Scaling Auctions as Insurance: A Case Study in Infrastructure Procurement

Valentin Bolotnyy [*]   Shoshana Vasserman [†]

September 17, 2019

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 bidders' risk aversion, the risk in each auction, and the distribution of bidders' private costs. 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. Email: svass@stanford.edu. We are indebted to 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 Nick Boffi, Chris Coey, and Chris Rackauckas—with a special thanks to Jim Savage and Charles Margossian—all of whose support and advice was instrumental. We are grateful to Steve Poftak and Jack Moran, Frank Kucharski, Michael McGrath, and Naresh Chetpelly, among other generous public servants at MassDOT. Their help was invaluable. We thank Eric Maskin, Ali Yurukoglu, Andrei Skrzypacz, Shengwu Li, Ashvin Gandhi and Cory Smith for their many insightful and thought provoking comments.

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 might 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 36 other states, MassDOT uses a *scaling auction*, whereby 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 bear 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 given the distribution of 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. Equilibrium unit bids trade off higher expected profits from high bids on items predicted to over-run against higher risk from under-bidding on other items in order to keep to the equilibrium score.

Using a detailed data set obtained through a partnership with MassDOT, we establish the patterns of bidding behavior that motivate our approach and estimate the key parameters that drive bidding under our model. 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

[1] According to the CBO, infrastructure spending 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.

quantity of each item, a blue book 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 on scaling auctions, we show that contractors *skew* their bids, placing high unit bids on items they believe will over-run the DOT quantity estimates and low unit bids on items they believe will 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.[2]

Furthermore, our data suggest that contractors are risk averse. As noted in Athey and Levin (2001), risk neutral bidders should 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.[3] We show that while contractors bid higher on items predicted to overrun, holding all else fixed, they also bid lower on items that are more uncertain. This suggests that contractors optimize not only with respect to expected profits, but also with respect to the risk that any given expectation will turn out to be wrong.

We estimate a structural model for uncertainty and optimal bidding in our data. 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 using the heterogeneity in optimal unit bids at the auction-bidder-item level to estimate each bidder's private cost for completing a project that she bid on, as well as the degree of her risk aversion.

We use our structural estimates to simulate the equilibrium outcomes of counterfactual DOT policies akin to those used in other forms of public procurement. First, we consider a counterfactual in which the DOT invests in producing more accurate item quantity estimates. To evaluate the maximum possible impact, we consider the extreme case in which the DOT is able to completely eliminate uncertainty about the realization of item quantities. Second, we consider a counterfactual in which the DOT requires bidders to commit to (part of) their payment at the time of bidding. To retain the bidding structure of the scaling auction, we consider a mechanism in which the DOT compensates bidders for $\mu \in [0, 1]$ times the DOT engineer's estimated quantity and $1 - \mu$ times the realized quantity of each item used. When $\mu = 1$, this is equivalent to a *lump sum* auction—the procurement mechanism currently used

[2]Similar observations were made in the context of timber (Athey and Levin (2001)) and highway (Bajari, Houghton, and Tadelis (2014)) procurement auctions.

[3]In a few rare instances, the DOT responded to suspicious bids by scrapping the auction all together and revising the specification for the project before putting it up for auction again. In these instances, the same bidders were able to participate, and so any cost incurred was minimal.

by the subway and transit division of the DOT.[4] Finally, we consider a policy of incentivizing additional bidder entry by subsidizing the cost of putting together a bid for all participants.[5]

Risk aversion and uncertainty have significant implications for the efficacy of these policies. When bidders are risk neutral, we show theoretically that none of these policies have any impact on equilibrium spending. Risk neutral bidders compete with respect to their expected profits alone and are indifferent to variance in their estimates. Thus, their bids sum to the same amount (ex-post) with or without uncertainty, and under a scaling auction, a lump sum auction, or anything in between. By contrast, risk averse bidders internalize a utility cost from uncertainty. High variance and lump sum auction rules both increase bidders' exposure to profit risk. Risk averse bidders must therefore submit higher overall bids in order to insure themselves.[6] Similarly, in contrast to the risk neutral case, risk averse bidders undervalue the uncertain benefit of entering an auction relative to the certain cost of putting together a bid. Thus, additional bidders may profitably be compelled to submit bids by an up-front subsidy.

Key to the empirical assessment of these policies, our model allows us to predict equilibrium bids at the item-bidder-auction level. To estimate bidder types, we construct moments using the second ("inner") condition of the equilibrium characterization, evaluated at the observed ("outer") score. This allows us to leverage heterogeneity across the different unit bids (67 in the median auction) submitted by each bidder in each auction in order to estimate a bidder-auction specific cost type, as well as a common coefficient of absolute risk aversion.[7]

This approach enables our counterfactual policy evaluations in two ways. First, it allows us to compute equilibrium bids for each type of bidder, against opponents whose private costs are believed to be drawn from a distribution of estimated costs in each auction. The equilibrium maps each bidder cost type to a competitive score and the vector of unit bids that sum to this score and maximize the bidder's certainty equivalent of winning under the given auction characteristics. Second, it allows us to compute the expected payment that the DOT would make to the winning bidder under each policy. Each counterfactual setting implies a different value for each possible vector of bids (and for each possible score) for each bidder. Thus, each counterfactual results in a different composition of item unit bids in equilibrium, and consequently a different expected payment by the DOT.

---

[4]A *lump sum* auction is one in which bidders bid a total project price at the start and are responsible for all realized costs.

[5]We discuss additional counterfactual policies in the online appendix.

[6]This analysis precludes ex-post hold up problems and the like, in which the bidder might sue the DOT for additional compensation. Considerations of this sort would further increase the costs to the DOT, and so our analysis serves as a conservative estimate of the total effect.

[7]By contrast, previous papers develop estimators using the "outer" first order condition and use only heterogeneity at the bidder-auction level.

Toward the first counterfactual, we compare the DOT's average expected payment under the baseline (status quo) equilibrium to the equilibrium that would arise if all uncertainty regarding item quantities were eliminated. The difference between these expected payments offers a measure of the cost that the DOT incurs from the existence of uncertainty. As such, it constitutes the maximum amount that the DOT may be willing to pay in order to reduce uncertainty.[8] We find that the DOT's cost in the counterfactual auction with no uncertainty is only $2,039—or 0.68%—lower, on average, than in the baseline auction.

This estimate reflects the changes in both risk and prediction that are induced by the counterfactual: by eliminating uncertainty, the DOT endows bidders with the exact right quantities needed to extract the maximum payment from their bids. To isolate the effect of risk itself, we repeat this counterfactual exercise under the assumption that bidders' quantity predictions are correct in the baseline as well. In this case, there is no bidder mis-optimization to correct in the baseline and so the DOT strictly saves money from eliminating risk: $173,237 (13.67%) on average.

Toward the second counterfactual, we compute the equilibrium of an alternative mechanism in which bidders' reimbursements are fixed at the engineer's quantity estimates (regardless of the quantities that are ultimately used). As payments to the winner are fixed at the time of bidding, this is payoff equivalent to a lump sum auction and there is no incentive to skew.[9] However, this effectively shifts all risk from the DOT onto the bidders: if an item comes to be used at a much higher quantity than had been anticipated, the winning bidder is liable to procure it without reimbursement.[10] For risk averse bidders, this additional risk is particularly costly, and so they submit much higher, less aggressive bids. We estimate that switching to a lump sum auction would increase DOT spending by 133% on average (85% on median).[11] The losses do not scale linearly with the amount of risk sharing, how-

[8]Programs to reduce uncertainty might involve investments in training and technological support tools for DOT engineers and designers. The Federal Highway Administration has a number of initiatives toward this end (DeWitt et al. (2005)).

[9]Lump sum auctions may also be preferred when contractors determine the materials to be used. In this case, lump sum auctions incentivize the winner to be economical. The Massachusetts Bay Transit Authority switched from scaling auctions to lump sum auctions in the mid-2000s. In conversations with the DOT, we were told that a major reason was that subway construction projects are much more difficult to spec out prior to the start of construction (and so, more difficult to minimize contractors' agency in what/how much is used) than the road and bridge construction projects in our study.

[10]All changes to the original quantity spec must be approved by an on-site DOT manager or engineer, and contractors' ability to influence the item quantities used is very limited. In line with conversations at the DOT, we interpret deviations from the engineer's spec as stemming from unanticipated, inherent project needs. This precludes consideration of moral hazard (by which winning contractors use higher quantities of items that they submitted high unit bids for). While we cannot empirically rule out the possibility of moral hazard, we believe that this reasonably approximates the projects in our study.

[11]The distribution is quite fat tailed and so the median is more representative than the mean. See Figure 10a for the full distribution of outcomes and Figure 27 for outcomes across a range of CARA coefficients.

ever. We estimate that DOT spending would increase by only 6.83% on average (3.64% on median) under an intermediate auction design in which the winning bidder is paid her unit bid multiplied by a 50-50 split of the engineer's estimate and realized quantity for each item.

Finally, we compute the counterfactual equilibrium of each auction in our sample such that the number of participating bidders is increased or decreased by one. We measure the value of increasing competition by comparing the DOT's average expected spending when the number of bidders is increased by one to that in the status quo. We then measure the cost of incentivizing an additional bidder to participate by computing bounds on the expenses that each bidder expects to incur in order to prepare a bid.[12] We estimate that adding an additional bidder to each auction would result in an average DOT saving of $88,562 (8.72%), but the average cost of entry is less than $2,600.[13] Thus, we argue that a policy aimed at incentivizing higher bidder participation could produce substantially larger savings to the DOT than fundamental changes to the auction environment such as those discussed in our first two counterfactuals.[14]

2 Related Literature

Our paper follows a rich literature on strategic manipulation in scoring auctions and is closest in spirit to Athey and Levin (2001) and Bajari, Houghton, and Tadelis (2014). Studying data from US timber auctions, Athey and Levin (2001) demonstrate evidence of bid skewing and argue that the bids they observe are indicative of bidders who are (similarly) better informed than the auctioneer—and are risk averse. Bajari, Houghton, and Tadelis (2014) ("BHT") study a setting similar to ours: scaling auctions used to procure highway paving contracts in California. However, the study's main focus is on the role of renegotiation in generating adaptation costs, and they estimate a model of risk neutral bidding with ex-post negotiation to evaluate this.[15]

[12]Bounds for these expenses—which we refer to as the cost of entry—are given by the expected returns to participating in each auction such that no more than the (status quo, observed) equilibrium number of bidders could profitably enter, but any fewer than the observed number of bidders would leave money on the table.

[13]The average lower bound is $2,336, while the average upper bound is $2,583. The 25th, 50th and 75th percentiles of the number of bidders in each auction are 4, 6 and 9 respectively.

[14]Our counterfactuals also bear implications for project design. Designing smaller, more compact projects may facilitate reducing the amount of uncertainty about item quantities. This would reduce the amount of risk involved for each segment of work, increasing the value of winning the contract to implement it and leading to more aggressive bids if the same bidders participate. However, it would decrease the value of entering any given auction, as winning would guarantee a smaller amount of work. If bidders incur a substantial fixed cost of entering each auction (regardless of size), this may result in less entry and fewer competitive bids overall. Our results suggest that the latter concern may dominate in practice.

[15]By contrast, there is essentially no renegotiation in our setting, and we focus on the role of auction design in exposing bidders to profit risk given exogenous uncertainty regarding item quantities.

Several recent papers have exploited variation in state-specific DOT policies to assess directions for cost-cutting. De Silva, Dunne, Kosmopoulou, and Lamarche (2016) apply a framework similar to BHT to examine the effects of a Texas DOT policy to reduce the scope of project changes and curb renegotiations. Luo and Takahashi (2019) study the Florida DOT, in which a project manager decides whether to use a scaling or lump sum auction after each project has been designed.[16]

More generally, our paper relates to the literature on auctions that feature multidimensional bids or risk aversion. Che (1993) and Asker and Cantillon (2008) study the equilibria of scoring auctions—of which scaling auctions are a particular example—with risk neutral bidders. 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 *optimal* mechanism by applying the characterization in Maskin and Riley (1984) and Matthews (1987) to our framework.[17] Our paper also relates to the literature on structural estimation in auctions. Most closely related are papers on auctions with risk aversion—such as Guerre, Perrigne, and Vuong (2009) and Campo (2012)—and on share auctions, such as Hortaçsu and McAdams (2010) and Kastl (2011).[18]

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 the engineers' estimates of the quantity of each item that will be needed, and a market unit rate for its cost.[19] The itemized list of quantities is then advertised to prospective bidders.[20]

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

[16]Luo and Takahashi find additional reduced form evidence of risk aversion using variation in the uptake of lump sum and scaling auctions. They then estimate a model of risk averse bidding to conclude that lump sum auctions are preferable to scaling auctions for projects with lower cost risk and vice versa. While there are some notable differences in our modeling and identification strategy, Luo and Takahashi's conclusion is consistent with our results.

[17]These papers study a standard independent private values (IPV) setting with risk averse bidders.

[18]Most recently, Häfner (2019) studies discriminatory share auctions with risk averse bidders. While his setting is quite different from ours, he also uses the multidimensionality of bids in a discriminatory share auction for identification.

[19]Labor costs are not bid as a separate item, and are expected to be integrated into each item's unit bid.

[20]MassDOT's estimate of market rates are not advertised to prospective bidders, and are used primarily for internal budgeting purposes.

work required, given their staff and equipment. Notably, it generally does not depend on past performance.[21] In order to submit a bid, a contractor posts a per-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 36 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 the rights to implement 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%).[22]

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.[23] MassDOT has entertained other proposals to curtail bid skewing, such as a 2017 push to require a minimum unit price on each item. However, this proposal was defeated after bidder protests.[24]

4 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. The DOT 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 $\sigma_r^2 = 1$.[25] We assume that the actual quantities are exogenous to

[21] DOT resident engineers are asked to submit an evaluation of the contractor at the end of each project. In extreme cases, a particularly low score can limit the contractor's ability to prequalify for future projects. However, assessment scores are often not submitted, and no score that would impact prequalification is observed in our data.

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

[23] See Section A.2 in the Online Appendix for a detailed discussion.

[24] We evaluate this proposal using our estimates in Section A.1 of the Online Appendix.

[25] As we discuss in Section 6, 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

the bidding process, and do not depend on who wins the auction in any way. Furthermore, we assume that the bidders' expectations are identical across both bidders.[26]

The bidders differ in their private costs for implementing the project. They have access to the same vendors for 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 the market-rate cost estimate for each item: $c_c = \$8$/ton for each ton of concrete and $c_r = \$12$/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.

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

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.

[26]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.

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:[27]

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

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

$$\underbrace{\mathbb{E}[q_c^a] \cdot (b_c^i - \alpha^i c_c) + \mathbb{E}[q_r^a] \cdot (b_r^i - \alpha^i c_r)}_{\text{Expection of Profits}} - \underbrace{\left[\frac{\gamma \sigma_c^2}{2} \cdot (b_c^i - \alpha^i c_c)^2 + \frac{\gamma \sigma_r^2}{2} \cdot (b_r^i - \alpha^i c_r)^2 \right]}_{\text{Variance of Profits}}. \quad (1)$$

Furthermore, as we discuss in Section 6, 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 $\text{CE}(\{b_c, b_r\}, \alpha^i, \mathbf{c}, \mathbf{q}^a)$, 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 s, bidder i considers the set of bid vectors that sum up to s under the DOT's score formula, and chooses the one that would maximize her certainty equivalent of profits conditional on winning. She then chooses the score $s^*(\alpha^i)$ that maximizes her total expected utility.

The level of skewness in equilibrium bids depends not only on the DOT's quantity estimates and the bidders' expectations of what the actual quantities will be, but also on the bidders' uncertainty regarding each prediction, the coefficient of bidders' risk aversion and the level of competition in the auction (determined by the distribution of bidder types

[27] See section 6 and Appendix A for a detailed derivation.

and the number of participants).[28] As an illustration, suppose that the common CARA coefficient is $\gamma = 0.05$, and consider a bidder in this auction who has type $\alpha^i = 1.5$.[29]

First, suppose that the bidder has decided to submit a total score of \$500. There are a number of ways in which the bidder could compose this score. For instance, she could bid her cost on concrete, $b^i_c = \$12$, and a dollar mark-up on traffic cones: $b^i_r = (\$500 - \$12 \times 10)/20 = \$19$. Alternatively, she could bid her cost on traffic cones, $b^i_r = \$18$, and a two-dollar mark-up on traffic cones: $b^i_c = (\$500 - \$18 \times 20)/10 = \$14$. Both of these bids would result in the same score, and so give the bidder the same chance of winning the auction. However, they yield very different expected utilities. Plugging each set of bids into Equation (1), 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 that sum to a score of \$500 are $b^i_c = \$47.78$ and $b^i_r = \$1.12$, yielding a certainty equivalent of \$87.98.

Intuitively, the equilibrium bids at a given score optimize the trade-off between increasing expected profits and increasing risk. The bidder predicts that concrete will over-run in quantity—she predicts that 12 tons will be used, whereas the DOT estimates 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), 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. Consequently, the bidder optimizes by using most of the balance of her score on concrete, bidding far below cost on traffic cones.

However, the incentive to bid higher on items projected to over-run is counteracted by an incentive to keep the variance term low. When the coefficient of risk aversion is relatively high or when the variance of an item's ex-post quantity distribution is high, the variance-minimization incentive may overpower expected profit-maximization. In this case, the bidder

[28] By contrast, if bidders are risk neutral, then skewness is determined by the ratio of each item's DOT quantity estimate and the bidders' expectation for its actual quantity: whichever item has the smallest ratio is given the maximal bid resulting in each bidder's score, while all other items are bid at 0. See Online Appendix A.3 for an extended discussion of how equilibrium bids, scores and costs differ between auctions with risk averse bidders and with risk neutral bidders. See Online Appendix A.4 for a worked out derivation of the example with risk neutral bidders.

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

would optimally bid close to cost—even on an item that is predicted to greatly overrun.

Moreover, the variance component of the bidder's certainty equivalent is quadratic in unit mark-ups, and so highly skewed bids are increasingly costly at higher scores. The optimal unit bid for traffic cones under a score of \$500 is very small: $b_r^i = \$1.12$. However, under a score of \$1,000, the optimal bid is much higher: $b_r^i = \$23.33$. The reason for this is that a low bid on traffic cones implies a high bid on concrete. A high mark-up on concrete increases the bidder's certainty equivalent linearly through the expected profit term, but decreases it at a quadratic rate through the variance term. Thus, there is a finite maximal mark-up for each item. As the score gets higher, it is possible to reach maximum mark-up on more overrunning items, and so there is more of an incentive to spread mark-ups across items. As the score that each bidder submits in equilibrium depends on the level of competition in the auction, competition impacts the level of skewness in equilibrium bids as well.

4.1 Policy Implications of Risk Aversion

Accounting for risk aversion bears significant implications for policy design. If bidders are presumed to be risk neutral, then there is no value in policies to reduce variance in item quantity realizations, limit ex-post compensation for overruns, or subsidize entry: in equilibrium, DOT spending would remain unchanged in each case.[30] However, if bidders are presumed to be risk averse, such policies may be consequential. Risk averse bidders suffer a penalty for uncertainty. Consequently, they bid more aggressively when uncertainty is reduced (so that the value for winning is higher) and bid less aggressively when uncertainty is increased.

To illustrate this, we demonstrate the equilibrium savings from eliminating uncertainty about item quantity realizations in the toy example above.[31] In Table 1, we present the expected cost incurred by the DOT in the equilibria of four auction settings. First we consider the baseline example discussed above when bidders are risk averse with CARA coefficient $\gamma = 0.05$ and when bidders are risk neutral (e.g. $\gamma = 0$).[32] Next, we consider each case under a counterfactual in which uncertainty regarding quantities is eliminated as follows: the DOT is able to discern the precise quantities that will be used and advertises the project with the ex-post quantities, rather than with 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. $\sigma_c^2 = \sigma_r^2 = 0$).

[30] See Online Appendix A.4 and Online Appendix A.5 for a full discussion.

[31] This is analogous to the counterfactual discussed in Section 9.1.

[32] To hone in on the effects of risk in particular, and not mis-estimation, we 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).

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

Table 1: Comparison of Expected DOT Costs

In each auction setting, we compute the equilibrium bids that each type of bidder would submit. Table 1 presents the expectation of the amount that the DOT would pay the winning bidder (e.g., $q_c^a b_c^w + q_r^a b_r^w$) at the equilibrium bidding strategy in each setting, taken with respect to the distribution of the type of the lowest (winning) bidder.[33] Whereas eliminating item quantity uncertainty reduces DOT costs by 7% with risk averse bidders, it has no impact with risk neutral bidders. This highlights the role of uncertainty in driving costs. Although the DOT quantity estimates—which are used to compose the score formula—change when uncertainty is eliminated, the bidders' estimates, $\mathbb{E}[q_c^a]$ and $\mathbb{E}[q_r^a]$—which determine the revenue that they expect from winning—are unchanged. Consequently risk neutral bidders, who are indifferent to the variance in their estimates, charge the DOT the same amount in equilibrium.[34]

CARA Coeff	Noisy Estimates	Perfect Estimates	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 2: Comparison of expected DOT costs under different levels of bidder risk aversion

Furthermore, the degree of the counterfactual savings with risk averse bidders depends on the baseline level of uncertainty in each project, the degree of bidders' risk aversion, and the level of competition in each auction. To illustrate this, we repeat the exercise summarized in Table 1 over different degrees of risk aversion. In Table 2, we present the expected DOT cost, under a range of CARA coefficients, the baseline example with noisy DOT estimates

[33] 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.

[34] Note that the unit bids underlying these expected costs change both in the risk neutral and risk averse case. See Online Appendix A.3 and Online Appendix A.4 for a detailed discussion.

and under the counterfactual in which the DOT eliminates quantity risk. The bolded row with a CARA coefficient of 0.05 corresponds to the right hand column of Table 1.[35]

5 Data and Reduced Form Results

5.1 Data

Our data come from MassDOT and cover highway and bridge construction and maintenance projects undertaken by the state from 1998 to 2015. We are limited by the extent of MassDOT's collection and storage of data on its projects. 4,294 construction and maintenance projects are in the DOT's digital records, although the coverage is sparse prior to the early 2000s. Using only the 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, we are left with 2,513 projects, 440 of which are related to bridge work. We focus on bridge projects alone for this paper, as these projects are particularly prone to item quantity adjustments. Coverage is especially poor in the first few years of the available data and is especially good since 2008, when MassHighway became MassDOT and a push to improve digital records went into effect.[36]

MassDOT began using an online procurement service, called Bid Express, in April 2011. Prior to Bid Express, each contractor submitted his bids in paper form and MassDOT personnel then manually entered the bid data into an internal data set. The shift from a paper process to an online process thus likely helped data collection efforts and improved data accuracy.

The rules of the procurement process were the same, however, before and after April 2011. All bidders who participate in an auction have been 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. The Project Manager assigns a Resident Engineer to monitor work on a particular project out in the field and to

[35]We repeat this exercise across different magnitudes of prediction noise in Online Appendix A.3.

[36]See Table 16 for a breakdown of the number of projects in our data, by year.

be the first to decide whether to approve or reject under-runs, over-runs, and Extra Work Orders (EWOs).[37] 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 projects relating to bridge construction and maintenance, as this is a focal point of interest to the DOT, as well as an area with a fair amount of uncertainty for the bidders.

Statistic	Mean	St. Dev.	Pctl(25)	Median	Pctl(75)
Project Length (Estimated)	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	6.55	3.04	4	6	9
# Types of Items	67.80	36.64	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
Extra Work Orders	$298,796	$295,173	$78,775	$195,068	$431,188

Table 3: Summary Statistics

Table 3 provides summary statistics for the bridge projects in our data set. We measure the extent to which MassDOT overpays the projected project cost in two ways. First, we consider the difference between what the DOT ultimately pays the winning bidder (the sum of the actual quantities used, multiplied the winning bidder's unit bids) and the DOT's initial estimate (the sum of the DOT's quantity estimates, multiplied by the DOT's estimate for each item's unit cost). Summary statistics for this measure are presented in the "Net Over-Cost (DOT Quantities)" row of Table 3. While it seems as though the DOT is saving money on net, this is a misrepresentation of the costs of bid skewing. As we demonstrated in Section 4, the DOT's estimate, which can be thought of as the *score* evaluated using the DOT's unit costs as bids, is not representative of the ex-post amount to be paid at those bids. Rather, a more appropriate metric is to compare the amount ultimately spent against the dot product of the DOT's unit cost estimates and the actual quantities used. This is presented in the "Net Over-Cost (Ex-Post Quantities)" row of Table 3. The median over-payment by this metric is about $15,000, but the 25th and 75th percentiles are about -$210,000 and $275,000. 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.

[37]The full approval process of changes in the initial project design involves several layers of review.

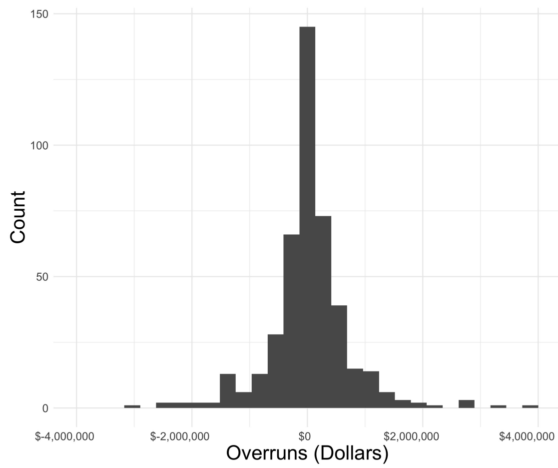

Figure 1: Net Over-Cost (Ex-Post Quantities) Across Bridge Projects

Description of Bidders

Across our data set, there are 2,883 unique project-bidder pairs (e.g. total bids submitted) across the 440 projects that were auctioned off. There are 116 unique firms that participate, albeit to different degrees. We distinguish firms that are rare participants by dividing firms into two groups: 'common' firms, which participate in at least 30 auctions within our data set, and 'rare firms', which participate in fewer than 30 auctions. We retain the 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.

Although there is little publicly available financial information about them, the firms in our data set are by and large relatively small, private, family-owned businesses.[38] Table 4 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 (e.g. the score) submitted 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 (the percent difference of the sum of unit bids multiplied by the ex-post quantities and the sum of blue book costs multiplied by the ex-post quantities) 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 than the common firms—$4.5 million and $4.2 million respectively, and this is reflected in a substantially larger ex-post over-run: 21.97% on average.

[38]Table 17 in the appendix presents the number of auctions participated in and won by each of the top 10 most common firms, as well as estimates of the number of full time employees on their payrolls.

	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 4: Comparison of Firms Participating in <30 vs 30+ Auctions

In addition to the firm's identity, there are a number of factors which 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 cost types across firms and auctions. One such factor is the firm's distance from the worksite.[39] 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 has placed a bid on within our dataset.[40] 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

[39] Although we do not observe precise locations for each project in our data, we observe which of the six geographic districts under MassDOT jurisdiction each project belongs to. We then geocode the headquarters of each firm by district, and compare districts for each project-bidder pair.

[40] Our data involve three distinct project types, according to the DOT taxonomy: Bridge Reconstruction/Rehabilitation projects, Bridge Replacement projects, and Structures Maintenance 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 its utilization—the share of the firm's capacity that is filled when it is bidding on the current project.[41] 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.[42]

Description of Quantity Estimates and Uncertainty

As we discuss in Section 4, scaling auctions improve social welfare 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 highly uncertain items (and higher bids on more predictable items).[43]

Our data set 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.[44]

[41] We measure capacity and utilization with respect to all projects with MassDOT recorded in our data—not just bridge projects.

[42] While we do not take dynamic considerations of capacity constraints into consideration, 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 its size.

[43] A number of different factors may influence the extent of item over/under-runs in a given project: the type of maintenance needed, the underlying state of the structure, the time since assessment and the skill of the project designer. While our dataset is insufficient to robustly estimate the causal effects of these features on over-runs, we account for observables such as the identity of the designer and project manager in the first stage of our estimation.

[44] Part of the reason that so many unique items appear so rarely in our data is that the DOT taxonomy is very specific. For instance, item 866100 – also known as "100 Mm Reflect. White Line (Thermoplastic)" – is distinct from item 867100 – "100 Mm Reflect. Yellow Line (Thermoplastic)," although clearly there is a relationship between them. In order to take these similarities into account, we project item-project pairs onto characteristic space constructed by natural language parsing of the item descriptions, as well as a number of numerical item-project features. We discuss this at greater length in the estimation section.

For each item, in every auction, we observe the quantity with which MassDOT 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 *blue book* DOT estimate for 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% 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%.

Furthermore, quantity over-runs 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 over-run 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.[45] 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.

5.2 Reduced Form Evidence for Risk Averse Bid Skewing

As in Athey and Levin (2001) and Bajari, Houghton, and Tadelis (2014), the bids in our dataset are consistent with a model of similarly informed bidders who bid strategically to maximize expected utility. In Figure 3, we plot the relationship between quantity over-runs and the percent by which each item was overbid above the blue book cost estimate by the winning bidder.[46] The binscatter is residualized. In order to obtain it, we first regress percent overbid 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 in red, 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

[45]The statement of majority here is with respect to items that appear multiple times.

[46]The percent overbid of an item is defined as $\frac{b_t - c_t}{c_t} \times 100$, where b_t is the bid on item t and c_t is the blue book unit cost 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.

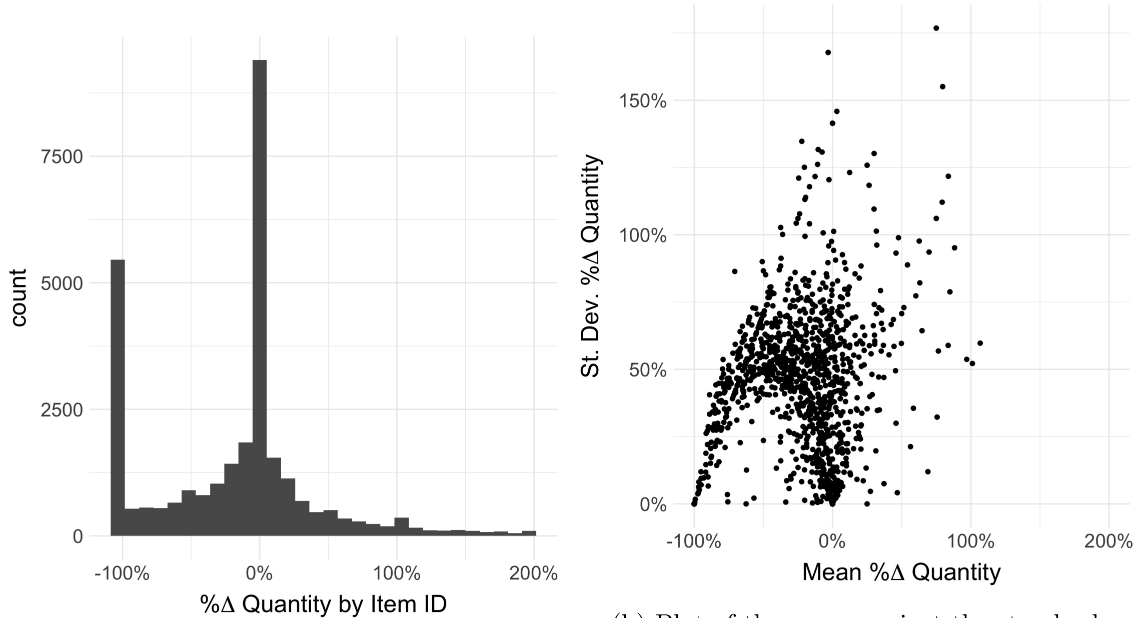

(a) Histogram of the percent quantity over-run across item-project pairs.

(b) Plot of the mean against the standard deviation of percent quantity over-runs for each unique item.

Figure 2

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 a very similar slope.[47] We use a similar procedure for all residualized binscatters in this section.[48]

As Figure 3 demonstrates, there is a significant positive relationship between percent quantity over-runs and percent overbids by the winning bidder. A 1% increase in quantity over-runs corresponds to a 0.085% increase in overbids on average. This suggests that the winning bidder is able to correctly predict which items will over-run on average. As in the example in Section 4, items predicted to over-run generally receive higher bids. Thus, as higher bids correspond to items that overran in our data, we conclude that bidders are informed beyond the DOT quantity estimates and are skewing strategically.

Furthermore, the bid skewing relationship is similar across bidders other than the winner. Figure 4a plots the residualized binscatter of percent overbids against percent quantity over-

[47]For each graph, we truncate observations at the top and bottom 1%. This is done for the purposes of clarity as outliers can distort the visibility of the general trends. We will include untruncated versions in an online appendix for robustness.

[48]As the axes in our graphs are in percentage terms for comparability across items, it may be difficult to see how significant different points are in net value. We thus present our main graphs with bid-level-weighted dot sizes in Appendix F.1.

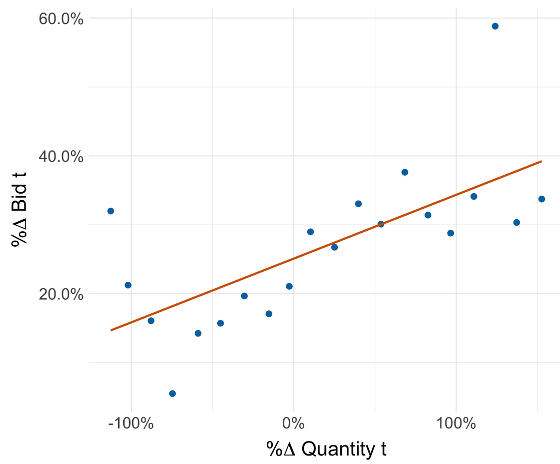

Figure 3: Residualized binscatter of item-level percent winner overbid against percent quantity over-run

runs for the winning bidder and the second-place bidder in each auction. With the exception of a few outlying points, the relationship between overbids and over-runs is very similar between the top two bidders. In the appendix, we show that this relationship is even stronger when we restrict the comparison to projects in which the first two bidder submit similar total scores.

Figure 4b plots a residualized binscatter of the winning bidder's unit bid for each item against the second place bidder's bid for the same item. Overall, the direction of skewing corresponds strongly between the top two bidders—a higher overbid by the winning bidder corresponds to a higher overbid by the second place bidder as well.[49] Together, these figures suggest that bidders have access to the same information regarding quantity over-runs.

While our data suggests 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 for 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. Athey and Levin (2001) argue that the interior bids observed in their data are suggestive of risk aversion among the bidders. While they acknowledge that other forces, such as fear of regulatory rebuke, may provide an alternative explanation for the lack of total bid skewing, they note that risk avoidance was the primary explanation given to them in interviews with professionals.

[49]Note that the percent overbids in Figure 4b appear to be substantially larger than those in Figure 4a. This is because while large overbids occur in the data, they are relatively rare and so are averaged down in the percent quantity over-run binning of Figure 4a.

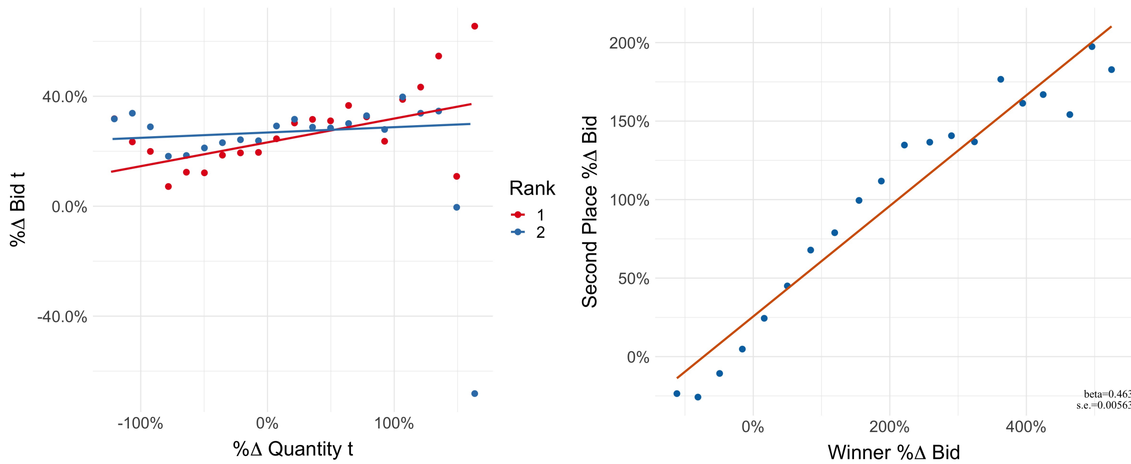

(a) Residualized binscatter of item-level percent overbid by the rank 1 (winning) and rank 2 bidder, against percent quantity over-run.

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

Figure 4

In addition to interior bids, risk aversion has several testable theoretical implications. As we discuss in Section 4, risk averse bidders balance the incentive to bid high on items that are projected to over-run with an incentive to bid lower on items that are uncertain. As such, we would expect bidders to bid lower on items that—everything else held fixed—have higher uncertainty. While we do not see observations of the same item in the same context with identifiably different uncertainty, we present the following suggestive evidence.

In Figures 5a and 5b, 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 7.[50] For every item, in every auction, our first stage gives us an estimate of the variance of the error on the best prediction of what the ex-post quantity of that item would be, given the information available at the time of bidding.

In Figure 5a, we plot a residualized binscatter of the winning bidder's absolute percent overbid on each item against the item's standard deviation—the square root of the estimated

[50] As we discuss in Section 7, we fit a model for the distribution of the ex-post quantity of each item in each auction. The model has two parts: first, we model the ex-post quantity of each item observation as a linear function of the DOT quantity estimate for that item and a vector of item-auction specific features, given a Gaussian error. Second, we model the variance of the Gaussian error in each observation to a lognormal distribution, the mean of which is also a linear function of the DOT quantity estimate and item-auction features. We fit this model jointly with a Hamiltonian Monte Carlo estimator, using the full history of item-auction observations in our data set. Intuitively, this is akin to projecting the ex-post quantity of each item observation onto its DOT estimate and feature vector, and then parametrically fitting the resulting residuals to a lognormal distribution.

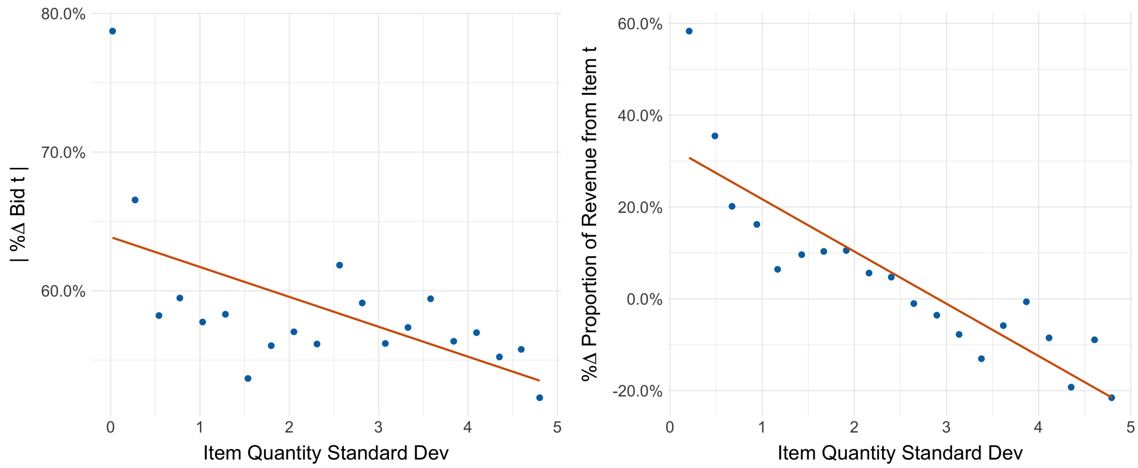

(a) Residualized binscatter of item-level percent absolute overbid against the square root of estimated item quantity variance.

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

Figure 5

prediction variance. The relationship is negative, suggesting that holding all else fixed, bidders bid closer to cost on items with higher variance, limiting their risk exposure.[51] Note, however, that this analysis does not directly account for the trade-off between quantity overruns and uncertainty. As in Equation (1), 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 that the winning bidder earned that was due to item t, and the proportion of the DOT's initial cost estimate that item t constituted.[52] In Figure 5b, we plot the residualized bin scatter 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, as our model of risk averse bidding predicts.

[51] To account for the impact of quantity expectations, we include $\%\Delta q_t$ as a control in the specification when residualizing. However, the qualitative negative relationship persists even if we exclude it. We present this in Appendix F.1 for completeness.

[52] 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/smaller role in the total cost of a project.

6 A Structural Model for Bidding With Risk Aversion

A procurement project is characterized by T items, each of which is needed in a different quantity. The DOT initiates an auction for the project by posting a list of the T items, along with a vector of *estimated* quantities $\mathbf{q}^e = \{q_1^e, \ldots, q_T^e\}$, with which it expects each item to be used. Once the auction is complete, the project is implemented in full by the winning bidder using the *actual* (ex-post) quantity q_t^a for each item t. The actual quantities $\mathbf{q}^a = \{q_1^a, \ldots, q_T^a\}$ are assumed to be fixed but unknown at the time of the auction. That is, from the perspective of the DOT and the bidders, the vector of actual quantities $\mathbf{q}^a$ is an exogenous random variable. The realization of $\mathbf{q}^a$ is independent of which bidder wins the auction, and at what price.[53]

The auction is simultaneous with sealed bids, but both the set of $m > 1$ participating bidders and the DOT's quantity estimates $\mathbf{q}^e$ are fixed and common knowledge to all participants at the start of the auction. In addition, prior to the auction, the bidders receive a symmetric noisy signal $\mathbf{q}^b = \{q_1^b, \ldots, q_T^b\}$ of what the ex post quantities for the project will be:

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

For simplicity, we assume that the signals are common across bidders. Thus, all bidders have the same expected value q_t^b for the actual quantity of item t, and the same variance σ_t^2, with which this estimate is off.[54]

Bidders differ in their private cost of production along a single dimensional *efficiency multiplier* α. At the time of the auction, every item t has a commonly-known market unit cost c_t. This cost represents the market price of the materials—generally items such as concrete, traffic cones, etc., which are standard and competitive—at the scale necessary for the project. However, the bidders vary in their labor and transportation costs, storage capacity, etc., yielding a multiplicative (dis)advantage over competitors. In particular, for every item t in the project, bidder i faces a unit cost of $\alpha^i c_t$, where α^i is the bidder's efficiency (multiplier) type. The efficiency type of each bidder i is drawn independently from a common, publicly known distribution with a well behaved density $f(\alpha^i)$ over a compact

[53]This assumption, which follows Bajari, Houghton, and Tadelis (2014) and Athey and Levin (2001), precludes the possibility of asymmetric moral hazard. In our reduced form section, we argue that the similarity in projected over-runs by the winning bidder and the runner-up suggests that if moral hazard affects bidding, its effects are anticipated symmetrically by bidders so that this assumption, too, will not affect our estimates greatly. It also precludes substitutability between items. While we cannot rule substitutions out, we argue that their scope is limited as only items on the DOT designer's project specification may be used for construction.

[54]It is not without loss of generality to assume that signals are common across bidders. However, we make this assumption for the sake of tractability.

subset $[\underline{\alpha}, \overline{\alpha}]$ of $\mathbb{R}_+$.[55] Each bidder privately observes only her own efficiency type prior to the auction, but the distribution of competitor types is common knowledge.

To participate, each bidder i submits a vector of unit prices $\mathbf{b}^i = \{b^i_1, \ldots, b^i_T\}$, setting the amount per unit that she will be paid for each item if she wins. The winner of the auction is determined according to a first-price scoring rule. Each bidder i is given a *score* based on her unit bids and the DOT quantity estimates:

$$s^i = \sum_{t=1}^{T} b^i_t q^e_t.$$

The bidder with the lowest score wins the contract and implements the project in full. Upon the completion of the project, the actual (ex-post) quantities $\mathbf{q}^a$ of the items are realized, and the winning bidder is paid her unit bid b^i_t multiplied by the ex-post quantity q^a_t for each item. The winning bidder is responsible for securing all of the materials and labor for the project privately, and so she also incurs a cost of $\alpha^i c_t$ multiplied by q^a_t for each item.[56]

Finally, we model the bidders as risk averse, with a standard CARA utility function over their earnings from the project and a common constant coefficient of absolute risk aversion γ:[57]

$$u(\pi) = 1 - \exp(-\gamma\pi). \tag{3}$$

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

$$\pi(\mathbf{b}^i, \alpha^i, \mathbf{c}, \mathbf{q}^a) = \sum_{t}^{T} q^a_t \cdot (b^i_t - \alpha^i c_t),$$

if she wins the auction.[58] As $\mathbf{q}^a$ is a random variable from the bidder's perspective at the time of bidding, her profit from winning is stochastic as well.

[55] The assumption that the distribution of efficiency types is common (e.g. not specific to individual bidders) is not critical to our analysis, and relaxing it would not substantially change our estimation method or results, although it might impact the counterfactuals.

[56] Only the winner of the auction incurs any costs. All losing bidders receive no further cost nor revenue from the project once the auction is complete.

[57] Equation (3) can be thought of as a normalization of the CARA utility function $u_i(\pi) = \exp(-\gamma_w) - \exp(-\gamma_w(w+\pi))$, where w is the bidder's wealth independent of the auction, and $\gamma_w = \frac{\gamma}{w}$ is the unnormalized CARA coefficient. When w is the same across all of the bidders in the auctions, this normalization is without loss of generality. While this is a strong assumption, we will maintain it throughout the main part of this paper for the purpose of tractability in this draft.

[58] We suppress the common auction characteristics $\mathbf{c}, \mathbf{q}^e, \mathbf{q}^a$ as arguments in the utility and profit functions for ease of exposition.

Bidder i chooses her bids so as to maximize her expected utility at the time of the auction:

$$\left(\underbrace{1-\mathbb{E}_{\mathbf{q}^a}\left[\exp\left(-\gamma\sum_{t=1}^{T}q_t^a\cdot(b_t^i-\alpha^i c_t)\right)\right]}_{\text{Expected utility conditional on winning}}\right)\cdot\underbrace{\left(\Pr\left\{s^i<s^j \text{ for all } j\neq i\right\}\right)}_{\text{Probability of winning with } s^i=\mathbf{b}^i\cdot\mathbf{q}^e}. \quad (4)$$

This is bidder i's expected utility over her profit 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. Bidders form their expectations for the expected utility of winning based on the posterior distribution of each q_t^a given by equation (2) at their signals q_t^b and σ_t^2. The expected utility of bidder i can therefore be rewritten:[59]

$$\left(1-\exp\left(-\gamma\sum_{t=1}^{T}q_t^b(b_t^i-\alpha^i c_t)-\frac{\gamma\sigma_t^2}{2}(b_t^i-\alpha^i c_t)^2\right)\right)\cdot\left(\Pr\left\{s^i<s^j \text{ for all } j\neq i\right\}\right).$$

6.1 Equilibrium Bidding Behavior

We now characterize the Bayesian Nash Equilibrium of the static first-price sealed bid scoring auction described in the previous section. Our setting is similar to Bajari, Houghton, and Tadelis (2014), which uses a special case of the Asker and Cantillon (2008) model, in which the project and its value to the DOT are fixed and independent of the winning bidder. We consider a linear scoring auction game with independent private values that can be characterized by a uni-dimensional "pseudo-type"—each bidder's *efficiency multiplier type* α.[60]

As in Bajari, Houghton, and Tadelis (2014) and Asker and Cantillon (2008), the optimal bidding problem in our setting can be decomposed into two parts: (1) given an efficiency type α, choose the optimal score s; (2) given a score s, choose the optimal bid vector $\mathbf{b}$ subject to the constraint that $\mathbf{b}\cdot\mathbf{q}^e=s$. As we describe below, the optimal choice of $\mathbf{b}$ conditional on a choice of s, a type α, and the auction characteristics, is deterministic and independent of competitive considerations. Therefore, at the optimum, the value of winning the auction to a bidder of type α who submits a score s—that is, the bidder's expected utility from winning the auction using the optimal vector of bids $\mathbf{b}$ that add up to a score of s—is determined entirely by her choice of s, and is monotonically increasing in s. Following a

[59]To get this, we rewrite $q_t^a=q_t^b-\varepsilon_t$, and take the expectation of the profit term with respect to the distribution of ε.

[60]Another related reference is Che (1993), which employs a uni-dimensional bidder type referred to as the bidder's "productive potential".

sub-case of Lebrun (2006), this game has a unique monotonic equilibrium in pure strategies.

We derive the equilibrium as follows for an arbitrary bidder i with efficiency type α^i:

1. Given a (winning) score s, we find the optimal bid vector $\mathbf{b}^i(s)$ s.t. $\sum_{t=1}^{T} b_t^i(s) q_t^e = s$.

   To do this, we solve the convex optimization program:

$$\max_{\mathbf{b}^i(s)} \left[1 - \exp\left(-\gamma \sum_{t=1}^{T} q_t^b (b_t^i(s) - \alpha^i c_t) - \frac{\gamma \sigma_t^2}{2} (b_t^i(s) - \alpha^i c_t)^2\right)\right] \tag{5}$$

$$\text{s.t.} \sum_{t=1}^{T} b_t^i(s) q_t^e = s.$$

   The objective function is separable in t and concave, and so this optimization problem will have a unique global maximum. Moreover, applying the monotone transformation $T(f(x)) = -log(-f(x) - 1)$, we can characterize the solution to (5) by the constrained quadratic program:

$$\max_{\mathbf{b}^i(s)} \left[\gamma \sum_{t=1}^{T} q_t^b (b_t^i(s) - \alpha^i c_t) - \frac{\gamma \sigma_t^2}{2} (b_t^i(s) - \alpha^i c_t)^2\right] \tag{6}$$

$$\text{s.t.} \sum_{t=1}^{T} b_t^i(s) q_t^e = s.$$

   The solution to this program is given by:[61]

$$b_{i,t}^*(s) = \alpha^i c_t + \frac{q_t^b}{\gamma \sigma_t^2} + \frac{q_t^e}{\sigma_t^2 \sum_{p=1}^{T} \left[\frac{(q_p^e)^2}{\sigma_p^2}\right]} \left(s - \sum_{p=1}^{T} \left[\alpha^i c_p q_p^e + \frac{q_p^b q_p^e}{\gamma \sigma_p^2}\right]\right). \tag{7}$$

2. Let $\mathbf{b}_i^*(s)$ be the optimal mapping from score to bid distribution for bidder i, as in equation (7). We find the optimal score for bidder i by maximizing her (unconditional) expected utility given the equilibrium distribution of opponent scores.

[61] This formulation of the optimal bid program does not explicitly constrain unit bids to be non-negative. This is not with loss of generality, and we apply the additional non-negativity constraint when computing counterfactual bids. However, as all observed bids are positive (meaning that the non-negativity constraint did not bind), this 'unconstrained' program serves as a very useful approximation to the solution of the fully constrained program. In particular, while the fully constrained program does not have a closed form solution and must be solved with interior point algorithms or the like, the 'unconstrained' version has a closed form solution that is linear in our parameters of interest. As we show in section 8, the bids predicted by our estimated model do quite well at matching the data.

Let $H^j(\cdot)$ be the CDF of contractor j's score. Then, by bidding a score of s, bidder i obtains an expected utility of:

$$\mathbb{E}[u_i(\pi_i(s)] = \underbrace{\left(1 - \exp\left(-\gamma \sum_{t=1}^{T} q_t^b (b_{i,t}^*(s) - \alpha^i c_t) - \frac{\gamma \sigma_t^2}{2}(b_{i,t}^*(s) - \alpha^i c_t)^2\right)\right)}_{\text{Expected utility conditional on winning}} \cdot \underbrace{\left(\prod_{k \neq i}(1 - H^k(s))\right)}_{\text{Prob of win w/ } s = \mathbf{b}_i^* \cdot \mathbf{q}^e} .$$

The first phrase in parentheses is i's expected utility from the total profit that she stands to make from winning the auction, and the second phrase is the probability that s is the lowest score given the equilibrium score distributions $H^j(\cdot)$ for competing contractors $j \neq i$. As the expected utility conditional on winning is increasing in s—due to the optimality of $\mathbf{b}_i^*(s)$—the unconditional expected utility function $\mathbb{E}[u_i(\pi_i(s)]$ is concave in s. Thus, the optimal strategy is described by the unique maximizer:

$$s^* = \arg\max \mathbb{E}[u_i(\pi_i(s)] \tag{8}$$

and the corresponding unit bid vector $\mathbf{b}_i^*(s^*)$.[62]

7 Econometric Model

We now present a two-step estimation procedure to estimate the model described in the previous section. 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 will use the full history of auctions in our data to build a statistical model of bidder expectations using publicly available project characteristics. However, it will not take into account bidding incentives in any particular auction. By contrast, the second stage will estimate the coefficient of risk aversion γ, as well as each bidder's efficiency type α for each auction that she participates in. In this stage, we take the first stage estimates as fixed and construct moments for GMM estimation using the optimality of observed bids submitted by each bidder i in auction n, given our model.

[62] As is standard in auction theoretic analysis (see Milgrom and Segal (2002), for example), the optimal strategy is described by the first order condition of Equation (8) with respect to s, taking into account the dependence of the optimal bid vector $\mathbf{b}_i^*$ on s. Previous work, such as Bajari et al. (2014), has used this first order condition as the primary source of identification of bidder types. In our analysis, we use the optimality of each unit bid at the optimal choice of s^* instead, and so we defer further discussion of equilibrium construction to Appendix A.

Stage 1: Estimating the Posterior Distribution of q_t^a

In the model presented in Section 6, we did not take a stance on what the signals in Equation (2) are based on. The reason for this was to emphasize the flexibility of our model with respect to possible signal structures: the only required assumption is that conditional on all of the information held at the time of bidding, the posterior distribution of each q_t^a can be approximated by a normal distribution with a commonly known mean and variance. In particular, it allows for correlations between items, as well as complicated forms of correlation between the bidders' beliefs and the DOT's expectations.[63]

For the purpose of estimation, however, we make an additional assumption. We assume that the posterior distribution of each q_t^a is given by a statistical model that conditions on q_t^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. This assumption can be thought of in several ways. It can be interpreted as an additional component of the structural model: the bidders use a statistical estimation procedure to assess the likelihood of item quantities, and consequently, the value of the project, 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, it is likely that firms do precisely this when forming their bids. There is a competitive industry of software for procurement bid management that touts sophisticated estimation of project input quantities and costs. Alternatively, this assumption could be thought of as the econometrician's model of each signal mean q_t^b and variance σ_t^2.

In particular, denote an auction by n and the items involved in auction n by $t \in \mathcal{T}(n)$. We model the realization of the actual quantity of item t in auction n by:

$$q_{t,n}^a = \widehat{q_{t,n}^b} + \eta_{t,n}, \text{ where } \eta_{t,n} \sim \mathcal{N}(0, \hat{\sigma}_{t,n}^2) \tag{9}$$

such that
$$\widehat{q_{t,n}^b} = \beta_{0,q} q_{t,n}^e + \vec{\beta}_q X_{t,n} \text{ and } \hat{\sigma}_{t,n} = \exp(\beta_{0,\sigma} q_{t,n}^e + \vec{\beta}_\sigma X_{t,n}). \tag{10}$$

Here, $\widehat{q_{t,n}^b}$ is the posterior mean of $q_{t,n}^a$ and $\hat{\sigma}_{t,n}$ 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 matrix of item-project characteristics $X_{t,n}$. We estimate this model with Hamiltonian Monte Carlo and use the posterior mode as a point estimate for the second stage of estimation.[64] We demonstrate the goodness of fit in Section 8.

[63]To be more precise, our model requires that item quantities q_t^a have independent Gaussian errors, but puts no restrictions on the means or standard deviations, which could be correlated across items, etc.

[64]We use Hamiltonian Monte Carlo as an efficient implementation of a likelihood method, optimized for a GLM. We discuss the flexibility of our approach to alternative models/estimators for the first stage in Appendix B.1.

Stage 2: Estimating Cost Types and the CARA Coefficient

We now discuss our econometric model for the estimation of the CARA coefficient of risk aversion γ, along with the set of bidder-auction efficiency types α_n^i. The key to our identification strategy lies in the heterogeneity of unit bids that we observe in our data. Our dataset contains a unit bid for every item, submitted by every participating bidder in every auction that we see. In particular, we have three main sources of heterogeneity: (1) bids submitted by different bidders in a given auction with the same project characteristics, item list, etc.; (2) bids submitted by the same bidders across different items and different auctions that have different project characteristics, etc.; (3) bids submitted for the same items by bidders across different auctions with different project characteristics, quantity projections, and participating bidders.

Denote auctions by n, the bidders participating in the auction by i, and the items involved in the auction by t. The model of optimal bidding described in Section 6 predicts that the optimal unit bid for item t for a bidder of type α_n^i in auction n is given by:

$$b^*_{t,i,n}(s^*_{i,n}) = \alpha_n^i c_{t,n} + \frac{q^b_{t,n}}{\gamma \sigma^2_{t,n}} + \frac{q^e_{t,n}}{\sigma^2_{t,n} \sum_{p=1}^{T_n} \left[\frac{(q^e_{p,n})^2}{\sigma^2_{p,n}} \right]} \left(s^*_{i,n} - \sum_{p=1}^{T_n} \left[\alpha^i c_{p,n} q^e_{p,n} + \frac{q^b_{p,n} q^e_{p,n}}{\gamma \sigma^2_{p,n}} \right] \right), \qquad (11)$$

where $s^*_{i,n}$ is the optimal score for this bidder, and by definition, $s^*_{i,n} = \sum_t^{T_n} q^e_{t,n} b^*_{t,i,n}(s^*_{i,n})$.

As discussed above, we identify $q^b_{t,n}$ and $\sigma^2_{t,n}$ in Stage 1 with a statistical model of ex-post quantities using the full history of item characteristics and quantities across the auctions in our data and independently of bids. To further reduce the dimensionality of our parameter space in the second stage, we model the bidder-auction efficiency type α_n^i as the sum of a bidder-specific fixed effect and a regression model of bidder-auction characteristics:

$$\alpha_n^i = \alpha^i + \vec{\beta}_\alpha X_{i,n}.$$

Finally, we make the following assumption to connect our first stage estimates to our bid data and close our model:

Assumption 1. *Let $b^d_{t,i,n}$ 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 $b^*_{t,i,n}$, subject to an IID, mean-zero measurement error $\nu_{t,i,n}$:*

$$b^d_{t,i,n} = b^*_{t,i,n} + \nu_{t,i,n}$$

where
$$\mathbb{E}[\nu_{t,i,n}] = 0 \text{ and } \nu_{t,i,n} \perp X_{t,n}, X_{i,n}.$$

Assumption 1 states that each unit bid observed in our data is given by the optimal bid implied by our model—at the true underlying parameters—subject to an idiosyncratic error that is independent across draws, and orthogonal to auction-item and auction-bidder characteristics. Such an error might come about because of rounding/smudging in the translation between the bidder's optimal bidding choice and the record available to the DOT (and consequently, to the econometrician). One might alternatively frame this error as an optimization error: the optimal choice of bids is a numerical solution to a constrained quadratic program that may not produce numbers that are convenient to report in currency.

To further see the need for Assumption 1, note that while an auction with T items and I bidders has $T \times I$ unit bids, our model allows for only T quantity predictions, T item variance terms, I bidder efficiency types, and 1 coefficient of risk aversion as free parameters to explain these bids. Absent an additional assumption, a model in which all $T \times I$ bids must match the bids in our data would be rejected in most cases. It is not, however, strictly necessary for our model to assert independence in error within bidder or project.[65]

Assumption 1 implies that the difference between each unit bid observed in our data and the optimal bid in Equation (11) has a mean of zero, and is uncorrelated with either the identity or the characteristics of the bidder bidding and the item being bid on. As such, we estimate the second stage parameters γ, α^i and $\vec{\beta}_\alpha$ through a just-identified system of moment conditions:
$$\mathbb{E}\left[\tilde{\nu}_{t,i,n} \cdot Z_{t,i,n} | X_{t,n}, X_{i,n}\right] = 0,$$

where Z is each of the following instruments:

- Indicator for unique firm IDs[66]
- Indicator for being a "top skewed item"
- The bidder-auction feature vectors that comprise $X_{i,n}$.

Here, each bid error observation $\tilde{\nu}_{t,i,n}$ is formed by subtracting the optimal bid $b^*_{t,i,n}$ for each item t, submitted by bidder i in auction n, as given by Equation (11) from the observed bid $b^d_{t,i,n}$. We write $\tilde{\nu}_{t,i,n}$ rather than $\nu_{t,i,n}$, as in the statement of Assumption 1, because

[65]The assumption that bid errors are independent affects our estimates through the procedure to compute standard errors. As we detail in Appendix B.1.2, we compute standard errors via a bootstrapping procedure in which we repeatedly draw auctions independently, at random, and with replacement.

[66]We include a unique ID for all firms involved in at least 10 auctions, and a grouped ID for all firms involved in 9 or fewer auctions. These correspond to unique α^i parameters.

Equation (11) invokes the optimal score $s^*_{i,n}$, which is not directly observed. Instead, we observe $s^d_{i,n} = \sum_t^{T_n} b^d_{i,t,n} q^e_{t,n}$, which is also noisy. As we explain in detail in the appendix, this additional noise cancels out in expectation when we plug $s^d_{i,n}$ into Equation (11) when forming our moments, and so does not bias estimation. For a detailed discussion on moment construction, estimation, and standard error computation, see Appendix B.1.2.

Note that our estimator uses *only* the optimality of item-level unit bids, evaluated at the *observed* equilibrium score, for identification. That is, we only use the solution to Step 1 from Section 6.1, evaluated at the solution to Step 2. As such, our estimator is agnostic to the equilibrium conditions that generate the observed scores themselves. It will remain valid under any model of equilibrium behavior in which the optimal way for individuals to choose unit bids conditional on a score is separable from the choice of the score itself.[67] In particular, our estimator is robust to allowing collusion or dynamic considerations, for instance, and could be adapted to allow for heterogeneous coefficients of risk aversion and multi-dimensional cost types.[68]

The three types of instruments above correspond to three types of moments. The first type of moment, constructed by interactions with firm ID dummies, can be interpreted as follows: the average bid error that a bidder with unique firm ID i submitted, across all auctions that i participated in, is asymptotically zero. There are 25 such moments, one for each unique bidder id i. These moments inform the fixed effects α^i, correspondingly.

The second moment focuses on items that were deemed as "top skew items" according to the DOT Engineering Office. These items are flagged as frequently being given noticeably high or low bids. According to our model, the variation in these bids is reflective of the level of bidders' responses to the uncertainty regarding the quantities of these items (in absolute terms and relative to the remainder of the project). As such, we focus on this set of items to identify the coefficient of risk aversion, γ. The moment can be interpreted as follows: the average of bid errors submitted on "top skew items" is asymptotically zero in the number of auctions in which these items are involved.

The third type of moment, which interacts bid errors with bidder-auction characteristics, can be interpreted as follows: the average bid error submitted in an auction n is orthogonal to each of the 14 bidder-auction features $X^j_{i,n}$, and is asymptotically zero in the number of auctions. There are 14 such moments, one for each column of the feature matrix $X_{i,n}$. We conjecture that the estimate of the coefficients β^j_α is primarily driven by these moments.

[67] This excludes settings in which the choice of unit bids plays a role in sustaining equilibrium. However, since bids are sealed and only the score determines the winner of a given auction, this is not very restrictive.

[68] We focus on the static, independent, uni-dimensional bidder type model in estimation and in the counterfactuals as this admits an easily interpretable unique equilibrium. While this may not capture all relevant strategic considerations, it allows us a clean and tractable way to address the ways in which the distribution of risk borne by different mechanism designs impacts equilibrium spending—the focus of our paper.

8 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. We present parameter estimates for the regression coefficients on the predicted quantity term $\widehat{q}^b_{t,n}$, as well as the variance term, $\widehat{\sigma}^2_{t,n}$ in Table 9 in the appendix. A histogram of the resulting variance terms themselves is plotted in Figure 6, 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.[69] In addition, we demonstrate the model fit of our first stage in Figure 13 and Table 8 in the appendix.

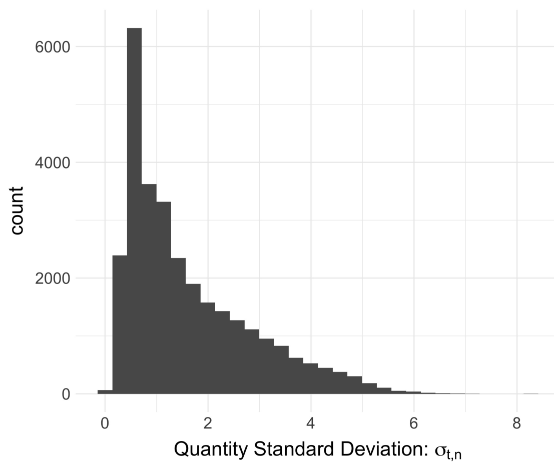

Figure 6: A histogram of standard deviation estimates for each item t in each project n

In the second stage, we estimate a common CARA coefficient γ, as well as a bidder-auction specific efficiency type $\alpha^i_n = \alpha^i + \beta_\alpha X_{i,n}$ for every bidder-auction pair in our data using the GMM estimator presented in Section 7. We summarize the results in Tables 5, 6, and 7. The full parameter estimates are presented in Table 10 in the appendix. The coefficient of risk aversion γ in our data is estimated to be about 0.046. An individual with this level of risk aversion would require a certain payment of \$23 to accept a 50-50 lottery to either win or lose \$1,000 with indifference, and \$2,223 to accept a 50-50 lottery to win or lose \$10,000.[70] As we report in Table 5, the 95% confidence interval around our estimate is

[69] Although we do not plot it here, in general, higher variances correspond to higher quantity predictions as well.

[70] The CARA coefficient we estimate here is only identified up to a dollar scaling. For numerical efficiency, we scaled all dollar values by \$1,000 in estimation and counterfactual simulation. Our results do not depend on the scaling, however. As we have verified, if we scale by an order of magnitude more (or less), the estimated CARA coefficient scales down (or up) by an order of magnitude correspondingly.

(0.032, 0.264). This interval is generated by a bootstrap, in which the data set of auctions is sampled (at the auction level) with replacement in each iteration.[71]

	Parameter Estimate	95Pct CI
$\widehat{\gamma}$	0.046	(0.032,0.264)

Table 5

		$\widehat{\alpha}_n^i$			
Project Type	Mean	St Dev	25%	Median	75%
All	0.975	0.261	0.822	0.949	1.139
Bridge Reconstruction/Rehab	1.019	0.25	0.85	1.005	1.225
Bridge Replacement	0.996	0.219	0.855	1.009	1.159
Structures Maintenance	0.919	0.312	0.782	0.873	0.978

Table 6: Summary statistics of α_n^i estimates by project type.

In Table 6, we present summary statistics of our estimates of bidder-auction efficiency types. We break down the results by project type to highlight the differences between different types of construction. An efficiency of 1 would suggest that the bidder faces costs exactly at the rates represented by MassDOT's blue book. Our results show that the median bidder overall has an efficiency type of 0.949, consistent with estimates of bidder costs by previous papers.[72] There is heterogeneity across project types, however. We estimate that the median bidder in a bridge rehabilitation project has an efficiency type of about 1.005, suggesting that she is about 0.5% less efficient than the DOT estimates. The median bidder in structures maintenance projects, however, has an efficiency type of about 0.873, suggesting that she is about 12.7% more efficient than the DOT estimates.

In Table 7, we present the ex-post markups for each winning bidder given their efficiency type:

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

This is the bidder's total ex-post profit from the project, normalized by her total cost. The numerator is given by the sum of the quantity of each item that was ultimately used $q_{t,n}^a$, multiplied by the bidder's profit from that item. The bidder's profit from an item is her unit

[71] At the moment, the bootstrap is only over the second stage, holding the first stage estimates fixed, due to the computational burden of bootstrapping the first stage as well.

[72] See Bajari, Houghton, and Tadelis (2014) and Bhattacharya, Roberts, and Sweeting (2014), for example.

bid $b_{t,i,n}$ minus her private cost for that item, given by her efficiency type α_n^i multiplied by the blue book market rate estimate $c_{t,n}$. The denominator is calculated similarly, summing only over the bidder's private costs.

		Bidder Markups			
Project Type	Mean	St Dev	25%	Median	75%
All	17.03%	60.88%	-12.84%	5.74%	27.53%
Bridge Reconstruction/Rehab	11.39%	35.88%	-15.61%	7.34%	23.07%
Bridge Replacement	12.8%	67.43%	-12.34%	1.43%	23.67%
Structures Maintenance	23.9%	62.12%	-9.66%	10.56%	39.13%

Table 7: Summary statistics of estimated winning bidders' markups given $\hat{\alpha}_n^i$

The median markup for a winning bidder in our data set, overall, is about 5.74%. There is heterogeneity across project types: the median within bridge replacement projects is 1.43%, for instance, while it is 10.56% for structures maintenance projects. Moreover, there is substantial variation within project types as well. The mean winner markups for bridge replacement and structures maintenance projects are 12.8% and 23.9%, respectively. This may be due to the heterogeneity in projects as well as the ex-post accuracy of bidders' quantity predictions. Furthermore, the 25th percentile of markups is negative for each of the projects as well. This may be due, in part, to inaccurate prediction of the ex-post quantities. However, the ex-post markup calculation does not take into account extra work orders. While we do not estimate profits on the extra work orders in our paper, and so cannot evaluate exactly how extra work orders would affect ex-post profits, this is a key component of BHT's estimation and likely make up the difference in mark-ups.

Finally, we demonstrate the fit of our structural model in Figures 15 and 16 and in Table 11, all in the appendix. Figure 15 plots the unit bids predicted by our model on the x-axis, and the unit bids observed in our data on the y-axis. Figure 16 plots a quantile-quantile plot of our model-predicted bids against the data bids. While bid predictions are not perfect, the correspondence between predictions and data is quite good. Table 11 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 0.879.

9 Counterfactuals

9.1 Perfectly Predicted DOT Quantities

In order to draw conclusions from our results, we return to the discussion in section 4. How much money *would* the DOT save if it were able to perfectly predict the actual quantities that are required for each project?

To answer this question, we solve for the equilibrium in each of the auctions in our bridge projects dataset, under the counterfactual setting in which the DOT perfectly predicts the actual quantities. We assume that the DOT's accuracy is common knowledge and so the bidders believe that the actual quantities will be equal to the DOT's projections with variance approaching zero when making their bidding decisions.[73]

Note that it is not sufficient to simply invert the econometric model of bidding described in Section 7 using our parameter estimates and the counterfactual conditions. The reason for this is that the distribution of competitors' scores is defined in equilibrium. As we demonstrated in Section 4, the score that a bidder with efficiency type α will submit in equilibrium depends on the DOT quantity estimates (as well as the bidders' beliefs and all other auction characteristics). It follows that the equilibrium score distribution itself depends on the DOT quantities, and so we need to solve for the equilibrium from auction primitives afresh in each setting.

An equilibrium of an auction in our setting is determined by the following primitives: the vector of DOT quantity estimates $\boldsymbol{q}^e$, the vector of bidder quantity model predictions, $\boldsymbol{q}^b$, the vector of bidder model variances, $\boldsymbol{\sigma}^2$, the vector of DOT cost estimates $\boldsymbol{c}$, the coefficient of risk aversion γ, and the distribution of the efficiency types of bidders participating in the auction. To evaluate our counterfactuals, we compute the equilibrium bids twice: first in the baseline setting and second in the counterfactual setting. The details of the equilibrium construction are presented in Appendix A.[74]

For the baseline setting, we use the DOT estimates $\boldsymbol{q}^e$ and $\boldsymbol{c}$ from the data, and the bidder quantity model parameters $\widehat{\boldsymbol{q}}^b$ and $\widehat{\boldsymbol{\sigma}}^2$ from the first stage of our estimation. For the distribution of bidder efficiency types, we use a parametric projection of the empirical distribution of the efficiency type estimates $\widehat{\alpha}^i_n$ from our second stage onto auction characteristics.[75] For

[73]In particular, we exclude considerations of short term gains that the DOT might make by accurately predicting actual quantities while the bidders use noisy signals. Since we assume that the bidders form their beliefs over actual quantities using statistics over historical data, any such gains would be short lived as the bidders would eventually realize that the DOT's quantities are accurate.

[74]As an additional measure of model fit, we plot the predicted winning scores from our baseline simulations against the actual scores in our data set in Figure 17 of Appendix D.

[75]Our model assumes that bidders in a given auction are ex-ante IID and so the distribution of bidder types must be auction, rather than bidder-auction, specific.

the coefficient of risk aversion, we use the estimate $\hat{\gamma} = 0.046$ from the second stage of our estimation. In Appendix F.2, we present the results of our main counterfactuals, replicated over a range of γ values that covers the 95% confidence interval of $\hat{\gamma}$ as well.

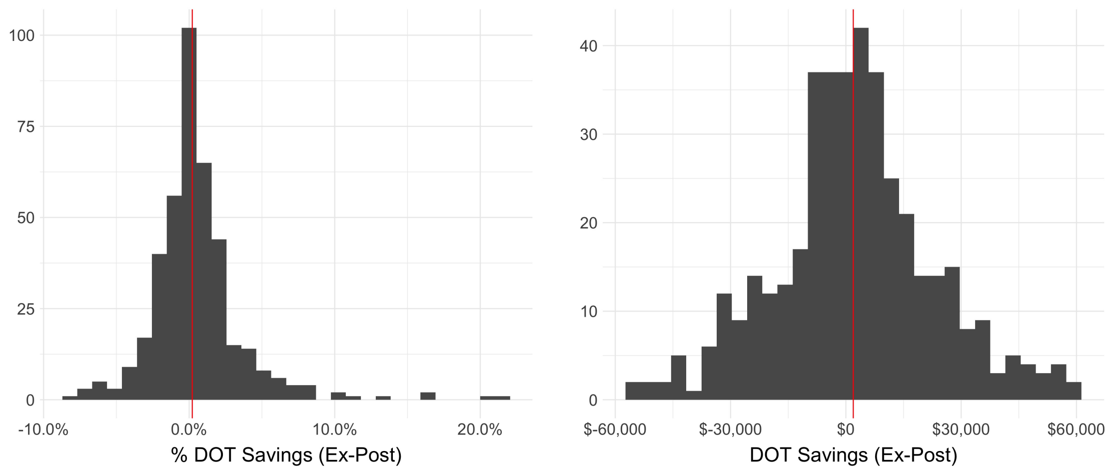

(a) DOT % Savings

(b) DOT Dollar Savings

Figure 7: Percent and dollar expected DOT ex-post Savings from a counterfactual in which risk is eliminated. The median is highlighted by a red vertical line in each case.

In Figure 7, we plot a histogram of the (a) percentage savings and (b) dollar savings to the DOT from the perfect quantity prediction counterfactual.[76] To calculate these savings, we compute the equilibrium bids for every efficiency type α twice: (1) under the baseline setting; and (2) under the counterfactual setting in which the DOT and bidder quantity estimates are equal to the true ex-post quantities, $\boldsymbol{q}^e = \boldsymbol{q}^b = \boldsymbol{q}^a$, and the bidders face no uncertainty, $\boldsymbol{\sigma}^2 \approx 0$.[77] In each case, we calculate the expected total amount that the DOT would pay the winning bidder in equilibrium: the expected value of the sum of the lowest efficiency type's unit bids multiplied by the ex-post item quantities $\boldsymbol{q}^a$.[78]

The dollar gains in Figure 7b are computed by taking the difference between the expected DOT cost under the baseline setting, and under the counterfactual setting for each auction. The percent gains in Figure 7a are given by dividing the dollar saving amount in each auction by the expected DOT cost under the baseline. Finally, we present the bidder utility gains from the counterfactual setting in Figure 8. We calculate bidder utility gains by taking

[76]The histograms in this section are truncated at the tails for the sake of readability. See Appendix E for summary statistics of the underlying data.

[77]We use $\boldsymbol{\sigma}^2 \approx 0$ rather than $\boldsymbol{\sigma}^2 = 0$ in order to avoid numerical overflow issues.

[78]More concretely, let $g(\alpha)$ and $G(\alpha)$ be the density and cumulative probability functions of bidders' efficiency types in a given auction. Let $g^1(\alpha) = Ng(\alpha)(1 - G(\alpha))^{N-1}$ be the density of the first order statistic of g—the density of the lowest type bidder, when there are N bidders in the auction. Denote $b_t^*(\alpha)$ as the equilibrium bid for item t for a bidder with efficiency type α in that auction. The expected DOT cost is given by $\int_{\underline{\alpha}}^{\overline{\alpha}} g^1(\tilde{\alpha}) \sum_t q_t^a b_t^*(\tilde{\alpha}) d\tilde{\alpha}$.

the difference between the (ex-ante) certainty equivalent of a bidder participating in each auction under the baseline and the analogous certainty equivalent under the counterfactual setting.[79] We present summary statistics for all three metrics in Table 12 in Appendix E.

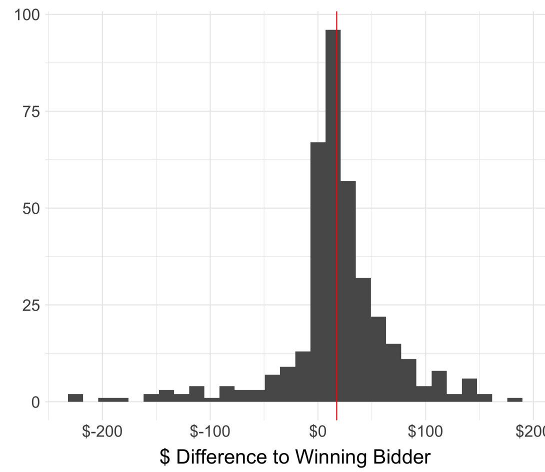

Figure 8: Winning bidder utility gains from a counterfactual in which risk is eliminated. The median is highlighted by a red vertical line.

We predict that the median expected saving to the DOT from eliminating uncertainty about ex-post quantities is about $2,171 or 0.22% of the baseline expected project cost. However, the standard deviation of savings is about $24,568 (4.23%) and the 25th and 75th percentiles are -$9,598 (-1.03%) and $13,887 (1.58%), respectively. This is reflective of the two opposing forces in effect when the DOT eliminates uncertainty. On the one hand, eliminating uncertainty drives bidder risk down, thereby increasing the value of the project to all of the bidders and causing them to bid more aggressively. On the other hand, the counterfactual allows bidders to optimize their bid choices with regard to the true quantities $\boldsymbol{q}^a$ that will be used in the project, whereas in the baseline, bidders optimize on the basis of quantity projections $\boldsymbol{q}^b$. That is, whereas in the baseline bidders optimize unit bids with regard to quantity predictions that may be inaccurate (and so, the bids may not be optimal with respect to the realized quantities, which the winner is ultimately paid for), in the counterfactual with no uncertainty, the bidders always optimize unit bids with respect to the actual quantities that will be used.

As a result, in the auctions where bidders "mis-optimized" under the baseline, the DOT bears a higher cost under the counterfactual. Notably, the ex-ante value of the auction to bidders does not change very much between the baseline and the counterfactual. The median increase in bidders' ex-ante certainty equivalents under the counterfactual is a mere $17.77,

[79] The certainty equivalent is defined as the amount of money that would make a bidder indifferent between participating in the auction or forgoing the auction to accept that amount with no uncertainty.

and the 25th and 75th percentiles are \$3.79 and \$43.95, respectively. This reflects the degree to which optimal bid selection in equilibrium allows bidders to insure themselves against risk. The value of the project rises in equilibrium, adding to the certainty equivalent, but this is offset by competition and an inability to profitably skew. Consequently, the certainty equivalent rises for some auctions, falls for others, but all in all stays much the same.

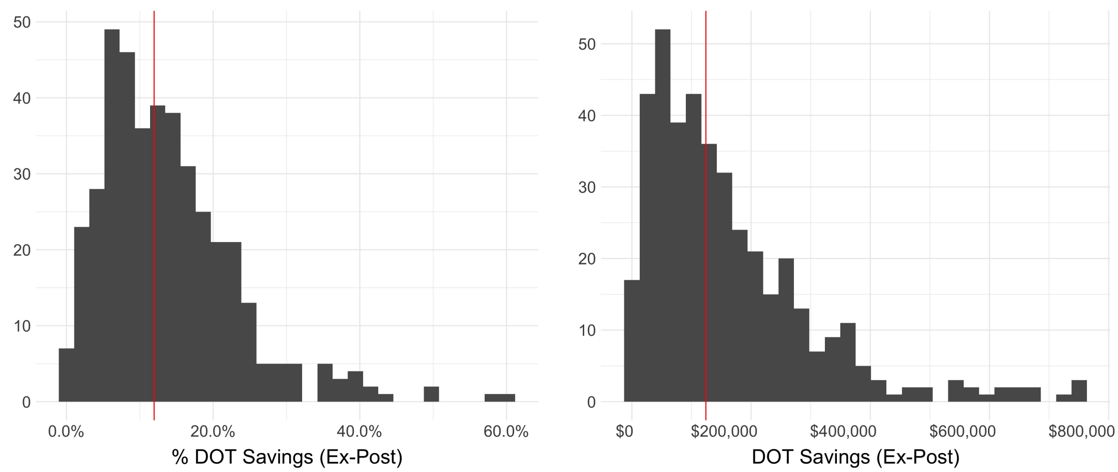

(a) DOT % Savings (b) DOT Dollar Savings

Figure 9: Percent and dollar expected DOT ex-post savings from a counterfactual in which risk is eliminated, relative to a baseline in which bidders accurately predict ex-post quantities, but believe their predictions to be noisy. The median is highlighted by a red vertical line in each case.

The projected expected DOT savings from eliminating risks detailed in Figure 7 reflect the two channels by which eliminating uncertainty changes the bidders' problem: (1) it eliminates risk, raising the value of the project and encouraging more aggressive bids; (2) it gives bidders access to the accurate ex-post quantities, allowing bidders to perfectly optimize their unit bids with respect to ex-post profits. In order to disentangle these two effects, we repeat the counterfactual exercise under the assumption that in the baseline, bidders' quantity projections $\boldsymbol{q}^b$ are equal to the ex-post quantities $\boldsymbol{q}^a$ (but that bidders still perceive the projections to be noisy with variance $\widehat{\boldsymbol{\sigma}}^2$). In this case, bidders always optimize correctly with respect to ex-post quantities, and so the second channel, by which eliminating risk can hurt DOT savings, is shut down.

The resulting expected DOT savings and bidder utility gains are reported in Figure 9. Absent bidder mis-optimization due to inaccuracies in their quantity projections, the median expected saving to the DOT is \$125,175 or 11.9% of the (adjusted) baseline expected cost. This can be thought of as an aggressive estimate of the potential savings from eliminating risk, whereas the previous estimate is a conservative estimate. Notably, the bidder ex-ante utility gains remain modest with a median certainty equivalent gain of \$4.89 from the counterfactual. This is because ex-ante utility is evaluated with respect to bidder beliefs—according to which equilibrium bids are optimized—rather than ex-post quantities. As such, the difference in baseline quantity predictions has little effect on the ex-ante total certainty equivalent of each auction (although it does change the particular choices of optimal bids across items). We present full summary statistics in Table 13 in Appendix E.

9.2 Alternative Risk Sharing Mechanisms: Lump Sum and μ-sharing Auctions

While highway and bridge procurement around the United States is predominately done through scaling auctions, public procurement in other American DOT departments, as well as in DOTs around the world, often employs auction mechanisms that place significantly more risk on contractors. The simplest example of this is a *lump sum* auction in which contractors submit a single total bid for completing the project. Subsequently, the winning contractor is responsible for all project costs incurred, independently of whether or not they exceed initial projections. Lump sum auctions have several properties that make them attractive to DOT officials. First, they require less detailed specifications from DOT engineers, as bidding does not require a comprehensive itemized list of tasks and materials.[80] Second, they incentivize the winning bidder to minimize costs (as all costs are privately incurred and not directly compensated), thereby reducing the scope for moral hazard.

However, lump sum auctions have worrisome incentive properties as well. First, because compensation is fixed at the time of bidding, projects that greatly exceed their scope are more likely to suffer from hold-up problems in which the winning contractor insists on negotiating additional payments before completing the project. Moreover, as we note in Section 4, lump sum auctions greatly increase contractors' exposure to risk. The increased risk exposure reduces the value of winning the auction and causes risk averse bidders to bid less aggressively, resulting in substantially higher costs to the DOT.

In this section, we evaluate the extent to which shifting risk exposure onto contractors, as in a lump sum auction, may be costly to the DOT. To hone in on the effect of risk exposure in particular, we maintain the main assumptions of our baseline model. Bidders are identical apart from a private, independently drawn, efficiency type α. The DOT advertises each

[80] Around 2007, the MBTA—the segment of MassDOT responsible for construction and maintenance of the public transportation system in the Boston area—switched from scaling auctions to lump sum auctions for the majority of its procurement. We spoke to officials about the decision for this transition in 2017. Chief among the reasons was the assertion that the scope of MBTA projects is much more difficult to define (and therefore spec out ex-ante) than the scope of highway and bridge projects. We interpreted this to mean that the difficulty/costs of producing a comprehensive list of items for MBTA projects was high. We sought data to compare costs after the switch, but were unable to obtain bidding or quantity records from before the switch.

project with a comprehensive list of items and (often inaccurate) quantity estimates $\boldsymbol{q}^e$. Bidders receive a common signal of what the ex-post quantities will be, which provides them with a vector of quantity projections $\boldsymbol{q}^b$ and a vector of variances of the projection noise $\boldsymbol{\sigma}^2$.

We define a μ-sharing auction for $\mu \in [0,1]$ as a scaling auction in which the winning bidder is paid

$$\sum_t (\mu q_t^a + (1-\mu) q_t^e) \cdot b_t,$$

upon completion of the project. That is, for every item t involved in the project, the winning bidder is paid her bid b_t multiplied by μ times the actual quantity of t used, plus $(1-\mu)$ times the ex-ante DOT estimate for the quantity of t. When $\mu = 0$, this is equivalent to a lump-sum auction, as the bidder is paid entirely based on her score, $\boldsymbol{b} \cdot \boldsymbol{q}^e$. When $\mu = 1$, this is a standard scaling auction, as in the baseline model. In general, the equilibrium bids for a bidder i with efficiency type α^i are characterized as in Section 6.1, with the following adjustment. The certainty equivalent in the constrained quadratic program to determine the optimal distribution of bids, conditional on a candidate score (as in Equation (6)), is replaced by its μ-sharing analog:

$$\gamma \sum_t \underbrace{(1-\mu) b_t q_t^e + (\mu b_t - \alpha c_t) q_t^b}_{\text{Expected Profits}} - \underbrace{\frac{\gamma \sigma_t^2}{2} (\mu b_t - \alpha c_t)^2}_{\text{Risk Term}}.$$

We defer a detailed derivation of the equilibrium to the appendix. As in the previous section, we calculate the change in expected DOT costs between a baseline auction in which bidders are paid according to the ex-post quantities $\boldsymbol{q}^a$ alone (e.g. $\mu = 1$) and a μ-sharing auction for $\mu \in (0,1]$. In each case, we use the DOT estimates $\boldsymbol{q}^e$, ex-post quantities $\boldsymbol{q}^a$, and blue book costs $\boldsymbol{c}$ from the data—as before—as well as our structural estimates for the CARA coefficient $\hat{\gamma}$ and the distribution of efficiency types conditional on auction characteristics. To focus in on the effect of the risk shifting alone, we shut down the bidder mis-optimization channel and assume that bidders' quantity projections $\boldsymbol{q}^b$ are equal to the actual ex-post quantities $\boldsymbol{q}^a$, but that bidders still perceive the projections to be noisy with variance $\widehat{\boldsymbol{\sigma}}^2$, from our first stage estimation.

We present the percent change in expected DOT costs under a lump sum auction in Figure 10a, and under a $\frac{1}{2}$-sharing auction in Figure 10b. The median expected loss from moving to a lump sum auction is 85.44%, while the median expected loss from a $\frac{1}{2}$-sharing auction is 3.64%. Both distributions exhibit fat tails. Summary statistics for each case are presented in Table 14 in Appendix E.[81]

[81] Small increases in risk may in fact reduce the DOT spending ex-post, as they may cause bidders to place larger bids on items with lower expected over-runså at a competitive score (even if the score itself rises).

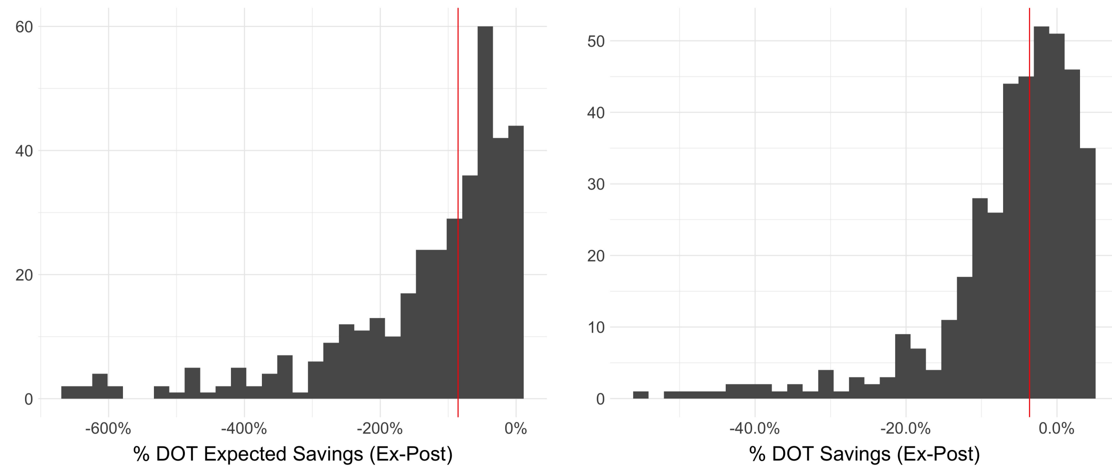

(a) DOT % Cost Change
Lump Sum ($\mu = 0$)

(b) DOT % Cost Change
$\mu = 1/2$

Figure 10: Histogram of expected DOT percent cost change from switching to a μ-sharing auction with $\mu = 0$ and $\mu = 1/2$. The median is highlighted by a red vertical line in each case.

10 Entry

It is well known that an increase to competition benefits an auctioneer. In this section we evaluate a policy to incentivize additional bidders to enter the auctions in our data. First, we estimate the expected amount that the DOT would save if an additional contractor were to enter each auction. We do this by computing the equilibrium bid function in each auction under the baseline (as in the counterfactuals described in Figure 7), and then under an extension of the baseline in which the number of bidders is increased by one. We calculate the expected cost savings in each auction by taking the difference between the expected amount paid by the DOT to the winning bidder in the baseline, and in the counterfactual with an additional bidder participating. Next, we estimate bounds on the cost of entry for a prospective bidder in a procedure akin to Pakes, Porter, Ho, and Ishii (2015), using the assumptions that bidders enter if they anticipate to profit more than the cost of entry and that total entry is set in equilibrium.

10.1 An Equilibrium Model of Entry

Each auction is advertised to a set of prospective (pre-approved) contractors. Upon receiving an advertisement, each prospective bidder observes the common auction characteristics: the location of the project, identity of involved DOT employees, the vector of DOT quantity estimates $\boldsymbol{q}^e$ and the blue book cost estimates, $\boldsymbol{c}$, as well as the refined quantity signals components $\boldsymbol{q}^b$ and $\boldsymbol{\sigma}^2$. Given this information, each bidder is also able to infer the distribution of efficiency types of the prospective contractors.[82] However, in order to discover her own (private) efficiency type, each bidder must invest a fixed amount K. For simplicity, we assume that K is common across bidders. The timeline of each prospective bidder's interaction with the auction is therefore as follows:

1. Bidder observes project characteristics and the entry cost
2. Bidder calculates the expected utility of entering and determines whether or not to participate
3. If she participates:
   - Bidder observes her private efficiency type α
   - Bidder chooses optimal unit bids given α, according to the equilibrium strategy

The expected utility of entry is as follows:

$$E[u(\pi)|N^*] = \int_{\underline{\alpha}}^{\overline{\alpha}} [E[u(\pi(s(\tilde{\alpha}), \tilde{\alpha})|N^*)] \cdot f(\tilde{\alpha})] \, d\tilde{\alpha},$$

where N^* is the equilibrium number of bidders participating in the auction, and $E[u(\pi(s(\tilde{\alpha}), \tilde{\alpha})|N^*)]$ is the expected utility from participating in the auction (and paying K) given efficiency type $\tilde{\alpha}$. In order for N^* to be the equilibrium number of bidders, it must be that the N^*th bidder found it profitable to enter, whereas the $N^* + 1$st bidder did not. That is:

$$E[u(\pi)|N^*] \geq 0 \geq E[u(\pi)|N^* + 1].$$

As such, the certainty equivalent of $E[u(\pi)|N^* + 1]$ (absent an entry cost) provides a lower bound on K, and the certainty equivalent of $E[u(\pi)|N^*]$ provides an upper bound on K.[83] We plot the distribution of upper and lower bounds on the cost of entry K in each auction in Figures 11a and 11b, respectively. In Figure 12, we plot the expected savings to the DOT from the entry of an additional bidder.

[82] As before, we assume that this distribution is the same for all prospective bidders conditional on auction characteristics.

[83] See Lemma 1 in the appendix for a formal proof.

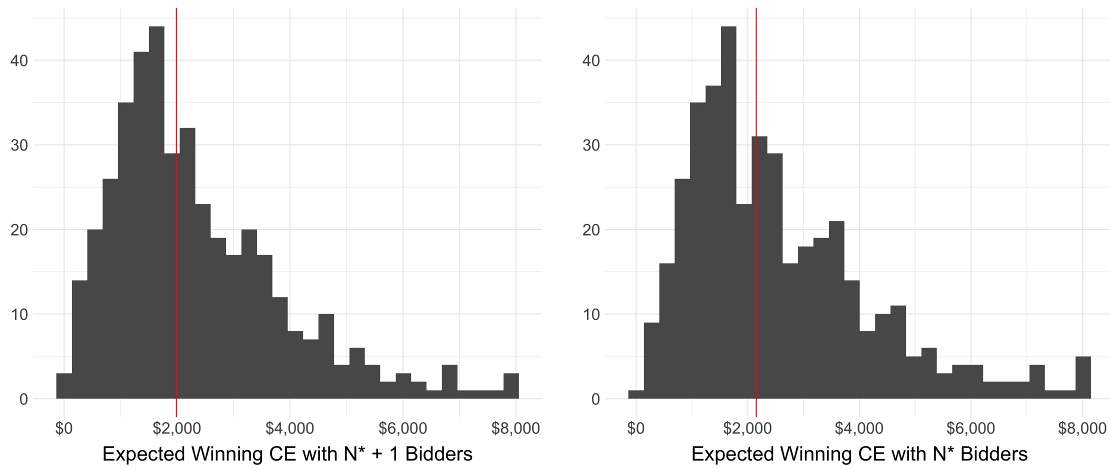

(a) Lower Bounds

(b) Upper Bounds

Figure 11: Distributions of upper and lower bounds on the cost of entry. The median is highlighted by a red vertical line in each panel.

The median lower (upper) bound on entry costs is $1,974 ($2,152), while the median DOT savings amount is $54,750. The distribution of DOT savings is quite fat tailed, however. While the mean lower (upper) bound on entry costs is $2,336 ($2,583), the mean DOT saving is $88,562. This suggests that there is substantial potential value to encouraging entry with a relatively modest guaranteed bonus payment to the winning bidder. Summary statistics are presented in Table 15 in Appendix E.

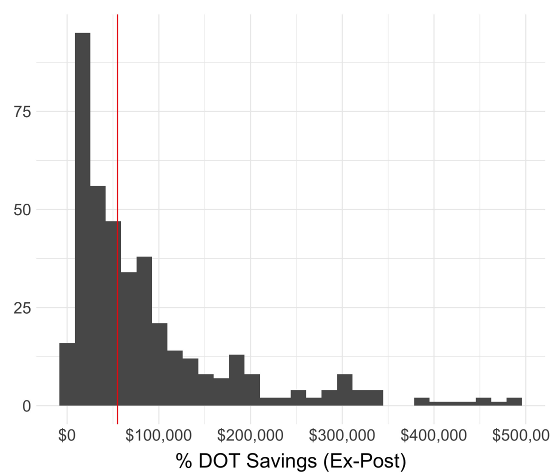

Figure 12: Distribution of the expected dollar savings to the DOT from an entry to each auction. The median is highlighted by a red vertical line.

11 Conclusion

This paper studies the bidding behavior of construction firms that participate in scaling procurement auctions hosted by the Massachusetts Department of Transportation. We develop a model of equilibrium bidding by risk averse bidders that are symmetrically 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 portfolio 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 lower bids. Furthermore, we argue that accounting for risk aversion has significant implications for policy design. Whereas common policies such as investment in improving DOT estimates, switching to a lump sum auction, or subsidizing entry costs would not change MassDOT spending in equilibrium if the bidders are risk neutral, they each have theoretically ambiguous, potentially large consequences for spending when bidders are risk averse.

To assess the effects of these different policies empirically in our context, we estimate the parameters underlying our model. We then simulate the equilibrium bid vector 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 large—it accounts for 12% of payments to the winning bidder in the median auction. However, an effort to reduce uncertainty may not reduce costs very much, as this would also improve bidders' ability to maximize their ex-post revenue. Moreover, switching to a lump sum auction would be very costly—85% more for the median project—as this 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. MassDOT could still do better, however. We find that incentivizing an additional entrant would save over 5% on the median auction. In the online appendix, we evaluate a proposal to enforce a minimum on the unit bids that bidders are allowed to submit and show that this has potential for producing substantial savings. Our framework enables evaluating farther-reaching policies as well, such as an adaptation of the 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.

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

Bhattacharya, V., J. W. Roberts, and A. Sweeting (2014). Regulating bidder participation in auctions. *The RAND Journal of Economics 45*(4), 675–704.

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

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).

Chamberlain, G. (1987). Asymptotic efficiency in estimation with conditional moment restrictions. *Journal of Econometrics 34*(3), 305–334.

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.

De Silva, D. G., T. Dunne, G. Kosmopoulou, and C. Lamarche (2016). Project modifications and bidding in highway procurement auctions. *Working Paper*.

DeWitt, S., G. Yakowenko, T. Bohuslav, T. Ferguson, E. Hoelker, K. Molenaar, G. Schiess, J. Smythe, J. Triplett, and R. Wagman (2005). Construction management practices in canada and europe. Technical report.

Dunning, I., J. Huchette, and M. Lubin (2017). Jump: A modeling language for mathematical optimization. *SIAM Review 59*(2), 295–320.

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.

Häfner, S. (2019). Determining risk aversion in share auctions. *Available at SSRN 3397027*.

Hortaçsu, A. and D. McAdams (2010). Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market. *Journal of Political Economy 118*(5), 833–865.

Kastl, J. (2011). Discrete bids and empirical inference in divisible good auctions. *The Review of Economic Studies 78*(3), 974–1014.

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

Lebrun, B. (2006). Uniqueness of the equilibrium in first-price auctions. *Games and Economic Behavior 55*(1), 131–151.

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. and I. Segal (2002). Envelope theorems for arbitrary choice sets. *Econometrica 70*(2), 583–601.

Pakes, A., J. Porter, K. Ho, and J. Ishii (2015). Moment inequalities and their application. *Econometrica 83*(1), 315–334.

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.

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

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

We construct the unique pure-strategy, monotonic equilibrium of a DOT procurement auction with DOT quantities $\boldsymbol{q}^e$, bidder quantity signals $\boldsymbol{q}^b$, variances $\boldsymbol{\sigma}^2$, DOT cost estimates $\boldsymbol{c}$, and I participating bidders. Each bidder has a privately observed efficiency type α^i that is publicly known to have been drawn from a well-behaved probability distribution over a bounded domain $[\underline{\alpha}, \overline{\alpha}]$. We denote the CDF and pdf of this distribution by $F(\alpha)$ and $f(\alpha)$, respectively.

In particular, for our counterfactual simulations, we assume that α^i is distributed according to a bounded log-normal distribution with a mean that depends on project characteristics, and a project-type-specific variance:

$$\alpha_n^i \sim \text{LogNormal}(\mu_n^\alpha, \sigma_n^{\alpha 2}) \tag{12}$$

where $\mu_n^\alpha = X_n \beta_\alpha$ and σ_n^α is project-type specific.[84] We estimate $\vec{\beta}_\alpha$ and $\vec{\sigma}_n^\alpha$ from the estimated distribution of α types, using Hamiltonian Monte Carlo with MC Stan. We continue to use $F(\cdot)$ and $f(\cdot)$ to refer to the CDF/PDF of this distribution for the remainder of the derivation for notational convenience.

The equilibrium assigns a unique equilibrium score $s(\alpha)$ to each efficiency type α. It is monotonic in the sense that $s(\cdot)$ is strictly increasing in α:

$$\alpha > \alpha' \iff s(\alpha) > s(\alpha'), \text{ for each pair } \alpha, \alpha' \in [\underline{\alpha}, \overline{\alpha}].$$

Under this condition, the probability that $s(\alpha^i)$ is smaller than $s(\alpha^j)$ in equilibrium is equal to the probability that α^i is smaller than α^j, for any α^i and α^j. We can therefore write the equilibrium expected utility of an arbitrary bidder i, using the distribution of α:

$$E[u(\pi(s(\alpha), \alpha))] = \underbrace{\left(1 - \exp\left(-\gamma \sum_{t=1}^{T} q_t^b (b_t^*(s(\alpha)) - \alpha c_t) - \frac{\gamma \sigma_t^2}{2} (b_t^*(s(\alpha)) - \alpha c_t)^2\right)\right)}_{\text{Expected utility conditional on winning}} \cdot \underbrace{(1 - F(\alpha))^{N-1}}_{\text{Prob of win w/ } s(\alpha) = \mathbf{b}^*(s(\alpha)) \cdot \mathbf{q}^e},$$

where N is the number of bidders participating in the auction. In order for $s(\cdot)$ to hold in equilibrium, it must be optimal for every bidder of efficiency type α to submit $s(\alpha)$ as her score. By the envelope theorem, this is ensured when the first order condition of expected

[84] It is not possible to statistically infer both the mean/variance parameters and the bounds of the lognormal domain. As such, we used 0.8 times the lowest observed (estimated) type in each auction as that auction's lower bound, and 1.2 times the highest observed type as the upper bound. For extreme cases, we truncated by 0.5 from below and 3.5 from above. These decisions do not have a qualitative bearing over our results, but do influence the precise counterfactual predictions.

utility with respect to $s(\alpha)$ holds:

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

Evaluating the derivative and rearranging, we characterize the equilibrium score function by the solution to the Ordinary Differential Equation:

$$s'(\alpha) \sum_{t=1}^{T} \left[\left(\gamma q_t^b - \gamma^2 \sigma^2 (b_t^*(s(\alpha)) - \alpha c_t) \right) \frac{\partial b_t^*(s(\alpha))}{\partial s} \right] = [\exp(\gamma \bar{\pi}(\alpha)) - 1] \sum_{k=1}^{N-1} \frac{f(\alpha)}{1 - F(\alpha)}, \quad (13)$$

where $\bar{\pi}(\alpha) = \sum_{t=1}^{T} q_t^b (b_t^*(s(\alpha)) - \alpha c_t) - \frac{\gamma \sigma_t^2}{2} (b_t^*(s(\alpha)) - \alpha^i c_t)^2$ and the bidding function $\mathbf{b}(s(\alpha))$ is optimal (given α). That is, given an equilibrium score $s(\alpha)$, the bidding function solves:

$$\max_{\mathbf{b}(s(\alpha))} \left[1 - \exp\left(-\gamma \sum_{t=1}^{T} q_t^b (b_t(s(\alpha)) - \alpha c_t) - \frac{\gamma \sigma_t^2}{2} (b_t(s(\alpha)) - \alpha c_t)^2 \right) \right] \quad (14)$$

$$\text{s.t.} \sum_{t=1}^{T} b_t(s(\alpha)) q_t^e = s$$

$$b_t(s(\alpha)) \geq 0 \text{ for each item } t.$$

For the counterfactual, we add the further restriction that the optimal bid vector be non-negative. In principle, this restriction should always hold, but we ignored it for the purpose of estimation as all observed bids are positive. For the counterfactual however, it is possible that the optimal unrestricted bids would be negative, and so it is important to include the restriction explicitly. With the additional non-negativity constraint, the convex programming problem in Equation (14) has no closed form solution and must be solved numerically. However, given a solution that determines which of the items have interior bids (rather than zero bids) at the optimum, the solution can be characterized as follows:

$$b_t^*(\cdot) = \max \begin{cases} \alpha c_t + \frac{q_t^b}{\gamma \sigma_t^2} + \frac{q_t^e}{\sigma_t^2 \sum_{t: b_t^*(\cdot)>0} \left[\frac{(q_t^e)^2}{\sigma_t^2} \right]} \left(s(\alpha) - \sum_{t: b_t^*(\cdot)>0} \left[\alpha c_t q_t^e + \frac{q_t^b}{\gamma \sigma_t^2} q_t^e \right] \right) \\ 0 \end{cases} \quad (15)$$

Note that when all items have interior bids, this is equivalent to Equation (7). We solve the ODE in Equation (13) numerically using a state-of-the-art stiff ODE solver using the

DifferentialEquations library in Julia.[85] At every evaluation of Equation (15) in the ODE solver, we compute the optimal bid vector at every score by numerically solving the program in Equation (14) using the IPOPT optimization suite through the JuMP framework.[86] We then compute the partial derivative $\frac{db_t^*(\cdot)}{s}$ using the (analytical) derivative of Equation (15), evaluated at the optimal bids found with the numerical solver.

This ODE is unique up to a boundary condition. To ensure that this indeed characterizes an equilibrium, we require that the highest possible efficiency type $\overline{\alpha}$ submits a score $s(\overline{\alpha})$ and that $s(\overline{\alpha})$ provides zero profit at the optimal bidding strategy. We compute $s(\overline{\alpha})$ numerically using this critereon directly and use it to initialize the ODE solver.

B Technical Details

B.1 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 $b^*_{t,i,n}(\theta)$, as defined by Equation (11). The subvector θ_1 refers to parameters estimated in the first stage, as detailed in Appendix B.1.1. The subvector θ_2 refers to parameters estimated in the second stage, as detailed in Appendix B.1.2. By Assumption 1, the residual of the optimal bid for each item-bidder-auction tuple with respect to its noisily observed bid: $\nu_{t,i,n} = b^d_{t,i,n} - b^*_{t,i,n}(\theta)$, 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.[87]

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 $b^*_{t,i,n}(\theta)$ at the true parameter vector θ.

Items $t \in \mathcal{T}(n)$ are characterized by a $P \times 1$ vector, $X_{t,n}$, of features. Bidders $i \in \mathcal{I}(n)$ are characterized by a $J \times 1$ vector, $X_{i,n}$, of features. The construction of $X_{t,n}$ and $X_{i,n}$ is

[85] See Rackauckas and Nie (2017) for details. We especially thank Chris Rackauckas, the lead developer of DifferentialEquations.jl, for helping us work through numerical issues in getting this to work.

[86] See Dunning, Huchette, and Lubin (2017) for more details.

[87] It is not strictly necessary to assume IIDness across bidders and items. However, allowing for further heterogeneity complicates estimation substantially and so we defer this to a robustness check using Bayesian methods in a future revision.

discussed in detail in Appendix B.2. Estimation proceeds in two stages. In the first stage, we estimate θ_1, the subvector 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.

B.1.1 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 a type of 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 B.2 below. For simplicity, we employ a linear model for the expected quantity of item t in auction n, $\widehat{q^b_{t,n}}$ as a function of the DOT quantity estimate $q^e_{t,n}$ and $X_{t,n}$.[88] In order to model the level of uncertainty in the projection $\widehat{q^b_{t,n}}$, we model the distribution of the quantity model fit residuals ($\eta_{t,n} = q^a_{t,n} - \widehat{q^b_{t,n}}$) 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) with the Stan probabilistic programming language.[89] We then take the posterior modes of the estimated distributions and use them as point estimates for the second stage.

$$q^a_{t,n} = \widehat{q^b_{t,n}} + \eta_{t,n} \text{ where } \eta_{t,n} \sim \mathcal{N}(0, \hat{\sigma}^2_{t,n}) \tag{16}$$

such that

$$\widehat{q^b_{t,n}} = \beta_{0,q} q^e_{t,n} + \vec{\beta}_q X_{t,n} \text{ and } \hat{\sigma}_{t,n} = \exp(\beta_{0,\sigma} q^e_{t,n} + \vec{\beta}_\sigma X_{t,n}). \tag{17}$$

Denote $\theta_1 = (\beta_{0,q}, \vec{\beta}_q, \beta_{0,\sigma}, \vec{\beta}_\sigma, \vec{\beta}_s, \vec{\sigma}_s)$ for the vector of first stage parameters and let $\hat{\theta}_1$ be the posterior modes 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, $\widehat{\sigma}^2_{t,n}$.

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

[89] See Carpenter, Gelman, Hoffman, Lee, Goodrich, Betancourt, Brubaker, Guo, Li, and Riddell (2017) for details on Stan.

B.1.2 Second Stage

Denote $\theta_2 = (\gamma, \alpha^1, \dots, \alpha^I, \beta_\alpha^1, \dots, \beta_\alpha^J)$ 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. Note that θ_2 is $(1+I+J)$-dimensional.

We estimate θ_2 in the second stage, using a GMM framework, evaluated at the first stage estimates $\hat{\theta}_1$:

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

where $g(\theta_2|\hat{\theta}_1)$ is a vector of moments, as a function of θ_2, evaluated at the estimates of θ_1 obtained in the first stage, and W is a weighting matrix. We make use of the following 3 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.[90] The second type of moment states that the average residual of a unit bid submitted for an item labeled as a "top skew item" by the DOT chief engineer's office is zero across auctions. There is one such moment. The third type of moment states that the average residual on a unit bid submitted in each auction is zero, independent of the auction-bidder characteristics of the bidder submitting the bid. There are J such moments—one for each of the auction-bidder characteristics. In total, there are $(1+I+J)$ moments, so that the GMM estimator is just identified. As such, the choice of W does not affect efficiency, and we weight each moment equally as a default.

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

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

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

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, and $\mathcal{T}_s$ as the subset of items that were labeled as "top skew items" by

[90] 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 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.

the DOT chief engineer's office. All moments are formed with respect to the *de-meaned* bid residual:

$$\tilde{\nu}_{t,i,n}(\theta_2|\hat{\theta}_1) = b^d_{t,i,n} - \boldsymbol{\alpha^i_n(\theta_2)} \left(c_{t,n} - \frac{q^e_{t,n}}{\widehat{\sigma}^2_{t,n} \sum_{p\in\mathcal{T}(n)} \left[\frac{(q^e_{p,n})^2}{\widehat{\sigma}^2_{p,n}}\right]} \left[\sum_{p\in\mathcal{T}(n)} c_{p,n} q^e_{p,n} \right] \right)$$
$$- \frac{1}{\boldsymbol{\gamma(\theta_2)}} \left(\frac{\widehat{q^b_{t,n}}}{\widehat{\sigma}^2_{t,n}} - \frac{q^e_{t,n}}{\widehat{\sigma}^2_{t,n} \sum_{p\in\mathcal{T}(n)} \left[\frac{(q^e_{p,n})^2}{\widehat{\sigma}^2_{p,n}}\right]} \left[\sum_{p\in\mathcal{T}(n)} \frac{\widehat{q^b_{p,n}} q^e_{p,n}}{\widehat{\sigma}^2_{p,n}} \right] \right)$$
$$- \frac{q^e_{t,n}}{\widehat{\sigma}^2_{t,n} \sum_{p\in\mathcal{T}(n)} \left[\frac{(q^e_{p,n})^2}{\widehat{\sigma}^2_{p,n}}\right]} \left[s^d_{i,n} \right],$$

where

$$\alpha^i_n(\theta_2) = \boldsymbol{\alpha^i(\theta_2)} + \boldsymbol{\beta_\alpha(\theta_2)} X_{i,n}.$$

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* optimal score, $s^*_{i,n}$. That is, since $s^d_{i,n}$ is composed of noisily observed unit bids, the de-meaned residual $\tilde{\nu}_{t,i,n}(\theta_2|\hat{\theta}_1)$ omits an unobserved score error term:

$$\tilde{\nu}_{t,i,n} = \nu_{i,t,n} - \frac{q^e_{t,n}}{\sigma^2_{t,n} \sum_{p\in\mathcal{T}(n)} \left[\frac{(q^e_{p,n})^2}{\sigma^2_{p,n}}\right]} \bar{\nu}_{i,n}, \tag{18}$$

where

$$\bar{\nu}_{i,n} = -\sum_{t=1}^{T_n} \nu_{t,i,n} q^e_{t,n}. \tag{19}$$

However, as bid residuals $\nu_{t,i,n}$ are assumed to be mean zero and independent of auction and item characteristics, $\mathbb{E}_n[\bar{\nu}_{i,n}]$, and the unobserved score error term is mean zero as well. Thus, the use of demeaned bid residuals does not pose a bias for our GMM estimation procedure.

We compute standard errors for θ_2 at the point estimates of θ_1, without accounting for the uncertainty in the point estimates themselves. In this case, the asymptotic variance of θ_2

follows the standard just-identified GMM form:[91]

$$\sqrt{n}(\hat{\theta_2} - \theta_2^0) \xrightarrow{\text{d}} \mathcal{N}(0, V)$$

where $V = (\Gamma\Delta\Gamma)'$, for

$$\Gamma = \mathbb{E}\left[\frac{\partial g}{\partial \theta_2}(\theta_2^0, \theta_1^0)\right] \text{ and } \Delta = \mathbb{E}\left[g(\theta_2^0, \theta_1^0)g(\theta_2^0, \theta_1^0)'\right].$$

However, as we detail below, we compute the standard errors presented in the text by bootstrap, rather than in-sample asymptotic approximation.

Estimation Procedure

To summarize, we estimate our parameters in a two-stage procedure. In the first stage, we estimate the informational parameters that model bidders' expectations over item quantities and competing scores. In the second stage, we use a two-step optimal GMM estimator to estimate the economic parameters:

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

2. Solve:

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

   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 $\alpha_n^i(\theta_2)$ be non-negative for every i and n.[92]

We calculate standard errors by a bootstrap procedure. In particular, we draw auctions at random with replacement from the total set of auctions in our sample, and repeat the step 2 optimization procedure. We repeat this 1000 times. The confidence interval presented in the results section corresponds to the 5th and 95th percentile of the parameter estimates across the bootstrap draws.

[91] Our standard error calculation does not at present account for uncertainty in the point estimates of θ_1. Accounting for this would modify the asymptotic variance formula by the standard two-step GMM sandwich formula as in Chamberlain (1987).

[92] 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). 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. We provide estimates without the non-negativity constraint as a robustness check. The results do not differ to an economically significant degree.

B.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.[93] 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, and dummies that indicate whether or not the item is frequently used in a single unit quantity, 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.[94] The resulting matrix is very high dimensional, and so we project

[93] 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.

[94] 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)

the matrix onto its principle components, and use the first 15.[95] 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 (10).

To estimate the efficiency type $\alpha_{i,n}$ for each bidder-auction pair, we combine each bidder's unique 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.[96] Our $X_{i,n}$ matrix has a total of 14 columns, and so we have a total of 66 efficiency-type parameters to identify. We use $X_{i,n}$ and the unique bidder ideas to model α_n^i in equation B.1.2.

Finally, we make use of a project-level characteristic matrix X_n in our counterfactuals, in order to parametrize the distribution of efficiency types in each auction. In principle, we could use the bidder-auction matrix $X_{i,n}$ here. However, this would require each bidder to know the identities of her competitors. For the purpose of our main counterfactuals, we focus on the simpler case in which the distribution of scores is homogenous across the bidders participating in a given auction. Therefore, we construct X_n by taking an average of X_n with respect to the bidders in auction n.

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

[96]We define "fringe" similarly to BHT, 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.

C Entry Cost Proofs

Lemma 1. *Consider an auction in which N^* bidders enter in equilibrium given an entry cost K. The cost of entry K is bounded from below by the certainty equivalent of participating in the auction, absent an entry cost, when $N^* + 1$ bidders participate. K is bounded from above by the certainty equivalent of participating in the auction absent an entry cost when N^* bidders participate.*

Proof. We break our proof into two steps. First, we argue that if a bidder of type α prefers to enter an auction at a cost of K, then:

$$(1 - \exp(-\gamma\bar{\pi}(\alpha))) \cdot (1 - F(\alpha))^{N^*-1} \geq 1 - \exp(-\gamma K) \tag{20}$$

where

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

is the bidder's certainty equivalent of profits conditional on winning the auction. This condition states that the bidder's expected utility of participating in the auction absent the entry cost K is at least as large as her utility of "keeping" K and not participating.

To see this, consider a bidder of type α and knows her type, but must still pay an entry fee of K in order to enter a given scaling auction, in which there are $N^* - 1$ opposing bidders. In order for the bidder to prefer to enter the auction, she must expect that her utility upon entering will be higher than her utility otherwise:

$$\mathbb{E}[u(\pi(s(\alpha), \alpha))] \geq 0, \tag{21}$$

$$\begin{aligned}\mathbb{E}[u(\pi(s(\alpha), \alpha))] = {} & \underbrace{(1 - \exp(\gamma K))}_{\text{Utility on entering and losing}} \cdot \underbrace{\left[1 - (1 - F(\alpha))^{N^*-1}\right]}_{\text{Prob of losing}} + \\ & \underbrace{\left(1 - \exp\left(\gamma K - \gamma \sum_{t=1}^{T} q_t^b(b_t^*(s(\alpha)) - \alpha c_t) - \frac{\gamma\sigma_t^2}{2}(b_t^*(s(\alpha)) - \alpha^i c_t)^2\right)\right)}_{\text{Expected utility conditional on entering and winning}} \cdot \underbrace{(1 - F(\alpha))^{N^*-1}}_{\text{Prob of win w/ } s(\alpha)}.\end{aligned}$$

Substituting and rearranging Inequality (21), we obtain that the bidder prefers to enter if and only if:

$$[1 - \exp(\gamma K) \cdot \exp(-\gamma\bar{\pi}(\alpha))] \cdot (1 - F(\alpha))^{N^*-1} + [1 - \exp(\gamma K)] \geq [1 - \exp(\gamma K)] \cdot (1 - F(\alpha))^{N^*-1}$$

Rearranging once more, we obtain:

$$1 - \exp(\gamma K)\left[1 - [1 - \exp(-\gamma \bar{\pi}(\alpha))] \cdot (1 - F(\alpha))^{N^*-1}\right] \geq 0$$

and so,

$$\exp(-\gamma K) \geq 1 - [1 - \exp(-\gamma \bar{\pi}(\alpha))] \cdot (1 - F(\alpha))^{N^*-1}$$

from which we obtain

$$[1 - \exp(-\gamma \bar{\pi}(\alpha))] \cdot (1 - F(\alpha))^{N^*-1} \geq 1 - \exp(-\gamma K).$$

as in Equation (20).

Lower Bound

We now derive a lower bound on K by considering the entry of the $N^* + 1$st bidder, where N^* is the equilibrium number of entrants to the auction given the entry cost K. By definition of N^*, it is unprofitable (in expectation) for the $N^* + 1$st bidder to enter. That is,

$$\int_{\underline{\alpha}}^{\overline{\alpha}} [\mathbb{E}[u(\pi(s(\tilde{\alpha}), \tilde{\alpha}))|N^* + 1] \cdot f(\tilde{\alpha})]\, d\tilde{\alpha} \leq 0,$$

where $\mathbb{E}[u(\pi(s(\tilde{\alpha}), \tilde{\alpha}))|N]$ is the bidder's expected utility from entering given N total entrants (including her) if she turns out to have type $\tilde{\alpha}$, as defined above.

We proceed as follows. Let $\mathbb{E}_\alpha[\cdot]$ denote the integral over α: $\int_{\underline{\alpha}}^{\overline{\alpha}} [\cdot]\, f(\tilde{\alpha}) d\tilde{\alpha}$.

$$\mathbb{E}_\alpha[\mathbb{E}[u(\pi(s(\tilde{\alpha}), \tilde{\alpha}))|N^* + 1]] =$$

$$\mathbb{E}_\alpha\left[(1 - \exp(\gamma K)) \cdot \left(1 - (1 - F(\alpha))^{N^*}\right)\right] + \mathbb{E}_\alpha\left[(1 - \exp(\gamma K)\exp(-\gamma \bar{\pi}(\alpha))) \cdot (1 - F(\alpha))^{N^*}\right]$$

$$= 1 - \exp(\gamma K) \cdot \left(1 - \mathbb{E}_\alpha\left[(1 - F(\alpha))^{N^*}\right] + \mathbb{E}_\alpha\left[\exp(-\gamma \bar{\pi}(\alpha)) \cdot (1 - F(\alpha))^{N^*}\right]\right). \tag{22}$$

Rearranging Equation (22), we have that if $\mathbb{E}_\alpha[\mathbb{E}[u(\pi(s(\tilde{\alpha}), \tilde{\alpha}))]|N^* + 1] \leq 0$, then:

$$1 - \exp(-\gamma K) \geq \mathbb{E}_\alpha\left[(1 - \exp(-\gamma \bar{\pi}(\alpha))) \cdot (1 - F(\alpha))^{N^*}\right]. \tag{23}$$

That is, the utility of having K dollars is greater than a bidder's expected utility of entering the auction at zero cost when there are $N^* + 1$ total entrants. Solving Inequality (23) for K, we obtain that the certainty equivalent of entering the auction at zero cost given $N + 1$

bidders provides a lower bound on the cost of entry.

Upper Bound

We now derive an upper bound on K by considering the entry of the N^*th bidder. By definition of N^* as the equilibrium number of entrants, it is profitable in expectation for this bidder to enter.

$$\int_{\underline{\alpha}}^{\overline{\alpha}} \left[\mathbb{E}[u(\pi(s(\tilde{\alpha}), \tilde{\alpha}))|N^*] \cdot f(\tilde{\alpha})\right] d\tilde{\alpha} \geq 0.$$

Writing $\mathbb{E}_\alpha[\cdot]$ for the integral over α: $\int_{\underline{\alpha}}^{\overline{\alpha}} [\cdot]\, f(\tilde{\alpha})d\tilde{\alpha}$ as before, and rearranging as before, we obtain that we have that if $\mathbb{E}_\alpha[\mathbb{E}[u(\pi(s(\tilde{\alpha}), \tilde{\alpha}))]|N^*] \geq 0$, then:

$$1 - \exp(-\gamma K) \leq \mathbb{E}_\alpha \left[(1 - \exp(-\gamma\bar{\pi}(\alpha))) \cdot (1 - F(\alpha))^{N^*-1}\right]. \tag{24}$$

That is, the utility of having K dollars is lower than a bidder's expected utility of entering the auction at zero cost when there are N^* total entrants. Solving Inequality (24) for K, we obtain that the certainty equivalent of entering the auction at zero cost given N bidders provides an upper bound on the cost of entry. □

D Estimation Results Tables

First Stage Model Fit

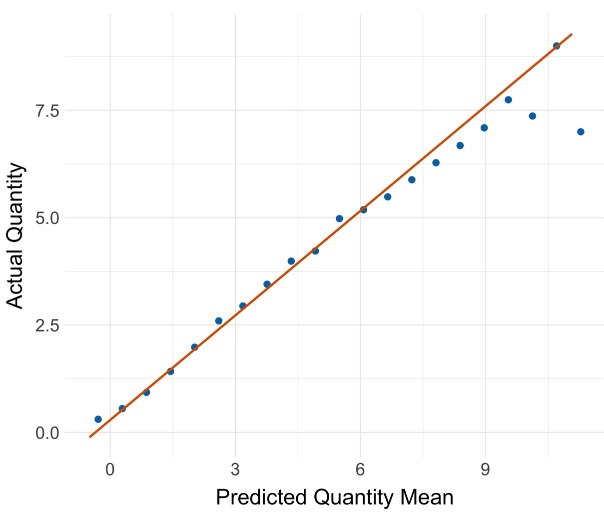

Figure 14: A bin scatter of actual quantities vs model predictions

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

Table 8: Regression report for figure 13

First Stage Parameter Estimates

Parameter	Rhat	n_eff	mean	sd	2.5%	50%	97.5%
$\beta_{0,\sigma}$	1.00	4000	-0.67	0.00	-0.67	-0.67	-0.66
$\beta_{\sigma}[1]$	1.00	1655	-0.05	0.01	-0.06	-0.05	-0.04
$\beta_{\sigma}[2]$	1.00	2120	0.02	0.00	0.01	0.02	0.03
$\beta_{\sigma}[3]$	1.00	3275	-0.02	0.00	-0.03	-0.02	-0.01
$\beta_{\sigma}[4]$	1.00	3516	0.00	0.00	-0.01	0.00	0.01
$\beta_{\sigma}[5]$	1.00	4000	0.02	0.00	0.01	0.02	0.03
$\beta_{\sigma}[6]$	1.00	3131	0.08	0.01	0.07	0.08	0.09
$\beta_{\sigma}[7]$	1.00	2275	0.03	0.01	0.02	0.03	0.04
$\beta_{\sigma}[8]$	1.00	1766	0.00	0.01	-0.01	0.00	0.01
$\beta_{\sigma}[9]$	1.00	1917	-0.01	0.01	-0.02	-0.01	0.00
$\beta_{\sigma}[10]$	1.00	1466	0.03	0.01	0.02	0.03	0.05
$\beta_{\sigma}[11]$	1.00	1952	-0.03	0.01	-0.04	-0.03	-0.02
$\beta_{\sigma}[12]$	1.00	2153	0.02	0.01	0.01	0.02	0.03
$\beta_{\sigma}[13]$	1.00	2590	0.04	0.01	0.03	0.04	0.05
$\beta_{\sigma}[14]$	1.00	2156	0.02	0.01	0.01	0.02	0.03
$\beta_{\sigma}[15]$	1.00	2992	0.00	0.00	-0.01	0.00	0.01
$\beta_{\sigma}[16]$	1.00	1856	-0.16	0.01	-0.18	-0.16	-0.15
$\beta_{\sigma}[17]$	1.00	4000	0.07	0.00	0.06	0.07	0.08
$\beta_{\sigma}[18]$	1.00	4000	0.02	0.00	0.02	0.02	0.03
$\beta_{0,q}$	1.00	4000	0.82	0.00	0.82	0.82	0.83
$\beta_{q}[1]$	1.00	3260	-0.02	0.00	-0.03	-0.02	-0.01
$\beta_{q}[2]$	1.00	4000	-0.01	0.00	-0.02	-0.01	-0.01
$\beta_{q}[3]$	1.00	4000	-0.03	0.00	-0.04	-0.03	-0.02
$\beta_{q}[4]$	1.00	4000	0.02	0.00	0.01	0.01	0.02
$\beta_{q}[5]$	1.00	4000	-0.02	0.00	-0.03	-0.02	-0.01
$\beta_{q}[6]$	1.00	4000	0.01	0.00	0.00	0.01	0.01
$\beta_{q}[7]$	1.00	4000	0.01	0.00	0.00	0.01	0.02
$\beta_{q}[8]$	1.00	2744	-0.03	0.00	-0.04	-0.03	-0.02
$\beta_{q}[9]$	1.00	4000	-0.03	0.00	-0.03	-0.03	-0.02
$\beta_{q}[10]$	1.00	2374	-0.02	0.00	-0.03	-0.02	-0.01
$\beta_{q}[11]$	1.00	4000	0.01	0.00	-0.00	0.01	0.01
$\beta_{q}[12]$	1.00	4000	-0.00	0.00	-0.01	-0.00	0.00
$\beta_{q}[13]$	1.00	4000	0.01	0.00	-0.00	0.01	0.01
$\beta_{q}[14]$	1.00	3366	0.03	0.00	0.02	0.03	0.03
$\beta_{q}[15]$	1.00	4000	0.01	0.00	0.00	0.01	0.02
$\beta_{q}[16]$	1.00	2890	0.01	0.00	0.01	0.01	0.02
$\beta_{q}[17]$	1.00	4000	-0.18	0.00	-0.19	-0.18	-0.17
$\beta_{q}[18]$	1.00	4000	-0.01	0.00	-0.02	-0.01	-0.00

Table 9: First Stage Parameter Estimates

Second Stage Parameter Estimates

	Parameter Estimate	95Pct CI
$\widehat{\gamma}$	0.046	(0.032,0.264)
$\widehat{\beta}[1]$	-0.011	(-0.167,0.137)
$\widehat{\beta}[2]$	-0.003	(-0.105,0.084)
$\widehat{\beta}[3]$	0.027	(-0.138,0.063)
$\widehat{\beta}[4]$	0.017	(-0.142,0.106)
$\widehat{\beta}[5]$	-0.055	(-0.014,0.214)
$\widehat{\beta}[6]$	0.021	(-0.014,0.175)
$\widehat{\beta}[7]$	0.017	(-0.153,0.259)
$\widehat{\beta}[8]$	0.051	(-0.025,0.079)
$\widehat{\beta}[9]$	-0.060	(-0.022,0.063)
$\widehat{\beta}[10]$	-0.006	(-0.151,0.037)
$\widehat{\beta}[11]$	-0.040	(-0.027,0.107)
$\widehat{\beta}[12]$	-0.023	(-0.161,0.152)
$\widehat{\beta}[13]$	0.097	(-0.09,0.233)
$\widehat{\beta}[14]$	0.085	(-0.242,0.176)

Table 10: Parameter estimates for the Second Stage GMM estimation

Second Stage Model Fit

Figure 15: A scatter plot of actual quantities vs model predictions.
Note: Unit bids are scaled so as to standardize quantities so exact dollar values are not representative.

Table 11: Regression report for figure 15

	Dependent variable:
	Data Bid
Predicted Bid	0.992***
	(0.001)
Constant	251.170
	(163.912)
Observations	215,332
R^2	0.879
Note:	*p<0.1; **p<0.05; ***p<0.01

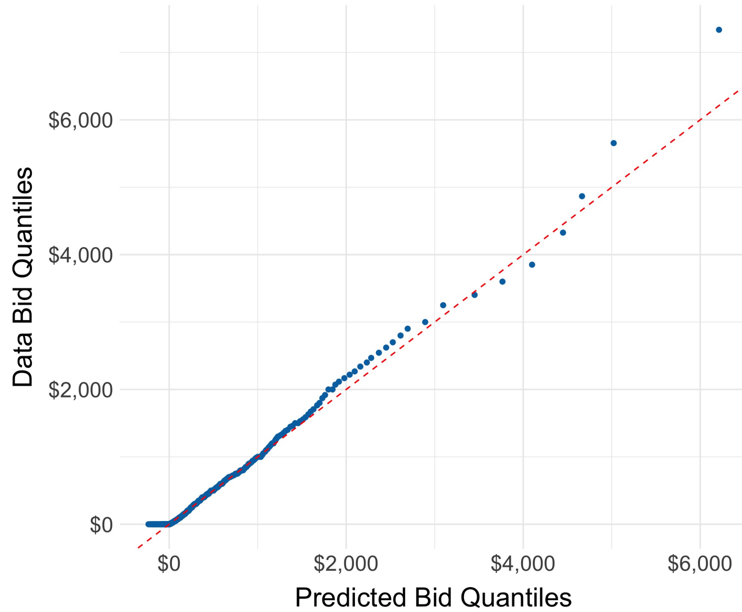

Figure 16: Quantile-Quantile plot of predicted bids against data bids. Quantiles are presented at the 0.0001 level and truncated at the top and bottom 0.01% for clarity. The 45-degree line is dashed in red for reference.
Note: Unit bids are scaled so as to standardize quantities so exact dollar values are not representative.

Equilibrium Winning Score Fit

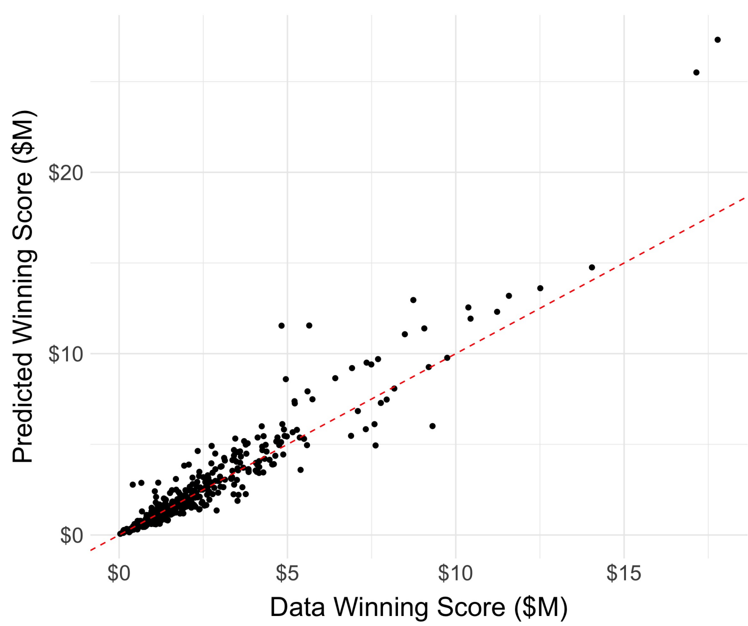

Figure 17: A scatter plot of actual winning scores against the winning scores predicted by our equilibrium simulation at the estimated parameters

E Counterfactual Results Tables

We report the summary statistics for the counterfactual results reported in Section 9. In each case, results are truncated at the top and bottom 1% to exclude extreme outliers from the mean/SD calculations.[97]

Summary Results for Perfect Quantity Predictions by the DOT

Statistic	Mean	St. Dev.	25%	Median	75%
Net DOT Savings	\$2,039.45	\$24,568.34	−\$9,598.27	\$2,171.15	\$13,886.84
% DOT Savings	0.68	4.23	−1.03	0.22	1.58
Bidder Gains	\$6.98	\$144.84	\$3.79	\$17.77	\$43.95

Table 12: Summary of expected DOT percent and dollar savings and bidder utility gains (in dollars) from the counterfactual setting in which the DOT reports perfectly accurate actual quantity estimates.

Summary Results for Perfect Quantity Predictions by the DOT, Shutting Down Baseline Mis-Prediction by the Bidders

Statistic	Mean	St. Dev.	25%	Median	75%
Net DOT Savings	\$173,237.20	\$166,191.20	\$60,630.78	\$125,174.80	\$226,748.60
% DOT Savings	13.67	9.09	7.06	11.90	18.26
Bidder Gains	\$21.18	\$126.79	−\$7.93	\$4.89	\$37.75

Table 13: Summary of expected DOT percent and dollar savings and bidder utility gains (in dollars) from the counterfactual setting in which the DOT reports perfectly accurate actual quantity estimates, relative to a baseline in which bidders accurately predict ex-post quantities, but believe their predictions to be noisy.

[97] These results exclude several projects in each case, for which the ODE solvers were unable to converge under the "standard" setup. Most of these projects are "extremal" in the sense that the imputed lower bound on the α type distribution (0.8 times the lowest estimated type in the auction) is below the truncation point of 0.5. We were able to get almost all of these to converge by changing the truncation point to 0.25. To maintain a clear comparison across projects, we exclude these from the main results. Including them does not make a significant difference—for instance, the mean percent savings in Table 12 changes from 13.67% to 13.74%—and we are happy to report the full results in whatever way is preferred.

Summary Results for a μ-Sharing Auction

DOT % Cost Change	Mean	St. Dev.	25%	Median	75%
Lump Sum	−132.86%	139.39%	−183.97%	−85.44%	−37.04%
$\mu = 1/2$	−6.83%	11.88%	−9.56%	−3.64%	0.29%

Table 14: Summary of expected DOT percent cost change from switching to a μ-sharing auction with $\mu = 0$ (lump sum) and $\mu = 1/2$.

Summary Results for Entry

Statistic	Mean	St. Dev.	25%	Median	75%
Net DOT Savings	$88,562.14	$95,686.34	$22,651.09	$54,749.68	$107,983.00
% DOT Savings	8.72%	8.18%	2.17%	5.67%	12.70%
Entry Cost Lower Bound	$2,336.00	$1,526.39	$1,267.97	$1,973.61	$3,166.82
Entry Cost Upper Bound	$2,583.35	$1,678.04	$1,374.68	$2,152.25	$3,455.62

Table 15: Summary of the welfare impacts of an additional bidder (not paying an entry cost) to each auction.

F Additional Tables and Figures

Shares of Projects with "Unbalanced" Bids

(a) Share of projects (x-axis) that have a particular share of their items breaking the MassDOT overbidding rule (y-axis)

(b) Share of projects (x-axis) that have a particular share of their items breaking the MassDOT underbidding rule (y-axis)

Figure 18

Distribution of Projects by Year in Our Data Estimated Number of Employees for

	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 16: Distribution of projects by year in our data

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 17: 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.

F.1 Robustness Checks

F.1.1 Alternative Specifications

For robustness, we replicate Figure 4a for projects where the top two bidders were within 10% of each other's scores.

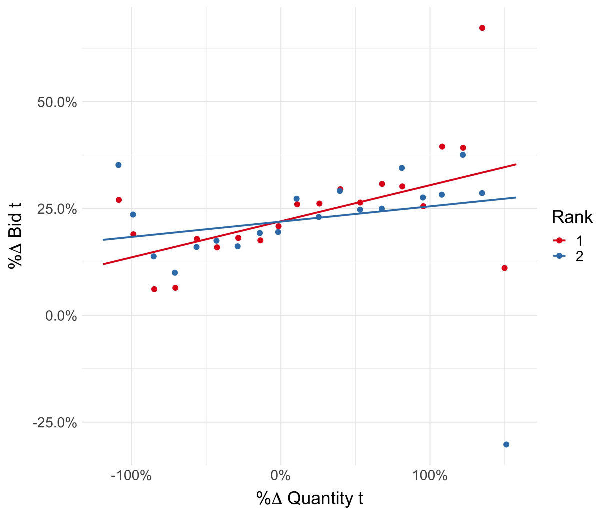

Figure 19: Residualized binscatter of item-level percent overbid by the rank 1 (winning) and rank 2 bidder, against percent quantity over-run—for projects where the top two scores are within 10% of each other

For robustness, we replicate Figure 5a, without controlling for $\%\Delta q_t$:

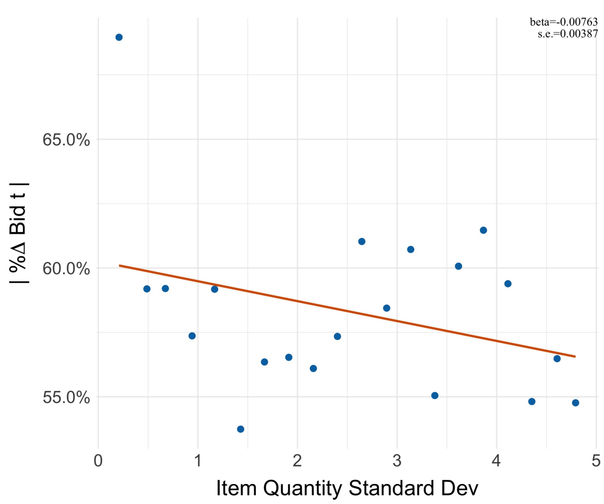

Figure 20: Residualized binscatter of item-level percent absolute overbid against the square root of estimated item quantity variance—without controlling for $\%\Delta q_t$

F.1.2 Bid Level-Weighted Bins

We replicate the main graphs from section 5, weighting the dots by the average bid levels in the bins that they represent. The purpose of this is to argue that outlier dots are generally relatively small, minor items.

Figure 3:

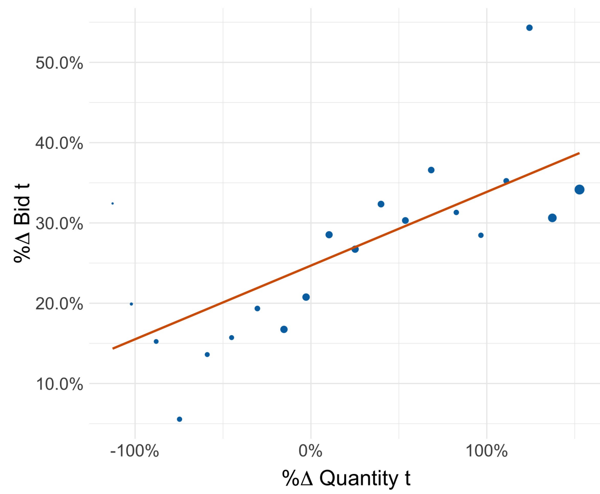

Figure 21: Residualized binscatter of item-level percent winner overbid against percent quantity over-run

Figure 4a:

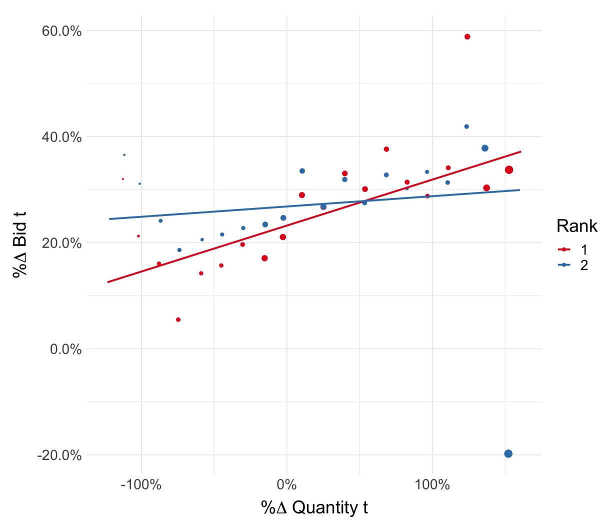

Figure 22: Residualized binscatter of item-level percent overbid by the rank 1 (winning) and rank 2 bidder, against percent quantity over-run

Figure 5a:

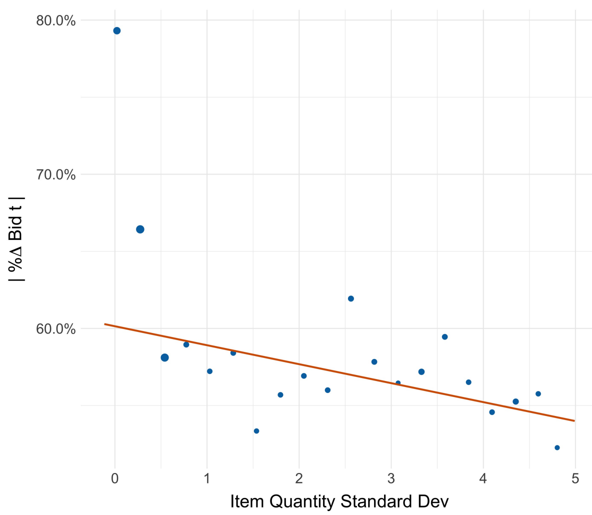

Figure 23: Residualized binscatter of item-level percent absolute overbid against the square root of estimated item quantity variance.

F.2 Counterfactuals Across Values of γ

We replicate our main counterfactuals from Section 9 over a range of CARA coefficients (γ) that covers the 95% interval around our estimate. We truncate at 1% in the scaling counterfactuals and at the 5% for the lump sum and entry counterfactuals or maximum visibility across the range CARA coefficients.[98]

Figure 7a:

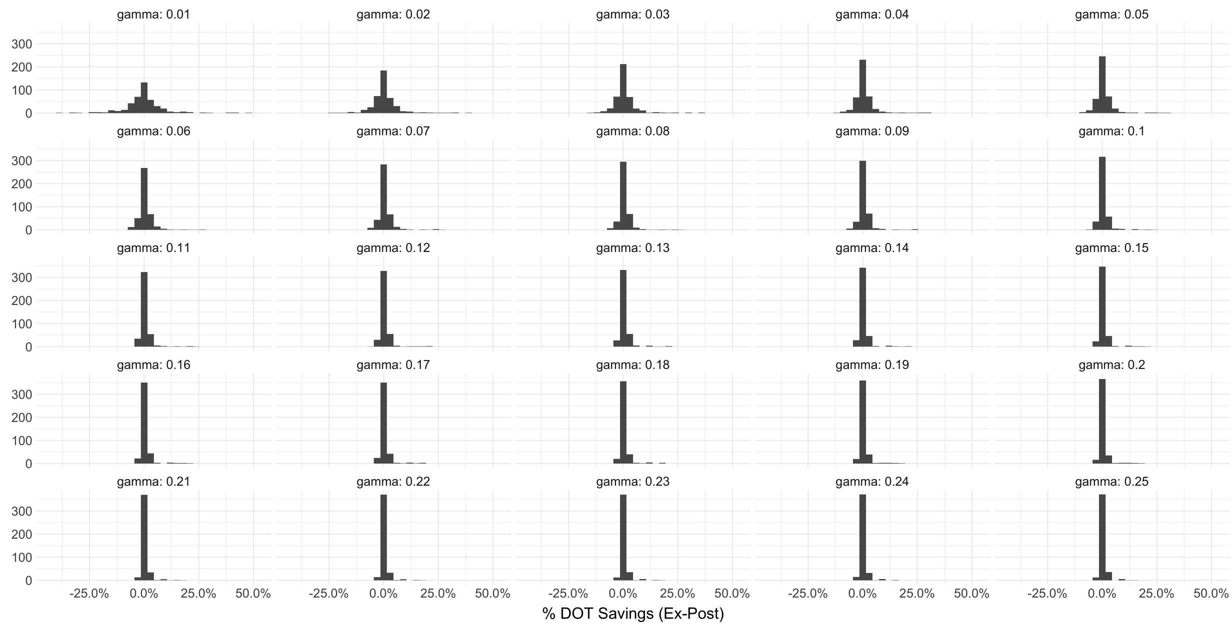

Figure 24: Percent expected DOT ex-post Savings from a counterfactual in which risk is eliminated.

[98]Computing each counterfactual requires solving several ODEs through numerical methods, which may occasionally fail under a "generic" set up. Given the number of ODEs solved for these robustness checks, we did not go through and fix each one. However we have a record of all failure points and can certainly improve this for publication.

Figure 9a:

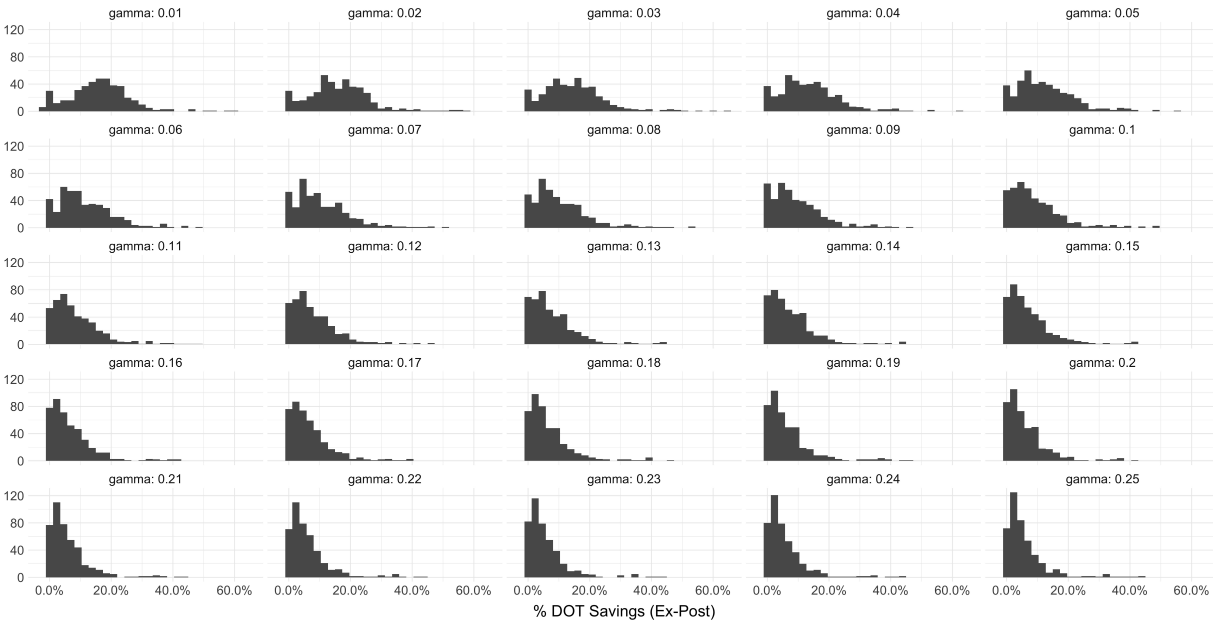

Figure 25: Percent expected DOT ex-post Savings from a counterfactual in which risk is eliminated, relative to a baseline in which bidders accurately predict ex-post quantities, but believe their predictions to be noisy.

Figure 10a with the same scale across values of γ for comparison:

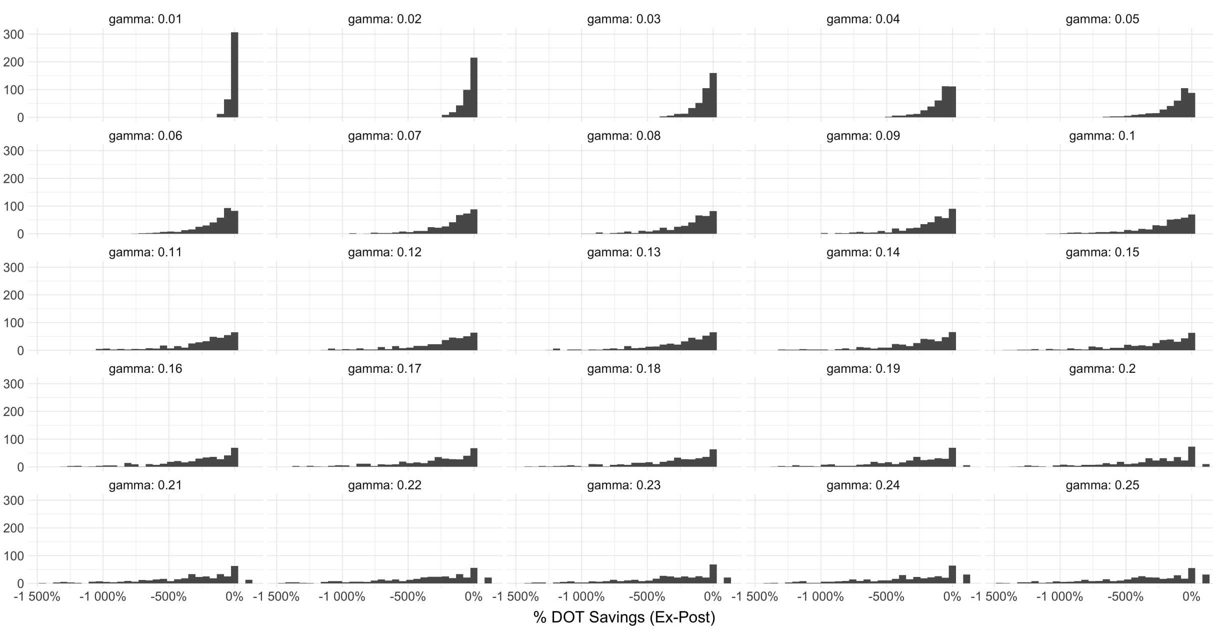

Figure 26: Percent expected DOT ex-post Savings from switching to a lump sum ($\mu = 0$) auction (when both the baseline and counterfactual assume that bidders accurately predict ex-post quantities, but believe their predictions to be noisy).

Figure 10a with the same different scales for visibility:

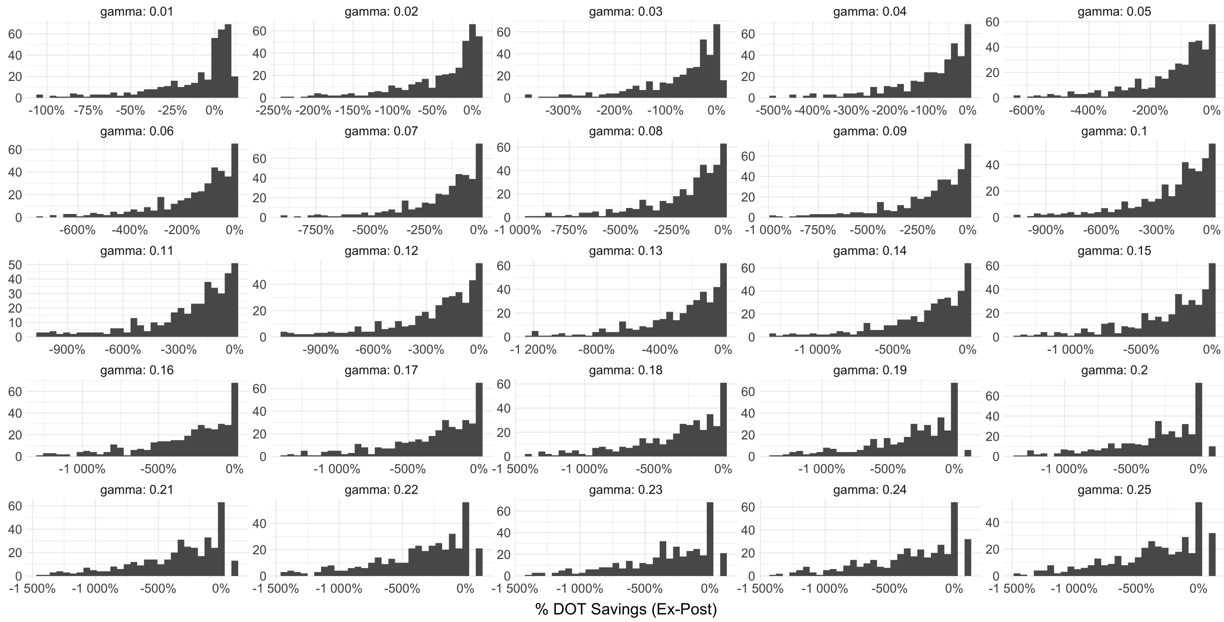

Figure 27: Percent expected DOT ex-post Savings from switching to a lump sum ($\mu = 0$) auction (when both the baseline and counterfactual assume that bidders accurately predict ex-post quantities, but believe their predictions to be noisy).

Figure 11a:

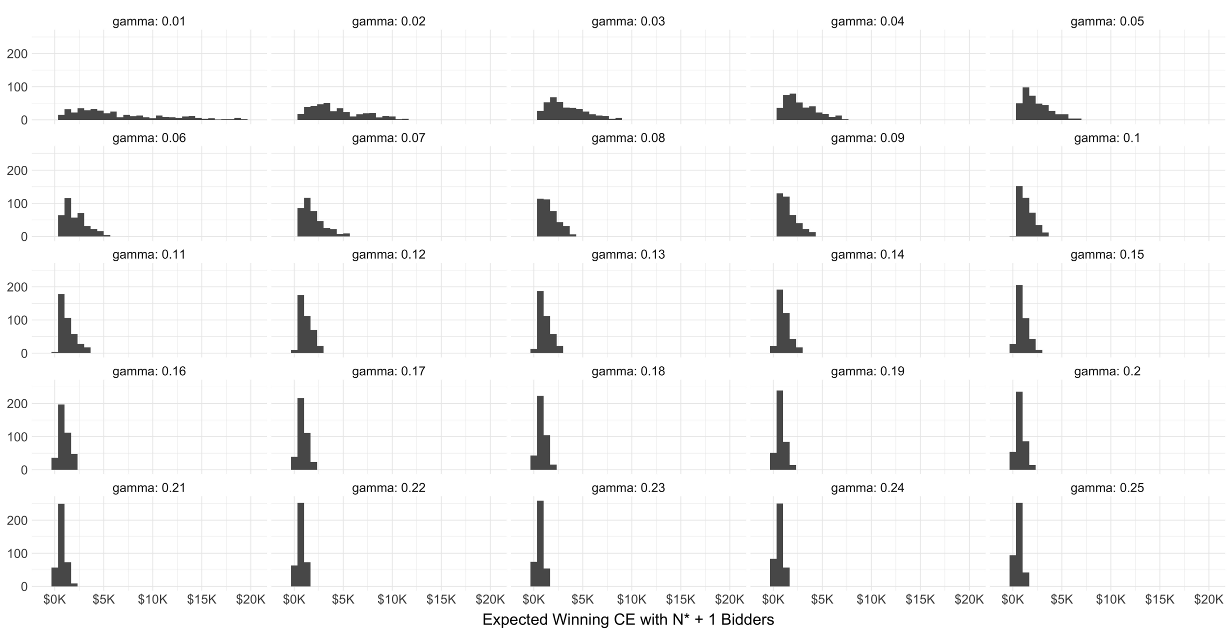

Figure 28: Distributions of lower bounds on the cost of entry

Figure 11b:

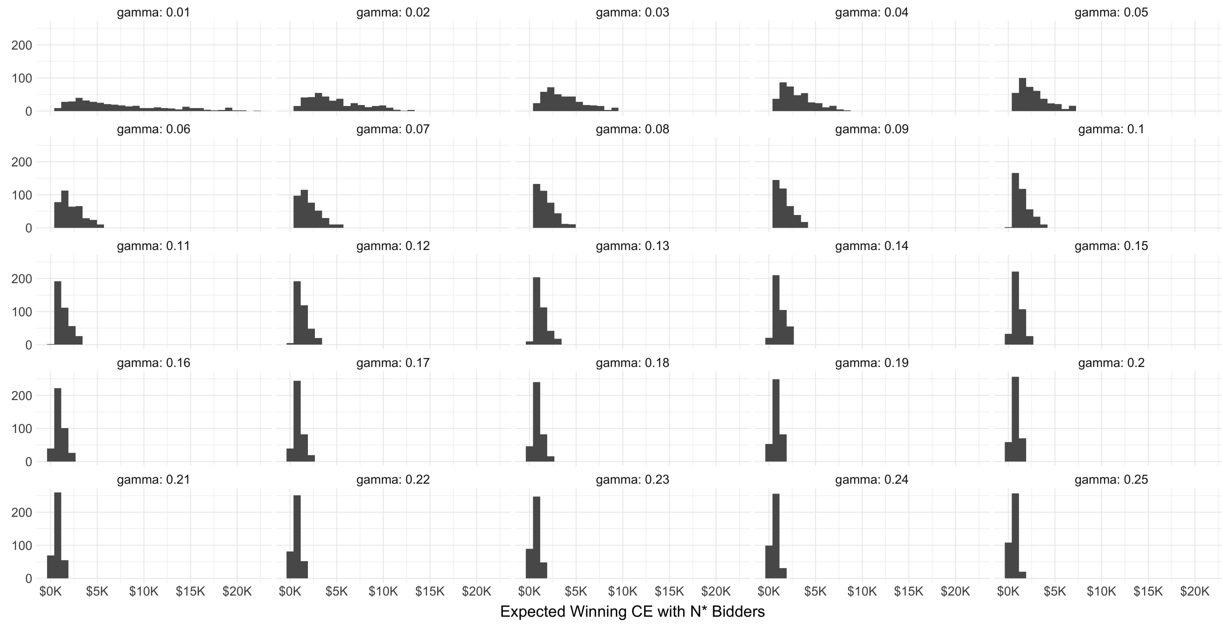

Figure 29: Distributions of upper bounds on the cost of entry

Figure 12:

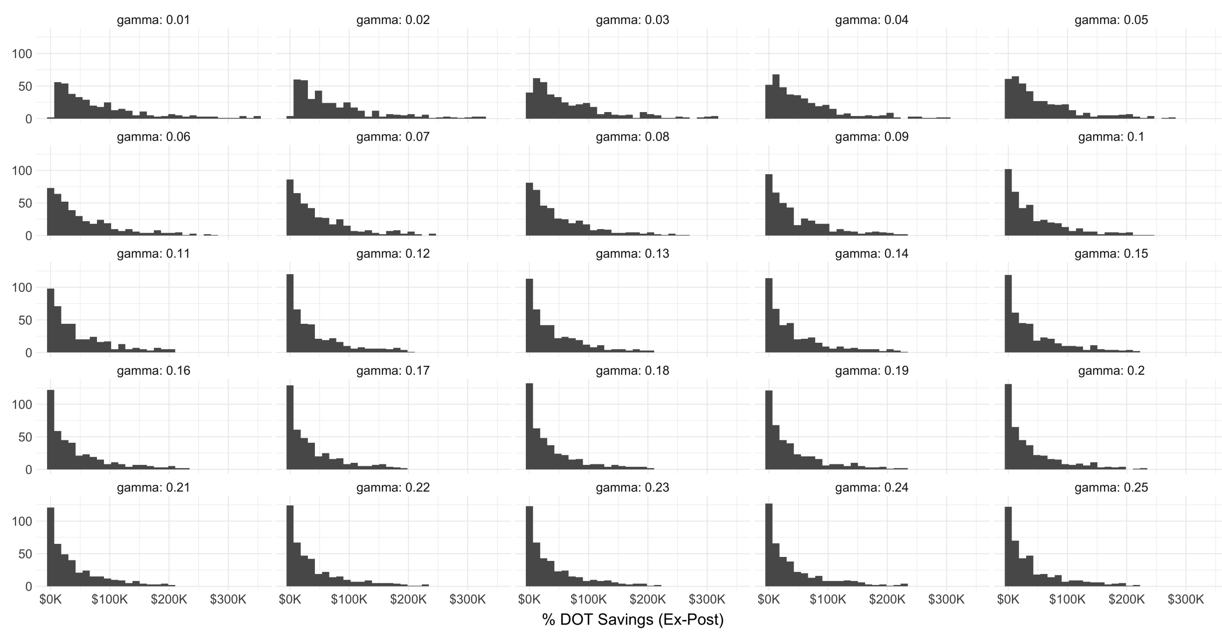

Figure 30: Distribution of the expected dollar savings to the DOT from an entry to each auction

A Online Appendix

A.1 Additional Mechanisms: Minimum Bid Restriction

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 this section, we evaluate how such a rule might impact equilibrium bidding, and ultimately, MassDOT costs. To do this, we simulate the equilibrium of every auction in our data set under the counterfactual condition that every unit bid must be no smaller than 25% of the blue book estimated market rate of that item.[99] Table 18 and Figure 31 present the changes in expected DOT spending as well as bidder surplus that would result from shifting from the (status quo) baseline to the minimum bid requirement counterfactual.

We find that for the median project, the DOT would expect to save $16,590 or 1.63% from imposing a minimum bid restriction of 25% of the market rate. While this result may seem counter-intuitive, an intuition for the mechanism that drives it is as follows. The minimum bid restriction only binds when there are items that bidders are willing to get minimal compensation for. For any such item, there is another item that the bidders overcharge for using the additional "slack" in summing up to the score that makes them competitive toward winning the auction. Furthermore, as in equilibrium, higher cost types (that is less competitive ones) bid weakly higher on each item, the minimum bid restriction binds more often for the lowest cost (most competitive) types.

Once the restriction is added, equilibrium scores do not increase enough to offset the increase in unit bids on "low value" items for competitive bidder types, and these bidders decrease their bids on the "high value" items to compensate. This is beneficial to the DOT on net because the composition of each bidder's bid profile is less risky: both "low value" and "high value" items are bid closer to cost. That is, although (loosely speaking) the certainty equivalent that each bidder gets in equilibrium does not change very much, the balance of expected profits to risk that comprises the (constrained) optimality of that certainty equivalent is de facto shifted to put more weight on minimizing risk and less weight on maximizing expected profits. Thus the DOT's expected cost (which are analogous to expected profit component) are lower.

A.1.1 Minimum Bid Restriction Equilibrium Construction Details

We consider the case in which the unit bid for item t by a bidder is bounded from below – that is, not allowed to be lower than—by a proportion l of the DOT-estimated market rate c_t. In this case, the equilibrium construction follows as in the standard case with one difference: the quadratic

[99] The choice of 25% is not particularly meaningful, and we could easily present the results for 10% or any other percentage.

Statistic	Mean	St. Dev.	Pctl(25)	Median	Pctl(75)
Net DOT Savings	\$26,413.11	\$27,188.83	\$7,343.50	\$16,590.74	\$34,434.38
% DOT Savings	2.81	3.17	0.61	1.63	3.88
Bidder Gains	−\$26.32	\$202.11	−\$27.74	−\$5.94	\$4.27

Table 18: Summary of Counterfactual Results from a Minimum Bid Restriction

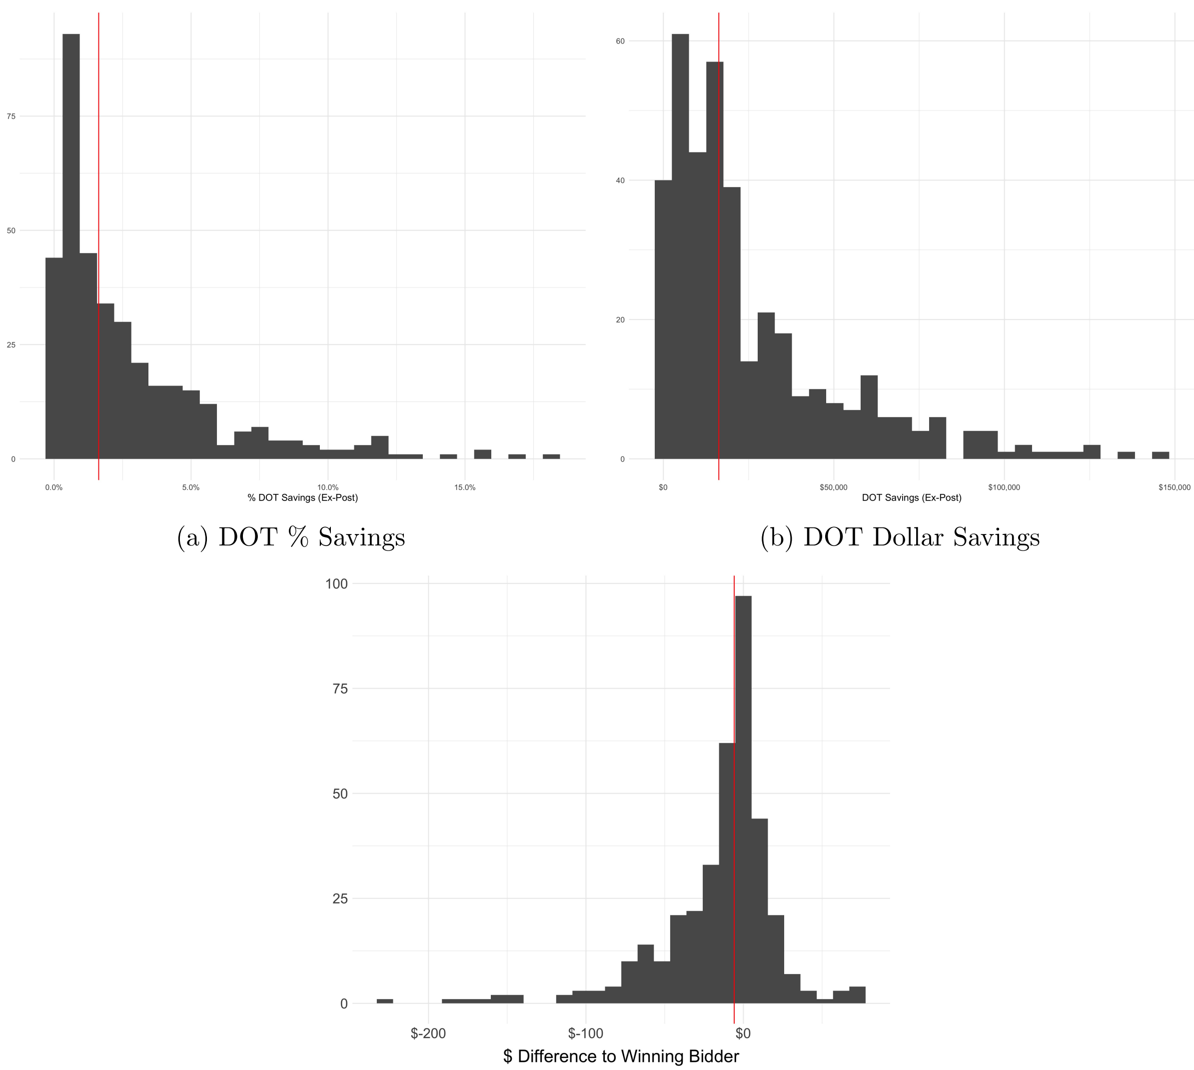

(a) DOT % Savings

(b) DOT Dollar Savings

(c) Winning Bidder Utility Gains

Figure 31: Percent and dollar expected DOT ex-post Savings, and bidder utility gains from a counterfactual in which unit bids are required to be at least 25% of the blue book market rate. The median is highlighted in red in each case.

program that determines the optimal bid vector at each score incorporates a boundary constraint on each element of the bid vector.

In particular, the optimal bid vector $\mathbf{b}$ for a bidder with efficiency type α, considering a score s solves the following program:

$$\max_{\mathbf{b}} \left[\sum_{t=1}^{T} \left(q_t^b (b_t - \alpha c_t) - \frac{\gamma \sigma_t^2}{2} (b_t - \alpha c_t)^2 \right) \right] \tag{25}$$

$$\text{s.t.} \sum_{t=1}^{T} b_t q_t^e = s.$$

$$b_t \geq 0 \text{ for each } t$$

$$b_t \geq l \cdot c_t \text{ for each } t$$

where l is a proportional lower bond limit on the unit bid for item t and c_t is the market rate for item t. Note that for any $l > 0$, the non-negativity constraint is redundant, and so we do not need to consider it separately.

We can formulate the program in system 25 with the following Lagrangian:

$$\mathcal{L} = \sum_{t=1}^{T} \left[\frac{\gamma^2 \sigma_t^2}{2} b_t^2 - \gamma (q_t^b + \gamma \sigma_t^2 \alpha c_t) b_t \right] - \lambda^{\mathcal{L}} \left(\sum_t q_t^e b_t - s \right) - \sum_t \omega_t^l (b_t - l c_t)$$

where $\lambda^{\mathcal{L}}$ is the Lagrangian multiplier on the equality constraint and ω_t^l is the Lagrangian multiplier on the constraint imposing that the bid is above the required limit, l multiplied by the market rate c_t.

The optimal bid vector $\mathbf{b}^* = \arg\min_{\mathbf{b}} \mathcal{L}$ is therefore implicitly defined by the first order conditions:

$$\frac{\partial \mathcal{L}}{\partial b_t}(\mathbf{b}^*) = \gamma^2 \sigma_t^2 b_t^* - \gamma (q_t^b + \gamma \sigma_t^2 \alpha c_t) - \lambda^{\mathcal{L}} q_t^e - \omega_t^l = 0, \text{ for each } t \tag{26}$$

$$\frac{\partial \mathcal{L}}{\partial \lambda^{\mathcal{L}}}(\mathbf{b}^*) = \sum_t q_t^e b_t - s = 0 \tag{27}$$

$$\frac{\partial \mathcal{L}}{\partial \omega_t^l}(\mathbf{b}^*) = -(b_t^* - l c_t) = 0, \text{ if } \omega_t^l > 0. \tag{28}$$

Rearranging equation (26):

$$b_t^* = \frac{(\gamma q_t^b + \gamma^2 \sigma_t^2 \alpha c_t) + \lambda^{\mathcal{L}} q_t^e + \omega_t^l}{\gamma^2 \sigma_t^2}$$

Plugging the definition of b_t^* into (27):

$$\lambda^{\mathcal{L}} = \frac{\gamma^2 s - \sum_t \left[\frac{q_t^e}{\sigma_t^2} (\gamma (q_t^b - \gamma \sigma_t^2 \alpha c_t) + \omega_t^l) \right]}{\sum_t \left[\frac{(q_t^e)^2}{\sigma_t^2} \right]}$$

Note that if $\omega_t^l > 0$ for some t, then $b_t^* = lc_t$. Thus whenever $b_t^* = lc_t$, we have:

$$\omega_t^l = -\gamma q_t^b - \gamma^2 \sigma_t^2 (\alpha - r) c_t) - \lambda^{\mathcal{L}} q_t^e$$

Using this definition of ω_t^l, we can rewrite $\lambda^{\mathcal{L}}$ as follows:

$$\lambda^{\mathcal{L}} = \frac{\gamma^2 s - \sum_{t:b_t^*=lc_t} \left(\gamma^2 q_t^e lc_t\right) - \sum_{t:b_t^*>lc_t} \left[\frac{\gamma q_t^b q_t^e}{\sigma_t^2} + q_t^e \gamma^2 \alpha c_t\right]}{\sum_{t:b_t^*>lc_t} \left[\frac{(q_t^e)^2}{\sigma_t^2}\right]}. \tag{29}$$

We can then plug the definition of $\lambda^{\mathcal{L}}$ from equation (29) into b_t^* to obtain a definition of b_t^* entirely in terms of data and outer variables.

$$b_t^*(\cdot) = \begin{cases} lc_t & \text{if } \omega_t^l = 0 \\ \alpha c_t + \frac{q_t^b}{\gamma \sigma_t^2} + \frac{q_t^e}{\sigma_t^2 \sum_{t:b_t^*>lc_t} \left[\frac{(q_t^e)^2}{\sigma_t^2}\right]} \left(s - \sum_{t:b_t^*=lc_t} (q_t^e lc_t) - \sum_{t:b_t^*>lc_t} \left[\alpha c_t q_t^e + \frac{q_t^b q_t^e}{\gamma \sigma_t^2}\right]\right) & \text{otherwise} \end{cases} \tag{30}$$

Note that the derivative of the optimal bid with respect to the score is given by:

$$\frac{\partial b_t^*(s)}{\partial s} = \frac{q_t^e}{\sigma_t^2 \sum_{t:b_t^*>lc_t} \left[\frac{(q_t^e)^2}{\sigma_t^2}\right]}.$$

A.2 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 underrun. 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*.[100]

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.[101]

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

[100] 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)).

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

work."[102] 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.[103] However, only 2.5% of projects have seen bidders rejected across all justifications, a handful of which were due to unbalanced bids.[104]

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 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).[105] 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 Section 4, there is a theoretical basis to question the relationship between

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

[103] See figures 18a and 18b in the appendix for more details.

[104] 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.

[105] 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.

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.

A.3 Extended Discussion of the Illustrative Example

In this section, we present an extended discussion of equilibrium behavior for the illustrative example discussed in Section 4. The parameters of the example are described in Table 19 below.

	DOT Estimates q^e	Bidders Expect $\mathbb{E}[q^a]$	Noise Var σ^2	Bidder Cost $\alpha \times c$
Concrete	10	12	2	12
Traffic Cones	20	16	1	18

Table 19: Auction parameters from the toy model.

As we discuss in Section 6, this auction game has a unique Bayes Nash Equilibrium. This equilibrium is characterized following the two-stage procedure described in Section 6.1: (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. $\sigma_c^2 = \sigma_r^2 = 0$).

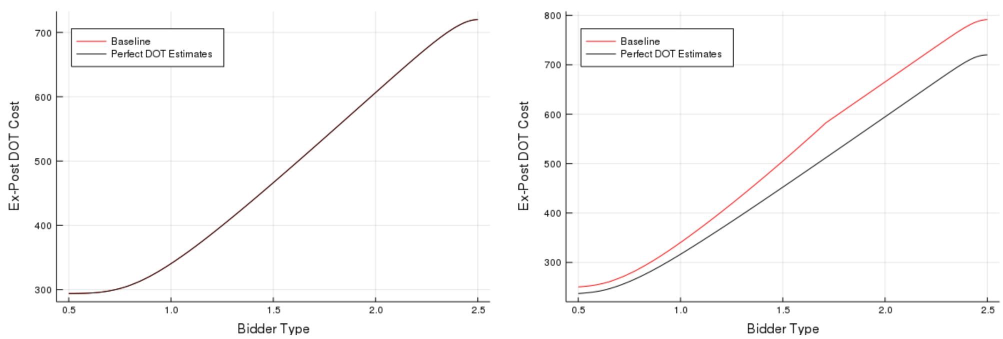

(a) Risk Neutral Bidders (b) Risk Averse Bidders: $\gamma = 0.05$

Figure 32: Equilibrium DOT Cost/Bidder Revenue by Bidder Type

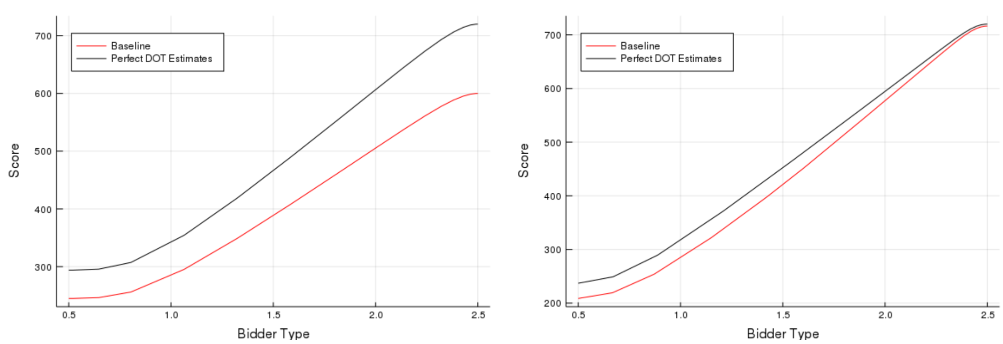

(a) Risk Neutral Bidders (b) Risk Averse Bidders: $\gamma = 0.05$

Figure 33: Equilibrium Score Functions by Bidder Type

In Figure 32a, 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 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 33a, 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.

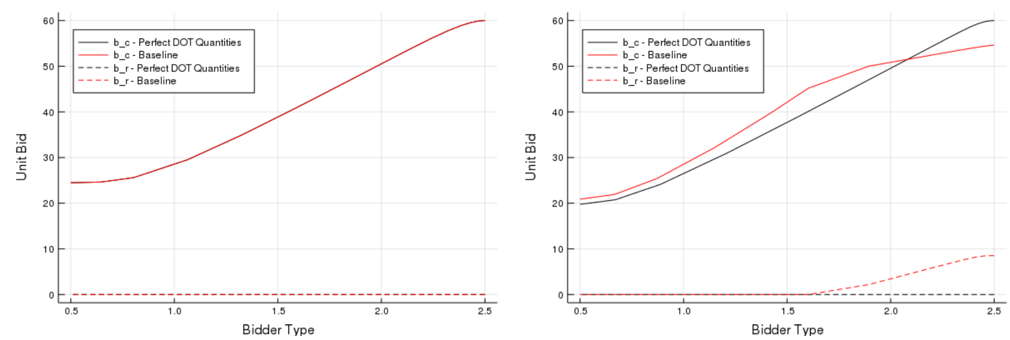

(a) Risk Neutral Bidders (b) Risk Averse Bidders: $\gamma = 0.05$

Figure 34: Equilibrium Unit Bids by Bidder Type

Figure 34a 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 Athey and Levin (2001), 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 Athey and Levin (2001) suggest that risk avoidance is a more likely dominant motive.

In Figures 32b, 33b and 34b, 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 over-bid on one item over the other.

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 1 over different degrees of risk aversion and different levels of uncertainty. In Table 2, we presented 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.[106]

In Table 20, we present the percent difference between the baseline and the counterfactual across CARA coefficients and the magnitude of the quantity noise variance. Each column corresponds to the percent savings to the DOT from the No Quantity Risk counterfactual when the baseline quantity variance term for each item is multiplied by the factor heading the column. For example, in the column labeled 0.5, the baseline equilibrium is computed with $\sigma_c^2 = 0.5 \times 2 = 1$ and $\sigma_r^2 = 0.5 \times 1 = 0.5$. Similarly, the bolded column corresponds to the last column of Table 2, and in the column labeled 2, the baseline equilibrium is computed with $\sigma_c^2 = 2 \times 2 = 4$ and $\sigma_r^2 = 2 \times 1 = 2$.[107]

[106] 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 $\sigma_c^2 = \sigma_r^2 = 0$.

[107] While the savings from eliminating risk are generally higher as prediction noise and risk aversion get higher, the relationship may not always be monotonic. This is because when risk and risk aversion in an auction is very high, bidders are incentivized to bid close to their costs across items so as to minimize their exposure. That is, the variance term overwhelms the prediction term. Note that this is, in part, a result of the CARA functional form.

	Magnitude of Prediction Noise			
CARA Coeff	0.1	0.5	**1**	2
0	0%	0%	**0%**	0%
0.001	0.01%	0.06%	**0.13%**	0.26%
0.005	0.06%	0.32%	**0.64%**	1.30%
0.01	0.13%	0.63%	**1.29%**	2.62%
0.05	0.60%	3.17%	**6.64%**	10.38%
0.10	1.19%	6.42%	**10.71%**	5.65%

Table 20: Percent DOT savings from eliminating quantity uncertainty under different levels of baseline uncertainty and bidder risk aversion

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:

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

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

$$E[\pi(s(\theta^i))|\theta^i] = \left(\frac{12}{10} \cdot s(\theta^i) - \theta^i\right) \cdot \text{Prob}\left(s(\theta^i) < s(\theta^j)\right) \tag{31}$$

$$= \left(\frac{12}{10} \cdot s(\theta^i) - \theta^i\right) \cdot \left(1 - F(\theta^i)\right), \tag{32}$$

where the second equality follows from the strict monotonicity of the equilibrium.[108]

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:

$$E[\pi(\theta^i)] = 12 \cdot \frac{s(\theta^i)}{10} - \theta^i \tag{33}$$

$$= \frac{\int_{\theta^i}^{\overline{\theta}} \left[1 - F(\tilde{\theta})\right] d\tilde{\theta}}{1 - F(\theta^i)}. \tag{34}$$

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

[108]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.

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:

$$s(\theta) = \min\left\{\frac{q_c^e}{q_c^b}, \frac{q_r^e}{q_r^b}\right\} \cdot \left[\theta + \frac{\int_{\theta}^{\overline{\theta}} \left[1 - F(\tilde{\theta})\right] d\tilde{\theta}}{1 - F(\theta)}\right]. \tag{35}$$

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 i wins, she will be paid $q_r^b \cdot b_r^*(\theta^i) = \left[\theta^i + \frac{\int_{\theta^i}^{\overline{\theta}}\left[1-F(\tilde{\theta})\right]d\tilde{\theta}}{1-F(\theta^i)}\right]$ 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 overrun 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.[109] 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.

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