z,N,s)} dG^{k(j)}(b|z,N,s)\Bigg) V_i(\omega(s,j))\Bigg] \\ + (1-\mathcal{P})E_{N_{\neg i}}\Bigg[\beta\sum_{j\in N}\Pr(j \text{ wins}|z,N,s;\sigma^*)V_i(\omega(s,j))\Bigg]\Bigg\}.$$

2.9 Appendix: Tables and Figures

Table 2.10: Probability of Participation Regressed on the Value of Past Awarded Contracts

	(1)	(2)
log(value of contracts within past 6 months)	0.008 ***	0.009 ***
	(0.001)	(0.002)
log(value of contracts 6 to 12 months in the past)	0.011 ***	0.011 ***
	(0.001)	(0.002)
log(value of contracts 12 to 18 months in the past)	0.004 ***	0.002
	(0.001)	(0.002)
log(value of contracts more than 18 months in the past)	0.008 ***	−0.006 ***
	(0.001)	(0.002)
log(Engineer's estimate)	−0.133 ***	−0.016 *
	(0.006)	(0.009)
Funding Source (Federal=1)	−0.046 ***	−0.019
	(0.017)	(0.028)
Number of Planholders - Regular	−0.038 ***	−0.100 ***
	(0.005)	(0.008)
Number of Planholders - New	−0.008 ***	−0.032 ***
	(0.001)	(0.002)
Number of Planholders - Fringe	0.001	0.011 ***
	(0.002)	(0.003)
Standardized backlog	−0.086 ***	−0.110 ***
	(0.012)	(0.018)
Year, Month and District Indicators	N	Y
N	30.297	30.297

Note: Standard errors provided in parentheses. All specifications also include a constant. Statistical significance: 1%(***), 5%(**), 10%(*).

Table 2.11: Log(bid) Regressed on the Value of Past Awarded Contracts

	(1)	(2)
log(value of contracts within past 6 months)	0.000	−0.001 ***
	(0.000)	(0.000)
log(value of contracts 6 to 12 months in the past)	−0.002 ***	−0.002 ***
	(0.000)	(0.000)
log(value of contracts 12 to 18 months in the past)	0.000	−0.001
	(0.000)	(0.000)
log(value of contracts more than 18 months in the past)	−0.005 ***	0.001 ***
	(0.000)	(0.000)
log(Engineer's estimate)	0.990 ***	0.995 ***
	(0.004)	(0.004)
Funding Source (Federal=1)	0.015	−0.022 **
	(0.011)	(0.011)
Number of Planholders - Regular	−0.034 ***	−0.016 ***
	(0.004)	(0.004)
Number of Planholders - New	−0.017 ***	−0.007 ***
	(0.002)	(0.002)
Number of Planholders - Fringe	−0.037 ***	−0.023 ***
	(0.007)	(0.005)
Standardized backlog	−0.004	0.006 *
	(0.004)	(0.004)
Year, Month and District Indicators	N	Y
N	12.782	12.782

Note: Standard errors provided in parentheses. All specifications also include a constant. Statistical significance: 1%(***), 5%(**), 10%(*).

Chapter 3

Subcontractor Experience in Procurement Auctions

Procurement contracting is a frequently used tool to allocate government expenditures, and often affects a large number of firms of various scales. While some of these firms can be relatively enormous companies participating in the market for decades, others are smaller businesses attempting to enter or establishing themselves in the market. Government procurement auctions often have incentives in place that provide advantage to certain businesses, such as small firms or firms owned by veterans or minorities. Sometimes these incentives come in the form of a direct discount on contractor's bid. Other times, they are specified as a percentage of contract value that should be sourced from such businesses; these auctions usually give the primary contractor the option to outsource part of the work to other firms through subcontracting. Subcontracting can therefore serve as an ideal route to gain experience and "get a foot in the door" of the main procurement market.

Evidence of learning-by-doing in procurement auction markets, including the results from Chapter 2 of this dissertation suggests that experience plays an important role in the future success of recent entrants. In Chapter 2, I quantified firms' learning-by-doing by considering how firms' cost distributions evolve over time with respect to their amount of prior experience. This is of particular importance for recent entrants to the market. In reality, however, these recent entrants are not new firms established for the sole purpose of participating in this particular market but often existing firms that are attempting to expand their business or enter a new segment of the market.

Many firms participate in the procurement market as subcontractors before actively submitting bids as primary contractors. Participation in the market as subcontractors - hired by the primary contractor for specific tasks - can provide a way to acquire bidding, negotiation or other relevant skills and resources. Serving as a potential subcontractor to one or more firms increases the probability of carrying out at least a portion of the project workload. While experience gained as a subcontractor might be more limited in scope and task-specific, it can still be a valuable stepping stone for firms that have only recently entered the market. This paper explores the significance of subcontractor experience accumulated before entering the procurement auction market as a primary contractor and examines the effect of past subcontracting experience on the future probability of success in the bidding process.

The data studied here comes from the same government highway procurement market in California as in Chapter 2, however, it considers an even earlier history of a firm - prior to submitting a bid and being observed as a primary contractor. Firms considering entering this market have a choice to bid independently (and possibly subcontract part of the project to other firms, if they don't have the necessary expertise or resources), or to be listed as a subcontractor on a bid submitted by another firm. While these two options are not mutually exclusive, and firms can potentially participate in the same auction both as the primary contractor and a subcontractor for another bidder, this situation is rarely observed in data.

3.1 Related Literature

While most of the attention in the literature on procurement auctions concentrates on the issue of bidders' capacity constraints and its effects on procurement costs and efficiency in the main market (Jofre-Bonet and Pesendorfer (2003),[1] Balat (2013),[2] Groeger (2013),[3] Saini (2012)[7]), lately there has been a shift in focus to bring subcontracting into the picture. The presence of subcontractors and the ability of primary contractors to outsource tasks softens the effects of capacity constraints and intensifies competition, since contractors facing unexpectedly high costs can use the subcontracting market to effectively lower their costs. On the other hand, strategically behaving firms serving in both the main as well as the subcontracting markets can face increased opportunity costs of losing an auction if submitting and winning a bid would imply a forgone opportunity to earn profit from subcontracting.

Jeziorski and Krasnokutskaya (2014)[12] use a dynamic model to extend the literature on the effects of capacity constraints to the subcontracting market. They consider a dynamic environment where contractors with capacity constraints can use the subcontractor market to re-sell part of the awarded project. In their model contractors' costs vary stochastically but also depend on the level of prior commitments (backlog) due to contractors' capacity constraints. They find that in the environment with subcontracting the burden of a larger projects is considerably reduced. Because the availability of subcontracting slows down accumulation of backlog, it leads to lower costs that in turn intensify competition in the market and reduce informational rents. Since these effects result in lower mark-ups and lower prices, the procurement costs are also lower than in an environment without subcontracting.

The literature on capacity constraints is mostly concerned with the primary contractor's decision to use the subcontracting market, and studies how the availability of subcontracting affects the outcome of the auction and the primary contractor's cost distribution. In contrast, this paper considers the other side of the coin, which includes the potential benefits of participating in the subcontracting market to the subcontractor – rather than the primary contractor.

Several other papers develop theoretical models of subcontracting in static environment. Wambach (2009)[13] shows that if subcontracting happens before the contract is awarded, it matters how subcontractors are chosen by the primary contractors. If subcontractors do not compete at different primary contractors they might take into account the fact that they can influence the ultimate competitiveness of "their" primary contractor through the terms negotiated between them with regards to the subcontracting work. Gale et al. (2000)[14] study subcontracting in sequential static procurement auctions with two symmetric bidders, and conclude that without the possibility of subcontracting the auction outcome is generally inefficient.

This paper is more closely related to the empirical papers on subcontracting. Miller (2014)[15] considers a smaller set of 32 bridge projects to study the impact of contractual incompleteness on subcontracting through ex-post revisions. Marion (2014)[16] investigates the issue of horizontal subcontracting using auction data from the California highway procurement market. He finds this occurs relatively rarely, with only about 10% of projects receiving a bid from a firm that was also listed as a subcontractor on another bid. He also examines the issues of affirmative action and subcontracting with the purpose of satisfying requirements regarding participation of disadvantaged businesses.

Perhaps the most related work by De Silva et al. (2011)[17] examines past subcontracting experience on business duration in data on road construction projects in Texas. The authors find that pre-entry experience in the form of early involvement as subcontractor increases chances of survival. They attribute prolonged firm duration to the lower risk subcontractor experience that improves firms' competitive advantage.

3.2 Market Description

The market studied in this paper is the same as in Chapter 2, and consists of road construction and repair projects awarded by the California Department of Transportation (Caltrans) in repeated first-price sealed-bid auctions.

Interested potential bidders can request additional materials providing further details on the location and description of the project and the bid letting date. The government engineer also specifies the list of items and their corresponding quantities required to complete the project, resulting in the engineer's estimate of the overall cost of the project.

Firms that decide to submit a bid are allowed to subcontract a portion[1] of the project to other firms (subcontractors) but must list on their bid all subcontractors accounting for at more than 0.5 percent of the contract value or $10,000, whichever is greater. Since bidders must provide the list of subcontractors as well as the price for each item, negotiations with potential subcontractors are conducted prior to bid submission and are completed before the outcome of the auction is realized.

As long as firms satisfy pre-qualification requirements, they are allowed to participate in the same auction both as a bidders as well as a subcontractor, even though

[1] The minimum percentage of the project value required to be completed by the primary contractor changes over time, together with the exact definition and inclusion of specialty projects. However, it is generally on the order of 50%.

this is not a common occurrence. Firms are also allowed to negotiate with multiple bidders, and be listed as a subcontractor on several competing bids.

3.3 Data and Summary Statistics

The data used in this paper covers the period from July 2003 through December 2009. During this period, there were 3,373 road and repair projects awarded, with about 540 auctions per year, with the number of projects generally rising over time.

The dataset consists of variables capturing the size and complexity of the projects (such as the engineer's estimate, number of workdays, and number of items, source of funding, location and type of project) together with a list of bids submitted for each project, including the identity of all bidding firms and their corresponding bid amounts. Furthermore, for each submitted bid (whether winning or losing) the data contains a list of subcontractors pre-arranged to carry out certain tasks conditional on the bidding firm being granted the project.

Overall, I observe 18,209 bids submitted by 1,079 different firms, with 5.4 bids submitted per projects on average. 468 of these firms won at least one contract during the period. On the subcontracting side, I observe 40,785 bidder-subcontractor pairs, with about 12 subcontractors listed on average per bid. Overall, 4,814 firms appear in the dataset during the period in the role of either subcontractor or bidder. Many of the firms (3,735 firms, corresponding to 77.6% of all firms) only serve as subcontractors – these are firms that often specialize in a particular narrow field of technical expertise, such as laying submarine cables or installing air-conditioning – and these firms might not have the technical expertise and resources to manage and carry out the whole project.

On the other hand, 612 (12.7%) firms are never listed as a subcontractor on another firm's bid, while the remaining 467 (10.7%) firms participate at least once as a subcontractor on a bid, while also submit at least one bid independently during the period. These are the firms of interest for which subcontracting can serve as another way of obtaining valuable experience in the market, through familiarizing themselves with the bidding process and bid preparation, and gaining experience in conducting pre-bidding negotiations and satisfying government requirements. On some occasions, the same firm participates in a particular auction both as an independent bidder as well as a subcontractor, although this occurs in less than 5% of the data. Marion (2014)[16] investigates the issue of horizontal subcontracting in this market in more detail.

In order to quantify the effects of subcontractor experience and investigate whether this type of experience improve future prospects of firms in this market, I compare bidders that enter this market without prior subcontracting experience with those that enter as after participating in the role of a subcontractor. I use the initial 6 month of data to establish the pool of incumbent bidders, and categorize firms that submit a bid for the first time after January 2004 as new entrants to the market.

Table 3.1 provides descriptive information for entrants with and without subcontracting. Firms with prior subcontracting experience participate on average more frequently and win more often.

Table 3.1: Summary Statistics for Recent Entrants by Subcontracting Experience

	Bidders with prior subcontracting experience		Bidders without prior subcontracting experience	
	Mean	Std. Dev.	Mean	Std. Dev.
Number of projects where listed as subcontractor	33.1	113.1	-	-
Number of winning bids where listed as subcontractor	7.5	4.5	-	-
Number of own bids	11.5	23.4	5.3	11.2
Number of own winning bids	2.1	5.3	0.8	2.3
Probability of winning	0.13	0.22	0.12	0.25

3.4 Empirical Analysis

The empirical analysis considers the question of the effect of past subcontractor experience on the future success of bidders, and addresses this question by studying firms as they enter the marketplace, with or without subcontracting experience. First, I investigate whether participation in the role of a subcontractor improves future chances of winning a project, and in particular, how does this effect depend on the number of past successes as a subcontractor, and on the amount of prior subcontracting work. Second, I examine the relationship between subcontracting experience and the level of bids that firms submit in future auctions. Again, I consider both the effect of the number of instances as well as the amount of subcontracting work performed in the past. I find that subcontracting experience improves the chances of bidders to succeed in future auctions. Furthermore, more recent subcontracting experience has a stronger effect than subcontracting experience further in the past.

3.4.1 Probability of Winning an Auction

In order to examine whether subcontractor experience has any lasting effect on the future of firms in this marketplace, and to quantify this effect, I compare similar bidders just entering the procurement market but differing in whether or not they have previously participated in the subcontractor market. There are several dimensions along which subcontracting experience could be measured: the number of bids (or projects) on which a firm is listed as a subcontractor; the number of only the winning bids on which a firm participates; or the dollar amount of work performed by the subcontractor. These measures capture different benefits that might arise. For example, the benefits of learning to negotiate prices accumulate from both successful and unsuccessful bids, as negotiations have to be completed prior to bid submission. On the other hand, construction cost reductions from learning-by-doing will only occur when the bid is successful and the project is carried out. Hence, I consider a number of different measures of subcontractor experience to capture a variety of gains that could result from subcontractor experience.

Table 3.2 presents the results of estimating a simple probit equation,

$$Probit(y|X, z) = \alpha + X'\beta + z\gamma$$

, regressing the dependent indicator variable for a winning bid y on the variable capturing any prior subcontractor experience, z, and a set of controls, X, on the sample of recent entrants. The control variables include project characteristics measuring the size and duration of projects such as the engineer's estimate and number of workdays (both after logarithmic transformation); auction characteristics measuring the competitiveness of the bidding environment such as the number of bidders and the number of planholders[2]; all regressions also include a constants and a set of indicator variables controlling for year and month effects.

The positive and statistically significant coefficient on the variable indicating prior subcontractor experience suggests that firms with such experience are more likely to win future auctions, even after controlling for the project characteristics. As might be expected since the sample only includes firms that have not been participating in the market for very long, these firms are less likely to win bidding for larger or more complex projects (even though these effects are not significant). Similarly, bidding against a larger number of competitors also decreases probability of winning. A less straightforward result is the positive coefficient on the number of planholders, which

[2]Planholders are all firms that requested detail bidding materials on a particular project. Bidders constitute a subset of the set of planholders, as not all planholders proceed to submit a bid.

might indicate "popularity" of projects.

Table 3.2: Probability of Winning for Recent Entrants

Dependent variable: Indicator for winner	
Indicator for a bidder with subcontractor experience	0.082 **
	(0.038)
Log (engineer's estimate)	−0.036 **
	(0.016)
Number of planholders	0.014 ***
	(0.003)
Number of bidders	−0.099 ***
	(0.007)
Log (number of workdays)	0.018
	(0.017)
Year and month indicators	Y
N	6,725

Note: Standard errors corrected for clustering by auction provided in parentheses. All specifications also include a constant. Statistical significance: 1%(***), 5%(**), 10%(*).

Just as firms can accumulate subcontractor experience gradually as they perform more projects, the relevant experience likely depreciates and become less valuable the further away it from an event. Over time, employees with the appropriate human capital might leave, or relationships with other firms and subcontractors might weaken. Since the data covers more than 5 years of auctions, in addition to the overall level of past experience I also consider finer measures that focus on the months immediately preceding a particular bid.

In Table 3.3, I consider the sample of recent entrants with subcontractor experience and examine three different measures. The first column uses the measure of total subcontractor experience, constructed as the count of all prior winning bids on which a firm was listed as a subcontractor. The second column limits the relevant window to the prior 12 months, and finally, the third column is the most restrictive, only accounting for the prior 6 months.

More recent subcontracting experience has a stronger effect that experience obtained further in the past. The positive and statistically significant coefficient indicates that firms which served as subcontractors on a larger number of projects are relatively more likely to win when submitting independent bids.

Table 3.3: Probability of Winning for Bidders with Subcontractor Experience

Dependent variable: Indicator for winner	(1)	(2)	(3)
Number of all prior winning bids	0.002		
on which listed as subcontractor	(0.001)		
Number of winning bids in the past 12 months		0.005 **	
on which listed as subcontractor		(0.002)	
Number of winning bids in the past 6 months			0.013 **
on which listed as subcontractor			(0.005)
Log (engineer's estimate)	−0.075 ***	−0.072 ***	−0.072 ***
	(0.024)	(0.024)	(0.024)
Number of planholders	0.017 ***	0.018 ***	0.018 ***
	(0.005)	(0.005)	(0.005)
Number of bidders	−0.090 ***	−0.088 ***	−0.088 ***
	(0.011)	(0.011)	(0.011)
Log (number of workdays)	0.008	0.011	0.012
	(0.025)	(0.025)	(0.025)
Year and month indicators	Y	Y	Y
N	3,910	3,910	3,910

Note: Standard errors corrected for clustering by auction provided in parentheses. All specifications also include a constant. Statistical significance: 1%(***), 5%(**), 10%(*).

The number of times that a firm has participated on projects as a subcontractor is a somewhat crude measure of its subcontractor experience. While this measure will capture benefits that accrue regardless of the project size, such as price negotiations, it will disregard important benefits that depend on the size, complexity or duration of a project. A subcontractor working on a small road repair will not gain the same experience as a subcontractor working on a construction of a new bridge. Therefore, I also consider the dollar value of prior subcontracting work (after logarithmic transformation).[3]

The results of a probit regression of the indicator for winning bid on the value of prior subcontractor experience are presented in Table 3.4. Larger dollar amount of past subcontracting work increases the future probability of winning an auction, and this relationship is preserved even after controlling for the amount of work performed as a primary contractor. The effects of the engineer's estimate and the number of competitors remain similar as before, when considering just the number of subcontracting occasions. Overall, the effects of subcontracting experience on the future probability of winning auctions suggest that the benefits of becoming familiar with the market and the government bidding procedures exceeds the potential effects of capacity constraints. If the backlogs accumulated by these firms due to subcontracting increased their construction costs for the foreseeable future, one would expect a negative effect on the probability of winning, and this effect would be more pronounced the shorter the period preceding the auction.

[3] While some bids specify the exact items that a subcontractor is going to perform, other provide limit detail. As an approximation, I construct the dollar value of work performed by a subcontractor as if the work was split evenly among all subcontractors.

Table 3.4: Probability of Winning for Bidders with Prior Subcontractor Experience

Dependent variable: Indicator for winner	(1)	(2)
Log (value of prior subcontractor work)	0.017 *	0.016 *
	(0.009)	(0.009)
Log (value of winning bids)		0.019 ***
		(0.004)
Log (engineer's estimate)	−0.077 ***	−0.094 ***
	(0.024)	(0.025)
Number of planholders	0.017 ***	0.020 ***
	(0.005)	(0.005)
Number of bidders	−0.090 ***	−0.088 ***
	(0.011)	(0.011)
Log (number of workdays)	0.010	0.027
	(0.025)	(0.025)
Year and month indicators	Y	Y
N	3,910	3,910

Note: Standard errors corrected for clustering by auction provided in parentheses. All specifications also include a constant. Statistical significance: 1%(***), 5%(**), 10%(*).

3.4.2 Level of Bid Submitted in a an Auction

Besides the probability of winning an auction, a second relevant measure for comparing the success of firms with different history of subcontracting experience is the level of their bids. In particular, the dependent variable that I use in the analysis below is the dollar amount submitted by a bidding firm divided by the engineer's estimate. The government engineer computes the estimate based on the list of items that need to be completed as well as the expected labor and material costs. The ability of a firm to submit a lower bid relative to the engineer's estimate will affect its probability of winning an auction, as these are first-price sealed-bid procurement auctions in which Caltrans awards the project to the lowest responsive qualified bidder.[4]

Similar to the analysis of probability of winning above, I start by comparing the overall effect of subcontractor experience on the relative bid. Table 3.5 presents the results of the OLS regression

$$\text{bid/engineer's estimate} = \alpha + X'\beta + z\gamma + \epsilon,$$

regressing the scaled bid on the variable indicating any past subcontracting experience, z, as well as the set of control variables contained in X. The results show that firms with past subcontractor experience are able to submit lower bids as compared to firms without any such experience, and this effect is statistically significant. Similarly as before, presence of a larger number of competitors has a downward effect on the relative bid, as firms need to compete more aggressively in order to win.

Focusing on recent entrants with subcontractor experience, I further investigate the relationship between the number of times a firm has completed work as a subcontractor and its relative bid level. In Table 3.6, the number of past winning bids on which a firm has been listed as a subcontractor has again been measured in three different ways: first, as the total number of events in the past; the number of events in the prior year; and the number of events in the prior 6 months. While subcontractor

[4]Depending on the source of funding and the time period, bidders can face additional requirements or goals such as sufficient participation of Disabled Veteran Business Enterprise firms on the contract. Some firms can receive an advantage through the small business bid preference program, which I studied in detail in Chapter 2. Marion has also examined affirmative action in auctions.

Table 3.5: Bid Amount Relative to Engineer's Estimate for Recent Entrants

Dependent variable: Bid/Engineer's estimate	
Indicator for bidder with subcontractor experience	−0.048 ***
	(0.007)
Number of planholders	0.002 **
	(0.001)
Number of bidders	−0.011 ***
	(0.002)
Log (number of workdays)	−0.002
	(0.006)
Year and month indicators	Y
N	6,725

Note: Standard errors corrected for clustering by auction are provided in parentheses. All specifications also include a constant. Statistical significance: 1%(***), 5%(**), 10%(*).

experience over the whole history of firm does not seem to have an effect on the bid, the recent experience has a significant negative effect, with the effect being stronger for the most recent experience.

Analogous to the study of probability of winning conducted earlier, I also examine the dollar value of past subcontractor experience. The pattern of this relationship again reveals that the amount of prior work performed in relatively recent past has a significantly negative effect on scaled bid. However, this time the effect is strongest when considering the full year prior to bidding, which suggests that firms performing subcontracting work in the recent past might be facing higher cost, which would tend to put upward pressure on their bids and counteract the benefits of obtaining such experience.

Overall, the analysis above has consistently shown that among firms entering this

Table 3.6: Bid Amount Relative to Engineer's Estimate for Recent Entrants with Subcontractor Experience

Dependent variable: Bid/Engineer's estimate	(1)	(2)	(3)
Number of all prior winning bids on which listed as subcontractor	0.000		
	(0.000)		
Number of winning bids in the past 12 months on which listed as subcontractor		−0.003 ***	
		(0.001)	
Number of winning bids in the past 6 months on which listed as subcontractor			−0.006 ***
			(0.001)
Number of planholders	0.003 **	0.002 *	0.002 *
	(0.001)	(0.001)	(0.001)
Number of bidders	−0.011 ***	−0.012 ***	−0.012 ***
	(0.002)	(0.002)	(0.002)
Log (number of workdays)	0.007	0.005	0.004
	(0.008)	(0.008)	(0.008)
Year and month indicators	Y	Y	Y
N	3,910	3,910	3,910

Note: Standard errors corrected for clustering by auction provided in parentheses. All specifications also include a constant. Statistical significance: 1%(***), 5%(**), 10%(*).

Table 3.7: Bid Amount Relative to Engineer's Estimate for Recent Entrants with Subcontractor Experience

Dependent variable: Bid/Engineer's estimate	(1)	(2)	(3)
Log (value of all prior subcontractor work)	−0.002		
	(0.001)		
Log (value of subcontractor work in the past 12 months)		−0.003 ***	
		(0.001)	
Log (value of subcontractor work in the past 6 months)			−0.002 ***
			(0.001)
Number of planholders	0.003 **	0.001	0.002
	(0.001)	(0.001)	(0.001)
Number of bidders	−0.011 ***	−0.011 ***	−0.011 ***
	(0.002)	(0.002)	(0.002)
Log (number of workdays)	0.006	0.006	0.006
	(0.008)	(0.008)	(0.008)
Year and month indicators	Y	Y	Y
N	3,910	3,910	3,910

Note: Standard errors corrected for clustering by auction provided in parentheses. All specifications also include a constant. Statistical significance: 1%(***), 5%(**), 10%(*).

market having subcontractor experience is associated with an increased probability of winning an auction and with lower bid relative to the engineer's estimate. These effects tend to be more pronounced in both magnitude and significance for the most recent experience, despite the possibility that firms participating on winning contracts shortly before submitting a bid might be facing higher costs due to increased backlog. Furthermore, both the number of past occasions on which a firm performed subcontracting work as well as the dollar value of the work seem to be informative. This might be related to the fact that subcontracting experience brings about benefits along various dimensions, including price negotiations, project coordination and management, or labor and material acquisition. While some of these benefits accumulate regardless of project size, others will tend to increase with the scale of the task.

Note that the analysis above does not attempt to explain firms' participation decisions, whether with respect to the decision to participate in the subcontractor market or to submit a bid. While these decision can be in part driven by the costs that firms face at a given point in time, and thus depend on the level of backlog, firms are likely taking into account many other factors, such as options available outside the government procurement process. Such factors are difficult to measure and are not available in the current dataset.

3.5 Conclusion

Despite a number of government programs in place to help firms establish themselves and remain in the procurement market for construction projects, bidder learning has garnered limited attention. In Chapter 2 of this dissertation, I presented evidence on learning-by-doing in the Caltrans procurement auction market suggesting that previous bidding experience plays an important role in the future success of recent entrants. I showed that as firms submit successful bids and accumulate experience, they slide down the learning curve and achieve average costs below those of inexperienced firms.

However, many firms participate in the closely intertwined subcontracting market either before or while actively submitting bids as primary contractors. Modeling the entire decision that such firms face would be a complex endeavor: not only do firms have to decide whether or not to submit a bid, and at what level but also whether or not to be listed as a subcontractor on one or more competitors' bids, and what price to negotiate for the pertinent tasks.

Firms in this situation face important tradeoff: serving as a potential subcontractor to one or more firms increases the probability of participating on the winning bid and carrying out at least a portion of the project workload; on the other hand, choosing to submit its own bid and not serve as a potential subcontractor might provide a cost advantage in bidding if the alternative subcontractors would be more costly to the competitors.

The decision gets even more involved in presence of learning-by-doing. If performing subcontracting work brings about long-run benefits the extend into the future and provide the subcontractor with a competitive advantage in subsequent auctions, the firm might find it optimal to participate in the subcontracting market prior to independent bidding.

While experience gained as a subcontractor is more limited in scope and task-specific, the empirical analysis in this chapter has shown that it can be a valuable stepping stone for firms that have only recently entered the market.

Chapter 4

Trading across Borders in Online Auctions

The clustering pattern of trade between countries has puzzled economists for a long time. Several theories have been proposed to explain this pattern, with gravity theory gaining the most attention. Gravity theory asserts that the volume of trade between countries is inversely proportional to the distance between these countries. Informally, this means that countries located geographically close to each other tend to trade more. The obvious rationale for this regularity is that since transaction costs are related to the distance close proximity implies lower trading costs.

We re-examine this issue using data from an online market for programming services. The market is organized around an online platform that allows buyers to solicit price quotes for their programming projects from a large number of affiliated sellers.

This market is international in nature: the buyers and sellers participating come from many different countries. At the same time, physical transaction costs are relatively low and should in theory be the same across buyer-seller country pairs.

We begin by presenting descriptive evidence documenting the clustering patterns in this online market. We then proceed to build and estimate a model that rationalizes project allocation decisions by allowing the distribution of buyers' tastes for sellers' attributes to vary across buyer countries while remaining constant across seller countries; similarly, this model allows for sellers' attributes to differ across seller countries but restricts them to remain constant across buyer countries. That is, our model does not allow for bilateral preferences. We use the estimation results to evaluate the importance of various potential channels that might contribute to clustering.

Transactions in this market are implemented in the form of multi-attribute auctions that allow buyers to deviate from allocation based solely on price (as is the case in standard auctions) in favor of service providers who are associated with higher buyer-specific value. We formalize the features of this market in a model where each project attracts a set of sellers who submit bids for buyer's consideration. The project is awarded to the seller who delivers the highest value over price as long as it exceeds buyer's outside option; otherwise, the project is not awarded. We model buyers' values as represented by the sum of a systematic component (seller's quality), which is the same for all buyers, and a stochastic buyer-seller-specific match component. We additionally allow for both strategic pricing and participation decisions of sellers.

We exploit the structure of our model to decompose observed clustering patterns of trade into various contributing components. We begin by separating the patterns associated with buyers' allocation decisions from those associated with sellers' participation decisions. Indeed, if the allocation decisions of a certain group of buyers demonstrate preference for a particular type of seller, this pattern may be further reinforced by the participation decisions of sellers aware of this preference. First, we isolate the allocation patterns based on buyers' decisions by eliminating differences in participation behavior. We then further evaluate the importance of buyers' weights, which are uniform across seller countries, relative to sellers' attributes, which are uniform across buyer countries. This analysis provides an insight into the nature of the "attraction" between particular sellers and buyers.

Our estimation methodology is tailored to accommodate the features of data that typically arise in service markets, and in the online service market especially. We have to specifically address the following features: (a) in this market, buyers are able to obtain a well-informed assessment of seller's quality through interpersonal communication and by reviewing examples of previous work; since such information is not available to the researcher, the buyer is likely to be better informed about seller's quality than the researcher; it is thus important to account in the analysis for the unobserved heterogeneity among sellers; (b) the rules applied by buyers to allocate their projects are not observed by the researcher and thus have to be modeled as random coefficients; (c) the choice set in this market tends to be buyer-specific; this complicates the analysis because the standard discrete choice methodologies rely on the fact that the probability of choosing a given individual option conditional on buyer's choice set can be precisely estimated from data – this property does not hold in our environment; (d) a large number of sellers stay in this market only for a brief period of time (transitory sellers).

We assume that transitory sellers are not able to provide buyers with extensive information about their qualities due to their short experience with the online platform. The buyer, therefore, evaluates transitory sellers on the basis of an average quality typical for such seller's country. In contrast, we assume that the quality of permanent sellers is fully observed by buyers. The large number of permanent sellers does not allow us to estimate quality at the level of the individual seller. Instead, we rely on the classification procedure proposed in Krasnokutskaya et al. (2014)[18] to recover the quality group structure for the set of permanent sellers. This procedure identifies sets of sellers that are characterized by equal quality levels, and thus substantially reduces the number of parameters that need to be recovered in estimation (each group of sellers is associated with the same quality level). We rely on the private variation in sellers' costs that remains after conditioning on sellers' characteristics to identify weights used by buyers in the allocation process.

We find that our model fits data quite well and is capable of explaining 80-90% of the magnitude of clustering observed in the data. Thus, bilateral preferences appear to play a limited role in this market. Furthermore, we establish that buyers' preferences for sellers' attributes are significantly reinforced by strategic participation which amplifies certain buyers' apparent preference for sellers from specific countries. In particular, sellers from any given country choose to participate more often with buyers who appear to give preference to their country. Finally, we decompose clustering in buyers' choices as accounted for by the heterogeneity in various seller attributes, such as the average reputation score, quality, and cost. We find that among the considered attributes differences in sellers' costs contribute most to the generation of the clustering pattern. Cost differences translate directly into price differences and when compounded by the differences in buyers' price sensitivity they become a significant source of the clustering allocation pattern.

The chapter is organized as follows. Section 4.1 summarizes the relevant literature, whereas Section 4.2 describes the market; the model is developed in Section 4.3. Section 4.4 discusses the empirical methodology, followed by Section 4.5 describing the data and Section 4.6 presenting the results of the descriptive analysis. The estimation results are provided in Section 4.7; Section 4.9 outlines the results of the clustering decomposition. Section 4.10 summarizes the findings.

4.1 Related Literature

Patterns of inter-country trade in physical goods have long been documented in the international trade literature (e.g. Dornbusch et al. (1977),[19] Bergstrand (1989),[20] Krugman (1991)[21]). Empirical analyses in this line of research have identified variables such as geographic proximity, shared borders, former colonial relationships, and common language as some of the important drivers of trade.

However, with the rise of the internet and global communication networks trade is no longer restricted to the physical networks of highways and shipping routes (Leamer and Storper (2001)[22]). Trading in professional services is becoming increasingly prevalent, with numerous examples in software development, evaluation of insurance claims, or provision of customer support. While the importance of historical relationships between countries will likely continue to drive patterns of trade, information and communications technology will erase many borders and put more weight on skills relevant in the markets for professional services. So far, however, there is little empirical evidence documenting these new trade patterns. Our work relates to two streams of prior literature: models explaining patterns of global trade and multi-attribute auction theory.

Many empirical studies analyzing bilateral trade flows have found the intensity of trade to be inversely related to the distance between the trade partners (Tinbergen (1962),[23] Poyhonen (1963),[24] Linnemann (1966)[25]). Moreover, the amount of trade is found to increase proportionally to the product of the GDPs of the two countries, with this relationship remaining remarkably robust across time as well as across different countries. This pattern of trade strongly resembles Newtonian physics, which postulates that the force between two objects is proportional to the product of the two masses divided by the square of the distance between the objects. For this reason, it is often referred to as *gravity theory*, and the basic model characterizes trade flows between countries i and j by the following equation

$$\text{Trade Flow}_{ij} = \alpha * (GDP_i * GDP_j)/\text{distance}(i, j).$$

Several empirical papers have augmented the basic gravity model to allow for effects of common borders (McCallum (1995)[26]), common language, former colonial relationships, and exchange rates (Frankel (1994)[27]).

In the market for online programming services that we analyze in this chapter, seller attributes other than price may affect buyer's choice of a provider. For this reason, our work also relates to a second stream of literature on multi-attribute auctions. Such auctions are commonly used in the procurement of non-homogeneous goods, where other attributes besides price such as quality or performance play a role in the allocation of the auction (such as procurement of a new weapons systems by the Department of Defense).

Unlike scoring auctions, analyzed by Laffont and Tirole (1987),[28] McAfee and McMillan (1987),[29] Che (1993),[30] and more recently by Asker and Cantillion (2006),[31] multi-attribute auctions have received relatively little attention in the literature. In contrast to a scoring rule auction, bidders participating in a multi-attribute auction cannot specify or choose characteristics of their services. Instead, buyers take into account long-run bidder characteristics that are not meant to be altered for the purpose of a specific auction. In addition, buyers in multi-attribute auctions do not announce the weights attached to different bidder characteristics.

Finally, the work most closely related to our paper is by Greenstein (1993).[32] Greenstein models and estimates the choice process of government computer procurement using a multinomial logit model of vendor choice. The probability of a contract being awarded is modeled as a function of a representative of the vendor being present at the corresponding government agency office, the extent of previous buyer-vendor interaction (including the installed base) and other compatibility related factors. The results establish a strong incumbent bias - even after controlling for factors capturing the vendor-buyer match, the buyer is much more likely to award the deal to its previous vendor.

4.2 Market Description

This paper studies an online market for programming services, in which a platform serves as an intermediary between buyers (the demand side) and potential sellers (the supply side). Buyers procure programming services such as platform programming, databases, graphics programming and website design by posting job announcements to which interested sellers can respond by submitting a quote for a price at which they would be willing to complete the task. While the majority of sellers attracted by this online market come from North America, Eastern Europe, and South and East Asia, the market serves buyers from a variety of countries around the world by providing them access to a diverse set of sellers who differ in cost and quality.

The intermediary company allocates jobs through multi-attribute auctions, allowing buyers to take into account multiple seller characteristics in addition to the price quote. As a result, the selected seller is not necessarily the one who submits the lowest quote, reflecting the importance of quality in addition to the cost of providing a service. An important feature of this allocation mechanism is that the award rule is not announced and thus remains unknown to other market participants.

The registry provides limited information on verifiable "outside" credentials as well as information about the on-site performance of the seller. The latter includes a history of performance-related measures such as reputation scores or ratings, buyers' numerical feedback about working with a given seller, as well as instances of delays and disputes. In the case of a dispute, the company provides professional arbitration services that ensure that a seller is paid if only if the completed job satisfies industry standards.

Sellers often communicate with buyers before posting price quotes, with an average of three messages exchanged between a seller and a buyer prior to submitting a price quote. Sellers can also attach an example of previous work or a sketch of the proposed code. The frequency and content of this communication are not observed by the other sellers. Hence, while the buyer has an opportunity to form an opinion about each seller's quality, competing sellers have much more limited knowledge of their competitors' levels of quality. In principle, competitors can infer seller's quality from his long-run rate of winning.

When a seller contacts a buyer for any reason, his code name appears on the project webpage. Therefore, at any point in time, a visitor to the page can see the list of all sellers who have contacted the buyer up to that point. This list generally does not coincide with the set of sellers who eventually submit price quotes, since sellers can contact buyers without submitting a quote. The list of sellers who have contacted a buyer likely reflects potential rather than actual competition in an auction, meaning that prospective sellers do not observe the exact set of their competitors.

Finally, most of the buyers in our dataset appear only once in this market. On the other hand, a large number of sellers stay in the market for a long time (permanent sellers), whereas a considerable fraction participate only in a few auctions before leaving the market (transitory sellers).

4.3 Model

The demand side of this market for programming services is represented by one-time buyers. Each buyer brings a single project to the market and seeks to hire a service provider who would complete the job. Project l is summarized by a set of characteristics, z_l, such as the date when the project is posted, and the contract terms (e.g., specification of the programming tasks, deadlines, etc.).[1] The buyer posts the description of the project and invites sellers to submit price quotes for the project.

The supply side of the market consists of sellers who are described by a vectors of characteristics, x, and a scalar characteristic, q, which admits a finite number of values, $q^1 < q^2 < \cdots < q^K$. We assume that characteristic q is vertical (or quality-like) in nature. Non-quality characteristics capture sellers' countries of origin as well as their recorded performance measures such as reputation scores, delays or instances of conflict. Sellers can be of two types: permanent, i.e. those who stay with the market for a long time; and transitory, who participate in the market only for a short time and leave after completing one or two projects.

4.3.1 Project Allocation

Buyers allocate projects through multi-attribute auctions that allow them to deviate from allocation based solely on price (as in standard auctions) in favor of service providers who are associated with higher buyer-specific value. We use A_l to denote the set of sellers who submit a bid for project l and refer to such sellers as *active* bidders. These sellers form the buyer's choice set.

Extensive communication between buyers and sellers as well the record of sellers' performance measures provided by the platform allows buyers to be well-informed about permanent sellers' attributes. We allow, however, that a researcher may not

[1] We use l to index a buyer as well as the job he posted.

observe permanent sellers' vertical characteristic q. Furthermore, both the researcher and the buyer are not informed about the vertical characteristic of transitory sellers (possibly because a seller does not have enough information to provide to the buyer).

Buyer l associates a (private) value, $\Delta_{l,i}$, with an active seller $i \in A_l$ and awards his project to the active seller with the highest level of $\Delta_{l,i} - b_{l,i}$ if this level exceeds buyer's outside option $U_{0,l}$; otherwise, he leaves the project unassigned. The buyer's value is a weighted average of seller's attributes with buyer-specific weights α_l, β_l and intercept $\epsilon_{i,l}$ (the residual value assigned by buyer to a match with a specific seller), i.e.,

$$\Delta_{i,l} = \alpha_l \tilde{q}^{k(i)}(x_i) + x_i\beta_l + \epsilon_{i,l}, \tag{4.1}$$

where $\tilde{q}$ refers to the residual quality after the mean quality associated with observed characteristics, $x_i\beta_l$, is netted out. We let $\epsilon_l = \{\epsilon_{1,l}, ...\epsilon_{|A|,l}\}$ and refer to (α_l, ϵ_l) as the vector of buyers' weights.

Note, that in the case when β_l is constant across buyers the expression for $\Delta_{i,l}$ can be re-written as

$$\Delta_{i,l} = \alpha_l q^{k(i)}(x_i) + \epsilon_{i,l}. \tag{4.2}$$

In keeping with the definition of a multi-attribute auction, sellers do not observe the weights or the outside option of a specific buyer and consider them to be a random draw from some joint distribution of weights and outside options characterizing the population of buyers.[2]

4.3.2 Seller's Strategies

Let N_l denote the set of potential bidders for a given auction l. This set is partitioned into subsets of permanent and transitory sellers, $N = N^p \cup N^t$. Recall that a buyer is informed about the full vector of permanent sellers' characteristics and of transitory sellers' x and $E[q|x]$.

During an auction for project l each potential bidder i observes some private signal, or entry costs, $E_{i,l}$, drawn from the distribution $F_E(.|(x_i, q_i))$ and is aware of N_l. More specifically, seller i's information set consists of $E_{i,l}$ and $I_{N,l}$, where the latter contains information on the number of potential permanent bidders by quality group, and the total number of potential transitory bidders. Given this information set, potential bidder i decides whether to participate in the auction or not. His entry strategy σ_i^E is a mapping from the supports of $E_{i,l}$ and $I_{N,l}$ into $\{0, 1\}$. We denote the entry decision (outcome) by $D_{i,l}$ ($D_{i,l} = 1$ if enters and $D_{i,l} = 0$ otherwise).

Upon entry, an active bidder observes a private cost $C_{i,l} \in \mathbb{R}_+$ for completing the project. The cost of seller i characterized by (x, q) is distributed according to $F_C^k(\cdot)$.

[2]Unobservability of $\epsilon_{l,i}$ rules out common (or even correlated between buyer and seller) understanding of the specific seller's suitability for a given project. This may appear to be restrictive in view of the extensive interaction between buyers and sellers. However, from the theoretical point of view, this component controls the size of (unsystematic part of) surplus generated by the specific buyer-seller match. The simulated solution of the game where $\epsilon_{l,i}$ is known indicates that the knowledge of $\epsilon_{l,i}$ allows bidders to extract a larger part of surplus. Thus, it is not in the interest of the buyer to share any information that would reveal $\epsilon_{l,i}$ to seller i. More specifically, the interview should be conducted in such a way as to elicit maximum information about seller's suitability without revealing the match value to the seller.

The bidder does not observe participation decisions of other potential bidders, and is thus unaware of the composition of the set of active bidders. He then submit a price $B_{i,l}$ based on his information set.

In line with the existing empirical auction literature, we assume that the observed outcomes reflect a type-symmetric pure strategy Bayesian Nash equilibrium (psBNE).[3] In such an equilibrium, participants who are *ex ante* identical (i.e. either "$i, j \in N_l^p$ and $q_i = q_j$, and $x_i = x_j$" or "$i, j \in N_l^t$ and $x_i = x_j$") adopt the same strategies. The bidding strategy for seller i who is characterized by (x, q) is denoted $\sigma^{r,(x,q)} : [\underline{c}, \overline{c}] \to \mathbb{R}_+$, and entrant i's expected profit from bidding b, i.e. submitting $\sigma^{r,(x,q)}(c) = b$ when the cost draw is equal to c is given by

$$\Pi^{r,(x,q)}(b, c; \sigma^{-i}) \equiv (b - c) \Pr(i \, \text{wins} \mid b, \, i \in (x, q); \sigma^{-i}),$$

where σ^{-i} denotes a profile of other sellers' strategies that they would use should they become active in a given project, and $\Pr(i \, \text{wins} \mid b, \, i \in (x, q); \sigma^{-i})$ the probability of seller i winning the auction by bidding b with the other active sellers' bids consistent with the strategies σ^{-i}. Notice that expression in $\Pr(i \, \text{wins} \mid b, \, i \in (x, q); \sigma^{-i})$ implicitly includes integration over possible sets of active competitors since information about active competitors is not available to seller i at the time when he decides on

[3] We abstract away from sellers' dynamic incentives associated with reputation building. This assumption does not impact the methodology for the estimation of buyers' weights or sellers' qualities. We provide detailed discussion of this issue in the empirical section of this paper when we recover the distribution of sellers' costs.

his bid.

Thus, a psBNE [4] is a profile of strategies $\{\sigma^{r,(x,q)}\}_{r\in\{p,t\},q\in\{q_1,\ldots,q_K\}}$ such that

$$\sigma^{r,(x,q)}(c) = \arg\max_b \Pi^{r,(x,q)}(b, c; \sigma^{-i}) \text{ for all } c.$$

4.4 Empirical Methodology

The large number of buyers and sellers participating in this market, sellers' self-selection into buyer-specific choice sets, as well as potential unobservability of some of the sellers' attributes make use of standard discrete choice techniques impractical in this setting. Instead, we rely on the two-step methodology developed in Krasnokutskaya et al. (2014)[18] which in the first step recovers sellers' groupings according to the unobserved (discrete) characteristic and then in the second step implements GMM procedure to recover sellers' quality levels and the distribution of buyers' weights. We summarize the details of the estimation procedure in this section.

4.4.1 Step 1: Classification Procedure

Recall that each seller i is characterized by a vector of observable characteristics X_i and a scalar measure of quality q_i which is not observed in the data. We allow for the distribution of q to be x-dependent, i.e. to vary with x. Then, the support

[4]See Krasnokutskaya et al. (2014)[18] for the proof of existence of such equilibrium.

of the distribution of qualities among sellers with $X_i = x$ is given by $\{q^k(x) : 1 \leq k \leq K_x\}$ where K_x is the cardinality of the support given x. Characteristics (x, q) induce partitioning of the set of sellers into a finite collection of groups, $\{S^{(x,q)}\}$, corresponding to different values of x and q and such that sellers within a group are characterized by the same value of x and q.[5] The first step of our procedure exploits the differences in the probability of winning across sellers of different groups to recover the group structure of the set of permanent sellers. The classification procedure is implemented conditional on the value of observable characteristic x. That is why in what follows we suppress reference to x and focus only on sellers' differences in q.

In this step, we make use of a non-parametric classification procedure proposed in Krasnokutskaya et al. (2014).[18] This procedure is based on the pairwise testing of inequality restrictions which relies on the proposition below. It exploits differences in the probability of winning across sellers with different levels of the unobserved characteristic.

Intuitively, if i and j participate in two separate but ex-ante identical auctions (in terms of the realized set of competitors) and submit the same price then the seller with the higher value of q has a higher chance of winning. Note that the winner is not deterministic in the presence of uncertainty about buyers' weights. The ranking of winning probabilities is preserved when aggregated over different sets of competitors as long as the probability of encountering a given set of competitors is the same for

[5]Such partition arises naturally when x and q are discrete. It requires discretizing xs if they are continuous.

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

To formulate the result more formally we need the following two assumptions:

(A1) *Sellers' private costs $C_{i,l}$ and the events of being active are independent across all $i \in S$ and across l. For each seller i with $q_i = q^k$, his cost in each auction is an independent draw from continuous distribution F_C^k with a density positive over support $[\underline{c}_k, \overline{c}_k]$.*[6] *The events of being active are independent across projects and sellers.*

(A2) *The three random vectors $(\alpha_l, U_{0,l})$, ϵ_l and C_l are mutually independent; match components $\epsilon_{i,l}$ are i.i.d. across i's; and $\epsilon_{i,l}$ and $(\alpha_l, U_{0,l})$ are continuously distributed with a density positive over $[\underline{\varepsilon}, \overline{\varepsilon}]$ and over $[0, \overline{\alpha}] \times [\underline{u}_0, \overline{u}_0]$ respectively.*[7]

For the remainder of this section we drop subscript l (rhw index for auctions/buyers) to simplify notation. Let $\mathcal{B}_i$ denote the support of prices submitted by seller i in a psBNE. For any $b \in \mathcal{B}_i \cap \mathcal{B}_j$, define a pair-specific index:

$$r_{i,j}(b) \equiv \Pr(i \,\text{wins} \mid B_i = b,\ i \in A,\ j \notin A). \tag{4.3}$$

[6]This assumption does not allow for a persistent unobserved seller-specific cost component in addition to quality. This excludes, for example, differences in opportunity costs associated with sellers' location (e.g. urban vs. rural) if it is not observed in the data. It might be possible to separately account for this type of unobserved seller heterogeneity since our current strategy identifies unobserved quality from buyers' choices whereas unobserved cost persistence might be identified from the additional correlation in prices (unaccounted for by quality) over time. We leave this extension to future research.

[7]Notice that we require that ϵ_l is orthogonal to $(\alpha_l, U_{0,l})$, whereas α_l and $U_{0,l}$ are allowed to be dependent. Such restriction appears to be plausible since we can think of α_l and $U_{0,l}$ as buyers' permanent tastes whereas ϵ_l characterizes active sellers' idiosyncratic suitabilities for the project which should not be related to buyer's outside option.

This index reflects the probability that seller i wins an auction when submitting bid b and when the set of his direct competitors does not include j. The proposition below establishes pairwise ranking of bidders i and j on the basis of indices $r_{i,j}(b)$ and $r_{j,i}(b)$.[8,9]

Proposition 1 *Under (A1)-(A2),*

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

for any pair of permanent sellers i, j and all b in the interior of $\mathcal{B}_i \cap \mathcal{B}_j$.[10]

Since in our model sellers' ordering with respect to q is transitive this result applied to a sufficiently large dataset allows arranging all sellers in the order of (weakly) increasing quality. In other words, we are able to identify the quality group structure and group affiliations of permanent sellers.

The main issue that we need to overcome in order to translate this identification strategy into a viable estimation method is that, while ordering with respect to q is transitive in our model, the estimation based on pairwise comparisons may result in estimates that violate transitivity in small samples, even when there are only two quality groups in the population.

[8] The index with restriction $\{i, j \in A\}$ is not monotone in bidders' quality. In fact, under such restriction the ranking of $r_{i,j}(b)$ and $r_{j,i}(b)$ depends on the distributions of buyers' weights.

[9] Proposition 1 also holds if we relax (A2) to allow dependence between α and ϵ and only require ϵ_i to be independent conditional on α.

[10] Here $sign(x) \equiv 1\{x > 0\} - 1\{x < 0\}$ for all $x \in \mathbf{R}$.

The nonparametric classification procedure we use proposes a method to estimate the whole group structure at once in a way that satisfies transitivity. Below, we provide a heuristic summary of how this is achieved in a simple case of two groups (corresponding to high and low values of q).

The idea is to divide the set of sellers into two groups such that sellers within each group are "closer" to each other than to sellers from the other group according to some metric which is based on index $r_{i,j}$. More specifically, for each seller i, we first divide the other sellers into two groups, one with sellers more likely to have higher quality than i and the other with sellers more likely to have lower quality than i. This division is implemented by comparing p-values from a pairwise bootstrap test of the inequality restrictions $r_{i,j}(b) \geq r_{j,i}(b)$ for all b. Next, we check whether seller i is more likely to belong to the first group or to the second group. Thus for each seller i, we estimate a separate group structure. We choose one of these structures to be our estimate of the underlying quality group structure so that the chosen structure has most empirical support (in terms of average p-values). The formal exposition of the classification method for a more general case of multiple quality groups can be found in Krasnokutskaya et al. (2014).[18]

The true number of the quality groups is usually unknown. Thus in estimation we use a consistent group number selection procedure which utilizes the following regularity.[11] When the sample size is large, misspecifying the number of quality

[11]The fact that the number of groups is recovered from data ensures that the assumption of quality discreteness is not overly restrictive. Indeed, any continuous distribution can be approximated

groups to be smaller than the true number of groups results in the weak empirical support of seller homogeneity for some of the estimated groups. On the other hand, when the number of quality groups is misspecified to be larger than or equal to the true number of groups, the group estimation does not show any sign of misspecification bias.

Note that once the quality group affiliations of sellers are known their identities are no longer important. For example, to condition on the realization of the set of active bidders no longer means to fix the list of bidders identities. Rather, it means to require that the set of active bidders should include specific numbers of sellers by type and quality group. Note that the quality group affiliation is not publicly observed for transitory sellers and thus only information on the overall number of transitory active bidders should be specified. In what follows we will use notation I_A to reflect such information about a given set of sellers A.

4.4.2 Step 2: GMM Estimation

The estimation is based on two sets of moment conditions. The first set of moments relates the probability that a permanent seller wins when the sets of actual and potential permanent bidders satisfy certain restrictions. The second set links

by a sequence of discrete distributions with finite supports. Therefore, one can obtain as good approximation of the continuous distribution of qualities by a discrete random variable as information in the data would allow if the support of the discrete random variable is not restricted. Note that modeling unobserved heterogeneity using a discrete distribution is common in empirical studies. For examples, see Heckman and Singer (1984),[33] Keane and Wolpin (1997),[34] and Crawford and Shum (2005)[35] to name just a few.

transitory and permanent sellers' empirical distribution of bids and participation frequencies to their theoretical counterparts.

The first set of moments consists of three subsets:

(1a) Moments that are based on permanent seller's probability of winning in an auction where two or more active permanent bidders belong to the same group. In these moment conditions, we compute expectations of the following functions: a constant (equal to one), the difference between the winning bid and a bid submitted by a permanent bidder from the same group, and the squared difference between the winning bid and a bid submitted by a bidder from the same group.

(1b) Moments that are based on the permanent seller's probability of winning in an auction where he competes with one or more active permanent bidders belonging to a different group. In these moment conditions, we compute expectations of the following functions: a constant (equal to one), the difference between the winning bid and a bid submitted by a seller from a different group, the squared difference between the winning bid and a bid submitted by a seller from a different group, respectively. We include moments for all possible pairs of different groups.

(1c) Moments that are based on the permanent seller's probability of winning in an auction where he competes with one or more active permanent bidders belonging to a different group, and at least one transitory active bidder belonging to a specific country group. In these moment conditions, we compute expectations of the following functions: the product of transitory bid and the differences between the winning bid and the bid of a permanent seller from a different group, the product of transitory bidder's characteristics other than price (e.g., the number of available scores, and a current average of scores) and the differences between the winning bid and the bid of a permanent seller from a different group. We include moments for all possible pairs of different permanent sellers' groups and all possible country groups of transitory sellers.

We supplement moment conditions in the first set by the following moment condition:

(1d) The probability that project is not allocated.

The second set of moments matches the following empirical moments to their theoretical counterparts: the mean and the variance of the transitory bid distributions, as well as the frequencies with which potential transitory sellers submit a bid aggregated to the level observed in the data. We include such moments for every permanent seller group, or correspondingly every combination of transitory seller's country group, the number of his ratings and the current average of his reputation scores. We additionally imposed the expected profit conditions that summarize the optimal participation decision of transitory bidders for each group of transitory sellers. These conditions impose the restriction that in equilibrium only potential bidders with entry costs below the ax-ante expected profit should participate.[12]

4.5 The Data

The data covers the initial 6 years of the company's operations in the online market for programming services, and contains information on close to 600,000 projects involving participation of around 50,000 different sellers. For every project, we observe the type of work, the approximate size of the project, the time requirements, and the geographic location of the buyer. We also observe all submitted bids and the geographic locations of all bidders, as well as the identity of the winner and measures of the winner's subsequent performance, such as arbitration or delays.

The projects range in type and size from very small (below $100) to larger ones (above $1,000), and typically last fewer than three weeks. Projects can include specific requirements, such as a need to use a particular programming language. The median number of bids submitted per project is six. The majority of buyers are one-time participants, with fewer than 2% of them returning for another project. Therefore, buyers rarely interact with the same seller more than once. Krasnokutskaya et al. (2014)[18] provide additional details describing the dataset.

Due to sample size considerations, we focus on graphics-related programming projects of medium size (between $200 to $700). Projects of this type involve pro-

[12]The identifying assumption here is that the distribution of the costs of entry is the same for all groups.

gramming computer games, computer-generated animation, and media-related programming. For each seller we discard the first year of his tenure and only use observations that correspond to the later years of his career with the online market.

Following the methodology developed in Krasnokutskaya et al. (2014),[18] we categorize *permanent* sellers as those with tenure longer than one year and *transitory* sellers as those with tenure less than one year. Tenure is defined as the length of time elapsed between the dates of seller's first and last bid. We apply the classification algorithm to allocate permanent sellers into unobservable quality groups. The classification algorithm allows sellers to differ in quality, with the buyers assumed to be able to observe only permanent sellers' levels of quality.

Sellers' observable characteristics used in the empirical analysis consist of country affiliation and average score (for permanent bidders only). Seller's country affiliation may provide a measure of convenience due to time difference, language proficiency, or work attitude. Sellers are grouped by geographic proximity and similarity of language and economic conditions into the following seven groups: North America (USA and Canada), Latin America, Western Europe, Eastern Europe, Middle East and Africa, South and East Asia, Australia (grouped with New Zealand). Majority of sellers originates from either North America, Eastern Europe or South and East Asia, making these three country groups the focus of our analysis.

In a similar approach buyers are also grouped by geographic proximity and the analysis focuses on the seven largest groups: North America (USA and Canada), United Kingdom, Western Europe[13], Southern Europe[14], Eastern Europe[15] , Oceania (Australia and New Zealand), and South and East Asia (India, Pakistan, and Singapore).

The long-run average reputation score for a permanent bidder is determined from his history of reputation scores, with the individual scores ranging between 0 and 10. Sellers are divided into three groups: average score less than 9.7 (low score), average score above 9.7 and below 9.9 (medium score), and average score above 9.9 (high score). This definition allocates approximately 30%, 30%, and 40% across the three seller groups.

The data is characterized by significant amount of international trade, with buyers frequently procuring services from sellers residing outside their own geographic locations. The welfare gains from trade accrued to U.S. buyers in this market have previously been examined in Krasnokutskaya et al. (2014),[18] finding relatively large gains due to access to lower-cost providers, and smaller gains due to access to sellers of different quality levels. However, these results do not directly translate into the overall welfare gains in the international context. Not only do buyers from various geographic regions differ in price sensitivity and the value of their outside option but their access to sellers of different quality levels in their national markets varies

[13]Western Europe consists of: Austria, Belgium, Denmark, Finland, Germany, Iceland, Ireland, Luxembourg, Malta, Netherlands, Norway, Sweden, and Switzerland.

[14]Southern Europe consists of: France, Greece, Italy, Portugal, and Spain.

[15]Eastern Europe consists of: Belarus, Bosnia and Herzegovina, Bulgaria, Croatia, Czech Republic, Estonia, Hungary, Latvia, Lithuania, Moldova, Poland, Romania, Russia, Serbia, Slovakia, Slovenia, and Ukraine.

substantially.

We therefore take advantage of the market design and rich dataset to explain and decompose clustering in international trading patterns by controlling for the country-specific differences in the quality of sellers, as well as the country-specific differences in buyers' price sensitivity and their outside options. The dataset used in the analysis includes 11,170 projects, for which 73,303 unique bids were submitted.

4.6 Descriptive Analysis

The following summary tables document the observed patterns in international trade among buyers and sellers from different geographic regions and provide an initial insight into the forces driving these patterns.

Tables 4.1 through 4.3 provide a summary overview of the bidding patterns across buyer-seller country pairs. Table 4.1 considers the bidding patterns from the perspective of buyers, and shows the percentage of projects advertised by buyers from a particular country that attract at least one bid from a particular seller country. We observe a clear tendency for local trade, with a larger percentage of projects receiving bids from local bidders. For example, 33% of projects submitted by a North American bidder receive at least one bid from a North American seller, while only 20% of projects initiated from South and East Asia do. Similar trend is observed for Eastern European sellers, who participate in over 80% of projects advertised by

Eastern European buyers, while only in 60% of projects from South and East Asia. The tendency for localization is somewhat less prominent with Asian sellers, who participate very frequently on all projects. However, South and East Asian projects still receive their bids relatively most frequently.

This pattern is confirmed when looking at the relative contribution of seller countries to the total volume of bids across buyer countries, presented in Table 4.2. This table emphasizes the prominence of Asian sellers in the market: 62% of all bids come from Asian sellers, 31% from Eastern European sellers, with only the remaining less than 7% coming from North American bidders.

Table 4.1: Percentage of Projects with Bids from a Seller Country Conditional on Buyer's Country

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	33.25	71.23	87.62
	UK	32.37	74.31	87.97
	Western Europe	28.18	75.03	85.91
	Southern Europe	27.30	72.34	88.30
	Eastern Europe	22.33	81.40	89.30
	Oceania	27.84	74.86	90.57
	South and East Asia	20.65	60.14	92.75

Note: Element (c1,c2) shows the percentage of projects submitted by a buyer from country c1 that involve at least one seller from country c2 bidding for the project.

Table 4.3 offers the view of the market from the perspective of sellers, and shows the percentage of projects by buyer country for which sellers submitted bids, relative to the overall frequency of all projects by buyer country. Hence the table shows

Table 4.2: Percentage of Bids Coming from a Particular Seller Country, Conditional on Buyer Country

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	7.1	30.8	62.1
	UK	6.5	31.6	61.9
	Western Europe	6.0	32.7	61.3
	Southern Europe	5.9	30.9	63.2
	Eastern Europe	4.3	30.1	65.6
	Oceania	5.0	31.0	64.0
	South and East Asia	4.9	25.3	69.8
Overall		6.6	31.0	62.4

whether sellers participate more or less frequently on projects from a particular buyer country than would be expected if they chose projects by chance. North American sellers bid for North America and UK projects slightly more frequently but bid for Eastern European and Asian projects much less frequently than would be expected under random participation. Eastern European sellers participate relatively very often on projects from their own region, and relatively less frequently on projects from South and East Asia. Finally, the plentiful Asian sellers participate on projects without much regard to the buyer's country, with slightly increased relative participation on projects by Asian buyers.

Since not every submitted bid is successful, it is informative to examine the pattern of winning. The differences in the rates of winning and bidding can be driven either by differences in sellers' bidding behavior or their success rate. Table 4.4 shows the percentage of winning bids out of all submitted bids for particular buyer-seller country

Table 4.3: Relative Proportion of Projects by Buyer/Seller Country Pairs Conditional on Seller's Country

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	1.05	0.99	1.00
	UK	1.02	1.03	1.00
	Western Europe	0.89	1.04	0.98
	Southern Europe	0.86	1.00	1.00
	Eastern Europe	0.71	1.13	1.02
	Oceania	0.88	1.04	1.03
	South and East Asia	0.65	0.83	1.05

Note: Element (c1,c2) shows the percentage of projects submitted by a buyer from country c1 among all projects for which sellers from country c2 submitted bids, relative to the overall probability of a project coming from a buyer from country c1.

pairs. The rate of success varies considerably from the low of 21% for North American sellers competing for projects from Eastern or Western Europe to the high of 71% for Asian bidders competing for projects originating in Asia. The success rate also varies across buyer countries with the same seller group. The general pattern follows the tendency for localization observed from the bidding patterns, with bids resulting in success more often on projects from a geographically nearby region.

Tables 4.5 and 4.6 further examine the probability of winning from the perspective of buyers and sellers, respectively. First, we present the percentage of projects from a particular buyer awarded to sellers from different countries. Again, there is a distinct localization pattern in which buyers choose to award their projects to sellers from the same region. North American buyers choose North American sellers more frequently

Table 4.4: Percentage of Winning Bids out of All Bids Submitted by Buyer/Seller Country Pairs

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	33.09	47.64	61.99
	UK	26.96	49.81	60.88
	Western Europe	20.83	55.05	60.62
	Southern Europe	25.97	52.94	61.04
	Eastern Europe	20.83	52.57	57.81
	Oceania	23.35	46.45	64.23
	South and East Asia	24.56	46.39	71.48

Note: Element (c1,c2) shows the percentage of winning bids among the bids submitted by sellers from country c2 for projects by a buyer from country c1.

than any other buyer group. Similarly, Eastern European and Asian buyers choose sellers from their own country more frequently than do buyers from other regions. However, because of the prominence of Asian sellers in this market, they receive more than 50% of projects regardless of the project's country of origin.

Table 4.6 provides the proportions of projects awarded to a seller group by a certain buyer relative to the overall percentage of projects awarded to the seller group by any buyer. Again, the pattern confirms the geographic localization of projects, and the buyer's willingness to award projects to sellers located in a nearby region.

Finally, table 4.7 captures the differences in the size of the projects by considering the average winning bid (scaled by the project's median bid) across buyer/seller country pairs. Even though sellers from North America and Eastern Europe win less frequently and fewer total projects than sellers from South and East Asia, their scaled

Table 4.5: Percentage of Awarded Projects by Buyer/Seller Country Pairs Conditional on Buyer's Country

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	11.01	33.94	54.31
	UK	8.73	37.02	53.55
	Western Europe	5.87	41.30	52.08
	Southern Europe	7.09	38.30	53.90
	Eastern Europe	4.65	42.79	51.63
	Oceania	6.50	34.78	58.18
	South and East Asia	5.07	27.90	66.30

Note: Element (c1,c2) shows the percentage of projects awarded to a seller from country c2 among all projects submitted by buyers from country c1.

Table 4.6: Relative Proportion of Projects Won by Buyer/Seller Country Pairs Conditional on Seller's Country

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	1.14	0.96	0.99
	UK	0.91	1.05	0.97
	Western Europe	0.61	1.17	0.95
	Southern Europe	0.74	1.08	0.98
	Eastern Europe	0.48	1.21	0.94
	Oceania	0.68	0.98	1.06
	South and East Asia	0.53	0.79	1.21

Note: Element (c1,c2) shows the percentage of projects submitted by a buyer from country c1 among all projects for which sellers from country c2 submitted winning bids.

average winning bids often exceed those for sellers from South and East Asia. Interestingly, sellers from North America constitute the smallest part of this market but their relative winning bids exceed those of other sellers in all cases except South and East Asia, where their winning bid is actually the lowest. This could be explained by sellers from North America being of higher quality than sellers from Eastern Europe, who in turn are of higher quality than the sellers from South and East Asia. Alternately, these differences could be explained by sellers from some regions bidding on larger projects, or bidding on projects with fewer competitors.

Table 4.7: Average Winning Bid (Scaled by Median Bid) across Buyer/Seller Country Pairs

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	0.99	0.94	0.91
	UK	1.01	0.97	0.92
	Western Europe	1.05	1.01	0.95
	Southern Europe	1.10	0.91	0.95
	Eastern Europe	1.13	0.87	0.95
	Oceania	1.02	0.94	0.88
	South and East Asia	0.84	0.91	0.89

Note: Element (c1,c2) shows the average of the winning bids across projects that involve seller from country c2 winning a project submitted by buyer from country c1.

4.7 Estimation Results

In this section we present results of the structural analysis. We assume that buyer's value depends on seller's group as defined by his observable characteristics (country and long-run average score) and quality. Buyers may in turn differ in the weight they put on quality as well as in the value of their outside option. The stochastic components of buyers' values (ϵ) as well as the stochastic part of outside option are assumed to be distributed according the extreme value distribution. The weight buyers put on quality is distributed according to the normal distribution.[16] We use the results of classification from Krasnokutskaya et al. (2014).[18] They are obtained on the basis of the projects auctioned by US buyers that are in $400-$700 range in terms of the size and have a duration of 2-3 weeks. The details of the estimation as well as the robustness check can be found in Krasnokutskaya et al. (2014).[18] In this paper we maintain the assumption that the classification of the quality groups is uniform across buyer countries and rely on the classification recovered for US buyers.

Table 4.8 reports the estimation results for three specifications. The first specification assumes that the distribution of buyers' weights for quality as well the distribution of their outside options is the same across all buyer countries. The second specification allows the mean of the distribution of the outside option to vary across buyer countries. The last specification additionally allows for the mean of the dis-

[16] To be absolutely correct this distribution should have a positive support. However, the mean and the standard deviation of this distribution are sufficiently small so that the distributional assumption does not cause problems.

tribution of buyers' weights for quality to differ across buyer countries. We would expect the fit to improve as we go from Specification 1 to Specification 3 since if buyers' choices differ across buyer countries in a systematic way then controlling for this observable buyer characteristic should help to rationalize larger fraction of choices. Notice, however, that all three specification assume unilateral tastes, i.e. buyers from a specific country may be willing to pay more or less for quality but not for a specific country of a bidder.

Results of the estimation are broadly consistent across specifications. The estimated variances of unobservables are of reasonable magnitude whereas the estimated mean of the outside option is below the expected value from most of "inside" options. This is not surprising since buyers rarely prefer outside option to the "inside" ones. The specifications that allow for observed heterogeneity in buyers' weights and outside options indicate that non-trivial differences exist across buyer countries. Thus buyers from South Europe, Eastern Europe and South and East Asia appear to be more price sensitive than buyers from other countries. The difference appears to be substantial (recall that the expected price sensitivity for US buyers is fixed at one). The estimated quality levels, while monotone within country-score group, appear to be higher in richer specifications.

Next, we investigate the fit of these specifications to the data. In particular, for every buyer country we compute the share of projects that are allocated to sellers from North America, Eastern Europe, South and East Asia or to the outside option,

Table 4.8: Buyers' Tastes and Quality Levels

Specification			(1)	(2)	(3)
log (σ_ϵ)			fixed=0	−0.069	−0.223
μ_{v_0}			−2.521	−2.098	−9.667
log (σ_{v_0})			log (σ_ϵ)	log (σ_ϵ)	log (σ_ϵ)
μ_α			1.142	fixed=1	fixed=1
log (σ_α)			−0.49	−2.781	fixed=0
$\mu_{\alpha_{c(i)}}$					
UK					−0.504
Western Europe					−0.778
Southern Europe					0.332
Eastern Europe					0.347
Oceania					0.212
South and East Asia					0.98
Buyer-country fixed effects					
North America				1.415	−0.678
UK				−0.68	−10.274
Western Europe				−0.37	0.274
Southern Europe				−0.825	−3.823
Eastern Europe				0.834	8.132
Oceania				−0.34	20.439
South and East Asia				−1.262	−0.142
Quality levels					
North America	low score	Q=1	0.426	−0.077	0.511
North America	low score	Q=2	1.04	0.604	0.977
North America	medium score	Q=1	0.698	0.37	0.666
North America	medium score	Q=2	1.974	1.502	1.732
North America	high score	Q=1	1.284	0.819	1.204
North America	high score	Q=2	1.423	0.739	1.123
Eastern Europe	low score	Q=1	−0.138	−0.039	0.408
Eastern Europe	low score	Q=2	1.027	0.668	1.053
Eastern Europe	medium score	Q=1	0.595	0.306	0.712
Eastern Europe	medium score	Q=2	1.909	1.448	1.723
Eastern Europe	medium score	Q=3	1.919	1.421	1.747
Eastern Europe	high score	Q=1	−0.274	−0.369	−0.217
Eastern Europe	high score	Q=2	1.522	1.057	1.37
Eastern Europe	high score	Q=3	1.888	1.34	1.707
South and East Asia	low score	Q=1	fixed=0	fixed=0	fixed=0
South and East Asia	low score	Q=2	0.864	0.432	0.823
South and East Asia	low score	Q=3	1.835	1.303	1.647
South and East Asia	medium score	Q=1	−0.697	−0.798	−0.203
South and East Asia	medium score	Q=2	0.65	0.267	0.943
South and East Asia	medium score	Q=3	1.572	1.022	1.417
South and East Asia	high score	Q=1	1.244	−0.1	1.218
South and East Asia	high score	Q=2	2.366	1.896	1.827

Note: The results are based on the dataset consisting of 11,170 projects. The quality level for South and East Asia, low score, Q=1, is normalized to be equal to zero.

respectively. The results for each specification are reported in Tables 4.9 through 4.11, respectively, which are analogous to Table 4.5 which is based on the observed data. As expected, the fit improves from Specification 1 to Specification 3. Interestingly, Specifications 1 and 2 tend to underpredict the allocation towards sellers from South and East Asia in favor of the outside option. Specification 3 does quite well with the exception of the cell corresponding to UK buyer's choices. It somewhat underpredicts allocation of projects towards Eastern European or South and East Asian sellers in favor of the outside option. However, overall the fit appears to be quite good: an average discrepancy per cell is one percentage point.

Table 4.9: Probability of Winning (Participation Effects Included) Based on Specification 1

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	10.61	34.47	52.82
	UK	9.43	35.42	52.95
	Western Europe	8.13	38.36	51.27
	Southern Europe	8.15	37.24	52.59
	Eastern Europe	6.54	38.98	52.98
	Oceania	7.74	35.20	55.17
	South and East Asia	6.67	27.00	63.88

The subsequent analysis relies on re-solving the model under a variety of circumstances. For this, we need to recover the distribution of sellers' private project costs as well as the entry costs. Our methodology for performing this step is explained in the following section.

Table 4.10: Probability of Winning (Participation Effects Included) Based on Specification 2

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	10.84	37.85	50.94
	UK	8.90	36.67	45.96
	Western Europe	7.95	39.95	46.28
	Southern Europe	7.66	37.56	45.05
	Eastern Europe	6.85	42.71	49.79
	Oceania	7.63	37.74	49.62
	South and East Asia	5.71	25.00	50.74

Table 4.11: Probability of Winning (Participation Effects Included) Based on Specification 3

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	10.26	34.55	55.19
	UK	6.70	27.23	39.03
	Western Europe	8.23	39.36	52.41
	Southern Europe	7.62	37.50	54.88
	Eastern Europe	6.44	39.20	54.36
	Oceania	7.39	34.74	57.87
	South and East Asia	6.26	25.22	68.52

4.8 Recovering the Distributions of Sellers' Private Costs

We recover the distributions of the sellers' costs conditional on sellers' attributes by combining the bid distributions of permanent sellers with the corresponding inverse bid functions:

$$F_C(c|(q,x)) = F_B(\xi^{-1}(c|(q,x))|(q,x)).$$

Here the inverse bid function, $\xi(b|(q,x))$, is derived from the first order condition of the corresponding permanent seller's optimization problem:

$$\xi(b_i|(q,x)_i) = b_i - \frac{P(i\,\text{wins} \mid b_i; \sigma^E_{-i}, \sigma^B_{-i})}{\frac{\partial}{\partial b}P(i\,\text{wins} \mid b; \sigma^E_{-i}, \sigma^B_{-i})|_{b=b_i}}.$$

Next, we assess the magnitude of the cost of entering the auction using the model of strategic participation. Under this model, the observed probability of participation satisfies the following equation

$$F_E(E[\pi^p(x,k)]) = \Pr(i \in A^{p,(x,k)}),$$

where $F_E(.)$ is the distribution of the entry costs and $\pi(x,k)$ is the ex-ante expected profit. We estimate the mean and standard deviation of the entry costs distribution by fitting truncated normal distribution (truncated at 0) to the set of points implied by the ex-ante expected profit and the probability of participation values for various covariate cells and quality groups.

The recovered inverse bid functions and the distributions of sellers' private costs are depicted in Figures 4.1 and 4.2, respectively.

The estimates indicate substantial heterogeneity in sellers' costs across countries. The estimated project cost distributions 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. This indicates substantial cost heterogeneity unrelated to the quality that characterizes the participants in this market.

The estimated value for the mean and standard deviation of the entry costs are 0.082 and 0.077, respectively. That is, entry costs roughly correspond to 7% of the project cost on average. This number is slightly higher than that documented in other markets.[17] The relatively large entry costs estimated in this market may reflect 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 seller's reputation score as given and ignore the possible dynamic considerations associated with reputation building. To mitigate this

[17]Studies of the US highway procurement market have estimated entry costs to be around $2-5\%$ of the engineer's estimate.

Figure 4.1: Estimated Bid Functions

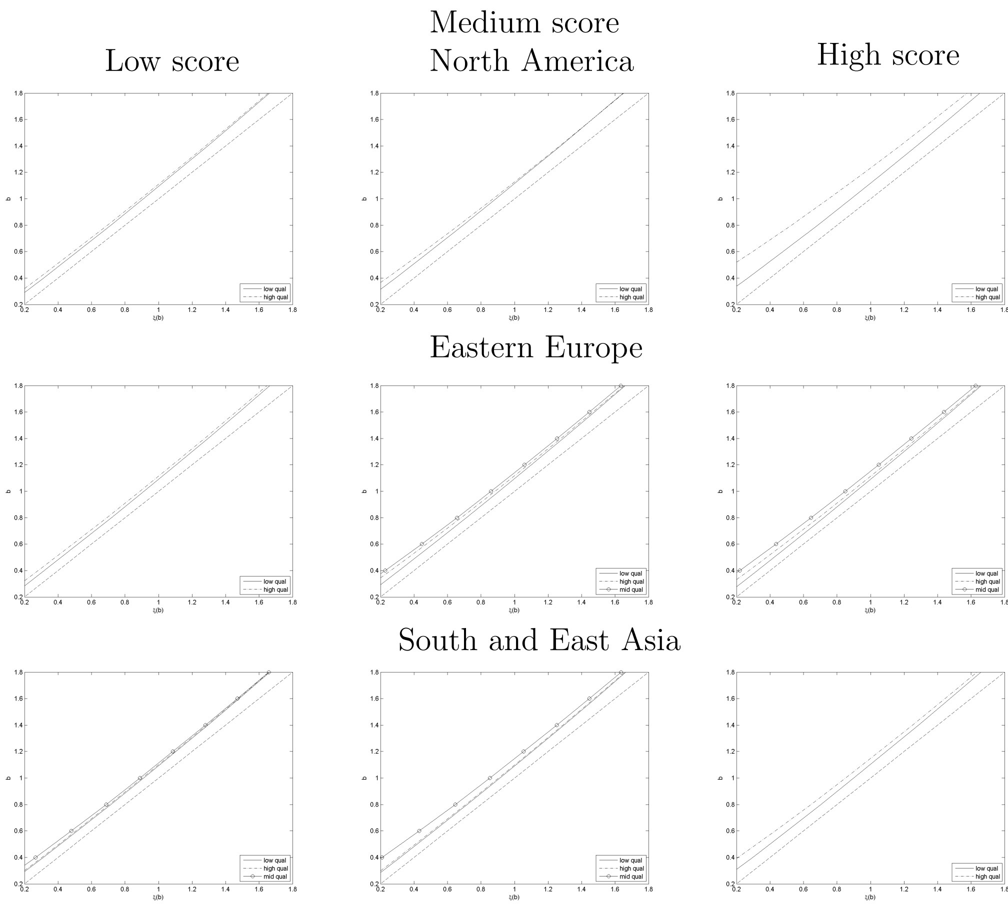

Figure 4.2: Estimated Density of Seller's Cost Distributions

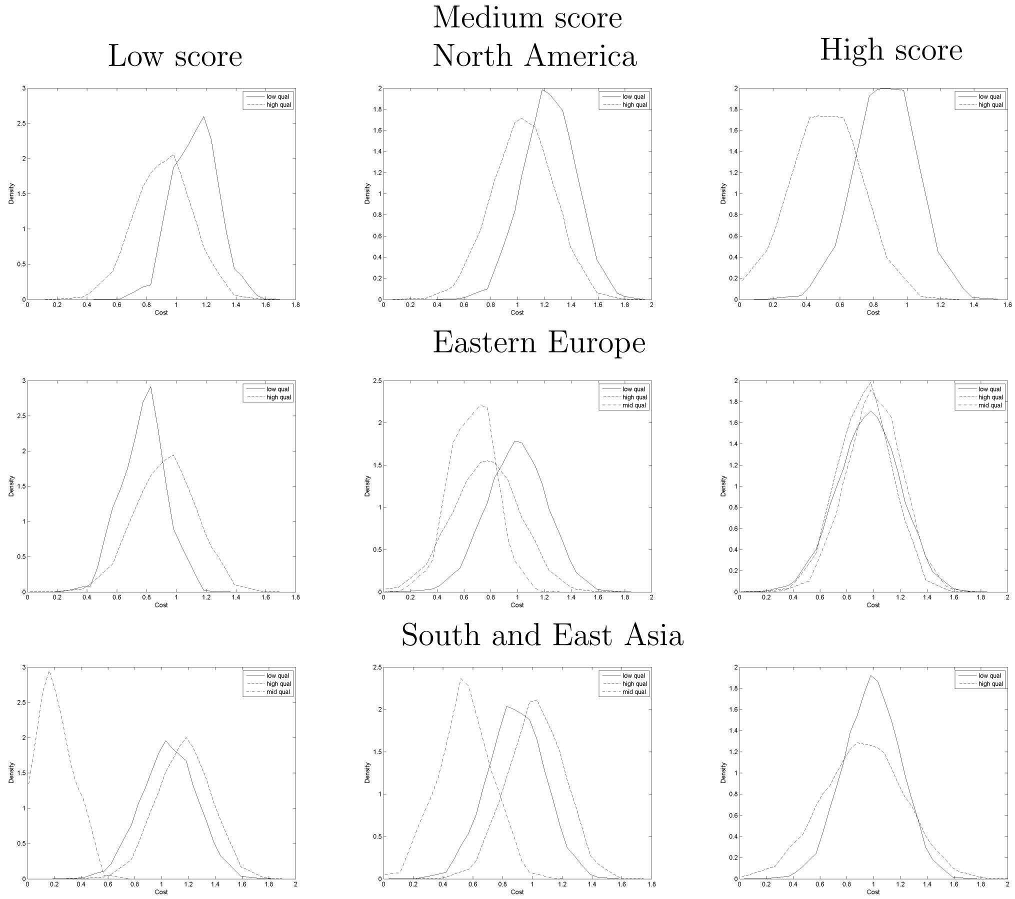

concern, we base our estimation of the distribution of sellers' costs on the optimization problem of a permanent seller. While permanent sellers may still take reputation-related concerns into account, the incentives associated with these concerns are likely to be quite weak. A single score does not make a large impact on average reputation score once a seller has completed ten 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.

4.9 Analysis of Clustering Decomposition

Next, we investigate the buyer-seller country pair trading patterns. We base this analysis on Specification 3 from Table 4.8.

Clustering in the Data

First, we document clustering in the data by computing the average deviation of the shares of projects allocated to different seller countries from the average allocation across buyer countries. This computation is reported in the first column of Table 4.15. We find that the largest deviations correspond to Western European, Eastern European, and South and East Asian buyers, and somewhat less pronounced for the North American buyers. As Table 4.13 demonstrates, North American buyers show slight preference for North American sellers away from Eastern European or South East Asian sellers, Western European buyers show preference toward Eastern

European sellers away from South East Asian sellers, Eastern European buyers towards Eastern European sellers away from South and East Asian sellers, and buyers from South and East Asia for South and East Asian sellers away from East European sellers. Overall, the bias appears to be regional with the exception of UK and Southern Europe which do not exhibit the regional bias.

Clustering generated by the model

We perform similar calculations using the fit table for Specification 3, with the results presented in Table 4.14. We find that Specification 3 is capable of generating 75% of the distortion away from the average pattern for all country groups except for UK and Eastern Europe. In the case of UK the model predicts excessive allocation towards the outside option at the expense of Eastern European and South East Asian sellers; whereas in the case of Eastern European buyers the model cannot generate sufficient bias towards Eastern European sellers. Interestingly, the model overpredicts the bias in the case of South and East Asian buyers – it indicates that according to their tastes they should be even more biased than they actually are. It is worth emphasizing that the model does a good job reproducing more than 75% of the clustering pattern without relying on bilateral buyers' preferences or country-pair-specific sellers' costs.

Removing participation effects

Our next step is to separate clustering associated with sellers' self-selection into participation from the clustering driven by the differences in buyers' preferences for seller's attributes as well as sellers' heterogeneity in those attributes. We achieve this through the following exercise: We generate a new set of participants for every auction by randomly reallocating sellers who have submitted bids in any given week to other projects auctioned off in this week; we then recompute bids that sellers reallocated to a different project would have submitted in the "new" auctions, and determine the winner under this new set of participants. The aggregated winning patterns under random participation are presented in Table 4.12, which is again analogous to Table 4.5 but purging the participation effects. To simulate the "new" bids we extend the numerical algorithm proposed in Marshall et al. (1994)[36] to our setting.

Table 4.12: Probability of Winning (After Removing Participation Effects) Based on Specification 3

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	7.72	33.16	59.12
	UK	4.14	15.75	26.72
	Western Europe	7.65	33.97	58.38
	Southern Europe	6.81	34.17	58.77
	Eastern Europe	6.16	34.21	59.63
	Oceania	6.16	30.10	63.74
	South and East Asia	6.11	25.83	67.02

We repeat the analysis of clustering decomposition for the case when clustering is defined relative to the shares of sellers' countries with respect to the number of all bids submitted, regardless of whether they are successful or not. The corresponding results are presented in Tables 4.16 through 4.18, starting with clustering in the data,

Table 4.13: Clustering Bias Relative to the Average Shares of Awarded Projects (Based on Data)

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	0.040	-0.026	-0.014
	UK	0.017	0.004	-0.022
	Western Europe	-0.011	0.047	-0.036
	Southern Europe	0.001	0.017	-0.018
	Eastern Europe	-0.023	0.062	-0.041
	Oceania	-0.005	-0.018	0.025
	South and East Asia	-0.019	-0.087	0.106

Note: Clustering bias is computed relative to the average market share of awarded projects by seller's country, with the averages computed across buyer countries.

Table 4.14: Clustering Bias Relative to the Average Shares of Awarded Projects (Based on Specification 3)

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	0.033	-0.020	-0.005
	UK	-0.003	-0.093	-0.167
	Western Europe	0.012	0.028	-0.033
	Southern Europe	0.006	0.009	-0.008
	Eastern Europe	-0.006	0.026	-0.013
	Oceania	0.004	-0.018	0.022
	South and East Asia	-0.007	-0.114	0.128

Note: Clustering bias is computed relative to the average market share of awarded projects by seller's country, with the averages computed across buyer countries.

Table 4.15: Average Deviation of the Shares of Allocated Projects

		Data	Model	Diff (%)
	North America	0.050	0.039	0.225
	UK	0.028	0.191	-5.809
Buyer	Western Europe	0.061	0.045	0.259
Country	Southern Europe	0.025	0.014	0.445
	Eastern Europe	0.078	0.030	0.615
	Oceania	0.031	0.029	0.073
	South and East Asia	0.138	0.171	-0.240

Note: Presented are the average deviations of the shares of projects allocated to different seller countries from the average allocation across buyer countries. Results are based on specification 3.

followed by clustering generated by the model with and without participation effects. For the clustering in the data, we find that Eastern European sellers tend to win more frequently while South and East Asian sellers win less frequently than their respective shares would suggest; this pattern is consistent across all buyer countries with the exception of the South and East Asian buyers, in which case the pattern is reversed, with South and East Asian buyers winning disproportionably more often, while Easter European seller winning relatively less often.

As before, we find that the model explains more than 90% of the clustering bias for most countries with the exception of UK and Eastern Europe. In the case of UK buyers, the model underpredicts clustering across all buyer countries, while for Eastern European buyers, frequency of winning is underpredicted for Eastern European sellers and overpredicted for the other sellers. The model reproduces the overall patterns of clustering for the Eastern European and South and East Asian sellers, with the former winning relatively more frequently than predicted by their share of bids, while the latter are winning less frequently; the model also captures the reversal of this pattern when the buyer originates from South and East Asia.

Finally, accounting for the selection into participation and removing the participation effect by reallocating bidders into auctions randomly within the same week generates approximately half of the bias, with the exception of UK, where the clustering bias is underpredicted even further. Thus it appears that selection into participation contributes significantly to generation of clustering pattern.

Table 4.16: Clustering Bias Relative to the Shares of Submitted Bids (Based on Data)

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	0.044	0.029	-0.081
	UK	0.021	0.060	-0.089
	Western Europe	-0.007	0.103	-0.103
	Southern Europe	0.005	0.073	-0.085
	Eastern Europe	-0.020	0.118	-0.108
	Oceania	-0.001	0.038	-0.042
	South and East Asia	-0.015	-0.031	0.039

Note: Clustering bias is computed relative to the overall market share of submitted bids by seller's country.

Decomposing Buyer's Choice

In this section we decompose buyers' choices to identify the mechanism that drives clustering in the model. Specifically, we homogenize sellers from different countries with respect to the importance of their average reputation scores, their quality levels,

Table 4.17: Clustering Bias Relative to the Shares of Submitted Bids (Based on Specification 3)

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	0.037	0.036	-0.072
	UK	0.001	-0.038	-0.234
	Western Europe	0.016	0.084	-0.100
	Southern Europe	0.010	0.065	-0.075
	Eastern Europe	-0.002	0.082	-0.080
	Oceania	0.008	0.037	-0.045
	South and East Asia	-0.003	-0.058	0.061

Note: Clustering bias is computed relative to the overall market share of submitted bids by seller's country.

Table 4.18: Clustering Bias Relative to the Shares of Submitted Bids after Removing the Participation Effect (Based on Specification 3)

		Seller Country		
		North America	Eastern Europe	South and East Asia
Buyer Country	North America	0.011	0.022	-0.033
	UK	-0.025	-0.152	-0.357
	Western Europe	0.011	0.030	-0.040
	Southern Europe	0.002	0.032	-0.036
	Eastern Europe	-0.004	0.032	-0.028
	Oceania	-0.004	-0.009	0.013
	South and East Asia	-0.005	-0.052	0.046

Note: Clustering bias is computed relative to the overall market share of submitted bids by seller's country.

Table 4.19: Average Deviation of the Shares of Submitted Bids

		Data	Model	Diff (%)
	North America	0.097	0.088	0.086
	UK	0.109	0.237	-1.169
Buyer	Western Europe	0.146	0.131	0.101
Country	Southern Europe	0.112	0.100	0.110
	Eastern Europe	0.161	0.115	0.286
	Oceania	0.057	0.059	-0.046
	South and East Asia	0.052	0.084	-0.617

Note: Presented are the average deviations of the shares of submitted bids by different seller countries from the average allocation across buyer countries. Results are based on specification 3.

and their cost distributions. At each step, we first re-solve the bidding functions under the new set of primitives and then recompute the market shares of sellers' countries for each buyer country. In this exercise we maintain participation patterns in terms of the probability of a seller from country c_1 submitting a bid to an auction held by a buyer from country c_2, as well as the probability of a seller from quality group q_1 competing against a seller from quality group q_2 as they are observed in the data. Results of this analysis are reported in Tables 4.20 through 4.23.

Note that these simulation results differ from the model fit tables presented above, as they are based on a smaller number of bidders per auction. Due to computational constraints, we are not resolving the equilibria for all possible configurations of bidders.

Taking the distribution of average scores, quality levels and costs as they are recovered from the data, we simulate the market shares of sellers from different countries for each particular buyer, as shown in Table 4.20. Similar to its empirical counterpart, Table 4.5, these results point to a distinct clustering pattern. More importantly, the findings from the simulations are consistent with the analysis of the patterns of clustering bias discussed above, where Eastern European sellers won more frequently than the relative share of their bids (or awarded projects) on auctions allocated by buyers from every country except South and East Asia, while South and East Asian sellers experience the opposite effect.

The deviation from the average market shares (where the average is taken over buyers' countries) is strongest for South and East Asian buyers, as well as Eastern European, Western European and buyers from the UK. Notice that the magnitudes are not exactly the same, as the participation pattern is somewhat different from the data. The deviation from the average shares mostly occurs due to the difference in allocation between Eastern European and South and East Asian sellers across buyer countries. In particular, UK and Western Europe tend to allocate to Eastern European sellers more frequently than their average market shares would suggest at the expense of Asian sellers; whereas Asian buyers allocate more frequently to Asian sellers and outside option at the expense of Eastern European sellers. This is consistent with the reversal of the clustering pattern for the case of Asian buyers that we have observed in our earlier analysis of the model predictions.

In the first step presented in Table 4.21 we homogenize the contribution of the average reputation score to seller's quality across countries. Homogenization of reputation scores actually strengthens the clustering pattern for the Eastern European and Asian buyers while leaving deviation from average shares for other buyer countries unaffected. The main pattern reflects reallocation of some projects away from Asian and towards Eastern European sellers with the exception of Eastern European buyers, who reallocate projects towards North American sellers and away from Eastern European sellers; and Asian buyers who reallocate projects from Eastern European sellers and outside option towards Asian sellers.

Next, in addition to the average score we also homogenize quality levels across seller countries. The results of this analysis are reported in Table 4.22. The changes in allocations are small and mostly due to the movement from the outside option to North American or Eastern European sellers (as in UK and Western Europe); from Eastern European towards Asian sellers (for South and Eastern Europe); or from Asian sellers towards North American or Eastern European sellers. Thus homogenization of qualities works to reduce clustering for Eastern European and Asian buyers but the effects are rather small.

Finally, in the last step we homogenize the distributions of costs conditional on quality and average score across sellers countries. This induces large reallocation of projects away from the clustering pattern. Thus, it appears that clustering is primarily generated by the differences in costs which translate into differences in prices. Heterogeneity of sellers' cost combined with buyer heterogeneity in price sensitivity is the driver generating the clustering pattern.

Table 4.20: Percentage of Projects Awarded to Seller from a Given Country Conditional on Buyer's Country (Based on Simulation)

		Seller Country			
		North America	Eastern Europe	South and East Asia	Outside Option
Buyer Country	North America	11.68	44.48	39.64	4.19
	UK	11.75	44.63	39.27	4.35
	Western Europe	11.50	45.02	39.11	4.37
	Southern Europe	12.24	42.53	41.41	3.82
	Eastern Europe	9.63	43.02	42.82	4.54
	Oceania	11.66	43.89	40.44	4.00
	South and East Asia	10.19	39.81	43.16	6.84

Note: Element (c1,c2) shows the percentage of projects awarded to a seller from country c2 among all projects submitted by buyers from country c1. Results are based on simulation.

Table 4.21: Percentage of Projects Awarded to Seller from a Given Country Conditional on Buyer's Country (Based on Simulation after Removing the Effect of the Average Score)

		Seller Country			
		North America	Eastern Europe	South and East Asia	Outside Option
Buyer Country	North America	10.64	45.00	40.50	3.86
	UK	10.68	45.09	40.29	3.94
	Western Europe	10.44	45.38	40.21	3.97
	Southern Europe	11.32	42.83	42.20	3.65
	Eastern Europe	8.83	42.95	44.09	4.13
	Oceania	10.73	44.32	41.13	3.82
	South and East Asia	9.30	38.86	45.94	5.89

Note: Element (c1,c2) shows the percentage of projects awarded to a seller from country c2 among all projects submitted by buyers from country c1. Results are based on simulation after purging differences in the distribution of average scores.

Table 4.22: Percentage of Projects Awarded to Seller from a Given Country Conditional on Buyer's Country (Based on Simulation, after Removing the Effect of Differences in the Average Score and Quality)

		Seller Country			
		North America	Eastern Europe	South and East Asia	Outside Option
Buyer Country	North America	10.72	46.10	39.34	3.84
	UK	10.75	46.21	39.15	3.89
	Western Europe	10.50	46.45	39.15	3.90
	Southern Europe	11.37	43.64	41.32	3.67
	Eastern Europe	8.93	43.82	43.13	4.11
	Oceania	10.83	45.30	40.06	3.81
	South and East Asia	9.47	39.99	44.70	5.84

Note: Element (c1,c2) shows the percentage of projects awarded to a seller from country c2 among all projects submitted by buyers from country c1. Results are based on simulation after purging differences in both the distributions of average scores and quality levels.

Table 4.23: Percentage of Projects Awarded to Seller from a Given Country Conditional on Buyer's Country (Based on Simulation, after Removing the Effects of Differences in the Average Score, Quality, and Cost)

		Seller Country			
		North America	Eastern Europe	South and East Asia	Outside Option
Buyer Country	North America	10.71	45.10	40.57	3.65
	UK	10.78	45.41	40.19	3.65
	Western Europe	10.62	45.60	40.14	3.66
	Southern Europe	10.97	41.27	44.29	3.47
	Eastern Europe	9.00	43.04	44.11	3.87
	Oceania	10.78	43.74	41.92	3.58
	South and East Asia	9.63	39.79	45.31	5.27

Note: Element (c1,c2) shows the percentage of projects awarded to a seller from country c2 among all projects submitted by buyers from country c1. Results are based on simulation after purging differences in the distributions of average scores, quality levels, as well as the distribution of costs.

4.10 Conclusion

This paper contributes to the literature on international trade by investigating bilateral patterns of trade in the online market for programming services. Geographic clustering of trade appears to persist in this online market despite seemingly low transaction costs that should not in theory depend on geographical proximity. We construct a model that rationalizes buyers' allocative decisions in this market through the country-based heterogeneity in buyers' preferences and sellers' attributes. The primitives of the model are recovered by applying novel methodology developed in Krasnokutskaya at al. (2014)[18] that accommodates the features of the data structures typically observed in online markets. We find that the model works well in reproducing clustering patterns observed in the data despite the lack of buyer-seller country pair preference component or buyer-seller country pair specific costs. In particular, we find that clustering is driven primarily by the difference in costs (and thus prices) of sellers and the differences in buyers' price sensitivities. Additionally, we show that clustering induced by the difference in these primitives is further amplified by strategic participation on the part of the sellers.

Bibliography

[1] M. Jofre-Bonet and M. Pesendorfer, "Estimation of a dynamic auction game," *Econometrica*, vol. 71, no. 5, pp. 1443–89, 2003.

[2] J. Balat, "Highway procurement and the stimulus package: Identification and estimation of dynamic auctions with unobserved heterogeneity," 2013, working paper.

[3] J. Groeger, "A study of participation in dynamic auctions," 2013, forthcoming in International Economic Review.

[4] S. Athey, D. Coey, and J. Levin, "Set-asides and subsidies in auctions," *American Economic Journal: Microeconomics*, vol. 5, no. 1, pp. 1–27, 2013.

[5] D. Levin and J. L. Smith, "Equilibrium in auctions with entry," *The American Economic Review*, vol. 84, pp. 585–99, 1994.

[6] J. Roberts and A. Sweeting, "Bailouts and the preservation of competition," 2010, working paper NBER #166650.

[7] V. Saini, “Endogenous asymmetry in a dynamic procurement auction,” *The RAND Journal of Economics*, vol. 43, no. 4, pp. 726–760, 2012.

[8] L. C. Benkard, “Learning and forgetting: The dynamics of aircraft production,” *The American Economic Review*, vol. 90, no. 4, pp. 1034–54, 2000.

[9] E. D. Darr, L. Argote, and D. Epple, “The acquisition, transfer, and depreciation of knowledge in service organizations: Productivity in franchises,” *Management Science*, vol. 41, no. 11, pp. 1750–62, 1995.

[10] J. Marion, “Are bid preferences benign? the effect of small business subsidies in highway procurement auctions,” *Journal of Public Economics*, vol. 91, no. 7-8, pp. 1591–624, 2007.

[11] E. Krasnokutskaya and K. Seim, “Bid preference programs and participation in highway procurement auctions,” *The American Economic Review*, vol. 101, no. 6, pp. 2653–2686, 2011.

[12] P. Jeziorski and E. Krasnokutskaya, “Dynamic auction environment with subcontracting,” 2014, working paper.

[13] A. Wambach, “How to subcontract,” *Economic Letters*, vol. 105, no. 2, pp. 152–155, 2009.

[14] I. L. Gale, D. B. Hausch, and M. Stegeman, “Sequential auctions with subcontracting,” *International Economic Review*, vol. 41, no. 4, pp. 12–35, 2000.

[15] D. P. Miller, "Subcontracting and competitive bidding on incomplete procurement contracts," 2014, forthcoming in the RAND Journal of Economics.

[16] J. Marion, "Sourcing from the enemy: Horizontal subcontracting in highway procurement," 2014, forthcoming in the Journal of Industrial Economics.

[17] D. G. De Silva, G. Kosmopoulou, and C. Lamarche, "Survival of contractors with previous subcontracting experience," *Economic Letters*, vol. 117, no. 1, pp. 1–9, 2011.

[18] E. Krasnokutskaya, K. Song, and X. Tang, "Uncovering unobserved agent heterogeneity through pairwise comparisons," University of Pennsylvania, working paper, 2014.

[19] R. Dornbusch, S. Fischer, and P. A. Samuelson, "Comparative advantage, trade and payments in a ricardian model with a continuum of goods," *The American Economic Review*, vol. 67, no. 5, pp. 829–839, 1977.

[20] J. H. Bergstrand, "The generalized gravity equation, monopolistic competition, and the factor proportion theory in international trade," *The Review of Economics and Statistics*, vol. 71, no. 1, pp. 143–153, 1989.

[21] P. R. Krugman, *Geography and Trade.* Cambridge, MA: MIT Press, 1991.

[22] E. E. Leamer and M. Storper, "The economic geography of the internet age," *Journal of International Business Studies*, vol. 32, no. 4, pp. 641–665, 2001.

[23] J. Tinbergen, *Shaping the World Economy: Suggestions for an International Economic Policy.* New York: The Twentieth Century Fund, 1962.

[24] P. Poyhonen, "A tentative model for the volume of trade between countries," *Weltwirtschaftliches Archiv*, vol. 90, no. 1, pp. 93–100, 1963.

[25] H. Linnemann, *An Econometric Study of International Trade Flows.* Amsterdam: North Holland Publishing Company, 1966.

[26] J. McCallum, "National borders matter: Canada-u.s. regional trade patterns," *The American Economic Review*, vol. 85, pp. 615–623, 1995.

[27] J. Frankel, "Emerging currency blocs," in *The Future of the International Monetary System and Its Institutions*, H. Genberg, Ed. Geneva: International Center for Monetary and Banking Studies, 1994.

[28] J. J. Laffont and J. Tirole, "Using cost observations to regulate firms," *The Journal of Political Economy*, vol. 94, pp. 614–641, 1986.

[29] R. McAfee and J. McMillan, "Competition for agency contracts," *The RAND Journal of Economics*, vol. 18, pp. 296–307, 1987.

[30] Y.-K. Che, "Design competition through multidimensional auctions," *The RAND Journal of Economics*, vol. 24, no. 4, pp. 668–680, 1993.

[31] J. Asker and E. Cantillon, "Properties of scoring auctions," *The RAND Journal of Economics*, vol. 39, no. 1, pp. 69–85, 2008.

[32] S. Greenstein, "Did installed base give an incumbent any (measurable) advantages in federal computer procurement?" *The RAND Journal of Economics*, vol. 24, no. 1, pp. 19–39, 1993.

[33] J. Heckman and B. Singer, "A method of minimizing the impact of distributional assumptions in econometric models for duration data," *Econometrica*, vol. 52, pp. 271–320, 1984.

[34] M. P. Keane and K. I. Wolpin, "The career decisions of young men," *Journal of Political Economy*, vol. 105, pp. 473–522, 1997.

[35] G. Crawford and M. Shum, "Uncertainty and learning in pharmaceutical demand," *Econometrica*, vol. 73, pp. 1137–1173, 2005.

[36] R. C. Marshall, M. J. Meurer, J.-F. Richard, and W. Stromquist, "Numerical analysis of asymmetric first price auctions," *Games and Economic Behavior*, vol. 7, no. 2, pp. 193–220, 1994.

[37] L. Argote, S. L. Beckman, and D. Epple, "The persistence and transfer of learning in industrial settings," *Management Science*, vol. 36, no. 2, pp. 140–154, 1990.

[38] T. Bates and D. Williams, "Do preferential procurement programs benefit minorty business?" *The American Economic Review*, vol. 86, no. 2, pp. 294–7, 1996.

[39] J. H. Bergstrand, "The gravity equation in international trade: Some microeconomic foundations and empirical evidence," *The Review of Economics and Statistics*, vol. 67, no. 3, pp. 474–481, 1985.

[40] D. Besanko, U. Doraszelski, Y. Kryukov, and M. Satterthwaite, "Learning-by-doing, organizational forgetting, and industry dynamics," *Econometrica*, vol. 78, no. 2, pp. 453–508, 2010.

[41] A. Corns and A. Schotter, "Can affirmative action be cost effective? an experimental examination of price-preference auctions," *The American Economic Review*, vol. 89, no. 1, pp. 291–305, 1999.

[42] P. Dasgupta and J. Stiglitz, "Learning-by-doing, market structure and industrial and trade policies," *Oxford Economic Papers, New Series*, vol. 40, no. 2, pp. 246–268, 1988.

[43] A. V. Deardorff, "Determinants of bilateral trade: Does gravity work in a neoclassical world?" in *In The Regionalization of the World Economy*, J. A. Frankel, Ed. Chicago: University of Chicago Press, 1998, pp. 7–32.

[44] E. Guerre, I. Perrigne, and Q. Vuong, "Optimal nonparametric estimation of first-price auctions," *Econometrica*, vol. 68, no. 3, pp. 525–74, 2000.

[45] H. Holzer and D. Neumark, "Assessing affirmative auction," *Journal of Economic Literature*, vol. 38, pp. 483–568, 2000.

[46] V. J. Hotz and R. A. Miller, "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, vol. 60, pp. 497–529, 1993.

[47] V. J. Hotz, R. A. Miller, S. Sanders, and J. Smith, "A simulation estimator for dynamic models of discrete choice," *The Review of Economic Studies*, vol. 61, no. 2, pp. 265–89, 1994.

[48] T. P. Hubbard and H. J. Paarsch, "Investigating bid preferences at low-price, sealed-bid auctions with endogenous participation," *International Journal of Indusrial Organization*, vol. 27, pp. 1–14, 2009.

[49] M. Jofre-Bonet and M. Pesendorfer, "Bidding behavior in a repeated procurement auction: A summary," *European Economic Review*, vol. 44, pp. 1006–1020, 2000.

[50] E. Krasnokutskaya, "Identification and estimation of auction models with unobserved heterogeneity," *The Review of Economic Studies*, vol. 78, no. 1, pp. 293–327, 2011.

[51] P. R. Krugman and A. J. Venables, "Globalization and the inequality of nations," *Quarterly Journal of Economics*, vol. 110, no. 4, pp. 857–880, 1995.

[52] T. Li and X. Zheng, "Entry and competition effects in first-price auctions: Theory and evidence from procurement auctions," *The Review of Economic Studies*, vol. 76, pp. 1397–1429, 2009.

[53] J. Marion, “How costly is affirmative action? government contracting and california’s proposition 209,” *The Review of Economics and Statistics*, vol. 91, no. 3, pp. 503–22, 2009.

[54] M. Pesendorfer and P. Schmidt-Dengler, “Asymptotic least squares estimators for dynamic games,” *The Review of Economic Studies*, vol. 75, pp. 901–928, 2008.

[55] D. G. D. Silva, T. D. Jeitschko, and G. Kosmopoulou, “Stochastic synergies in sequential auctions,” *International Journal of Indusrial Organization*, vol. 23, no. 3-4, pp. 183–201, 2005.

[56] P. Somaini, “Competition and interdependent costs in highway procurement,” 2011, working paper.

Vita

Lucia Tiererova received the B. A. degree *summa cum laude* in Economics and Mathematics from Bates College in 2006. Upon graduation, she was inducted into the Phi Beta Kappa and Sigma Xi honor societies, and was awarded the Stangle Award in Economics for outstanding thesis. In 2008 she enrolled in the Economics Ph.D. program at Johns Hopkins University. Her research is in the field of empirical industrial organization, and in particular she focuses on the empirical study of auctions. She received the George Owen Fellowship as well as the Dean's Teaching Fellowship to independently teach an undergraduate seminar-style course.