Scaling Auctions as Insurance: A Case Study in Infrastructure Procurement

Valentin Bolotnyy *   Shoshana Vasserman †

March 2023

Abstract

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.

*Hoover Institution, Stanford University. Email: vbolotnyy@stanford.edu

†Stanford Graduate School of Business and NBER. Email: svass@stanford.edu. We are grateful to the editor and anonymous referees for their many thoughtful suggestions for making the paper better. 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, as well as to Steve Poftak and Jack Moran, Frank Kucharski, Michael McGrath, and Naresh Chetpelly, among other generous public servants at MassDOT. We are also very grateful to Ali Yurukoglu, Andrzej Skrzypacz, Paulo Somaini, Peter Reiss, Susan Athey and Zi Yang Kang for their many insightful comments and suggestions, as well as to Charles Margossian, Philip Greengard, and especially Cameron Pfiffer for the boost that their help gave to our computational capacity.

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 (MassDOT or "the DOT"). Along with 40 other states, MassDOT uses a *scaling auction*, where bidders submit unit price bids 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 (that is, 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 over-run 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

[1] According 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.

exploited such separability to succinctly characterize equilibrium play. Studying auctions with quasi-linear 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 over-ran—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 dataset 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 ex-ante estimates and ex-post realizations 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 over-run the DOT quantity estimates and low unit bids on items that tend to under-run. 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 over-run by the largest amount. By contrast, the vast majority of unit bids observed in our data are interior (that is, non-extremal), even though no significant penalty for penny bidding has ever been exercised. We show that while contractors bid higher on items predicted to over-run, 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 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 from 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: under-runs and over-runs 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) over-bids and (unit) over-runs in auction data could be interpreted as evidence that bidders are able to predict which components of their bids will over-run. 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 et al. (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 over-runs, the ultimate payment to the winning bidder would be the same. Bajari et al. (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 et al. (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 over-bids and over-runs 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), 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 out-competing 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 non-identification results detailed in Guerre et al. (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 pre-commit 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 (2019), a contemporary paper that studies infrastructure procurement by the Florida DOT. Like us, Luo and Takahashi (2019) considers risk averse bidders and compares scaling auctions against lump sum auctions. However, this paper follows a GPV/Campo-style approach for identification and reduces the project components that receive unit bids (which number 67 for the median auction in our dataset) into two aggregates—one aggregate certain component and one aggregate uncertain component—for estimation. As such, while Luo and Takahashi (2019) 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 pre-qualified by MassDOT. Pre-qualification 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 on-site MassDOT manager. The winning contractor is ultimately paid the sum of her unit price bid multiplied by the *actual* quantity of each item 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.

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

[3]See Section B in the Online Appendix for a detailed discussion.

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. 2,513 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 pre-qualified 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 under-runs, over-runs, and Extra Work Orders (EWOs). The full approval process of changes to the initial project design involves several layers of review. Under-runs and over-runs, 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 actual item quantities used and the estimated item quantities. 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 under-runs and over-runs in bridge construction and maintenance projects.

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 Over-Cost (DOT Quantities)	−\$286,245	\$2.12 million	−\$480,487	−\$119,950	\$167,933
Net Over-Cost (Ex-Post Quantities)	−\$26,990	\$1.36 million	−\$208,554	\$15,653	\$275,219
Percent Over-Cost (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

Table 1: Summary statistics for the 440 MassDOT bridge projects auctioned during 1998-2015

Table 1 provides summary statistics for the bridge projects in our dataset. 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 under "Net Over-Cost (DOT Quantities)" in Table 1.

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 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 Over-Cost (Ex-Post Quantities)" row of Table 1, and in percent of the final cost paid out in "Percent Over-Cost (Ex-Post Quantities)".[4] The median over-payment 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 over-payment across projects. As we will show in our counterfactual section, the distribution of over-payment corresponds to the potential savings from the elimination of risk.

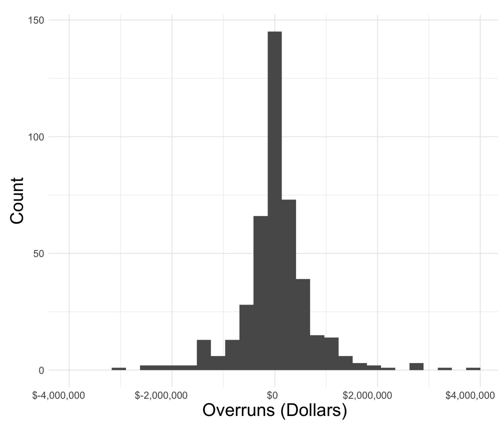

Figure 1: Net Over-Cost (Ex-Post Quantities) Across Bridge Projects

Bidder Characteristics There are 2,883 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-

[4]Note that the average "Percent Over-Cost (Ex-Post Quantities)" in Table 1 is the average percent of costs, rather than the ratio of the average net over-cost to average total cost.

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 purposes of estimation. Common firms constitute 2,263 (78%) of total bids submitted and 351 (80%) of auction victories.

	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 Over-run 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

Table 2: Comparison of firms participating in <30 vs. 30+ auctions

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 2 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 over-run 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 over-run: 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 dataset. 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 dataset includes records of 2,985 unique items, as per MassDOT's internal taxonomy.

Spread across 440 projects, these items constitute 29,834 unique item-project pairs. Of the 2,985 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, 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 engineers' 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 2a presents a histogram of the percent quantity over-run across item observations. The percent quantity over-run is defined as the difference of the ex-post quantity of an item observation and its DOT quantity estimates, normalized by the DOT estimate: $\frac{q_t^a - q_t^e}{q_t^e}$. In addition to the 23.3% of item-project observations in which quantity over-runs are 0%, another 18% involve items that are not used at all (so that the over-run is equal to -100%). The remaining over-runs 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 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 over-runs 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 of 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 over-runs often vary across observations of the same item in different auctions. Figure 2b plots the mean percent quantity over-run 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 over-run 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.

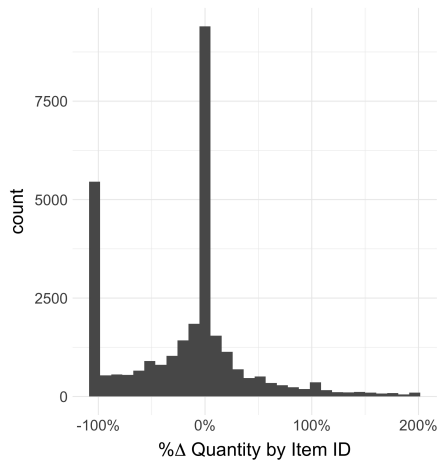

(a) Histogram of the percent quantity over-run across item-project pairs.

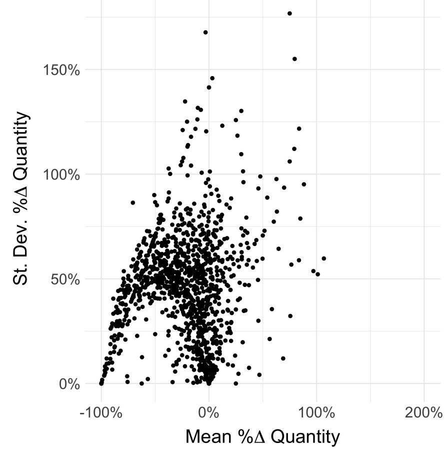

(b) Scatterplot of standard deviation vs. mean quantity over-run for each item.

Figure 2

Reduced Form Evidence for Risk Averse Bid Skewing As in Athey and Levin (2001) and Bajari et al. (2014), the bids in our data are consistent with a model of similarly informed bidders who bid strategically to maximize expected utility. In Figure 3a, we plot the relationship between quantity over-runs and the percent by which each item was over-bid above the blue book cost estimate. We do this for both the winning bidder and the second place bidder.[5] The binscatter is residualized. In order to obtain it, we first regress percent over-bid on a range of controls and obtain residuals. We then regress percent over-run 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.[6] We use a similar procedure for all residualized binscatters in this section.

As Figure 3a demonstrates, there is a significant positive relationship between percent quantity over-runs and percent over-bids by the winning bidder. A 1% increase in quantity

[5]The percent over-bid 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 over-run is similarly defined as $\frac{q_t^a - q_t^e}{q_t^e} \times 100$, where q_t^a is the amount of item t that was ultimately used and q_t^e is the DOT quantity estimate for item t that is used to calculate bidder scores.

[6]For each graph, we truncate observations at the top and bottom 1% to make the trends easier to see.

over-runs corresponds to a 0.086% increase in over-bids on average. Higher bids on over-running 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 over-run the DOT estimates on average, and to skew accordingly.

Furthermore, Figure 3a 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 over-bid on items that wound up over-running on average.[7] This suggests that the winning and second place bidder are similarly able to predict over-runs. In Appendix E, we show that this pattern holds for bidders ranked 3 and 4 as well, and that the average percent over-bids 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.

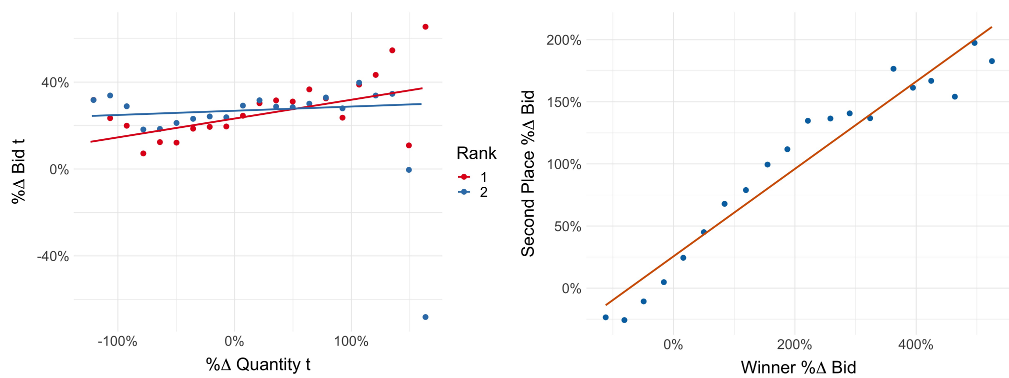

(a) Residualized binscatter of item-level percent over-bid against percent quantity over-run by the top 2 bidders.

(b) Binscatter of item-level percent over-bids by the rank 2 bidder against the rank 1 (winning) bidder.

Figure 3

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 "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 the 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 the 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

[7] As Figure 3a 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.

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 over-run 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 under-run.

This observation is consistent with the pattern demonstrated in Figure 3a. Here, the second place bidder—who submitted a higher overall score by definition—generally exhibits less severe skewing: a 1% increase in quantity over-runs corresponds to only a 0.019% increase in over-bids on average. Figure 3b, 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 over-bid by the winning bidder corresponds to a higher over-bid 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 4a and 4b, 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 4a, we plot a residualized binscatter of the winning bidder's absolute percent over-bid on each item against the item's standard deviation—the square root of the estimated prediction variance. This captures the uncertainty of each item 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 deviation—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

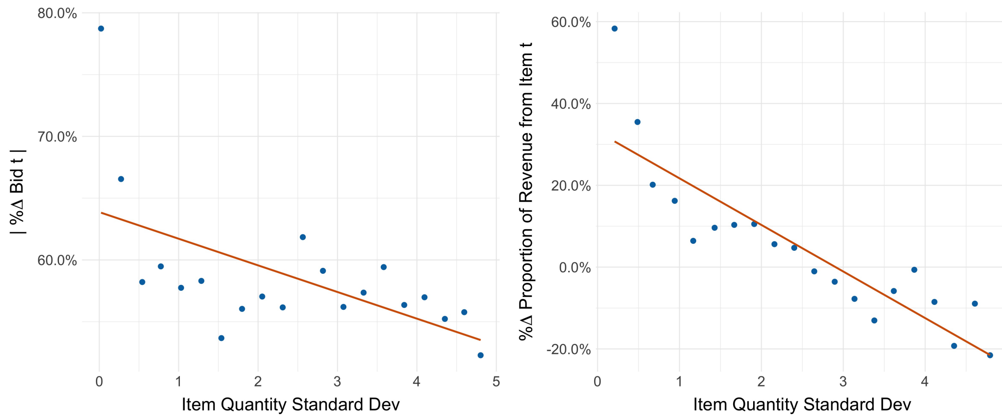

(a) Residualized binscatter of item-level percent absolute over-bid against the square root of estimated item quantity variance.

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

Figure 4

conditions.[8] Items with the highest standard deviations—such 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 4a demonstrates, the relationship between item-level absolute over-bids 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 over-runs 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 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 4b, we

[8]While 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.

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 independent private values 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:

- $\mathbf{q}^e = (q_1^e, \ldots, q_T^e)$: DOT estimates based on underlying conditions at the project site
- $\mathbf{q}^b = \left(q_1^b, \ldots, q_T^b\right)$: Common refined estimates based on public information
- $\mathbf{q}^a = (q_1^a, \ldots, q_T^a)$: 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:

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

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:

$$u_i(\pi) = 1 - \exp(-\gamma_i \pi). \tag{2}$$

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 $[\underline{\alpha}, \overline{\alpha}] \times [\underline{\gamma}, \overline{\gamma}]$ 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: $\sum_t^T q_t^a \cdot (b_{t,i} - \alpha_i c_t)$. 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—that does not depend on bids or quantities. Combining these components, bidder i's expected utility from participating in the auction is given by:

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

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:

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

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

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.[9] This maximization—which we call 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 multi-dimensional, 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

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

to the following two-stage problem when the support of feasible scores is sufficiently high.[10] 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:[11]

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

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:

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

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 by all

[10] See Appendix C.1 for details.

[11] See Campbell (2017) for a survey.

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 10b 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 over-use 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 extra work order 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 the 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 the bidders' efficiency types α and risk 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_t^a can be approximated by a normal 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. In particular, we model the realization of the actual quantity of item t in auction n as:

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

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

Here, $q_{t,n}^b$ is the posterior mean of $q_{t,n}^a$ and $\sigma_{t,n}$ is the square root of its posterior 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, denoted by $\widehat{\theta}_1$, as a point estimate for the second stage of estimation. We summarize the estimates of $\widehat{\theta}_1$ 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, a consistent correlation is in general difficult to identify and estimate.

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 dataset 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 $\widehat{q}_{t,n}^b$ and variances $\widehat{\sigma}_{t,n}^2$ 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 $s_{i,n}^*$, 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 $b_{t,i,n}^*(s_{i,n}^*|\widehat{\theta}_1, \alpha_{i,n}, \gamma_{i,n})$, subject to an IID mean-zero measurement error $\nu_{t,i,n}$:*

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

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 $\widehat{\theta}_1$, 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 $b^d_{t,i,n}$ 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 co-vary 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 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 $\widehat{q}^b_{t,n}$ and variance $\widehat{\sigma}^2_{t,n}$ terms in Table 7 in Appendix D. In addition, we demonstrate the model fit for the first stage in Figure 6 and Table 9. A histogram of the variance estimates $\widehat{\sigma}^2_{t,n}$ themselves is plotted in Figure 5a, below. 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 variance terms are between 0 and 3, with a trailing number of higher values.

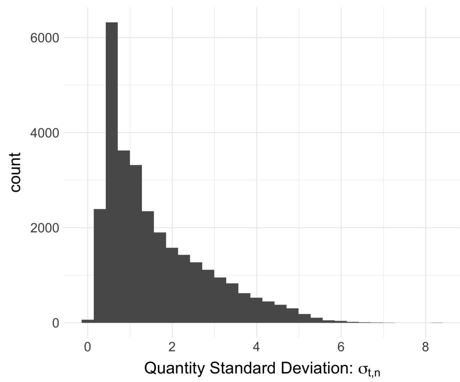

(a) A histogram of standard deviation estimates for each item t in each project n.

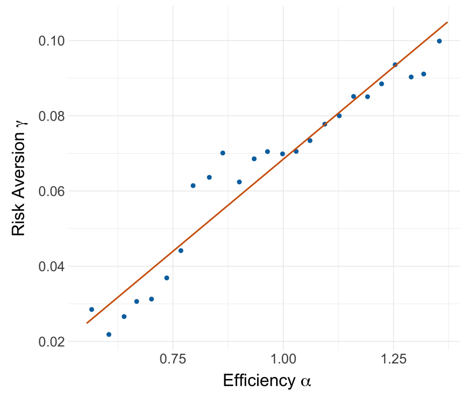

(b) Binscatter of estimated efficiency types $\alpha_{i,n}$ against risk aversion types $\gamma_{i,n}$.

Figure 5

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 3 and 4. Bootstrapped standard errors and confidence intervals are presented in Table 8 in Appendix D.

In Table 3, we present summary statistics of our estimates of bidder-auction risk aversion and efficiency types. The median coefficient of risk aversion $\widehat{\gamma}_{i,n}$ is estimated to be 0.061 when dollar values are scaled by \$1,000. 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 \$1,000 with indifference, and \$2,878 to accept a 50-50 lottery to win or lose \$10,000.[12] The median efficiency type $\widehat{\alpha}_{i,n}$ 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 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 \$2,665 to accept a lottery for \$10,000 (or \$28 to accept a 50-50 lottery for \$1,000), the median among structures maintenance projects would require a certain payment of \$3,444

[12]It is not unusual for managers to exhibit high levels of risk aversion at high stakes. For instance, a McKinsey survey of 1,500 managers found that most required a certainty equivalent of \$328M to accept a risk of losing an investment worth \$100M (Lovallo, Koller, Uhlaner, and Kahneman, 2020).

for the same lottery (or \$37 for a 50-50 lottery for \$1,000). 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 \$1,000 or \$2,041 (\$4,205) 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.

Project Type	Mean	SD	25%	Median	75%
			$\widehat{\gamma}_{i,n}$		
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
			$\widehat{\alpha}_{i,n}$		
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

Table 3: Summary statistics of $\gamma_{i,n}$ and $\alpha_{i,n}$ estimates by project type.

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

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

In Table 4, we summarize the distribution of ex-post markups implied by our estimates.

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 cost:

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

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 conditional on the equilibrium score that is observed in our data.

	Estimated Bidder Markups				
Num Bidders in Auction	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%

Table 4: Summary statistics of markups evaluated at $\widehat{\alpha}_{i,n}$ by number of participating bidders

We defer a demonstration of the goodness of fit of our structural model to Appendix D. Figure 7 presents a scatter plot and Figure 8a 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. Table 10 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 meta-type $\tau_{i,n}$ 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 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.[13] 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 meta-type $\tau_{i,n}$, 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 meta-types $\tau_{i,n}$ 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.

[13]Our definition of potential bidders is similar to that of Kong (2020). The maximum number of potential bidders is 20, while 5 auctions were grouped alone and excluded. See Appendix C.4 for details.

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 extra work orders that were needed.

Equilibrium Bidding Strategies Following the literature (Athey, Levin, and Seira (2011); Roberts and Sweeting (2013)), we construct a type-symmetric Bayes Nash 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 positive value $V_n(\tau_{i,n})$ from participating in the auction:

$$V_n(\tau_{i,n}) = \sum_{m=1}^{M} \begin{pmatrix} M-1 \\ m-1 \end{pmatrix} F_n(\tau_n^*)^{m-1} \left(1 - F_n(\tau_n^*)\right)^{M-m} \mathrm{EU}_n \left(s_n^*(\tau_{i,n}) | \lambda_n, \kappa_n, m\right), \qquad (9)$$

where $F_n(\tau_n^*)$ is the probability that an independent draw of a competing potential bidder's type is below τ_n^* and $\mathrm{EU}_n \left(s_n^*(\tau_{i,n}) | \lambda_n, \kappa_n, m\right)$ 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 for a given score s submitted by a bidder with type $\tau_{i,n}$ is given by:

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

where we suppress the arguments λ_n, κ_n, m in the $\mathrm{EU}_n(\cdot)$ operator because they are held fixed when computing the equilibrium score mapping in each auction instance. Here, $\mathrm{CE}_n(\mathbf{b}_{i,n}^*(s)|\tau_{i,n})$ 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 extra work orders. We assume that realizations of extra work order 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 extra work orders in our data is the final sum paid out to the winning bidder, EWO_n, we model bidders' certainty equivalent from extra work orders 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: $(1 - F^*_n(s^{*-1}_n(\tilde{s})))^{m-1}$.

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 $\text{CE}_n(\mathbf{b}^*_\tau(s^*_n(\tau))|\tau)$ rather than ex-post DOT costs. See Appendix A for more details on equilibrium construction, and Figure 8b 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 extra work orders 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 below in Table 5. 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,

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

Table 5: Summary statistics of calibrated estimates of EWO coefficients and entry costs

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.[14] In addition, we find that extra work order 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.

$$\overbrace{\sum_t \underbrace{(b_t q_t^b - \alpha c_t q_t^b)}_{\text{Expected Profits}} - \underbrace{\frac{\gamma \sigma_t^2}{2}(b_t - \alpha c_t)^2}_{\text{Risk Term}}}^{CE(\mathbf{b}|\alpha,\gamma)\,:\,\text{Scaling Auction}} \qquad \overbrace{\sum_t \underbrace{(b_t q_t^e - \alpha c_t q_t^b)}_{\text{Expected Profits}} - \underbrace{\frac{\gamma \sigma_t^2}{2}(\alpha c_t)^2}_{\text{Risk Term}}}^{CE(\mathbf{b}|\alpha,\gamma)\,:\,\text{Lump Sum Auction}} \tag{10}$$

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

[14] Based 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 office estimated score.

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 dataset. 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:

$$\overbrace{\sum_t \underbrace{b_t q_t^e + \left(\mu b_t (q_t^b - q_t^e)\right) - \alpha c_t q_t^b}_{\text{Expected Profits}} - \underbrace{\frac{\gamma \sigma_t^2}{2} \left(\mu b_t - \alpha c_t\right)^2}_{\text{Risk Term}}}^{CE(\mathbf{b},\alpha,\gamma)\,:\,\text{Lump Auction with Renegotiation}}. \tag{11}$$

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 pre-committing 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 over-use 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 over-runs 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 $\mathbf{q}^e = \mathbf{q}^a$. All bidders know that these quantities are correct and so they do not anticipate any further uncertainty: $\mathbf{q}^b = \mathbf{q}^a$ and $\boldsymbol{\sigma}^2 = \mathbf{0}$.

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 $\boldsymbol{\sigma}^2$ is unchanged. In this case, eliminating risk reduces the baseline variances $\boldsymbol{\sigma}^2$ 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 $\widehat{\mathbf{q}}^b$ 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 over-run the DOT's quantity estimates and low bids on items they predict will under-run. 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 over-run 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 allow MassDOT 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 multi-stage optimal mechanism established by Maskin and Riley (1984) and Matthews (1987), we leave this for future work.

References

Asker, J. and E. Cantillon (2008). Properties of scoring auctions. *RAND Journal of Economics 39*(1), 69–85.

Athey, S. and J. Levin (2001). Information and competition in US forest service timber auctions. *Journal of Political Economy 109*(2), 375–417.

Athey, S., J. Levin, and E. Seira (2011). Comparing open and sealed bid auctions: Evidence from timber auctions. *The Quarterly Journal of Economics 126*(1), 207–257.

Bajari, P., S. Houghton, and S. Tadelis (2014). Bidding for incomplete contracts: An empirical analysis. *American Economic Review 104*(April), 1288–1319(32).

Betancourt, M. (2017). A conceptual introduction to hamiltonian monte carlo. *arXiv preprint arXiv:1701.02434*.

Campbell, J. Y. (2017). *Financial decisions and markets: a course in asset pricing.* Princeton University Press.

Campo, S. (2012). Risk aversion and asymmetry in procurement auctions: Identification, estimation and application to construction procurements. *Journal of Econometrics 168*(1), 96–107.

Campo, S., E. Guerre, I. Perrigne, and Q. Vuong (2011). Semiparametric estimation of first-price auctions with risk-averse bidders. *The Review of Economic Studies 78*(1), 112–147.

Carpenter, B., A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell (2017). Stan: A probabilistic programming language. *Journal of Statistical Software 76*(1).

Che, Y.-k. (1993). Design Competition Through Multidimensional Auctions Author. *RAND Journal of Economics 24*(4), 668–680.

Cohen Seglias Pallas Greenhall and Furman PC (2018, July). Unbalanced bidding.

Grundl, S. and Y. Zhu (2019). Identification and estimation of risk aversion in first-price auctions with unobserved auction heterogeneity. *Journal of Econometrics 210*(2), 363–378.

Guerre, E., I. Perrigne, and Q. Vuong (2009). Nonparametric identification of risk aversion in first-price auctions under exclusion restrictions. *Econometrica 77*(4), 1193–1227.

Hu, Y., D. McAdams, and M. Shum (2013). Identification of first-price auctions with non-separable unobserved heterogeneity. *Journal of Econometrics 174*(2), 186–193.

Hubbard, T. P. and R. Kirkegaard (2019). Bid-separation in asymmetric auctions.

Hubbard, T. P. and H. J. Paarsch (2014). On the numerical solution of equilibria in auction models with asymmetries within the private-values paradigm. In *Handbook of computational economics*, Volume 3, pp. 37–115. Elsevier.

Klemperer, P. (1999). Auction theory: A guide to the literature. *Journal of Economic Surveys 13*(3), 227–286.

Kong, Y. (2020). Not knowing the competition: evidence and implications for auction design. *The RAND Journal of Economics 51*(3), 840–867.

Krasnokutskaya, E. (2011). Identification and estimation of auction models with unobserved heterogeneity. *Review of Economic Studies 78*(1), 293–327.

Laffont, J.-J. and J. Tirole (1993). *A theory of incentives in procurement and regulation*. MIT press.

Lovallo, D., T. Koller, R. Uhlaner, and D. Kahneman (2020). Your company is too risk-averse; here's why and what to do about it. *Harvard Business Review 98*(2), 104–111.

Luo, Y. (2020). Unobserved heterogeneity in auctions under restricted stochastic dominance. *Journal of Econometrics 216*(2), 354–374.

Luo, Y. and H. Takahashi (2019). Bidding for contracts under uncertain demand: Skewed bidding and risk sharing. *Working Paper*.

Maskin, E. and J. Riley (1984). Optimal auctions with risk averse buyers. *Econometrica: Journal of the Econometric Society*, 1473–1518.

Matthews, S. (1987). Comparing auctions for risk averse buyers: A buyer's point of view. *Econometrica: Journal of the Econometric Society*, 633–646.

Milgrom, P. R. and R. J. Weber (1982). A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, 1089–1122.

Rackauckas, C. and Q. Nie (2017). Differentialequations.jl – a performant and feature-rich ecosystem for solving differential equations in julia. *Journal of Open Research Software 5*, 1.

Reny, P. J. (2011). On the existence of monotone pure-strategy equilibria in bayesian games. *Econometrica 79*(2), 499–553.

Reny, P. J. and S. Zamir (2004). On the existence of pure strategy monotone equilibria in asymmetric first-price auctions. *Econometrica 72*(4), 1105–1125.

Roberts, J. W. and A. Sweeting (2013). When should sellers use auctions? *American Economic Review 103*(5), 1830–61.

Samuelson, W. F. (1985). Competitive bidding with entry costs. *Economics Letters 17*(1), 53–57.

Skitmore, M. and D. Cattell (2013). On being balanced in an unbalanced world. *Journal of the Operational Research Society 64*(1), 138–146.

Somaini, P. (2020). Identification in auction models with interdependent costs. *Journal of Political Economy 128*(10), 3820–3871.

Stark, R. M. (1974). Unbalanced highway contract tendering. *Journal of the Operational Research Society 25*(3), 373–388.

Tait, W. (1971). The role of the civil engineer in the planning, design and construction of a modern highway. *Proceedings of the Institution of Civil Engineers 49*(2), 211–220.

A Scaling Equilibrium Construction

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

- $q^e_{t,n}$: DOT quantity estimates
- $q^b_{t,n}$: Bidder quantity predictions
- $\sigma^2_{t,n}$: Bidder predictions' variances
- $c_{t,n}$: 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 $q^e_{t,n}$, $c_{t,n}$ and I_n are taken directly from data provided by the DOT. The remaining objects are estimates derived from Section 7. The estimates of $q^b_{t,n}$ and $\sigma^2_{t,n}$ 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 $\widehat{\cdot}$ 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.[15] 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 $\mu^\alpha_n = \beta^\alpha X_n$ and a project-type-specific variance σ^α_n. We estimate β^α and σ^α_n from the distribution of estimated $\alpha_{i,n}$ parameters across bidders and auctions, using Hamiltonian Monte Carlo sampling.[16] The resulting fitted parameters fully characterize the distribution of bidder types in each auction.

[15] As 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 (Table 11) and a histogram of the residuals (Figure 9) for the log-Poisson model.

[16] Since efficiency types are drawn from a truncated distribution, within the same sampling procedure we also fit that maximum type $\overline{\alpha}_n$ in each auction, as well as a maximum possible type for each bin of auctions, $\overline{\overline{\alpha}}_{\mathcal{B}(n)}$, to jointly rationalize the distribution of entrants in each bin. We compute $\overline{\alpha}_n$ 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:

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

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

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

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

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

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 have no trouble with this, we use a custom algorithm specified to our problem as detailed in Appendix C for computational efficiency.

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: $\varphi : [\underline{\tau}, \overline{\tau}] \to \mathbb{R}$, we can rewrite the probability that i will win the auction under bid

s as follows:

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

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

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

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:

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

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.[17] 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:

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

where

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

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

$$V(\varphi(\tau)) = 1 - \exp[-\gamma(\tau) \cdot \widetilde{\mathrm{CE}}(\varphi(\tau))]$$

$$\frac{\partial}{\partial s}V(\varphi(\tau)) = \gamma(\tau) \cdot \frac{\partial}{\partial s}\widetilde{\mathrm{CE}}(\varphi(\tau)) \times \exp[-\gamma(\tau) \cdot \widetilde{\mathrm{CE}}(\varphi(\tau))],$$

$$\frac{\partial}{\partial s}\widetilde{\mathrm{CE}}(\varphi(\tau)) = \sum_{t=1}^{T}\left[\frac{\partial b_t^*(\varphi(\tau))}{\partial s}\left(q_t^b - \gamma(\tau) \cdot \sigma^2(b_t^*(\varphi(\tau)) - \alpha(\tau) \cdot c_t)\right)\right].$$

Here the derivative of $b_t^*(\varphi(\tau))$ is taken with respect to the solution of the portfolio problem

[17] Here $a_0(\tau) = 1 - \exp(\gamma(\tau) \cdot \kappa)$ 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.

in Equation (14), and can be computed exactly through forward-mode auto-differentiation 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) 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 choice of the bidder's score:

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

With μ-renegotiation:

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

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:

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

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 $\mathrm{EU}\left(\varphi^*(\tau)|\tau, m\right)$ 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.[18] 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.[19] 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 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 $\varphi(\cdot)$ 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:

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

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

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.

[18] Note 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.

[19] Although 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.

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 over-using items she had bid high on. Knowing this, she might bid higher on items that will be easier to over-use.

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 over-running on profitable items. That is, for each auction n, we take the set of items for which the winning bidder made a profit from over-runs. For each such item t, we define the maximum allowable overuse level: $\bar{y}_{t,n} = (q^a_{t,n} - q^e_{t,n})/q^e_{t,n}$. We assume that all profitable over-runs was intentional in our data, so that the variance on these items, $\bar{\sigma}_{t,n}$, is 0.[20].

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 item in each auction format. To do so, we modify Equation (5) as follows:

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

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

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 similarly to the baseline case, where $q^b_{t,n}$ can be inflated by $y_{t,n} \times 100$ percent, up to the maximum overuse level $\bar{y}_{t,n}$ for each item.

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

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 over-use an item t, it must be optimal to over-use 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.[21] 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 6, 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.

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

Table 6: A comparison of counterfactual savings with and without moral hazard[a]

[a]Due 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.

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

[21]Given 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.

(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.[22]

Proposition 1. *Suppose that each bidder's type (α, γ) is drawn from a 2-dimensional Euclidean cube $[\underline{\alpha}, \overline{\alpha}] \times [\underline{\gamma}, \overline{\gamma}]$. 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:*

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

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 non-zero bids is known:

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

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

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

is sufficient to show increasing differences for the certainty equivalent function:

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

Taking derivatives, we obtain:

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

Noting that

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

does not depend on α or γ, and

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

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

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

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 $s > \overline{\alpha} \sum\limits_t \Big[q_t^e c_t \Big]$, then $\frac{\partial^2 CE(s,\alpha,\gamma)}{\partial s \partial \gamma} < 0$ for all s, α and γ. Reparameterizing bidder risk aversion types types according to $\tilde{\gamma} = -\gamma$, we obtain the increasing differences condition: $\frac{\partial^2 CE(s,\alpha,\tilde{\gamma})}{\partial s \partial \tilde{\gamma}} > 0$.

□

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

Our dataset consists of 440 bridge projects with a total of 218,110 unit bid observations. Of these, there are 2,883 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 2,985 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.[23] 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 spec 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.[24] The resulting matrix is very high dimensional, and so we project the matrix onto its principle components, and use the first 15.[25] In addition, we added 3 stand-alone project features: a dummy variable indicating whether the item is often given a single unit quantity (indicating that its quantity is particularly discrete), 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 $X_{t,n}$, 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

[23] That 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.

[24] There are 11 items that have been flagged at our request by the cheif 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)

[25] We have tried replicating this using more/less principle components and the results are very stable.

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 less than 10 auctions in our data set into a single firm ID. This results in 52 unique bidder IDs: 51 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.[26] Our $X_{i,n}$ matrix has a total of 14 columns, and so we have a total of 132 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 Appendix C.3.2.

C.3 Econometric Details

Let $b^d_{t,i,n}$ 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 parameterize the model prediction for each bid $b^*_{t,i,n}(s^*_{i,n}|\theta)$, as defined by Equation (6). The sub-vector θ_1 refers to parameters estimated in the first stage, as detailed in Appendix C.3.1. The sub-vector 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: $\nu_{t,i,n} = b^d_{t,i,n} - b^*_{t,i,n}(s^*_{i,n}|\theta)$, is distributed identically and independently with a mean of zero across items, bidders and auctions. Furthermore, $\nu_{t,i,n}$ 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 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,

[26] We define "fringe" similarly to Bajari et al. (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 dataset.

we estimate θ_1, the sub-vector of parameters that governs 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 dataset 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 below. For simplicity, we employ a linear model for the expected quantity of item t in auction n, $q^b_{t,n}$ as a function of the DOT quantity estimate $q^e_{t,n}$ and $X_{t,n}$.[27] In order to model the level of uncertainty in the projection $q^b_{t,n}$, we model the distribution of the quantity model fit residuals ($\eta_{t,n} = q^a_{t,n} - q^b_{t,n}$) with a LogNormal regression function of $q^e_{t,n}$ and $X_{t,n}$ as well. The full model specification is below. 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) (Betancourt, 2017) with the Stan probabilistic programming language.[28] We then take the posterior means of the estimated distributions and use them as point estimates for the second stage.

$$q^a_{t,n} = q^b_{t,n} + \eta_{t,n} \text{ where } \eta_{t,n} \sim \mathcal{N}(0, \sigma^2_{t,n}) \tag{24}$$

such that $$q^b_{t,n} = \beta_{0,q} q^e_{t,n} + \beta_q X_{t,n} \text{ and } \sigma_{t,n} = \exp(\beta_{0,\sigma} q^e_{t,n} + \beta_\sigma X_{t,n}). \tag{25}$$

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: $\widehat{q}^b_{t,n}$ as well as the variance of that prediction, $\widehat{\sigma}^2_{t,n}$.

C.3.2 Second Stage

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

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

[28] See Carpenter, Gelman, Hoffman, Lee, Goodrich, Betancourt, Brubaker, Guo, Li, and Riddell (2017).

score of $s_{i,n}$ is given by:

$$b^*_{t,i,n}(s_{i,n}) = \alpha_{i,n} c_{t,n} + \frac{q^b_{t,n}/\sigma^2_{t,n}}{\gamma_{i,n}} + \frac{q^e_{t,n}/\sigma^2_{t,n}}{\sum\limits_{r:b^*_{r,i,n}(s)>0} \left[\frac{(q^e_{r,n})^2}{\sigma^2_{r,n}}\right]} \left(s_{i,n} - \sum_{r:b^*_{r,i,n}(s_{i,n})>0} q^e_{r,n} \left[\alpha_{i,n} c_{r,n} + \frac{q^b_{r,n}/\sigma^2_{r,n}}{\gamma_{i,n}} \right] \right),$$

where $\left(q^b_{1,n}, \ldots q^b_{T_n,n}\right)$ and $\left(\sigma^2_{1,n}, \ldots \sigma^2_{T_n,n}\right)$ 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 $s_{i,n}$:

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

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,

$$A_{t,i,n} = c_{t,n} - \frac{q^e_{t,n}/\sigma^2_{t,n}}{\sum\limits_{r:b^*_{r,i,n}(s)>0} \left[\frac{(q^e_{r,n})^2}{\sigma^2_{r,n}}\right]} \sum_{r:b^*_{r,i,n}(s_{i,n})>0} q^e_{r,n} c_{r,n}$$

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 $q^b_{t,n}/\sigma^2_{t,n}$ 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 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 $q^b_{t,n}$ and $\sigma^2_{t,n}$ are pinned down by our first stage estimates $\widehat{q}^b_{t,n}$ and $\widehat{\sigma}^2_{t,n}$; (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 over-determined 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:

$$(b^d_{t,i,n} - Z^d_{t,i,n}(\hat{\theta}_1) \cdot s^d_{i,n}) = A_{t,i,n}(\hat{\theta}_1) \cdot \alpha_{i,n} + G_{t,i,n}(\hat{\theta}_1) \cdot \frac{1}{\gamma_{i,n}} + \tilde{\nu}_{t,i,n}, \tag{27}$$

where $b^d_{t,i,n}$ and $s^d_{i,n}$ 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 (30) 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 $\mathbb{E}[\tilde{\nu}_{t,i,n}|A_{t,i,n}(\theta_1), G_{t,i,n}(\theta_1), Z^d_{t,i,n}(\theta_1)] = 0$ within each bidder-auction. However, as some auctions have relatively few items and the variance of bid errors may be large across item, we pool across auctions by projecting $\alpha_{i,n}$ and $\gamma_{i,n}$ on a vector of bidder-auction characteristics:

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

This projection induces heteroskedasticity 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, $\mathbb{E}[\tilde{\nu}_{t,i,n}|A_{t,i,n}(\theta_1), G_{t,i,n}(\theta_1), Z^d_{t,i,n}(\theta_1), X_{i,n}] = 0$ across bidders i and auctions n, that is guaranteed by our Assumption 1.

GMM Specification To efficiently estimate our second stage, we implement a GMM procedure that applies the orthogonality condition above across sub-samples 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 $\theta_2 = (\beta^1_{0,\gamma}, \ldots, \beta^I_{0,\gamma}, \beta^1_{X,\gamma}, \ldots, \beta^J_{X,\gamma}, \beta^1_{0,\alpha}, \ldots, \beta^I_{0,\alpha}, \beta^1_{X,\alpha}, \ldots, \beta^J_{X,\alpha})$ 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.[29] We estimate θ_2 in the second stage, using a GMM framework

[29]To simplify notation, we do not distinguish between 'unique' bidders—e.g. bidders who appear in 30+ auctions—and rare bidders, whom we group into a single unique bidder ID for the purposes of this econometrics section. For the latter group, we treat all observations of rare bidders as observations of the

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

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

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 4 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.[30] As there are both $2J + 2I$ second stage parameters and moments, we use an identity matrix for W.

$$m_i^1(\theta_2|\hat{\theta}_1) = \mathbb{E}_n \left[\frac{1}{|\mathcal{T}(n)|} \sum_{t \in \mathcal{T}(n)} \tilde{\nu}_{t,i,n}(\theta_2|\hat{\theta}_1) \cdot \mathbf{1}_{i \in \mathcal{I}(n)} \right]$$

$$m_j^2(\theta_2|\hat{\theta}_1) = \mathbb{E}_n \left[\frac{1}{|\mathcal{I}(n)| \cdot |\mathcal{T}(n)|} \sum_{i \in \mathcal{I}(n)} \sum_{t \in \mathcal{T}(n)} \tilde{\nu}_{t,i,n}(\theta_2|\hat{\theta}_1) \cdot \mathbf{1}_{i \in \mathcal{I}(n)} \cdot X_{i,n}^j \right]$$

$$m_i^3(\theta_2|\hat{\theta}_1) = \mathbb{E}_n \left[\frac{1}{|\mathcal{T}_s(n)|} \sum_{t \in \mathcal{T}(n)} \tilde{\nu}_{t,i,n}(\theta_2|\hat{\theta}_1) \cdot \mathbf{1}_{i \in \mathcal{I}(n)} \cdot \mathbf{1}_{t \in \mathcal{T}_s} \right]$$

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

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

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.

[30] "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-/under- bid items in our data.

moments are formed with respect to the *de-meaned* bid residual:

$$\begin{aligned}\tilde{\nu}_{t,i,n}(\theta_2|\hat{\theta}_1) = b^d_{t,i,n} - \ \alpha_{i,n}(\theta_2) & \left(c_{t,n} - \frac{q^e_{t,n}}{\widehat{\sigma}^2_{t,n} \sum_{p\in\mathcal{T}(n)} \left[\frac{(q^e_{p,n})^2}{\widehat{\sigma}^2_{p,n}}\right]} \left[\sum_{p\in\mathcal{T}(n)} c_{p,n} q^e_{p,n} \right] \right) \\ & - \frac{1}{\gamma_{i,n}(\theta_2)} \left(\frac{\widehat{q}^b_{t,n}}{\widehat{\sigma}^2_{t,n}} - \frac{q^e_{t,n}}{\widehat{\sigma}^2_{t,n} \sum_{p\in\mathcal{T}(n)} \left[\frac{(q^e_{p,n})^2}{\widehat{\sigma}^2_{p,n}}\right]} \left[\sum_{p\in\mathcal{T}(n)} \frac{\widehat{q}^b_{p,n} q^e_{p,n}}{\widehat{\sigma}^2_{p,n}} \right] \right) \\ & - \frac{q^e_{t,n}}{\widehat{\sigma}^2_{t,n} \sum_{p\in\mathcal{T}(n)} \left[\frac{(q^e_{p,n})^2}{\widehat{\sigma}^2_{p,n}}\right]} \left[s^d_{i,n} \right],\end{aligned} \tag{29}$$

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

The residual terms in the moments are *de-meaned* in the sense that they use the *observed* score $s^d_{i,n}$ 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 de-meaned 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 $\bar{\nu}_{\gamma,i,n}$.

$$\tilde{\nu}_{t,i,n} = \nu_{i,t,n} + \bar{\nu}_{s,i,n} + \bar{\nu}_{\alpha,i,n} + \bar{\nu}_{\gamma,i,n}, \tag{30}$$

where

$$\begin{aligned}\bar{\nu}_{s,i,n} &= \frac{q^e_{t,n}}{\widehat{\sigma}^2_{t,n} \sum_{p\in\mathcal{T}(n)} \left[\frac{(q^e_{p,n})^2}{\widehat{\sigma}^2_{p,n}}\right]} \sum_{t=1}^{T_n} \nu_{t,i,n} q^e_{t,n}, \\ \bar{\nu}_{\alpha,i,n} &= \nu_{\alpha,i,n} \left(c_{t,n} - \frac{q^e_{t,n}}{\widehat{\sigma}^2_{t,n} \sum_{p\in\mathcal{T}(n)} \left[\frac{(q^e_{p,n})^2}{\widehat{\sigma}^2_{p,n}}\right]} \left[\sum_{p\in\mathcal{T}(n)} c_{p,n} q^e_{p,n} \right] \right), \\ \bar{\nu}_{\gamma,i,n} &= \nu_{\gamma,i,n} \left(\frac{\widehat{q}^b_{t,n}}{\widehat{\sigma}^2_{t,n}} - \frac{q^e_{t,n}}{\widehat{\sigma}^2_{t,n} \sum_{p\in\mathcal{T}(n)} \left[\frac{(q^e_{p,n})^2}{\widehat{\sigma}^2_{p,n}}\right]} \left[\sum_{p\in\mathcal{T}(n)} \frac{\widehat{q}^b_{p,n} q^e_{p,n}}{\widehat{\sigma}^2_{p,n}} \right] \right).\end{aligned} \tag{31}$$

The formulas in Equation (31) 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 $s^d_{i,n}$ is observed with a sum of errors:

$$s^d_{i,n} = \sum_{t=1}^{T_n} (b^*_{t,i,n} + \nu_{t,i,n}) q^e_{t,n} = s^*_{i,n} + \sum_{t=1}^{T_n} \nu_{t,i,n} q^e_{t,n}.$$

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.

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:

$$\hat{\theta}_2 = \min_{\theta_2} \left\{ \frac{1}{I} \sum_i m^1_i(\theta_2|\hat{\theta}_1)^2 + \frac{1}{J} \sum_{j=1}^{J} m^2_j(\theta_2|\hat{\theta}_1)^2 + \frac{1}{I} \sum_i m^3_i(\theta_2|\hat{\theta}_1)^2 + \frac{1}{J} \sum_{j=1}^{J} m^4_j(\theta_2|\hat{\theta}_1)^2 \right\}$$

   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.[31]

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.[32] Next, we draw 100 auctions at random with replacement from the

[31]This 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 non-negativity are non-differentiable functions. The results do not differ to an economically significant degree.

[32]Our 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 dsitribution of which is summarized in Table 7. 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

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 8 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 et al. (2011); Roberts and Sweeting (2013)) and more generally in first price auctions with risk averse bidders (e.g. Guerre et al. (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 bridge 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.[33]

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

parameter, we instead take 100 random samples, and plug each draw into each iteration of the second stage estimator.

[33]In 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.

a multiplicative shock $\frac{1}{u^m} > 0$, we can rewrite Equation (3) as:

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

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:

$$(b_{t,i,n}^d - Z_{t,i,n}^d(\hat{\theta}_1) \cdot s_{i,n}^d) = A_{t,i,n}(\hat{\theta}_1) \cdot \alpha_{i,n} + G_{t,i,n}(\hat{\theta}_1) \cdot \frac{u_n^m}{\gamma_{i,n}} + \tilde{\nu}_{t,i,n}. \tag{32}$$

As Equation (32) 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 shock 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 estimates for $\gamma_{i,n}$ that captures the ratio $\frac{\gamma_{i,n}}{u_n^m}$, which plays the same functional role in counterfactuals.[34] If u_n^m is an inflation adjustment, $\frac{\gamma_{i,n}}{u_n^m}$ 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 $\sigma_{t,n}^2 + \overline{\sigma_n^2}$ or $\sigma_{t,n}^2 \times \overline{\sigma_n^2}$ for an unobservable factor $\overline{\sigma_n^2}$. 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.

[34]Note 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 5b 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.

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.[35]

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.[36]

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

[35] See Table 1 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 4 for the distribution of ex-post markups (without accounting for EWOs or the cost of uncertainty).

[36] Note that although $\lambda \cdot$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.

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.[37] 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 $\widehat{\tau_n^*}(\kappa, \lambda)$ that obtains zero expected utility of entry under the empirical entry frequency $\hat{q}_n$:

$$V_n(\widehat{\tau_n^*}(\kappa, \lambda)) = \sum_{m=1}^{M} \begin{pmatrix} M-1 \\ m-1 \end{pmatrix} \hat{q}_n^{m-1}(1-\hat{q}_n)^{M-m} \cdot \mathrm{EU}_n\left(\varphi^*(\widehat{\tau_n^*}(\kappa, \lambda))|\kappa, \lambda, m\right) = 0. \tag{33}$$

We then compute two statistics of the entry model: the predicted entry rate $q_n^p(\kappa, \lambda)$, and the threshold type quantile $F_n(\widehat{\tau_n^*}(\kappa, \lambda))$. To obtain $F_n(\widehat{\tau_n^*}(\kappa, \lambda))$, we evaluate the CDF of the distribution of potential types in auction n at $\widehat{\tau_n^*}(\kappa, \lambda)$. In equilibrium, only bidders whose types are below $\widehat{\tau_n^*}(\kappa, \lambda)$ enter the auction, and so $F_n(\widehat{\tau_n^*}(\kappa, \lambda))$ yields the probability of entry in Equation (33). Comparing $F_n(\widehat{\tau_n^*}(\kappa, \lambda))$ 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:

$$\hat{\kappa}_{\mathcal{B}(n)}, \hat{\lambda}_{\mathcal{B}(n)} = \arg\min_{\kappa,\lambda} \Big\{ \sum_{n\in\mathcal{B}n} (\hat{q}_n - q_n^p(\kappa, \lambda))^2 + (\hat{q}_n - F_n(\widehat{\tau^*}(\kappa, \lambda)))^2 \Big\}. \tag{34}$$

[37] An 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.

D Estimation Results Tables

First Stage Parameter Estimates

Parameter	Rhat	n_eff	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

Table 7: First Stage Parameter Estimates

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 estimtaion procedure. In each iteration, we redraw 438 auctions at random with replacement and reestimate our second stage GMM model.[38] Table 8 presents the resulting 95% confidence interval for the headline estimates in Tables 3 and 4, as well as their standard deviations both across all draws, and within the 95% confidence interval.

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

Table 8: Second Stage Bootstrap Errors and Quantiles

First Stage Model Fit

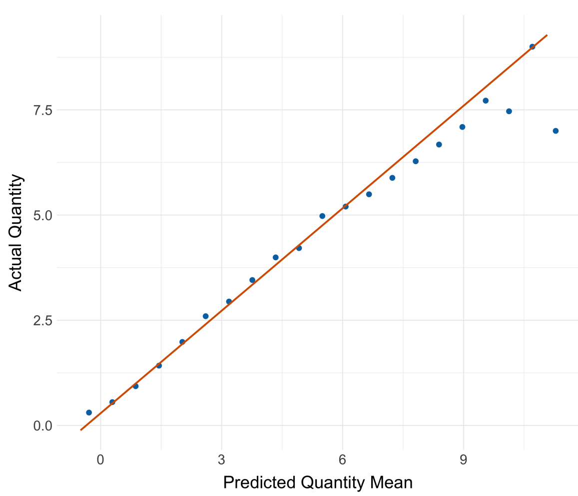

Figure 6: Bin scatter of actual quantities vs predictions from the first stage model fit.

	Dependent variable:
	Actual Quantity
Predicted Quantity	0.812***
	(0.005)
Constant	0.291***
	(0.015)
Observations	29,834
R^2	0.476

Table 9: Regression report for Figure 6

[38]We exclude 2 auctions with outlying high costs from the second stage of estimation, leaving 438 auctions.

Second Stage Model Fit

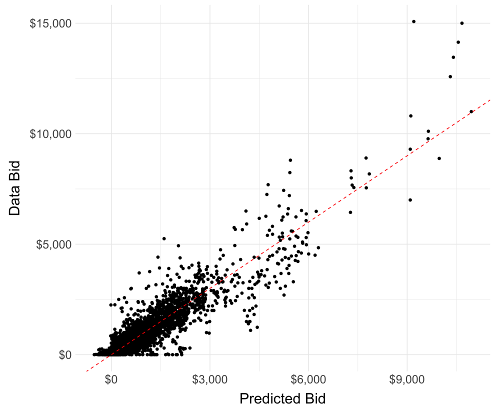

Figure 7: Scatter plot of observed unit bids vs fitted bids from the second stage model. Note: Unit bids are scaled so as to standardize quantities so exact dollar values are not representative.

	Dependent variable:
	Data Bid
Predicted Bid	0.967***
	(0.001)
Constant	989.557***
	(162.154)
Observations	215,332
R^2	0.881

Table 10: Regression report for Figure 7

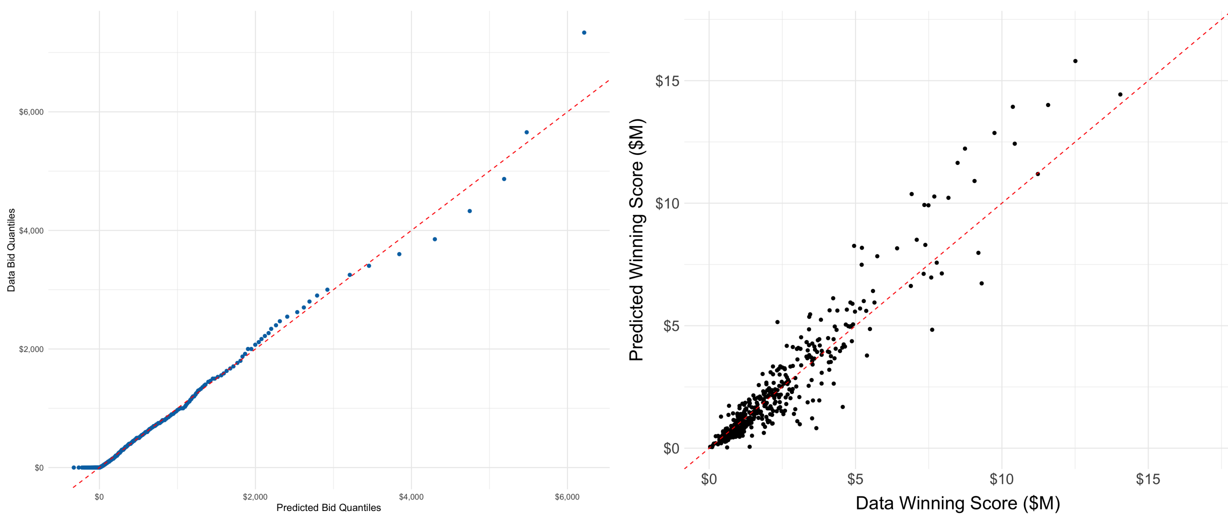

(a) Quantile-Quantile plot of predicted bids against data bids. Quantiles are presented at the 0.0001 level and truncated at the top and bottom 0.01%.

(b) Scatter plot of actual winning scores against the winning scores predicted by our equilibrium simulation at the estimated parameters.

Figure 8: Fit plots for bids and scores with the 45-degree line, dashed in red, for reference.

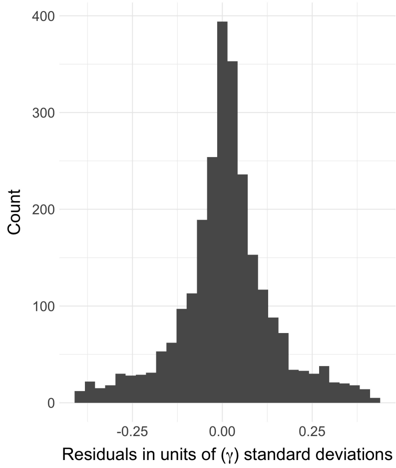

Figure 9: Histogram of residuals from the Poisson regression model discussed in Appendix A in units of γ standard deviations, truncated at top/bottom 5% for visibility.

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

Table 11: Regression report for the prediction fit of the Poisson regression model for γ discussed in Appendix A.

E Additional Figures

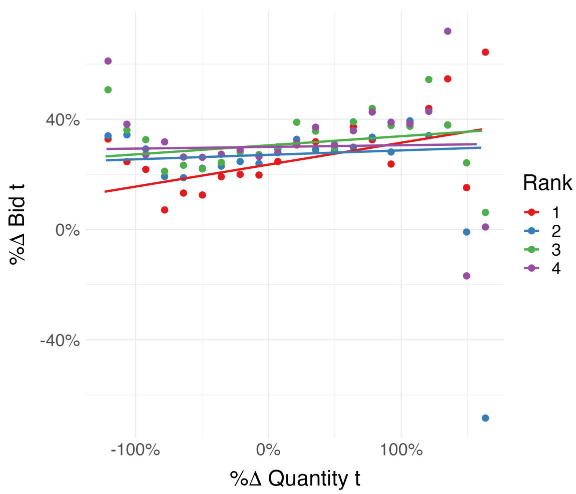

(a) Replication of Figure 3a with bidders of rank 1-4.

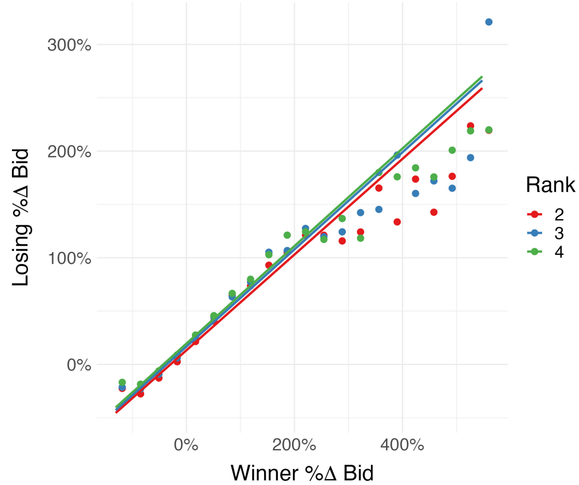

(b) Replication of Figure 3b with bidders of rank 1-4.

Figure 10

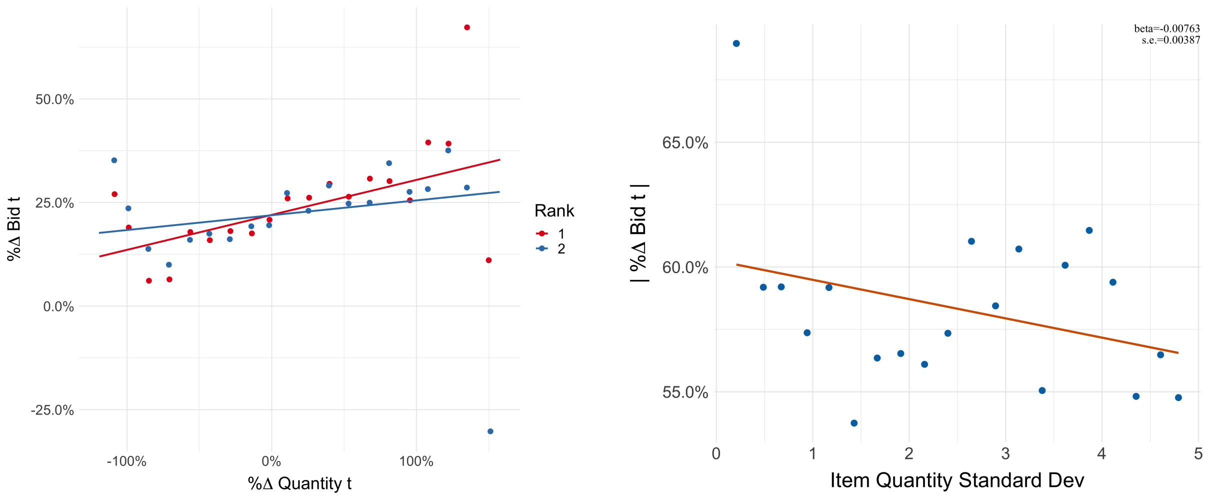

(a) Replication of Figure 3a when the top two bidders' scores are within 10% of each other.

(b) Replication of Figure 4a, without controlling for $\%\Delta q_t$.

Figure 11

F Counterfactual Results Tables

We report the summary statistics for the counterfactual results reported in Section 8.[39]

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

Table 12: Summary Statistics for CF Threshold Type Changes under Endogenous Entry

[39] Each 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.

CF Type	Outcome	Mean	SD	25%	50%	75%
With Endogenous Entry						
Lump Sum	% DOT Savings	-97.8	190.8	-102.8	-42.2	-17.7
Lump Sum w 2:1 Negotiation	% DOT Savings	-17.3	28.6	-28.6	-13.8	-2.9
Lump Sum w 50/50 Negotiation	% DOT Savings	-13.7	24.1	-23.5	-8.5	-0.2
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 Negotiation	$ DOT Savings	-321,549.6	567,099.9	-402,865.9	-124,194.6	-16,320.3
Lump Sum w 50/50 Negotiation	$ DOT Savings	-224,092.1	431,945.0	-261,757.2	-92,431.1	-2,481.1
No Risk (Correct q)	$ DOT Savings	138,945.4	599,836.5	6,510.8	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 Negotiation	% Bidder Gain	166.2	387.2	20.5	48.4	149.9
Lump Sum w 50/50 Negotiation	% Bidder Gain	147.7	283.0	11.9	31.6	135.5
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	9,719.6	37,593.9	87,450.3
Lump Sum w 2:1 Negotiation	$ Bidder Gain	105,305.1	218,136.7	9,911.0	25,754.0	84,472.3
Lump Sum w 50/50 Negotiation	$ Bidder Gain	83,017.4	153,886.3	6,332.6	17,229.4	92,870.8
No Risk (Correct q)	$ Bidder Gain	116,567.3	470,277.3	-24,138.8	-4,951.4	87,635.4
No Risk (Estimated q)	$ Bidder Gain	-188,641.5	462,137.2	-131,162.9	-6,954.3	1,088.3
Holding Entry Probabilities Fixed						
Lump Sum	% DOT Savings	-221.6	379.7	-226.1	-95.8	-44.1
Lump Sum w 2:1 Negotiation	% DOT Savings	-19.0	28.6	-33.0	-14.6	-4.6
Lump Sum w 50/50 Negotiation	% DOT Savings	-8.8	17.6	-16.2	-6.4	-0.8
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 Negotiation	$ DOT Savings	-406,632.3	724,732.9	-471,577.3	-135,711.1	-30,123.9
Lump Sum w 50/50 Negotiation	$ DOT Savings	-192,411.7	361,492.8	-229,326.5	-60,254.4	-5,489.5
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	3,528.1	36,913.2
Lump Sum	% Bidder Gain	59.9	38.4	36.8	55.4	81.5
Lump Sum w 2:1 Negotiation	% Bidder Gain	40.3	35.8	20.4	34.5	54.4
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	8,492.3	17,776.4	31,684.2
Lump Sum w 50/50 Negotiation	$ Bidder Gain	12,895.1	23,537.1	4,488.2	9,993.6	18,917.0
No Risk (Correct q)	$ Bidder Gain	9,333.1	134,422.0	-17,502.0	-9,346.5	-3,353.8
No Risk (Estimated q)	$ Bidder Gain	-16,629.8	134,947.4	-1,355.1	332.4	3,264.4

Table 13: Summary Statistics for CF Outcomes With and Without Endogenous Entry

A Supplemental Appendix (Not for Publication)

A.1 Additional Tables and Figures

Distribution of Projects by Year in Our Data

	Year	Num Projects	Percent	Cumul Percent
1	1998	1	0.227	0.227
2	1999	5	1.136	1.364
3	2000	5	1.136	2.500
4	2001	20	4.545	7.045
5	2002	27	6.136	13.182
6	2003	26	5.909	19.091
7	2004	25	5.682	24.773
8	2005	37	8.409	33.182
9	2006	21	4.773	37.955
10	2007	32	7.273	45.227
11	2008	53	12.045	57.273
12	2009	46	10.455	67.727
13	2010	61	13.864	81.591
14	2011	32	7.273	88.864
15	2012	24	5.455	94.318
16	2013	19	4.318	98.636
17	2014	6	1.364	100

Table 14: Distribution of projects by year in our data

Estimated Number of Employees for Most Common Firms

Bidder Name	# Employees	# Auctions Bid	# Auctions Won
MIG Corporation	80	297	38
Northern Constr Services LLC	80	286	26
SPS New England Inc	75	210	58
ET&L Corp	1	201	26
B&E Construction Corp	9	118	16
NEL Corporation	68	116	36
Construction Dynamics Inc	22	113	10
S&R Corporation	20	111	16
New England Infrastructure	35	95	6
James A Gross Inc	7	78	7

Table 15: All 24 most common firms in our sample are privately owned, and so there is no publicly available, verifiable information on their revenues or expenses. The numbers of employees presented here were drawn from Manta, an online directory of small businesses, and cross-referenced with LinkedIn, on which a subset of these firms list a range of their employee counts, as of November 2018. Note that there is some ambiguity as to who "counts" as an employee, as such firms often hire additional construction laborers on a project-by-project basis. The "family owned" label is drawn from the firms' self-descriptions on their websites.

Bid Level-Weighted Bins

We replicate the main graphs from Section 4, weighting the dots by the average bid levels in the bins that they represent. This demonstrates that outlier dots are generally relatively small, minor items, and that overestimated items are not rare.

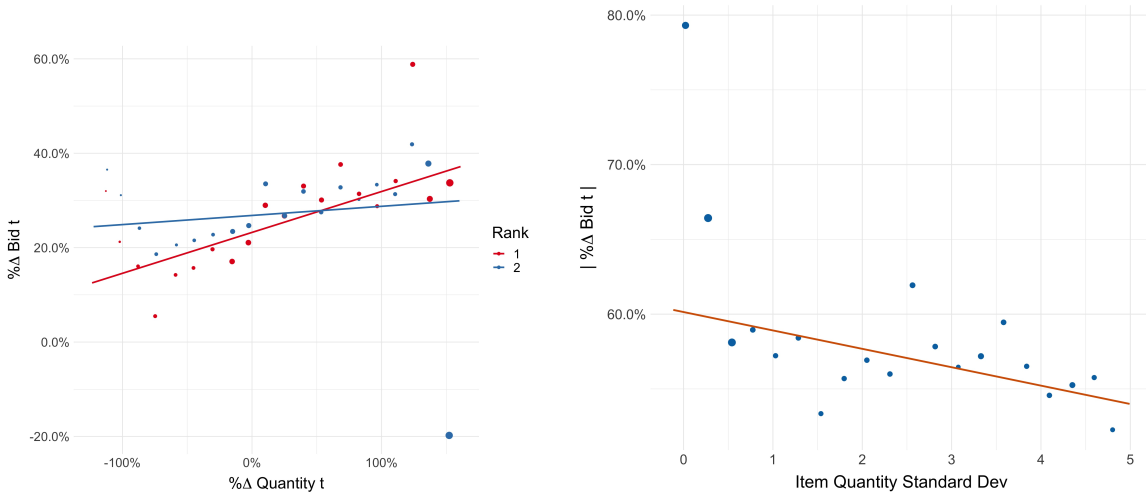

(a) Replication of Figure 3a with weighted bins. (b) Replication of Figure 4a with weighted bins.

Figure 12

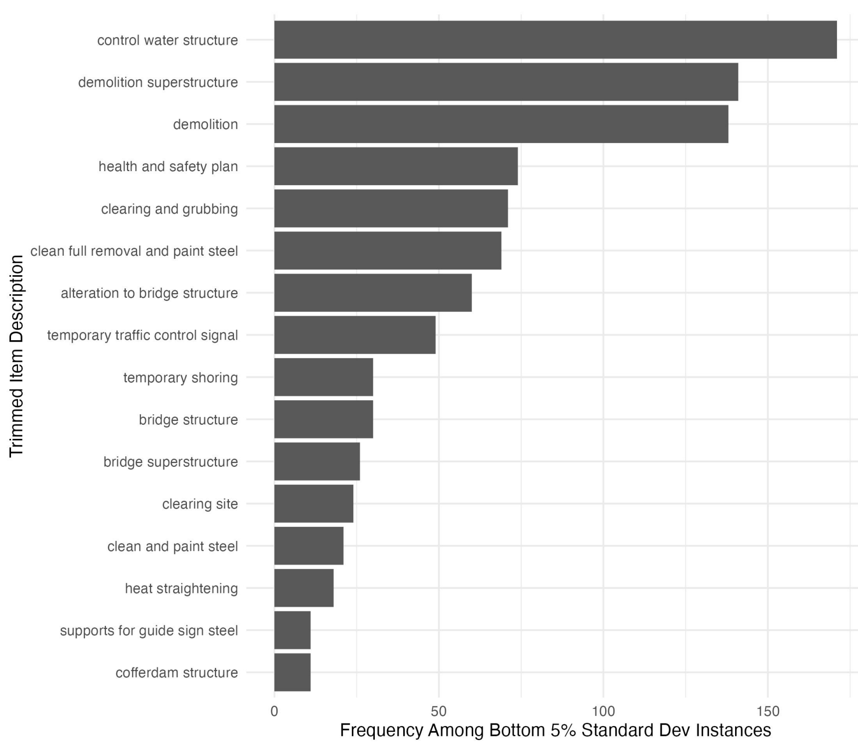

Figure 13: Trimmed description of items with at least 10 instances among the 5% lowest $\sigma_{t,n}$

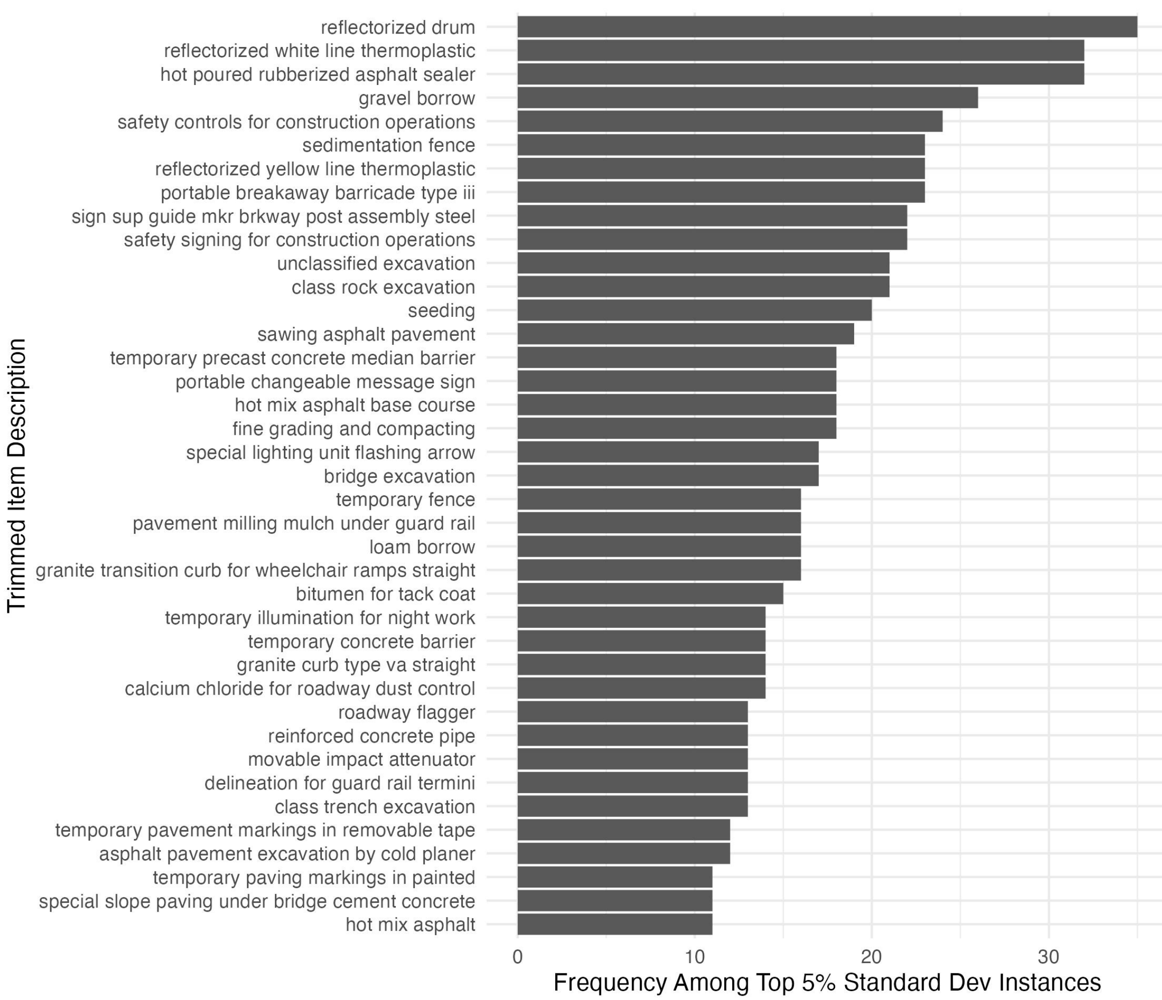

Figure 14: Trimmed description of items with at least 10 instances among the 5% highest $\sigma_{t,n}$

A.2 Illustrative Example

Consider the following simple example of infrastructure procurement bidding. Two bidders compete for a project that requires two types of inputs to complete: concrete and traffic cones. MassDOT ("the DOT" for short) estimates that 10 tons of concrete and 20 traffic cones will be necessary to complete the project. Upon inspection, the bidders determine that the actual quantities of each item that will be used – random variables that we will denote q_c^a and q_r^a for concrete and traffic cones, respectively – are normally distributed with means $\mathbb{E}[q_c^a] = 12$ and $\mathbb{E}[q_r^a] = 16$ and variances $\sigma_c^2 = 2$ and $\sigma_r^2 = 1$.[40] We assume that the actual quantities are exogenous to the bidding process, and do not depend on who wins the auction in any way. Furthermore, we will assume that the bidders' expectations are identical across both bidders.[41]

The bidders differ in their private costs for implementing the project. They have access to the same vendors for the raw materials, but differ in the cost of storing and transporting the materials to the site of construction as well as the cost of labor, depending on the site's location, the state of their caseload at the time and firm-level idiosyncrasies. We therefore describe each bidder's cost as a multiplicative factor α of market-rate cost estimate for each item: $c_c = \$8/\text{ton}$ for each ton of Concrete and $c_r = \$12/\text{pack}$ for each pack of 100 traffic cones. Each bidder i knows her own type α^i at the time of bidding, as well as the distribution (but not realization) of her opponent's type.

To participate in the auction, each bidder i submits a unit bid for each of the items: b_c^i and b_r^i. The winner of the auction is then chosen on the basis of her *score*: the sum of her unit bids multiplied by the DOT's quantity estimates:

$$s^i = 10b_c^i + 20b_r^i.$$

Once a winner is selected, she will implement the project and earn net profits of her unit bids, less the unit costs of each item, multiplied by the *realized* quantities of each item that are ultimately used. At the time of bidding, these quantities are unrealized samples of random variables.

Bidders are endowed with a standard CARA utility function over their earnings from the project with a common constant coefficient of absolute risk aversion γ:

$$u(\pi) = 1 - \exp(-\gamma\pi).$$

Bidders are exposed to two sources of risk: (1) uncertainty over winning the auction; (2) uncertainty over the profits that they would earn at the realized ex-post quantity of each item.

[40] As we discuss in section 5, we assume that the distributions of q_c^a and q_r^a are independent conditional on available information regarding the auction. This assumption, as well as the assumption that the quantity distributions are not truncated at 0 (so that quantities cannot be negative) are made for the purpose of computational traceability in our structural model. If item quantities are correlated, bidders' risk exposure is higher, and so our results can be seen as a conservative estimate of this case.

[41] These assumptions align with the characterization of highway and bridge projects in practice: the projects are highly standardized and all decisions regarding quantity changes must be approved by an on-site DOT official, thereby limiting contractors' ability to influence ex-post quantities.

The profit π that bidder i earns is either 0, if she loses the auction, or

$$\pi(\mathbf{b}^i, \alpha^i, \mathbf{c}, \mathbf{q}^a) = q_c^a \cdot (b_c^i - \alpha^i c_c) + q_r^a \cdot (b_r^i - \alpha^i c_r),$$

if she wins the auction. Bidder i's expected utility at the time of the auction is therefore given by:

$$\mathbb{E}[u(\pi(\mathbf{b}^i, \alpha^i, \mathbf{c}, \mathbf{q}^a))] = \left(\underbrace{1 - \mathbb{E}_{\mathbf{q}^a}\left[\exp\left(-\gamma \cdot \pi(\mathbf{b}^i, \alpha^i, \mathbf{c}, \mathbf{q}^a)\right)\right]}_{\text{Expected utility conditional on winning}} \right) \times \underbrace{(\Pr\{s^i < s^j\}).}_{\text{Probability of winning with } s^i = 10b_c^i + 20b_r^i}$$

That is, bidder i's expected utility from submitting a set of bids b_c^i and b_r^i is the product of the utility that she expects to get (given those bids) if she were to win the auction, and the probability that she will win the auction at those bids. The expectation of utility conditional on winning is with respect to the realizations of the item quantities q_c^a and q_r^a, entirely.

As the ex-post quantities are distributed as independent Gaussians, the expected utility term above can be rewritten in terms of the certainty equivalent of bidder i's profits conditional on winning:[42]

$$1 - \exp\left(-\gamma \cdot \mathrm{CE}(\mathbf{b}^i, \alpha^i, \mathbf{c}, \mathbf{q}^a)\right),$$

where the certainty equivalent of profits $\mathrm{CE}(\mathbf{b}^i, \alpha^i, \mathbf{c}, \mathbf{q}^a)$ is given by:

$$\underbrace{\mathbb{E}[q_c^a] \cdot (b_c^i - \alpha^i c_c) + \mathbb{E}[q_r^a] \cdot (b_r^i - \alpha^i c_r)}_{\text{Expection of Profits}} - \underbrace{\left[\frac{\gamma \sigma_c^2}{2} \cdot (b_c^i - \alpha^i c_c)^2 + \frac{\gamma \sigma_r^2}{2} \cdot (b_r^i - \alpha^i c_r)^2\right]}_{\text{Variance of Profits}}. \tag{35}$$

Furthermore, as we discuss in Section 5, the optimal selection of bids for each bidder i can be described as the solution to a two-stage problem:

Inner: For each possible score s, choose the bids b_c and b_r that maximize $\mathrm{CE}(\{b_c, b_r\}, \alpha^i, \mathbf{c}, \mathbf{q}^a)$, subject to the score constraint: $10b_c + 20b_r = s$.

Outer: Choose the score $s^*(\alpha^i)$ that maximizes expected utility $\mathbb{E}[u(\pi(\mathbf{b}^i(s), \alpha^i))]$, where $\mathbf{b}^i(s)$ is the solution to the inner step, evaluated at s.

That is, at every possible score that bidder i might consider, she chooses the bids that sum to s for the purpose of the DOT's evaluation of who will win the auction, and maximize her certainty equivalent of profits conditional on winning. She then chooses the score that maximizes her total expected utility.

To see how this decision process can generate bids that appear mathematically unbalanced, suppose, for example, that the common CARA coefficient is $\gamma = 0.05$, and consider a bidder in this auction who has type $\alpha^i = 1.5$.[43] Suppose, furthermore, that the bidder has decided to submit

[42]See Section 5 and the Appendix for a detailed derivation.

[43]That is, for each ton of concrete that will be used will cost, the bidder incur a cost of $\alpha^i \times c_c = 1.5 \times \$8 = \$12$, and for each pack of traffic cones that will be used, she will incur a cost of $\alpha^i \times c_r = 1.5 \times \$12 = \$18$.

a total score of \$500. There are a number of ways in which the bidder could construct a score of \$500. For instance, she could bid her cost on concrete, $b_c^i = \$12$, and a dollar mark-up on traffic cones: $b_r^i = (\$500 - \$12 \times 10)/20 = \$19$. Alternatively, she could bid her cost on traffic cones, $b_r^i = \$18$, and a two-dollar mark-up on traffic cones: $b_c^i = (\$500 - \$18 \times 20)/10 = \$14$. Both of these bids would result in the same score, and so give the bidder the same chances of winning the auction. However, they yield very different expected utilities to the bidder. Plugging each set of bids into equation (35), we find that the first set of bids produces a certainty equivalent of:

$$12 \times (\$0) + 16 \times (\$1) - \frac{0.05 \times 2}{2} \times (\$0)^2 - \frac{0.05 \times 1}{2} \times (\$1)^2 = \$15.98,$$

whereas the second set of bids produces a certainty equivalent of

$$12 \times (\$2) + 16 \times (\$0) - \frac{0.05 \times 2}{2} \times (\$2)^2 - \frac{0.05 \times 1}{2} \times (\$0)^2 = \$23.80.$$

In fact, further inspection shows that the optimal bids giving a score of \$500 are $b_c^i = \$47.78$ and $b_r^i = \$1.12$, yielding a certainty equivalent of \$87.98. The intuition for this is precisely that described by **?**, and the contractors cited by Stark (1974): the bidder predicts that concrete will over-run in quantity – she predicts that 12 tons will be used, whereas the DOT estimated only 10 – and that traffic cones will under-run – she predicts that 16 will be used, rather than the DOT's estimate of 20. When the variance terms aren't too large (relatively), the interpretation is quite simple: every additional dollar bid on concrete is worth approximately 12/10 in expectation, whereas every additional dollar bid on traffic cones is worth only 16/20.

However, the incentive to bid higher on items projected to over-run is dampened when the variance term is relatively large. This can arise when the coefficient of risk aversion is relatively high or when the variance of an item's ex-post quantity distribution is high. More generally, as demonstrated in equation (35), the certainty equivalent of profits is increasing in the expected quantity of each item, $\mathbb{E}[q_c^a]$ and $\mathbb{E}[q_r^a]$, but decreasing in the variance of each item σ_c^2 and σ_r^2.

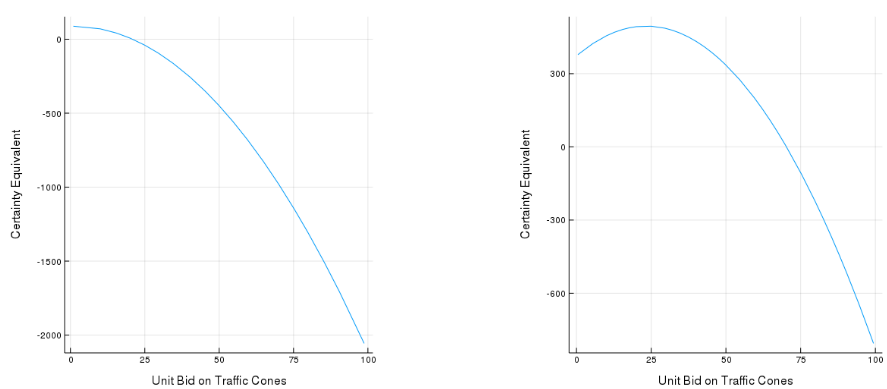

(a) Score = \$500 (b) Score = \$1000

Figure 15: Certainty equivalent as a function of her unit bid on traffic cones, for the example bidder submitting a score of \$500 or \$1,000.

Moreover, the extent of bid skewing can depend on the level of competition in the auction. Figure 15 plots the bidder's certainty equivalent as a function of her unit bid on traffic cones when she chooses to submit a total score of (a) \$500, and when she chooses to submit a score of (b) \$1,000. In the first case, the bid that optimizes the certainty equivalent is very small, $b_r^i = \$1.12$. In the second case, however, the optimal bid is much higher at $b_r^i = \$23.33$. The reason for this is that a low bid on traffic cones implies a high bid on concrete. A high markup on concrete decreases the bidder's certainty equivalent at a quadratic rate. Thus, as the score gets higher, there is more of an incentive to spread markups across items, rather than bidding very high on select items, and very low on others.

A.3 Bid Skewing in Equilibrium

As we discuss in Section 5, the auction game described above has a unique Bayes Nash Equilibrium. This equilibrium is characterized following the two-stage procedure described in Appendix A: (1) given an equilibrium score $s(\alpha)$, each bidder of type α submits the vector of unit bids that maximizes her certainty equivalent conditional on winning, and sums to $s(\alpha)$; (2) The equilibrium score is chosen optimally, such that there does not exist a type α and an alternative score $\tilde{s}$, so that a bidder of type α can attain a higher expected utility with the score $\tilde{s}$ than with $s(\alpha)$.

The optimal selection of bids given an equilibrium score depends on the bidders' expectations over ex-post quantities and the DOT's posted estimates, as well as on the coefficient of risk aversion and the level of uncertainty in the bidders' expectations. High over-runs cause bidders to produce more heavily skewed bids, whereas high risk aversion and high levels of uncertainty push bidders to produce more balanced bids.

In addition to influencing the relative skewness of bids, these factors also have a level effect on bidder utility. Higher expectations of ex-post quantities raise the certainty equivalent conditional on winning for every bidder. Higher levels of uncertainty (and a higher degree of risk aversion), however, induce a cost for bidders that lowers the certainty equivalent. Consequently, higher levels of uncertainty lower the value of participating for every bidder and result in less aggressive bidding behavior, and higher costs to the DOT in equilibrium.

To demonstrate this, we plot the equilibrium score, unit-bid distribution and ex-post revenue for every bidder type α in our example. To illustrate the effects of risk and risk aversion on bidder behavior and DOT costs, we compare the equilibria in four cases. First, we compute the equilibrium in our example when bidders are risk averse with CARA coefficient $\gamma = 0.05$, and when bidders are risk neutral (e.g. $\gamma = 0$). To hone in on the effects of risk in particular, and not mis-estimation, we will assume that the bidders' expectations of ex-post quantities are perfectly correct (e.g. the realization of q_c^a is equal to $\mathbb{E}[q_c^a]$, although the bidders do not know this ex-ante, and still assume their estimates are noisy with Gaussian error).

Next, we compute the equilibrium in each case under the counterfactual in which uncertainty regarding quantities is eliminated. In particular, we consider a setting in which the DOT is able to discern the precise quantities that will be used, and advertise the project with the ex-post quantities, rather than imprecise estimates. The DOT's accuracy is common knowledge, and so upon seeing the DOT numbers in this counterfactual, the bidders are certain of what the ex-post quantities will be (e.g. $\sigma_c^2 = \sigma_r^2 = 0$).

	Risk Neutral Bidders	Risk Averse Bidders
Noisy Quantity Estimates	\$326.76	\$317.32
Perfect Quantity Estimates	\$326.76	\$296.26

Table 16: Comparison of Expected DOT Costs

In Table 16, we present the expected (ex-post) DOT cost in each case. This is the expectation of the amount that the DOT would pay the winning bidder $q_c^a b_c^w + q_r^a b_r^w$ at the equilibrium bidding strategy in each setting, taken with respect to the distribution of the type of the lowest (winning) bidder.[44] When bidders are risk neutral ($\gamma = 0$), the equilibrium cost to the DOT does not change when the DOT improves its quantity estimates. The reason for this is that since $\gamma = 0$, the variance term in Equation (35) is zero regardless of the level of the noise in quantity predictions. As the bidders' quantity expectations $\mathbb{E}[q_c^a]$ and $\mathbb{E}[q_r^a]$ are unchanged, the expected revenue of the winning bidder (corresponding to the expected cost to the DOT) is unchanged as well.

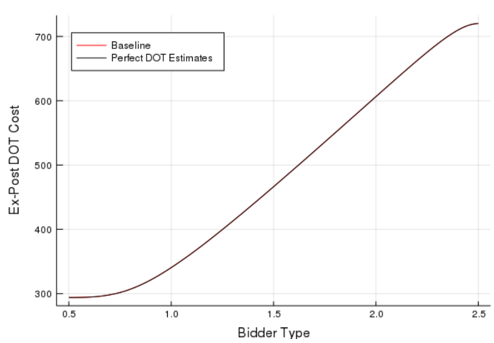

(a) Risk Neutral Bidders

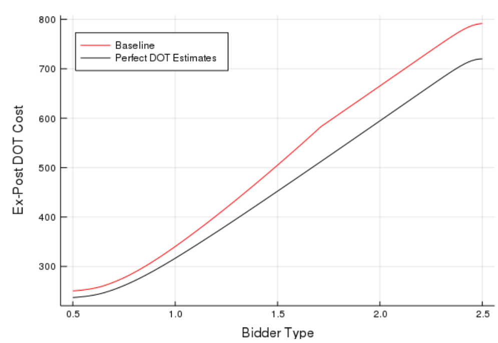

(b) Risk Averse Bidders: $\gamma = 0.05$

Figure 16: Equilibrium DOT Cost/Bidder Revenue by Bidder Type

In Figure 16a, we plot the revenue that each type of bidder expects to get in equilibrium when bidders are risk neutral. The red line corresponds to the baseline setting, in which the DOT underestimates the ex-post quantity of concrete, and overestimates the ex-post quantity of traffic cones. The black line corresponds to the counterfactual in which both quantities are precisely estimated, and bidders have no residual uncertainty about what the quantities will be. While the

[44]In order to simulate equilibria, we need to assume a distribution of bidder types. For this example, we assume that bidder types are distributed according to a truncated lognormal distribution, $\alpha \sim \text{LogNormal}(0, 0.2)$ that is bounded from above by 2.5. There is nothing special about this particular choice, and we could easily have made others with similar results.

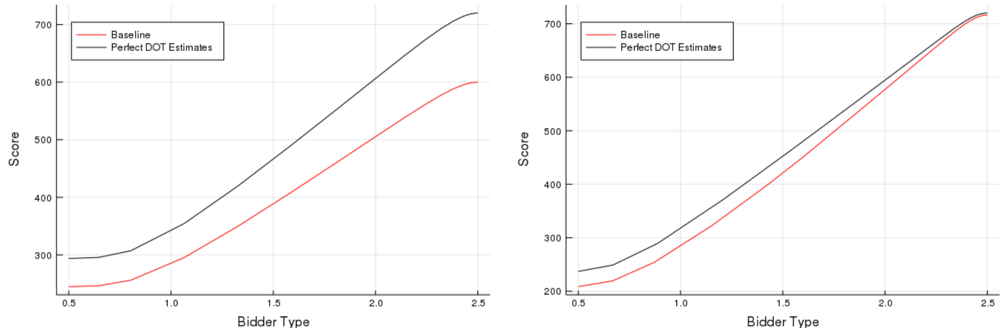

(a) Risk Neutral Bidders (b) Risk Averse Bidders: $\gamma = 0.05$

Figure 17: Equilibrium Score Functions by Bidder Type

ex-post cost to the DOT is the same whether or not the DOT quantity estimates are correct, the unit bids and resulting scores that bidders submit are different. In Figure 17a, we plot the equilibrium score for each bidder type when bidders are risk neutral. The score at every bidder type is smaller under the baseline than under the counterfactual in which the DOT discerns ex-post quantities. This is because the scores in the counterfactual correspond to the bidders' expected revenues, while the scores in the baseline multiply bids that are skewed to up-weight over-running items by their under-estimated DOT quantities. See the Appendix for a full derivation and discussion of the risk neutral case.

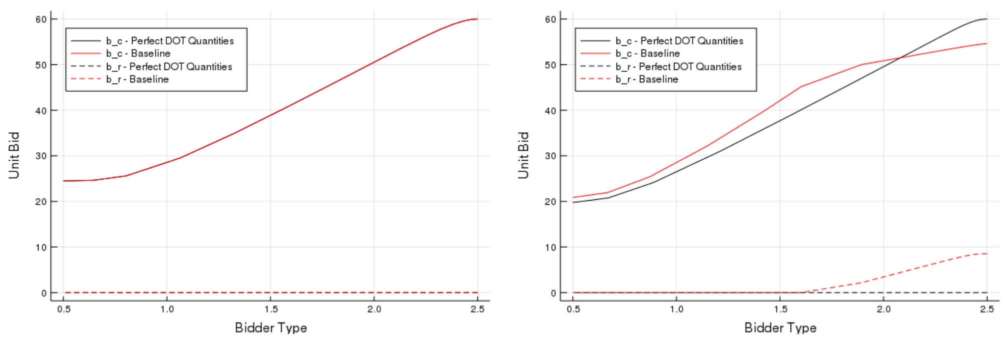

(a) Risk Neutral Bidders (b) Risk Averse Bidders: $\gamma = 0.05$

Figure 18: Equilibrium Unit Bids by Bidder Type

Figure 18a plots the unit bid that each type of bidder submits in equilibrium when bidders are risk neutral. As before, the red lines correspond to the baseline setting in which the DOT mis-estimates quantities, whereas the black lines correspond to the counterfactual setting in which the DOT discerns ex-post quantities perfectly. The solid line in each case corresponds to the unit bid for concrete $b_c(\alpha)$ that each α type of bidder submits in equilibrium. The dashed line corresponds to the equilibrium unit bid for traffic cones $b_r(\alpha)$ for each bidder type. Notably, in every case, the optimal bid for each bidder puts the maximum possible amount (conditional on the bidder's equilibrium score) on the item that is predicted to over-run the most, and \$0 on the other item. This is a direct implication of optimal bidding by risk neutral bidders, absent an external impetus to do otherwise. As noted by **?**, this suggests that the observations of *interior* or *intermediately-skewed* bids in our data, as well as in Athey and Levin's, are inconsistent with a model of risk neutral bidders. Other work, such as Bajari, Houghton, and Tadelis (2014) have rationalized interior bids by modeling a heuristic penalty for extreme skewing that represents a fear of regulatory rebuke. However, no significant regulatory enforcement against bid skewing has ever been exercised by MassDOT, and discussions of bidding incentives in related papers as well as in **?** suggest that risk avoidance is a more likely dominant motive.

In Figures 16b, 17b and 18b, we plot the equilibrium revenue, score and bid for every bidder type, when bidders are risk averse with the CARA coefficient $\gamma = 0.05$. Unlike the risk-neutral case, the DOT's elimination of uncertainty regarding quantities has a tangible impact on DOT costs. When the DOT eliminates quantity risk for the bidders, it substantially increases the value of the project for all of the bidders, causing more competitive bidding behavior. Seen another way, uncertainty regarding ex-post quantities imposes a cost to the bidders, on top of the cost of implementing the project upon winning. In equilibrium, bidders submit bids that allow them to recover all of their costs (plus a mark-up). When uncertainty is eliminated, the cost of the project decreases, and so the total revenue needed to recover each bidder's costs decreases as well. Note, also, that the elimination of uncertainty results in different levels of skewing across the unit bids of different items. Whereas under the baseline, bidders with types $\alpha > 1.6$ place increasing interior bids on traffic cones, when risk is eliminated, this is no longer the case. However, this is subject to a tie breaking rule – when the DOT perfectly predicts ex-post quantities, there are no over-runs, and so there is no meaningful different to overbid on one item over the other. The analysis of the optimal bid (conditional on a score) here is analogous to that under risk neutrality, and so we defer details to the appendix.

CARA Coeff	Baseline	No Quantity Risk	Pct Diff
0	\$326.76	\$326.76	0%
0.001	\$326.04	\$325.62	0.13%
0.005	\$323.49	\$321.41	0.64%
0.01	\$321.01	\$316.88	1.29%
0.05	**\$317.32**	**\$296.26**	**6.64%**
0.10	\$319.83	\$285.57	10.71%

Table 17: Comparison of expected DOT costs under different levels of bidder risk aversion

While the general observation that reducing uncertainty may result in meaningful cost savings to the DOT, the degree of these savings depends on the baseline level of uncertainty in each project, as well as the degree of bidders' risk aversion and the level of competition in each auction (constituted by the distribution of cost types and the number of participating bidders). To illustrate this, we repeat the exercise summarized in Table 16 over different degrees of risk aversion and different levels of uncertainty. In Table 17, we present the expected DOT cost under the baseline example and under the counterfactual in which the DOT eliminates quantity risk, as well as the percent difference between the two, for a range of CARA coefficients.[45] The bolded row with a CARA coefficient of 0.05 corresponds to the right hand column of Table 16.

A.4 Worked Out Example of Risk Neutral Bidding

Two risk-neutral bidders compete for a project that requires two types of inputs to complete: concrete and traffic cones. The DOT estimates that 10 tons of concrete and 20 traffic cones will be necessary to complete the project. However, the bidders (both) anticipate that the actual quantities that will be used – random variables that we will denote q_c^a and q_r^a for concrete and traffic cones, respectively – are distributed with means $\mathbb{E}[q_c^a] = 12$ and $\mathbb{E}[q_r^a] = 10$. We will assume that the actual quantities are exogenous to the bidding process, and do not depend on who wins the auction in any way.

The bidders differ in their private costs for the materials (including overhead, etc.): each bidder i incurs a privately known flat unit cost c_c^i for each unit of concrete and c_r^i for each traffic cone used. Thus, at the time of bidding, each bidder i expects to incur a total cost

$$\theta^i \equiv \mathbb{E}\left[q_c^a c_c^i + q_r^a c_r^i\right] = 12c_c^i + 10c_r^i,$$

if she were to win the auction. Each bidder i submits a unit bid for each of the items: b_c^i and b_r^i. The winner of the auction is then chosen on the basis of her *score*: the sum of her unit bids multiplied the DOT's quantity estimates:

$$s^i = 10b_c^i + 20b_r^i.$$

Once a winner is selected, she will implement the project and earn net profits of her unit bids, less the unit costs of each item, multiplied by the *realized* quantities of each item that are ultimately used. At the time of bidding, these quantities are unrealized samples of random variables. However, as the bidders are risk-neutral, they consider the expected value of profits to make their bidding decisions:

$$\begin{aligned} E[\pi(b_c^i, b_r^i)|c_c^i, c_r^i] &= \underbrace{\mathbb{E}\left[(q_c^a b_c^i + q_r^a b_r^i) - (q_c^a c_c^i + q_r^a c_r^i)\right]}_{\text{Expected profits conditional on winning}} \times \underbrace{\text{Prob}(s^i < s^j)}_{\text{Probability of winning}} \\ &= \left((12b_c^i + 10b_r^i) - \theta^i\right) \times \text{Prob}\left((10b_c^i + 20b_r^i) < (10b_c^j + 20b_r^j)\right). \end{aligned}$$

The key intuition for bid skewing is as follows. Suppose that the bidders' expectations of the actual

[45] That is, in the baseline, the DOT posts quantity estimates $q_c^e = 10$ and $q_r^e = 20$, while bidders predict that $\mathbb{E}[q_c^a] = 12$ and $\mathbb{E}[q_r^a] = 18$ with $\sigma_c^2 = 2$ and $\sigma_r^2 = 1$. In the No Quantity Risk counterfactual, the DOT discerns that $q_c^e = q_c^a = 12$ and $q_r^e = q_r^a = 18$, so that $\sigma_c^2 = \sigma_r^2 = 0$.

quantities to be used are accurate. Then for any score s that bidder i deems competitive, she can construct unit bids that maximize her ex-post profits if she wins the auction. For example, suppose that bidder i has unit costs $c_c^i = \$70$ and $c_r^i = \$3$, and she has decided to submit a score of \$1000. She could bid her costs with a \$5 markup on concrete and a \$9.50 markup on traffic cones: $b_c^i = \$75$ and $b_r^i = \$12.50$, yielding a net profit of \$155. However, if instead, she bids $b_c^i = \$99.98$ and $b_r^i = \$0.01$, bidder i could submit the same score, but earn a profit of nearly \$330 if she wins.

This logic suggests that the DOT's inaccurate estimates of item quantities enable bidders to extract surplus profits without ceding a competitive edge. If the DOT were able to predict the actual quantities correctly, it would eliminate the possibility of bid skewing. In order for bidder i to submit a score of \$1000 in this case, she would need to choose unit bids such that $12b_c^i + 20b_r^i = \$1000$—the exact revenue that she would earn upon winning the auction. She could still bid $b_r^i = \$0.01$, for example, but then she would need to bid $b_c^i = \$83.33$, resulting in a revenue of \$1000 and a profit of \$130 if she wins the auction. A quick inspection shows that no choice of b_c^i and b_r^i could improve her expected revenue at the same score.

It would follow that when bidders have more accurate assessments of what the actual item quantities will be – as is generally considered to be the case – bids with apparent skewing *are materially* more costly to the DOT. If the bidders were to share their expectations truthfully with the DOT, it appears that a lower total cost might be incurred without affecting the level of competition.

However, this intuition does not take into account the equilibrium effect that a change in DOT quantity estimates would have on the competitive choice of score. It is not true that if a score of \$1000 is optimal for bidder i under inaccurate DOT quantity estimates, then it will remain optimal under accurate DOT estimates as well. As we demonstrate below, when equilibrium score selection is taken into consideration, the apparent possibility of extracting higher revenues by skewing unit bids is shut down entirely.

To illustrate this point, we derive the equilibrium bidding strategy for each bidder in our example. In order to close the model, we need to make an assumption about the bidders' beliefs over their opponents' costs. Bidder i's expected total cost for the project θ^i is fixed at the time of bidding, and does not depend on her unit bids. For simplicity, we will assume that these expected total costs are distributed according to some commonly known distribution: $\theta \sim F[\underline{\theta}, \overline{\theta}]$.

By application of Asker and Cantillon (2010), there is a unique (up to payoff equivalence) monotonic equilibrium in which each bidder of type θ submits a unique equilibrium score $s(\theta)$, using unit bids that maximize her expected profits conditional on winning, and add up to $s(\theta)$. That is, in equilibrium, each bidder i submits a vector of bids $\{b_c(\theta^i), b_r(\theta^i)\}$ such that:

$$\{b_c(\theta^i), b_r(\theta^i)\} = \arg \max_{\{b_c, b_r\}} \left\{12b_c + 40b_r - \theta^i\right\} \text{ s.t. } 10b_c + 50b_r = s(\theta^i).$$

Solving this, we quickly see that at the optimum, $b_r(\theta^i) = 0$ and $b_c(\theta^i) = s(\theta^i)/10$ (to see this, note that if $b_r = 0$, then the bidder earns a revenue of $\frac{12}{10} \cdot s(\theta^i)$ whereas if $b_c = 0$, then the bidder earns a revenue of $\frac{40}{50} \cdot s(\theta^i)$.)

The equilibrium can therefore be characterized by the optimality of $s(\theta)$ with respect to the expected profits of a bidder with expected total cost θ:

$$E[\pi(s(\theta^i))|\theta^i] = \left(\frac{12}{10} \cdot s(\theta^i) - \theta^i\right) \cdot \text{Prob}\left(s(\theta^i) < s(\theta^j)\right) \tag{36}$$

$$= \left(\frac{12}{10} \cdot s(\theta^i) - \theta^i\right) \cdot \left(1 - F(\theta^i)\right), \tag{37}$$

where the second equality follows from the strict monotonicity of the equilibrium.[46]

As in a standard first price auction, the optimality of the score mapping is characterized by the first order condition of expected profits with respect to $s(\theta)$:

$$\frac{\partial \mathbb{E}[\pi(\tilde{s}, \theta)]}{\partial \tilde{s}}|_{\tilde{s}=s(\theta)} = 0.$$

Solving the resulting differential equation, we obtain:

$$s(\theta) = \frac{10}{12}\left[\theta + \frac{\int_{\theta}^{\overline{\theta}} \left[1 - F(\tilde{\theta})\right] d\tilde{\theta}}{1 - F(\theta)}\right].$$

Thus, each bidder i will bid $b_c(\theta^i) = \frac{s(\theta^i)}{10}$ and $b_r(\theta) = 0$. If bidder i wins the auction, she expects to earn a markup of:

$$E[\pi(\theta^i)] = 12 \cdot \frac{s(\theta^i)}{10} - \theta^i \tag{38}$$

$$= \frac{\int_{\theta^i}^{\overline{\theta}} \left[1 - F(\tilde{\theta})\right] d\tilde{\theta}}{1 - F(\theta^i)}. \tag{39}$$

More generally, no matter *what* the quantities projected by the DOT are – entirely correct or wildly inaccurate – the winner of the auction and the markup that she will earn in equilibrium will be the same.

In particular, writing q_c^e and q_r^e for the DOT's quantity projections (so that a bidder's score is given by $s = b_c q_c^e + b_r q_r^e$) and q_c^b and q_r^b for the bidders' expectations for the actual quantities, the equilibrium score function can be written:

$$s(\theta) = \min\left\{\frac{q_c^e}{q_c^b}, \frac{q_r^e}{q_r^b}\right\} \cdot \left[\theta + \frac{\int_{\theta}^{\overline{\theta}} \left[1 - F(\tilde{\theta})\right] d\tilde{\theta}}{1 - F(\theta)}\right]. \tag{40}$$

Suppose that $\frac{q_r^e}{q_r^b} \leq \frac{q_c^e}{q_c^b}$. Then bidder i will bid $b_r^*(\theta^i) = \frac{s(\theta^i)}{q_r^e}$ and $b_c^*(\theta^i) = 0$. Consequently, if bidder

[46] More concretely, a monotonic equilibrium requires that for any $\theta' > \theta$, $s(\theta') > s(\theta)$. Therefore, the probability that $s(\theta^i)$ is lower than $s(\theta^j)$ is equal to the probability that θ^i is lower than θ^j.

i wins, she will be paid $q_r^b \cdot b_r^*(\theta^i) = \left[\theta^i + \frac{\int_{\theta^i}^{\bar{\theta}} \left[1 - F(\tilde{\theta})\right] d\tilde{\theta}}{1 - F(\theta^i)}\right]$ as in our example.

The probability of winning is determined by the probability of having the lowest cost type, in equilibrium, and so this too is unaffected by the DOT's quantity estimates. That is, the level of competition and the degree of markups extracted by the bidders is determined entirely by the density of the distribution of expected total costs among the competitors. The more likely it is that bidders have similar costs, the lower the markups that the bidders can extract. However, regardless of whether the DOT posts accurate quantity estimates—in which case, bidders cannot benefit from skewing their unit bids at any score—or not, the expected cost of the project to the DOT will be the same in equilibrium. Therefore, a mathematically unbalanced bid, while indicative of a discrepancy in the quantity estimates made by the bidders and the DOT, is not indicative of a material loss to the government.

A.5 Discussion of Policy Inefficacy If Bidders are Risk Neutral

In the body of our paper, we present three counterfactual policy proposals: (1) reducing the latent uncertainty about item quantities; (2) switching to a lump sum auction; (3) subsidizing entry costs in order to incentivize additional entry. We show that when bidders are risk averse (and in particular, under the estimated level of risk aversion), these policies each have a significant impact on DOT spending in equilibrium. In this section, we argue formally that risk aversion is key to these results. In particular, if bidders are instead risk neutral, then the effect of all of these policies is unambiguously null in equilibrium.

The key intuition to these results is as follows. The equilibrium construction in Appendix A.4 would be almost identical with T items, rather than 2. For risk neutral bidders, the choice of bid vector that maximizes the "inner" optimization problem conditional on a score is independent of the bidder's type: at the optimum, each bidder bids her entire score (normalized by the DOT engineer's quantity projection) on the item that will over-run the most in expectation, and zero on every other item.

Thus, bidding is effectively one-dimensional, and all of the properties of standard single-item first price auctions with risk neutral bidders apply. In particular, not only would the policies to reduce uncertainty or switch to a lump sum auction be ineffective (which follows directly from the model as "risk" does not enter into bidder preferences), but so would a policy to subsidize bidders. This result is a consequence of revenue equivalence: the decrease in expected costs from the entry of an additional bidder is equal to the expected profit that this bidder would earn upon entering.[47] As such, the equilibrium number of entrants to a given auction is necessarily efficient: incentivizing an additional bidder would cost the full additional surplus that this bidder would bring. By contrast, revenue equivalence does not apply with risk averse bidders, and the additional bidder's expected utility from entering may be lower than the decrease in expected costs if she enters.

[47]See Klemperer (1999) for a fuller but still intuitive discussion of this.

B Views of Bid Skewing by Contractors and MassDOT Managers

Bid Skewing Among Contractors

The practice of *unbalanced bidding*—or *bid skewing*—in scaling auctions appears, in the words of one review, "to be ubiquitous" (Skitmore and Cattell (2013)). References to bid skewing in operations research and construction management journals date as far back as 1935 and as recently as 2010. A key component of skewing is the bidders' ability to predict quantity over/under-runs and optimize accordingly. Stark (1974), for instance, characterizes contemporary accounts of bidding:

> *Knowledgeable contractors independently assess quantities searching for items apt to seriously under-run. By setting modest unit bids for these items they can considerably enhance the competitiveness of their total bid.*

Uncertainty regarding the quantities that will ultimately be used presents a challenge to optimal bid-skewing, however. In an overview of "modern" highway construction planning, Tait (1971) writes:

> *...there is a risk in manipulating rates independently of true cost, for the quantities schedule in the bill of quantities are only estimates and significant differences may be found in the actual quantities measured in the works and on which payment would be based.*

In order to manage the complexities of bid selection, contractors often employ experts and software geared for statistical prediction and optimization. Discussing the use of his algorithm for optimal bidding in consulting for a large construction firm, Stark (1974) notes a manager's prediction that such software would soon become widespread—reducing asymmetries between bidders and increasing allocative efficiency in the industry.

> *...since the model was public and others might find it useful as well, it had the longer term promise of eroding some uncertainties and irrelevancies in the tendering process. Their elimination...increased the likelihood that fewer contracts would be awarded by chance and that his firm would be a beneficiary.*

Since then, an assortment of decision support tools for estimating item quantities and optimizing bids has become widely available. A search on Capterra, a web platform that facilitates research for business software buyers, yields 181 distinct results. In a survey on construction management software trends, Capterra estimates that contractors spend an average $2,700 annually on software. The top 3 platforms command a market share of 36% and surveyed firms report having used their current software for about 2 years—suggesting a competitive environment. Asked what was most improved by the software, a leading 21% of respondents said, "estimating accuracy", while 14% (in third place) said “bidding”.

MassDOT Challenges to Bid Skewing

Concerns that sophisticated bidding strategies may allow contractors to extract excessively large payments have led to a number of lawsuits about MassDOT's right to reject suspicious bids. The Federal Highway Administration (FHWA) has explicit policies that allow officials to reject bids that are deemed manipulative. However, the legal burden of proof for a manipulative bid is quite high. In order for a bid to be legally rejected, it must be proven to be *materially unbalanced*.[48]

> *A bid is materially unbalanced if there is a reasonable doubt that award to the bidder ... will result in the lowest ultimate cost to the Government. Consequently, a materially unbalanced bid may not be accepted.*[49]

However, as the definition for material unbalancedness is very broad, FHWA statute requires that a bid be *mathematically* unbalanced as a precondition. A *mathematically unbalanced* bid is defined as one, “structured on the basis of nominal prices for some work and inflated prices for other work.”[50] In other words, it is a bid that appears to be strategically skewed. In order to discourage bid skewing, many regional DOTs use concrete criteria to define mathematically unbalanced bids. In Massachusetts, a bid is considered mathematically unbalanced if it contains any line-item for which the unit bid is (1) over (under) the office cost estimate and (2) over (under) the average unit bid of bidders ranked 2-5 by more than 25%.

In principle, a mathematically unbalanced bid elicits a flag for MassDOT officials to examine the possibility of material unbalancedness. However, in practice, such bids are ubiquitous, and substantial challenges by MassDOT are very rare. In our data, only about 20% of projects do not have a single item breaking MassDOT's overbidding rule, and only about 10% of projects do not have a single item breaking the underbidding rule. Indeed, most projects have a substantial portion of unit bids that should trigger a mathematical unbalancedness flag. However, only 2.5% of projects have seen bidders rejected across all justifications, a handful of which were due to unbalanced bids.[51]

The Difficulty of Determining ‘Materially Unbalanced’ Bids

A primary reason that so few mathematically unbalanced bids are penalized is that material unbalancedness is very hard to prove. In a precedent-setting 1984 case, the Boston Water and Sewer Commission was sued by the second-lowest bidder for awarding a contract to R.J. Longo Construction Co., Inc., a contractor who had the lowest total bid along with a penny bid. The Massachusetts Superior Court ruled that the Commission acted correctly, since the Commission saw no evidence

[48] See Federal Acquisition Regulations, Sec. 14.201-6(e)(2) for sealed bids in general and Sec. 36.205(d) for construction specifically (Cohen Seglias Pallas Greenhall and Furman PC (2018)).

[49] Matter of: Crown Laundry and Dry Cleaners, Comp. Gen. B-208795.2, April 22, 1983.

[50] Matter of: Howell Construction, Comp. Gen. B-225766 (1987)

[51] MassDOT does not reject individual bidders, but rather withdraws the project from auction and possibly resubmits it for auction after a revision of the project spec.

that the penny bid would generate losses for the state. More specifically, no convincing evidence was presented that if the penny bid did generate losses, the losses would exceed the premium on construction that the second-lowest bidder wanted to charge (Mass Superior Court, 1984).[52] In January 2017, MassDOT attempted to require a minimum bid for every unit price item in a various locations contract due to bid skewing concerns. SPS New England, Inc. protested, arguing that such rules preclude the project from being awarded to the lowest responsible bidder. The Massachusetts Assistant Attorney General ruled in favor of the contractor on August 1, 2017.

In fact, as we show in Appendix A.2, there is a theoretical basis to question the relationship between mathematical and material unbalancedness. As we demonstrate, bid skewing plays dual roles in bidders' strategic behavior. On the one hand, bidders extract higher ex-post profits by placing higher bids on items that they predict will over-run in quantity. On the other hand, bidders reduce ex-ante risk by placing lower bids on items, regarding which they are particularly uncertain. Moreover, when bidders are similarly informed regarding ex-post quantities, the profits from predicting over-runs are largely competed away in equilibrium, but the reduction in ex-ante risk is passed on to MassDOT in the form of cost-savings.

C Solving Portfolio Problems

We present a fast, deterministic algorithm to solve the constrained quadratic programs found in the bidders' portfolio problems. Given the independence of items within a project, each problem can be represented in the form:

$$\max_{\{x_i\}_i} \sum_i a_i x_i - b_i x_i^2 \text{ subject to } \begin{cases} \sum_i q_i x_i = s \\ x_i \geq 0 \text{ for each } i \end{cases}$$

The primal formulation of this problem is

$$\max_{\{x_i\}_i} \min_v \left\{ \sum_i f_i(x) + v\left(q'x - s\right) \right\}$$

$$\text{where } f_i(x) = a_i x - b_i x^2 + p_i(x)$$

[52] In response to this case, MassDOT inserted the following clause into Subsection 4.06 of the MassDOT Standard Specifications for Highways and Bridges: "No adjustment will be made for any item of work identified as having an unrealistic unit price as described in Subsection 4.04." This clause, inserted in the Supplemental Specifications dated December 11, 2002, made it difficult for contractors to renegotiate the unit price of penny bid items during the course of construction. An internal MassDOT memo from the time shows that Construction Industries of Massachusetts (CIM) requested that this clause be removed. One MassDOT engineer disagreed, writing that "if it is determined that MHD should modify Subsection 4.06 as requested by CIM it should be noted that the Department may not necessarily be awarding the contract to the lowest responsible bidder as required." The clause was removed from Subsection 4.06 in the June 15, 2012 Supplemental Specifications.

$$\text{and } p_i(x) = \begin{cases} \infty & \text{if } x_i < 0 \\ 0 & \text{otherwise} \end{cases}$$

where v is the Lagrange multiplier on the linear constraint. By well known results, we can instead solve the dual problem:

$$\min_v \underbrace{\max_{\{x_i\}_i} \left\{ \sum_i f_i(x) + v\left(q'x - s\right) \right\}}_{g(x)}$$

The First Order Conditions of $g(x)$ are given by $\frac{\nabla_i g(x)}{\partial x_i} = a_i - 2bx_i + vq_i = 0$ and so, at the optimum:

$$x_i^* = \max\left\{ \frac{a_i + vq_i}{2b_i}, 0 \right\}.$$

Substituting x^* into the dual objective, we obtain:

$$\min_v \left\{ \sum_i h_i(x_i^*) - v\left(s\right) \right\} \tag{41}$$

$$\text{where: } h_i(x_i^*) = \begin{cases} \frac{(a_i+vq_i)^2}{2b_i} - b_i\left(\frac{a_i+vq_i}{2b_i}\right)^2 & \text{if } \frac{a_i+vq_i}{2b_i} > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Simplifying this further,

$$h_i(x_i^*) = \begin{cases} \frac{1}{4b_i}\left(a_i + vq_i\right)^2 & \text{if } \frac{a_i+vq_i}{2b_i} > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Thus, the solution to the original problem is the v_k^* that minimizes Equation (41) with k non-zero components with the form $x_i^* = \frac{a_i + v_k^* q_i}{2b_i}$. Noting that $\frac{a_i+vq_i}{2b_i} > 0 \iff v > \frac{-a_i}{q_i}$, we propose the following algorithm for solving the problem:

1. Rank $\{i\}$ in order of $\frac{-a_i}{q_i}$ (lowest to highest). Note that under this sorting, for any v, if $v \leq \frac{-a_j}{q_j}$ for some j, then $v \leq \frac{-a_k}{q_k}$ for all $k > j$. Consequently if $v \leq \frac{-a_j}{q_j}$ for some j, then $h_k(x_k^*(v)) = 0$ for all $k > j$ as well, so that we do not need to consider the contributions of elements $k > j$ in the objective.

2. For each k, let $\tilde{v}_k$ to be the value of v that minimizes Equation (41) on the interval $(\frac{-a_k}{q_k}, \frac{-a_{k+1}}{q_{k+1}}]$. Iteratively search over indices k, computing $\tilde{v}_k$, and compare them to find the global minimizer among them.

Note that for any k in Step 2., there is a closed form solution to $\tilde{v}_k$:

$$\tilde{v}_k = \arg \min_{v \in \left(\frac{-a_k}{q_k}, \frac{-a_{k+1}}{q_{k+1}}\right]} \left\{ \left[\sum_{i \leq k} \frac{1}{4b_i} \left(a_i + \tilde{v}_k q_i\right)^2 \right] - \tilde{v}_k s \right\}.$$

This is a sum of quadratics (e.g. a quadratic), and so we find the optimum by taking the FOC:

$$\tilde{v}_k^* = \min \left\{ \frac{2s - \sum_{i \leq k} \frac{a_i q_i}{b_i}}{\sum_{i \leq k} \frac{1}{b_i} q_i^2}, \frac{-a_{k+1}}{q_{k+1}} \right\}. \tag{42}$$

Thus, we can compute $\tilde{v}_k^*$ for each sequential index k with two operations by considering just the kth contribution to Equation (42). Finally, we compare each $\tilde{v}_k^*$ across indices k to find the global minimizer.

Edge Cases This algorithm above will work so long as $\frac{a_i + v q_i}{2b_i}$ is well defined – that is so long as $b_i > 0$. When there is (at least one) element i such that $b_i = 0$ (and so it is linear), the optimal solution will stop propagating v_k's as soon as it hits the first linear element in the $-a_i/q_i$ rank order. At that point (say the linear element is the kth one): $v_k = -a_k/q_k$ and $x_k = \frac{s - \sum_{i \leq k} q_i x_i^*}{q_k}$.

Adding Item-Level Constraints Suppose that we add item-specific constraints, so that our problem is:

$$\max_{\{x_i\}_i} \sum_i a_i x_i - b_i x_i^2; \text{ subject to } \begin{cases} \sum_i q_i x_i = s \\ x_i \geq r_i \text{ for each } i \end{cases}$$

where $r_i > 0$ is some (known) number for each component i.

To use our algorithm above, we simply transform x into a new variable: $y = x - r$

$$\max_{\{y_i\}_i} \sum_i a_i (y_i + r_i) - b_i (y_i + r_i)^2; \text{ subject to } \begin{cases} \sum_i q_i (y_i + r_i) = s \\ y_i \geq 0 \text{ for each } i \end{cases}$$

Simplifying, we see that this fits right into our previous framework:

$$\max_{\{y_i\}_i} \sum_i \tilde{a}_i y_i - \tilde{b}_i y_i^2 + \tilde{C}_i; \text{ subject to } \begin{cases} \sum_i q_i y_i = \tilde{s} \\ y_i \geq 0 \text{ for each } i \end{cases} \text{ where: } \begin{cases} \tilde{a}_i = a_i - 2b_i r_i \\ \tilde{b}_i = b_i \\ \tilde{C}_i = a_i r_i - b_i r_i^2 \\ \tilde{s} = s - \sum_i q_i r_i \end{cases}$$

Note that $\tilde{C}$ is a constant and so does not affect optimization.

D Model Extensions

D.1 Constant Relative Risk Aversion

Our baseline model assumed that bidders are endowed with a CARA utility function. While CARA is commonly used as a local approximation of general risk aversion—and does reasonably well at matching data in our simulations—it is likely to mischaracterize preferences when the stakes are an order of magnitude or more higher, as in our lump sum counterfactual. In this section, we discuss how our baseline model may be extended to a model of CRRA utility.

Primitives Suppose that bidders are instead risk averse with a standard CRRA utility function over their earnings from the project and a common constant coefficient of relative risk aversion γ:

$$u(\pi) = \frac{\pi^{1-\gamma}}{1-\gamma}. \tag{43}$$

As in our baseline, we will assume that all bidders in the same auction have the same level of (CRRA) risk aversion γ, but that the bidders are differentiated along their cost efficiency type α. Similarly, bidder expectations for the quantities with which different items will be needed are noisy to different degrees.

Given the functional form of the CRRA model, a model of quantity uncertainty with Gaussian noise would be intractable. However, as is common in the asset pricing literature, we can make use of several identities and approximation techniques. The first such tool is the following fact, which shows that the expected utility of a lognormal random variable has a closed form. Suppose that $\log(x) \sim \text{Normal}(\mu, \sigma^2)$. Then:

$$E[u(x)] = \frac{\exp\left[(1-\gamma)\mu + \frac{1}{2}(1-\gamma)^2\sigma^2\right]}{1-\gamma}. \tag{44}$$

In order to use this identity, we employ several assumptions. First, we assume that the ratio of ex-post to (DOT-predicted) ex-ante quantities is distributed lognormally:

$$\frac{q_t^a}{q_t^e} \sim \text{LogNormal}\left(\bar{p}_t, \sigma_t^2\right).$$

Modeling Expected Utility Considering a bidder with efficiency type α, we write the bidder's ex-post profit function:

$$\pi(\mathbf{b}) = \sum_t q_t^a \cdot (b_t - \alpha c_t) = \sum_t q_t^e \cdot \left(\frac{q_t^a}{q_t^e}\right) \cdot (b_t - \alpha c_t). \tag{45}$$

Normalizing by $(s - \alpha \sum_k q_k^e c_k)$—which we denote by $R(s)$—and reparametrizing, we can further write this:

$$\pi(\mathbf{b}) = \left(s - \alpha \sum_k q_k^e c_k\right) \cdot \left[\sum_t \left(\frac{q_t^a}{q_t^e}\right) \cdot \frac{q_t^e \cdot (b_t - \alpha c_t)}{s - \alpha \sum_k q_k^e c_k}\right] \tag{46}$$

$$= R(s) \cdot \left[\sum_t \left(\frac{q_t^a}{q_t^e}\right) \cdot \frac{q_t^e \cdot (b_t - \alpha c_t)}{R(s)}\right] \tag{47}$$

$$= R(s) \cdot \left[\sum_t w_t \exp(p_t)\right], \tag{48}$$

where we use $w_t \equiv \frac{q_t^e \cdot (b_t - \alpha c_t)}{R}$ and $p_t \equiv \log\left(\frac{q_t^a}{q_t^e}\right)$ in the latter equality.

While Equation (48) looks like a portfolio, it is not a lognormal random variable. However, taking logs and a second order Taylor approximation thereof, we achieve:

$$\log(\pi(\mathbf{b})) \approx \log(R(s)) + \log\left(\sum_t w_t\right) + \sum_t \frac{w_t p_t}{\sum_k w_k} + \frac{1}{2}\sum_t \frac{w_t p_t^2}{\sum_k w_k} - \frac{1}{2}\sum_t \sum_r \frac{w_t w_t p_t p_r}{\left(\sum w_k\right)^2}$$

To account for higher order terms, we apply a final technique from asset pricing. Using a continuous time approximation, we arrive at:

$$\log(\pi(\mathbf{b})) \approx \text{Normal}\left(\log(R(s)) + \log\left(\textstyle\sum_t w_t\right) + \sum_t \frac{w_t \bar{p}_t}{\sum_k w_k} + \frac{1}{2}\sum_t \frac{w_t \sigma_t^2}{\sum_k w_k} - \frac{1}{2}\sum_t \frac{w_t^2 \sigma_t^2}{(\sum_k w_k)^2}, \mathbf{w}'\Sigma\mathbf{w}\right)$$

where Σ is a diagonal matrix of variances, σ^2 and $w_t = q_t^e \cdot (b_t - \alpha c_t)$. Given this lognormal approximation of $\pi(\mathbf{b})$, we can now compute expected utility. Plugging in the relevant terms and simplifying, we find:

$$\begin{aligned}\log(E[u(\pi(\mathbf{b}))]) \approx \log\left(\frac{[R(s)]^{1-\gamma}}{1-\gamma}\right) + \\ \frac{1-\gamma}{[R(s)]^2} \cdot \left[\sum_t R(s) \cdot q_t^e \cdot (b_t - \alpha c_t) \cdot \left(\bar{p}_t + \frac{1}{2}\sigma_t^2\right) - \frac{\gamma}{2}\sum_t (q_t^e)^2 \sigma_t^2 \cdot (b_t - \alpha c_t)^2\right].\end{aligned} \tag{49}$$

Computing Optimal Bids As in the CARA case, the expected utility of profits upon winning in Equation (49) yields a portfolio problem. For a given score s, the bidder will choose a vector of bids $b_t(s)$ that maximize Equation (49) subject to the condition that they are all non-negative and sum to s. Given the non-negativity constraints, there is again no closed form solution for optimal unit bids. However, at a given solution, every non-zero optimal unit bid has the following form:

$$b_t^* = \alpha \cdot \left[c_t - \frac{\sum\limits_{k:b_k>0} q_k^e c_k}{q_t^e \sigma_t^2 \sum\limits_{k:b_k>0} \frac{1}{\sigma_k^2}}\right] + \frac{R(s)}{\gamma} \cdot \left[\frac{\frac{\bar{p}_t}{\sigma_t^2} + \frac{1}{2}}{q_t^e} - \frac{\sum\limits_{k:b_k>0} \frac{\bar{p}_k}{\sigma_k^2} + \frac{1}{2}}{q_t^e \sigma_t^2 \sum\limits_{k:b_k>0} \frac{1}{\sigma_k^2}}\right] + \left[\frac{s}{q_t^e \sigma_t^2 \sum\limits_{k:b_k>0} \frac{1}{\sigma_k^2}}\right]. \tag{50}$$

Estimation Our estimation procedure mirrors that of the CARA model, making adjustments as necessary for the alternative set of assumptions. First, we estimate a first stage model of log-quantity-ratio predictions and variances. Denoting an auction by n, we assume that the posterior distribution of each $\log(q^a_{t,n}/q^e_{t,n})$ is given by a statistical model that conditions on 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. In particular, we model the realization of the actual quantity of item t in auction n as:

$$\log(q^a_{t,n}/q^e_{t,n}) = \widehat{p^b_{t,n}} + \eta_{t,n}, \text{ where } \eta_{t,n} \sim \mathcal{N}(0, \hat{\sigma}^2_{t,n}) \tag{51}$$

$$\text{such that} \quad \widehat{p^b_{t,n}} = \vec{\beta}_p X_{t,n} \text{ and } \hat{\sigma}_{t,n} = \exp(\vec{\beta}_\sigma X_{t,n}). \tag{52}$$

As in the CARA case, we estimate this model with Hamiltonian Monte Carlo and use the estimates $\widehat{p^b_{t,n}}$ and $\hat{\sigma}_{t,n}$ directly in the second stage of our procedure. Plugging the first stage estimates into Equation (50), we obtain a prediction of each bidder i's unit bid for each item t given her observed score $s_{i,n}$ in auction n. We therefore adapt Assumption 1 once more and use the analogous moment conditions to the CARA case to estimate a project-wide CRRA parameter γ_n for every auction and an efficiency coefficient for every bidder-auction pair $\alpha_{i,n}$. Given the structure of CRRA, bidder wealth is a relevant part of the utility function, that influences the optimal portfolio choice. While this again makes use of several approximations, we account for wealth by augmenting $R(s) \equiv W + s - \alpha \sum_k q^e_k c_k$ throughout our model. We allow wealth W to vary from project to project, and estimate it by projecting onto the matrix of average project-bidder characteristics X_n. We present our estimates in Table 18 below. Notably, our estimates of α are higher than in the CARA case, so that implied markups are negative in many cases. However, the CRRA coefficients are mostly between 0.6 and 0.8, which is within the region found in related studies. See pg 133 of Campo, Guerre, Perrigne, and Vuong (2011) for a discussion.

Parameter	Mean	SD	25%	50%	75%
α	1.215	0.163	1.088	1.185	1.353
γ	0.704	0.101	0.638	0.692	0.756
W	243.228	135.568	140.988	235.571	333.681

Table 18: CRRA Estimate Summary Statistics Across All Projects

Computing Equilibrium Outcomes We construct equilibria analogously to the CARA case. Here the analog of Equation (18) is given by:

$$\begin{aligned} \mathrm{EU}(\sigma(\alpha_i), \alpha_i) = \exp\left[(1-\gamma) \cdot V(\sigma(\alpha_i), \alpha_i)\right] \cdot (1 - F(\sigma^{-1}(\sigma(\alpha_i))) \\ + W^{1-\gamma} \cdot F(\sigma^{-1}(\sigma(\alpha_i)), \end{aligned} \tag{53}$$

CF Type	Δ Cost Units	Mean	SD	25%	50%	75%
Lump Sum	%	-45.5	39.2	-50.2	-34	-22.2
Lump Sum	$	-845,095	741,921	-1,149,156	-661,338	-341,346
No Risk (Correct q)	%	-19.8	18.1	-26.2	-14.4	-7.6
No Risk (Correct q)	$	-307,528	248,509	-421,562	-242,390	-147,097
No Risk (Estimated q)	%	$-41.7^{\dagger}$	$61.2^{\dagger}$	-45.7	-24.4	-14.9
No Risk (Estimated q)	$	$-465,041^{\dagger}$	$315,305^{\dagger}$	-595,512	-368,224	-234,416

Note: † refers to samples truncated by 5% to exclude extreme values

Table 19: Summary of Counterfactual Changes in Expected DOT Spending

where

$$V(\sigma(\alpha_i), \alpha_i) = \log\left(R(\sigma(\alpha_i), \alpha_i)\right) + \frac{1}{R(\sigma(\alpha_i), \alpha_i)} \cdot \left[\sum_t q_t^e \cdot (b_t(\sigma(\alpha_i)) - \alpha_i c_t) \cdot \left(\bar{p}_t + \frac{1}{2}\sigma_t^2\right)\right] - \frac{\gamma/2}{R(\sigma(\alpha_i), \alpha_i)^2} \cdot \left[\sum_t (q_t^e)^2 \sigma_t^2 \cdot (b_t(\sigma(\alpha_i)) - \alpha_i c_t)^2\right], \quad (54)$$

and $b_t(\sigma(\alpha_i))$ is chosen according to the portfolio problem, with $R = W + \sigma(\alpha_i) - \alpha_i \sum_r q_r^e c_r$.

The equilibrium function $\sigma(\alpha)$ is thus given by the solution to the differential equation:

$$\sigma'(\alpha) = \frac{f(\alpha)}{1 - F(\alpha)} \frac{1}{(1-\gamma)\frac{\partial V(\sigma(\alpha),\alpha)}{\partial s}} \left[1 - \frac{W^{1-\gamma}}{\exp\left[(1-\gamma) \cdot V(\sigma(\alpha), \alpha)\right]}\right] \quad (55)$$

subject to the boundary condition that the highest α type is indifferent between participating in the auction or not: $V(\sigma(\overline{\alpha}), \overline{\alpha}) = \log(W)$. We present summary results for the lump sum and no-risk counterfactuals in Table 19.

D.2 Asymmetric Types

Our baseline model assumes that bidders are entirely homogeneous except for their private efficiency type α. While this greatly simplifies the analysis of counterfactual policies, our estimation strategy is able to accommodate models with multidimensional bidders fairly easy. In this section, we discuss an extension to a model in which bidding firms of different sizes may have different CARA coefficients in each auction.

Modeling and Estimation While the focus of our paper is on the portfolio risk that bidders face *within* an auction, their exposure to that risk may depend on how large a role a given auction plays in their total profits. That is, firms that participate in many auctions and have larger portfolios may be less sensitive to the risk in a given project than firms that rely on a small number of projects. To account for this, we split our 25 unique bidders by the number of auctions that they participated in throughout our data. This splits the bidders roughly evenly: 14 firms are labeled as "frequent" bidders because they participated in 50 or more auctions, whereas the rest are "rare". However, the distributions of participation are quite different: the average "frequent" firm participated in 123 auctions, whereas the average "rare" firm participated in 20.

As each bidder's choice of optimal unit bids conditional on her equilibrium score does not depend on the types of any other bidders, Equation (6) and its empirical analogue remain unchanged if bidders have different values of γ within an auction. Thus, for the purpose of estimation, we need only modify the empirical model for γ in Section 6 to allow for variation by bidder size type:

$$\gamma_{i,n} = \gamma_{g(i),n} + \vec{\beta}_{\gamma} X_n, \tag{56}$$

where $g(i) \in \{\mathrm{f}, \mathrm{r}\}$ denotes whether the bidder is a frequent bidder.

Parameter	Bidder Type	Mean	SD	25%	50%	75%
α	Frequent	0.946	0.231	0.846	0.978	1.092
α	Rare	0.944	0.235	0.777	0.976	1.127
γ	Frequent	0.053	0.043	0.024	0.043	0.066
γ	Rare	0.074	0.092	0.025	0.051	0.086

Table 20: 2-Type Estimate Summary Statistics Across All Projects

Our results, summarized in Table 20, show that frequent bidders are not generally more cost efficient than rare bidders. While the median cost multiplier is about the same in both groups, the 25th percentile of rare firms is about 8% more efficient than the corresponding percentile of frequent firms, but the 75th percentile of rare firms is 3% less efficient than frequent firms. On the other hand, we find that frequent bidders exhibit slightly less risk aversion than rarer bidders. Neither of these results are surprising: frequent bidders likely specialize in work for the DOT. As such, they participate in auctions where they do not have a particularly strong efficiency advantage. But their risks are also spread across more auctions, and so the weight of uncertainty in any particular project is lower.

Computing Equilibrium Outcomes While the estimation procedure is largely unchanged, computing equilibria when types are multidimensional is more demanding. We assume that bidder efficiency types α are drawn IID from the auction-wide distribution, and that bidder-type-specific CARA coefficients are publicly known. As such, there is an asymmetric equilibrium in monotonic strategies such that each bidder type bids according to the monotone function $\sigma_g : [\underline{\alpha}, \overline{\alpha}] \to \mathbb{R}$.[53] Under a set of candidate strategies, the probability that i will win the auction under bid s is given by:

$$\prod_{j \neq i} [1 - H_j(s)] = \prod_{j \neq i} [Pr(s < \sigma_j(\alpha_j))] \tag{57}$$

$$= \prod_{j \neq i} \left[1 - F(\sigma_j^{-1}(s))\right] \tag{58}$$

$$= \prod_{j \neq i} [1 - F(\varphi_j(s))] \tag{59}$$

where the final equality relabels the inverse bid function for convenience. Applying logic analogous to our baseline specification to this more general framework, we derive the following set of differential equations:

$$\frac{\partial \varphi_f(\tilde{s})}{\partial s} = \frac{1 - F(\varphi_f(\tilde{s}))}{f(\varphi_f(\tilde{s}))} \cdot \left[\frac{1}{(N_f + N_r - 1)} \left(N_r \cdot \frac{\frac{\partial}{\partial s} V(\varphi_r(\tilde{s}))}{V(\varphi_r(\tilde{s}))} - (N_r - 1) \cdot \frac{\frac{\partial}{\partial s} V(\varphi_f(\tilde{s}))}{V(\varphi_f(\tilde{s}))}\right)\right]$$

$$\frac{\partial \varphi_r(\tilde{s})}{\partial s} = \frac{1 - F(\varphi_r(\tilde{s}))}{f(\varphi_r(\tilde{s}))} \cdot \left[\frac{1}{(N_f + N_r - 1)} \left(N_f \cdot \frac{\frac{\partial}{\partial s} V(\varphi_r(\tilde{s}))}{V(\varphi_r(\tilde{s}))} - (N_f - 1) \cdot \frac{\frac{\partial}{\partial s} V(\varphi_f(\tilde{s}))}{V(\varphi_f(\tilde{s}))}\right)\right]$$

$$\text{where} \quad V(\varphi_g(\tilde{s})) = 1 - \exp(-\gamma_g CE(\tilde{s}, \gamma_g, \varphi_g(\tilde{s}))) \tag{60}$$

$$\text{and} \quad CE(\tilde{s}, \gamma_g, \varphi_g(\tilde{s})) = \sum_{t=1}^{T} q_t^b \left(b_t^*(\tilde{s}) - \varphi_g(\tilde{s}) c_t\right) - \frac{\gamma_g \sigma_t^2}{2} \left(b_t^*(\tilde{s}) - \varphi_g(\tilde{s}) c_t\right)^2. \tag{61}$$

Here, the optimal bids $b_t^*(\tilde{s})$ are derived from the solution to the bidder's portfolio problem in Equation (5), as in the baseline model. This set of differential equations characterizes an equilibrium subject to several boundary conditions. The particular boundary conditions that apply to a given auction depend on the auction format, the number of bidders of each frequency type that participate, and the magnitude of the difference between the bidders' CARA coefficients.

There are three possible cases. The "standard" case[54] requires that the highest and lowest α types of both frequency groups submit the same score, and that the highest α type of the more risk averse group (who has no chance of winning) earn a certainty equivalent of zero. This applies when there is only one bidder of the less risk averse group, or if project portfolio risk is irrelevant as

[53]The existence of an equilibrium follows from Reny and Zamir (2004). Uniqueness may not be guaranteed and so our construction procedure may be thought of as imposing an equilibrium selection criterion.

[54]See Hubbard and Paarsch (2014) for a comprehensive survey.

in the no-risk counterfactual. However, as noted in Hubbard and Kirkegaard (2019), the standard conditions do not generally induce an equilibrium in an asymmetric auction with more than 2 bidders when the bidders have different supports for their value distributions. In our case, while both frequency groups have the same support of cost efficiency types α, their different CARA coefficients induce different (overlapping) supports on the certainty equivalents of their portfolios at each score.

As such, when there are two or more bidders of each frequency group, different boundary conditions at either end apply. Suppose, for example that $\gamma_f < \gamma_r$, and let $\overline{s_f}$ be the score that generates a zero certainty equivalent for the least competitive (highest α) frequent bidder. No bidder submitting a higher score has a chance of winning, and so in equilibrium, any rare bidder submitting $\overline{s_f}$ or higher earns a zero certainty equivalent as well. Thus, there is some cutoff $\overline{a_f}$ such that every rare bidder with $\alpha > \overline{a_f}$ submits a score that provides her a certainty equivalent of zero. On the left hand boundary, there are two possibilities. If the certainty equivalent distributions of the frequent and rare groups are sufficiently similar, then the lowest α types of each group will submit the same score as in the standard case. However, if the distributions are too far apart, the bidding functions might be bifurcated as in Hubbard and Kirkegaard (2019). That is, there may be a set of α types from the frequent group that compete against each other at scores too low for rare bidders to be willing to participate. In this case, there are also two different starting scores $\underline{s_f} < \underline{s_r}$ such that $\varphi_f(\underline{s_f}) = \underline{\alpha}$ and $\varphi_f(\underline{s_r}) = \underline{\alpha}$.

To compute equilibria for each auction, we solve a boundary value problem with a shooting algorithm based on Hubbard and Paarsch (2014). To account for bifurcation, we allow for two regions in the ODE solution: a homogeneous region in which only one group of bidders compete against each other, and a heterogeneous region that follows the equations described in this section. We find the beginning of the heterogeneous region by iteratively checking for the first score at which both groups of bidders have a type willing to participate in bidding.[55] While shooting methods are known to be highly sensitive to the step size of their integration method, we found that a modified shooting method with high-step-size Euler integration worked most consistently for our problem. However, as the number of steps required for convergence increases with numerical complexity, we focused on a subset of auctions for comparison with our baseline model.

In Table 21, we present a comparison of our main counterfactuals under the baseline model and under the asymmetric model for a sample of auctions in our data. The sample accounts for roughly half of the projects in our dataset and is similar in distribution to projects of the Bridge Reconstruction and Rehabilitation category. Among the projects included, the average γ in our baseline (homogeneous) specification is 0.49 (with standard deviation 0.083 and median 0.029). By contrast, the average frequent type γ is 0.040 (s.d., 0.065; median, 0.028) and the average rare type γ is 0.048 (s.d., 0.048; median, 0.027).[56]

[55]This procedure is similar to the check for "active" bidders in Somaini (2020).

[56]We chose projects to include in the sample on the basis of computational efficiency. This tends to select for projects with lower overall risk aversion and costs, as these projects exhibit more stable numerical

In general, we find that the results between the 1-type and 2-type cases are within a few percentage points of each other. For the median auction in this sample, the 2-type model predicts a risk premium that is 0.4 percentage points higher and a lump sum cost that is 0.2 percentage points lower than the 1-type model. A similar pattern holds at the average and every quartile as well. This suggests that the value of lowering risk by using a scaling format or by reducing uncertainty may be a bit lower if bidders are asymmetric, but that our baseline 1-type model is capturing the overall magnitude of the counterfactual effects well.

CF Type	Outcome	Mean	SD	25%	50%	75%
No Risk (Correct q)	% DOT Savings 1-Type	17.8	11.9	10.3	16.9	22.5
No Risk (Correct q)	% DOT Savings 2-Type	18.2	11.9	10.9	17.3	23.5
No Risk (Correct q)	$ DOT Savings 1-Type	217,539	200,062	80,474	161,121	293,246
No Risk (Correct q)	$ DOT Savings 2-Type	225,817	213,274	85,708	169,542	298,545
Lump Sum (Correct q)	% DOT Savings 1-Type	-90.7	245.8	-77.0	-21.6	4.1
Lump Sum (Correct q)	% DOT Savings 2-Type	-84.0	218.4	-73.6	-21.4	3.2
Lump Sum (Correct q)	$ DOT Savings 1-Type	-1,696,121	7,843,739	-929,132	-242,667	13,323
Lump Sum (Correct q)	$ DOT Savings 2-Type	-1,600,511	6,827,165	-926,175	-222,249	8,240

Table 21: Comparisons of 1-type vs 2-type Counterfactuals for a Sample of Auctions

behavior at the tails. The total sample is 248 projects. However as we were not able to get convergence on all projects in all auction formats, our comparison includes 237 projects in the no risk counterfactual and 204 projects in the lump sum counterfactual.