Dynamic Oligopoly in Procurement Auction Markets

Dakshina De Silva* Matthew Gentry † Pasquale Schiraldi‡

March 2024

Abstract

This study develops a model of dynamic oligopoly competition in large procurement auction markets and applies this to Texas Department of Transportation (TxDOT) highway construction and maintenance contracts. Forward-looking firms enter, exit, grow, and gain experience while competing in a sequence of low-price, sealed-bid auctions. Firms play a Moment-based Markovian Equilibrium (MME) similar in spirit to Ifrach and Weintraub (2017), rendering analysis tractable even in markets with hundreds of firms. Structural estimates based on TxDOT data from 2000-2012 highlight the significant role of dynamic factors like firm experience in shaping market outcomes. The research intends to evaluate the impact of procurement mechanism design, including reserve price policies and eligibility requirements, through counterfactual analysis of long-run industry responses to alternative mechanisms.

1 Introduction

Suppliers in many procurement markets, such as for highway construction and maintenance contracts, are specialized firms. In such markets, the rules of the procurement mechanism can endogenously affect not only how currently active firms bid in specific auctions, but (in the long run) how many firms enter or exit the industry, as well as the composition of the firms that remain. For example, while the classic Myerson (1981) optimal mechanism may minimize procurement costs in the short run, low reserve prices could also decrease profits and thereby increase long-run industry exit. More fundamentally, in markets where profitability

*Department of Economics, Lancaster University Management School; email: d.desilva@lancaster.ac.uk
†Department of Economics, Florida State University; email: mgentry@fsu.edu
‡Department of Economics, London School of Economics and Political Science and CEPR; email: p.schiraldi@lse.ac.uk

depends, at least in part, on dynamic considerations such as investment or learning by doing, the procurer may reap long-run savings from policies that allow new firms to win contracts and develop expertise, even if these increase short-run procurement costs. Theoretical work has clearly established that mechanism design can have large impacts on long-run industry composition. Very little is known empirically, however, about these long-run channels of procurement policy design.

This study explores long-run industry dynamics, and the role of procurement policy in shaping these, within a major low-bid procurement auction market: that for Texas Department of Transportation (TxDOT) highway construction and maintenance contracts. Using data from January 2000 to December 2013, we explore long-run entry, exit, and firm evolution patterns in the TxDOT market. While the overall number of prime contractors is relatively stable over time, there is substantial churn in the set of active firms, with 10-15 percent of previously active prime contractors exiting and another 10-15 percent entering the market in any given year.[1] Of these new entrants, approximately one third enter through waived projects, while the remainder enter through unwaived projects. More than half of firms entering the prime contractor market through waived projects eventually become prequalified prime contractors, often with a lag of several years between first entry and full prequalification. We then turn to a series of reduced-form regressions exploring how firms' bidding decisions depend on firm experience, firm size, qualification status, and backlog, finding patterns consistent with learning by doing in the TxDOT market.

Motivated by these empirical patterns, we develop a dynamic structural model that embeds a model of spot market competition via low-price auctions within an overarching dynamic oligopoly model. Firms in the industry are forward-looking and differ on four state dimensions: experience, backlog, capacity, and prequalification status. Following Arcidiacono, Bayer, Blevins, and Ellickson (2016), we frame our overarching dynamic model in continuous time, with firms receiving qualification opportunities, bidding opportunities, and

[1]Note that entry and exit from the prime contractor market does not imply firm-level entry and exit, since firms may remain active as subcontractors on TxDOT contracts or in other construction activities.

industry exit opportunities at exogenous Poisson rates. Upon receiving a bidding opportunity, firms observe their completion costs and compete in a low-price sealed-bid auction against other participating bidders. Firm-level state variables evolve based on firm choices, auction outcomes, and exogenous transition events. A pool of potential entrants receive entry opportunities at exogenous Poisson rates and make forward-looking entry decisions, with entering firms remaining active until they exit the industry.

The TxDOT market involves a large number of active firms—on average about 700 per year over our sample. Consequently, a standard Markov Perfect Equilibrium (MPE), which assumes that each firm exactly tracks each rival's state, is both behaviorally implausible and computationally infeasible. We, therefore, adopt the Moment-based Markov Equilibrium (MME) solution concept of Ifrach and Weintraub (2017).[2] Adapted to our context, the key simplifying assumption defining a MME is that, when forecasting its profits from future auctions, each firm tracks only its own state and a vector of aggregate industry moments, not the whole industry state. This provides a vastly more tractable and, in our view, more plausible model than MPE in markets with many firms. Importantly, however, this simplification applies only to future auctions; once paired with a known set of rivals in a specific auction, each firm conditions its bidding strategy on the specific states of each rival bidder. We thereby embed the rich strategic auction-level interactions typical in the empirical auction literature within a MME model of dynamic oligopoly which is tractable even in very large markets. It seems realistic to think that players are likely to only have access to or consider summary statistics, or a subset of monitoring data regarding the state variables in the game when making entry/exit decisions, but account for more information at the bidding stage as they need to monitor only a few players.

Drawing on classic results from the empirical auction literature (e.g. Guerre, Perrigne, and Vuong (2000)), we show that our model yields a remarkably simple conditional choice probability (CCP) representation of MME continuation values. We estimate our model using

[2]Other recent papers also make use of this approach, see for example Jeon (2022); Corbae and D'Erasmo (2021); Gowrisankaran et al. (2023); Gerarden (2023); Sears at al. 2022.

TxDOT data from 2000-2012, finding that dynamic firm-level factors such as experience and capacity have substantial impacts on both project completion costs and continuation values. This in turn suggests that procurement design could have important long-run effects. In work in progress, we are implementing several counterfactuals which aim to quantify these long-run effects, including for example changes to TxDOT's prequalification waiver program, implementation of a static Myerson optimal mechanism, and switching from first-price auctions to second-price auctions (noting, as pointed out by Saini (2012), that it is typically harder for weak bidders to win and gain experience in second-price auctions). We emphasize, however, that our model can also be applied to study the long-run industry effects of a wide range of other potential counterfactual policies, almost all of which would be informative since such long-run effects are presently very poorly understood.

This study aims to contribute to a small but growing literature on dynamic auctions, including Jofre-Bonet and Pesendorfer (2003), Groeger (2014), Balat (2015), and Saini (2012) among others. With the exception of Saini (2012), who explores Markov-perfect dynamics in a computational simulation exercise, these prior studies principally analyze bidding behavior among a fixed (exogenous) set of industry participants, accounting for dynamic factors affecting bidding and entry, such as capacity constraints or recent entry experience. In contrast, we aim to analyze the long-run determinants of industry structure, accounting for these factors as well as firm-level entry, exit, and experience. To our knowledge, this long-run industry view has never been applied to a procurement auction market, and will provide new insight into many counterfactual questions that have been unexplored heretofore.

This preliminary draft is organized as follows. Section 2 provides institutional background on the TxDOT market and describes patterns of industry-level entrantsy and exit, auction-level bidding, and firm-level evolution within it. Section 3 develops our dynamic oligopoly model and derives the CCP representation of continuation values which is the key to our empirical strategy. Section 4 provides estimates of structural parameters, continuation values, and completion costs derived from the model. Section 5 concludes with a brief dis-

cussion of intended counterfactuals. Appendix A provides further estimation details, while Appendix B provides detailed variable definitions.

2 Background, data, and descriptive statistics

TxDOT is responsible for building and maintaining all roadways in Texas. TxDOT strives to partner with contractors who achieve high-quality results for the best value to create a safe and reliable transportation system for Texas. Hence, for most projects, TxDOT requires firms to become pre-qualified before holding plans or submitting bids. Pre-qualification is intended to ensure that only financially sound companies with proven track records compete for TxDOT projects, and requires the qualifying firm to provide an independently audited financial statement along with a completed 'Confidential Questionnaire' and other required supporting documents. TxDOT waives the full prequalification requirement, however, for some small construction, maintenance, and special projects. To compete for 'waived' projects, firms need only complete a simpler 'Bidder's Questionnaire,' with no requirement for a financial audit. Conversations with TxDOT officials suggest that one goal of 'waived' projects is to provide an avenue for currently non-qualified bidders to gain experience bidding for TxDOT projects.

TxDOT announces projects to be let at least 28 days in advance. The project listing specifies a detailed description of the project, the number of tasks needed to complete the project, work type, location of the project, days allowed to complete the project, funding agency (state or federal), and the engineer's cost estimate (ECE) in addition to whether the project is 'waived'. To bid for a project, eligible firms must request plans, after which they may prepare and submit itemized bids. Once the bidding period expires, TxDOT tabulates bids and determines the winning bidder. The winning bidder is determined solely by price—all awards are to the lowest eligible bidder, although TxDOT also reserves the right to reject all bids. After the project is awarded, TxDOT discloses the identities of all

firms that tendered a bid and the amount of each bid for each contract.

2.1 Data and Descriptive Statistics

We have data for all construction projects auctioned by TxDOT from January 1998 to December 2013. For each project, we observe all information in the TxDOT project listing, including in particular the ECE, the type and location of work, and whether or not qualification requirements were waived. We also observe all submitted bids, the identities of the bidding firms, and whether or not each project was ultimately awarded.

We categorize projects into six types based on the material shares of a project. The "Standard Specifications for Construction and Maintenance of Highways, Streets, and Bridges" code book adopted by TxDOT describes the six material groups for projects based on bid items. These six material cost shares are constructed from detailed information on bid items and the project's overall engineering cost estimate. These include: 1) asphalt surface work (i.e., hot-mix asphalt); 2) earth work (i.e., excavation); 3) miscellaneous work (i.e., mobilization); 4) structures (bridges); 5) subgrade (i.e., proof rolling); and 6) lighting and signaling work (i.e., highway sign lighting fixtures). We also categorize projects into five zones identified by TxDOT in order to control for physical features of areas as they require different grades of material to complete a project.

We also construct the following additional firm-level variables. We measure each firm's scope of work outstanding with TxDOT via its backlog, which we construct by summing across the non-completed value of existing contracts. For each firm-project pair, we construct the distance between project location and firm location, which we include as a proxy for equipment transportation costs to the project. We define a firm as entering the TxDOT market immediately before the date of its first observed activity. We define a firm as exiting the TxDOT market if it goes 12 consecutive months without activity, where activity is defined as holding plans, submitting bids, or having non-zero work outstanding (backlog). In this case, we interpret the firm as exiting immediately following its last activity. As a proxy for

firm capacity, we consider the maximum, over the past 18 months, of a monthly capacity usage measure equal to the firm's initial backlog plus the sum of ECEs among projects for which the firm held plans within the month.[3] We normalize this firm-level capacity variable by the sum of capacities for all firms within the month to obtain a proxy for the market share of each firm over time. Finally, we classify each firm as CQ (qualified) or not according to the first date it holds plans on a project with normal qualification requirements. If this is also the date of first activity, we classify the firm as entering with CQ status; otherwise, we classify it as non-CQ until the date of the first non-waived auction for which it holds plans.[4]

Our analysis is based on data from January 2000 onward. We use data from 1998 and 1999 to construct firm-level historical variables such as backlog, activity dates, qualification status, and win counts.[5] Backlog, bids and ECE are expressed in January 2000 constant-dollar values.

Table 1 provides summary statistics for projects let each year according to qualification status. For each year, we report the average number of plan holders, bidders, engineer's cost estimate, relative bid, and winning bid for those contracts. When considering waived projects, we further partition the number of potential and actual bidders by their prequalification status. For most years, unwaived contracts see slightly more plan holders and bidders than waived contracts. There are about four non-prequalified plan holders and one

[3]Under TxDOT rules, a contractor may request project details only for projects whose total ECEs do not exceed pre-qualified bidding capacity less the firm's backlog. This institutional feature motivates us to measure within month used capacity as the firm's backlog plus the sum of ECEs among projects for which the firm held plans within the month. Since each monthly used capacity must be less than the firm's overall capacity, we then take the maximum over 18 months to better approximate actual capacity.

[4]CQ status is granted to a firm once it reaches certain requirements in terms of experience and resources upon submitting a simple documentation. While we do not observe the exact date of qualification, we observe the first time the firm held plans requiring qualification. While under TxDOT rules the certification is only valid for 12 months, we consider qualification, once achieved, to be a persistent state, since the main barrier to qualification is demonstrating the required experience and financial resources rather than the paperwork needed to obtain the formal certification. In our sample, 362 firms were required to requalify, and 343 firms were requalified within three months. The remaining 19 firms qualified within a maximum of eight months.

[5]Note that we have identified 247 out-of-state firms that participate in TxDOT auctions. Among these firms, 214 participate as prequalified firms and, on average, bid on projects worth about $10 million. Since these large firms have a history of winning projects in multiple states, we set their initial win count to 28, corresponding to the average win count among Texas-based prequalified firms. The remaining 32 non-prequalified firms bid on projects worth about $215,000. We consider them inexperienced firms.

Table 1: Summary statistics by project types

Year	Waived projects						Unwaived projects				
	Projects	ECE ($)[a]	Relative		n	$n_{q=1}$	Projects	ECE ($)[a]	Relative		n
			Bid	Win					Bid	Win	
2000	156	195,748.90	1.163	0.979	4.397	3.551	914	2,763,536.00	1.049	0.942	4.544
2001	131	360,249.50	1.146	0.987	4.168	3.687	713	3,425,191.00	1.060	0.953	4.306
2002	111	281,485.10	1.064	0.897	5.730	4.973	651	4,207,879.00	1.006	0.896	5.048
2003	153	184,828.50	1.128	0.923	5.301	4.699	645	5,092,556.00	1.035	0.934	4.859
2004	154	182,233.80	1.116	0.957	4.740	4.292	701	4,637,363.00	1.082	0.982	4.317
2005	129	128,479.20	1.217	1.032	3.891	3.636	755	3,292,533.00	1.169	1.049	3.913
2006	110	171,611.00	1.163	0.990	3.891	3.491	785	3,031,770.00	1.140	1.026	3.908
2007	86	125,947.50	1.172	0.991	4.477	4.140	671	3,333,504.00	1.053	0.928	4.441
2008	44	117,954.10	1.140	0.973	5.636	4.886	500	2,655,836.00	1.076	0.947	5.488
2009	87	142,693.90	1.004	0.818	7.207	6.632	741	2,973,292.00	0.940	0.803	6.744
2010	153	150,868.00	1.145	0.954	5.739	5.510	801	2,646,686.00	1.045	0.923	5.542
2011	138	120,375.20	1.162	1.014	4.978	4.725	652	3,966,039.00	1.089	0.971	5.147
2012	87	103,464.80	1.173	0.978	5.092	4.816	642	2,830,953.00	1.119	1.004	4.533
2013	68	102,144.10	1.176	0.979	5.294	4.927	626	2,843,199.00	1.090	0.992	4.438
All	1,607	177,490.30	1.136	0.964	4.958	4.492	9,797	3,386,572.00	1.061	0.954	4.787

[a]ECE is in January 2000 constant-dollar value.

non-prequalified bidders on average present in a waived contract. On average, the ECE for waived projects is about $177,500 while the ECE for unwaived projects is about $3.39 million. The average relative bid is about 1.136 for waived projects, while, for unwaived projects, the relative bid is about 1.061. However, there is not much difference between waived relative winning bids (0.964) and unwaived relative winning bids (0.954).

In Table 2, we present summary statistics by firm types. During our sample period, we observe 2,483 unique firms in the TxDOT market, who in total submitted 54,864 bids. The average ECE was about $2.9 million, while the average bid relative to the ECE was about 1.07. The winning bid is about 0.96 relative to the ECE. On average, all firms have won about 25 contracts in the past, while their market share is about 0.3 percent. On average, a firm's backlog is worth about $9.8 million, and the distance to a project location from its location is about 234 miles. On average, one can observe firms in the TxDOT market for about 55 months.

We observe 1,906 currently prequalified firms. Note that these firms may initially enter as non-prequalified firms. In our sample period, 348 of 911 non-prequalified firms transitioned into qualified status. As we can see, prequalified firms bid on projects that are more than ten times larger than projects bid by an average non-prequalified bidder. Importantly, on average, prequalified bidders have won about 25 more projects than non-prequalified bidders.

Table 2: Summary statistics by firm types

Variable	Current firm type		
	All	Prequalified	Non-prequalified
	(1)	(2)	(3)
Number of firms	2,483	1,906	911
Number of bids	54,866	54,118	748
Average ECE[a]	2.934	2.974	0.257
	(8.126)	(8.178)	(0.770)
Average relative bid	1.072	1.071	1.091
	(0.245)	(0.242)	(0.363)
Average relative winning bid	0.955	0.955	0.934
	(0.175)	(0.175)	(0.213)
Experience (past win counts)	24.881	26.882	1.165
	(41.059)	(42.132)	(4.288)
Firm's average market share	0.003	0.003	0.00004
	(0.007)	(0.007)	(0.0001)
Average backlog[a]	9.807	10.629	0.067
	(34.692)	(36.004)	(1.138)
Average distance to the project location	234.369	234.410	232.604
	(335.750)	(337.164)	(267.258)
Months in the market (in months)	55.351	57.171	24.740
	(44.022)	(44.355)	(25.573)

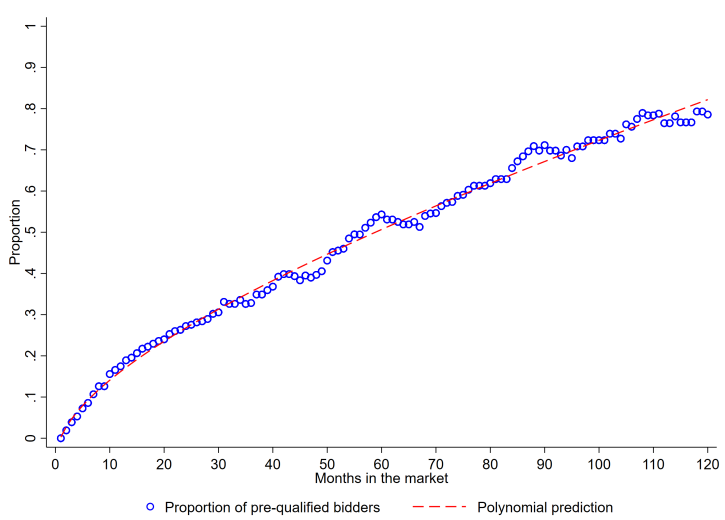
[a]Is in January 2000 constant-dollar (in millions) value.
Standard deviations are in parentheses.

While the distance to a project location is similar between prequalified and non-prequalified bidders, prequalified bidders' average backlog is about $10 million more than non-prequalified bidders, whose backlog is only about $67,000. Further, an average non-prequalified firm survives only about 25 months in the TxDOT market, which is less than half the time an average prequalified bidder stays in the TxDOT market (57 months).

2.2 Analysis of Firm Dynamics and Project Outcomes

We next explore the mechanisms underlying the patterns documented above, examining firm evolution, survival patterns, and the causal relationships between winning contracts and firm development. We begin by documenting how firms evolve and survive differentially based on qualification status. We then provide causal evidence from regression discontinuity analysis of close auctions, followed by analysis of the trade-offs between fostering competition and ensuring project execution quality.

Figure 1: Fraction of non-qualified entrants becoming qualified over time in the market



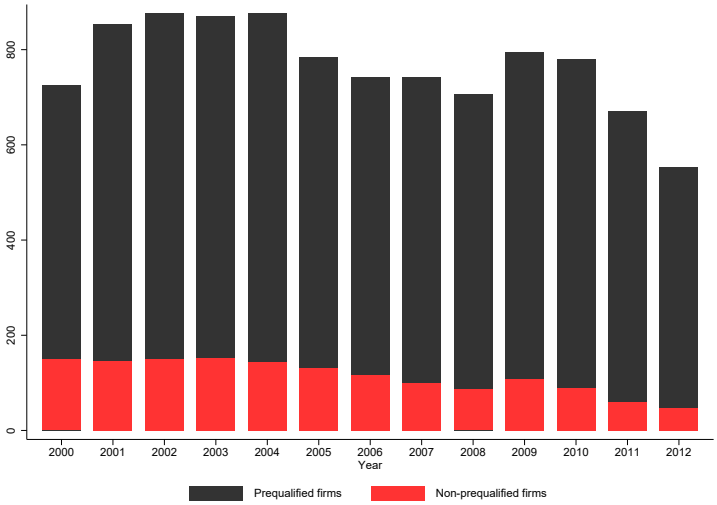
2.2.1 Patterns of firm evolution and survival

Conversations with TxDOT officials indicate that TxDOT considers waived projects to be an avenue by which small firms can enter the TxDOT market, gain experience and, subsequently, grow and prequalify while increasing competition. Figure 1 shows how non-prequalified firms progress into prequalified firms as they age. When constructing this figure, we consider only non-prequalified firms that entered the TxDOT market since January 2000. The hollow circles denote the actual proportion of firms that entered as non-prequalified firms and became prequalified. We also plot a fractional polynomial prediction of this status evolution and show it with a dashed line.[6]

We next explore patterns of entry, exit and evolution in the TxDOT market. We observe about 2,500 distinct firms ever active in the TxDOT market. Figure 2 plots the number of distinct firms active in each year of our sample, as well as the number of these which had yet to qualify at some point in the year. On average, between 600 and 850 firms are active each year. While most active firms are prequalified firms, in each year a non-negligible minority

[6]See Royston and Altman (1994) and Royston and Sauerbrei (2008) for details on fractional polynomial estimation.

Figure 2: Number of active firms per year



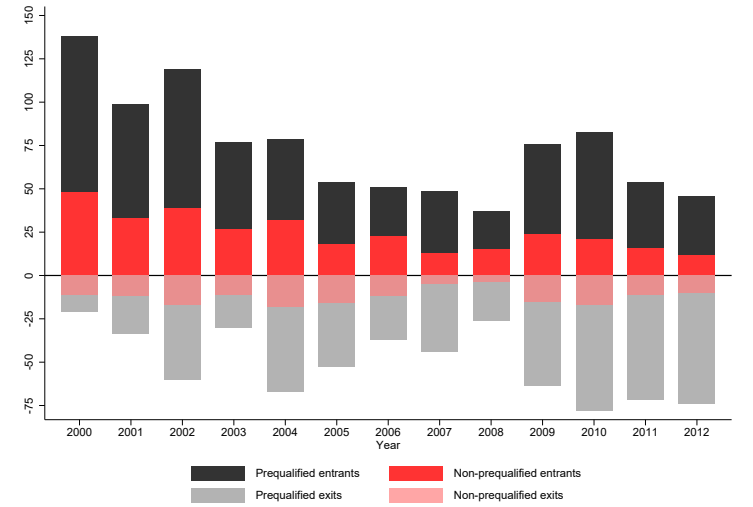
of active firms spend at least part of the year non-qualified. Moreover, as we show next, these not-yet-qualified firms make up a larger share of new entrants than firms overall.

The TxDOT market also has substantial churn, with more entry and exit than the net changes illustrated in Figure 2 would suggest.[7] Figure 3 plots the number of firms entering and exiting the market each year in our sample, broken down by prequalification status at the time of entry / exit. In most years, the sum of entering and exiting firms exceeds one hundred, which is large relative to the number of active firms. Note further that some firms that enter as non-prequalified will eventually become qualified; thus, the non-prequalified exits illustrated in Figure 3 should be interpreted as firms which never qualified.

Finally, we explore firm-level survival patterns in the TxDOT market. Toward this end, we estimate a simple Cox proportional hazard model, where the dependent variable is a firm's exit date. In Figure 4, we then plot survival probabilities by time in the market and prequalification status. As can be seen in Figure 4, firms that enter the market and remain

[7]Note that a firm may change its name if it moves from a 'partnership' to a 'limited liability company.' For example, 'DGS Construction' could be renamed as 'DGS Construction Ltd.' In this case, the firm will receive a new TxDOT identification number. In consecutive months, we find 19 instances where a firm changed its name to a similar name as mentioned above when located in the exact geocoded location. However, in our setting, we consider these two names to belong to the same entity.

Figure 3: Number of firms entering and exiting the TxDOT market per year



unqualified tend to exit the market relatively quickly: almost 80 percent exit within the first year. Meanwhile, firms that become qualified tend to remain in the market much longer, with more than half of firms active for at least two years. Even these firms, however, experience significant attrition, with only about 35 percent remaining active for at least five years.

These patterns of firm evolution and differential survival motivate our empirical analysis below, which seeks to understand the mechanisms driving these dynamics.

2.2.2 Causal evidence from regression discontinuity analysis

To isolate the causal effect of winning contracts on firm trajectories, we implement a regression discontinuity (RD) design that exploits quasi-random variation in close auction outcomes. We define the running variable as the bid margin: for losing bidders, the difference between their bid and the winning bid normalized by the engineer's cost estimate (ECE); for winning bidders, the difference between their bid and the second-lowest bid normalized by ECE. The treatment—winning the auction—exhibits a sharp discontinuity at a bid margin of zero. In very close auctions, whether a firm wins or loses is plausibly random conditional on bidding near the threshold.

Figure 4: Survival probabilities for prequalified and non-prequalified firms

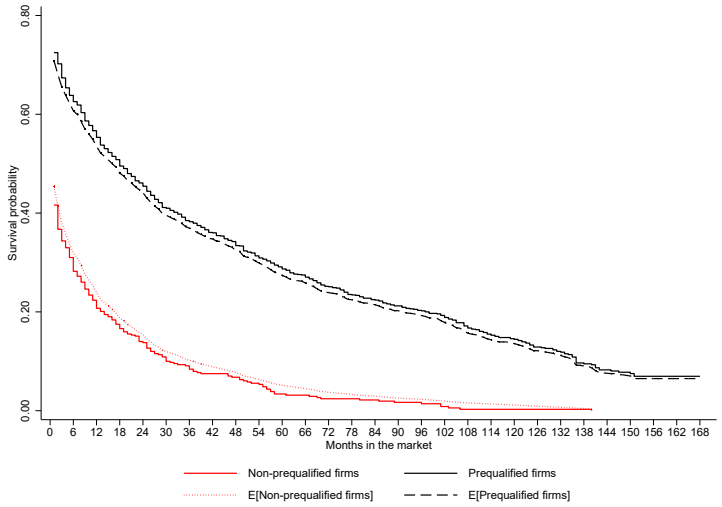


Table 3 presents results from local linear regressions using a 3% bandwidth around the winning threshold, including both direct effects and interactions with log prior wins to test for heterogeneous treatment effects.

The results reveal several important findings. First, winning has a strong causal effect on firm survival. Winning a single contract increases the three-year survival probability by 8.4 percentage points ($p < 0.001$) for completely inexperienced firms. This effect diminishes significantly with experience: the interaction term of -0.025 implies that for a firm with approximately 28 prior wins, the survival benefit of an additional win becomes negligible.

Second, we find suggestive evidence that winning increases the value of future contracts. For inexperienced firms, winning increases the log value of contracts won over the next three years by 0.89 ($p = 0.079$), implying an approximate 89% increase. While only marginally significant at the 10% level, the economic magnitude is substantial. The interaction term of -0.26 ($p = 0.056$) suggests the effect is strongest for new entrants and diminishes with experience.

Third, while the point estimate for the effect on future win counts is positive (0.098), it is not statistically significant at conventional levels. This may reflect the fact that winning

Table 3: Regression Discontinuity Estimates: Effect of Winning on Firm Outcomes

	(1) Log Size of Future Wins	(2) Survive 3 Years	(3) Log Number of Future Wins
Win	0.889*	0.084***	0.098
	(0.506)	(0.018)	(0.081)
Win × Log(Prior Wins)	−0.262*	−0.025***	−0.006
	(0.137)	(0.005)	(0.022)
Log(Prior Wins)	2.268***	0.025***	0.614***
	(0.117)	(0.005)	(0.019)
Bid Margin	−11.436	0.418	−2.413
	(17.706)	(0.455)	(3.410)
Bid Margin × Win	12.669	−0.420	2.470
	(17.734)	(0.455)	(3.414)
Constant	8.969***	0.914***	0.744***
	(0.435)	(0.018)	(0.070)
Observations	3,514	3,514	3,514
R-squared	0.280	0.071	0.465

Notes: This table reports regression discontinuity estimates of the effect of winning an auction on future firm outcomes. The running variable is the bid margin, defined as (bid - winning bid)/ECE for losers and (bid - second bid)/ECE for winners. All specifications use local linear regression with a 3% bandwidth around the winning threshold and robust standard errors (in parentheses). The sample is restricted to firms where wincount_1 $\leq$ bidcount_1 to exclude out-of-state firms with imputed experience.
*** p<0.01, ** p<0.05, * p<0.10

Table 4: Regression results for auction outcomes

Variable	Additional days	Log(final pay)		Number of bidders
	OLS			PPML
	(1)	(2)	(3)	(4)
Experience (log of past win counts)	-1.065	-0.040	-0.033	
	(0.317)	(0.010)	(0.009)	
Firm's market share	1.255	-0.001	0.003	
	(1.219)	(0.019)	(0.015)	
Log of backlog	0.206	0.060	0.046	
	(0.462)	(0.015)	(0.014)	
Prequalified firm	-3.066	-0.052	-0.031	
	(2.015)	(0.028)	(0.021)	
Waived auction	0.899	0.049	0.051	0.067
	(0.918)	(0.019)	(0.016)	(0.024)
Log of ECE	3.369	0.977		0.062
	(0.759)	(0.014)		(0.021)
Log of bid			1.004	
			(0.010)	
Log of number of tasks to complete the project	0.588	0.070	-0.014	-0.048
	(0.589)	(0.012)	(0.010)	(0.015)
Log of number of bidders	-2.106	-0.153	-0.005	
	(0.722)	(0.015)	(0.012)	
Log sum of market share of all plan holders in the auction	-9.498	1.953	0.802	
	(25.559)	(0.550)	(0.436)	
Log of total value of a new building permits in Texas	Yes	Yes	Yes	Yes
Fed/State effect	Yes	Yes	Yes	Yes
Project type and zone effects	Yes	Yes	Yes	Yes
Observations	2,290	2,290	2,290	2,290
R^2	0.044	0.776	0.838	
Log pseudolikelihood				-4,771

Robust standard errors are in parentheses.

affects the intensive margin (size of contracts won) more than the extensive margin (number of contracts won), or simply that we lack power to detect this effect.[8]

2.2.3 Project execution trade-offs

While the regression discontinuity results demonstrate the importance of early wins for firm development, facilitating such wins through policies like qualification waivers involves potential trade-offs. Table 4 explores these trade-offs by examining how winner characteristics affect immediate project execution.

From TxDOT's perspective, waiving qualification requirements has two potential benefits: in the short run, it increases the number of competitors in the waived auction (column

[8]We conduct robustness checks using a wider 5% bandwidth around the winning threshold, which yields qualitatively similar results across all specifications.

4), and in the long run, by allowing small firms to gain experience, it may make these firms more effective competitors. However, offsetting these benefits is the risk that awarding to less-experienced firms may lead to greater cost overruns or delays.

Using comparable waived and unwaived auctions (ECE less than $300,000), we find that awarding to more experienced firms significantly reduces both project delays (column 1) and cost overruns (columns 2 and 3). For final payments, this holds whether we condition on the log ECE (column 2) or the winning bid (column 3). Importantly, however, the waiver status itself does not significantly affect final costs or delays after controlling for winner characteristics, suggesting that performance issues stem from inexperience rather than the waiver mechanism per se. Award to a qualified firm may also reduce both delays and final payments, but this effect is weakly significant at best.

2.2.4 Key findings

Our analysis reveals several critical insights about firm dynamics in procurement markets. The patterns documented in Section 2.2.1 show dramatic differences in firm trajectories based on qualification status: while 80% of non-qualified firms exit within their first year, over half of qualified firms that entered as non-qualified eventually become qualified, demonstrating the potential for firm development.

The regression discontinuity evidence establishes that winning contracts causally drives these divergent trajectories. Winning increases three-year survival probability by 8.4 percentage points for inexperienced firms, with effects diminishing as firms accumulate experience. This heterogeneity directly explains the survival patterns observed in Figure 4—early wins are crucial for establishing viability, transforming firms from the high-mortality non-qualified group into sustainable market participants.

Yet the project outcomes analysis reveals the tension inherent in fostering such development: inexperienced winners deliver projects with more delays and cost overruns. This trade-off between dynamic efficiency (building a competitive supplier base through firm de-

velopment) and static efficiency (minimizing immediate project costs) lies at the heart of procurement policy design.

The evidence suggests qualification waivers successfully navigate this trade-off. By facilitating entry and early wins for new firms, waivers enable the firm development documented in Figure 1, ultimately expanding the competitive supplier base. The substantial magnitude of the long-term benefits—particularly the survival effects—suggests these dynamic gains justify the short-term execution costs. This empirical foundation motivates our structural model, which will quantify these trade-offs more precisely and evaluate alternative policy designs.

3 A dynamic oligopoly model of the TxDOT market

Motivated by the patterns documented above, we consider an infinite-horizon continuous time dynamic game in the spirit of Arcidiacono et al. (2016). Players are prime contracting firms in the bidding market for TxDOT highway construction contracts. In what follows, we refer to the TxDOT highway procurement market as "the market," and set of firms active in this market as "the industry." Time is continuous and indexed by t. All players discount the future at common (known) exponential rate ρ.

3.1 Model setup

Let $\mathcal{J}$ denote the set of active firms in the industry in a particular moment in time, with $J = |\mathcal{J}|$ the number of active firms. There are L locations in the market, indexed by $\ell \in \{1, ..., L\} \equiv \mathcal{L}$. Each firm $i \in \mathcal{J}$ is endowed with a single time-invariant location $\ell_i \in \mathcal{L}$. For future reference, let $\mathcal{J}_\ell \subset \mathcal{J}$ denote the set of firms active at location $\ell \in \mathcal{L}$ (where by definition $\mathcal{J}_1, ..., \mathcal{J}_\ell$) form a partition of $\mathcal{J}$). In addition to their locations, firms $i \in \mathcal{J}$ are characterized by three state variables which evolve over time: experience, indexed by $w_i \in \{0, ..., W\} \equiv \mathcal{W}$, backlog, indexed by $z_i \in \{0, ..., Z\} \equiv \mathcal{Z}$, and qualification status

$q_i \in \{0,1\}$. Let the firm state vector $s_i = (\ell_i, w_i, z_i, q_i)$ characterize both time-invariant and time-varying states of firm i, and let $\mathcal{S} = \mathcal{L} \times \mathcal{W} \times \mathcal{Z} \times \{0,1\}$ be the set of possible values for the firm-specific state s_i, $S = |\mathcal{S}| = 2L(W+1)(Z+1)$ be the cardinality of $\mathcal{S}$, and $\iota = 1, ..., S$ be any scalar-valued index of the elements of $\mathcal{S}$. Firms play type-symmetric strategies, so that outcomes depend only on the numbers, not identities, of firms at each state. Following Ericson and Pakes (1995), we therefore define the industry state $s^{\mathcal{J}}$ as an S-element vector counting the number of active firms at each point in the firm-level state space $\mathcal{S}$: $s^{\mathcal{J}} = (s_\iota^{\mathcal{J}})_{\iota=1}^{S}$, where $s_\iota^{\mathcal{J}} = |\{j \in \mathcal{J} : s_j = \mathcal{S}_\iota\}|$. Finally, since our data include the 2008 recession, we also allow for an aggregate state $a \in \{1, ..., A\} = \mathcal{A}$ representing other external factors relevant for the market (in our application, outside construction opportunities).

Active firms compete via low-price sealed-bid auctions to win a sequence of heterogenous highway construction projects $k \in \{1, 2, ...\}$ generated over time by TxDOT. Each project k is described by a vector of observable characteristics $x_k \in \mathcal{X}$ drawn from distribution F_x. Each active firm incurs a flow cost $c_f(s_i, a)$ while it remains in the market. Firms enter and exit the industry, participate and win auctions, gain experience becoming more productive, increase and decrease in size, and complete projects as described in detail below.

Due to the large size of the TxDOT market (with 600-800 firms active per year), the set of possible industry states $s^{\mathcal{J}}$ in our application is huge (larger than can be represented using standard floating-point numbers). In our empirical analysis, we will thus ultimately model firms as tracking their own states s_i, the aggregate state a, and the exact states of competitors for each auction in which they participate, but not the full industry state $s^{\mathcal{J}}$. Rather, we model each firm i as tracking two low-dimensional functions of $s^{\mathcal{J}}$: a vector of industry moments $\hat{s}_0(s^{\mathcal{J}})$ summarizing characteristics of all firms active in the industry, and a vector of local moments $\hat{s}_{\ell_i}(s^{\mathcal{J}_{\ell_i}})$ summarizing characteristics of firms $j \in \mathcal{J}_{\ell_i}$ active at the firm's own location ℓ_i. We will ultimately close our model using a location-based extension of the MME solution concept of Ifrach and Weintraub (2017), defined formally in Section 3.2 below. This equilibrium concept is particularly appropriate to our setting,

which includes the large aggregate shocks of the Great Recession. Importantly, however, our empirical framework can also straightforwardly accommodate a number of other recently-proposed large-market solution concepts, an advantage to which we return after defining our location-based MME in Section 3.2.

Projects, auctions, and matching Active firms compete in low-price sealed-bid auctions for TxDOT construction projects, bidding opportunities for which arise according to an exogenous Poisson process. Specifically, new projects are generated by TxDOT with Poisson rate λ_0. Upon arrival of a project k, its characteristics x_k are drawn from F_x. Each active firm $i \in \mathcal{J}$ is then instantaneously matched to project k according to a (common knowledge) match probability function $m(x_k, s_i, \bar{s})$ which depends on characteristics of the project x_ℓ, the firm's state s_i, and the market moments $\bar{s}$.[9] If matched to project ℓ, firm i becomes a bidder for project ℓ. We include the market moments $\bar{s}$ in $m(x_\ell, s_i, \bar{s})$ to allow for potential nonlinear or non-monotone relationships between the number of active firms and auction-level outcomes, as can be induced for example by endogenous auction-level entry as in Li and Zheng (2009).[10]

Bidding Let $\mathcal{J}_\ell^b \subset \mathcal{J}_\ell$ denote the subset of firms $j \in \mathcal{J}_\ell$ matched to auction ℓ. Upon being matched to auction ℓ, each bidder $i \in \mathcal{J}_\ell$ observes auction characteristics x_ℓ and the states of all rivals $j \in \mathcal{J}_\ell$ and draws their private completion cost $c_{i\ell}$ from a distribution $F_c(\cdot|s_i, x_\ell, a)$, with costs drawn independently across bidders conditional on observables and

[9] In particular, if $q_i = 0$ but the auction requires a pre-qualification status, i.e. $x_{\ell,q} = 1$, then firm i is not eligible to become a bidder for project ℓ and thus $m(s_i, x_\ell, \bar{s}) = 0$.

[10] In preliminary versions of this paper, we modeled matching between active firms and projects at the planholder level, with matched planholders observing the set of rival planholders and making endogenous auction-level entry decisions based on private entry costs as in Li and Zheng (2009) and Groeger (2014). Our dynamic framework extends straightforwardly to incorporate this auction-level entry stage, and indeed we estimated several early versions including an additional auction-level entry choice. Ultimately, however, we felt that this additional entry layer distracted from what we see as our main contribution, development of a framework suitable for analysis of the dynamic evolution of large-scale procurement auction markets. For purposes of this paper, therefore, we elected to focus on a simpler model with matching directly to the bidding stage. We believe, however, that extension of our framework to include an auction-level entry stage represents a natural and interesting avenue for future research.

cost distributions common knowledge to bidders. Note that, when firms play type-symmetric strategies, only rival types, not rival identities, are relevant for forecasting bidding behavior. Paralleling our definition of the industry state $s^{\mathcal{J}}$, we therefore define the competition state s^b_ℓ for auction ℓ as a $S \times 1$ vector whose elements count the number of bidders $j \in \mathcal{J}^b_\ell$ at each element of the firm-level state space $\mathcal{S}$.

Based on their private cost realization $c_{i\ell}$, project characteristics x_ℓ, the competition state s^b_ℓ, and tracked market moments $\bar{s}$, each bidder $i \in \mathcal{J}^b_\ell$ submits a bid $b_{i\ell}$. TxDOT sets a secret reserve price R_ℓ, drawn from distribution $G_0(\cdot|x_\ell)$. The bidder submitting the minimum bid wins, subject to this bid being below the secret reserve price.

If bidder i wins auction ℓ, they receive their bid $b_{i\ell}$, incur their completion cost $c_{i\ell}$, and augment their backlog by the size of project ℓ: $z_i' = z_i + z_\ell$. In addition, their experience level ω_i increments by one, up to the maximum experience level W: $\omega_i' = \min\{\omega_i + 1, W\}$. States for non-winners are unchanged. Matching, bidding and winning occur instantaneously.

Project completion, capacity, and qualification transitions In addition to endogenous transitions through winning auctions, firms' states also transition through the following processes. First, firm i with state s_i and positive backlog ($z_i > 0$) transitions to each lower backlog level $z_i' < z_i$ according to a Poisson jump process with rate $\delta_z(z_i', s_i)$, where $\delta_z(z_i', s_i)$ represents expected backlog transitions through project completion.[11]

Second, firm i with current state s_i transitions between capacity levels according to a Poisson jump process, transitioning to each new capacity level $\kappa_i' \in \mathcal{K}, \kappa_i' \neq \kappa_i$ with rate $\delta_\kappa(\kappa_i', s_i)$. Conceptually, it is straightforward to extend our framework to model capacity as a dynamic decision. Empirically, however, we observe only an approximation to firm capacity, for which we believe our simple Poisson process is more appropriate. Note also that the transition rate $\delta_\kappa(\kappa_i', s_i)$ depends on all aspects of the firm's current state. Our

[11] This is, of course, a stylization of the actual completion process. We could in principle model the actual completion rates for each of the firm's current projects, but this would require tracking many more state variables per firm: set of projects won, size remaining and scheduled completion date on each project. The simple Poisson process here parsimoniously approximates the actual completion process without increasing the state space unduly. We use actual completion data to estimate this process as described later.

formulation therefore allows auction outcomes to endogenously affect capacity evolution on the equilibrium path. The main limitation is that we implicitly hold capacity transition rates conditional on state fixed in counterfactuals, whereas with fully dynamic policies these rates would counterfactually evolve.

Finally, recalling that eligibility for prequalification is determined primarily by a firm's accumulation of the requisite experience and capacity, we model firm i's transition from non-qualified status, $q_i = 0$, to qualified status, $q_i = 1$, as an exogenous Poisson rate $\delta_q(s_i)$ which depends on firm i's current state. As above, we emphasize that this framing allows qualification to depend endogenously on auction outcomes, insofar as these affect the firm's experience, backlog, and capacity. The main caveat is that qualification rates conditional on state are held fixed in counterfactuals, although we see this as less of a limitation than for capacity since the state s_i directly controls for the main factors determining qualification.

Industry exit Each active firm $i \in \mathcal{J}$ with backlog $z_i = 0$ receives opportunities to exit the industry at Poisson rate λ_χ. Upon receiving an exit opportunity, a firm may choose either to exit ($\chi = 1$) or not ($\chi = 0$). If the firm chooses to exit, it receives instantaneous payoff $\Psi_\chi(s_i) + \sigma_\chi(\kappa_i)\epsilon_{\chi,i}$ and takes no further actions, where $\Psi_\chi(s_i)$ is the average scrap value for a firm with state s_i, $\epsilon_{\chi,i}$ is drawn i.i.d. from a mean-zero logistic distribution, and $\sigma_\chi(\kappa_i)$ is a scale parameter which potentially depends on firm capacity (allowing, for example, smaller firms to experience smaller shocks). Otherwise, the firm continues in its current state.

Industry entry New potential entrants arrive in the market at Poisson rate λ_e. If a new potential entrant chooses to enter, it incurs instantaneous payoff $-\Psi_e(a) + \epsilon_{e,i}$, where $\Psi_e(a)$ is an average entry cost and $\epsilon_{e,i}$ is an idiosyncratic component drawn i.i.d. from a mean-zero logistic distribution with scale parameter σ_e. Firm i then becomes a member of the set of active firms with initial states $\omega_i = 0$, $z_i = 0$, $\kappa_i = 1$, and $q_i \in \{0, 1\}$ drawn from distribution F_q^e. Otherwise, potential entrant i receives a net value of zero and takes no further action.

In addition to new potential entrants described above, established prime contractors

may sometimes cross over to the TxDOT market from other states. Since these established external firms are not the focus of our model, we model these established prime contractors as entering exogenously at Poisson rate $\lambda_e^e(\bar{s})$ which depends on the market moments $\bar{s}$. Since we do not have detailed data on prior activities of all non-Texas firms, we model established external firms as arriving with qualification $q_i = 1$ and experience $\omega_i = \bar{\omega}$, where $\bar{\omega}$ is calibrated at average experience of qualified firms in Texas.[12]

3.2 Moment-based Markovian Equilibrium (MME)

In markets with few firms, it is typical to analyze dynamic oligopoly behavior using the Markov Perfect Equilibrium (MPE) solution concept (Maskin and Tirole (1988a,b), Ericson and Pakes (1995)), in which each active firm is modeled as conditioning strategies and beliefs on the full industry state. Unfortunately, in industries with more than a handful of firms, MPE quickly becomes unwieldy, as the cardinality of the industry state grows exponentially in the number of firms. This curse of dimensionality renders MPE both infeasible and (in our view) behaviorally implausible in markets such as TxDOT with hundreds of active firms.[13] Motivated by this problem, recent work on large oligopolies has instead relied on several alternative, closely related, solution concepts tailored to large-scale dynamic games: Oblivious Equilibrium (OE, Weintraub, Benkard, and Roy (2008)), Extended Oblivious Equilibrium, (EOE, Weintraub, Benkard, and Roy (2010)) , Stationary Equilibrium (SE, Adlakha, Johari, and Weintraub (2015)), and Moment-based Markovian Equilibrium (MME, Ifrach and Weintraub (2017)). In brief, these solution concepts obtain tractability by modeling firms as tracking their own states plus realizations of at most a few other relevant states (e.g. aggregate shocks), with all other aspects of the industry state summarized via low-dimensional moments (MME), long-run averages (OE, EOE), or long-run distributions (SE). They can

[12] Since we interpret backlog and capacity as specific to the TxDOT market, we model backlog and capacity for established outside firms in the same way as new entrants.

[13] For example, in our TxDOT application, there are typically more than 600 active firms, each of which can be in one of more than 500 distinct states, implying a set of industry states with cardinality larger than the maximum value, approximately 2×10^{308}, which can be represented in standard floating-point arithmetic.

be viewed either as approximations to MPE or as models of firm behavior in their own right.

Depending on the application in question, our model can be closed under any of these four solution concepts, with only minor modifications to our subsequent analysis. In practice, the key tradeoff is whether one wishes to prioritize allowing for aggregate shocks (MME, EOE), or conditions under which convergence to MPE can be guaranteed (OE, SE); since industry states will depend on the full history of aggregate shocks, models which allow for the former typically cannot guarantee the latter. Since our data include the 2008 recession, we here close the model building on the MME solution concept of Ifrach and Weintraub (2017), which allows straightforwardly for aggregate shocks.[14] Intuitively, our MME solution models firms as forecasting payoffs from future auctions based on their own states s_i, the aggregate state a, and the low-dimensional vector of industry moments $\hat{s}(s^{\mathcal{J}})$ introduced above, but not the full industry state $s^{\mathcal{J}}$. Under conditions outlined below, this allows us to replace the high-dimensional industry state $s^{\mathcal{J}}$ with the low-dimensional industry moments $\hat{s}(s^{\mathcal{J}})$ in each firm's dynamic programming problem. Crucially, however, this simplification applies only to forecasts of competition in future auctions: once matched to a specific auction, bidders further observe and condition their bidding strategies on the exact states of auction rivals. Our MME solution thereby embeds strategically rich auction-level "spot market" competition within an overarching model of dynamic oligopoly which remains tractable even in markets with hundreds of firms.

To formalize this MME solution concept, we first discuss how firms' beliefs are formed over the states $\bar{s}$ and s^b_ℓ and formally state them in Assumptions 1 and 2. Let $\tilde{g}^b$ denote firm i's belief about the competition state s^b_ℓ for auction ℓ with characteristics x_ℓ, formed upon being matched with project ℓ but before observing rivals or their states. Let $\mathcal{H}^{\mathcal{J}}$ denote any history of the industry up to the present. We model $\tilde{g}^b$ as depending on exact values of x_ℓ, the firm's own state s_i, and the aggregate state a, but as depending on the industry state $s^{\mathcal{J}}$

[14] Ifrach and Weintraub (2017) motivate MME as a tool to analyze markets with dominant firms. From the perspective of other players, however, the state of the dominant firm is isomorphic to an aggregate shock, implying that Ifrach and Weintraub (2017)'s definitions simplify straightforwardly to markets with aggregate shocks instead of dominant firms.

and other aspects of the history $\mathcal{H}^{\mathcal{J}}$ only through the industry moments $\hat{s}(s^{\mathcal{J}})$. Conditional on x_ℓ, s_i, and $\bar{s} = (\hat{s}(s^{\mathcal{J}}), a)$, however, firms have correct beliefs on average:

Assumption 1. *For any auction characteristics x_ℓ and any histories $\mathcal{H}_1^{\mathcal{J}}, \mathcal{H}_2^{\mathcal{J}}$ such that $s_{i,1} = s_{i,2} = s_i$ and $\bar{s}_1 = (\hat{s}(s_1^{\mathcal{J}}), a_1) = \bar{s}_2 = (\hat{s}(s_2^{\mathcal{J}}), a_2) = \bar{s}$,*

$$\tilde{g}^b(s_\ell^b | x_\ell, \mathcal{H}_1^{\mathcal{J}}) = \tilde{g}^b(s_\ell^b | x_\ell, \mathcal{H}_2^{\mathcal{J}}) \equiv \tilde{g}^b(s_\ell^b | x_\ell, s_i, \bar{s}) \quad \forall \quad s_\ell^b.$$

Moreover, $\tilde{g}^b(s_\ell^b | x_\ell, s_i, \bar{s})$ is equal to the long-run distribution of s_ℓ^b conditional on $(a, x_\ell, s_i, \bar{s})$.

We next turn to firm beliefs over $\bar{s} = (\hat{s}(s^{\mathcal{J}}), a)$. In practice, we discretize $\bar{s}$ into M possible values: $\bar{s} \in \{\bar{s}_1, ..., \bar{s}_M\}$. We further assume that each firm i perceives $\bar{s}$ to evolve according to a perceived transition kernel which is a Markov jump process with transition rates consistent with long-run observed transitions:

Assumption 2. *Each firm perceives $\bar{s} = (\hat{s}(s^{\mathcal{J}}), a)$ to follow a Markov jump process in which, for each distinct $m, n \in \{1, ..., M\}$, the perceived transition rate from $\bar{s} = \bar{s}_m$ to $\bar{s} = \bar{s}_n$ is a constant, denoted γ_{mn}, equal to the long-run average transition rate from $\bar{s}_m$ to $\bar{s}_n$ implied by the true stationary distribution of $(s^{\mathcal{J}}, a)$.*

Finally, consistent with this structure on beliefs, we model firms as playing moment-based strategies. For an incumbent, a strategy must specify rules for bidding and industry exit, while for a potential entrant, a strategy must specify a rule for industry entry. For incumbents, a moment-based bidding rule is a function $\beta(c_{i\ell} | x_\ell, s_i, s_\ell^b, \bar{s})$ such that, conditional on being matched to auction ℓ with characteristics x_ℓ, drawing private completion cost $c_{i\ell}$, and facing rival bidders described by the bidding state s_ℓ^b and aggregate conditions summarized by the market moments $\bar{s}$, firm i with state s_i submits bid $b_{i\ell} = \beta(c_{i\ell} | x_\ell, s_i, s_\ell^b, \bar{s})$. A moment-based exit rule can be described by a real-valued function $\bar{\chi}(s_i, \bar{s})$ such that, upon receiving an exit opportunity, firm i with state s_i facing market moments $\bar{s}$ exits whenever $\epsilon_{\chi,i} \geq \bar{\chi}(s_i, \bar{s})$. Finally, recalling that potential entrants are ex ante identical, a moment-based

entry rule can be described by a real-valued function $\bar{\eta}(\bar{s})$ such that, upon receiving an entry opportunity facing market moments $\bar{s}$, firm i enters whenever $\epsilon_{e,i} \geq \bar{\eta}(\bar{s})$.[15]

Following Ifrach and Weintraub (2017), we define a MME as an equilibrium in moment-based strategies when firms have moment-based beliefs satisfying Assumptions 1-2. We further assume that firm behavior in the TxDOT market is well-described by such an MME:

Definition 1. *A MME of our model is a tuple of moment-based bidding, industry exit, and industry entry rules $(\beta, \bar{\chi}, \bar{\eta})$, together with beliefs over s^b_ℓ, rival bids, and $\bar{s}$, such that:*

1. *Taking firm beliefs and rival strategies as given, incumbent strategies $(\beta, \bar{e}, \bar{\chi})$ maximize incumbents' perceived continuation values within the class of moment-based strategies;*
2. *Taking firm beliefs and rival strategies as given, entrant strategies $\bar{\eta}$ maximize potential entrants' continuation values within the class of moment-based entry strategies;*
3. *Beliefs over s^b_ℓ satisfy Assumption 1;*
4. *Beliefs over rival bids given $(x_\ell, s_i, s^b_\ell, \bar{s})$ are consistent with $\beta(\cdot|x_\ell, s_i, s^b_\ell, \bar{s})$;*
5. *Beliefs over $\bar{s}$ are described by a perceived transition kernel satisfying Assumption 2.*

Assumption 3. *Firm behavior is consistent with an MME of our model.*

Note that Assumptions 1-3 restrict how the realized aggregate state $s^{\mathcal{J}}$ enters beliefs and strategies, but not how a firm conditions on other payoff-relevant information. In particular, once a firm has been matched with a specific project ℓ, it will condition its bidding decisions on the exact states of rivals faced in the auction. Consequently, when rivals play MME strategies, the exact industry state $s^{\mathcal{J}}$ is relevant for firm i's payoffs only insofar as it influences the distribution of rivals faced in future auctions; the firm is already conditioning on all states relevant for auction ℓ. The core restriction embodied by Assumptions 1-3 is thus

[15] Our model focuses on entry decisions by small, not-yet-established firms. In practice, we also observe some large, established firms entering from outside states. We model these established firms as arriving at an exogenous Poisson rate $\lambda^e_e(\bar{s})$ which is determined outside of equilibrium.

that forecasts of future payoffs based on the full industry state $s^{\mathcal{J}}$ are well-approximated by forecasts based only on the tracked moments $\hat{s}(s^{\mathcal{J}})$. For large procurement markets such as TxDOT, we view this as a very plausible model of behavior, preserving the richness of strategic interation within individual auction "spot markets" while nevertheless rendering the full dynamic oligopoly model tractable.

3.3 Continuation values in Moment-based Markovian Equilibrium

Suppose that firms in the industry play an MME as defined above. Consider an active firm i with state s_i which, at some instant, is not matched with an auction. For this firm, states of rival firms are relevant for future profit only through expected competition in future auctions, which firm i forecasts conditional only on the aggregate moments $\bar{s}$. Furthermore, firm i perceives $\bar{s}$ to evolve according to a Markov jump process. We can therefore describe firm i's perceived continuation value with a value function $V(s_i, \bar{s})$ which depends on the high-dimensional industry state $s^{\mathcal{J}}$ only through the tracked moments $\bar{s}$.

To characterize $V(s_i, \bar{s})$, first observe that in any short interval of time h, a presently unmatched active firm i may experience any of the following events:

Auction arrival: With Poisson rate λ_0, a new project ℓ may arrive, in which case project characteristics X_ℓ will be drawn from F_x and firm i becomes a matched bidder with probability $m(X_\ell, s_i, \bar{s})$;

Exit opportunity: With Poisson rate λ_χ, firm i with $z_i = 0$ may receive an exit opportunity;

Qualification: With Poisson rate $\lambda_q(s_i)$, firm i with $q_i = 0$ may transition to $q_i = 1$;

Capacity: With Poisson rate $\delta_\kappa(\kappa_i', s_i)$, firm i's capacity may transition from κ_i to κ_i';

Completion: With Poisson rate $\delta_z(z_i', s_i)$, firm i's backlog may transition from z_i to $z_i' < z_i$;

Aggregate moment change: With Poisson rate γ_{mn}, $\bar{s}$ may transition from $\bar{s}_m$ to $\bar{s}_n$.

While qualification, capacity, completion, and aggregate moment changes lead to simple state transitions, auction arrival and exit opportunity events lead to equilibrium choices which depend on ex ante unknown private information. For future reference, let $\Pi^0(s_i, \bar{s})$ be the ex ante expected change in i's continuation value upon arrival of a new project (before observing project characteristics, matched bidders, or private completion costs); we derive $\Pi^0(s_i, \bar{s})$ in detail in Section 3.4 below. Similarly, let $\Pi^\chi(s_i, \bar{s})$ be the ex ante expected change in firm i's continuation value upon receiving an exit opportunity (before observing private exit payoffs), which we derive in detail in Section 3.5. We can then characterize $V(s_i, \bar{s})$ as the unique solution to the following Bellman equation: for all $s_i \in \mathcal{S}$ and $\bar{s}_m \in \{\bar{s}_1, ..., \bar{s}_M\}$,

$$\begin{aligned}\rho V(s_i; \bar{s}_m) = -c_f(s_i, a) + \lambda_0 \Pi^0(s_i; \bar{s}_m) + \lambda_\chi \mathbb{I}[z_i = 0]\Pi^\chi(s_i, \bar{s}) \\ + \lambda_q(s_i)\Big(V(s_i' = (\omega_i, z_i, \kappa_i, 1); \bar{s}) - V(s_i, \bar{s}_m)\Big) \\ + \sum_{z_i' < z_i} \delta_z(z_i', s_i)\Big(V(s_i' = (\omega_i, z_i', \kappa_i, q_i); \bar{s}_m) - V(s_i, \bar{s}_m)\Big) \\ + \sum_{\kappa_i'} \delta_\kappa(\kappa_i', s_i)\Big(V(s_i' = (\omega_i, z_i, \kappa_i', q_i); \bar{s}_m) - V(s_i, \bar{s}_m)\Big) \\ + \sum_{n=1}^{M} \gamma_{mn}\Big(V(s_i; \bar{s}_n) - V(s_i, \bar{s}_m)\Big). \quad (1)\end{aligned}$$

The left-hand side of (1) represents the flow value a bidder receives from being in state $(s_i, \bar{s}_m)$. By definition, this flow value must equal to the sum of expected flow surplus from each possible event above, which the right-hand side computes. Specifically, $\lambda_0 \Pi_0(s_i; \bar{s}_m)$ represents i's net surplus from arrival of new projects, $\lambda_\chi \mathbb{I}[z_i = 0]\Pi^\chi(s_i, \bar{s})$ represents i's net surplus from arrival of exit opportunities conditional on being exit-eligible, and the final four lines represent changes in continuation value induced by state transitions through qualification, project completion, capacity changes, and aggregate moments respectively.

We next derive a computationally tractable Conditional Choice Probability (CCP) representation of $V(s_i, \bar{s})$. Toward this end, we first show that $\Pi^0(s_i, \bar{s})$ can be recovered directly

from equilibrium bidding and matching outcomes. We then derive a CCP representation of $\Pi^x(s_i, \bar{s})$ which is identified up to an expression which is linear in unknown model parameters. Combining these expressions, we obtain a simple linear-in-parameters representation of $V(s_i, \bar{s})$, which renders estimation of our model simple even in large markets.

3.4 Identification of auction arrival profit $\Pi^0(s_i, \bar{s})$

Bidding equilibrium and ex ante bidder profit First consider firm i's bidding decision upon being matched with auction ℓ, after observing project characteristics x_ℓ, the bidding state s^b_ℓ, and their own completion cost $c_{i\ell}$. If firm i wins the auction with bid b_i, it receives spot profit $b_i - c_{i\ell}$ and continuation value $V(s_i + \Delta s_i, \bar{s})$, where Δs_i is the (deterministic) change in s_i induced by winning project ℓ. Meanwhile, if firm i loses auction ℓ, it receives continuation value $V(s_i, \bar{s})$.[16] Based on the public state $(x_\ell, s_i, s^b_\ell, \bar{s})$, as well as its private completion cost $c_{i\ell}$, bidder i thus chooses its bid b_i to maximize

$$\max_{b_i} \left\{ \Big(b_i - c_{i\ell} + V(s_i + \Delta s_i, \bar{s}) - V(s_i, \bar{s}) \Big) \times P_i(b_i | x_\ell, s_i, s^b_\ell, \bar{s}) + V(s_i, \bar{s}) \right\}, \tag{2}$$

where $P_i(b_i | x_\ell, s_i, s^b_\ell, \bar{s})$ is the equilibrium probability that firm i wins with bid b_i conditional on the public state $(x_\ell, s_i, s^b_\ell, \bar{s})$. In any MME, we can express $P_i(b_i | x_\ell, s_i, s^b_\ell, \bar{s})$ as

$$P_i(b_i | x_\ell, s_i, s^b_\ell, \bar{s}) = [1 - G_0(b_i | x_\ell)] \prod_{j \in \mathcal{J}^b_\ell, j \neq i} [1 - G_j(b_i | x_\ell, s_i, s^b_\ell, \bar{s})],$$

where $G_j(\cdot | x_\ell, s_i, s^b_\ell, \bar{s})$ is the c.d.f. of equilibrium bids submitted by bidder $j \in \mathcal{J}^b_\ell$ given state $(x_\ell, s_i, s^b_\ell, \bar{s})$, and as above $G_0(b_i | x_\ell)$ is the c.d.f. of TxDOT's secret reserve price.[17]

If bid $b_{i\ell}$ is an optimal bid for bidder i, then it follows from (2) that $b_{i\ell}$ must (almost

[16] In expressing firm i's continuation payoffs, we implicitly make use of the fact that, from i's perspective, the aggregate moments $\bar{s}$ follow an exogenous process which is not affected by the outcome of auction ℓ.

[17] We assume that the secret reserve price distribution only depends on auction characteristics.

surely) satisfy the necessary Guerre et al. (2000)-type first-order condition

$$\Big(b_{i\ell} - c_{i\ell} + V(s_i + \Delta s_i, \bar{s}) - V(s_i, \bar{s})\Big) = \frac{P_i(b_{i\ell}|x_\ell, s_i, s^b_\ell, \bar{s})}{-p_i(b_{i\ell}|x_\ell, s_i, s^b_\ell, \bar{s})}, \tag{3}$$

where $p_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s}) = \frac{\partial P_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})}{\partial b_i}$ is the derivative of i's winning probability with respect to i's bid. The left-hand side is the change in i's continuation value induced by winning auction ℓ at the observed bid $b_{i\ell}$. The right-hand side is the (negative) inverse hazard rate of the minimum bid among i's rivals, evaluated at i's bid $b_{i\ell}$ conditional on s^b_ℓ. This term is identified directly from equilibrium bidding data.

Let $\Pi^b_i(x_\ell, s_i, s^b_\ell, \bar{s})$ denote the expected change in i's continuation value induced by bidding in auction ℓ, taken conditional on $(x_\ell, s_i, s^b_\ell, \bar{s})$ with respect to i's ex ante unknown completion cost c_{il}. In view of (3), we may express $\Pi^b(x_\ell, s_i, s^b_\ell, \bar{s})$ as

$$\Pi^b_i(x_\ell, s_i, s^b_\ell, \bar{s}) = \int \left[\frac{P_i(B|x_\ell, s_i, s^b_\ell, \bar{s})}{-p_i(B|x_\ell, s_i, s^b_\ell, \bar{s})}\right] P_i(B|x_\ell, s_i, s^b_\ell, \bar{s})\, g_i(B|x_\ell, s_i, s^b_\ell, \bar{s})\, dB; \tag{4}$$

i.e., the expected change in i's continuation value assuming optimal bidding times the probability of winning at each optimal bid, taken with respect to i's equilibrium bid density $g_i(\cdot|x_\ell, s_i, s^b_\ell, \bar{s})$. Since all objects on the RHS are identified, $\Pi^b_i(x_\ell, s_i, s^b_\ell, \bar{s})$ is identified.

Expected match surplus Let $\Pi^m(x_\ell, s_i, \bar{s})$ denote firm i's perceived expected surplus upon becoming a matched bidder for a project with characteristics x_ℓ, evaluated after observing x_ℓ but before observing the set of rival bidders also matched with project ℓ. In any MME, we can express $\Pi^m(x_\ell, s_i, \bar{s})$ as

$$\Pi^m(x_\ell, s_i, \bar{s}) = \sum_{S^b_{-i}} \Pi^b_i(x_\ell, S^e_\ell, \bar{s})\, \tilde{g}^b(S^e_\ell|x_\ell, s_i, \bar{s}), \tag{5}$$

where, as in Assumption 1, $\tilde{g}_b(s^b_\ell|x_\ell, s_i, \bar{s})$ is the long-run average distribution of bidder states observed when a firm with state s_i is matched to a project with characteristics x_ℓ given the

aggregate state moments $\bar{s}$. We can therefore directly identify $\Pi^m(x_\ell, s_i, \bar{s})$ as the average of $\Pi^b(s_\ell^e)$ across observed matches with characteristics $(x_\ell, s_i, \bar{s})$.

Ex ante surplus from new project arrival Finally, we return to $\Pi^0(s_i, \bar{s})$, bidder i's surplus upon arrival of a new project, before observing either project characteristics x_ℓ or matching outcomes. Conditional on arrival of a project with characteristics x_ℓ, bidder i becomes a matched bidder with probability $m(x_\ell, s_i, \bar{s})$. If firm i is matched, they then receive expected surplus $\Pi^m(x_\ell, s_i, \bar{s})$ characterized above. $\Pi^0(s_i, \bar{s})$ is therefore given by

$$\Pi^0(s_i, \bar{s}) = E_{X_\ell}\Big[\Pi^m(X_\ell, s_i, \bar{s}) m(X_\ell, s_i, \bar{s})\Big]. \tag{6}$$

We can directly identify $\Pi^m(x_\ell, s_i, \bar{s})$ as above and $m(x_\ell, s_i, \bar{s})$ from observed matching rates, implying that $\Pi^0(s_i, \bar{s})$ is also directly identified.

3.5 Ex ante exit opportunity surplus $\Pi^\chi(s_i, \bar{s})$

If active firm i has backlog $z_i = 0$, then it is eligible to exit the industry and receives exit opportunities with Poisson rate λ_χ. When an exit opportunity arrives, firm i can either choose to exit and receive scrap value $\Psi_\chi(s_i) + \sigma_\chi(\kappa_i)\epsilon_{\chi,i}$, or not and receive continuation value $V(s_i, \bar{s}_m)$. For $d \in \{0, 1\}$, let $P_{\chi,d}(s_i; \bar{s}_m)$ denote the CCP that an exit-eligible firm i optimally makes exit decision d conditional upon receiving an exit opportunity. The expected net surplus firm i receives from an optimal exit choice in this event can be written as

$$\begin{aligned}\Pi^\chi(s_i, \bar{s}) &= E \max\left\{V(s_i, \bar{s}), \Psi_\chi(s_i) + \sigma_\chi(\kappa_i)\epsilon_{\chi,i}\right\} - V(s_i; \bar{s}) \\ &= -\sigma_\chi(\kappa_i) \log P_{\chi,0}(s_i, \bar{s}),\end{aligned} \tag{7}$$

Again, the CCPs on the right-hand side of (7) are directly identified, implying a straightforward expression of $\Pi^\chi(s_i, \bar{s})$ which is identified up to σ_x.

3.6 CCP representation of $V(s_i, \bar{s}_m)$

Finally, we return to our main interest: the value function $V(s_i; \bar{s}_m)$ describing the expected continuation value of active firm i with current state s_i, facing an industry state described by tracked moments $\bar{s} = \bar{s}_m$. Let $N = 2WZKM$ be the total number of states $(s_i, \bar{s}_m)$, and let $\mathbf{V}$ be the $N \times 1$ vector whose elements stack i's continuation values $V(s_i, \bar{s}_m)$ in each state $(s_i, \bar{s}_m)$. We can then express $\mathbb{V}$ compactly in matrix form as follows. Let $\mathbf{c}_f$, $\mathbf{\Pi}^0$, and $\mathbf{P}_{\chi,0}$ be $N \times 1$ vectors which stack, respectively, the state-specific terms $c_f(s_i, a)$, $\Pi^0(s_i, \bar{s}_m)$, and $P_{\chi,0}(s_i, \bar{s}_m)$ across states $(s_i, \bar{s}_m)$, and let $\mathbb{I}_\chi$ and Σ_χ be $N \times N$ diagonal matrices with diagonal elements $\mathbb{I}[z_i = 0]$ and $\sigma_\chi(\kappa_i)$ respectively, all ordered in the same way as $\mathbf{V}$. Collecting all terms involving $\mathbf{V}$ on the left-hand side of (1), we can then express (1) in matrix form as

$$\mathbf{TV} = -\mathbf{c}_f + \lambda_0 \mathbf{\Pi^0} + \lambda_\chi \mathbb{I}_\chi \mathbf{P}_{\chi,0} \Sigma_\chi, \tag{8}$$

where the $N \times N$ matrix $\mathbf{T}$ collects known and / or directly estimated transition rates multiplying each element of $\mathbf{V}$ in the resulting system of equations. Furthermoare, as usual in Bellman equations, the matrix $\mathbf{T}$ will be diagonally dominant and thus invertible. We can therefore solve (8) to obtain our final CCP representation for $\mathbf{V}$:

$$\mathbf{V} = -\mathbf{T}^{-1}\mathbf{c}_f + \lambda_0 \mathbf{T}^{-1}\mathbf{\Pi^0} + \lambda_\chi \mathbf{T}^{-1}\mathbb{I}_\chi \mathbb{P}_{\chi,0} \Sigma_\chi. \tag{9}$$

We can therefore express continuation values $\mathbf{V}$ compactly in a form that is linear in the parameters $\mathbf{c}_f$ and $\sigma_\chi(\kappa_i)$. Linear-in-parameters expressions for $\mathbf{V}$ are not unusual; see for example Hotz and Miller (1993), Aguirregabiria and Mira (2007), Pesendorfer and Schmidt-Dengler (2008), and Arcidiacono et al. (2016) among others. Given the strategic richness embedded in our auction spot market, however, the ability to encapsulate dynamic effects of future auctions in the single directly identified term Π^0 is both striking and elegant. This in turn renders estimation within our framework surprisingly tractable, as we describe next.

4 Estimation Methods and Results

In taking our model to data, we interpret firms as playing an MME with three tracked industry moments $\hat{s}(s^{\mathcal{J}})$: the log number of active firms, the log number of prequalified firms, and the log value of total backlog among active firms. We take the aggregate state a to be the log value of new housing construction in Texas, which proxies for changes in the outside option over our sample period (in particular due to the 2008 financial crisis). The market moment vector $\bar{s} = (\hat{s}(s^{\mathcal{J}}), a)$ thus includes four components overall. We take the base unit of time as a day, and assume that firms discount the future at an annualized rate of 0.95 (which, given our daily base unit, implies $\rho \approx 0.00014$). We then estimate our model in several steps, paralleling the steps involved in constructing the value function above.

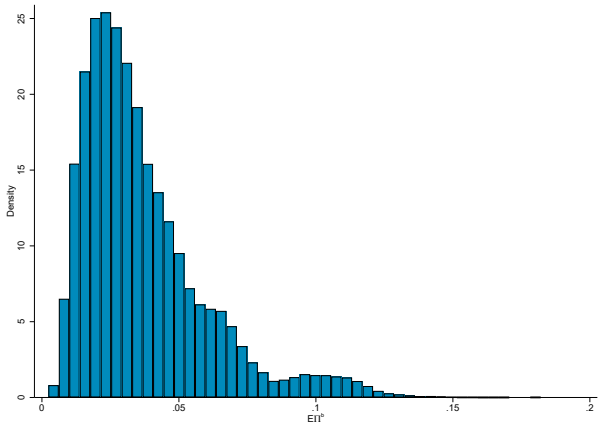
In implementing our dynamic methods, which assume a finite state space, we discretize state variables as follows. We discretize firm states s_i on a grid with $Z = 11$ backlog levels, $W = 6$ experience levels, $K = 5$ capacity levels, and two qualification levels, for a total of $S = 660$ distinct firm states. We similarly discretize the four elements of $\bar{s}$ on a grid with three values (low, medium, and high) in each dimension, for a total of $M = 81$ distinct aggregate states.[18] For each state variable, we first identify cut-points based on evenly spaced percentiles of the distribution of the variable in question in the relevant data (bidding data for firm states, time series data for aggregate variables). We then bin observations based on these cut-points, where to each binned observation we assign a value equal to the mean of the variable in question conditional on being within the bin.

4.1 Bid Distribution and Bidding Profit

We first estimate the equilibrium bidding distribution $G_i(\cdot|\cdot)$. In practice, following Jofre-Bonet and Pesendorfer (2003) and Athey et al. (2011), we specify a parametric first-step

[18] In practice, we observe aggregate states only monthly. We thus assign dates of aggregate state changes to the start of each month. In principle, we could simulate time paths of aggregate state changes within months and estimate all steps below using simulated MLE. Since aggregate states evolve very slowly, however, we expect the loss associated with start-of-month date assignment to be negligible.

Figure 5: Distribution of $\Pi_i^b(s_\ell^b)$ relative to engineer estimates



approximation of $G_i(\cdot|\cdot)$. In particular, we specify $G_i(\cdot|s_\ell^b)$ as a Weibull distribution:

$$G_i(\cdot|x_\ell, s_i, s_\ell^b, \bar{s}) = 1 - exp\left(-\frac{1}{\lambda_\ell}\left(\frac{b_{i\ell}}{ece_\ell} - lb_\ell\right)\right)^{k_\ell} \tag{10}$$

where ece_ℓ is the engineer's estimate for auction ℓ and lb_ℓ is the infimum support of scale-normalized bids in auction ℓ.[19] We parameterize the infimum support lb_ℓ, the scale parameter λ_ℓ, and the shape parameter k_ℓ as linear functions of the covariates in s_ℓ^b. We estimate the parameters in the equation above using MLE, subject to the constraint that $\frac{b_i}{ece_\ell} - lb_\ell \geq 0$.

With estimates $\hat{G}_i(\cdot)$ of $G_i(\cdot)$ for each bidder in hand, we plug in to equation 4 to construct an estimate $\hat{\Pi}_i^b(\cdot)$ for $\Pi_i^b(\cdot)$, the ex ante expected change in i's continuation value induced by bidding in an auction. Figure 5 reports the distribution, across auctions and bidders, of Π_i^b normalized by to the engineer's estimate ece_ℓ. On average, we find that the expected change in continuation value associated with bidding in an auction is equal to 4.08% of the engineer's estimate, a magnitude we view as empirically quite plausible.

[19] Following Jofre-Bonet and Pesendorfer (2003), we assume that the infimum support of bids in each auction the same for all bidders.

4.2 Expected match surplus

By definition, match surplus $\Pi^m(x_\ell, s_i, \bar{s})$ is the expectation of bid-stage surplus $\Pi^b(\cdot)$ across observed matches with characteristics $(x_\ell, s_i, \bar{s})$. Thus, to estimate $\Pi^m(x_\ell, s_i, \bar{s})$, we simply run a linear regression of $\hat{\Pi}^b_{i\ell}$ on $(x_\ell, s_i, \bar{s})$ and compute its predicted value $\hat{\Pi}^m_{i\ell}$. We include in this regression all auction characteristics, bidder states, and market states, as well as pairwise interactions between log experience, log backlog, and firm capacity. Results are reported in Table 5. Experience increases surplus, particularly for larger projects and larger firms, while distance decrease surplus. Backlog effects are non-monotone, with median-backlog firms having lower surplus than those with zero or very high backlogs. Large firms have lower baseline surplus but relative advantages for large projects and with high backlog. Note that these results pertain to bid-stage surplus inclusive of changes in continuation value induced by winning, not to structural completion costs. In Section 4.10, we recover estimates of completion costs after differencing out continuation values.

Table 5: Match profit regression: $\hat{\Pi}^b_{i\ell}$ on $(x_\ell, s_i, \bar{s})$

VARIABLES	COEF	SE
ECE (in millions)	0.0175	(0.00181)
Log Distance to Project	-0.00760	(0.000596)
Log Experience	0.00720	(0.00145)
Log Experience X ECE (millions)	0.000614	(0.000618)
Log Experience X Firm Capacity	0.224	(0.286)
Firm Capacity	-5.949	(2.003)
Firm Capacity X ECE (millions)	0.179	(0.121)
Firm Capacity X Log Backlog	0.128	(0.0400)
Log Backlog	0.00271	(0.00140)
I[Zero Backlog]	0.0446	(0.0186)
Prequalified Firm	-0.00954	(0.00253)
Log Total Prequalified Competitors in Market	-0.242	(0.0213)
Log Total Competitors in Market	0.228	(0.0195)
Log Total Value of Backlog in Market	-0.000702	(0.000397)
Log Value of New Housing Construction	0.0482	(0.00181)
Constant	-0.953	(0.0446)
Project Type and Zone Effects	Yes	
Observations	54,864	
R-squared	0.659	

Robust standard errors in parentheses

Table 6: Match probability $m(x_\ell, s_i, \bar{s})$: logit coefficients

VARIABLES	COEF	SE
Log Distance to Project	-0.257	(0.00336)
Log Experience X Log ECE	-0.135	(0.00304)
Log Backlog X Log Experience	-0.00392	(0.000827)
I[Firm Capacity = 1]	12.95	(0.406)
I[Firm Capacity = 2]	6.330	(0.415)
I[Firm Capacity = 3]	1.222	(0.423)
I[Firm Capacity = 4]	-3.810	(0.438)
I[Firm Capacity = 5]	-9.554	(0.455)
Log ECE X I[Firm Capacity = 1]	-0.461	(0.0102)
Log ECE X I[Firm Capacity = 2]	0.0510	(0.0114)
Log ECE X I[Firm Capacity = 3]	0.424	(0.0122)
Log ECE X I[Firm Capacity = 4]	0.755	(0.0137)
Log ECE X I[Firm Capacity = 5]	1.097	(0.0147)
Log Experience X I[Firm Capacity = 1]	1.921	(0.0402)
Log Experience X I[Firm Capacity = 2]	2.056	(0.0438)
Log Experience X I[Firm Capacity = 3]	2.144	(0.0463)
Log Experience X I[Firm Capacity = 4]	2.439	(0.0481)
Log Experience X I[Firm Capacity = 5]	2.714	(0.0507)
Log Backlog X I[Firm Capacity = 1]	0.0789	(0.00220)
Log Backlog X I[Firm Capacity = 2]	0.0797	(0.00274)
Log Backlog X I[Firm Capacity = 3]	0.0665	(0.00352)
Log Backlog X I[Firm Capacity = 4]	0.0423	(0.00398)
Log Backlog X I[Firm Capacity = 5]	0.0478	(0.00482)
Project Type and Zone Effects	Yes	
Observations	5,842,571	

Standard errors in parentheses

4.3 Match probability function

We parameterize the match probability function, $m(x_\ell, s_i, \bar{s})$, as a logit function whose index is a linear-in-parameters function of $(x_\ell, s_i, \bar{s})$. We estimate the parameters of this logit function via MLE, where observations are all valid potential matches between firms and projects (subject to the restriction that non-qualified firms cannot be matched with non-waived auctions) and the outcome is an indicator for whether each valid potential match was realized. The resulting coefficients are reported in Table 6. On average, larger firms match more frequently and with larger projects, with experience and backlog both increasing match rates (especially for larger firms). The latter effect is particularly interesting, and suggests that high current activity may generate efficiencies in future auction participation.

4.4 Auction arrival rate and auction arrival surplus

We choose the auction arrival rate, λ_0, to match the average auctions per day in our data. Since auction arrival is Poisson, this simply requires setting $\hat{\lambda}_0 = \frac{\text{\# days in sample}}{\text{\#auctions in sample}}$.

We estimate the auction arrival surplus $\Pi^0(s_i, \bar{s})$, by simulation. We draw a sample of $R = 1000$ auctions $\{x_r\}_{r=1}^R$ from the empirical distribution. For each combination of states $(s_i, \bar{s}_m)$ and each sampled auction x_ℓ, we predict expected match surplus $\hat{\Pi}^m(x_r, s_i, \bar{s})$ and match probabilities $\hat{m}(x_r, s_i, \bar{s})$ using the parameters estimated above. We then estimate $\Pi^0(s_i, \bar{s})$ by the simulated average $\hat{\Pi}^0(s_i, \bar{s}) = \frac{1}{R}\sum_{r=1}^R \hat{\Pi}^m(x_r, s_i, \bar{s})\hat{m}(x_r, s_i, \bar{s})$.

4.5 State transition rates

We parameterize the state-specific Poison rates associated with Markov jump processes for firm-level completion, qualification, and capacity events as follows:

Completion Backlog transition rates, $\delta_z(z_i', s_i)$, are specified as the product of a base rate λ_z representing the arrival of a completion event, times a truncated Poisson pdf (truncated between 1 and z_i) representing the number of steps taken in the event a completion event occurs. We take λ_z as a parameter to be estimated, and parameterize the rate in the truncated Poisson distribution as a log-linear-in-parameters function of s_i.

Qualification Qualification rates, $\lambda_q(s_i)$, are parameterized as a log-linear-in-parameters function of s_i.

Capacity We model capacity as evolving in single steps, to either one size larger or one size smaller. We specify the rate $\delta_\kappa(\kappa_i', s_i)$ associated with each of these transitions as a log-linear-in-parameters function of s_i.

We then estimate the parameters of each transition rate function using CMLE based on the relevant observed firm-level state transitions, as described in detail in Appendix B.

4.6 Transition kernel for aggregate state moments

As an empirical MME transition kernel, we assume that firms perceive each element $\bar{s}_j$ of $\bar{s}$, $j \in \{1, ..., 4\}$, as evolving in single steps according to a Markov jump process. Rates of increase (if feasible) and decrease (if feasible) in the aggregate state a, new housing construction in Texas, depend only on indicators for current value of a, reflecting our expectation that a is primarily driven by broader macro-economic conditions. Rates of increase (if feasible) and decrease (if feasible) for each of the industry moments $\hat{s}(s^{\mathcal{J}})$ are modeled as separate log-linear-in-parameters functions of indicators for the current level of the moment in question plus linear terms for other elements of $\bar{s}$. We then estimate parameters in each rate function using CMLE based on observed aggregate transitions as described in Appendix A.2.

4.7 First-step Exit CCPs

To construct first-step CCP estimates of exit policies $P_{\chi,d}(s_i, \bar{s})$, we parameterize exit probabilities as a logit function of $(s_i, \bar{s})$ whose odds ratio is linear in parameters. We calibrate the exit opportunity rate λ_χ such that exit-eligible firms (with backlog $z_{it} = 0$) receive, on average, one exit opportunity per month. For exit-elibile firms, we define the exit rate $q_\chi^{CCP}(s_i, \bar{s}) = \lambda_\chi P_{\chi,1}^{CCP}(s_i, \bar{s})$, and estimate parameters of the policy function $P_{\chi,1}(s_i, \bar{s})$ via CMLE as described in Appendix A.1. We then plug estimated first-step non-exit probability $\hat{P}_{\chi,1}(s_i, \bar{s})$ into (7) to obtain our CCP representation of $\hat{\Pi}_\chi$.

4.8 Structural parameter estimation

Finally, we turn to our model's structural parameters governing flow opportunity costs ($c_f(s_i, a)$), the distribution of scrap values ($\Psi_\chi(s_i)$ and σ_χ), the distribution of industry entry costs ($\Psi_e(a)$ and σ_e), and the distribution of project completion costs ($F_c(\cdot|x_\ell, s_i, a)$).

We parameterize mean scrap values, the logit error scale $\sigma_\chi(s_i)$, and flow opportunity costs as linear in parameters: $\Psi_\chi(s_i) = X_\chi(\kappa_i, \omega_i, q_i)\theta_\chi$ and $\sigma_\chi(s_i) = X_{\sigma_\chi}(\kappa_i)\theta_{\sigma_\chi}$, and

$c_f(s_i, a) = X_f(s_i, a)\theta_f$. Plugging these forms for $\Psi_\chi(s_i)$, $\sigma_\chi(s_i)$, and $c_f(s_i, a)$, plus estimates for $\hat{\mathbb{T}}$, $\hat{\Pi}^0$, $\hat{P}_{\chi,0}$, and Poisson rates constructed as above, into (9) yields a linear-in-parameters CCP representation of $V(s_i, \bar{s})$:

$$\hat{\mathbf{V}}(\theta_f, \theta_\chi, \theta_{\sigma_\chi}) = \hat{\mathbf{T}}^{-1}\mathbf{X}_f\theta_f + \hat{\lambda}_0\hat{\mathbf{T}}^{-1}\hat{\mathbf{\Pi}}^0 + \lambda_\chi\hat{\mathbf{T}}^{-1}\mathbb{I}_\chi\hat{P}_{\chi,0}\mathbf{X}_{\sigma_\chi}\theta_{\sigma_\chi}. \tag{11}$$

Conditional on receiving an exit opportunity, the structural probability that an incumbent chooses to exit can be expressed as

$$P_{\chi,1}(s_i, \bar{s}; \theta_f, \theta_\chi, \theta_{\sigma_\chi}) = \frac{\exp([X_\chi\theta_\chi - \hat{V}(s_i, \bar{s}; \theta_f, \theta_\chi, \theta_{\sigma_\chi})]/[X_{\sigma_\chi}\theta_{\sigma_\chi}])}{1 + \exp([X_\chi\theta_\chi - \hat{V}(s_i, \bar{s}; \theta_f, \theta_\chi, \sigma_\chi)]/[X_{\sigma_\chi}\theta_{\sigma_\chi}])}. \tag{12}$$

The structural exit rate for an exit-eligible firm is thus $q_\chi(s_i, \bar{s}; \theta_f, \theta_\chi, \theta_{\sigma_\chi}) = \lambda_\chi P_{\chi,1}(\theta_f, \theta_\chi, \theta_{\sigma_\chi})$.

We specify mean entry costs as $\Psi_e(a) = X_e(a)\theta_e$. Conditional on receiving an entry opportunity, the structural probability that a potential entrant chooses to enter is

$$\begin{aligned} &P_{e,1}(\bar{s}; \theta_f, \theta_\chi, \theta_{\sigma_\chi}, \theta_e, \sigma_e) \\ &\quad = \frac{\exp([p^0_{e,q}\hat{V}(s^0_e, \bar{s}; \theta_f, \theta_\chi, \theta_{\sigma_\chi}) + p^1_{e,q}\hat{V}(s^1_e, \bar{s}; \theta_f, \theta_\chi, \theta_{\sigma_\chi}) - X_e\theta_e]/\sigma_e)}{1 + \exp([p^0_{e,q}\hat{V}(s^0_e, \bar{s}; \theta_f, \theta_\chi, \theta_{\sigma_\chi}) + p^1_{e,q}\hat{V}(s^1_e, \bar{s}; \theta_f, \theta_\chi, \theta_{\sigma_\chi}) - X_\chi\theta_e]/\sigma_e)}, \end{aligned} \tag{13}$$

where $p^0_{e,q}$ and $p^1_{e,q}$ are the (exogenous) probabilities of realizing non-qualified or qualified states upon entry, and $s^0_e = (0, 0, 1, 0)$ and $s^1_e = (0, 0, 1, 1)$ are the associated post-entry states. We calibrate the number of potential entrants N^e as the average monthly number of non-prime subcontractors over our sample, and the entry opportunity rate λ_e such that each potential entrant receives, on average, one entry opportunity per quarter. The structural rate of new firm entry is therefore $N_e\lambda_e P_{e,1}(\bar{s}; \theta_f, \theta_\chi, \theta_{\sigma_\chi}, \theta_e, \sigma_e)$.[20]

Lastly, we parameterize the distribution $F_c(\cdot|x_\ell, s_i, a)$ of completion costs $c_{i\ell}$ based on

[20]For simplicity, we assume that that the pool of potential entrants does not change over time (i.e., that each entering firm is immediately replaced). Consequently, there is no loss in aggregating entry outcomes to spells defined by constant values of $\bar{s}$. We could alternatively track each potential entrant separately; this is straightforward in principle, although it would require adding yet another layer to our model.

the following model for ECE-normalized completion costs $c_{i\ell}/ece_\ell$:

$$\frac{c_{i\ell}}{ece_\ell} = \mu(x_\ell, s_i, a) + \sigma_c(x_\ell)\epsilon_{c,i\ell}, \tag{14}$$

where $\mu(x_\ell, s_i, a) = X_\mu(x_\ell, s_i, a)\theta_\mu$ parameterizes the mean of $c_{i\ell}/ece_\ell$, $\sigma_c(x_\ell)$ parameterizes the standard deviation of $c_{i\ell}/ece_\ell$, and $\epsilon_{c,i\ell}$ is an i.i.d. standard normal error. We also explored other parameterizations, such as log-normal models for $c_{i\ell}$, but found that the simple specification (14) provided a better representation of the data.[21]

Substituting from the bid first-order condition (3) and the CCP representation (11), we can form an empirical pseudo-likelihood from (14) based on[22]

$$\hat{c}_{i\ell} = \frac{b_{i\ell}}{ece_\ell} + \frac{\Delta_{i\ell}\hat{V}(s_i, \bar{s}; \theta_f, \theta_\chi, \theta_{\sigma_\chi})}{ece_\ell} + \frac{\hat{P}_i(b_{i\ell}|x_\ell, s_i, s^b_\ell, \bar{s})}{ece_\ell \hat{p}_i(b_{i\ell}|x_\ell, s_i, s^b_\ell, \bar{s})} \sim N(\mu(x_\ell, s_i, a), \sigma_c(x_\ell)). \tag{15}$$

In addition to identifying $\mu(x_\ell, s_i, a)$ and $\sigma_c(x_\ell)$, (15) also helps to identify the parameters $(\theta_f, \theta_\chi, \sigma_\chi)$ entering the CCP representation $\hat{V}$. This additional source of identification, from auction-level FOCs, allows us to recover parameters (such as flow opportunity costs of firm backlog) that would not be identified using exit decisions only.

We estimate $(\theta_f, \theta_\chi, \theta_{\sigma_\chi}, \theta_e, \sigma_e, \theta_c, \theta_\sigma)$ simultaneously using pseudo-MLE estimation (Aguirregabiria and Mira (2007)) based on observed industry exit decisions, observed industry entry outcomes, and implied firm completion costs for all projects with engineers' estimates of at least \$100,000 (95% of projects in our data). We form pseudo log-likelihoods for industry exit decisions based on (12), industry entry outcomes based on (13), and completion costs in the estimation sample based on (15), as described in detail in Appendix A. We then maximize

[21] While our normal specification for $\epsilon_{c,i\ell}$ does permit negative completion costs, these are not practically important; at our baseline estimates reported below, only about 0.003 of estimated costs are negative. In estimation, we have also explored modifying (14) to involve a normal distribution truncated at zero, finding very similar results. Introducing this truncation would, however, substantially complicate the sparse-grid strategy we aim to employ for counterfactuals, requiring us to double the number of grid dimensions. Given the large costs of the curse of dimensionality, we prefer the simple parameterization (14).

[22] In practice, we evaluate $\Delta_{i\ell}\hat{V}(\cdot)$ by interpolating i's pre- and post-win backlog and experience levels linearly on the discretized grid $\mathcal{Z} \times \mathcal{W}$, holding other states fixed at their discretized bid-stage values.

the sum of these (pseudo) log-likelihoods subject to the restriction that both scrap values and error scales be non-negative. We report estimates for dynamic parameters $(\theta_f, \theta_\chi, \sigma_\chi, \theta_e, \sigma_e)$ and completion costs $(\theta_c, \theta_\sigma)$ in the next two subsections 4.9 and 4.10 respectively.

4.9 Dynamic parameter estimates

Table 7 present estimates of our models' dynamic parameters $(\hat{\theta}_f, \hat{\theta}_\chi, \hat{\theta}_{\sigma_\chi}, \hat{\theta}_e, \hat{\sigma}_e)$ derived from the pseudo-MLE procedure above. Panel (a) presents scrap value parameters θ_χ and θ_{σ_χ}, panel (b) presents entry cost parameters θ_e and σ_e, and panel (c) presents flow opportunity cost parameters θ_f. Estimated error scale parameters $\hat{\theta}_{\sigma_\chi}$ suggest that larger firms have substantially larger idiosyncratic scrap value shocks, which is natural since these firms are much larger overall. Estimated entry costs are somewhat larger than expected, approximately \$5.12m, and increasing in outside construction activity. The estimated idiosyncratic $\hat{\sigma}_e$ are both small relative to their respective means, suggesting that our model can rationalize our data with relatively small idiosyncratic shocks. Since the dynamic portion of our model is relatively parsimonious, we view this as highly encouraging.

Figure 6 illustrates the MME value function $\hat{V}(s_i, \bar{s})$ implied by these dynamic estimates, focusing on variation in $\hat{V}(s_i, \bar{s})$ across firm-specific states s_i. Panels (a) and (b) illustrate how continuation values vary with experience and backlog for low-capacity ($\kappa_i = 1$) and high-capacity ($\kappa_i = K = 5$) firms respectively. Meanwhile, panels (c) and (d) illustrate how continuation values vary with capacity and log backlog for firms with low experience ($\omega_i = 1$) and high experience ($\omega_i = W = 5$) respectively. Experience and capacity both increase continuation values, with strong complementarities between the two. Such complementarities suggest that policies which influence the way firms gain experience—for example, qualification waivers—have substantial scope to influence long-run market composition.

Table 7: Dynamic scrap value, flow opportunity costs, and entry cost parameters

(a) Scrap value parameters (millions of dollars, positive: higher cost)

Parameter	Estimate	SE
θ_χ: $\mathbb{I}$[Firm Size = 1]	2.7781	0.0221
θ_χ: $\mathbb{I}$[Firm Size = 2]	3.6308	0.144
θ_χ: $\mathbb{I}$[Firm Size = 3]	5.1718	0.3069
θ_χ: $\mathbb{I}$[Firm Size = 4]	2.453	1.3768
θ_χ: $\mathbb{I}$[Firm Size = 5]	0	–
θ_χ: Std Log Exp X $\mathbb{I}$[Firm Size = 1]	-1.388	0.2596
θ_χ: Std Log Exp X $\mathbb{I}$[Firm Size = 2]	-1.3426	0.8749
θ_χ: Std Log Exp X $\mathbb{I}$[Firm Size = 3]	-5.3717	1.6648
θ_χ: Std Log Exp X $\mathbb{I}$[Firm Size = 4]	1.6785	8.0123
θ_χ: Std Log Exp X $\mathbb{I}$[Firm Size = 5]	-7.6913	30.6866
θ_χ: (Std Log Exp)^2 X $\mathbb{I}$[Firm Size = 1]	4.2333	0.4108
θ_χ: (Std Log Exp)^2 X $\mathbb{I}$[Firm Size = 2]	5.2382	1.1117
θ_χ: (Std Log Exp)^2 X $\mathbb{I}$[Firm Size = 3]	11.7775	1.8936
θ_χ: (Std Log Exp)^2 X $\mathbb{I}$[Firm Size = 4]	15.5031	10.0667
θ_χ: (Std Log Exp)^2 X $\mathbb{I}$[Firm Size = 5]	25.0649	33.4837
θ_{σ_χ}: Intercept	0.445	0.0052
θ_{σ_χ}: Std Firm Size	5.8288	0.081

(b) Entry cost parameters (millions of dollars, positive: higher cost)

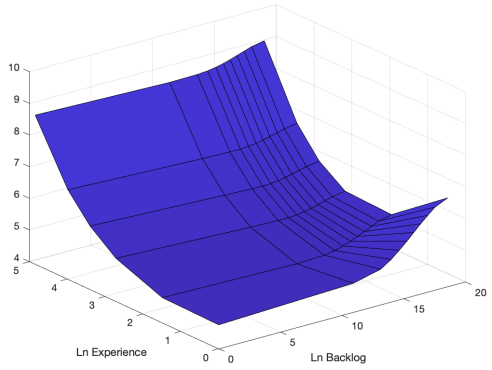
Parameter	Estimate	SE
θ_e: Constant	5.1248	0.1011
θ_e: Standardized Log New Housing	0.1725	0.0364
σ_e: Logistic error scale	0.1445	0.0283

(c) Flow cost parameters (millions of dollars, positive: higher cost)

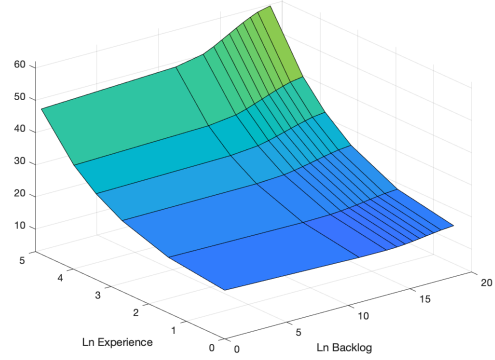
Parameter	Estimate	SE
θ_f: Standardized Log Backlog	-1.2599	0.145
θ_f: Std Log Experience X Std Log Backlog	26.8796	0.1482
θ_f: Std Log Backlog X Std Log New Housing	0.6189	0.2272

Notes: Covariates multiplying $(\theta_f, \theta_\chi, \theta_e)$ are standardized so that units are standard deviations around the mean of the relevant data (firm spells for θ_f and θ_χ and aggregate entry spells for θ_e). θ_f, θ_χ, and θ_e can thus be interpreted as changes (in millions of dollars) induced by a one-standard-deviation change in the characteristic of interest.

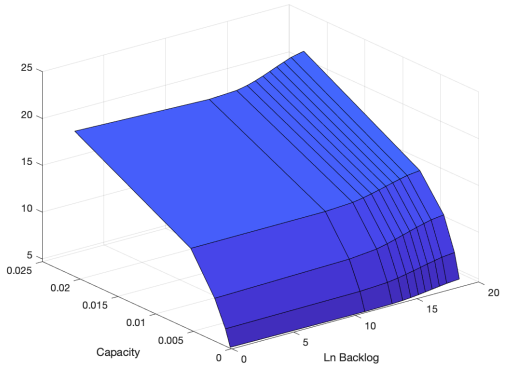
Figure 6: Estimated MME value function $\hat{V}(s_i, \bar{s})$ for a qualified firm as a function of experience, backlog, and capacity, evaluating aggregate moments $\bar{s}$ at their midpoint levels.



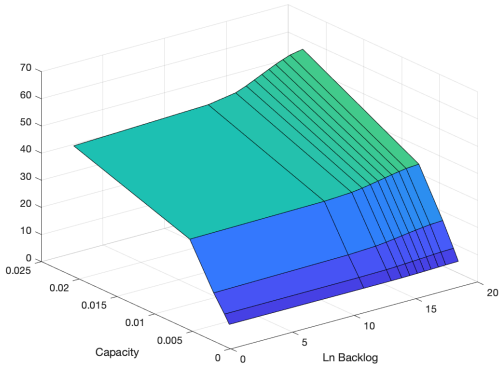
(a) Value function in experience versus log backlog for low capacity firm ($\kappa_i = 1$)



(b) Value function in experience versus log backlog for high capacity firm ($\kappa_i = K = 6$)



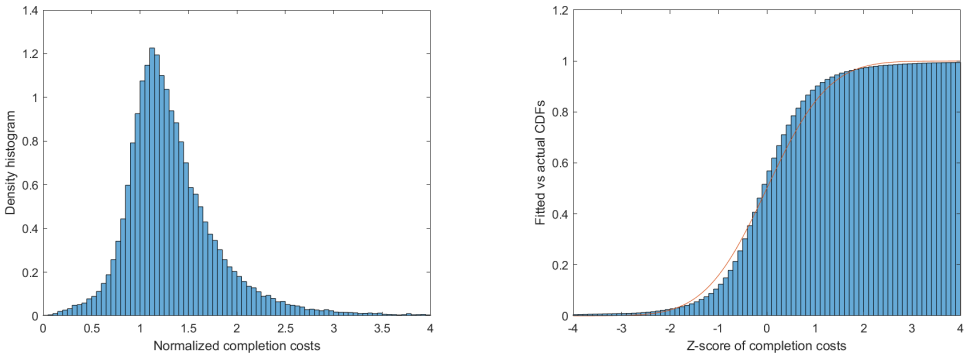
(c) Value function in capacity versus log backlog for low experience firm ($\omega_i = 1$)



(d) Value function in capacity versus backlog for high experience firm ($\omega_i = W = 5$)

Figure 7: Distribution and fit of $\hat{c}_{i\ell}/ece_\ell$ for projects with $ece_\ell > \$100k$

(a) Estimates: Histogram of $\hat{c}_{i\ell}/ece_\ell$ (b) Fit: Z-scores of $\hat{c}_{i\ell}/ece_\ell$



4.10 Completion cost estimates

Figure 7 illustrates the distribution and fit of our completion cost estimates within our estimation sample (the 95% of projects with engineer's estimates of at least $100,000). Panel (a) illustrates the histogram of $\hat{c}_{i\ell}/ece_\ell$ across all bidders for projects with $ece_\ell \geq$.[23] The distribution of $\hat{c}_{i\ell}/ece_\ell$ is slightly asymmetric, with a first decile value of 0.71, a median of 1.05, a mean of 1.08, and a ninth decile value of 1.48, all magnitudes we view as plausible. Meanwhile, Panel (b) plots the cumulative histogram of implied z-scores for these cost estimates relative to their predicted CDF, that of a standard normal distribution. Encouragingly, Panel (b) suggests that the fit of our baseline specification (14) is quite good, which we view as highly encouraging given the complexity of our model.

Next, in Table 8, we report estimates for parameters governing the mean and standard deviation of normalized completion costs $c_{i\ell}/ece_\ell$. Log standard deviation coefficients, reported in Panel (b), imply that project size substantially impacts the dispersion of $c_{i\ell}/ece_\ell$; for projects with $ece_\ell = \$100k$, $ece_\ell = \$1m$, and $ece_\ell = \$10m$, implied standard deviations of $c_{i\ell}/ece_\ell$ are approximately 0.546, 0.357, and 0.233 respectively. Mean coefficients, reported in Panel (a), suggest that experience substantially decreases expected completion costs, while distance and complexity both increase completion costs. While capacity patterns are some-

[23] For clarity, we illustrate this histogram for $0 \leq \hat{c}_{i\ell}/ece_\ell \leq 4$, which captures more than 99% of observations.

what noisy, larger firms broadly have higher baseline costs but greater economies of scale associated with high backlog. More non-highway construction increases completion costs, which is natural since such construction likely increases both opportunity and input costs for TxDOT work. Recalling that the dependent variable is $c_{i\ell}/ece_\ell$, the small coefficient on the engineer's estimate confirms that average costs scale nearly one-for-one in project size. As with continuation values, the relatively strong effects of dynamic factors such as experience and capacity on firm costs suggests that policies which influence the long-run composition of the TxDOT industry can have large effects on long-run market outcomes.

5 Conclusion and planned counterfactuals

This paper proposes a tractable dynamic oligopoly framework within which to analyze long-run industry composition in large procurement markets. Combining recent innovations in analysis of dynamic oligopoly models with many firms (Ifrach and Weintraub (2017)) and in continuous time (Arcidiacono et al. (2016)) with identification insights from the empirical auction literature, we obtain a simple CCP estimator for MME continuation values which renders estimation tractable even in markets with hundreds or thousands of firms. Estimation using 2000-2012 TxDOT data suggests that our model provides a good representation of the TxDOT market, yielding natural estimates of completion costs, flow costs, scrap values and entry costs accounting for forward-looking firm behavior. These estimates confirm that dynamic factors such as experience and capacity have substantial impacts on firms' completion costs, participation rates, and continuation values. This in turn suggests that procurement policies which impact long-run industry composition have substantial scope to affect long-run market outcomes, a mechanism about which little is known empirically.

In work in progress, we aim to explore several counterfactuals which will allow us to quantify such long-run impacts of procurement policy design in our TxDOT application. In practice, the main challenge in implementing these counterfactuals is that spot market

Table 8: Structural estimates: completion costs

(a) Completion costs: mean parameters (millions of dollars)

Parameter	Estimate	SE
θ_μ: Constant	-1.9173	0.1279
θ_μ: Log Experience	-0.0349	0.0015
θ_μ: $\mathbb{I}$[Firm Size = 2]	2.6978	0.0773
θ_μ: $\mathbb{I}$[Firm Size = 3]	4.155	0.0751
θ_μ: $\mathbb{I}$[Firm Size = 4]	4.2175	0.0763
θ_μ: $\mathbb{I}$[Firm Size = 5]	4.699	0.074
θ_μ: Log Engineer's Estimate X $\mathbb{I}$[Firm Size = 1]	0.0312	0.0048
θ_μ: Log Engineer's Estimate X $\mathbb{I}$[Firm Size = 2]	-0.1329	0.004
θ_μ: Log Engineer's Estimate X $\mathbb{I}$[Firm Size = 3]	-0.2101	0.003
θ_μ: Log Engineer's Estimate X $\mathbb{I}$[Firm Size = 4]	-0.2114	0.0025
θ_μ: Log Engineer's Estimate X $\mathbb{I}$[Firm Size = 5]	-0.1343	0.0018
θ_μ: Log Backlog X $\mathbb{I}$[Firm Size = 1]	0.0044	0.0009
θ_μ: Log Backlog X $\mathbb{I}$[Firm Size = 2]	-0.0185	0.0008
θ_μ: Log Backlog X $\mathbb{I}$[Firm Size = 3]	-0.0326	0.0007
θ_μ: Log Backlog X $\mathbb{I}$[Firm Size = 4]	-0.0276	0.0006
θ_μ: Log Backlog X $\mathbb{I}$[Firm Size = 5]	-0.114	0.0006
θ_μ: Log Distance	0.0146	0.0011
θ_μ: Log Number of Tasks	0.0834	0.0021
θ_μ: Log New TX Housing	0.0967	0.0051
Contract Type Controls	YES	
Contract Zone Controls	YES	

(b) Completion costs: Log standard deviation parameters

Parameter	Estimate	SE
θ_σ: Constant	1.5242	0.0193
θ_σ: Log Engineer's Estimate	-0.1849	0.0014

competition is through asymmetric first-price auctions, whose solutions are numerically difficult. Leveraging the fact that mean completion costs and continuation values enter the bidding problem isomorphically, however, we plan to apply recent developments in sparse grid interpolation methods (e.g. Brumm and Scheidegger (2017)) to solve for auction-level payoffs outside of the dynamic oligopoly equilibrium problem. With this accomplished, we can solve for counterfactual MME outcomes relatively simply as in Ifrach and Weintraub (2017). As examples, we plan to quantify the long-run effects of a static optimal reserve price, of switching from first to second-price auctions, and of revising or eliminating qualification waivers. Our model, however, also allows analysis of many other policy changes whose long-run effects are at present largely unknown.

References

Adlakha, S., R. Johari, and G. Y. Weintraub (2015): "Equilibria of dynamic games with many players: Existence, approximation, and market structure," Journal of Economic Theory, 156, 269–316, computer Science and Economic Theory.

Aguirregabiria, V. and P. Mira (2007): "Sequential Estimation of Dynamic Discrete Games," Econometrica, 75, 1–53.

Arcidiacono, P., P. J. Bayer, J. R. Blevins, and P. Ellickson (2016): "Estimation of Dynamic Discrete Choice Models in Continuous Time with an Application to Retail Competition," Review of Economic Studies, 83, 889–931.

Athey, S., J. Levin, and E. Seira (2011): "Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions," The Quarterly Journal of Economics, 126, 207–257.

Balat, J. (2015): "Highway Procurement and the Stimulus Package: Identification and Estimation of Dynamic Auctions with Unobserved Heterogeneity," Working paper, University of Texas at Austin.

Brumm, J. and S. Scheidegger (2017): "Using adaptive sparse grids to solve high-

dimensional dynamic models," Econometrica, 85, 1575–1612.

Ericson, R. and A. Pakes (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," The Review of Economic Studies, 62, 53–82.

Groeger, J. R. (2014): "A Study of Participation in Dynamic Auctions," International Economic Review, 55, 1129–1154.

Guerre, E., I. Perrigne, and Q. Vuong (2000): "Optimal Nonparametric Estimation of First-Price Auctions," Econometrica, 68, 525–574.

Hotz, V. J. and R. A. Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," The Review of Economic Studies, 60, 497–529.

Ifrach, B. and G. Y. Weintraub (2017): "A Framework for Dynamic Oligopoly in Concentrated Industries," Review of Economic Studies, 84, 1106–1150.

Jofre-Bonet, M. and M. Pesendorfer (2003): "Estimation of a Dynamic Auction Game," Econometrica, 71, 1443–1489.

Li, T. and X. Zheng (2009): "Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions," Review of Economic Studies, 76, 1397–1429.

Maskin, E. and J. Tirole (1988a): "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs," Econometrica, 56, 549–569.

——— (1988b): "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles," Econometrica, 56, 571–599.

Myerson, R. (1981): "Optimal auction design," Mathematics of Operations Research, 6, 58–73.

Pesendorfer, M. and P. Schmidt-Dengler (2008): "Asymptotic Least Squares Estimators for Dynamic Games," The Review of Economic Studies, 75, 901–928.

Saini, V. (2012): "Endogenous Asymmetry in a Dynamic Procurement Auction," RAND Journal of Economics, 43, 726–760.

Weintraub, G. Y., C. L. Benkard, and B. V. Roy (2008): "Markov Perfect Industry

Dynamics with Many Firms," Econometrica, 76, 1375–1411.

——— (2010): "Computational Methods for Oblivious Equilibrium," Operations Research, 58, 1247–1265.

Appendix A: Estimation details

A.1: Firm-specific transition rates

To estimate firm-specific transition rates, we partition our data into firm-specific spells defined by changes in the state variables $(s_i, \bar{s})$, where any change in either s_i or $\bar{s}$ defines the start of a new spell. For each spell t for firm i, let Δ_{it} denote the holding time (length) of the spell and s_{it} and $\bar{s}_{it}$ denote states prevailing during the spell. For each transition type τ, let d_{it}^{τ} be an indicator equal to one if spell it ended through a type-τ transition and zero if spell it ended through some other event. For each firm state $s_i' \in \mathcal{S}_i$, let $q_\tau(s'|s_{it}, \bar{s}_{it}, \theta_\tau)$ denote the Poisson rate of transitioning from state s_{it} to state s' through transition τ, where θ_τ are unknown parameters governing the associated rate. Finally, let $Q_\tau(s_{it}, \bar{s}_{it}, \theta^\tau) = \sum_{s' \in \in \mathcal{S}_i \cup \not\vdash, s' \neq s_{it}} q_\tau(s'|s_{it}, \bar{s}_{it}, \theta^\tau)$ denote the sum of rates of moving to all possible new states $s' \neq s_{it}$ through transition τ in spell it. We can then express the probability that spell it ends with no transition through τ ($d_{it}^{\tau} = 0$) after holding time Δ_{it} as the survival function of an exponential distribution with rate $Q_\tau(s_{it}, \bar{s}_{it}, \theta^\tau)$:

$$P(\Delta_{it}, d_{it}^{\tau} = 0 | s_{it}, \bar{s}_{it}, \theta_\tau) = \exp(-Q^\tau(s_{it}, \bar{s}_{it}, \theta^\tau)\Delta_{it}).$$

Meanwhile, the probability that spell it ends with a transition through τ ($d_{it}^{\tau} = 1$) to distinct state $s' \neq s_{it}$ after holding time Δ_{it} can be expressed as

$$\begin{aligned} P(\Delta_{it}, d_{it}^{\tau} = 1, s' | s_{it}, \bar{s}_{it}, \theta^\tau) &= Q_\tau(s_{it}, \bar{s}_{it}, \theta_\tau) \exp(-Q^\tau(s_{it}, \bar{s}_{it}, \theta^\tau)\Delta_{it}) \times \frac{q_\tau(s'|s_{it}, \bar{s}_{it}, \theta^\tau)}{Q_\tau(s_{it}, \bar{s}_{it}\theta_\tau)} \\ &= \exp(-Q_\tau(s_{it}, \bar{s}_{it}, \theta_\tau)\Delta_{it}) \times q_\tau(s'|s_{it}, \bar{s}_{it}, \theta_\tau), \end{aligned}$$

where, in the first line, the first term represents the p.d.f. of a transition through τ at holding time Δ_{it}, and the second term represents the probability of moving to s' conditional on a transition through τ. Combining cases, we can express the conditional log likelihood of spell outcome $(\Delta_{it}, d_{it}^{\tau}, s')$ through transition τ given parameters θ^τ as

$$\ell_{it}^{\tau}(\theta_\tau) = -Q_\tau(s_{it}, \bar{s}_{it}, \theta_\tau)\Delta_{it} + d_{it}^{\tau} \ln q_\tau(s'|s_{it}, \bar{s}_{it}, \theta_\tau). \tag{16}$$

We estimate θ^τ for each firm-specific transition τ by conditional maximum likelihood estimation (CMLE) based on the conditional likelihood $\mathcal{L}^\tau(\theta^\tau) = \sum_{i,t} \ell(\theta^\tau)$.

A.2: Perceived aggregate moment transition rates

We assume that firms perceive each element $\bar{s}_j$, $j = 1, ..., 4$, of the aggregate moment vector $\bar{s}$ as evolving in single steps according to a Markov jump process. Specifically, $\bar{s}_j$ increases by one unit (if feasible) with rate $q_j^+(\bar{s}, \beta_j^+) = \exp(X_j(\bar{s})\beta_j^+)$ and decreasing by one unit (if feasible) with rate $q_j^-(\bar{s}, \beta_j^-) = \exp(X_j(\bar{s})\beta_j^-)$, where $X_j(\bar{s})$ is a known function of $\bar{s}$ and β_j^+ and β_j^- are parameters to be estimated. In practice, for our three aggregate moments measuring industry composition, $X_j(\bar{s})$ includes separate indicators for each level of $\bar{s}_j$ plus linear terms for the other elements of $\bar{s}$. For our final aggregate moment, which measures the log value of new housing construction in Texas, $X_j(\bar{s})$ includes only indicators for the

current level of $\bar{s}_j$. This reflects our expectation that new housing construction is driven more by the broader macro-economy than by the TxDOT market specifically.

We partition our data on $\bar{s}$ into T distinct spells, indexed $t = 1, ...T$, between changes in $\bar{s}$. Let $\bar{s}_t$ denote the value of $\bar{s}$ during spell t and Δ_t be the length of spell t. For each element $\bar{s}_j$ of $\bar{s}$, we estimate β_j^+ and β_j^- based on a (pseudo) CMLE estimator paralleling (16) above:

$$(\hat{\beta}_j^+, \hat{\beta}_j^-) = \arg\max \left\{ \sum_{t=1}^{T} -\Delta_t \cdot [q_j^+(\bar{s}_t) + q_j^-(\bar{s}_t)] + \sum_{t=1}^{T} d_{jt}^+ \log q_j^+(\bar{s}_t) + \sum_{t=1}^{T} d_{jt}^- \log q_j^-(\bar{s}_t) \right\},$$

where d_{jt}^+ and d_{jt}^- are indicators for whether spell t ended with an increase in $\bar{s}_j$ or a decrease in $\bar{s}_j$ (or neither, in which case both indicators are zero). For each distinct $m, n \in \{1, ..., M\}$, we then derive the implied rate at which $\bar{s}$ transitions from $\bar{s}_m$ to $\bar{s}_n$, which we take as an estimate for the MME firm belief γ_{mn} over the corresponding transition.[24]

Appendix B: Dominant firms

This Appendix extends our baseline framework to account for the possibility of dominant firms. Following IW (2017), suppose that a subset of firm states $s_i \in S^d \subset S$ correspond to "dominant" firm states, while the remaining states $s_i \in S^f = S \backslash S^d$ correspond to fringe states. For example, in our application, "dominant" states could be those for qualified firms with maximal capacity ($q_i = 1$, $\kappa_i = K$), or with both maximal capacity and maximal experience ($q_i = 1$, $\kappa_i = K$, ω_i). Let the dominant firm state s^d be a $|S^d| \times 1$ vector whose elements count the number of firms at each element of the dominant state set S^d, and the industry fringe state s^f be a $|S^f| \times 1$ vector whose elements count the number of firms at each element of the fringe state S^f. For future reference, let $\mathcal{J}^d \subset \mathcal{J}$ denote the set of dominant firms active at a given instant in time.

Following IW (2017), we now model each firm i as tracking its own state s_i, the dominant firm state s^d, a vector of moments $\hat{s}(s^f)$ of the fringe state s^f, and the aggregate state a. We extend the MME industry state $\bar{s} = (s^d, \hat{s}(s^f), a)$ to include the dominant firms state s_d in addition to the tracked fringe moments $\hat{s}(s^f)$, and the aggregate state a. We define moment-based bidding rules $\beta(c_{i\ell}|x_\ell, s_i, s_\ell^b, \bar{s})$, exit rules $\bar{\chi}(s_i, \bar{s})$, and entry rules $\bar{\eta}(\bar{s})$ as in the main text, bearing in mind that the MME industry state $\bar{s} = (s^d, \hat{s}(s^f), a)$ now exactly tracks the dominant firms state s^d. We extend Assumptions 1-2 and Definition 1 to accommodate this richer information structure as follows:

Assumption 4. *For any auction characteristics x_ℓ and any histories $\mathcal{H}_1^{\mathcal{J}}, \mathcal{H}_2^{\mathcal{J}}$ such that $s_{i,1} = s_{i,2} = s_i$ and $\bar{s}_1 = (s_1^d, \hat{s}(s_1^f), a_1) = \bar{s}_2 = (s_2^d, \hat{s}(s_2^f), a_2) = \bar{s}$,*

$$\tilde{g}^b(s_\ell^b|x_\ell, \mathcal{H}_1^{\mathcal{J}}) = \tilde{g}^b(s_\ell^b|x_\ell, \mathcal{H}_2^{\mathcal{J}}) \equiv \tilde{g}^b(s_\ell^b|x_\ell, s_i, \bar{s}) \quad \forall \quad s_\ell^b.$$

[24] In view of our parameterizations, γ_{mn} will be nonzero only for states $\bar{s}_m$, $\bar{s}_n$ involving an increase or decrease of one unit in exactly element $\bar{s}_j$ of $\bar{s}$, in which case γ_{mn} will equal either $q_j^+(\bar{s}_m)$ or $q_j^-(\bar{s}_m)$ respectively.

f Moreover, $\tilde{g}^b(s^b_\ell|x_\ell, s_i, \bar{s})$ is equal to the long-run distribution of s^b_ℓ conditional on $(x_\ell, s_i, \bar{s})$.

Assumption 5. *For each firm i, beliefs over state transitions satisfy the following properties:*

1. *For each dominant rival $j \in \mathcal{J}^d$, firm i correctly perceives j's transition process between dominant states $s_j, s'_j \in S^d$;*

2. *Firm i correctly perceives transition rates in s^d induced by a currently dominant firm becoming non-dominant, with associated changes in $\hat{s}$ equal to corresponding long-run averages;*

3. *Firm i perceives transition rates in $\bar{s}$ induced by currently non-dominant rivals becoming dominant to follow a Markov jump process in which the perceived transition rate from s^d to $s^{d,\prime}$ given $\bar{s}$, denoted $\gamma_{fd}(\bar{s}'|\bar{s})$, equals the corresponding long-run average transition rate from $\bar{s}$ to $\bar{s}$ given $\bar{s}$;*

4. *Excluding transitions from or to dominance, firm i perceives the fringe moments $\hat{s}$ to follow a Markov jump process in which the perceived transition rate from $\hat{s}$ to $\hat{s}'$ given $\bar{s}$, denoted $\gamma_f(\hat{s}'|\bar{s})$, equals the corresponding long-run average transition rate from $\hat{s}$ to $\hat{s}'$ given $\bar{s}$.*

We then define an MME with dominant firms paralleling Definition 1, but substituting Assumptions 4-5 for Assumptions 1-2. We could also, or alternatively, adapt our baseline solution concept to allow, for example, firms to predict the impacts of their choices on the aggregate moments $\hat{s}$ as in Gowrisankaran et al. (2023). Such extensions would be implemented similarly to the dominant-firms model described here, although with some differences in detail due to differences in perceived transition rates.

B.1: Bidding equilibrium

First consider firm i's bidding decision upon being matched with auction ℓ, after observing project characteristics x_ℓ, the bidding state s^b_ℓ, and their own completion cost $c_{i\ell}$. Let $\Delta_{i\ell}$ be the deterministic change in i's state induced by winning auction ℓ, and for each and for each $j \in \cup\mathcal{J}_\ell$, let $\bar{\Delta}_{j\ell}$ be the (deterministic) change in the MME state $\bar{s}$ induced by j winning auction ℓ (where, under the assumptions above, $\bar{\Delta}_{j\ell} = 0$ if $j \neq \mathcal{J}^d$). Firm i now chooses its bid to maximize its change in continuation payoff accounting for changes in the industry state associated with dominant firms winning auction ℓ, i.e. to solve:

$$\max_{b_i} \Bigg\{ \Big(b_i - c_{i\ell} + V(s_i + \Delta_{i\ell}, \bar{s} + \bar{\Delta}_{i\ell}) - V(s_i, \bar{s}) \Big) \times P_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s}) \\ + \sum_{j \in \mathcal{J}^d_\ell, j \neq i} \Big(V(s_i, \bar{s} + \bar{\Delta}_{j\ell}) - V(s_i, \bar{s}) \Big) \times P_j(b_i|x_\ell, s_i, s^b_\ell, \bar{s}) \Bigg\},$$

where for each $j \in \mathcal{J}_\ell$, $P_j(b_i|x_\ell, s_i, s^b_\ell, \bar{s})$ denotes the probability that firm j wins the auction as a function of i's bid b_i. The FOC for i's bidding problem is

$$P_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s}) + \Big(b_i - c_{i\ell} + V(s_i + \Delta_{i\ell}, \bar{s} + \bar{\Delta}_i) - V(s_i, \bar{s})\Big) \times p_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s}) + \sum_{j \in \mathcal{J}^d_\ell, j \neq i} \Big(V(s_i, \bar{s} + \bar{\Delta}_{j\ell}) - V(s_i, \bar{s})\Big) \times p_j(b_i|x_\ell, s_i, s^b_\ell, \bar{s}) = 0, \quad (17)$$

which rearranged implies that i's optimal bid b_i must satisfy

$$\Big(b_i - c_{i\ell} + V(s_i + \Delta_{i\ell}, \bar{s} + \bar{\Delta}_{i\ell}) - V(s_i, \bar{s})\Big) = -\frac{P_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})}{p_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})} - \sum_{j \in \mathcal{J}^d_\ell, j \neq i} \Big(V(s_i, \bar{s} + \bar{\Delta}_{j\ell}) - V(s_i, \bar{s})\Big) \times \frac{p_j(b_i|x_\ell, s_i, s^b_\ell, \bar{s})}{p_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})}. \quad (18)$$

Substituting into (17), we can therefore express i's interim bid-stage surplus as

$$\pi^b_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s}) = -\frac{P_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})}{p_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})} P_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s}) + \sum_{j \in \mathcal{J}^d_\ell, j \neq i} \Big(V(s_i, \bar{s} + \bar{\Delta}_{j\ell}) - V(s_i, \bar{s})\Big) \times \left[P_j(b_i|x_\ell, s_i, s^b_\ell, \bar{s}) - \frac{p_j(b_i|x_\ell, s_i, s^b_\ell, \bar{s})}{p_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})} P_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})\right].$$

Consequently, integrating across realizations of b_i, we can express i's ex ante bid-stage surplus $\Pi^b_i(x_\ell, s_i, s^b_\ell, \bar{s})$ in a form paralleling Jofre-Bonet and Pesendorfer (2003):

$$\Pi^b_i(x_\ell, s_i, s^b_\ell, \bar{s}) = R^b_i(x_\ell, s_i, s^b_\ell, \bar{s}) + \sum_{j \in \mathcal{J}^d_\ell, j \neq i} \Big(V(s_i, \bar{s} + \bar{\Delta}_{j\ell}) - V(s_i, \bar{s})\Big) \times \big[P^b_{j,i}(x_\ell, s_i, s^b_\ell, \bar{s}) - H^b_{ij}(x_\ell, s_i, s^b_\ell)\big], \quad (19)$$

where $P^b_{j,i}(x_\ell, s_i, s^b_\ell, \bar{s})$ denotes the ex ante bid-stage probability that j wins an auction in which i participates,

$$P^b_{j,i}(x_\ell, s_i, s^b_\ell, \bar{s}) = \int_{b_i} P_j(b_i|x_\ell, s_i, s^b_\ell, \bar{s})\, dG_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s}),$$

and $R^b_i(x_\ell, s_i, s^b_\ell, \bar{s})$ and $H^b_{ij}(x_\ell, s_i, s^b_\ell, \bar{s})$ are defined by

$$R^b_i(x_\ell, s_i, s^b_\ell, \bar{s}) = \int_{b_i} \frac{P_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})}{p_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})} P_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})\, dG_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})$$

$$H^b_{ij}(x_\ell, s_i, s^b_\ell, \bar{s}) = \int_{b_i} \frac{p_j(b_i|x_\ell, s_i, s^b_\ell, \bar{s})}{p_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})} P_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s})\, dG_i(b_i|x_\ell, s_i, s^b_\ell, \bar{s}).$$

Note that $P^b_{j,i}(x_\ell, s_i, s^b_\ell, \bar{s})$, $R^b_i(x_\ell, s_i, s^b_\ell, \bar{s})$ and $H^b_{j,i}(x_\ell, s_i, s^b_\ell, \bar{s})$ are all directly identified.

B.2: Match surplus

Next consider active firm i with state s_i matched to auction ℓ with characteristics x_ℓ, before observing the competition state s_ℓ^b for auction ℓ. Taking expectations with respect to unknown rival match realizations, we can express match-stage surplus for this firm i as

$$\Pi^m(x_\ell, s_i, \bar{s}) = R_i^m(x_\ell, s_i, \bar{s}) + \sum_{j \in \mathcal{J}^d, j \neq i} \Big(V(s_i, \bar{s} + \bar{\Delta}_j) - V(s_i, \bar{s}) \Big) \times [P_{j,i}^m(x_\ell, s_i, \bar{s}) - H_{ij}^m(x_\ell, s_i, \bar{s})],$$

where, letting $\mathbb{I}_{j\ell}$ be an indicator for whether dominant firm $j \in \mathcal{J}^d$ is matched to auction ℓ, we define

$$\begin{aligned} R_i^m(x_\ell, s_i, \bar{s}) &= E[R_i^b(x_\ell, s_i, S_\ell^b, \bar{s}) | x_\ell, s_i, \bar{s}] \\ P_{j,i}^m(x_\ell, s_i, \bar{s}) &= E[\mathbb{I}_{j\ell} P_i^b(x_\ell, s_i, S_\ell^b, \bar{s}) | x_\ell, s_i, \bar{s}] \\ H_{j,i}^m(x_\ell, s_i, \bar{s}) &= E[\mathbb{I}_{j\ell} H_{j,i}^b(x_\ell, s_i, S_\ell^b, \bar{s}) | x_\ell, s_i, \bar{s}]. \end{aligned}$$

Observe that, by definition, $P_{j,i}^m(x_\ell, s_i, \bar{s})$ is the ex ante probability that dominant rival $j \in \mathcal{J}^d$ wins auction ℓ, conditional on firm i being matched to ℓ. Note further that $R_i^m(x_\ell, s_i, \bar{s})$, $P_{j,i}^m(x_\ell, s_i, \bar{s})$, and $H_{j,i}^m(x_\ell, s_i, \bar{s})$ are all directly identified.

B.3: Ex ante auction arrival surplus

Finally, consider the expected surplus of an active firm with state s_i facing MME state $\bar{s}$ upon learning that an auction has arrived, but before observing auction characteristics or match outcomes. If firm i is matched with auction ℓ, it receives expected net surplus $\Pi^m(x_\ell, s_i, \bar{s})$ characterized above. On the other hand, if firm i is not matched with auction ℓ, its continuation value will stil change by $\Big(V(s_i, \bar{s} + \bar{\Delta}_{j\ell}) - V(s_i, \bar{s}) \Big)$ if dominant rival $j \in \mathcal{J}^d$ wins the auction. For each $j \in \mathcal{J}^d$, let $P_{j,-i}(x_\ell, s_i, \bar{s})$ denote the ex ante probability that j wins an auction with characteristics x_ℓ conditional on the event that i is not matched. We can then express i's ex ante surplus from auction arrival as

$$\begin{aligned} <<<<<<<< \textit{LocalChanges} \Pi^0(s_i, \bar{s}) = E[\Pi^m(X_\ell, s_i, \bar{s}) m(X_\ell, s_i, \bar{s})] + \sum_{j \in \mathcal{J}^d, j \neq i} \Big(V(s_i, \bar{s} + \bar{\Delta}_{j\ell}) - V(s_i, \bar{s}) \Big) = \\ - E\left[\sum_{j \in \mathcal{J}^d, j \neq i} \Big(V(s_i, \bar{s} + \bar{\Delta}_{j\ell}) - V(s_i, \bar{s}) \Big) \times H_{j,i}^m(X_\ell, s_i, \bar{s}) \times m(X_\ell, s_i, \bar{s} \right. \\ + E\left[\sum_{j \in \mathcal{J}^d, j \neq i} \Big(V(s_i, \bar{s} + \bar{\Delta}_{j\ell}) - V(s_i, \bar{s}) \Big) \times P_{j,i}^m(X_\ell, s_i, \bar{s}) \times m(X_\ell, s_i, \bar{s} \right. \\ + E\left[\sum_{j \in \mathcal{J}^d, j \neq i} \big(V(s_i, \bar{s} + \bar{\Delta}_{j\ell}) - V(s_i, \bar{s}) \big) \times P_{j,-i}(X_\ell, s_i, \bar{s}) \times (1 - m(X \right. \end{aligned}$$

Let $P_j(x_\ell, s_i, \bar{s})$ denote the probability, before observing any match outcomes, that dominant rival $j \in \mathcal{J}^d$ wins auction ℓ with characteristics x_ℓ given i's tracked information $(s_i, \bar{s})$. Note that, by definition,

$$P_j(x_\ell, s_i, \bar{s}) = P_{j,i}(x_\ell, s_i, \bar{s})m(x_\ell, s_i, \bar{s}) + P_{j,-i}(x_\ell, s_i, \bar{s})(1 - m(x_\ell, s_i, \bar{s})).$$

Letting $R^0(s_i, \bar{s}) = E[R_i^m(X_\ell, s_i, \bar{s})m(X_\ell, s_i, \bar{s})]$, we can therefore simplify $\Pi^0(s_i, \bar{s})$ to

$$\begin{aligned}\Pi^0(s_i, \bar{s}) &= R^0(s_i, \bar{s}) \\ &+ E\left[\sum_{j \in \mathcal{J}^d, j \neq i} \Big(V(s_i, \bar{s} + \bar{\Delta}_{j\ell}) - V(s_i, \bar{s})\Big) \times \big[P_j(X_\ell, s_i, \bar{s}) - H_{j,i}^m(X_\ell, s_i, \bar{s})m(X_\ell, s_i, \bar{s})\big]\right].\end{aligned}$$

In general, the change $\bar{\Delta}_{j\ell}$ in the aggregate state $\bar{s}$ induced by firm j winning auction ℓ will depend both on the firm's current state s_j and on auction characteristics x_ℓ. Importantly, however, any combination of (j, x_ℓ) inducing the same next industry state $\bar{s}' = \bar{s} + \bar{\Delta}_{j\ell}$ will also induce the same perceived continuation value $V(s_i, \bar{s}')$ for firm i. Consequently, letting $\bar{S}(s_i, \bar{s})$ be the set of possible next industry states that could arise following some dominant rival $j \in \mathcal{J}^d, j \neq i$ winning an auction, we can further simplify $\Pi^0(s_i, \bar{s})$ to

$$\Pi^0(s_i, \bar{s}) = R^0(s_i, \bar{s}) + \sum_{\bar{s}' \in \bar{S}(s_i, \bar{s})} V(s_i, \bar{s}')H^0(\bar{s}', s_i, \bar{s}), \tag{20}$$

where $H^0(\bar{s}', s_i, \bar{s})$ is a 'FOC-adjusted transition rate' defined by

$$H^0(\bar{s}', s_i, \bar{s}) = \sum_{j \in \mathcal{J}^d, j \neq i} E\Big[\mathbb{I}[\bar{s} + \bar{\Delta}_{j\ell} = \bar{s}'] \times \big[P_j(X_\ell, s_i, \bar{s}) - H_{j,i}^m(X_\ell, s_i, \bar{s})m(X_\ell, s_i, \bar{s})\big]\Big]. \tag{21}$$

Note that $H^0(\bar{s}', s_i, \bar{s})$ is not the actual transition rate in $\bar{s}'$ induced by dominant rivals winning arriving auctions. Rather, it is this actual transition rate adjusted by the term

$$- \sum_{j \in \mathcal{J}^d, j \neq i} E\Big[\mathbb{I}[\bar{s} + \bar{\Delta}_{j\ell} = \bar{s}']H_{j,i}^m(X_\ell, s_i, \bar{s})m(X_\ell, s_i, \bar{s})\Big],$$

which reflects the ex ante extent to which i internalizes their impact on aggregate transitions through the bid-stage FOC for optimal bidding in auctions with which they are matched. Crucially, however, all objects on the right-hand side of (21) are identified, which implies that $H^0(\bar{s}'; s_i, \bar{s})$ is directly identified. Since $R^0(s_i, \bar{s})$ is also directly identified, it follows that we can express ex ante auction arrival surplus $\Pi^0(s_i, \bar{s})$ as an identified linear function of $V(s_i, \bar{s})$. This is expected in view of Jofre-Bonet and Pesendorfer (2003); the derivation above simply adapts the details of that in Jofre-Bonet and Pesendorfer (2003) to our case.

B.4: CCP representation of $V(s_i, \bar{s})$

Finally, we return to the Bellman equation characterizing the value function $V(s_i, \bar{s})$. In any instant in time, in addition to the events described in the main text, firm i must account for transitions in $\bar{s}$ induced by its own transitions from, to, or among dominant states, as well as (i) transitions in the states of dominant rivals, (ii) transitions of dominant rivals to the fringe, and (iii) transitions of fringe firms to dominance. For each $s_i \in S_i$ and $\bar{s}', \bar{s} \in \bar{S}$, let $\gamma(\bar{s}', s_i, \bar{s})$ denote the perceived aggregate equilibrium transition rate from $(s_i, \bar{s})$ to $\bar{s}'$ induced by transitions in all rivals' states through channels excluding auctions:

$$\gamma(\bar{s}', s_i, \bar{s}) = \gamma_{dd}(\bar{s}', s_i, \bar{s}) + \gamma_{df}(\bar{s}', s_i, \bar{s}) + \gamma_{fd}(\bar{s}', s_i, \bar{s}) + \gamma_{ff}(\bar{s}', s_i, \bar{s}),$$

where γ_{dd} γ_{df}, γ_{fd}, and γ_{ff} denote, respectively, firm i's beliefs about transition rates in $\bar{s}$ induced by rival dominant-to-dominant, dominant-to-fringe, fringe-to-dominant, and fringe-to-fringe transitions through channels excluding auctions. Similarly, let $\bar{\Delta}_i(s_i', s_i, \bar{s})$ denote the change in the MME state $\bar{s}$ associated with a transition by i from s_i to s_i' (with $\bar{\Delta}_i(s_i', s_i, \bar{s}) = 0$ if neither $s_i \in S^d$ or $s_i' \in S^d$). We can then express the Bellman equation defining $V(s_i, \bar{s})$ as

$$\begin{aligned}
\rho V(s_i; \bar{s}) = -c_f(s_i, a) + \lambda_0 R^0(s_i; \bar{s}) + \lambda_\chi \mathbb{I}[z_i = 0]\Pi^\chi(s_i, \bar{s}) \\
+ \lambda_q(s_i)\Big(V(s_i'; \bar{s} + \bar{\Delta}_i(s_i', s_i, \bar{s})) - V(s_i, \bar{s})\Big) \\
+ \sum_{z_i' < z_i} \delta_z(z_i', s_i)\Big(V(s_i'; \bar{s} + \bar{\Delta}_i(s_i', s_i, \bar{s})) - V(s_i, \bar{s}_m)\Big) \\
+ \sum_{\kappa_i'} \delta_\kappa(\kappa_i', s_i)\Big(V(s_i'; \bar{s} + \bar{\Delta}_i(s_i', s_i, \bar{s})) - V(s_i, \bar{s}_m)\Big) \\
+ \sum_{\bar{s}' \in \bar{S}} [\lambda_0 H^0(\bar{s}', s_i, \bar{s}) + \gamma(\bar{s}', s_i, \bar{s})]\Big(V(s_i; \bar{s}') - V(s_i, \bar{s})\Big). \quad (22)
\end{aligned}$$

This again defines a linear system in $V(s_i, \bar{s})$. Collecting terms involving V on the right hand side, we can again express this system in compact matrix form as:

$$\mathbf{TV} = -\mathbf{c}_f + \lambda_0 \mathbf{R^0} + \lambda_\chi \mathbb{I}_\chi \mathbf{\Psi}_\chi + \lambda_\chi \mathbb{I}_\chi \mathcal{E}_\chi \sigma_\chi, \quad (23)$$

where $\mathbf{T}$ is a matrix of known or directly estimable effective transition rates, and the right-hand side is directly identified up to the parameters $\mathbf{c}_f$, $\mathbf{\Psi}_\chi$, σ_χ. The matrix T will again be diagonally dominant, implying that we can solve for $\mathbf{V}$ as a linear function of parameters.

Note that although the form of (23) appears almost identical to the main text, the matrix $\mathbf{T}$ will be substantially different (and typically less sparse). In particular, in addition to actual / perceived equilibrium transition rates, $\mathbf{T}$ will include an adjustment for optimal bidding in auctions in which i participates, captured through the function H^0 defined above. Importantly, however, $\mathbf{T}$, $\mathbf{R^0}$ and $\mathcal{E}_\chi$ can still be consistently estimated, after which estimation of structural parameters can proceed as in the main text.

Appendix C: Solving counterfactuals

C.1: Contraction representation of value function

Although yielding a simple characterization of $V(s_i, \bar{s})$, the right-hand side of equation (1) need not be a contraction mapping. In this subsection, we derive an alternative characterization of $V(s_i, \bar{s})$ which is a contraction mapping. Although this can be obtained by a simple rearrangement of (1), we here provide an alternative derivation which highlights connections between our continuous-time setting and typical discrete-time models.

Toward this end, consider firm i in MME state $(s_i, \bar{s})$. With some abuse of notation, let $\mathcal{Y}$ denote the set of possible events which could change the MME state as perceived by firm i: i.e., auction arrival events, exit opportunity events, each possible own state transition, and each possible aggregate state transition. For each event type $y \in \mathcal{Y}$, let $q(y|s_i, \bar{s})$ denote the arrival rate of the associated event type in state $(s_i, \bar{s})$ from the perspective of firm i; e.g., if y is an auction arrival event, $q(y|s_i, \bar{s}) = \lambda_0$, if y is an exit opportunity event, $q(y|s_i, \bar{s}) = \lambda_\chi \mathbb{I}[z_i = 0]$, if y is a transition the MME state from $\bar{s}_m$ to $\bar{s}_m$, $q(y|s_i, \bar{s}) = \gamma_{mn}$, and similarly for all other firm state and aggregate state transition events. Let $Q(s_i, \bar{s}) = \sum_{y \in \mathcal{Y}} q(y|s_i, \bar{s})$ denote the sum of (perceived) arrival rates for all events which could affect firm i's continuation payoffs. Finally, let $W(y, s_i, \bar{s})$ denote i's expected continuation value from the instant event $y \in \mathcal{Y}$ occurs onward. For example, for an auction arrival event, $W(y, s_i, \bar{s}) = V(s_i, \bar{s}) + \Pi^0(s_i, \bar{s})$, while for a transition from $\bar{s}$ to $\bar{s}'$, $W(y, s_i, \bar{s}) = V(s_i, \bar{s}')$.

Now return to firm i's continuation value $V(s_i, \bar{s})$. Let τ, a random variable, denote the time until the next event $y \in \mathcal{Y}$ as perceived by firm i. In general, we can express $V(s_i, \bar{s})$ as the sum of two components: expected discounted flow payoffs until the time τ until the next event $Y \in \mathcal{Y}$, plus the discounted value from this event on:

$$V(s_i, \bar{s}) = c_f(s_i, \bar{s}) E\left[\int_0^\tau e^{-\rho t} dt | s_i, \bar{s}\right] + E\left[e^{-\rho\tau} W(Y, s_i, \bar{s}) | s_i, \bar{s}\right]. \tag{24}$$

Furthermore, since firm i perceives $\bar{s}$ to follow a Markov jump process, in each instant firm i will perceive τ to follow a Poisson distribution with rate parameter $Q(s_i, \bar{s})$. Moreover, conditional on an event occurring in any instant t, firm i perceives this event to be of each possible type $y \in \mathcal{Y}$ with probability equal to $q(y|s_i, \bar{s})/Q(s_i, \bar{s})$. Importantly, these perceived probabilities do not depend on τ. Bearing these properties in mind, we can simplify (24) to

$$V(s_i, \bar{s}) = c_f(s_i, \bar{s}) E\left[\int_0^\tau e^{-\rho t} dt | s_i, \bar{s}\right] + E[e^{-\rho\tau} | s_i, \bar{s}] \times E\left[W(Y, s_i, \bar{s}) \middle| s_i, \bar{s}\right].$$

Furthermore, again applying properties of Markov jump processes, we have

$$E\left[\int_0^\tau e^{-\rho t} dt \,\middle|\, s_i, \bar{s}\right] = \frac{1}{\rho + Q(s_i, \bar{s})},$$
$$E[e^{-\rho\tau} | s_i, \bar{s}] = \frac{Q(s_i, \bar{s})}{\rho + Q(s_i, \bar{s})},$$

and

$$E\left[W(y, s_i, \bar{s})|s_i, \bar{s}\right] = \sum_{y \in \mathcal{Y}} W(y, s_i, \bar{s}) \frac{q(y|s_i, \bar{s})}{Q(s_i, \bar{s})}. \tag{25}$$

We can thus rewrite the Bellman equation (24) above in a form similar to a discrete-time Bellman equation:

$$\begin{aligned} V(s_i, \bar{s}) &= \frac{c_f(s_i, \bar{s})}{\rho + Q(s_i, \bar{s})} + \frac{Q(s_i, \bar{s})}{\rho + Q(s_i, \bar{s})} \times E\left[W(Y, s_i, \bar{s})|s_i, \bar{s}\right] \\ &= \frac{c_f(s_i, \bar{s})}{\rho + Q(s_i, \bar{s})} + \frac{Q(s_i, \bar{s})}{\rho + Q(s_i, \bar{s})} \times \left(\sum_{y \in \mathcal{Y}} W(y, s_i, \bar{s}) \frac{q(y|s_i, \bar{s})}{Q(s_i, \bar{s})} \right). \end{aligned} \tag{26}$$

The right-hand side of (26) equals a flow payoff term plus a state-specific "discount factor" times an expected continuation value from the next event on. The "discount factor" $\frac{Q(s_i,\bar{s})}{\rho+Q(s_i,\bar{s})}$ represents the expected discounted time until the next event, with more rapid event arrival rates (larger $Q(s_i, \bar{s})$) implying shorter "periods" and thus larger effective discount factors. Furthermore, since V appears only in the final expected continuation value, it is straightforward to show that (26) satisfies Blackwell's sufficient conditions for a contraction mapping, with modulus equal to the maximum state-specific "discount factor":

$$\max_{s_i, \bar{s}} E[e^{-\rho\tau}|s_i, \bar{s}] = \max_{s_i, \bar{s}} \frac{Q(s_i, \bar{s})}{\rho + Q(s_i, \bar{s})} < 1.$$

Iteration on (26) will therefore yield $V(s_i, \bar{s})$ as the unique fixed point.

C.2: Potential solution algorithms

We take as given a function for calculating expected bid-stage profit $\pi_i^b(x_\ell, s_i, s_\ell^b, \bar{s})$ given continuation values. For a Vickrey auction, this expected bid-stage profit calculation is relatively simple; for a low-price auction it would require solving an asymmetric low-bid auction which is a challenging problem in itself.

The derivations in Section D.1 suggest several potential iterative approaches to solving for a MME. The simplest would likely be a simulation-based "real-time dynamic programming" algorithm based on (24), in which transition probabilities and continuation values are updated by (i) simulating τ and Y from the current state given existing policies, (ii) simulating or calculating endogenous instantaneous payoffs associated with these draws, (ii) updating continuation value estimates based on (24), and (iv) updating policies based on updated continuation values. This algorithm would essentially be a continuous-time MME version of the algorithms proposed in Pakes and McGuire (2001) and Ifrach and Weintraub (2017)'s Algorithm 2, and appropriately implemented should eventually converge to an MME.

One disadvantage of this simple approach is that updating (24) based on random draws of τ and Y is effectively using random draws to approximate integrals with closed-form analytic solutions. An alternative, perhaps superior approach would be to instead proceed as follows: (i) simulate τ and Y from current policies, (ii) update estimates of $\Pi^0(s_i, \bar{s})$ and MME transition rates γ based on these draws, (iii) update $V(s_i, \bar{s})$ using one or more

iterations of (26), and (iv) update dynamic policies based on updated continuation values. This approach would still estimate the equilibrium-determined objects $\Pi^0(s_i, \bar{s})$ and γ via simulation, but would use the exact forms for expectations in (26) to update V rather than approximating these via simulation draws of continuation payoffs. Note that, in contrast to typical discrete-time models, the sums in our continuous-time Bellman equation (26) involve relatively few terms. I conjecture that the cost of such a simulation / Bellman algorithm would not be much larger than the pure recursive simulation approach above, while avoiding simulation error due to approximation of expectations with simulation draws.

A still more precise, though likely much more computationally costly, approach would be to additionally calculate $\Pi^0(s_i, \bar{s})$ exactly at every iteration (approximating expectations over auction characteristics with a fixed sample of observed auction x_ℓ's). Using this approach, the only MME object estimated by stochastic simulation would be MME transition probabilities, as these have no analytic closed form. Again, the tradeoff would be greater precision per iteration versus greater iteration cost.

A fourth approach, paralleling Ifrach and Weintraub (2017)'s Algorithm 1, would be to iterate the following steps: (i) Fully solve value functions given beliefs over MME transition probabilities, (ii) given policies, simulate a long time path of the model, and (iii) update MME transition probabilities based on simulated histories. This approach would probably be the most precise, but would involve the highest cost per iteration. For this reason, Ifrach and Weintraub (2017) suggest using a real-time dynamic programming algorithm to find a solution and this full solution approach to verify the solution has converged.

In addition, we may wish to maintain the exact belief parameterizations imposed in estimation (e.g., that beliefs over match profit are formed via a linear regression, and that beliefs over aggregate state transition rates take a particular parametric form). In this case, it should be possible to augment the simulation steps above with belief update steps in which regression coefficients and MME transition parameters are iteratively updated based on new simulation draws. For linear regression, where parameter estimates are given in closed form, this iterative update is relatively easy (one need only update the sums appearing in the standard OLS formula to reflect new observations). For MME transition beliefs, the update would be more complicated as the MLE solution would depend on all simulation draws (not just recent draws). I think, however, that one could borrow techniques from machine learning such as batch gradient descent with averaging to update parameters using information contained only in recent simulation draws.

For the moment, we focus on relatively simpler algorithms which do not aim to maintain the same parametric structure on beliefs in the counterfactual. With these in hand, I conjecture it will be relatively simple to extend the solver to incorporate particular parametric structure on beliefs following the ideas above.

Appendix D: Description of Variables

Experience (past win counts)
We proxy a firm's experience with its past number of TxDOT projects won. Note that here we use all data from 1998 and build histories for firms that enter before January 2000. Each

firm's win count history is updated through our analysis period. Note that 214 out of state firms participate as prequalified firms, and, on average, they bid on projects worth about $10 million. These large firms have a history of winning projects in multiple states, and we assume their initial win count is 25 when they join the TxDOT market for the first time. This win count of 25 corresponds to the average count of Texas based prequalified firms' past win counts. The remaining 32 non-prequalified firms bid on projects worth about $215,000, and we consider them inexperienced firms.

Backlog
A backlog variable is constructed for each month for all planholders in the data set. At the start of the panel in 1997 or for new entrants, each bidder's backlog is initialized to zero. As a bidder wins projects, the dollar value of the project is added to the backlog of the bidder in the month of the bid letting. As the project commences, the backlog is worked off by subtracting the incoming data on project payments. The length of the project is constructed using the calendar day variable. The substantial number of years of data available prior to the analysis sample (1998-2000) allows us to initialize the backlog series with two years of data. Backlog values are converted to January 2000 constant-dollar (in millions) values using TxDOT HCI.

Firm's market share
To measure the size of each firm over time, we construct a proxy for firm capacity.Note that a contractor may request project details only for projects whose total ECEs do not exceed bidding capacity less the firm's backlog. Motivated by this institutional feature, for each firm and month, we first define a proxy for within-month used capacity as the firm's backlog plus the sum of ECEs among projects for which the firm held plans within the month. We then construct a proxy for the firm's capacity in a given month as the maximum, over the past 18 months, of these within month used capacities. We normalize this firm-level capacity variable by the sum of capacities for all firms within the month. We use the resulting variable as a proxy for the market share of each firm over time.

Fully prequalified firm
A firm that has provided an audited financial statement prepared by an independent Certified Public Accountant along with a completed Texas Department of Transportation (TxDOT) 'Confidential Questionnaire' and other required supporting documents.

Registered firm but not fully prequalified
All bidders must at least complete the 'Bidder's Questionnaire' (BQ) form and provide all additional requested information in order to comply with the requirements for bidding. Firms that only fill the BQ can submit bids only on 'waived' project. TxDOT waives the prequalification requirement for small construction, maintenance, and special projects that are valued at less than $300,000.

Firm's distance to the project location
For each contractor, we mapped their address into longitudinal and latitude coordinates and used them when calculating the distance to a project. For the project location, the coordinates of the centroid of the county where the project is listed is used. The distance variable is constructed using the 'vincenty' stata code that calculates distances based on geodesic differences between two points.

Waived auction
A project that does not have the prequalification requirement is categorized as a waived auction. The variable 'waived' takes the value of one for a waived auction and zero otherwise.

Bid

This is the total bid submitted by a bidder for a project. The bid consists of all values for tasks to be completed and sum of these tasks is the total bid.

Engineer's cost estimate
TxDOT engineers' cost estimate for the project. This is the estimated value of all tasks (number of units multiplied by unit price) to be performed in order to complete the project. All estimates are converted to January 2000 constant-dollar (in millions) value using highway cost index (HCI) provided by TxDOT.

Number of tasks to complete the project
The item-level information in a project design plan is a list of all items defining the tasks to be completed. An individual item could be the amount and type of asphalt to be used on a paving project, the length of the guardrail to be installed, the number and type of signage to be placed at a project, the square yards of earth that need to be moved, or the type, size and quantity of trees to be planted. Each item is described by an 8-digit code, a detailed description of the item, a unit of measure for an item (e.g., square meters, linear feet, gallons) and the quantity to be installed.

Days to complete the project (calendar days)
TxDOT estimated the length of the project in calendar days provided with design plans.

Number of planholders
Plan holders are potential bidders who have requested design plans for a project from TxDOT. When a firm requests plans, they become registered planholders. This information is public and available prior to bid letting. Hence all potential bidders know the set of potential rivals. Total number of firms that requested plans are the number of planholders in an auction.

Number of bidders
The number of firms that submitted bids in an auction.

Federal project
Projects are funded by Federal or State funds. The federal project variable takes the value of one for federally funded projects and zero otherwise.

Project types
We categorize projects into six types based on the material shares of a project. The six material groups for projects based on bid items is described by the "Standard Specifications for Construction and Maintenance of Highways, Streets, and Bridges" code book adopted by TxDOT. These six material cost shares are constructed from detailed information on bid items and the project's overall engineering cost estimate. These include: 1) asphalt surface work (i.e., hot-mix asphalt); 2) earth work (i.e., excavation); 3) miscellaneous work (i.e., mobilization); 4)structures (bridges); 5) subgrade (i.e., Proof Rolling); and 6) lighting and signaling work (i.e., highway sign lighting fixtures).

Project zones
We also categorize projects into five zones identified by TxDOT in order to control for physical features of areas as they require different grades of material to complete a project.

Additional days
For every completed project, TxDOT reports the number of additional (or fewer) days taken to complete the project beyond the original estimated days.

Final pay

The final total value a winning firm received after completion of the project. This pay may differ from the winning bid submitted at bid letting due to under or overruns and project modifications.

The sum of past win counts of all plan holders in the auction
This variable is the sum of all planholders' win counts for a given auction.

The sum of the backlog of all planholders in the auction This is the sum of planholders' backlog in each auction.

The sum of market share of all plan holders in the auction This is the sum of the firm-level market share for a given auction.

Rivals' minimum distance to the project location
For each planholder in a given auction, we construct its rivals' minimum distance to the project location based on rivals' distance to the project location.

The total value of the backlog of all planholders in the market
This is the sum of all active firms' backlog in the market in each month.

The total number of competitors in the market
This is the sum of the number of active firms in the market for each month.

The total number of prequalified competitors in the market
This is the sum of the number of active prequalified firms in the market in a given month.

The total value of new building permits in Texas
This is the total value of new building permits approved for single-family, 2-4 family and 5-plus family units in Texas. These monthly data are provided by the Texas A & M University Real Estate Research Center (https://www.recenter.tamu.edu/data/building-permits)