

ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

<https://www.econometricsociety.org/>

Econometrica, Vol. 91, No. 4 (July, 2023), 1205–1259

SCALING AUCTIONS AS INSURANCE: A CASE STUDY IN INFRASTRUCTURE PROCUREMENT

VALENTIN BOLOTNYY

Hoover Institution, Stanford University

SHOSHANA VASSERMAN

Stanford Graduate School of Business and NBER

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

SCALING AUCTIONS AS INSURANCE: A CASE STUDY IN INFRASTRUCTURE PROCUREMENT

VALENTIN BOLOTNYY

Hoover Institution, Stanford University

SHOSHANA VASSERMAN

Stanford Graduate School of Business and NBER

Most U.S. government spending on highways and bridges is done through “scaling” procurement auctions, in which private construction firms submit unit price bids for each piece of material required to complete a project. Using data on bridge maintenance projects undertaken by the Massachusetts Department of Transportation (MassDOT), we present evidence that firm bidding behavior in this context is consistent with optimal skewing under risk aversion: firms limit their risk exposure by placing lower unit bids on items with greater uncertainty. We estimate the amount of uncertainty in each auction, and the distribution of bidders’ private costs and risk aversion. Simulating equilibrium item-level bids under counterfactual settings, we estimate the fraction of project spending that is due to risk and evaluate auction mechanisms under consideration by policymakers. We find that scaling auctions provide substantial savings relative to lump sum auctions and show how our framework can be used to evaluate alternative auction designs.

KEYWORDS: Scoring auctions, procurement, risk-averse bidders, multidimensional bids, empirical market design.

1. INTRODUCTION

INFRASTRUCTURE INVESTMENT underlies nearly every part of the American economy and constitutes hundreds of billions of dollars in public spending each year.¹ However, investments are often directed into complex projects that experience unexpected changes. Project uncertainty can be costly to the firms that implement construction—many of whose businesses are centered on public works—and to the government. The extent of firms’ risk exposure depends not only on project design, but also on the mechanism used to allocate contracts. Contracts with lower risk exposure may be more lucrative, and thus invite more competitive bids. As such, risk sharing between firms and the government can play a significant role in the effectiveness of policies meant to reduce taxpayer expenditures.

We study the mechanism by which contracts for construction work are allocated by the Highway and Bridge Division of the Massachusetts Department of Transportation

Valentin Bolotnny: vbolotnny@stanford.edu

Shoshana Vasserman: svass@stanford.edu

This paper was a chapter in our PhD dissertations. We are indebted to our advisors Ariel Pakes, Elie Tamer, Robin Lee, Edward Glaeser, Claudia Goldin, Nathaniel Hendren, Lawrence Katz, and Andrei Shleifer for their guidance and support. We are deeply grateful to Ali Yurukoglu, Andrzej Skrzypacz, Peter Reiss, Susan Athey, and especially Paulo Somaini and Zi Yang Kang for their insightful comments and suggestions, as well as to Charles Margossian, Philip Greengard, and especially Cameron Pfiffer for their invaluable help with computational aspects of this paper. Finally, we thank Steve Poflak, Jack Moran, Frank Kucharski, Michael McGrath, and Naresh Chetpelly, among other generous public servants at MassDOT for their patience and support throughout this project.

¹According to the CBO, infrastructure spending typically accounts for roughly \$416B or 2.4% of GDP annually across federal, state, and local levels. Of this, \$165B—40%—is spent on highways and bridges alone.

(MassDOT or “the DOT”). Along with 40 other states, MassDOT uses a *scaling auction*, in which bidders submit a unit price bid for each item in a comprehensive list of tasks and materials required to complete a project. The winning bidder is determined by the lowest sum of unit bids multiplied by item quantity estimates produced by MassDOT project designers. This winner is then paid based on the quantities ultimately used in completing the project.

Scaling auctions thus have several key features. First, they are widespread and common in public infrastructure procurement. Second, they collect bids over units (i.e., tasks and materials) that are standardized and comparable across auctions. Third, they implement a partial sharing of risk between the government and private contractors.

To study auction design in this setting, we specify and estimate a model of bidding in scaling auctions with risk-averse bidders. Our model characterizes equilibrium bids in two separable steps: an “outer” condition that ensures that a bidder’s *score*—the weighted sum of unit bids that is used to determine the winner of the auction—is optimally competitive with respect to the opposing bidders, and an “inner” condition that ensures that the unit bids chosen to sum up to the equilibrium score maximize the expected utility of winning. As first noted by [Athey and Levin \(2001\)](#) in the context of timber auctions, the “inner” condition constitutes a portfolio optimization problem for bidders: equilibrium unit bids distribute a bidder’s score across different items, trading off *higher expected profits* from high bids on items predicted to overrun against *higher risk* from low bids on other items.

The separability of the “inner” and “outer” problems yields a useful property: given an observation of a bidder’s equilibrium score, her equilibrium unit bids are fully specified by the characterization of her (“inner”) portfolio problem. Previous work on scoring auctions has exploited such separability to succinctly characterize equilibrium play. Studying auctions with quasilinear scoring rules and risk-neutral bidders, [Asker and Cantillon \(2008\)](#) show that equilibrium outcomes can be characterized through a one-dimensional auction over scores, even when bidder types are high-dimensional. Closer to our setting, [Athey and Levin \(2001\)](#) use separability to argue that observations of profitable *skewing*—placing higher bids on items that ultimately overrun—can be interpreted as evidence that bidders were better informed than the auctioneer. In this paper, we take this logic further and show that the solutions to bidders’ portfolio problems—subject to their scores—can be used to estimate the distribution of bidder types and simulate the outcomes of counterfactual DOT policies.

Using a detailed data set obtained through a partnership with MassDOT, we establish the patterns of bidding behavior that motivate our approach. For each auction in our study, we observe the full set of items involved in construction, along with the ex ante DOT estimate and ex post realization of the quantity of each item, a DOT estimate of the market unit rate for the item, and the unit price bid that each bidder who participated in the auction submitted. As in prior work, we show that contractors skew their bids, placing high unit bids on items that tend to overrun the DOT quantity estimates and low unit bids on items that tend to underrun. This suggests that contractors are generally able to predict the direction of ex post changes to project specifications, and bid so as to increase their ex post earnings.

Furthermore, our data suggest that contractors are risk averse. As noted in [Athey and Levin \(2001\)](#), risk-neutral bidders would be predicted to submit “penny” bids—unit bids of essentially zero—on all but the items that are predicted to overrun by the largest amount. By contrast, the vast majority of unit bids observed in our data are interior (i.e., nonextremal), even though no significant penalty for penny bidding has ever been exercised. We show that while contractors bid higher on items predicted to overrun, holding

all else fixed, they also bid lower on items that are more uncertain. This suggests that contractors optimize not only with respect to expected profits, but also with respect to the risk that any given expectation will turn out to be wrong.

Bidder risk aversion, combined with inherent project risk, has significant implications for DOT spending, as well as for the efficacy of policies to reduce it. Risk-averse bidders internalize a utility cost from uncertainty and require higher overall bids in order to insure themselves sufficiently to be willing to participate. As such, auction rules that decrease bidders' exposure to significant losses can be effective toward lowering overall bids and subsequently lowering DOT payments to the winning bidder.

In order to gauge the level of risk and risk aversion in our data, we estimate a structural model of uncertainty and optimal bidding. In the first stage of our estimation procedure, we use the history of predicted and realized item quantities to fit a model of bidder uncertainty over item quantity realizations. In the second stage, we construct a Generalized Method of Moments (GMM) estimator for bidders' costs and risk aversion in each project. Our estimator relies only on predictions of optimal unit bids at the auction-bidder-item level, evaluated from each bidder's portfolio problem subject to the constraint implied by her observed score. As such, our identification strategy leverages granular variation in project composition (e.g., which items are needed, at what market rate, and in what quantities), in addition to more standard project characteristics such as the identity of the designing engineer. As our predictions of optimal bids capture the bidders' competitive considerations entirely through their scores—which are taken as data—our estimation approach does not require strong assumptions about bidders' beliefs about their opponents, nor does it require exogenous variation in the composition of bidders across auctions.

We use our structural estimates to evaluate the cost of uncertainty to the DOT, as well as the performance of scaling auctions relative to alternatives used in other procurement settings. Using an independent private values (IPV) framework and a calibrated model of endogenous participation in the spirit of Samuelson (1985), we simulate equilibrium outcomes under a counterfactual setting in which uncertainty about item quantities is reduced to zero. When bidder predictions are held fixed—the only change is that uncertainty about these predictions is eliminated—we find that DOT spending decreases by 14.5% for the median auction. This suggests that project uncertainty contributes to a substantial risk premium.

However, scaling auctions perform quite well on the whole, given the level of uncertainty in these projects. The most common alternative mechanism for procurement is a *lump sum* auction, in which bidders commit to a total payment at the time of the auction and are liable for all implementation costs afterward. Lump sum auctions require less planning by the DOT, and they incentivize bidders to be economical when they can be. But for projects that are highly standardized and monitored—such as the bridge projects in our data—lump sum auctions primarily shift risk from the DOT onto the risk-averse bidders. Seen in this light, scaling auctions provide a powerful lever for the DOT to lower its costs: not only do scaling auctions provide insurance by reimbursing bidders for every item that is ultimately used, but they also allow bidders to hedge their risks through portfolio optimization.

Our simulations suggest that moving from the scaling format to a lump sum format would increase DOT spending by 42% for the median auction. However, this result compounds two opposing effects. Bidders in a lump sum auction need to bid higher overall in order to compensate for their increased liability. These higher bids translate to higher costs for the DOT. On the other hand, higher liability may also cause the least competitive

bidders to be priced out and choose not to participate in the auction at all. This reduces the number of participating bidders on average, but increases the competitiveness of the bidders who do participate. In our sample, the marginal bidder willing to participate under the lump sum format was 20% more cost-efficient and 28% less risk averse than the marginal bidder under the scaling format. This positive selection effect cuts the overall cost of moving to a lump sum format by *more than half*: were participation held fixed, lump sum auctions would be 96% more costly to the DOT than scaling auctions.

Given these results, we ask whether scaling auctions can be further improved through a policy that might reasonably be considered by the DOT. A hybrid format in which bidders commit to a fixed payment at the time of bidding, but are able to renegotiate for a higher payment *ex post*, eliminates most of the added DOT costs from lump sum liability. In our sample, renegotiation with 2:1 bargaining power after a lump sum auction reduces added DOT costs to 14%, while renegotiation with equal bargaining power reduces added costs to only 8.5%. Still, both renegotiation options increase costs relative to the baseline scaling auction, and neither affects the distribution of participants substantially. As we do not find evidence of sufficient moral hazard to overturn these results in our setting, we conclude that switching to any type of lump sum format is unlikely to improve upon the status quo.

Furthermore, while we find a substantial risk premium by eliminating uncertainty holding everything else fixed, a policy to reduce uncertainty—through training or directives, for instance—may not be very effective at reducing costs in practice. Uncertainty in our data is fairly symmetric: underruns and overruns both occur frequently, both in quantities and in DOT spending. As such, when we compare the no-uncertainty counterfactual against the status quo, DOT costs *increase* by nearly 2% for the median auction once changes in bidder participation are accounted for. This is because eliminating uncertainty gives bidders access to the exact quantities that will ultimately be used, allowing them to avoid making “mistakes” (from an *ex post* perspective) that had benefited the DOT under uncertainty. In many cases, the DOT cost from this difference in predictions counteracts the savings from the elimination of the risk premium. Thus, we also conclude that the benefit from policies to reduce uncertainty may not hold up in light of practical considerations.

2. RELATED LITERATURE

Strategic bid skewing in scaling auctions has been documented in various contexts where bidders may be better informed than the auctioneer. Studying U.S. timber auctions, [Athey and Levin \(2001\)](#) first established that positive correlations between (dollar) overbids and (unit) overruns in auction data could be interpreted as evidence that bidders are able to predict which components of their bids will overrun. [Bajari, Houghton, and Tadelis \(2014\)](#) made a similar observation in the context of highway paving procurement auctions in California. However, neither paper evaluates the welfare impact of bid-skewing or the underlying uncertainty that causes it.

Bidders who are risk-neutral, as in the model proposed by [Bajari, Houghton, and Tadelis \(2014\)](#), would be predicted to skew “completely”—that is, bid high on one component of the project and zero on all the others—unless they face an additional incentive not to do so. Moreover, absent such an incentive, there is no welfare cost to skewing whatsoever: were the government to perfectly predict quantities such that there are no overruns, the ultimate payment to the winning bidder would be the same. [Bajari, Houghton, and Tadelis \(2014\)](#) accounts for the lack of complete skewing in their data by imposing a

penalty on unit bids that increases in both the distance between the bid and the government's unit cost estimate and the distance between the ex ante unit quantity estimate and the ex post realization of that quantity. While this enables [Bajari, Houghton, and Tadelis \(2014\)](#) to estimate average adaptation cost multipliers and calibrate the cost of ex post renegotiation, the penalty function coefficient is not found to be significantly different from zero, and no bidder-specific types or counterfactual strategies are estimated.

As in [Athey and Levin \(2001\)](#), our paper argues that the absence of complete skewing is primarily driven by risk aversion. Our model of risk-averse bidding predicts that unit bids will be skewed both as a function of bidders' predictions of ex post quantities and the amount of uncertainty in each prediction. The heart of our paper rests in the resulting portfolio optimization problem. This problem determines the spread of unit bids for each score that a bidder submits, and consequently, both the bidder's private value for winning the auction and the government's ex post payment to the bidder if she wins—both of which differ from the score itself.

The portfolio characterization of bid skewing has several key implications for the analysis of scaling auctions. First, it allows us to construct reduced form correlation tests for risk aversion: much as a positive correlation between overbids and overruns is evidence of bidder information, a negative correlation between absolute markups and component-level uncertainty is indicative of risk aversion. Second, it provides a novel channel for identification of bidder and auction-level model parameters. Our identification strategy differs from the canonical approaches of [Guerre, Perrigne, and Vuong \(2009\)](#) (GPV), [Campo, Guerre, Perrigne, and Vuong \(2011\)](#) and [Campo \(2012\)](#). Like these papers, we make use of functional form assumptions such as the CARA utility function. However, whereas their approaches rely on the optimality of single-dimensional bids with respect to the probability of outcompeting other bidders—analogous to the first-order condition characterizing the optimal *score* in our model—our approach uses the optimality of the composition of unit bids to maximize the value of executing a contract *conditional* on each bidder's score.

This has important implications for the assumptions about equilibrium play that are required. The Campo and GPV approaches require bids to be interpreted as equilibrium outcomes of an explicit competitive bidding game—whether a symmetric IPV game or an asymmetric affiliated values game. By contrast, our identification approach is agnostic to the competitive conditions under which each bidder's score is chosen. Subject to comparatively weak conditions that guarantee the separability of the deterministic portfolio optimization problem from the equilibrium problem of choosing a score for each bidder, our identification strategy is robust to assumptions about the mapping between bidder types and equilibrium scores. These assumptions include correlation between cost efficiency and risk aversion, the possibility of dynamic considerations and even collusion. This does not mean that we circumvent the nonidentification results detailed in [Guerre, Perrigne, and Vuong \(2009\)](#): our approach applies a parametric characterization of bidders' utility and relies on exogenous variation in the distribution of contract values across auctions. However, the particular assumptions applied are different: instead of assumptions on bidders' beliefs about each other, we use assumptions on bidders' beliefs about project characteristics. This set of assumptions may be preferable in a highly standardized infrastructure procurement setting such as ours, where historical information is publicly available and bidders are often industry veterans, but where inherent uncertainty about the underlying physical conditions at each project site is high.

Finally, our portfolio approach facilitates counterfactual analyses of alternative auction rules. Because bidders are not paid the *score* that they compete with—but rather an ex

ante uncertain transformation of their unit bids—a prediction of counterfactual scores that does not model the relationship between scores and unit bids would be insufficient to generate predictions for government costs or welfare. Using our model, we evaluate policies to reduce uncertainty regarding item quantities and to precommit to payments at the time of bidding.

Our paper contributes to a substantial literature on the efficiency of infrastructure procurement auctions. Closest to us is [Luo and Takahashi \(2023\)](#), a contemporary paper that studies infrastructure procurement by the Florida DOT. Like us, [Luo and Takahashi \(2023\)](#) considers risk-averse bidders and compares scaling auctions against lump sum auctions. However, this paper follows a Campo/GPV-style approach for identification and reduces the project components that receive unit bids (which number 67 for the median auction in our data set) into two aggregates—one aggregate certain component and one aggregate uncertain component—for estimation. As such, while [Luo and Takahashi \(2023\)](#) offers novel evidence of risk aversion and the costliness of lump sum auctions in settings with high uncertainty, we view our analyses as complementary in methodology and contribution.

More generally, our paper builds on a rich literature on scoring auctions. While the theoretical results for risk-neutral bidders in [Che \(1993\)](#) and [Asker and Cantillon \(2008\)](#) do not apply to our model directly, the separability of equilibrium bidding into an “outer” score-setting stage and an “inner” portfolio-maximizing stage in our model is analogous to the separability of quality provision and bidding. Our paper also relates to the theoretical literature on optimal mechanism design. While we focus on “practical” mechanisms—ones that do not require knowledge of the bidder-type distribution, for instance—it is possible to characterize the theoretically *optimal* mechanism for our setting by applying the characterization in [Maskin and Riley \(1984\)](#) and [Matthews \(1987\)](#) to our framework.

3. SCALING AUCTIONS WITH MASSDOT

Like most other states, Massachusetts manages the construction and maintenance for its highways and bridges through its Department of Transportation. In order to develop a new project, MassDOT engineers assemble a detailed specification of what the project will entail. This includes an itemized list of every task and material (item) that is necessary to complete the project, along with estimates for the quantity with which it will be needed and a market unit rate for its cost. The itemized list of quantities is then advertised to prospective bidders.

In order to participate in an auction for a given project, a contractor must first be prequalified by MassDOT. Prequalification entails that the contractor is able to complete the work required, given their staff and equipment. Notably, it generally does not depend on past performance. In order to submit a bid, a contractor posts a unit price for each of the items specified by MassDOT. Since April 2011, all bids have been processed through an online platform, Bid Express, which is also used by 40 other state DOTs. All bids are private until the completion of the auction.

Once an auction is complete, each contractor is given a score, computed by the sum of the product of each item’s estimated quantity and the contractor’s unit-price bid for it. The bidder with the lowest score is then awarded a contract to execute the project in full. In the process of construction, it is common for items to be used in quantities that deviate from MassDOT specifications. All changes, however, must be approved by an onsite MassDOT manager. The winning contractor is ultimately paid the sum, across items, of the product of her unit price bid and the *actual* quantity of the item that was used. Unit

prices are almost never renegotiated. However, there is a mechanical price adjustment on certain commodities such as steel and gasoline if their market prices fluctuate beyond a predefined threshold (typically 5%).²

MassDOT reserves the right to reject bids that are heavily skewed. However, this has never been successfully enforced and most bids violate the condition that should trigger rejection.³ While MassDOT has entertained other proposals to curtail bid skewing, such as a 2017 push to require minimum unit bids, these efforts have thus far not been successful.

4. DATA AND REDUCED FORM RESULTS

Our data come from MassDOT and cover highway and bridge construction and maintenance projects undertaken by the state from 1998 to 2015. We work with projects for which MassDOT has digital records on (1) identities of the winning and losing bidders; (2) bids for the winning and losing bidders; and (3) data on the actual quantities used for each item. 2513 projects meet these criteria, 440 of which are related to bridge work. We focus on bridge projects for this paper, as these projects are particularly prone to item quantity adjustments.

All bidders who participate in auctions for these projects are able to see, *ex post*, how everyone bid on each item. In addition, all contractors have access to summary statistics on past bids for each item, across time, and location. Officially, all interested bidders find out about the specifications and expectations of each project at the same time, when the project is advertised (a short while before it opens up for bidding). Only those contractors who have been prequalified at the beginning of the year to do the work required by the project can bid on the project. Thus, contractors do not have a say in project designs, which are furnished either in-house by MassDOT or by an outside consultant.

Once a winning bidder is selected, project management moves under the purview of an engineer working in one of six MassDOT districts around the state. This Project Manager assigns a Resident Engineer to monitor work on a particular project out in the field and to be the first to decide whether to approve or reject underruns, overruns, and Extra Work Orders (EWOs). The full approval process of changes to the initial project design involves several layers of review. Underruns and overruns, as the DOT defines them and as we will define them here, apply to the items specified in the initial project design and refer to the difference between the actual item quantities that were used and the *ex ante* DOT estimates. EWOs refer to work done outside of the scope of the initial contract design and are most often negotiated as lump sum payments from the DOT to the contractor. For the purposes of our discussion and analyses, we will focus on underruns and overruns in bridge construction and maintenance projects.

Table I provides summary statistics for the bridge projects in our data set. We measure the extent to which MassDOT overpays projected project costs in two ways. First, we consider the difference between what the DOT ultimately pays the winning bidder and the DOT's initial estimate of what it will pay at the conclusion of the auction. Summary statistics for this measure are presented in the "Net Overcost (DOT Quantities)" row of Table I. While it seems as though the DOT is saving money on net, this is a misrepresentation of the costs of bid skewing. The initial estimate—which uses the DOT's *ex ante* quantity estimates and corresponds to the winning bidder's *score* in our model—is not

²See <https://www.mass.gov/service-details/massdot-special-provisions> for details.

³See Section E in the Online Appendix (Bolotny and Vasserman (2023)) for a detailed discussion.

TABLE I
SUMMARY STATISTICS FOR THE 440 MASSDOT BRIDGE PROJECTS AUCTIONED, 1998–2015.

Statistic	Mean	St. Dev.	Pctl(25)	Median	Pctl(75)
Project Length (Estimate)	1.53 years	0.89 years	0.88 years	1.48 years	2.01 years
Project Value (DOT estimate)	\$2.72 million	\$3.89 million	\$981,281	\$1.79 million	\$3.3 million
# Bidders	7	3	4	6	9
# Types of Items	68	37	37	67	92
Net Overcost (DOT quantities)	−\$286,245	\$2.12 million	−\$480,487	−\$119,950	\$167,933
Net Overcost (Ex post quantities)	−\$26,990	\$1.36 million	−\$208,554	\$15,653	\$275,219
Percent Overcost (Ex post quantities)	8.46%	36.14%	−12.35%	1.67%	23.28%
Extra Work Orders	298,796	295,173	78,775	195,068	431,188

necessarily representative of the payments that the bidder expects upon winning. Sophisticated bidders anticipate changes from the initial DOT estimates, and bid accordingly to maximize their ex post payments. As such, a more appropriate metric is to compare the amount that was ultimately spent in each project against the dot product of the DOT’s unit cost estimates and the actual quantities used. This is presented in dollars in the “Net Overcost (Ex post quantities)” row of Table I, and in percent of the final cost paid out in “Percent Overcost (Ex post quantities)”.⁴ The median overpayment by this metric is about \$15,000 (1.67%), but the 25th and 75th percentiles are about −\$210,000 (−12.35%) and \$275,000 (23.28%). Figure 1 shows the spread of overpayment across projects. As we will show in our counterfactual section, the distribution of overpayment corresponds to the potential savings from the elimination of risk.

Bidder Characteristics. There are 2883 unique project-bidder pairs (i.e., total bids submitted) across the 440 projects that were auctioned off. There are 116 unique firms that participate, albeit to different degrees. We divide them into two groups: “common” firms, which participate in at least 30 auctions within our dataset, and “rare firms,” which partic-




FIGURE 1.—Net overcost (ex post quantities) across bridge projects.

⁴Note that the average “Percent Overcost (Ex post quantities)” in Table I is the average percent of costs, rather than the ratio of the average net overcost to average total cost.

TABLE II
COMPARISON OF FIRMS PARTICIPATING IN <30 VS. 30+ AUCTIONS.

	Common Firm	Rare Firm
Number of Firms	24	92
Total Number of Bids Submitted	2263	620
Mean Number of Bids Submitted Per Firm	94.29	6.74
Median Number of Bids Submitted Per Firm	63.0	2.5
Total Number of Wins	351	89
Mean Number of Wins Per Firm	14.62	0.97
Median Number of Wins Per Firm	10	0
Mean Bid Submitted	\$2,774,941	\$4,535,310
Mean Ex post Cost of Bid	\$2,608,921	\$4,159,949
Mean Ex post Overrun of Bid	9.7%	21.97%
Percent of Bids on Projects in the Same District	28.19%	15.95%
Percent of Bids by Revenue Dominant Firms	51.67%	11.80%
Mean Specialization	24.44	2.51
Mean Capacity	10.38	2.75
Mean Utilization Ratio	53.05	25.50

ipate in fewer than 30 auctions. We retain individual identifiers for each of the 24 common firms, but group the 92 rare firms together for the purposes of estimation. Common firms constitute 2263 (78%) of total bids submitted and 351 (80%) of auction victories.

Although there is little publicly available financial information about them, the firms in our data are by and large relatively small, private, family-owned businesses. Table II presents summary statistics of the two firm groups. The mean (median) common firm submitted bids to 94.29 (63) auctions and won 14.62 (10) of them. The mean total bid (or score) is about \$2.8 million, while the mean ex post DOT cost implied by the firm's unit bids is \$2.6 million. The mean ex post cost overrun is 9.73%. By contrast, the mean (median) rare firm submitted bids to 6.74 (2.5) auctions and won 0.97 (0) of them. The mean total bid and ex post scores are quite a bit larger for rare firms—\$4.5 million and \$4.2 million, respectively. This is reflected in a substantially larger ex post overrun: 21.97% on average.

In addition to the firm's identity, there are a number of factors that may influence its competitiveness in a given auction. While we do not consider a structural interpretation for these factors in our model, we treat them as characteristics that help explain heterogeneity in costs and risk aversion across auctions and firms. One such factor is the firm's distance from the worksite. Although we do not observe precise locations for each project, we observe which of the six geographic districts under MassDOT jurisdiction each project belongs to, as well as the location of each bidder's headquarters. Using this, we proxy for distance by assigning each project-bidder pair an indicator for whether the project is located in the same district as the bidder's headquarters. Among common firms, 28.19% of bids were on projects that were located in the same district as the bidding firm's headquarters. By contrast, only 15.95% of bids among rare firms were in matching districts.

Another factor is specialization or experience with a particular type of project. We calculate the specialization of a project-bidder pair as the share of auctions of the same project type that the bidding firm bid on within our data set. Our data involve three project types, according to DOT taxonomy: Bridge Reconstruction/Rehabilitation, Bridge Replacement, and Structures Maintenance. The mean specialization of a common firm is 24.44%, while the mean specialization of a rare firm is 2.51%. As projects have varying

sizes, we compute a measure of specialization in terms of project revenue as well. We define a revenue-dominant firm (within a project-type) as a firm that has been awarded more than 1% of the total money spent by the DOT across projects of that project type. Among common firms, 51.67% of bids submitted were by firms that were revenue dominant in the relevant project type; among rare firms, the proportion of bids by revenue dominant firms is 11.8%.

A third factor of competitiveness is each firm's capacity—the maximum number of DOT projects that the firm has ever had open while bidding on another project—and a fourth factor is its utilization—the share of the firm's capacity that is filled when it is bidding on the current project. We measure capacity and utilization with respect to all MassDOT projects recorded in our data—not just bridge projects. The mean capacity is 10.38 projects among common firms and 2.75 projects among rare firms. This suggests that rare firms generally have less business with the DOT, either because they are smaller in size, or because the DOT constitutes a smaller portion of their operations. The mean utilization ratio, however, is 53.05% for common firms and 25.5% for rare firms. This suggests that firms in our data are likely to have ongoing business with the DOT at the time of bidding and are likely to have spare capacity during adjacent auctions that they did not participate in. While we do not model dynamic considerations regarding capacity constraints directly, we find our measure of capacity to be a useful metric of the extent of a firm's dealings with the DOT, as well as of its size.


Quantity Estimates and Uncertainty. As we discuss in Section 8, scaling auctions mitigate DOT costs by enabling risk-averse bidders to insure themselves against uncertainty about the item quantities that will ultimately be used for each project. The welfare benefit is particularly strong if the uncertainty regarding ex post quantities varies across items within a project, and especially so if there are a few items that have particularly high variance. When this is the case, bidders in a scaling auction can greatly reduce the risk that they face by placing minimal bids on the uncertain items (and higher bids on more predictable items).

Our data set includes records of 2985 unique items, as per MassDOT's internal taxonomy. Spread across 440 projects, these items constitute 29,834 unique item-project pairs. Of the 2985 unique items, 50% appear in only one project. The 75th, 90th, and 95th percentiles of unique items by number of appearances in our data are 4, 16, and 45 auctions, respectively.

For each item t , in every auction, we observe the quantity with which the DOT predicted it would be used at the time of the auction— q_t^e in our model—the quantity with which the item was ultimately used— q_t^a —and a DOT engineer's estimate of the market rate for the unit cost of the item. The DOT quantities are typically inaccurate: 76.7% of item observations in our data had ex post quantities that deviated from the DOT estimates.

Figure 2(a) presents a histogram of the percent quantity overrun across item observations. The percent quantity overrun is defined as the difference of the ex post quantity of an item observation and its DOT quantity estimate, normalized by the DOT estimate: $\frac{q_t^a - q_t^e}{q_t^e} \times 100$. In addition to the 23.3% of item-project observations in which quantity overruns are 0%, another 18% involve items that are not used at all (so that the overrun is equal to -100%). The remaining overruns are distributed more or less symmetrically around 0%.

Ex post deviations from DOT quantity estimates are caused by a number of different mechanisms. Some deviations arise from standard procedures. For instance, as ex ante DOT estimates are used for budgeting purposes, there may be reason for adjusting the



(a) Histogram of the percent quantity overrun across item-project pairs.

(b) Scatterplot of standard deviation vs. mean quantity overrun by item.

FIGURE 2.—Descriptives of quantity overruns across items.

quantities of certain items after the design stage. One example is concrete, which is heavily used, has quantities that are difficult to predict precisely, and often overruns in our data. It is also common for the DOT to include certain items that are unlikely to be used at all—just in case—in order to support its policy of avoiding ex post renegotiation. Prominent examples of such items include flashing arrows and illumination for night work. While mechanisms of this sort are largely systemic, there remains a substantial amount of variation in ex post quantities simply due to the inherent uncertainty entailed in construction. A large fraction of Massachusetts bridges are structurally deficient, making it difficult to ascertain the exact severity of their condition prior to construction. Based on our conversations with DOT engineers, qualified bidders are all aware of these mechanisms, and are generally thought to have better specialized knowledge and quality predictive software than the DOT.

Quantity overruns often vary across observations of the same item in different auctions. Figure 2(b) plots the mean percent quantity overrun for each unique item with at least 2 observations against its standard deviation. While a few items have standard deviations close to 0, the majority of items have standard deviations that are as large or larger than the absolute value of their means. That is, the percent overrun of the majority of unique items varies substantially across observations. While this is a coarse approximation of the uncertainty that bidders face with regard to each item—it does not take item or project characteristics into account, for example—it is suggestive of the scope of risk in each auction.

Reduced Form Evidence for Risk-Averse Bid Skewing. As in [Athey and Levin \(2001\)](#) and [Bajari, Houghton, and Tadelis \(2014\)](#), the bids in our data are consistent with a model of similarly informed bidders who bid strategically to maximize expected utility. In Figure 3(a), we plot the relationship between quantity overruns and the percent by which each item was overbid above the DOT market rate (“blue book”) cost estimate.⁵ We do this for both the winning bidder and the second place bidder.⁶ The binscatter is residu-

⁵In public procurement, the term “blue book” is commonly used to refer to industry standard prices (see, e.g., <https://www.transit.dot.gov/funding/procurement/third-party-procurement/subcontracts>).

⁶The percent overbid of an item is defined as $\frac{b_t - c_t}{c_t} \times 100$, where b_t is the bid on item t and c_t is the DOT market rate estimate of item t . The percent quantity overrun is defined as in Figure 2(a).

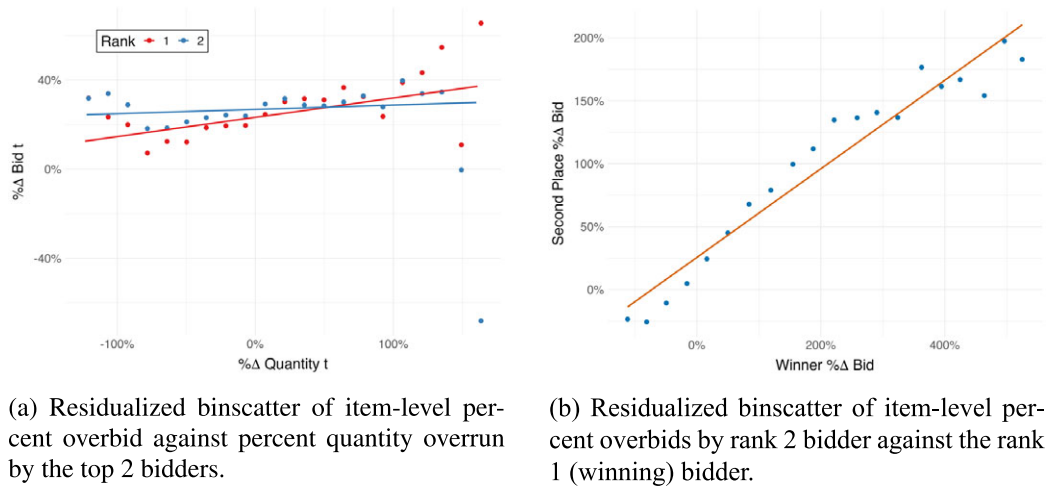


FIGURE 3.—Comparison of item-level overbids by the top two bidders.

alized. In order to obtain it, we first regress percent overbid on a range of controls and obtain residuals. We then regress percent overrun on the same controls and obtain residuals. Finally, to obtain the slope, we regress the residuals from the first regression on the residuals from the second. Controls include the DOT estimate of total project cost, the initially stated project length in days, and the number of participating bidders, as well as fixed effects for: item IDs, the year in which the project was opened for bidding, the project type, resident engineer, project manager, and project designer. Specifications that exclude item fixed effects or include an array of additional controls produce very similar slopes.⁷ We use a similar procedure for all residualized binscatters in this section.

As Figure 3(a) demonstrates, there is a significant positive relationship between percent quantity overruns and percent overbids by the winning bidder. A 1% increase in quantity overruns corresponds to a 0.086% increase in overbids on average. Higher bids on overrunning items correspond to higher earnings *ex post*. Thus, as higher bids correspond to items that overran in our data, we conclude that the winning bidder is able to correctly predict which items will overrun the DOT estimates on average, and to skew accordingly.

Furthermore, Figure 3(a) shows that losing bidders generally skew their bids in the same direction as winning bidders. With the exception of a few outlying points, the top two bidders both overbid on items that wound up overrunning on average.⁸ This suggests that the winning and second place bidder are similarly able to predict overruns. In Appendix E, we show that this pattern holds for bidders ranked 3 and 4 as well, and that the average percent overbids in each bin are even closer together when we restrict our comparison to projects in which the top two bidders submit similar total scores, and thus likely have similar private costs and risk tolerance.

While our data suggest that bidders do engage in bid skewing, there is no evidence of *total* bid skewing, in which a few items are given very high unit bids and the rest are given

⁷For each graph, we truncate observations at the top and bottom 1% to make the trends easier to see.
⁸As Figure 3(a) shows, the top two bidders diverge on items that overran by more than 100% on average. This is suggestive of moral hazard: the winning bidder profitably bid high on these items, while the losing bidder bid low on them. To account for this possibility, we bound the impact of moral hazard on our results in Appendix B.


“penny bids.” The average number of unit bids worth \$0.10 or less by the winning bidder is 0.51—or 0.7% of the items in an auction. The average number of unit bids worth \$0.50, \$1.00, and \$10.00, respectively, is 1.68, 2.85, and 13.91, corresponding to 2.8%, 4.73%, and 23.29% of the items in an auction. This observation is consistent with previous studies of bidding in scaling auctions. While some studies have cited other forces, such as fear of regulatory rebuke, as alternative explanations for the lack of total bid skewing, others like [Athey and Levin \(2001\)](#) have argued on the basis of interviews with professionals that risk management is a primary concern driving this bidding behavior.

The absence of total bid skewing is not the only testable implication of bidders’ risk aversion. Risk-averse bidders balance the incentive to bid high on items that are projected to overrun with an incentive to bid close to cost on items that are highly uncertain. As our model in Section 5 shows, bidders with higher costs and higher scores face larger amounts of risk from extremal bids. As such, they are less willing to skew strongly or bid far below cost on items predicted to underrun.


This observation is consistent with the pattern demonstrated in Figure 3(a). Here, the second place bidder—who submitted a higher overall score by definition—generally exhibits less severe skewing: a 1% increase in quantity overruns corresponds to only a 0.019% increase in overbids on average. Figure 3(b), which plots a residualized binscatter of the second place bidder’s unit bid for each item against the winning bidder’s unit bid for the same item, shows a similar pattern. While the direction of skewing corresponds strongly between the top two bidders—a higher overbid by the winning bidder corresponds to a higher overbid by the second place bidder as well—the second place bidder’s skewing is more subdued.

The bids in our data also exhibit more direct evidence of risk aversion. We would expect risk-averse bidders to bid lower markups on items that—everything else held fixed—have higher uncertainty. While we do not see observations of the same item in the exact same context with different uncertainty, we present the following suggestive evidence that such behavior is occurring. In Figures 4(a) and 4(b), we plot the relationship between the unit bid for each item in each auction by the winning bidder, and an estimate of the level of uncertainty regarding the ex post quantity of that item (in the context of the particular auction). To calculate the level of uncertainty for each item, we use the results of our first-stage estimation, discussed in Section 6. For every item, in every auction, our first stage gives us an estimate of the variance of the error for the best prediction of what the ex post quantity of that item would be, given information available at the time of bidding. In Figure 4(a), we plot a residualized binscatter of the winning bidder’s absolute percent overbid on each item against the item’s standard deviation—the square root of the estimated prediction variance. This captures the uncertainty of each item’s quantity prediction across auctions in which it may appear with different DOT expectations and project compositions. While the exact numbers may vary across projects with different characteristics and staffing, the types of items that often have the highest or lowest standard deviations are indicative of the range of uncertainties that bidders face when composing their portfolios. Items with the lowest standard deviations—such as “control water structure” and “clearing and grubbing”—tend to be used in lumpy units or fractions of a unit, and are unlikely to depend on unforeseen conditions.⁹ Items with the highest standard deviations—such

⁹While we do see items for which q^e and q^a are typically a single unit, there are many instances in which q^e is a unit, but q^a is a fraction (above or below 1). As we cannot cleanly distinguish which items this might apply to, we model quantities for items with unit q^e in the same way as others. Items with little variation between q^e and q^a are then estimated to have a small variance parameter.



(a) Residualized binscatter of item-level percent absolute overbid against the square root of estimated quantity variance.



(b) Residualized binscatter of item-level percent difference in revenue contribution against the square root of estimated quantity variance.

FIGURE 4.—Measures of item-level overbids against quantity uncertainty.

as various types of fencing, asphalt and excavation—often depend on the scope and the depth of the maintenance site, and so are more difficult to predict accurately before construction is started.

As Figure 4(a) demonstrates, the relationship between item-level absolute overbids and standard deviations is negative. This suggests that holding all else fixed, bidders bid closer to cost on items with higher variance, and thus limit their risk exposure. Note, however, that this analysis does not directly account for the trade-off between quantity overruns and uncertainty. Under our model, a bidder's certainty equivalent increases in the predicted quantity of each item, but decreases in the item's quantity variance. To account for this trade-off, we consider the following alternative metric for bidding high on an item:

$$\% \Delta \text{ Revenue Contribution from } t = \frac{\frac{b_t q_t^a}{\sum_p b_p q_p^a} - \frac{c_t q_t^e}{\sum_p c_t q_p^e}}{\frac{c_t q_t^e}{\sum_p c_t q_p^e}} \times 100.$$

This is the percentage difference in the proportion of the total revenue earned by the winning bidder from item t , and the proportion of the DOT's initial total cost estimate that item t constituted. We take the percent difference between the item's revenue contribution to the bidder and its cost contribution to the DOT's total estimate in order to normalize across items that inherently play a bigger or smaller role in a project's total cost. In Figure 4(b), we plot the residualized binscatter of the %Δ Revenue Contribution due to each item against the item's quantity standard deviation. The negative relationship here is particularly pronounced, providing further evidence that bidders allocate proportionally less weight in their expected revenue to items with high variance. Our model of risk-averse bidding predicts exactly this kind of relationship.

5. A STRUCTURAL MODEL FOR BIDDING WITH RISK AVERSION

In this section, we present the theoretical framework underlying our empirical exercise. We present a parsimonious model of risk-averse bidding in a scaling auction and characterize the optimal spread of unit bids across project components under mild assumptions about the distribution of bidder types and bidder beliefs about their competitors. In Section 8, we augment this characterization with a model of endogenous participation and score selection under an IPV framework motivated by the estimates in Section 7.

The bidding stage of a procurement auction consists of N qualified bidders who compete for a contract to complete a single construction project. A project is characterized by the T items listed in the DOT project specification. Prior to bidding, bidders observe a DOT estimate q_t^e for each item t 's quantity, as well as an additional noisy public signal q_t^b . Although we do not model this explicitly until Section 6, the public signal should be thought of as a refinement of q_t^e that incorporates further public information, such as the identity of the design engineer and historical trends for similar projects. Upon completion of construction, the *actual* quantity q_t^a of each item is realized, independently of which bidder won the auction and at what price. To summarize, there are three kinds of quantity objects:

- $\mathbf{q}^e = (q_1^e, \dots, q_T^e)$: DOT estimates based on underlying conditions at the project site
- $\mathbf{q}^b = (q_1^b, \dots, q_T^b)$: Common refined estimates based on public information
- $\mathbf{q}^a = (q_1^a, \dots, q_T^a)$: Actual quantities, realized ex post independently of the auction.

To compete in the auction, each bidder i must submit a unit bid $b_{t,i}$ for every item t involved in the auction. Bids are simultaneous and sealed until the conclusion of the auction. To determine a winner, each bidder i is given a *score* equal to the sum of her unit bids weighted by the DOT quantity estimates: $s_i = \sum_{t=1}^T b_{t,i} q_t^e$. The bidder with the lowest score wins the contract and executes the project in full. Once the project is complete, the winning bidder is paid her unit bid $b_{t,i}$ multiplied by the actual quantity of item t that was needed, q_t^a .

Bidder Efficiency. The winner of a procurement auction is responsible for securing all of the items required to complete construction. The majority of these items—such as concrete and traffic cones—are standard, competitive goods that have a commonly known market unit cost c_t at the time of the auction. However, bidders differ in their labor, storage, and transportation costs across different projects. To capture this, we assume that bidders differ along a single-dimensional *efficiency multiplier* α . That is, for every item t required for a project, bidder i faces a unit cost of $\alpha_i c_t$, where α_i is the bidder's efficiency type.

Uncertainty and Risk Aversion. Bidders' expectations for how many units of different items will be needed for a project are noisy to different degrees. For tractability, we assume that the bidders' public signal for each item t is normally distributed around the actual quantity of t , with an item-specific variance parameter:

$$q_t^b = q_t^a + \epsilon_t, \quad \text{where } \epsilon_t \sim \mathcal{N}(0, \sigma_t^2). \quad (1)$$

In addition, we assume that bidders are risk averse with a standard CARA utility function over earnings from the project and a private constant coefficient of absolute risk aversion γ_i :

$$u_i(\pi) = 1 - \exp(-\gamma_i \pi). \quad (2)$$

Timing and Information. Prior to the auction, all bidders observe the market cost c_t , DOT estimate q_t^e , public signal q_t^b , and variance σ_t^2 for each item. Bidders' private types (α_i, γ_i) are drawn independently from a publicly known distribution over a compact subset $[\underline{\alpha}, \bar{\alpha}] \times [\underline{\gamma}, \bar{\gamma}]$ of \mathbb{R}_+^2 . Before submitting her bid, each bidder observes her private type as well as the number of competitors that are participating in the auction.

Bidder Payoffs. If a bidder loses the auction, she does not pay or earn anything regardless of her bid. If bidder i wins the auction with bid vector \mathbf{b}_i , she profits the difference between her unit bid and her unit cost for each item, multiplied by the quantity with which the item is ultimately used: $\sum_{t=1}^T q_t^a \cdot (b_{t,i} - \alpha_i c_t)$. As the realization of \mathbf{q}^a is unknown at the time of bidding, bidders face two sources of uncertainty in bidding: uncertainty about their probability of winning and uncertainty about the profits they would earn upon winning. In addition, the winning bidder may anticipate earning an additional lump sum payment ξ —such as an extra work order (EWO)—that does not depend on bids or quantities. Combining these components, bidder i 's expected utility from participating in the auction is given by

$$\underbrace{\left(1 - \mathbb{E}_{\mathbf{q}^a} \left[\exp \left(-\gamma_i \left(\xi + \sum_{t=1}^T q_t^a \cdot (b_{t,i} - \alpha_i c_t) \right) \right) \right] \right)}_{\text{Expected utility conditional on winning with } \mathbf{b}_i} \times \underbrace{\left(\Pr \{ \mathbf{b}_i \cdot \mathbf{q}^e < s_j \text{ for all } j \neq i \} \right)}_{\text{Probability of winning with } \mathbf{b}_i}.$$

This is bidder i 's expected utility from the profit she would earn if she were to win the auction, multiplied by the probability that her score—at the chosen unit bids—will be the lowest one offered, so that she will win. Substituting the bidders' Gaussian signal from Equation (1) and taking the expectation, bidder i 's expected utility is given by

$$\left(1 - \exp \left(-\gamma_i \left(\xi + \sum_{t=1}^T q_t^b (b_{t,i} - \alpha_i c_t) - \frac{\gamma_i \sigma_t^2}{2} (b_{t,i} - \alpha_i c_t)^2 \right) \right) \right) \quad (3)$$

$$\times \left(\Pr \{ \mathbf{b}_i \cdot \mathbf{q}^e < s_j \text{ for all } j \neq i \} \right). \quad (4)$$

Separability of the Bidder's Problem. Notably, bidder i 's expected utility from participating in the auction is *separable* in the following two ways. (i) The probability that bidder i will win (Equation (4)) is entirely determined by the score $s_i = \mathbf{b}_i \cdot \mathbf{q}^e$ and the distribution of competing scores. Thus, any selection of unit bids that sums to the same score yields the same probability of winning. (ii) The expected utility conditional on winning for bidder i (Equation (3)) depends only on the selection of unit bids submitted by i , and is independent of any other bidder's bids.

The separability property implies that a bidder's score is *payoff-sufficient* for her choice of unit bids: in any equilibrium, the vector of unit bids submitted by each bidder must maximize the bidder's expected utility from winning, conditional on the constraint that the unit bids sum to the bidder's equilibrium score.¹⁰ This maximization—which we call

¹⁰This condition is guaranteed to hold in any equilibria where every type of bidder has a nonzero chance of winning the auction. As such, it will hold in every symmetric equilibrium, but does not require symmetry.

the bidder's *portfolio problem*—is a deterministic unilateral optimization problem: it does not depend on the bidder's beliefs about her competition. Instead, all equilibrium considerations are channeled through the bidder's choice of her equilibrium score, which disciplines the portfolio problem through a linear constraint on feasible unit bids.

Characterizing Equilibrium. As bidder types are multidimensional, it may not be possible to specify a unique equilibrium in scores without further assumptions. However, by [Reny \(2011\)](#), there exists a monotone pure strategy equilibrium characterized by the solution to the following two-stage problem when the support of feasible scores is sufficiently high.¹¹ In the first stage, each bidder i chooses a score $s^*(\alpha_i, \gamma_i)$ based on her private type (α_i, γ_i) . This determines the bidder's probability of winning and constrains the second stage of her bidding strategy. In the second stage, bidder i chooses a vector of unit bids \mathbf{b}_i that solves her portfolio problem, subject to the constraint that $\mathbf{b}_i \cdot \mathbf{q}^e = s^*(\alpha_i, \gamma_i)$.

In order for the bids to constitute an equilibrium, \mathbf{b}_i must maximize bidder i 's expected utility conditional on winning subject to the score constraint. This optimization problem is strictly convex, and so it has a unique global maximum for any given score. Furthermore, applying a monotone transformation to Equation (3), this problem reduces to a constrained quadratic program, similar to those studied in standard asset pricing texts:¹²

$$\begin{aligned} \mathbf{b}_i^*(s) = \arg \max_{\mathbf{b}_i} & \left[\sum_{t=1}^T q_t^b (b_{t,i} - \alpha_i c_t) - \frac{\gamma_i \sigma_i^2}{2} (b_{t,i} - \alpha_i c_t)^2 \right] \\ \text{s.t. } & \sum_{t=1}^T b_{t,i} q_t^e = s \quad \text{and} \quad b_{t,i} \geq 0 \quad \text{for all } t. \end{aligned} \quad (5)$$

As unit bids cannot be negative, the portfolio problem in Equation (5) does not have a closed-form solution, and must be solved numerically. However, the optimal unit bid for each item receiving positive weight in the portfolio has the following form:

$$b_{t,i}^*(s) = \alpha_i c_t + \frac{q_t^b / \sigma_t^2}{\gamma_i} + \frac{q_t^e / \sigma_t^2}{\sum_{r: b_{r,i}^*(s) > 0} \left[\frac{(q_r^e)^2}{\sigma_r^2} \right]} \left(s - \sum_{r: b_{r,i}^*(s) > 0} q_r^e \left[\alpha_i c_r + \frac{q_r^b / \sigma_r^2}{\gamma_i} \right] \right). \quad (6)$$

Note that the optimal bid for each item is not only a function of that item's own unit cost and expected quantity-to-variance ratio, but also of the costs, expectations, and variances of the other items receiving positive weight in the optimal portfolio—as well as the bidder's score. As such, variation in the composition of project needs and uncertainty would induce variation in unit bids even if the competitive structure (e.g., the participating bidders and the distribution of their private costs) were the same. This variation is the driver of our identification strategy for estimating bidder cost efficiency and risk-aversion types.

Discussion. It is worth pausing to highlight where our model imposes strong restrictions on bidder responses to uncertainty, and where it does not. A key assumption of our model is that item quantity realizations are (i) fully exogenous and (ii) equally anticipated

¹¹We restrict attention throughout to Bayes–Nash equilibria; see Appendix C.1 for details.

¹²See [Campbell \(2017\)](#) for a survey.

by all bidders. We motivate this assumption in two ways. First, as we noted in Section 3, onsite MassDOT managers—and not contractors—are the primary parties responsible for updating item quantity needs. Second, not only the top two, but all bidders skew their bids in the same direction on average. For instance, Figure E.1(b) in Appendix E shows that the relationship of overbids by bidders of rank 2, 3, and 4 relative to those of the winning bidder all have nearly the same slopes. Still, this assumption shuts down two channels of potential inefficiency from scaling auctions. First, it does not allow for “moral hazard” (Laffont and Tirole (1993)), by which winning bidders overuse items that they had bid high on. Second, it does not allow for a “winner’s curse” (Milgrom and Weber (1982)), by which winning bidders may be adversely selected based on their optimism about quantity realizations. Toward the first concern, we consider an extension of our model with moral hazard in Appendix B. However, we acknowledge the second concern as a limitation of our current work.

On the other hand, our model is flexible with respect to the distribution of bidder types, and additional payments that are not directly bid upon. Our characterization of optimal unit bids does not impose any structure on the correlation between a bidder’s efficiency type α_i and her risk-aversion type γ_i . While different distributions of α and γ may allow for different mappings of bidder types to equilibrium scores, Equation (6) uniquely characterizes the mapping of bidder types to unit bids for each score across all such equilibria. In Section 6, we take advantage of this observation to estimate bidder types in each auction and find that α and γ are strongly positively correlated. Similarly, while many other auction-level considerations—both observed and unobserved—may shape the ways in which equilibrium scores are determined, these considerations do not affect the validity of Equation (6) so long as unit-level costs and quantity expectations are unaffected. As such, our estimation procedure for α and γ is agnostic to bidders’ expectations over EWO payments and entry costs—features that we calibrate under a more stylized model in Section 8.

6. ECONOMETRIC MODEL

We now present a two-step estimation procedure to estimate the primitives of our baseline model. We split our parameters into two categories: (1) statistical/historical parameters, which we estimate in the first stage and (2) economic parameters, which we estimate in the second stage. The first set of parameters characterizes bidders’ beliefs over the distribution of actual quantities. The estimation procedure for this stage employs the full history of auctions in our data to build a statistical model of bidder expectations using publicly available project and item characteristics. However, it does not take into account information on bids or bidders in any auction. The second-stage estimates bidders’ efficiency types α and risk-aversion parameters γ in each auction. For this stage, we take the first-stage estimates as fixed and construct moments for GMM estimation based on idiosyncratic deviations between observed unit bids and optimal unit bids given by Equation (6).

Stage 1: Estimating the Distribution of the Quantity Signals. In the model presented in Section 5, we did not take a stance on what the signals in Equation (1) are based on. The reason for this was to emphasize the flexibility of our model with respect to possible signal structures: the only assumption used is that—conditional on all of the information held at the time of bidding—the bidders’ common belief of the posterior distribution of each q_i^a can be approximated by a Gaussian distribution with a commonly known mean and variance.

For the purpose of estimation, however, we make an additional assumption. Denoting an auction by n , we assume that the posterior distribution of each $q_{t,n}^a$ is given by a statistical model that conditions on $q_{t,n}^e$, item characteristics (e.g., the item's type classification), observable project characteristics (e.g., the project's location, project manager, designer, etc.), and the history of DOT projects.¹³ In particular, we model the realization of the actual quantity of item t in auction n as

$$q_{t,n}^a = q_{t,n}^b + \eta_{t,n}, \quad \text{where } \eta_{t,n} \sim \mathcal{N}(0, \sigma_{t,n}^2) \quad (7)$$

$$\text{such that } q_{t,n}^b = \beta_{0,q} q_{t,n}^e + \beta_q X_{t,n} \quad \text{and} \quad \sigma_{t,n} = \exp(\beta_{0,\sigma} q_{t,n}^e + \beta_\sigma X_{t,n}). \quad (8)$$

Here, $q_{t,n}^b$ is the predicted mean of $q_{t,n}^a$ and $\sigma_{t,n}$ is the square root of its predicted variance—linear and log-linear functions of the DOT estimate for item t 's quantity $q_{t,n}^e$ and a vector of item-project characteristics $X_{t,n}$. We estimate this model with Hamiltonian Monte Carlo and use the posterior mean of the post-warmup draws, denoted by $\hat{\theta}_1$, as a point estimate for the second stage of estimation. We summarize the estimates of $\hat{\theta}_1$ and demonstrate the goodness-of-fit in Appendix D.

Note that our model allows for correlations between item means $q_{t,n}^b$ and variances $\sigma_{t,n}^2$ through observables, but assumes that deviations $\eta_{t,n}$ from the means are independent across items within an auction. This is not a binding constraint from a theoretical perspective: in principle, our approach could accommodate correlations across $\eta_{t,n}$ as well. However, in contrast to the asset pricing literature, each observation of a “portfolio” in our data is composed of a different basket of items. As such, it would be difficult to identify and estimate consistent correlations across items in our setting.

Our model of bidder quantity signals can be thought of in several ways. It can be interpreted as an additional component of the structural model: the bidders use our method as a statistical estimation procedure to assess the likelihood of item quantities prior to bidding. The DOT quantities, item, and project characteristics are indeed all publicly known at the time of bidding, as are historical records of DOT projections and ex post quantities. Furthermore, there is a mature industry of software for procurement bid management that touts sophisticated estimation of project input quantities and costs. It is thus likely that firms use similar off-the-shelf tools to forecast project needs. Alternatively, this assumption could be thought of as the econometrician's model of each signal mean q_t^b and variance σ_t^2 .

Stage 2: Estimating Efficiency Types and Risk Aversion. Our data set contains a unit bid for every item submitted by every participating bidder in each auction in our sample. Every auction carries a different set of project characteristics, a different composition of items to be bid, different quantity expectations and variances, and a different set of participating bidders who may have different concurrent advantages in efficiency and risk aversion. These features collectively determine the optimal unit bid for every item-bidder-auction (t, i, n) observation in our sample. In order to estimate the bidder types underlying each portfolio of unit bids, we make two main assumptions. First, we assume that the optimal unit bid for item t by bidder i in auction n is determined by the formula in Equation (6), given the predicted quantity means $\hat{q}_{t,n}^b$ and variances $\hat{\sigma}_{t,n}^2$ from our first stage, and a bidder-auction efficiency $\alpha_{i,n}$ and risk-aversion $\gamma_{i,n}$ parameter. Second, we assume that the optimal bid for each (t, i, n) tuple, evaluated at the equilibrium score $s_{i,n}^*$,

¹³We index auctions and projects interchangeably with n .

is observed by the econometrician with an idiosyncratic mean-zero measurement error. We summarize these assumptions as follows.

ASSUMPTION 1: Let $b_{t,i,n}^d$ denote the unit bid for item t submitted by bidder i in auction n , as observed in our data. Each observed unit bid is equal to the optimal bid $b_{t,i,n}^*(s_{i,n}^*|\hat{\theta}_1, \alpha_{i,n}, \gamma_{i,n})$, subject to an IID mean-zero measurement error $\nu_{t,i,n}$:

$$b_{t,i,n}^d = b_{t,i,n}^*(s_{i,n}^*|\hat{\theta}_1, \alpha_{i,n}, \gamma_{i,n}) + \nu_{t,i,n} \quad \text{where } \mathbb{E}[\nu_{t,i,n}] = 0 \text{ and } \nu_{t,i,n} \perp X_{t,n}, X_{i,n}, X_n.$$

Assumption 1 states that the optimal unit bid for each (t, i, n) is observed in our data with an idiosyncratic error that is independent across draws, and orthogonal to auction-item, auction-bidder, and auction-wide characteristics. We interpret these errors as measurement or rounding errors: the results of rounding or smudging in the translation between the bidder's optimal bidding choice and the record available to the DOT.

Note that while we do not directly observe the equilibrium score $s_{i,n}^*$ for each (i, n) pair, our observations of unit bids provide a noisy signal of it: $s_{i,n}^d = \sum_t b_{t,i,n}^d q_{t,i,n}^e = s_{i,n}^* + \bar{\nu}_{i,n}$, where $\mathbb{E}[\bar{\nu}_{i,n} \cdot X_{i,n}] = 0$ by Assumption 1. As we discussed in Section 5, $s_{i,n}^*$ is a sufficient statistic for bidder i 's competitive considerations in auction n . As such, our formula for $b_{t,i,n}^*$ accounts for bidders' beliefs about their opponents through $s_{i,n}^*$ and explains the residual systematic variation in unit bids through the weights that bidders' efficiency and risk-aversion parameters place on items with different balances of expected quantities and uncertainty.

The intuition for identification is as follows: given the estimates of the first-stage parameters $\hat{\theta}_1$, the formula for $b_{t,i,n}^*$ can be written as a linear projection of $\alpha_{i,n}$, $\frac{1}{\gamma_{i,n}}$ and $s_{i,n}^*$. Under Assumption 1, the vector of observed bids $b_{t,i,n}^d$ can therefore be expressed as a system of regression equations with $\alpha_{i,n}$ and $\gamma_{i,n}$ as the coefficients of observed weights that capture the relative value, in terms of cost and uncertainty respectively, of bidding higher on each item (t, i, n) within bidder i 's portfolio in auction n , and an orthogonal residual term. The vector of $\alpha_{i,n}$ and $\gamma_{i,n}$ is thus identified by the orthogonality of bid residuals with respect to the (t, i, n) characteristics defining each item's weights, as in a standard OLS setting.

Because our characterization of optimal bids does not require assumptions on the distribution of bidder types, we allow both $\alpha_{i,n}$ and $\gamma_{i,n}$ to vary flexibly at the bidder-auction level. Variation in $\alpha_{i,n}$ reflects differences in bidder size, capacity at the time of bidding, and specialization for the particular project at hand. For instance, if a bidder's headquarters is closer to the project site, her transportation costs for all of the items involved will be lower than for an equivalent project further away. Variation in $\gamma_{i,n}$ reflects a CARA approximation of bidders' risk aversion with respect to the stakes involved in each auction. This can differ across bidders with different financial situations and may covary with time and with bidders' capacity. To capture these relationships parsimoniously in our limited sample, we project $\alpha_{i,n}$ and $\gamma_{i,n}$ onto a bidder fixed-effect and a vector of bidder-auction characteristics $X_{i,n}$. This specification imposes a correlation between the efficiency and risk aversion of each bidder-auction through realizations of the characteristics $X_{i,n}$. However, the correlation may, in general, take on different signs and vary across types of bidders and auctions.

To estimate our second-stage parameters efficiently, we apply a GMM procedure leveraging the orthogonality of the bid-level residuals implied by Assumption 1. Each moment condition corresponds to the weighted expectation of residuals across a fixed slice

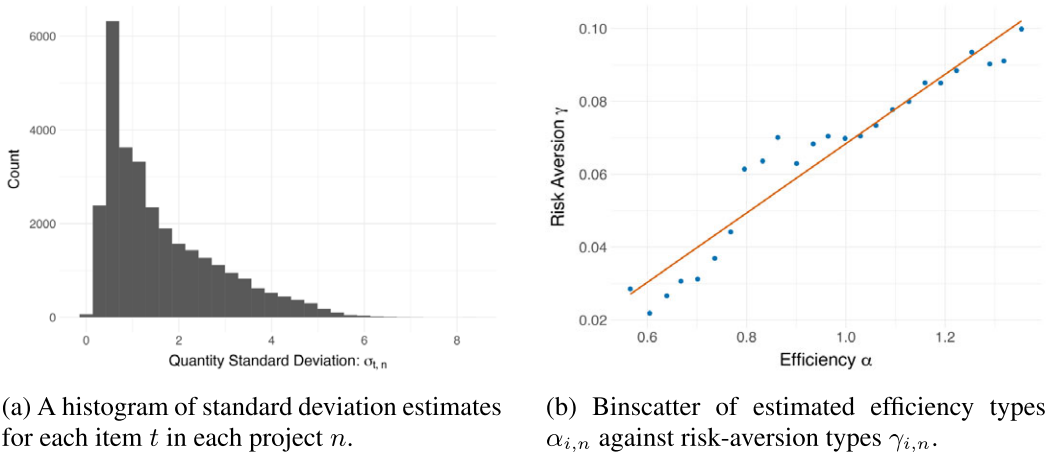


FIGURE 5.—Descriptives of parameter estimates from the first and second stage.

of bidder-auction pairs in a sample of auction draws that asymptotically approaches infinity. We describe identification and estimation for the first- and second-stage parameters in detail in Appendix C.3.1 and Appendix C.3.2, respectively. In Appendix C.3.3, we discuss robustness to unobserved auction heterogeneity.

7. ESTIMATION RESULTS

Our structural estimation procedure consists of two parts. In the first stage, we estimate the distribution of the ex post quantity of each item conditional on its item-auction characteristics using Hamiltonian Monte Carlo sampling. We present parameter estimates for the regression coefficients underlying the predicted quantity $\hat{q}_{t,n}^b$ and variance $\hat{\sigma}_{t,n}^2$ terms in Table D.I in Appendix D. In addition, we demonstrate the model fit for the first stage in Figure D.1. A histogram of the standard deviation estimates $\hat{\sigma}_{t,n}$ themselves is plotted in Figure 5(a). Prior to estimation, all item quantities were scaled so as to be of comparable value between 0 and 10. As demonstrated in the histogram, the majority of standard deviation terms are between 0 and 3, with a trailing number of higher values.

In the second stage, we estimate an efficiency type $\alpha_{i,n}$ and CARA coefficient $\gamma_{i,n}$ for every bidder-auction pair in our data using the GMM estimator presented in Section 6. We summarize the results in Tables III and IV. Bootstrapped standard errors and confidence intervals are presented in Table D.II in Appendix D.

In Table III, we present summary statistics of our estimates of bidder-auction risk aversion and efficiency types. The median coefficient of risk-aversion $\hat{\gamma}_{i,n}$ is estimated to be 0.061 when dollar values are scaled by \$1000. An individual with this level of risk aversion would require a certain payment of \$30 to accept a 50–50 lottery to either win or lose \$1000 with indifference, and \$2878 to accept a 50–50 lottery to win or lose \$10,000.¹⁴ The median efficiency type $\hat{\alpha}_{i,n}$ is estimated to be 1.053. An efficiency of 1 would suggest that the bidder faces costs exactly at the rates represented by MassDOT’s estimates. Our

¹⁴It is not unusual for managers to exhibit high levels of risk aversion at high stakes. For instance, a McKinsey survey of 1500 managers found that most required a certainty equivalent of \$328M to accept a risk of losing an investment worth \$100M (LovoalKollerUhlenerKahneman (2020)).

TABLE III
SUMMARY STATISTICS OF $\gamma_{i,n}$ AND $\alpha_{i,n}$ ESTIMATES BY PROJECT TYPE.

Project Type	Mean	SD	25%	Median	75%
$\hat{\gamma}_{i,n}$					
All	0.088	0.083	0.042	0.061	0.096
Bridge Reconstruction/Rehab	0.088	0.082	0.046	0.062	0.098
Bridge Replacement	0.080	0.082	0.041	0.056	0.079
Structures Maintenance	0.100	0.084	0.043	0.075	0.129
$\hat{\alpha}_{i,n}$					
All	1.033	0.22	0.953	1.053	1.175
Bridge Reconstruction/Rehab	1.085	0.231	0.978	1.095	1.275
Bridge Replacement	1.044	0.214	0.949	1.058	1.206
Structures Maintenance	0.985	0.213	0.941	1.030	1.110

results thus suggest that the material costs for the median bidder-auction are 5% higher than the DOT’s estimates.

However, there is substantial heterogeneity across bidders and projects. While the median risk-aversion parameter among bridge replacement projects would require a certain payment of \$2665 to accept a lottery for \$10,000 (or \$28 to accept a 50–50 lottery for \$1000), the median among structures maintenance projects would require a certain payment of \$3444 for the same lottery (or \$37 for a 50–50 lottery for \$1000). Looking across project types, the 25th percentile (75th percentile) of bidder-auction risk-aversion parameters across project types would require a certain payment of \$20 (\$48) to accept a 50–50 lottery for \$1000 or \$2041 (\$4205) to accept a lottery for \$10,000. There is also substantial heterogeneity in cost efficiency. For instance, the 25th-percentile bidder across project types has an auction-specific cost multiplier of 0.953, suggesting that she obtains costs that are 5% lower than the DOT cost estimates, while the 75th-percentile bidder has costs about 17.5% higher than the DOT cost estimates.

Our estimates allow us to examine the relationship between bidders’ cost efficiency and risk aversion. Figure 5(b) plots a binscatter of estimated efficiency types $\hat{\alpha}_{i,n}$ against their corresponding CARA parameters $\hat{\gamma}_{i,n}$. To smoothly control for heterogeneity across auctions, we first take the log of $\hat{\gamma}_{i,n}$ and demean each estimate by subtracting the auction-level average $\hat{\alpha}_{i,n}$ and $\log(\hat{\gamma}_{i,n})$, respectively. To make the binscatter easier to read, we then add back the cross-auction average $\hat{\alpha}_{i,n}$ and $\log(\hat{\gamma}_{i,n})$ to each value, and exponentiate $\log(\hat{\gamma}_{i,n})$.

As the figure shows, $\hat{\alpha}_{i,n}$ and $\hat{\gamma}_{i,n}$ are strongly positively correlated. The corresponding log-linear regression suggests that a 10-percentage point increase in $\hat{\alpha}_{i,n}$ corresponds to a 20% increase in $\hat{\gamma}_{i,n}$ after accounting for auction-level fixed effects. That is, a bidder who is 10% less efficient is also 20% more risk averse. Moreover, $\hat{\gamma}_{i,n}$ is well predicted by a monotonic function of $\hat{\alpha}_{i,n}$. The log-linear fixed effects regression above explains 80% of the variation in $\hat{\gamma}_{i,n}$ across bidder-auction pairs. In Section 8, we build on this relationship to project bidder efficiency and risk aversion onto a single-dimensional metatype for counterfactuals.

In Table IV, we summarize the distribution of ex post markups implied by our estimates of $\hat{\alpha}_{i,n}$. The markup for bidder i in auction n is given by the total ex post profit that the bidder would obtain from completing the project given her bids, normalized by her total

TABLE IV
SUMMARY STATISTICS OF MARKUP ESTIMATES BY NUMBER OF PARTICIPATING BIDDERS.

Num Bidders in Auction	Estimated Bidder Markups				
	Mean	SD	25%	Median	75%
All	21%	49%	−9%	10%	36%
2–3 Bidders	41%	76%	0%	18%	54%
4–6 Bidders	26%	58%	−6%	13%	40%
7+ Bidders	16%	39%	−10%	7%	31%

cost:

$$\text{Markup} = \frac{\sum_t q_{t,n}^a \cdot (b_{t,i,n} - \alpha_{i,n} c_{t,n})}{\sum_t q_{t,n}^a \cdot (\alpha_{i,n} c_{t,n})}.$$

The median estimated markup in our sample is 10% and the mean is 21%. Rather than summarize estimated markups by project type, we split projects by the number of participating bidders in each auction. Although there is substantial heterogeneity within each group, markups are generally decreasing with the amount of competition. Note that this markup measure does not account for extra work orders, as our baseline model does not identify the cost of fulfilling them. While this may help explain why some of the markup estimates on the lower tail are negative, it does not imply that our parameter estimates are biased. Any additional payments that may have been anticipated at the time of bidding would have influenced the bidders' choice of equilibrium score. However, because these payments were not bid upon, their presence would not change the solution to the bid portfolio optimization problem conditioned on the equilibrium score that is observed in our data.

We defer a demonstration of the goodness-of-fit of our structural model to Appendix D. Figure D.2(a) presents a scatterplot and Figure D.3(a) presents a quantile-quantile plot, both of the unit bids predicted by our model against the unit bids observed in our data. While the bid predictions are not perfect, the correspondence between predictions and data is quite good. Figure D.2(b) presents a regression analysis of the predictiveness of our model fit on the observed data. Our model fit predicts data bids with an R-squared of 0.881.

8. COUNTERFACTUALS

In Section 4, we argued that the bids observed in MassDOT bridge procurement auctions are suggestive of risk aversion. In Sections 5 and 6, we showed that under risk aversion, observations of unit bids alone could be used to identify parameters for bidder efficiency and risk aversion, independently of strong assumptions about the competitive environment. In this section, we quantify the levels of risk and risk aversion exhibited in these auctions and evaluate the effectiveness of several counterfactual policies to lower DOT costs.

In order to evaluate counterfactual equilibrium outcomes, we require further assumptions about the strategic environment facing bidders. In estimating bidder-type parameters, we interpreted the scores submitted by bidders in our data as equilibrium outcomes.

We were able to remain agnostic about how the scores came about because portfolio optimization, subject to a given score, is unaffected by other bidders' behavior. However, if the auction format were to change, the scores—and perhaps even the bidders who participate—would likely change in equilibrium.

In order to account for such changes, we consider a stylized model of bidder beliefs and participation decisions. Leveraging our observation that estimated risk-aversion types are well approximated by a monotonically increasing function of efficiency types within an auction, we model bidder-auction types $(\alpha_{i,n}, \gamma_{i,n})$ as deterministic increasing transformations of a private single-dimensional metatype $\tau_{i,n}$ that is drawn independently from a commonly known distribution for each potential bidder in each auction. To allow for endogenous changes in the set of participating bidders, we adopt a model of selected entry in the spirit of Samuelson (1985). Using this model, we calibrate the entry costs and extra work order (EWO) expectations that best rationalize the observed rates of entry in our data. We then compute equilibrium entry and bidding strategies under each counterfactual policy, and compare the resulting expected DOT expenditures to the status quo policy.

Timing, Beliefs, and Endogenous Entry. Each auction has a limited set of potential bidders who may consider submitting a bid. For simplicity, we assume that bidders who bid on an auction within the same project type, geographic region, and year could have participated in any other such auction.¹⁵ We assume that all potential bidders consider one auction at a time and know the number and type distribution of their potential competitors.

Bidder types vary across auctions through a combination of bidder-auction characteristics. For simplicity, we project these characteristics onto a single-dimensional metatype $\tau_{i,n}$, such that for each auction n , each potential bidder i 's efficiency type $\alpha_n(\tau_{i,n})$ and risk-aversion type $\gamma_n(\tau_{i,n})$ are fully determined by the realization of $\tau_{i,n}$. We assume that metatypes $\tau_{i,n}$ are drawn IID from a common auction-specific distribution with CDF F_n that is calibrated by fitting to the estimates from Section 7, as described in Appendix A.

The timing of bidding is as follows. Once an auction n is advertised, all potential bidders observe their types $(\alpha_n(\tau_{i,n}), \gamma_n(\tau_{i,n}))$ and evaluate whether they would like to participate. Preparing a bid in order to participate is costly. We assume that all bidders who choose to enter an auction n incur a common entry cost κ_n , regardless of whether they win. Once participation decisions are realized, the number of bidders who paid the entry cost is publicized and each participating bidder submits a vector of unit bids and a corresponding total score. The bidder with the lowest total score is awarded the contract and fulfills it in full, incurring all fulfillment costs. Upon completion of the contract, the winning bidder is reimbursed per unit of each item that is ultimately used according to her unit bid for that item, along with a fixed payment for any EWOs that were needed.

Equilibrium Bidding Strategies. Following the literature (Athey, Levin, and Seira (2011), Roberts and Sweeting (2013)), we construct a type-symmetric equilibrium in monotone pure strategies in three stages. In the first stage, all types $\tau_{i,n}$ below a threshold τ_n^* enter the auction. For this to be an equilibrium, every type $\tau_{i,n} < \tau_n^*$ must expect a

¹⁵Our definition of potential bidders is similar to that of Kong (2020). The maximum number of potential bidders is 20, while five auctions were grouped alone and excluded; see Appendix C.4 for details.

positive value $V_n(\tau_{i,n})$ from participating in the auction:

$$V_n(\tau_{i,n}) = \sum_{m=1}^M \binom{M-1}{m-1} F_n(\tau_n^*)^{m-1} (1 - F_n(\tau_n^*))^{M-m} \text{EU}_n(s_n^*(\tau_{i,n}) | \lambda_n, \kappa_n, m), \quad (9)$$

where $F_n(\tau_n^*)$ is the probability that an independent draw of a competing potential bidder's type is below τ_n^* and $\text{EU}_n(s_n^*(\tau_{i,n}) | \lambda_n, \kappa_n, m)$ is the expected utility that the bidder expects to earn under her equilibrium bidding strategy $s_n^*(\tau_{i,n})$ if $m - 1$ competing bidders participate. The marginal type τ_n^* expects to earn no profit, and so equilibrium entry probabilities are defined by the equation $V_n(\tau_n^*) = 0$.

In the second stage, the bidder observes the number of competing bidders and chooses her equilibrium strategy $s_n^*(\tau_{i,n})$ so as to maximize her expected utility from participating in the auction. Building on the discussion in Section 5, the expected utility $\text{EU}_n(s | \tau_{i,n})$ for a given score s submitted by a bidder with type $\tau_{i,n}$ is given by

$$\begin{aligned} & [1 - \exp(-\gamma_n(\tau_{i,n}) \cdot [\text{CE}_n(\mathbf{b}_{i,n}^*(s) | \tau_{i,n}) + \lambda_n \cdot \text{EWO}_n - \kappa_n])] \times \prod_{j \neq i} [1 - H_{j,n}(s)] \\ & + [1 - \exp(\gamma_n(\tau_{i,n}) \cdot \kappa_n)] \times \left(1 - \prod_{j \neq i} [1 - H_{j,n}(s)] \right), \end{aligned}$$

where we suppress the arguments λ_n, κ_n, m in the $\text{EU}_n(\cdot)$ operator because they are held fixed when computing the equilibrium score mapping in each auction instance. Here, $\text{CE}_n(\mathbf{b}_{i,n}^*(s) | \tau_{i,n})$ is the certainty equivalent of profits from the portfolio of items involved in auction n as defined in Equation (3), and the vector of optimal unit bids $\mathbf{b}_{i,n}^*(s)$ is given by the solution to the portfolio problem in Equation (5). In addition to optimal portfolio profits, the bidder anticipates earning a profit from EWOs. We assume that realizations of EWO payments are exogenous to the outcome of the auction and that expectations of extra profits are common to all bidders. As the only information available about EWOs in our data is the final sum paid out to the winning bidder, EWO_n , we model bidders' certainty equivalent from EWOs in reduced form by $\lambda_n \cdot \text{EWO}_n$, where λ_n is an exogenous scalar that is common and known to all bidders in the auction.

The term $\prod_{j \neq i} [1 - H_{j,n}(s)]$ corresponds to the probability that s is the lowest score submitted under the distribution of opponent scores. Note that the entry cost κ_n is paid independently of the outcome of the auction. Given our model of participation, each bidder believes that her opponents' types $\tau_{j,n}$ are distributed IID according to a truncated CDF F_n^* , such that $F_n^*(\tau_n^*) = 1$. As such, under the unique monotonic equilibrium, the probability of winning with score \tilde{s} given $m - 1$ opponents is equal to the probability that the unique bidder-type $\tilde{\tau}$ who submits \tilde{s} under the equilibrium scoring strategy $s_n^*(\cdot)$ is also the lowest (e.g., most competitive) type participating in the auction: $(1 - F_n^*(s_n^{*-1}(\tilde{s})))^{m-1}$.

The function that maps bidder-types τ to equilibrium scores $s_n^*(\tau)$ in each auction is characterized by the first-order condition establishing the optimality of $\text{EU}_n(s_n^*(\tau) | \tau_{i,n})$ with respect to s . To evaluate the outcome of each auction upon entry, we solve the resulting differential equation in each auction n with respect to each possible number of potential bidders m ; solve the portfolio problem for every possible $(\tau, s_n^*(\tau))$ pair to obtain the equilibrium bid vector $\mathbf{b}_\tau^*(s_n^*(\tau))$; and integrate the resulting ex post DOT cost function $(\mathbf{b}_\tau(s_n^*(\tau)) \cdot \mathbf{q}^a)$ with respect to the first-order statistic of the τ distribution. To evaluate bidder welfare, we integrate over the bidder certainty equivalent function $\text{CE}_n(\mathbf{b}_\tau^*(s_n^*(\tau)) | \tau)$

TABLE V
SUMMARY STATISTICS OF CALIBRATED ESTIMATES OF EWO COEFFICIENTS AND ENTRY COSTS.

Parameter	Mean	SD	25%	50%	75%
κ_n	\$332	\$834	\$0.21	\$23	\$127
λ_n	0.44	0.42	0	0.35	0.9

rather than the DOT cost function. See Appendix A for more details on equilibrium construction, and Figure D.3(b) for a comparison of the equilibrium scores predicted according to this model against those observed in the data. A full tabulation of counterfactual results is presented in Appendix F.

Calibrating Expected EWO Profits and Entry Costs. Although EWOs and entry costs do not affect the optimal spread of unit bids once a score has been chosen, they may impact the equilibrium mapping of bidder types to scores. To account for this in our counterfactual simulations, we first calibrate the entry cost κ_n and EWO coefficient λ_n in each auction using moments from our entry model. We describe our calibration procedure in detail in Appendix C.4 and present summary statistics for the calibrated parameters in Table V. For most auctions, we are only able to rationalize a very small entry fee. This is because the certainty equivalent of profits for the threshold-type τ_n^* is typically very low, given the large number of potential bidders and constraints on the magnitude of bids that are allowable in the case that a single bidder participates.¹⁶ In addition, we find that EWO coefficients are often quite high: the median λ_n suggests that bidders anticipate a certainty equivalent of 35% of the EWO amount that was paid out.

Scaling Auctions as Insurance. While scaling auctions are widely used in many parts of public procurement, they are not ubiquitous. Even within MassDOT, there is heterogeneity: in 2007, the division responsible for public transportation switched from scaling auctions to a *lump sum* format in which contractors submit a single total bid for completing the project under auction. Lump sum auctions have some attractive properties. They may require less detailed specification plans from DOT engineers and they pass the incentive to minimize costs onto the contractor, thereby reducing the scope for moral hazard.

However, in the context of bridge maintenance projects—where DOT officials are able to monitor work effectively enough to eliminate moral hazard concerns—scaling auctions provide an important mechanism for containing DOT costs. Lump sum auctions require bidders to pre-commit to a payment at the time of bidding, leaving them liable for all unforeseen changes. By contrast, scaling auctions compensate bidders for whatever item quantities are actually used, and they allow bidders to hedge their bets through portfolio optimization. Equation (10) compares the certainty equivalent that a bidder submitting the same bid vector \mathbf{b} would expect under a scaling auction and under a lump sum auction.

CE($\mathbf{b}|\alpha, \gamma$): Scaling Auction

$$\sum_t \underbrace{(b_t q_t^b - \alpha c_t q_t^b)}_{\text{Expected Profits}} - \underbrace{\frac{\gamma \sigma_t^2}{2} (b_t - \alpha c_t)^2}_{\text{Risk Term}}$$

CE($\mathbf{b}|\alpha, \gamma$): Lump Sum Auction

$$\sum_t \underbrace{(b_t q_t^e - \alpha c_t q_t^b)}_{\text{Expected Profits}} - \underbrace{\frac{\gamma \sigma_t^2}{2} (\alpha c_t)^2}_{\text{Risk Term}}$$

(10)

¹⁶Based on the guidelines in the MassDOT Standard Operating Procedures, we assume that single-bidder bids are allowable if they amount to no more than a 25% markup over the DOT-estimated score.

Whereas lump sum auctions force bidders to internalize the full cost of uncertainty for each item, scaling auctions allow bidders to temper their risk exposure by bidding close to their cost on items with high variance. In this sense, scaling auctions provide insurance against project uncertainty that is unavailable in lump sum auctions. They allow bidders to sacrifice higher expected profits—for instance, through higher mark-ups on items expected to overrun—in exchange for lower risk. In settings with high levels of uncertainty such as ours, this may make it possible for bidders to obtain the same certainty equivalent with a lower score than under the lump sum format, incentivizing each firm to bid more aggressively.

Our simulations show that the amount of insurance provisioned by MassDOT bridge auctions is substantial. Moving to a lump sum format would increase DOT payments to the winning bidder by over 42% in the median auction in our data set. Given the scope of the projects in our data, this amounts to additional spending of over \$300,000 per auction.

Moreover, this difference in spending compounds the effects of two competing equilibrium forces. On the one hand, the added risk exposure generated by lump sum auctions acts to increase DOT spending in two ways. First, as we noted above, bidders require higher guaranteed payments in order to be willing to participate. This is seen most clearly through the threshold bidder-type τ^* , who must bid higher in order to break even. Subsequently, more competitive bidders may need to bid higher as well, in order to satisfy incentive compatibility. Second, the threshold-type τ^* itself may need to decrease in order to make breaking even feasible. In this case, the ex ante probability of participation decreases, as does the expected number of competing bidders.

On the other hand, if moving to a lump sum format causes the threshold-type τ^* to decrease, then the bidders who do participate are more competitive on average. As efficiency and risk aversion are positively correlated in our sample, this selection effect is amplified: selected bidders are both more efficient and less risk averse. As such, conditional on the number of bidders, the competitive pressure intensifies and bidders are pushed to submit lower bids—compensating, in part, for the decrease in overall bidder participation.

In our simulations, switching to a lump sum format reduces the median threshold type by 20% in cost efficiency and 28% in risk aversion. If participation levels were held fixed so that τ^* remained at the baseline level, switching to a lump sum format would increase DOT spending by 96% for the median auction. Selected participation therefore compensates for over half of the added DOT spending from increased risk exposure in lump sum auctions.

Lump Sums With Renegotiation. The lump sum simulations above assume that—as is the case in scaling auctions—there is no ex post renegotiation. As such, bidders are liable for the entirety of unforeseen project costs no matter how large they become. In practice, however, bidders may be able to recoup some of their costs by negotiating for additional payments based on ex post quantity realizations. To account for this possibility, we consider a model in which bidders expect to be able to recoup a percentage μ of earnings lost due to unforeseen project changes—for instance, through a Nash bargaining negotiation. In order to credibly convey their lost value from the project, the bidders are forced to show that their claimed ex post project total could be generated by unit prices that are consistent with their initial lump-bid. As such, each bidder expects to be paid her score (as in the basic lump sum case), plus μ of the ex post quantity differential of each item multiplied by the item's unit price. Bidders internalize the additional ex post payments at the time of bidding and so they choose unit prices to maximize their certainty

equivalent:

$$\overbrace{\sum_t b_i q_i^e + (\mu b_i (q_i^b - q_i^e)) - \alpha c_i q_i^b}^{\text{Expected Profits}} - \underbrace{\frac{\gamma \sigma_i^2}{2} (\mu b_i - \alpha c_i)^2}_{\text{Risk Term}}. \quad (11)$$

Comparing Equation (11) to Equation (10), it is clear that renegotiation makes it possible for bidders to reduce the amount of risk that they are exposed to using their unit bids. This is not only because bidders are partially reimbursed for every item, but also because precommitting to unit prices allows the bidders to optimize their balance of risks as they would in a scaling auction. In our simulations, we find that moving from a scaling auction format to a lump sum format with 2:1 renegotiation ($\mu = 0.33$) would only increase DOT costs by 14% or \$124,195 for the median auction. If the bidder is able to bargain with equal weight ($\mu = 0.5$), the increase to median DOT costs is only 8.5% or \$92,431.

The substantial reduction in DOT costs incurred from adding a renegotiation stage to the lump sum format suggests that providing even a small amount of insurance to bidders may allow them to significantly cut their risk exposure. This dynamic bears out in the realization of counterfactual threshold types as well. We find that the median threshold type decreases by only 0.09% in cost efficiency and 0.13% in risk aversion when moving from a scaling auction to a lump sum auction with 2:1 renegotiation. Under 50–50 renegotiation, the threshold type is approximately unchanged altogether. Thus, renegotiation obviates the majority of the participation and selection effects induced by the lump sum format, and DOT costs are nearly the same whether or not participation adjusts.

Moral Hazard. As renegotiation eliminates much of the added cost of moving to a lump sum format, a natural question is whether lump sum auctions with renegotiation might be preferred to scaling auctions under some circumstances. In Appendix B, we relax the assumption that bidders cannot affect the ex post realization of item quantities. In this case, switching from a scaling format to a lump sum format would change the quantities that are realized in equilibrium, as the winning bidder would no longer have an incentive to overuse items with high bids (or, in the case of lump sum with renegotiation, have a smaller incentive). While we cannot identify the extent to which item quantities are manipulable within our framework, we consider a bounding exercise in which we assume that *all* profitable overruns observed in our data are manipulations. Treating this as an upper bound for the cost savings in lump sum auctions under moral hazard, we find that the median cost of switching to a lump sum format decreases by at most 30% without renegotiation and does not decrease at all with renegotiation. This suggests that our qualitative conclusions would likely hold under a level of moral hazard that is consistent with our data.

The Cost of Uncertainty Under Scaling Auctions. Our results above suggest not only that scaling auctions provide substantive insurance for the bidders in our data, but also that the amount of project uncertainty that bidders face is large. As such, a direct method to reduce DOT costs may be to simply lower ex ante uncertainty—for instance, by improving inspection directives and engineer training. To test the potential for a policy of this sort to be effective, we consider an extreme counterfactual in which the DOT is able to perfectly predict exactly what quantity of each item will be required to fulfill each project. The ex post correct quantities are then posted publicly from the beginning, so that $q^e = q^a$. All

bidders know that these quantities are correct and so they do not anticipate any further uncertainty: $\mathbf{q}^b = \mathbf{q}^a$ and $\sigma^2 = \mathbf{0}$.

Disclosing ex post quantities to bidders at the start of an auction has two effects. First, it trivializes the portfolio problem: with no project uncertainty, there is no need to hedge or skew. Second, it grants bidders access to the exact ex post quantities—rather than sophisticated estimates thereof—allowing the bidders to perfectly maximize their earnings from an ex post perspective at the DOT's expense. As we are primarily interested in quantifying the cost of uncertainty itself, we first compare the no-risk counterfactual against a baseline in which the bidders' ex ante predictions align with the ex post quantities ($\mathbf{q}^b = \mathbf{q}^a$), but the level of uncertainty σ^2 is unchanged. In this case, eliminating risk reduces the baseline variances σ^2 to zero, but does not affect bidders' quantity expectations. We find substantial reductions: DOT spending would decrease by about 14.5% or \$145,920 in the median auction. This suggests that—holding all else fixed—the level of uncertainty about item quantities plays a substantial role in determining DOT costs.

However, a policy of simply reducing risk may not hold up to practical considerations. When we compare the no-risk counterfactual to a baseline with the estimated predictions $\hat{\mathbf{q}}^b$ from Section 7, the DOT cost for the median auction *increases* by 1.9% or \$18,782. This suggests that the cost of allowing bidders to optimize their bids with respect to the realized quantities—as opposed to noisy predictions that may induce errors that are beneficial from the DOT's ex post perspective—may balance out the savings from eliminating risk. As such, we conclude that policies to curtail risk directly would be unlikely to improve upon the status quo, given the insurance already conferred to bidders by the scaling format.

9. CONCLUSION

This paper studies the bidding behavior of construction firms that participate in scaling procurement auctions run by the Massachusetts Department of Transportation. We develop a model of equilibrium bidding by risk-averse bidders that are collectively better informed than the auctioneer. As noted previously in the literature, informed bidders are incentivized to strategically *skew* their bids, placing high bids on items they predict will overrun the DOT's quantity estimates and low bids on items they predict will underrun. Risk-averse bidders go further—by balancing their bid portfolios across items with different levels of uncertainty, they limit their exposure to the risk of unexpected changes in the quantities ultimately needed to complete a project.

We present evidence that bidding in our setting is consistent with these predictions: holding all else fixed, items that overrun MassDOT's predictions have higher bids on average, while items that bear higher uncertainty have bids that are closer to their unit costs. Furthermore, we argue that accounting for risk aversion has significant implications for policy design. If the bidders were risk neutral, common policies such as switching to a lump sum format or investing in engineer training to reduce uncertainty would not change MassDOT spending in equilibrium. If the bidders are risk averse, however, these policies have theoretically ambiguous, potentially large consequences.

To assess the cost due to risk in our context and evaluate the effectiveness of these different policies empirically, we estimate the parameters underlying our model. We then simulate equilibrium entry and bidding decisions at the item-bidder-auction level for every type of bidder in each of our auctions under the aforementioned counterfactual policies. We estimate that the level of uncertainty in our setting is substantial—it entails a premium of up to 14.5% on payments to the winning bidder in the median auction. However,

an effort to reduce uncertainty may not reduce costs very much in practice, as this would also improve bidders' ability to maximize their ex post revenue. Moreover, switching to a lump sum auction would be very costly—42% more for the median project—because it would expose bidders to full liability for unexpected changes in the project specification.

Viewed in this light, scaling auctions provide MassDOT with an effective mechanism to insure bidders against inevitable shocks due to underlying conditions that are unearthed at the time of construction. While our framework enables evaluating farther-reaching policies as well, such as an adaptation of the multistage optimal mechanism established by Maskin and Riley (1984) and Matthews (1987), we leave this for future work.

APPENDIX A: SCALING EQUILIBRIUM CONSTRUCTION

Inputs. For each counterfactual, we take the following auction-specific objects as inputs:

- $q_{i,n}^e$: DOT quantity estimates
- $q_{i,n}^b$: Bidder quantity predictions
- $\sigma_{i,n}^2$: Bidder predictions' variances
- $c_{i,n}$: DOT cost estimates
- I_n : The number of participating bidders
- $F_n(\tau)$ and $f_n(\tau)$: CDF and PDF of the distribution of bidder types.

The objects $q_{i,n}^e$, $c_{i,n}$ and I_n are taken directly from data provided by the DOT. The remaining objects are estimates derived from Section 7. The estimates of $q_{i,n}^b$ and $\sigma_{i,n}^2$ are taken from the first stage of our estimation procedure. To obtain $F_n(\tau)$ and $f_n(\tau)$, we fit the estimates of bidder-auction types $\alpha_{i,n}$ and $\gamma_{i,n}$ from the second stage as described below. For simplicity of exposition, we omit the $\hat{\cdot}$ mark when referring to these estimated parameters.

To estimate the distribution of bidder efficiency and risk-aversion types, we first project our bidder-auction type parameters $\alpha_{i,n}$ and $\gamma_{i,n}$ onto a single dimension $\tau_{i,n}$. For simplicity, we normalize with respect to efficiency types, so that $\tau_{i,n} = \alpha_{i,n}$ for each (i, n) pair. We then fit the risk-aversion types $\gamma_{i,n}$ to a Poisson regression model of $\log(\alpha_{i,n})$ with auction fixed effects.¹⁷ Finally, we apply the fitted regression model to project a unique risk-aversion type $\gamma_n(\alpha)$ for each α in each auction n . In summary, we obtain the following map from τ to efficiency and risk aversion: $\alpha_n(\tau) = \tau$ and $\gamma_n(\tau) = g_n(\alpha_n(\tau))$ where g_n is the regression model fit for a bidder with efficiency $\alpha_n(\tau)$ participating in auction n . Note that because the regression model includes auction fixed effects, the projection $g_n(\cdot)$ will generally differ across auctions.

Next, we fit the distribution of estimated efficiency types $\alpha_{i,n}$ to a truncated log-normal distribution with a mean that depends on project characteristics according to $\mu_n^\alpha = \beta^\alpha X_n$ and a project-type-specific variance σ_n^α . We estimate β^α and σ_n^α from the distribution of estimated $\alpha_{i,n}$ parameters across bidders and auctions, using Hamiltonian Monte Carlo sampling.¹⁸ The resulting fitted parameters fully characterize the distribution of bidder types in each auction.

¹⁷As we discuss in Section 7, a log-linear fixed-effects regression of $\log(\gamma_{i,n})$ on $\alpha_{i,n}$ yields an R^2 of 0.80. The log-Poisson regression model improves the fit and increases R^2 to 0.86. In Appendix D, we present a prediction fit regression (Figure D.4(b)) and a histogram of the residuals (Figure D.4(a)) for the log-Poisson model.

¹⁸Since efficiency types are drawn from a truncated distribution, within the same sampling procedure we also fit that maximum type $\bar{\alpha}_n$ in each auction, as well as a maximum possible type for each bin of auctions, $\bar{\alpha}_{B(n)}$, to jointly rationalize the distribution of entrants in each bin. We compute $\bar{\alpha}_n$ again when calibrating entry costs and extra work-order coefficients, using the rational entry condition to add precision.

Equilibrium Construction. Bidders first choose whether or not to enter the auction; then, if they choose to enter, they choose their bids according to a symmetric equilibrium in which the distribution of bidder types and the number of competing bidders is common knowledge. To characterize equilibrium outcomes, we proceed by backward induction. We first characterize equilibrium bidding upon entry; then we characterize equilibrium entry strategies. Since each auction is considered independently under our model, we suppress the auction-specific marker n in notation below for expositional simplicity.

Equilibrium Bidding. As discussed in Section 8, the expected utility that a bidder i receives for participating in an auction with m bidders is given by

$$\begin{aligned} \text{EU}(s|\tau_i, m) = & [1 - \exp(-\gamma(\tau_i) \cdot [\text{CE}(\mathbf{b}_i^*(s)|\tau_i) + \lambda \cdot \text{EWO} - \kappa])] \times \prod_{j \neq i} [1 - H_j(s)] \\ & + [1 - \exp(\gamma(\tau_i) \cdot \kappa)] \times \left(1 - \prod_{j \neq i} [1 - H_j(s)]\right), \end{aligned} \quad (12)$$

where $\text{CE}(\mathbf{b}_i^*(s)|\tau_i)$ is given by

$$\text{CE}(\mathbf{b}_i^*(s)|\tau_i) = \sum_{t=1}^T q_t^b (b_{t,i}^*(s) - \alpha(\tau_i) \cdot c_t) - \frac{\gamma(\tau_i) \cdot \sigma_t^2}{2} \cdot (b_{t,i}^*(s) - \alpha(\tau_i) \cdot c_t)^2, \quad (13)$$

and the vector $\mathbf{b}_i^*(s)$ is given by the solution to the portfolio problem:

$$\begin{aligned} \mathbf{b}_i^*(s) = & \arg \max_{\mathbf{b}} \left[\sum_{t=1}^T q_t^b (b_t - \alpha(\tau_i) \cdot c_t) - \frac{\gamma(\tau_i) \cdot \sigma_t^2}{2} \cdot (b_t - \alpha(\tau_i) \cdot c_t)^2 \right] \\ \text{s.t. } & \sum_{t=1}^T b_t q_t^e = s \quad \text{and} \quad b_t \geq 0 \quad \text{for all } t. \end{aligned} \quad (14)$$

Note that as unit bids cannot be negative, the portfolio problem in Equation (14) does not have a closed-form solution, and must be solved numerically. In every instance that optimal bids must be evaluated, we compute them through a constrained quadratic programming solver. While standard quadratic solvers should also work well, for computational efficiency we use a custom algorithm specified to our problem as detailed in Online Appendix F.

We assume that bidder types τ are drawn IID from the auction-wide distribution and construct a symmetric equilibrium in monotone strategies. Writing the equilibrium bidding function: $\varphi : [\underline{\tau}, \bar{\tau}] \rightarrow \mathbb{R}$, we can rewrite the probability that i will win the auction under bid s as follows:

$$\prod_{j \neq i} [1 - H_j(s)] = \prod_{j \neq i} [Pr(s < \varphi_j(\tau_j))] \quad (15)$$

$$= \prod_{j \neq i} [1 - F(\varphi_j^{-1}(s))] \quad (16)$$

$$= [1 - F(\varphi^{-1}(s))]^{m-1}, \quad (17)$$

where Equation (16) follows from the monotonicity of the equilibrium bidding function, and Equation (17) follows from symmetry, by which all bidders use the same equilibrium bidding function. To simplify notation, we rewrite Equation (12) as follows, dropping the i subscript:

$$\text{EU}(\varphi(\tau)|\tau, m) = a_0(\tau) + a_1(\tau)(V(\varphi(\tau)) \times [1 - F(\varphi^{-1}(\varphi(\tau)))]^{m-1}), \quad (18)$$

where $V(\varphi(\tau))$ is the expected utility of winning the auction under the strategy φ and type τ , and $a_0(\tau)$ and $a_1(\tau)$ are constants that do not vary with the score s .¹⁹ We proceed following an adaption of the procedure detailed in [Hubbard and Paarsch \(2014\)](#). Differentiating with respect to s , we obtain the following first-order condition:

$$\frac{\partial \text{EU}(\varphi(\tau)|\tau, m)}{\partial s} = 0, \quad (19)$$

where

$$\begin{aligned} \frac{\partial \text{EU}(\varphi(\tau)|\tau, m)}{\partial s} &= a_1(\tau) \left(\frac{\partial}{\partial s} V(\varphi(\tau)) \times [1 - F(\varphi^{-1}(\varphi(\tau)))]^{m-1} \right) \\ &\quad + a_1(\tau) \left(V(\varphi(\tau)) \times \frac{\partial}{\partial s} [1 - F(\varphi^{-1}(\varphi(\tau)))]^{m-1} \right). \end{aligned}$$

Writing $\widetilde{\text{CE}}(\varphi(\tau))$ as shorthand for $(\text{CE}(\mathbf{b}(\varphi(\tau))|\tau) + \lambda \cdot \text{EWO})$, we obtain $\frac{\partial}{\partial s} V(\varphi(\tau))$ by

$$\begin{aligned} V(\varphi(\tau)) &= 1 - \exp[-\gamma(\tau) \cdot \widetilde{\text{CE}}(\varphi(\tau))], \\ \frac{\partial}{\partial s} V(\varphi(\tau)) &= \gamma(\tau) \cdot \frac{\partial}{\partial s} \widetilde{\text{CE}}(\varphi(\tau)) \times \exp[-\gamma(\tau) \cdot \widetilde{\text{CE}}(\varphi(\tau))], \\ \frac{\partial}{\partial s} \widetilde{\text{CE}}(\varphi(\tau)) &= \sum_{t=1}^T \left[\frac{\partial b_t^*(\varphi(\tau))}{\partial s} (q_t^b - \gamma(\tau) \cdot \sigma^2(b_t^*(\varphi(\tau)) - \alpha(\tau) \cdot c_t)) \right]. \end{aligned}$$

Here, the derivative of $b_t^*(\varphi(\tau))$ is taken with respect to the solution of the portfolio problem in Equation (14), and can be computed exactly through forward-mode autodifferentiation of our portfolio optimization algorithm. The derivative of the second part of Equation (18) is given through the product rule and

$$\begin{aligned} \frac{\partial}{\partial s} [1 - F(\varphi^{-1}(\varphi(\tau)))] &= [-f(\varphi^{-1}(\varphi(\tau)))] \times \frac{1}{\varphi'(\varphi^{-1}(\varphi(\tau)))} \\ &= \frac{-f(\tau)}{\varphi'(\tau)}. \end{aligned}$$

To find the equilibrium bidding function, we solve the differential equation in Equation (19) using stiff ODE methods implemented by [Rackauckas and Nie \(2017\)](#). The ODE is defined with respect to an initial boundary condition in which the threshold (highest

¹⁹Here, $a_0(\tau) = 1 - \exp(\gamma(\tau) \cdot \kappa)$ and $a_1(\tau) = \exp(\gamma(\tau) \cdot \kappa)$. As these do not depend on the bids submitted to the auction, they cancel out in Equation (19) and can be considered constants.

τ type receives zero utility upon winning. We find the score that generates this condition when the threshold type (like all others) uses our portfolio maximization program to choose her bids at any score.

Lump Sum Equilibria. As noted in Section 8, the lump sum problem is identical to the scaling auction problem except for the formulation of the certainty equivalent. Here, there is no portfolio problem and the certainty equivalent is given directly by the bidder's score:

$$\widetilde{\text{CE}}_{\text{Lump}}(\varphi(\tau)) = \varphi(\tau) - \left[\sum_{t=1}^T q_t^b(\alpha(\tau) \cdot c_t) + \frac{\gamma(\tau) \cdot \sigma_t^2}{2} \cdot (\alpha(\tau) \cdot c_t)^2 \right] + \lambda \cdot \text{EWO}.$$

With μ -renegotiation,

$$\begin{aligned} \widetilde{\text{CE}}_{\mu}(\varphi(\tau)) = & \lambda \cdot \text{EWO} + \sum_{t=1}^T (\mu q_t^b + (1 - \mu) q_t^e) \cdot b_t^*(\varphi(\tau)) - q_t^b(\alpha(\tau) \cdot c_t) \\ & - \frac{\gamma(\tau) \cdot \sigma_t^2}{2} \cdot (\mu \cdot b_t^*(\varphi(\tau)) - \alpha(\tau) \cdot c_t)^2. \end{aligned}$$

Equilibrium Entry. Prior to choosing whether or not to enter an auction, each bidder observes her type τ and decides whether it would be profitable in expectation to enter. We construct a monotone pure strategy equilibrium such that all types below a threshold τ^* enter, and all types above the threshold stay out of the auction. That is, for every type $\tau < \tau^*$, the expected value of entry is positive: $V(\tau) > 0$, where

$$V(\tau) = \sum_{m=1}^M \binom{M-1}{m-1} F(\tau^*)^{m-1} (1 - F(\tau^*))^{M-m} \cdot \text{EU}(\varphi^*(\tau)|\tau, m). \quad (20)$$

Here, $(1 - F(\tau^*))$ is the probability that an independent draw of a bidder type is below τ^* under the distribution of types in the auction, and $\text{EU}(\varphi^*(\tau)|\tau, m)$ is the expected utility that the bidder expects to earn under her equilibrium bidding strategy $\varphi^*(\tau)$ as defined above, in the case that $m - 1$ competing bidders participate.²⁰ In the case that there are no other competing bidders (e.g., $m = 1$), we assume that the bidder submits the maximum allowable amount as her score and optimizes her unit bid spread subject to this total.²¹ By construction, the marginal (threshold)-type τ^* expects to earn no profit, and so equilibrium entry probabilities are defined by the equation $V(\tau^*) = 0$.

To find the threshold type in each auction, we numerically solve for the root of $V(\tau^*)$. The solution then provides both the probability that each number of entrants would be realized, and the worst (highest) type of bidder in the auction, with respect to which the equilibrium in Equation (19) is defined. Note that the threshold type plays two roles in

²⁰Note that although we did not demarcate this explicitly, the equilibrium bidding strategy $\varphi^*(\tau)$ itself also depends on τ^* as the distribution of competing bidder types affects the competitiveness of bidding.

²¹Although there is no explicit maximum allowable bid, MassDOT guidelines state that monopolist bidders with scores more than 5% over the DOT estimate warrant scrutiny, and allow rejecting bids with items that are priced 25% over the DOT estimate under certain conditions. For simplicity, we take 125% of the DOT score as a maximum upper bound on the score that a monopolist can submit. In simulations, we found that increasing or decreasing this bound does not substantially change results.

the equilibrium definition: it both (a) provides the boundary no-profit condition that pins down a unique solution to the ODE in Equation (19); and (b) determines the truncation of the CDF of types among bidders participating in the auction.

Computing Equilibrium Outcomes. To evaluate counterfactual equilibrium outcomes, we first compute the threshold-type τ^* and the equilibrium mapping $\varphi(\cdot)$ for each auction as described above. We then compute the expected (1) DOT cost and (2) bidder certainty equivalent by integrating over the distribution of winning types, across all combinations of bidder entries:

$$\overline{\text{COST}} = \sum_{m=1}^M \binom{M-1}{m-1} F(\tau^*)^{m-1} (1 - F(\tau^*))^{M-m} \int_{\underline{\tau}}^{\tau^*} \sum_t (q_t^a \cdot b_t^*(\varphi_m^*(\tilde{\tau}))) \cdot dF_m^1(\tilde{\tau}) d\tilde{\tau},$$

$$\overline{\text{CE}} = \sum_{m=1}^M \binom{M-1}{m-1} F(\tau^*)^{m-1} (1 - F(\tau^*))^{M-m} \int_{\underline{\tau}}^{\tau^*} \sum_t \widetilde{\text{CE}}(\varphi_m^*(\tilde{\tau})) \cdot dF_m^1(\tilde{\tau}) d\tilde{\tau},$$

where $\varphi_m^*(\cdot)$ is shorthand for $\varphi^*(\cdot|m)$ and F_m^1 is the CDF of the first-order statistic of the bidder-type distribution in the auction when m total bidders are present.

APPENDIX B: ROBUSTNESS TO MORAL HAZARD

Although lump sum auctions induce more risk for bidders, they also provide an incentive for the winning bidder to minimize her costs. As MassDOT bridge procurement is heavily monitored, our baseline model assumes that ex post bidder costs are exogenously determined by the quantity distributions involved in each project. However, if bidders *were* able to influence the ex post realization of item quantities, then the winning bidder would be able to make additional profit by overusing items she had bid high on. Knowing this, she might bid higher on items that will be easier to overuse.

Although we cannot identify the presence of such moral hazard from our data, our framework can be extended to account for it in counterfactual simulations. In this section, we consider a model of moral hazard in which bidders are able to choose which item quantities to augment (and by how much) before they choose optimal unit bids. Our model is aimed to capture the most extreme version of moral hazard that may be consistent with our data: not only can bidders overuse items that they bid high on, but they can also strategically choose higher bids on items that they intend to overuse at the bidding stage.

In order to capture the constraints that limit bidders from overusing profitable items ad infinitum in a conservative way, we assume that the extent of bidders' quantity adjustments are bounded by the observed levels of overrunning on profitable items. That is, for each auction n , we take the set of items for which the winning bidder made a profit from overruns. For each such item t , we define the maximum allowable overuse level: $\bar{y}_{t,n} = (q_{t,n}^a - q_{t,n}^e)/q_{t,n}^e$. We assume that all profitable overruns in our data were intentional, so that the variance on these items, $\bar{\sigma}_{t,n}$, is 0.²²

To evaluate the extent to which our counterfactuals could change under moral hazard, we solve for the equilibrium quantity, unit bid and score for each auction, bidder and

²²For all other items, we set $\bar{y}_{t,n} = 0$ and keep the variance estimates as in our baseline model.

TABLE B.I

A COMPARISON OF COUNTERFACTUAL SAVINGS WITH AND WITHOUT MORAL HAZARD^a.

CF Type	Moral Hazard?	Outcome	Mean	SD	25%	50%	75%
Lump Sum	No	% DOT Savings	-264.39	421.65	-273.59	-104.92	-46.99
Lump Sum	Yes	% DOT Savings	-209.35	391.22	-204.13	-73.59	-21.23
50-50 Renegotiation	No	% DOT Savings	-10.81	16.05	-17.27	-6.78	-1.99
50-50 Renegotiation	Yes	% DOT Savings	-13.13	14.95	-20.61	-10.52	-4.98
2:1 Renegotiation	No	% DOT Savings	-24.90	30.45	-39.98	-17.73	-6.56
2:1 Renegotiation	Yes	% DOT Savings	-26.01	28.49	-40.81	-20.15	-9.81

^aDue to increased numerical errors, the sample of auctions being compared in this exercise is a bit smaller (95% for lump sum and 50-50 negotiation; 85% for 2:1 negotiation) than in our main counterfactual results.

item in each auction format. To do so, we modify Equation (5) to solve for the optimal bids $\mathbf{b}_n^*(s|\tau)$ and overuses $\mathbf{y}_n^*(s|\tau)$ as follows:

$$\arg \max_{\{\mathbf{b}, \mathbf{y}\}} \left[\sum_{t=1}^T (1 + y_{t,n}) \cdot q_{t,n}^b \cdot (b_t - \alpha_n(\tau) \cdot c_{t,n}) - \frac{\gamma_n(\tau) \cdot \bar{\sigma}_{t,n}^2}{2} (b_t - \alpha_n(\tau) \cdot c_{t,n})^2 \right]$$

$$\text{s.t. } \sum_{t=1}^T b_t \cdot q_{t,n}^e = s, b_t \geq 0, \quad \text{and} \quad 0 \leq y_t \leq \bar{y}_{t,n} \quad \text{for all } t.$$

For pure lump sum auctions, there is no incentive to overuse quantities, and so no items are overused (although their variance is still assumed to be zero). For lump sum auctions with renegotiation, Equation (11) is adjusted similar to the baseline case, where $q_{t,n}^b$ can be inflated by $y_{t,n} \times 100$ percent, up to the maximum overuse level $\bar{y}_{t,n}$ for each item.

To solve for optimal bids and quantities under moral hazard, we recast the augmented optimization problem in \mathbf{b}_n and \mathbf{y}_n as a mixed-integer program, using the observation that if it is optimal to overuse an item t , it must be optimal to overuse it as much as possible. Given the added complexity of these problems, we compute equilibrium savings under the observed number of bidders in each auction only, and do not account for endogenous entry.²³ However, given the magnitudes of outcomes with and without moral hazard, we do not expect that the results would change qualitatively under the full endogenous entry model.

In Table B.I, we present a comparison of the percent DOT savings under each of the lump sum counterfactuals with and without moral hazard. Our results suggest that the possibility of moral hazard would not substantially change our analysis. For lump sum auctions with renegotiation, the median auction is almost unaffected even under our conservative definition of moral hazard. In fact, the cost of moving from the baseline to a lump sum format increases by a few percentage points on median with 50-50 and 2:1 renegotiation. The cost of moving to a lump sum auction without renegotiation is more affected: the cost of a median auction decreases by about 30-percentage points. Nevertheless, lump sum auctions remain much more costly than scaling auctions and our qualitative conclusions hold.

²³Given the added complexity of mixed-integer programming, we solve the augmented bid optimization problems with a custom solver provided by Gurobi. While this works well, Gurobi licensing restrictions limit the number of calls that can be made at a time.

APPENDIX C: TECHNICAL DETAILS

C.1. *Proof of Equilibrium Existence*

While our counterfactual equilibrium construction projects bidder efficiency and risk-aversion types on a single dimension, our estimation procedure allows for arbitrary correlations between α and γ across and within bidders. For any equilibrium in which a bidder with type (α, γ) submits a score s , the optimal vector of unit bids is characterized by Equation (5). While we cannot guarantee the uniqueness of an equilibrium without further assumptions, the application of [Reny \(2011\)](#) below implies that we can interpret the distribution of scores observed in our data as *an* equilibrium under some conditions.²⁴

PROPOSITION 1: *Suppose that each bidder's type (α, γ) is drawn from a 2-dimensional Euclidean cube $[\underline{\alpha}, \bar{\alpha}] \times [\underline{\gamma}, \bar{\gamma}]$. Then when the support of feasible scores is sufficiently high, there exists a monotone pure-strategy equilibrium in which each type (α, γ) submits a score $s^*(\alpha, \gamma)$, and a vector of unit bids $\mathbf{b}(s^*(\alpha, \gamma), \alpha, \gamma)$, characterized by the solution to the quadratic program:*

$$\begin{aligned} \mathbf{b}^*(s, \alpha, \gamma) = \arg \max_{\mathbf{b}} & \left[\gamma_i \sum_{t=1}^T q_t^b (b_t(s, \alpha, \gamma) - \alpha c_t) - \frac{\gamma \sigma_t^2}{2} (b_t(s, \alpha, \gamma) - \alpha c_t)^2 \right] \\ \text{s.t. } & \sum_{t=1}^T b_t(s, \alpha, \gamma) \cdot q_t^e = s \quad \text{and} \quad b_t(s, \alpha, \gamma) \geq 0 \quad \text{for all } t. \end{aligned} \quad (21)$$

PROOF: Consider the normal form representation of the game described in Section 5, such that each bidder's action corresponds to a score s , and for a given score s , a bidder of type (α, γ) obtains the expected utility of winning generated by the vector of unit bids $\mathbf{b}^*(s, \alpha, \gamma)$ defined in Equation (21). This game immediately satisfies conditions (i) and (ii) of Proposition 3.1 in [Reny \(2011\)](#), as the support of the type space is a compact subset of Euclidean space and the action space is one-dimensional. By Corollary 4.2 of [Reny \(2011\)](#), there exists an equilibrium in monotone pure strategies if for each bidder i and for every monotone joint pure strategy σ_{-i} of other players, bidder i 's expected utility from winning, $W(\cdot, \sigma_{-i})$, satisfies increasing differences in each dimension of the bidder-type space.

Note that the solution to Equation (21) can be characterized in closed form once the set of nonzero bids is known:

$$b^*(s, \alpha, \gamma) = \alpha c_t + \frac{q_t^b / \sigma_t^2}{\gamma} + \frac{q_t^e / \sigma_t^2}{\sum_{r: b_r^*(s, \cdot) > 0} \left[\frac{(q_r^e)^2}{\sigma_r^2} \right]} \left(s - \sum_{r: b_r^*(s, \cdot) > 0} q_r^e \left[\alpha c_r + \frac{q_r^b / \sigma_r^2}{\gamma} \right] \right). \quad (22)$$

As a bidder's expected utility of winning with a score s does not depend on the scores of her opponents, we can drop the dependency of W on σ_{-i} without loss. As $\underline{\alpha} > 0$ and $\underline{\gamma} > 0$, it is sufficient to show increasing differences for the certainty equivalent function:

$$CE(s, \alpha, \gamma) = \sum_{t=1}^T q_t^b (b_t^*(s, \alpha, \gamma) - \alpha c_t) - \frac{\gamma \sigma_t^2}{2} (b_t^*(s, \alpha, \gamma) - \alpha c_t)^2.$$

²⁴We are grateful to Paulo Somaini for pointing out the connection to [Reny \(2011\)](#) underlying this proof.

Taking derivatives, we obtain

$$\frac{\partial CE(s, \alpha, \gamma)}{\partial s} = \sum_{t=1}^T \frac{\partial b_t^*(s, \alpha, \gamma)}{\partial s} \cdot [q_t^b - \gamma \sigma_t^2 (b_t^*(s, \alpha, \gamma) - \alpha c_t)].$$

Noting that

$$\frac{\partial b_t^*(s, \alpha, \gamma)}{\partial s} = \frac{q_t^e / \sigma_t^2}{\sum_{r: b_r^*(s, \cdot) > 0} \left[\frac{(q_r^e)^2}{\sigma_r^2} \right]} > 0$$

does not depend on α or γ , and

$$\frac{\partial (b_t^*(s, \alpha, \gamma) - \alpha c_t)}{\partial \alpha} = - \frac{q_t^e / \sigma_t^2 \cdot \sum_{r: b_r^*(s, \cdot) > 0} q_r^e c_r}{\sum_{r: b_r^*(s, \cdot) > 0} \left[\frac{(q_r^e)^2}{\sigma_r^2} \right]} < 0,$$

it follows that $\frac{\partial^2 CE(s, \alpha, \gamma)}{\partial s \partial \alpha} > 0$ for all s and α . Considering γ , we obtain

$$\frac{\partial^2 CE(s, \alpha, \gamma)}{\partial s \partial \gamma} = \frac{\sum_{r: b_r^*(s, \cdot) > 0} [q_r^e \cdot \alpha c_r] - s}{\sum_{r: b_r^*(s, \cdot) > 0} \left[\frac{(q_r^e)^2}{\sigma_r^2} \right]}. \quad (23)$$

Note that the numerator of Equation (23) is negative whenever the score exceeds the ex ante (DOT-predicted) cost of completing the auction. Thus, when the support of feasible scores is high enough that $s > \bar{\alpha} \sum_i [q_i^e c_i]$, we have $\frac{\partial^2 CE(s, \alpha, \gamma)}{\partial s \partial \gamma} < 0$ for all s, α and γ . Reparameterizing bidder risk-aversion types according to $\tilde{\gamma} = -\gamma$, we obtain the increasing differences condition: $\frac{\partial^2 CE(s, \alpha, \tilde{\gamma})}{\partial s \partial \tilde{\gamma}} > 0$. *Q.E.D.*

C.2. Projecting Items and Bidder-Auction Pairs Onto Characteristic Space

Our data set consists of 440 bridge projects with a total of 218,110 unit bid observations. Of these, there are 2883 unique bidder-project pairs and 29,834 unique item-project pairs. Each auction has an average of 6.55 bidders and 67.8 items. Of these, there are 116 unique bidders and 2985 unique items (as per the DOT's internal taxonomy). In order to keep the computational burden of our estimator within a manageable range, while still capturing heterogeneity across bidders and items within and across projects, we project item-project and bidder-project pairs onto characteristic space.

We first build a characteristic-space model of items as follows. The DOT codes each item observation in two ways: a 6-digit item ID, and a text description of what the item is. Each item ID comprises a hierarchical taxonomy of item classification. That is, the more digits two items have in common (from left to right), the closer the two items are. For

example, item 866100—also known as “100 Mm Reflect. White Line (Thermoplastic)” —is much closer to item 867100—“100 Mm Reflect. Yellow Line (Thermoplastic)” —than it is to item 853100—“Portable Breakaway Barricade Type Iii”—and even farther from item 701000—“Concrete Sidewalk.” To leverage the information in both the item IDs and the description, we break the IDs into digits, and tokenize the item description.²⁵ We then add summary statistics for each item: the relative commonness with which the item is used in projects, the average DOT cost estimate for that item, dummies that indicate whether or not the item is frequently used as a single unit, and whether the item is often ultimately not used at all.

We create an item-project level characteristic matrix by combining the item characteristic matrix with project-level characteristics: the project category, the identities of the project manager, designer, and engineer, the district in which the project is located, the project duration, the number of items in the project specification that the engineer has flagged for us as “commonly skewed,” and the share of projects administered by the manager and engineer that over/under-ran.²⁶ The resulting matrix is very high-dimensional, and so we project the matrix onto its principal components, and use the first 15.²⁷ In addition, we added 3 stand-alone item features: a dummy variable indicating whether the item is often given a single unit quantity, the historical share of observations of that item in which it was not used at all, and an indicator for whether or not the item itself is a “commonly skewed” item. The result is the matrix $X_{i,n}$, used in the estimation in Equation (24).

To estimate bidder-types $\alpha_{i,n}$ and $\gamma_{i,n}$ for each bidder-auction pair, we combine each bidder’s firm ID with the matrix of project characteristics described above, and a matrix of project-bidder specific features. As a number of bidders only participate in a few auctions, we combine all bidders who appear in fewer than 30 auctions in our data set into a single firm ID. This results in 25 unique bidder IDs: 24 unique firms and one aggregate group. For project-bidder characteristics, we compute the bidder’s *specialization* in each project type—the share of projects of the same type as the current project that the bidder has bid on, the bidder’s *capacity*—the maximum number of DOT projects that the DOT has ever had open while bidding on another project, and the bidder’s *utilization*—the share of the bidder’s capacity that is filled when she is bidding on the current project. We also include dummies for whether or not the bidder is a *fringe* bidder, and whether or not the bidder’s headquarters is located in the same district as the project at hand.²⁸ Our $X_{i,n}$ matrix has a total of 14 columns, and so we have a total of 78 bidder-type parameters to identify. We use $X_{i,n}$ and the unique bidder IDs to model $\alpha_{i,n}$ and $\gamma_{i,n}$ in Equation (30).

²⁵That is, we split each description up by words, clean them up and remove common “stop” words. Then we create a large dummy matrix in which entry i, j is 1 if the unique item indexed at i contains the word indexed by j in its description. We owe a big thanks to Jim Savage for suggesting this approach.

²⁶There are 11 items that have been flagged at our request by the chief engineer: 120100: Unclassified Excavation; 129600: Bridge Pavement Excavation; 220000: Drainage Structure Adjusted; 450900: Contractor Quality Control; 464000: Bitumen For Tack Coat; 472000: Hot Mix Asphalt For Miscellaneous Work; 624100: Steel Thrie Beam Highway Guard (Double Faced); 851000: Safety Controls For Construction Operations (Traffic Cones For Traffic Management); 853200: Temporary Concrete Barrier; 853403: Movable Impact Attenuator; 853800: Temporary Illumination For Work Zone (Temporary Illumination For Night Work).

²⁷We have tried replicating this using more/fewer principal components and the results are very stable.

²⁸We define “fringe” similar to Bajari, Houghton, and Tadelis (2014), as a firm that receives less than 1% of the total funds spent by the DOT on projects within the same project type as the auction being considered, within the scope of our data set.

C.3. Econometric Details

Let $b_{t,i,n}^d$ denote the unit bid observed by the econometrician for item t by bidder i in auction n . Let $\theta = (\theta_1, \theta_2)$ be the vector of variables that parametrize the model prediction for each bid $b_{t,i,n}^*(s_{i,n}^*|\theta)$, as defined by Equation (6). The subvector θ_1 refers to parameters estimated in the first stage, as detailed in Appendix C.3.1. The subvector θ_2 refers to parameters estimated in the second stage, as detailed in Appendix C.3.2. By Assumption 1, the residual of the optimal bid for each item-bidder-auction tuple with respect to its noisily observed bid, $v_{t,i,n} = b_{t,i,n}^d - b_{t,i,n}^*(s_{i,n}^*|\theta)$, is distributed identically and independently with a mean of zero across items, bidders, and auctions. Furthermore, $v_{t,i,n}$ is orthogonal to the identity and characteristics of each item, bidder, and auction.

Our estimation procedure treats each auction n as a random sample from some unknown distribution. As such, auctions are exchangeable. Each auction n has an associated set of bidders who participate in the auction, $\mathcal{I}(n)$, as well as an associated set of items that receive bids in the auction, $\mathcal{T}(n)$. $\mathcal{I}(n)$ and $\mathcal{T}(n)$ are characteristics of auction n and so are drawn according to the underlying distribution over auctions themselves. For each bidder $i \in \mathcal{I}(n)$ and item $t \in \mathcal{T}(n)$, our model assigns a unique true bid $b_{t,i,n}^*(s_{i,n}^*|\theta)$ at the true parameter vector θ .

Items $t \in \mathcal{T}(n)$ are characterized by a $P \times 1$ vector, $X_{t,n}$, of features. Bidders $i \in \mathcal{I}(n)$ are characterized by a $J \times 1$ vector, $X_{i,n}$, of features. The construction of $X_{t,n}$ and $X_{i,n}$ is discussed in detail in Appendix C.2. Estimation proceeds in two stages. In the first stage, we estimate θ_1 , the set of parameters that govern bidders' beliefs over ex post item quantities, using a best-predictor model estimated with Hamiltonian Monte Carlo. In the second stage, we estimate θ_2 , which characterizes bidders' risk aversion and cost types, using a GMM estimator.

C.3.1. First Stage

In the first stage, we use the full data set of auctions available to us in order to estimate a best-predictor model of expected item quantities conditional on DOT estimates and project-item characteristics, as well as the level of uncertainty that characterizes each projection.

Each observation is an instance of an item t , being used in an auctioned project n . Each observation (t, n) is associated with a vector of item-auction characteristic features $X_{t,n}$, the construction of which is discussed in Appendix C.2. For simplicity, we employ a linear model for $q_{t,n}^b$, the expected quantity of item t in auction n , as a function of the DOT quantity estimate $q_{t,n}^e$ and $X_{t,n}$.²⁹ In order to model the level of uncertainty in the projection $q_{t,n}^b$, we model the distribution of the quantity model fit residuals ($\eta_{t,n} = q_{t,n}^a - q_{t,n}^b$) with a LogNormal regression function of $q_{t,n}^e$ and $X_{t,n}$. The full model specification is below:

$$q_{t,n}^a = q_{t,n}^b + \eta_{t,n} \quad \text{where } \eta_{t,n} \sim \mathcal{N}(0, \sigma_{t,n}^2) \quad (24)$$

such that

$$q_{t,n}^b = \beta_{0,q} q_{t,n}^e + \beta_q X_{t,n} \quad \text{and} \quad \sigma_{t,n} = \exp(\beta_{0,\sigma} q_{t,n}^e + \beta_\sigma X_{t,n}). \quad (25)$$

While we could fit this in two stages (first, fit the expected quantity and then fit the distribution of the residuals), we do this jointly using Hamiltonian Monte Carlo (HMC)

²⁹In principle, any statistical model (not necessarily a linear one) would be sound.

(Betancourt (2017)) with the Stan probabilistic programming language.³⁰ We then take the posterior means of the estimated distributions and use them as point estimates for the second stage. Denote $\theta_1 = (\beta_{0,q}, \beta_q, \beta_{0,\sigma}, \beta_\sigma)$ for the vector of first-stage parameters and let $\hat{\theta}_1$ be the posterior means of θ_1 , produced by the first-stage HMC estimation. Thus, $\hat{\theta}_1$ specifies, for each item $t \in \mathcal{T}(n)$ in each auction n , the model estimate of bidders' predictions for the item's quantity, $\hat{q}_{t,n}^b$, as well as the variance of that prediction, $\hat{\sigma}_{t,n}^2$.

C.3.2. Second Stage

C.3.2.1. Identification. Our model of equilibrium bidding in Section 5 states that the optimal bid vector for a bidder with efficiency-type $\alpha_{i,n}$ and risk-aversion type $\gamma_{i,n}$, submitting a total score of $s_{i,n}$ is given by

$$b_{t,i,n}^*(s_{i,n}) = \alpha_{i,n} c_{t,n} + \frac{q_{t,n}^b / \sigma_{t,n}^2}{\gamma_{i,n}} + \frac{q_{t,n}^e / \sigma_{t,n}^2}{\sum_{r: b_{r,i,n}^*(s_{i,n}) > 0} \left[\frac{(q_{r,n}^e)^2}{\sigma_{r,n}^2} \right]} \left(s_{i,n} - \sum_{r: b_{r,i,n}^*(s_{i,n}) > 0} q_{r,n}^e \left[\alpha_{i,n} c_{r,n} + \frac{q_{r,n}^b / \sigma_{r,n}^2}{\gamma_{i,n}} \right] \right),$$

where $(q_{1,n}^b, \dots, q_{T_n,n}^b)$ and $(\sigma_{1,n}^2, \dots, \sigma_{T_n,n}^2)$ are exogenously fixed and commonly observed by all bidders. Collecting exogenous terms, we can rewrite $b_{t,i,n}^*(s_{i,n})$ as a linear projection of $\alpha_{i,n}$, $\frac{1}{\gamma_{i,n}}$ and $s_{i,n}$:

$$b_{t,i,n}^*(s_{i,n}) = A_{t,i,n} \cdot \alpha_{i,n} + G_{t,i,n} \cdot \frac{1}{\gamma_{i,n}} + Z_{t,i,n} \cdot s_{i,n}. \quad (26)$$

Here, $A_{t,i,n}$, $G_{t,i,n}$, and $Z_{t,i,n}$ capture the variation in the relative value of bidding higher on item (t, i, n) relative to the other items in bidder i 's portfolio in auction n . For instance,

$$A_{t,i,n} = c_{t,n} - \frac{q_{t,n}^e / \sigma_{t,n}^2}{\sum_{r: b_{r,i,n}^*(s_{i,n}) > 0} \left[\frac{(q_{r,n}^e)^2}{\sigma_{r,n}^2} \right]} \sum_{r: b_{r,i,n}^*(s_{i,n}) > 0} q_{r,n}^e c_{r,n}$$

corresponds to the difference between the unit market rate of item (t, i, n) and the expected sum of market rates (given DOT quantity estimates) of all of the items in (i, n) 's portfolio, weighted by the relative DOT-quantity-to-uncertainty ratio of item (t, i, n) with respect to the rest of the portfolio. Similarly, $G_{t,i,n}$ corresponds to the difference in the bidders' expected quantity-to-uncertainty ratio of item (t, i, n) to the sum of bidder quantity-to-uncertainty ratios among the items in (i, n) 's portfolio, weighted by the relative DOT quantity-to-uncertainty ratio. All else held fixed, a higher individual market rate $c_{t,i,n}$ or bidder quantity-to-uncertainty ratio $q_{t,n}^b / \sigma_{t,n}^2$ corresponds to a higher unit bid $b_{t,i,n}^*$. However, the trade-off by which items with higher costs are compensated with higher bids

³⁰See Carpenter et al. (2017).

depends on the weight of relative risks and returns across the portfolio, as moderated by the cost efficiency and risk-aversion coefficients $\alpha_{i,n}$ and $\gamma_{i,n}$. As such, our model predicts that the variation in unit bids observed in our data is driven by variation in how the distribution of market rates and quantity/uncertainty predictions across different items and auctions generates optimal portfolios for different bidder types.

We make two further assumptions: (1) bidder beliefs over $q_{t,n}^b$ and $\sigma_{t,n}^2$ are pinned down by our first-stage estimates $\hat{q}_{t,n}^b$ and $\hat{\sigma}_{t,n}^2$; (2) optimal unit bids are observed with an exogenous mean-zero measurement error $\nu_{t,i,n}$. Note that without the second assumption, Equation (26) would result in an overdetermined system: within each auction-bidder pair, there are $T_n - 1$ equations and 2 unknowns. With measurement error, however, Equation (26) becomes a system of regression equations:

$$(b_{t,i,n}^d - Z_{t,i,n}^d(\hat{\theta}_1) \cdot s_{i,n}^d) = A_{t,i,n}(\hat{\theta}_1) \cdot \alpha_{i,n} + G_{t,i,n}(\hat{\theta}_1) \cdot \frac{1}{\gamma_{i,n}} + \tilde{\nu}_{t,i,n}, \quad (27)$$

where $b_{t,i,n}^d$ and $s_{i,n}^d$ are unit bids and scores, respectively, as they are observed in the data, $\hat{\theta}_1$ refers to the quantity model parameters estimated in the first stage and $\tilde{\nu}_{t,i,n}$ is an orthogonal mean-zero bid error. We formally define $\tilde{\nu}_{t,i,n}$ in Equation (31) in the next section.

If the number of items in each auction could be taken to infinity, the second-stage parameters $\alpha_{i,n}$ and $\gamma_{i,n}$ would be consistently estimated within each auction-bidder pair under a standard orthogonality condition $\mathbb{E}[\tilde{\nu}_{t,i,n} | A_{t,i,n}(\theta_1), G_{t,i,n}(\theta_1), Z_{t,i,n}^d(\theta_1)] = 0$ within each bidder-auction. However, as some auctions have relatively few items and the variance of bid errors may be large across items, we pool across auctions by projecting $\alpha_{i,n}$ and $\gamma_{i,n}$ on a vector of bidder-auction characteristics:

$$\alpha_{i,n} = \beta_{0,\alpha}^i + \beta_{X,\alpha} X_{i,n} + \nu_{\alpha,i,n} \quad \text{and} \quad \frac{1}{\gamma_{i,n}} = \beta_{0,\gamma}^i + \beta_{X,\gamma} X_{i,n} + \nu_{\gamma,i,n}. \quad (28)$$

This projection induces heteroscedasticity and requires a somewhat stronger assumption on the exogenous distribution of measurement errors: not only must they be uncorrelated with item characteristics within a single bidder-auction pair, but across bidders and auctions as well. However, we argue that under our measurement error interpretation, this is a reasonable extension of the same logic: bidders have the same propensity to round or misreport optimal unit bids in all auctions within our time frame. The coefficients β_α , β_γ may thus be consistently estimated under the augmented orthogonality condition, $\mathbb{E}[\tilde{\nu}_{t,i,n} | A_{t,i,n}(\theta_1), G_{t,i,n}(\theta_1), Z_{t,i,n}^d(\theta_1), X_{i,n}] = 0$ across bidders i and auctions n , that is guaranteed by our Assumption 1.

C.3.2.2. GMM Specification. To efficiently estimate our second stage, we implement a GMM procedure that applies the orthogonality condition above across subsamples of our data. Each moment corresponds to the weighted expectation of bid errors across a fixed slice of bidder-auction pairs in a sample of auction draws. In this sense, our GMM procedure may be thought of as a weighted OLS procedure building on Equation (27).

Denote $\theta_2 = (\beta_{0,\gamma}^1, \dots, \beta_{0,\gamma}^I, \beta_{X,\gamma}^1, \dots, \beta_{X,\gamma}^I, \beta_{0,\alpha}^1, \dots, \beta_{0,\alpha}^I, \beta_{X,\alpha}^1, \dots, \beta_{X,\alpha}^I)$ for the vector of second-stage parameters, where I is the number of unique firm IDs and J is the number of auction-bidder features.³¹ We estimate θ_2 in the second stage, using a GMM

³¹To simplify notation, we do not distinguish between “unique” bidders—for example, bidders who appear in 30+ auctions—and rare bidders—whom we group into a single unique bidder ID—for the purposes of this

framework evaluated at the first-stage estimates $\hat{\theta}_1$:

$$\hat{\theta}_2 = \arg \min_{\theta_2} \mathbb{E}_n [g(\theta_2, \hat{\theta}_1)' W g(\theta_2, \hat{\theta}_1)],$$

where $g(\theta_2, \hat{\theta}_1)$ is a vector of moments given a candidate θ_2 and W is a weighting matrix. We make use of the following four types of moments, asymptotic in the number of auctions N . The first type of moment states that the average residual of a unit bid submitted by each (unique) bidder i is zero across auctions. There are I such moments, where I is the number of unique bidders. The second type of moment states that the average residual on a unit bid submitted in each auction is zero, independent of the auction-specific characteristics of the bidder submitting the bid. There are J such moments—one for each of the auction-bidder characteristics. The third and fourth types of moments focus on items likely to be subject to high variance in risk exposure by interacting the bidder-level (type 1) and characteristic-level (type 2) unit bid residuals with an indicator for whether the item being bid on was labeled as a “top skew item” by the DOT.³² As there are both $2J + 2I$ second-stage parameters and moments, we use an identity matrix for W :

$$\begin{aligned} m_i^1(\theta_2 | \hat{\theta}_1) &= \mathbb{E}_n \left[\frac{1}{|\mathcal{I}(n)|} \sum_{t \in \mathcal{T}(n)} \tilde{v}_{t,i,n}(\theta_2 | \hat{\theta}_1) \cdot \mathbf{1}_{i \in \mathcal{I}(n)} \right], \\ m_j^2(\theta_2 | \hat{\theta}_1) &= \mathbb{E}_n \left[\frac{1}{|\mathcal{I}(n)| \cdot |\mathcal{T}(n)|} \sum_{i \in \mathcal{I}(n)} \sum_{t \in \mathcal{T}(n)} \tilde{v}_{t,i,n}(\theta_2 | \hat{\theta}_1) \cdot \mathbf{1}_{i \in \mathcal{I}(n)} \cdot X_{i,n}^j \right], \\ m_i^3(\theta_2 | \hat{\theta}_1) &= \mathbb{E}_n \left[\frac{1}{|\mathcal{T}_s(n)|} \sum_{t \in \mathcal{T}(n)} \tilde{v}_{t,i,n}(\theta_2 | \hat{\theta}_1) \cdot \mathbf{1}_{i \in \mathcal{I}(n)} \cdot \mathbf{1}_{t \in \mathcal{T}_s} \right], \\ m_j^4(\theta_2 | \hat{\theta}_1) &= \mathbb{E}_n \left[\frac{1}{|\mathcal{I}(n)| \cdot |\mathcal{T}_s(n)|} \sum_{i \in \mathcal{I}(n)} \sum_{t \in \mathcal{T}(n)} \tilde{v}_{t,i,n}(\theta_2 | \hat{\theta}_1) \cdot \mathbf{1}_{i \in \mathcal{I}(n)} \cdot \mathbf{1}_{t \in \mathcal{T}_s} \cdot X_{i,n}^j \right]. \end{aligned}$$

For each auction n , we denote $\mathcal{I}(n)$ as the set of bidders involved in n , $\mathcal{T}(n)$ as the set of items used in n , \mathcal{T}_s as the subset of all items that were labeled as “top skew items” by the DOT chief engineer’s office, and $\mathcal{T}_s(n)$ as the set of “top skew items” in auction n . All moments are formed with respect to the *demeaned* bid residual:

$$\tilde{v}_{t,i,n}(\theta_2 | \hat{\theta}_1) = b_{t,i,n}^d - \alpha_{i,n}(\theta_2) \left(c_{t,n} - \frac{q_{t,n}^e}{\hat{\sigma}_{t,n}^2 \sum_{p \in \mathcal{T}(n)} \left[\frac{(q_{p,n}^e)^2}{\hat{\sigma}_{p,n}^2} \right]} \left[\sum_{p \in \mathcal{T}(n)} c_{p,n} q_{p,n}^e \right] \right)$$

econometrics section. For the latter group, we treat all observations of rare bidders as observations of the same single bidder, who may enter a given auction more than once, with a different draw of auction-bidder characteristics but the same bidder fixed effect determining her efficiency type.

³²“Top skewed items” are items that were flagged by the DOT’s Engineering Office as being prone to having especially high or low bids. These items were cited as often incurring systemic fluctuations in ex post quantities. The list of these items largely corresponds to the most frequent strongly over/underbid items in our data.

$$\begin{aligned}
& - \frac{1}{\gamma_{i,n}(\theta_2)} \left(\frac{\hat{q}_{t,n}^b}{\hat{\sigma}_{t,n}^2} - \frac{q_{t,n}^e}{\hat{\sigma}_{t,n}^2 \sum_{p \in \mathcal{T}(n)} \left[\frac{(q_{p,n}^e)^2}{\hat{\sigma}_{p,n}^2} \right]} \left[\sum_{p \in \mathcal{T}(n)} \frac{\hat{q}_{p,n}^b q_{p,n}^e}{\hat{\sigma}_{p,n}^2} \right] \right) \\
& - \frac{q_{t,n}^e}{\hat{\sigma}_{t,n}^2 \sum_{p \in \mathcal{T}(n)} \left[\frac{(q_{p,n}^e)^2}{\hat{\sigma}_{p,n}^2} \right]} [s_{i,n}^d], \tag{29}
\end{aligned}$$

where

$$\alpha_{i,n}(\theta_2) = \beta_{0,\alpha}^i(\theta_2) + \beta_{X,\alpha}(\theta_2)X_{i,n} \quad \text{and} \quad \frac{1}{\gamma_{i,n}(\theta_2)} = \beta_{0,\gamma}^i(\theta_2) + \beta_{X,\gamma}(\theta_2)X_{i,n}. \tag{30}$$

The residual terms in the moments are *demeaned* in the sense that they use the *observed* score $s_{i,n}^d$ in the formulation of the optimal bid for (t, i, n) (rather than the *true* equilibrium score, $s_{i,n}^*$) and the projections of $\alpha_{i,n}$ and $\gamma_{i,n}$ (without the residuals $\nu_{\alpha,i,n}$ and $\nu_{\gamma,i,n}$). That is, the demeaned residual $\tilde{\nu}_{t,i,n}(\theta_2|\hat{\theta}_1)$ omits an unobserved score error term $\bar{\nu}_{s,i,n}$, along with projection error terms $\bar{\nu}_{\alpha,i,n}$ and $\bar{\nu}_{\gamma,i,n}$,

$$\tilde{\nu}_{t,i,n} = \nu_{t,i,n} + \bar{\nu}_{s,i,n} + \bar{\nu}_{\alpha,i,n} + \bar{\nu}_{\gamma,i,n}, \tag{31}$$

where

$$\begin{aligned}
\bar{\nu}_{s,i,n} &= \frac{q_{t,n}^e}{\hat{\sigma}_{t,n}^2 \sum_{p \in \mathcal{T}(n)} \left[\frac{(q_{p,n}^e)^2}{\hat{\sigma}_{p,n}^2} \right]} \sum_{t=1}^{T_n} \nu_{t,i,n} q_{t,n}^e, \\
\bar{\nu}_{\alpha,i,n} &= \nu_{\alpha,i,n} \left(c_{t,n} - \frac{q_{t,n}^e}{\hat{\sigma}_{t,n}^2 \sum_{p \in \mathcal{T}(n)} \left[\frac{(q_{p,n}^e)^2}{\hat{\sigma}_{p,n}^2} \right]} \left[\sum_{p \in \mathcal{T}(n)} c_{p,n} q_{p,n}^e \right] \right), \\
\bar{\nu}_{\gamma,i,n} &= \nu_{\gamma,i,n} \left(\frac{\hat{q}_{t,n}^b}{\hat{\sigma}_{t,n}^2} - \frac{q_{t,n}^e}{\hat{\sigma}_{t,n}^2 \sum_{p \in \mathcal{T}(n)} \left[\frac{(q_{p,n}^e)^2}{\hat{\sigma}_{p,n}^2} \right]} \left[\sum_{p \in \mathcal{T}(n)} \frac{\hat{q}_{p,n}^b q_{p,n}^e}{\hat{\sigma}_{p,n}^2} \right] \right). \tag{32}
\end{aligned}$$

The formulas in Equation (32) are derived by plugging the optimal score $s_{i,n}^*$, and the $\alpha_{i,n}$ and $\gamma_{i,n}$ parameters defined in Equation (28) into the optimal bid equation given by Equation (6), and moving the residual terms to the left-hand side. For instance, because each unit bid is observed with an error $\nu_{t,i,n}$, the total score $s_{i,n}^d$ is observed with a sum of

errors:

$$s_{i,n}^d = \sum_{t=1}^{T_n} (b_{t,i,n}^* + \nu_{t,i,n}) q_{t,n}^e = s_{i,n}^* + \sum_{t=1}^{T_n} \nu_{t,i,n} q_{t,n}^e.$$

Note, however, that the orthogonality of $\nu_{t,i,n}$ given Assumption 1 implies that $\mathbb{E}_n[\bar{\nu}_{s,i,n}]$, $\mathbb{E}_n[\bar{\nu}_{\alpha,i,n}]$ and $\mathbb{E}_n[\bar{\nu}_{\gamma,i,n}]$ asymptote to zero and are orthogonal to bidder and auction characteristics as well. Consequently, $\tilde{\nu}_{t,i,n}$ satisfies the necessary orthogonality constraints.

C.3.2.3. Estimation Procedure. To summarize, we estimate our parameters in a two-stage procedure. In the first stage, we estimate the parameters that model bidders' expectations over item quantities. In the second stage, we use an optimal GMM estimator to estimate the parameters governing bidder types:

1. Estimate $\hat{\theta}_1 = (\hat{\beta}_{0,q}, \hat{\beta}_q, \hat{\beta}_{0,\sigma}, \hat{\beta}_\sigma)$ and initialize θ_2
2. Solve

$$\hat{\theta}_2 = \min_{\theta_2} \left\{ \frac{1}{I} \sum_{i=1}^I (m_i^1(\theta_2 | \hat{\theta}_1)^2 + m_i^3(\theta_2 | \hat{\theta}_1)^2) + \frac{1}{J} \sum_{j=1}^J (m_j^2(\theta_2 | \hat{\theta}_1)^2 + m_j^4(\theta_2 | \hat{\theta}_1)^2) \right\},$$

where I is the set of unique firm IDs and J is the number of columns in $X_{i,n}$. This optimization problem is solved subject to the constraint that $\alpha_{i,n}(\theta_2)$ and $\gamma_{i,n}(\theta_2)$ be within a reasonable range for every i and n .³³

We calculate standard errors by a two-step bootstrap procedure. First, we take 100 draws from the posterior distribution of the quantity model parameters θ_1 in stage 1 of our estimation procedure.³⁴ Next, we draw 100 auctions at random with replacement from the total set of auctions in our sample, and repeat the step 2 optimization procedure for each combination of a sample from the θ_1 distribution, and a sample of auctions. The confidence intervals presented in Table D.II in Appendix D correspond to the 2.5th and 97.5th percentile of parameter estimates across the resulting 10,000 bootstrap draws.

C.3.3. Robustness to Unobserved Auction Heterogeneity

A large literature has considered the role of unobserved auction-level heterogeneity in the identification of bidders' values in timber and procurement auctions (e.g., Krasnokutskaya (2011), Athey, Levin, and Seira (2011), Roberts and Sweeting (2013)) and in first

³³This is a computationally efficient approach to impose the theoretical restriction that bidder costs are positive (so that bidders do not gain money from using materials). To calibrate reasonable boundary values for α and γ , we take two standard deviations above and below the unconstrained estimates of the parameters estimated under a simpler model with one γ for all bidders and no constraints. We find that this constraint does not bind for the vast majority of observations. One could alternatively impose this through an additional moment condition. However, this would add a substantial computational burden as indicators for nonnegativity are nondifferentiable functions. The results do not differ to an economically significant degree.

³⁴Our first-stage model is estimated using a Hamiltonian Monte Carlo procedure (Betancourt (2017)), which, like other Markov chain Monte Carlo methods, draws a sequence of samples that (after convergence) is distributed according to the "target" distribution (e.g., the distribution of parameters governing the expected quantities and variances for each item in each auction). The result of the procedure is a vector of "posterior draws" for each parameter in θ_1 , the distribution of which is summarized in Table D.I. Our main estimates $\hat{\theta}_1$ (e.g., those plugged into the second stage) are the means of these (post-convergence) posterior draws. To generate our bootstrap estimates, instead of taking the mean of the posterior draws for each parameter, we instead take 100 random samples, and plug each draw into each iteration of the second-stage estimator.

price auctions with risk-averse bidders (e.g., Guerre, Perrigne, and Vuong (2009), Hu, McAdams, and Shum (2013), Grundl and Zhu (2019), Luo (2020)).

To discuss the role of unobserved heterogeneity in our setting, we consider several ways in which unobserved auction-level shocks may enter into bidders' considerations. First, we consider an additive profit shock u^a : an additively-separable auction-level profit or cost that bidders anticipate at the time of bidding, but that is not bid upon and is not observable to the econometrician. We already consider one type of such shock in the form of the extra work order (EWO) payment ξ in Section 5. As with EWOs, auction-level shocks such as bonus payments or lump-sum costs for setting up work on particular project sites may affect the distribution of bidders willing to enter into each auction. However, as these shocks do not affect the portfolio optimization problem in Equation (5) once a score is chosen, our identification of bidder types is unbiased by them. Thus, while our second-stage estimation approach does not allow us to back out u^a , neither our estimates of bidder types nor their interpretation need change to account for them.³⁵

Second, we consider unobserved heterogeneity that affects profits multiplicatively, for instance, if bidders anticipate inflation that will devalue dollars at the time that payments are made. Suppressing the auction identifier n , and including both an additive shock u^a and a multiplicative shock $\frac{1}{u^m} > 0$, we can rewrite Equation (3) as

$$\begin{aligned} 1 - \mathbb{E}_{q^a} \left[\exp \left(-\frac{\gamma_i}{u^m} \left(u^a + \xi + \sum_{t=1}^T q_t^a \cdot (b_{i,t} - \alpha_i c_t) \right) \right) \right] \\ = 1 - \exp \left(-\frac{\gamma_i}{u^m} \left(u^a + \xi + \sum_{t=1}^T q_t^b (b_{i,t} - \alpha_i c_t) - \frac{\gamma_i \sigma_i^2}{2u^m} (b_{i,t} - \alpha_i c_t)^2 \right) \right). \end{aligned}$$

As in the previous case, extra work payments ξ and other additive shocks u^a do not affect the structure of the portfolio optimization problem in Equation (5), and so they do not pose a problem for identifying bidder types. However, the multiplicative shock $\frac{1}{u^m}$ affects the relative weight that bidders place on risk. In this case, $\frac{1}{u^m}$ affects optimal bidding and Equation (27) becomes

$$(b_{t,i,n}^d - Z_{t,i,n}^d(\hat{\theta}_1) \cdot s_{i,n}^d) = A_{t,i,n}(\hat{\theta}_1) \cdot \alpha_{i,n} + G_{t,i,n}(\hat{\theta}_1) \cdot \frac{u_n^m}{\gamma_{i,n}} + \tilde{v}_{t,i,n}. \quad (33)$$

As Equation (33) shows, our estimates of efficiency-types $\alpha_{i,n}$ remain unbiased, but the estimated risk-aversion parameters $\gamma_{i,n}$ are not separately identified from the auction shocks u_n^m without further restrictions. However, as our analysis relies on bidder-auction estimates of $\gamma_{i,n}$ (rather than a cross-auction bidder-level parameter, for instance), accounting for u_n^m may be treated as a reparametrization of the estimates for $\gamma_{i,n}$. That is, each estimate of $\gamma_{i,n}$ may be thought of as the ratio $\frac{\gamma_{i,n}}{u_n^m}$, which plays the same functional role

³⁵In Section 8, we calibrate a parametric model of entry costs and EWO payment expectations based on bidders' entry decisions under an IPV framework parametrized by our first- and second-stage estimates. If additional additive shocks u^a were relevant to bidders' decisions, then our estimates of ξ would incorporate them. However, our reduced-form model of ξ may be more likely to be misspecified in this case.

in counterfactuals.³⁶ If u_n^m is an inflation adjustment, this ratio may be thought of as the unit-adjusted CARA coefficient for bidder i in auction n .

Finally, we consider unobservable shocks that may affect costs or quantity estimates, but not bids. If bidders anticipate additional costs or quantity variance because, for instance, they observe that a project site is in especially bad shape, then the cost of each item t for a bidder i in auction n might better be represented as $(\alpha_{i,n}c_{t,n} + \bar{c}_n)$ or $\alpha_{i,n}(c_{t,n} \times \bar{c}_n)$ for some auction-level amount \bar{c}_n that is unobservable to the econometrician. Similarly, the variance of an object t might better be represented as $\sigma_{t,n}^2 + \bar{\sigma}_n^2$ or $\sigma_{t,n}^2 \times \bar{\sigma}_n^2$ for an unobservable factor $\bar{\sigma}_n^2$. In each of these cases, the portfolio optimization problem in Equation (5) would be misspecified and our estimates of $\alpha_{i,n}$ and $\gamma_{i,n}$ might both be biased. As such, our estimates are *not robust* to cost- or variance- specific unobservable heterogeneity of this sort. While it may be possible to account for such considerations with an additional set of parametric assumptions, we leave this for future work.

C.4. Entry Parameter Calibration Details

In this section, we describe the procedure by which we calibrate the entry cost κ and extra work order multiplier λ in each auction. At a high level, both κ and λ help explain the patterns of entry observed in our data. While there is substantial heterogeneity across projects, entry into auctions in our sample is generally quite high: the median auction has 9 potential bidders and 6 participating (e.g., entering) bidders. At the same time, many participating bidders have relatively high (e.g., inefficient and risk averse) types τ , and the profit margins implied by our estimates are often small.³⁷

Holding all else fixed, the entry cost κ explains why not all potential bidders enter into each auction. As κ is incurred upon participation—irrespective of winning or losing the auction—it does not affect bidding after entry decisions are realized. However, it raises a trade-off for the entry decision itself: only bidders whose expected utility of participating (the expected utility of winning multiplied by the probability of winning, integrated over all possible numbers of entrants) is higher than the certain utility cost of entry will participate. For each auction—with its costs and uncertainties, its distribution of potential bidder types, its EWO amount, and its λ and κ —the threshold bidder type is the type for whom this tradeoff is balanced. A higher entry cost κ implies that fewer types of bidders will find it profitable to participate, and predicts a lower entry rate. Conversely, a lower entry cost κ predicts a higher entry rate.

Unlike the entry cost, EWO earnings (scaled by λ) are only earned if a bidder wins the auction. As such, λ impacts both the probability of entry and the choice of equilibrium score (and hence, optimal portfolio bidding) upon entry. Holding all else fixed, a higher λ reduces the break-even point for potential threshold bidders, and rationalizes entry by higher τ types. A higher λ may also rationalize lower equilibrium mark-ups based on item bids, as bidders account for EWO earnings when considering the expected utility of winning.³⁸

³⁶Note that our parametrization of the relationship between $\alpha_{i,n}$ and $\gamma_{i,n}$ within each auction accounts for a multiplicative auction-level fixed effect. For instance, the positive correlation between the estimates of $\alpha_{i,n}$ and $\gamma_{i,n}$ in Figure 5(b) is plotted *after* subtracting the auction-level mean of $\log(\gamma_{i,n})$ within each auction (then, adding the cross-auction mean of $\log(\gamma_{i,n})$ and exponentiating); see Section 7 for a full description.

³⁷See Table I for a summary of the number of realized bidders and the magnitude of extra work orders. To get a sense of the profit margins for participating bidders, see Table IV for the distribution of ex post markups (without accounting for EWOs or the cost of uncertainty).

³⁸Note that although $\lambda \cdot \text{EWO}$ is assumed to be homogeneous across bidders in a given auction, the additional profits that this term adds to winning are not fully competed away under CARA utility.

To calibrate κ and λ , we compare the theoretical predictions for the entry probability and threshold-type quantile in each auction against their empirical analogs in our data. To generate predictions under each choice of κ and λ , we simulate the equilibrium entry game detailed in Appendix A for each auction. The resulting predictions are a function of not only κ and λ , but also the distribution of potential bidder types, the item quantities, uncertainties, and market rates, and the EWO amount in each auction.

To construct groups of potential bidders, we classify all auctions with the same project type, year and geographic region (a binary split of the 6 districts defined by MassDOT) into a distinct bin $\mathcal{B}(n)$. Each bin $\mathcal{B}(n)$ represents a set of comparable auctions, with similar qualifications and bidder availability. We define the number of potential bidders in each auction as the maximum number of bidders seen in any auction within the same bin.³⁹ In addition, we assume that κ and λ are homogeneous within each bin, reflecting the idea that auctions within the same bin involve similar costs for preparing a bid, as well as similar levels of uncertainty over the magnitude of the EWO earnings that will be realized. Finally, we define the empirical frequency of entry \hat{q}_n for each auction n as the ratio of the number of bidders who participated in auction n to the number of potential bidders in the bin $\mathcal{B}(n)$.

Our calibration procedure applies a grid search across possible values of κ and λ . For each pair of parameters along the grid and auction n , we first find the threshold-type $\hat{\tau}_n^*(\kappa, \lambda)$ that obtains zero expected utility of entry under the empirical entry frequency \hat{q}_n :

$$V_n(\hat{\tau}_n^*(\kappa, \lambda)) = \sum_{m=1}^M \binom{M-1}{m-1} \hat{q}_n^{m-1} (1 - \hat{q}_n)^{M-m} \cdot \text{EU}_n(\varphi^*(\hat{\tau}_n^*(\kappa, \lambda)) | \kappa, \lambda, m) = 0. \quad (34)$$

We then compute two statistics of the entry model: the predicted entry rate $q_n^p(\kappa, \lambda)$, and the threshold-type quantile $F_n(\hat{\tau}_n^*(\kappa, \lambda))$. To obtain $F_n(\hat{\tau}_n^*(\kappa, \lambda))$, we evaluate the CDF of the distribution of potential types in auction n at $\hat{\tau}_n^*(\kappa, \lambda)$. In equilibrium, only bidders whose types are below $\hat{\tau}_n^*(\kappa, \lambda)$ enter the auction, and so $F_n(\hat{\tau}_n^*(\kappa, \lambda))$ yields the probability of entry in Equation (34). Comparing $F_n(\hat{\tau}_n^*(\kappa, \lambda))$ to the empirical frequency \hat{q}_n thus provides an ex ante measure-of-fit between the model prediction at λ and κ and the empirical entry rate.

To obtain $q_n^p(\kappa, \lambda)$, we compute the equilibrium of each auction n given κ , λ , and $\tau_n^*(\kappa, \lambda)$ as described in Appendix A. We then simulate 1000 draws of τ according to the distribution of potential types in n . For each τ draw, we compute the value of participating in the auction, $V_n(\tau)$, according to Equation (20), and subsequently, compute the share of draws in which $V_n(\tau)$ is positive, so that bidders of type τ would choose to enter. In equilibrium, the share of entries equals the empirical entry frequency \hat{q}_n as well. Comparing $q_n^p(\kappa, \lambda)$ to \hat{q}_n thus provides an ex post measure-of-fit for the model prediction at λ and κ .

Finally, since κ and λ vary at the bin-level, we find the best fit by selecting the $\hat{\kappa}_{\mathcal{B}(n)}$ and $\hat{\lambda}_{\mathcal{B}(n)}$ that minimize the the sum of squared deviations of each comparison within each bin:

$$\hat{\kappa}_{\mathcal{B}(n)}, \hat{\lambda}_{\mathcal{B}(n)} = \arg \min_{\kappa, \lambda} \left\{ \sum_{n \in \mathcal{B}(n)} (\hat{q}_n - q_n^p(\kappa, \lambda))^2 + (\hat{q}_n - F_n(\hat{\tau}_n^*(\kappa, \lambda)))^2 \right\}. \quad (35)$$

³⁹An alternative method would be to consider each unique bidder ever seen within a bin as a potential bidder. As there are many small bidders who participate, this can generate quite a large number of potential bidders, yielding very small entry probabilities. While we prefer our specification and find it more realistic, this alternative is feasible within our framework as well.

APPENDIX D: ESTIMATION RESULTS TABLES

First-Stage Parameter Estimates.

TABLE D.I

FIRST-STAGE PARAMETER ESTIMATES. \hat{R} VALUES CORRESPOND TO A CONVERGENCE DIAGNOSTIC FOR MARKOV CHAIN MONTE CARLO. SEE <https://mc-stan.org/posterior/reference/rhat.html> FOR DETAILS.

Parameter	\hat{R}	# Effective Samples	Mean	SD	2.5%	50%	97.5%
$\beta_q[1]$	1.00	8643	0.82	0.00	0.82	0.82	0.83
$\beta_q[2]$	1.00	4404	-0.02	0.00	-0.03	-0.02	-0.01
$\beta_q[3]$	1.00	5066	-0.01	0.00	-0.02	-0.01	-0.01
$\beta_q[4]$	1.00	6002	-0.03	0.00	-0.04	-0.03	-0.02
$\beta_q[5]$	1.00	6477	0.01	0.00	0.01	0.01	0.02
$\beta_q[6]$	1.00	5454	-0.02	0.00	-0.03	-0.02	-0.01
$\beta_q[7]$	1.00	5552	0.01	0.00	0.00	0.01	0.02
$\beta_q[8]$	1.00	5772	0.01	0.00	0.00	0.01	0.02
$\beta_q[9]$	1.00	3502	-0.03	0.00	-0.04	-0.03	-0.02
$\beta_q[10]$	1.00	4293	-0.03	0.00	-0.03	-0.03	-0.02
$\beta_q[11]$	1.00	3160	-0.02	0.00	-0.03	-0.02	-0.01
$\beta_q[12]$	1.00	3383	0.01	0.00	-0.00	0.01	0.01
$\beta_q[13]$	1.00	3885	-0.00	0.00	-0.01	-0.00	0.00
$\beta_q[14]$	1.00	4879	0.01	0.00	-0.00	0.01	0.01
$\beta_q[15]$	1.00	3216	0.03	0.00	0.02	0.03	0.03
$\beta_q[16]$	1.00	7501	0.01	0.00	0.00	0.01	0.02
$\beta_q[17]$	1.00	4048	0.01	0.00	0.01	0.01	0.02
$\beta_q[18]$	1.00	6995	-0.18	0.00	-0.19	-0.18	-0.17
$\beta_q[19]$	1.00	6760	-0.01	0.00	-0.02	-0.01	-0.00
$\beta_\sigma[1]$	1.00	7025	-0.67	0.00	-0.67	-0.67	-0.66
$\beta_\sigma[2]$	1.00	1975	-0.05	0.01	-0.06	-0.05	-0.04
$\beta_\sigma[3]$	1.00	2931	0.02	0.00	0.01	0.02	0.03
$\beta_\sigma[4]$	1.00	4243	-0.02	0.00	-0.03	-0.02	-0.01
$\beta_\sigma[5]$	1.00	4284	0.00	0.00	-0.01	0.00	0.01
$\beta_\sigma[6]$	1.00	4056	0.02	0.00	0.01	0.02	0.03
$\beta_\sigma[7]$	1.00	3849	0.08	0.01	0.07	0.08	0.09
$\beta_\sigma[8]$	1.00	2301	0.03	0.01	0.02	0.03	0.04
$\beta_\sigma[9]$	1.00	1736	0.00	0.01	-0.01	0.00	0.01
$\beta_\sigma[10]$	1.00	1813	-0.01	0.01	-0.02	-0.01	0.00
$\beta_\sigma[11]$	1.00	1421	0.03	0.01	0.02	0.03	0.05
$\beta_\sigma[12]$	1.00	2158	-0.03	0.01	-0.04	-0.03	-0.02
$\beta_\sigma[13]$	1.00	2134	0.02	0.01	0.01	0.02	0.03
$\beta_\sigma[14]$	1.00	2789	0.04	0.01	0.03	0.04	0.05
$\beta_\sigma[15]$	1.00	2182	0.02	0.01	0.01	0.02	0.03
$\beta_\sigma[16]$	1.00	3493	0.00	0.00	-0.01	0.00	0.01
$\beta_\sigma[17]$	1.00	2109	-0.16	0.01	-0.18	-0.16	-0.15
$\beta_\sigma[18]$	1.00	5823	0.07	0.00	0.06	0.07	0.08
$\beta_\sigma[19]$	1.00	6423	0.02	0.00	0.02	0.02	0.03


Second-Stage Parameter Estimates. We obtain standard errors and confidence bounds through a two-step bootstrapping procedure. We first take 100 draws from the posterior distribution of the first-stage model. Then for each first-stage draw, we perform 100 bootstrap iterations of the second-stage estimation procedure. In each iteration, we re-draw 438 auctions at random with replacement and reestimate our second-stage GMM

model.⁴⁰ Table D.II presents the resulting 95% confidence interval for the headline estimates in Tables III and IV, as well as their standard deviations both across all draws, and within the 95% confidence interval.

TABLE D.II
SECOND-STAGE BOOTSTRAP ERRORS AND QUANTILES.

Parameter	Estimate	SD	SD Within CI	2.5%	97.5%
Mean α	1.033	0.036	0.031	0.958	1.100
25% α	0.953	0.057	0.047	0.793	1.028
50% α	1.053	0.044	0.038	0.965	1.132
75% α	1.175	0.042	0.036	1.106	1.275
Mean γ	0.088	0.021	0.016	0.065	0.142
25% γ	0.042	0.008	0.007	0.026	0.059
50% γ	0.061	0.016	0.013	0.041	0.102
75% γ	0.096	0.042	0.029	0.066	0.219
Mean Markup	0.209	0.048	0.041	0.116	0.317
25% Markup	-0.087	0.031	0.027	-0.146	-0.022
50% Markup	0.099	0.044	0.038	0.018	0.199
75% Markup	0.355	0.069	0.056	0.238	0.519

Model Fit Figures.



(a) Binscatter of actual quantities vs. predictions from the first-stage model fit.

<i>Dependent variable:</i>	
Actual Quantity	
Predicted Quantity	0.812 (0.005)
Constant	0.291 (0.015)
Observations	29,834
R ²	0.476

(b) Regression report for Figure D.1a.

FIGURE D.1.—First-stage model fit.

⁴⁰We exclude two auctions with outlying high costs from the second stage of estimation, leaving 438 auctions.




FIGURE D.2.—Second-stage model fit.

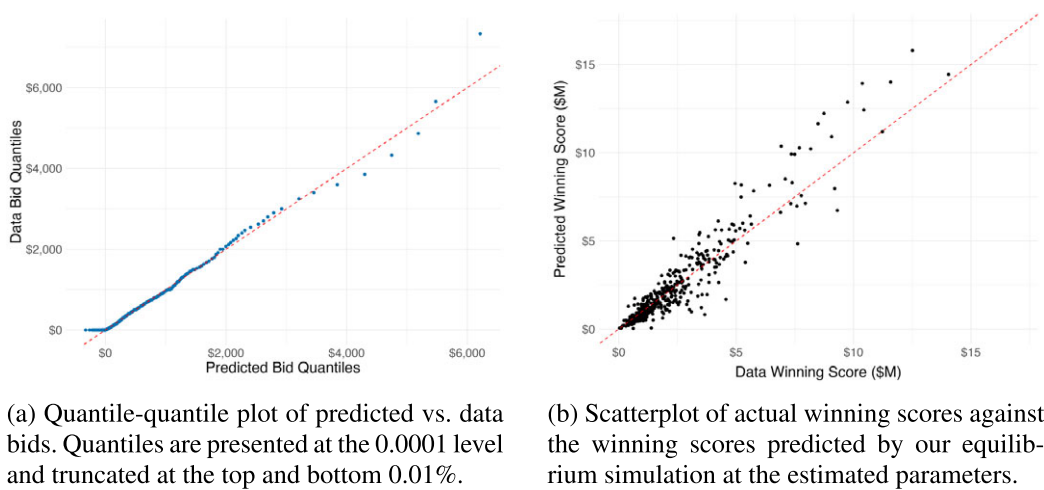
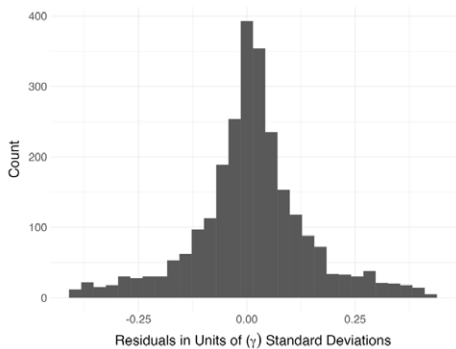


FIGURE D.3.—Fit plots for bids and scores with the 45-degree line, dashed in red, for reference.




(a) Histogram of residuals from the Poisson regression model discussed in [Appendix A](#) in units of γ standard deviations, truncated at the top and bottom 5% for visibility.

	<i>Dependent variable:</i>
	$\gamma_{i,n}$
Predicted $\gamma_n(\alpha_{i,n})$	0.999 (0.008)
Constant	0.0001 (0.001)
Observations	2,867
R ²	0.855


(b) Regression report for the prediction fit of the Poisson regression model for γ discussed in [Appendix A](#).

FIGURE D.4.—Fit plots for the Poisson regression model for γ discussed in [Appendix A](#).

APPENDIX E: ADDITIONAL FIGURES

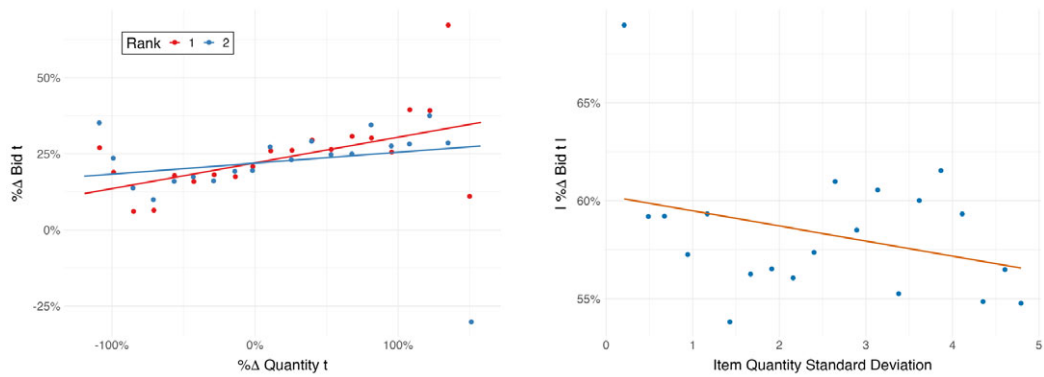


(a) Replication of [Figure 3a](#) with bidders ranked 1–4.



(b) Replication of [Figure 3b](#) with bidders ranked 1–4.

FIGURE E.1.—Replication of [Figure 3](#) with bidders 1–4.



(a) Replication of Figure 3a when the top two bidders' scores are within 10% of each other. (b) Replication of Figure 4a, without controlling for $\% \Delta q_t$.

FIGURE E.2.—Replication of Figures 3(a) and 4(a) with alternative specifications.

APPENDIX F: COUNTERFACTUAL RESULTS TABLES

We report the summary statistics for the counterfactual results reported in Section 8.⁴¹

TABLE FI
SUMMARY STATISTICS FOR CF THRESHOLD-TYPE CHANGES UNDER ENDOGENOUS ENTRY.

CF Type	% Change in Threshold	Mean	SD	25%	50%	75%
Lump Sum	Cost Efficiency (α)	−19.48	15.00	−30.65	−20.38	−7.52
1:2 Renegotiation	Cost Efficiency (α)	−6.08	9.14	−10.06	−0.09	0
2:1 Renegotiation	Cost Efficiency (α)	−0.93	3.05	0	0	0
50–50 Renegotiation	Cost Efficiency (α)	−2.41	5.39	−0.80	0	0
No Risk	Cost Efficiency (α)	−0.13	3.33	0	0	0
Lump Sum	Risk Aversion (γ)	−26.17	19.62	−41.24	−28.19	−10.74
1:2 Renegotiation	Risk Aversion (γ)	−8.42	12.44	−14.28	−0.13	0
2:1 Renegotiation	Risk Aversion (γ)	−1.32	4.26	0	0	0
50–50 Renegotiation	Risk Aversion (γ)	−3.38	7.48	−1.16	0	0
No Risk	Risk Aversion (γ)	−0.15	4.87	0	0	0

⁴¹Each outcome is winsorized by 1% to exclude extreme outliers from the mean/SD calculations. These results exclude a small number of projects for which the ODE solvers did not converge without special tuning.

TABLE F.II
SUMMARY STATISTICS FOR CF OUTCOMES WITH AND WITHOUT ENDOGENOUS ENTRY.

CF Type	Outcome	Mean	SD	25%	50%	75%
<i>With Endogenous Entry</i>						
Lump Sum	% DOT Savings	-97.8	190.8	-102.8	-42.2	-17.7
Lump Sum w 2:1	% DOT Savings	-17.3	28.6	-28.6	-13.8	-2.9
Negotiation						
Lump Sum w 50/50	% DOT Savings	-13.7	24.1	-23.5	-8.5	-0.2
Negotiation						
No Risk (Correct q)	% DOT Savings	8.0	36.3	0.7	14.5	25.4
No Risk (Estimated q)	% DOT Savings	-17.5	50.4	-17.0	-1.9	3.3
Lump Sum	\$ DOT Savings	-1,127,957.1	2,780,230.0	-734,758.8	-302,299.0	-97,254.3
Lump Sum w 2:1	\$ DOT Savings	-321,549.6	567,099.9	-402,865.9	-124,194.6	-16,320.3
Negotiation						
Lump Sum w 50/50	\$ DOT Savings	-224,092.1	431,945.0	-261,757.2	-92,431.1	-2481.1
Negotiation						
No Risk (Correct q)	\$ DOT Savings	138,945.4	599,836.5	6510.8	145,919.7	339,107.5
No Risk (Estimated q)	\$ DOT Savings	-190,488.4	514,707.8	-162,948.3	-18,782.2	30,373.2
Lump Sum	% Bidder Gain	114.9	144.9	27.9	75.8	168.7
Lump Sum w 2:1	% Bidder Gain	166.2	387.2	20.5	48.4	149.9
Negotiation						
Lump Sum w 50/50	% Bidder Gain	147.7	283.0	11.9	31.6	135.5
Negotiation						
No Risk (Correct q)	% Bidder Gain	166.2	497.9	-32.6	-8.9	100.2
No Risk (Estimated q)	% Bidder Gain	-305.2	762.4	-207.7	-13.5	2.0
Lump Sum	\$ Bidder Gain	63,641.1	96,905.4	9719.6	37,593.9	87,450.3
Lump Sum w 2:1	\$ Bidder Gain	105,305.1	218,136.7	9911.0	25,754.0	84,472.3
Negotiation						
Lump Sum w 50/50	\$ Bidder Gain	83,017.4	153,886.3	6332.6	17,229.4	92,870.8
Negotiation						
No Risk (Correct q)	\$ Bidder Gain	116,567.3	470,277.3	-24,138.8	-4951.4	87,635.4
No Risk (Estimated q)	\$ Bidder Gain	-188,641.5	462,137.2	-131,162.9	-6954.3	1088.3
<i>Holding Entry</i>						
<i>Probabilities Fixed</i>						
Lump Sum	% DOT Savings	-221.6	379.7	-226.1	-95.8	-44.1
Lump Sum w 2:1	% DOT Savings	-19.0	28.6	-33.0	-14.6	-4.6
Negotiation						
Lump Sum w 50/50	% DOT Savings	-8.8	17.6	-16.2	-6.4	-0.8
Negotiation						
No Risk (Correct q)	% DOT Savings	15.1	22.1	7.2	15.2	23.3
No Risk (Estimated q)	% DOT Savings	-1.3	29.4	-2.3	0.4	4.1
Lump Sum	\$ DOT Savings	-2,837,833.9	6,080,259.4	-1,967,597.4	-709,927.9	-245,497.9
Lump Sum w 2:1	\$ DOT Savings	-406,632.3	724,732.9	-471,577.3	-135,711.1	-30,123.9
Negotiation						
Lump Sum w 50/50	\$ DOT Savings	-192,411.7	361,492.8	-229,326.5	-60,254.4	-5489.5
Negotiation						
No Risk (Correct q)	\$ DOT Savings	193,577.9	255,820.1	81,278.6	161,773.9	287,398.3
No Risk (Estimated q)	\$ DOT Savings	-14,278.2	205,900.0	-28,301.2	3528.1	36,913.2
Lump Sum	% Bidder Gain	59.9	38.4	36.8	55.4	81.5
Lump Sum w 2:1	% Bidder Gain	40.3	35.8	20.4	34.5	54.4
Negotiation						

(Continues)

TABLE F.II
Continued.

CF Type	Outcome	Mean	SD	25%	50%	75%
Lump Sum w 50/50 Negotiation	% Bidder Gain	26.3	39.1	9.7	20.7	33.0
No Risk (Correct q)	% Bidder Gain	37.9	358.3	−24.8	−14.3	−6.7
No Risk (Estimated q)	% Bidder Gain	−47.0	327.7	−2.8	0.6	6.7
Lump Sum	\$ Bidder Gain	30,738.1	26,910.7	14,099.2	27,557.2	42,744.8
Lump Sum w 2:1 Negotiation	\$ Bidder Gain	21,294.9	25,479.0	8492.3	17,776.4	31,684.2
Lump Sum w 50/50 Negotiation	\$ Bidder Gain	12,895.1	23,537.1	4488.2	9993.6	18,917.0
No Risk (Correct q)	\$ Bidder Gain	9333.1	134,422.0	−17,502.0	−9346.5	−3353.8
No Risk (Estimated q)	\$ Bidder Gain	−16,629.8	134,947.4	−1355.1	332.4	3264.4

REFERENCES

ASKER, JOHN, AND ESTELLE CANTILLON (2008): “Properties of Scoring Auctions,” *RAND Journal of Economics*, 39 (1), 69–85. [1206,1210]

ATHEY, SUSAN, AND JONATHAN LEVIN (2001): “Information and Competition in US Forest Service Timber Auctions,” *Journal of Political Economy*, 109 (2), 375–417. [1206,1208,1209,1215,1217]

ATHEY, SUSAN, JONATHAN LEVIN, AND ENRIQUE SEIRA (2011): “Comparing Open and Sealed Bid Auctions: Evidence From Timber Auctions,” *The Quarterly Journal of Economics*, 126 (1), 207–257. [1228,1248]

BAJARI, PATRICK, STEPHANIE HOUGHTON, AND STEVEN TADELIS (2014): “Bidding for Incomplete Contracts: An Empirical Analysis,” *American Economic Review*, 104 (32), 1288–1319. [1208,1209,1215,1242]

BETANCOURT, MICHAEL (2017): “A Conceptual Introduction to Hamiltonian Monte Carlo,” arXiv preprint arXiv:1701.02434. [1244,1248]

BOLOTNYY, VALENTIN, AND SHOSHANA VASSERMAN (2023): “Supplement to ‘Scaling Auctions as Insurance: A Case Study in Infrastructure Procurement’,” *Econometrica Supplemental Material*, 91, <https://doi.org/10.3982/ECTA17673>. [1211]

CAMPBELL, JOHN Y. (2017): *Financial Decisions and Markets: A Course in Asset Pricing*. Princeton University Press. [1221]

CAMPO, SANDRA (2012): “Risk Aversion and Asymmetry in Procurement Auctions: Identification, Estimation and Application to Construction Procurements,” *Journal of Econometrics*, 168 (1), 96–107. [1209]

CAMPO, SANDRA, EMMANUEL GUERRE, ISABELLE PERRIGNE, AND QUANG VUONG (2011): “Semiparametric Estimation of First-Price Auctions With Risk-Averse Bidders,” *The Review of Economic Studies*, 78 (1), 112–147. [1209]

CARPENTER, BOB, ANDREW GELMAN, MATTHEW D. HOFFMAN, DANIEL LEE, BEN GOODRICH, MICHAEL BETANCOURT, MARCUS BRUBAKER, JIQIANG GUO, PETER LI, AND ALLEN RIDDELL (2017): “Stan: A Probabilistic Programming Language,” *Journal of Statistical Software*, 76 (1). [1244]

CHE, YEON-KOO (1993): “Design Competition Through Multidimensional Auctions Author,” *RAND Journal of Economics*, 24 (4), 668–680. [1210]

GRUNDL, SERAFIN, AND YU ZHU (2019): “Identification and Estimation of Risk Aversion in First-Price Auctions With Unobserved Auction Heterogeneity,” *Journal of Econometrics*, 210 (2), 363–378. [1249]

GUERRE, EMMANUEL, ISABELLE PERRIGNE, AND QUANG VUONG (2009): “Nonparametric Identification of Risk Aversion in First-Price Auctions Under Exclusion Restrictions,” *Econometrica*, 77 (4), 1193–1227. [1209,1249]

HU, YINGYAO, DAVID MCADAMS, AND MATTHEW SHUM (2013): “Identification of First-Price Auctions With Non-Separable Unobserved Heterogeneity,” *Journal of Econometrics*, 174 (2), 186–193. [1249]

HUBBARD, TIMOTHY P., AND HARRY J. PAARSCH (2014): “On the Numerical Solution of Equilibria in Auction Models With Asymmetries Within the Private-Values Paradigm,” in *Handbook of Computational Economics*, Vol. 3. Elsevier, 37–115. [1236]

KONG, YUNMI (2020): “Not Knowing the Competition: Evidence and Implications for Auction Design,” *The RAND Journal of Economics*, 51 (3), 840–867. [1228]

KRASNOKUTSKAYA, ELENA (2011): “Identification and Estimation of Auction Models With Unobserved Heterogeneity,” *Review of Economic Studies*, 78 (1), 293–327. [1248]

- LAFFONT, JEAN-JACQUES, AND JEAN TIROLE (1993): *A Theory of Incentives in Procurement and Regulation*. MIT press. [1222]
- LOVALLO, DAN, TIM KOLLER, ROBERT UHLANER, AND DANIEL KAHNEMAN (2020): “Your Company Is Too Risk-Averse; Here’s Why and What to Do About It,” *Harvard Business Review*, 98 (2), 104–111. [1225]
- LUO, YAO (2020): “Unobserved Heterogeneity in Auctions Under Restricted Stochastic Dominance,” *Journal of Econometrics*, 216 (2), 354–374. [1249]
- LUO, YAO, AND HIDENORI TAKAHASHI (2023): “Bidding for Contracts Under Uncertain Demand: Skewed Bidding and Risk Sharing,” *RAND Journal of Economics*. [1210]
- MASKIN, ERIC, AND JOHN RILEY (1984): “Optimal Auctions With Risk Averse Buyers,” *Econometrica: Journal of the Econometric Society*, 1473–1518. [1210,1234]
- MATTHEWS, STEVEN (1987): “Comparing Auctions for Risk Averse Buyers: A Buyer’s Point of View,” *Econometrica: Journal of the Econometric Society*, 633–646. [1210,1234]
- MILGROM, PAUL R., AND ROBERT J. WEBER (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica: Journal of the Econometric Society*, 1089–1122. [1222]
- RACKAUCKAS, CHRISTOPHER, AND QING NIE (2017): “DifferentialEquations.jl—a Performant and Feature-Rich Ecosystem for Solving Differential Equations in Julia,” *Journal of Open Research Software*, 5, 1. [1236]
- RENY, PHILIP J. (2011): “On the Existence of Monotone Pure-Strategy Equilibria in Bayesian Games,” *Econometrica*, 79 (2), 499–553. [1221,1240]
- ROBERTS, JAMES W., AND ANDREW SWEETING (2013): “When Should Sellers Use Auctions?” *American Economic Review*, 103 (5), 1830–1861. [1228,1248]
- SAMUELSON, WILLIAM F. (1985): “Competitive Bidding With Entry Costs,” *Economics Letters*, 17 (1), 53–57. [1207,1228]

Co-editors Aviv Nevo and Kate Ho handled this manuscript.

Manuscript received 17 September, 2019; final version accepted 13 March, 2023; available online 17 March, 2023.