Scaling Auctions as Insurance: A Case Study in Infrastructure Procurement

Click here for the latest version

Valentin Bolotnyy Shoshana Vasserman (JMP) *

November 15, 2018

Abstract

The U.S. government spends about $165B per year on highways and bridges, or about 1% of GDP. Much of it is spent through "scaling" procurement auctions, in which private construction firms submit unit price bids for each piece of material required to complete a project. The winner is determined by the lowest total cost—given government estimates of the amount of each material needed—but, critically, they are paid based on the realized quantities used. This creates an incentive for firms to skew their bids—bidding high when they believe the government is underestimating an item's quantity and vice versa—and raises concerns of rent-extraction among policymakers. For risk averse bidders, however, scaling auctions provide a distinctive way to generate surplus: they enable firms to limit their risk exposure by placing lower unit bids on items with greater uncertainty. To assess this effect empirically, we develop a structural model of scaling auctions with risk averse bidders. Using data on bridge maintenance projects undertaken by the Massachusetts Department of Transportation (MassDOT), we present evidence that bidding behavior is consistent with optimal skewing under risk aversion. We then estimate bidders' risk aversion, the risk in each auction, and the distribution of bidders' private costs. Finally, we simulate equilibrium item-level bids under counterfactual settings to estimate the fraction of MassDOT spending that is due to risk and evaluate alternative mechanisms under consideration by MassDOT. We find that scaling auctions provide substantial savings to MassDOT relative to lump sum auctions and suggest several policies that might improve on the status quo.

*Department of Economics, Harvard University. Email: vbolotnyy@fas.harvard.edu and svasserman@fas.harvard.edu. We are indebted to our advisors, Ariel Pakes, Elie Tamer, Robin Lee, Edward Glaeser, Nathaniel Hendren and Andrei Shleifer for their guidance and support, as well as Eric Maskin, Chris Avery, Shengwu Li, 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 Ashvin Gandhi, Cory Smith, Jonathan Libgober, Myrto Kalouptsidi, Isaiah Andrews, Jeffrey Miron, Hershel Macaulay and the attendants of Harvard seminars for their insightful 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] Infrastructure is also politically popular: voters and policy-makers alike support increasing spending on infrastructure projects by as much as 100% over the next decade.[2] However, infrastructure projects are often complex and subject to unexpected changes. Uncertainty can be costly to the firms implementing construction—many of whose businesses are centered on public works. The mechanisms used to procure construction work can play a key role in mitigating firms' exposure to risk. Limiting risk makes prospective contracts more lucrative to firms and increases competition, thereby reducing tax payer expenditures.

In this paper we study the mechanism by which contracts for construction work are allocated by the Highway and Bridge Division in the Massachusetts Department of Transportation (MassDOT). As in 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 DOT project designers. The winner is then paid based on the quantities ultimately used in completing the project.

A common concern among policy-makers is that bidders may extract rents from the DOT by "skewing" their bids: placing high unit bids on items that will over-run the DOT estimates and low unit bids on items that will under-run. Bid-skewing has been documented as far back as 1935, and referred to as commonplace as recently as 2009 (Skitmore and Cattell (2013)). Previous work on timber auctions (Athey and Levin (2001)) and highway construction (Bajari, Houghton, and Tadelis (2014)) has demonstrated evidence that bidders skew correctly on average and that the most competitive bidders skew in a similar way. This suggests that competitive bidders are similarly able to predict which items will over/under-run.

As we demonstrate, the markup charged to the DOT depends not only on the level of competition in the auction, but also on the uncertainty about the ultimate needs of a project—conditional on the DOT's specification—as well as the degree of risk aversion that contending bidders face. If bidders are risk neutral and equally informed, bid-skewing

[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.

[2] Recent polls have consistently shown around 70% of voters in support of increased infrastructure spending along the lines of the $1.5 trillion plan outlined by the Trump administration. See YouGov and Gallop for example. A major infrastructure bill is expected to entertain bi-partisan support following the 2018 election (Nilson (2018)). This is in addition to a 2015 bill passed with bi-partisan support to increase infrastructure spending by $305 billion over five years. ((Berman (2015)).

produces no additional cost to the DOT in equilibrium. Contractors choose their bids using refined quantity estimates, and any information rent is competed away. Risk averse bidders, however, use bid skewing to balance the uncertainty in a project across the different items involved. As in the risk neutral case, bidders generally submit higher bids for items they believe will over-run the DOT quantity estimates. However, the incentive to raise bids on items predicted to over-run is dampened by the level of noise in this prediction. Moreover, the risk lowers the value of a project to bidders, causing them to bid less aggressively and consequently extract higher payments from the DOT.

Notably, risk averse bidders will generally submit *interior bids*—unit bids that are above zero—whereas risk neutral bidders will submit "penny" bids—unit bids of essentially zero—on all but the items that are predicted to overrun by the largest amount, absent an external force to prevent this.[3] This matches the observations in our data, in which the vast majority of unit bids are interior, but no significant penalty for penny bidding has ever been exercised.[4]

Moreover, taking uncertainty and risk aversion into account has significant implications for comparisons across auctions. Risk neutral bidders would profit identically under a scaling auction, a lump sum auction – in which bidders bid a total project price and are responsible for all realized costs – or anything in between. Risk averse bidders, however, are sensitive to the differences in risk exposure under each of these mechanisms. Scaling auctions compensate bidders for every unit that is ultimately used. As such, the only risk that bidders are exposed to (upon winning the auction) is the risk that they "mis-optimized" in selecting their bid spread across items given the ex-post quantity realizations. Under a lump sum auction, however, bidders bear the entirety of the cost risk involved in the project. If the realized quantities are substantially larger than the predicted values used during bidding, the winning bidder is liable for the differences, with no further compensation.[5] In equilibrium, bidders will insure themselves against the risk that they face by submitting higher overall bids. Thus, scaling auctions, in which the level of risk from uncertainty about the ex-post quantities in a project can be minimized by the bidders, are predicted to produce substantially lower overall costs to the DOT.

Our contributions are three-fold. First, we construct a parsimonious model of competitive bidding in a scaling auction with risk averse bidders who shoulder uncertainty over the quantities that will ultimately be used. We show that risk aversion and uncertainty are

[3]See section 2 for a discussion of the model predictions under risk neutral and risk averse bidders.

[4]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.

[5]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.

sufficient to explain interior bids, in contrast to previous work, which relies on heuristic penalties on penny bids. We then provide reduced form evidence that the bidding behavior observed in our data is consistent with the predictions of our model. Furthermore, we demonstrate how our model can be used to evaluate the cost that the DOT incurs due to uncertainty in its project specifications. Second, we estimate a structural model for uncertainty and optimal bidding in our data. We employ a two-stage procedure to estimate the level of risk in each project, the degree of risk aversion, and the distribution of bidders' private costs. In the first stage, we estimate a model of bidder uncertainty using the history of predicted and realized item quantities. In the second stage, we use the specification of equilibrium unit bids implied by our model to construct a GMM estimator for risk aversion and bidder cost types. Third, we use our structural estimates to simulate counterfactual auction equilibria in which: (1) the DOT eliminates all uncertainty about item quantities; (2) the DOT employs a μ-risk-sharing auction in which it compensates bidders for μ times the prespecified estimated quantity and $1-\mu$ times the realized quantity of each item used. Finally, we calculate bounds on the cost of entry for an additional bidder to each auction, as well as the cost savings to the DOT from an additional entry.

Using the first counterfactual results, we assess the DOT's cost from uncertainty by taking the difference in the expected amount paid to the winning bidder in the baseline auction (the auction used in the status quo) and in the counterfactual setting with all uncertainty removed. We find that the DOT's cost in the baseline auction is only \$2,145—or 0.70%—higher, on average, than in the counterfactual auction with no uncertainty. However, this estimate reflects the sum of two opposing forces that are shifted by the counterfactual: risk and prediction. In the baseline, bidders use a best prediction (given available information) of the ultimate item quantities. These predictions may be inaccurate in-sample, and so the bids submitted may not be optimal (from the bidders' perspective, after observing the ex-post quantities). By contrast, in the counterfactual setting with all uncertainty eliminated, bidders know the exact quantities that will be used and optimize accordingly. Consequently, the DOT winds up paying more than in the baseline for some projects. To isolate the effect of risk itself, we repeat the counterfactual exercise under the assumption that bidders' quantity predictions are correct (but bidders still interpret these predictions as coming from noisy signals as before) in the baseline. In this case, there is no bidder mis-optimization in the baseline, and so the DOT strictly saves money from eliminating risk: \$172,513 (13.74%) on average.

Using the second set of counterfactual results, we assess the extent and direction to which DOT costs would change if the DOT switched from the scaling auction to an alternative in

which part or all of the amount paid to the winning bidder is fixed at the time of bidding.[6] A mechanism of this sort curbs bidders' ability to skew their bids: in the limiting case of a lump sum auction, bidders are paid the amount that they bid and so, there is no advantage to spreading unit bids across items in any particular way. It may also offer benefits to the DOT by reducing its burden in project specification and budgeting flexibility.[7] However, mechanisms of this sort effectively shift risk from the DOT to the bidders. As such, they lower the expected value of winning each auction and induce higher, less aggressive bids. We estimate that switching to a lump sum auction would increase DOT costs by 128% on average (85% on median). The losses do not scale linearly with the amount of risk sharing, however. We estimate that if the DOT were to pay the winning bidder her unit bid multiplied by a 50-50 split of the ex-ante and ex-post quantities for each item, costs would increase by 6.84% on average (3.47% on median).

Finally, while major improvements to quantity estimation may be difficult to achieve across the board, efforts to increase competition may offer an additional channel to improve DOT cost efficiency. We estimate that adding an additional bidder to each auction results in an average DOT savings of $82,583 (8.90%). Furthermore, our estimates of lower bounds on bidders' cost of entry suggest that an increased (guaranteed) payment of as little as $2,316 (on average) could incentivize an additional entry.

Our analysis is enabled by a rich and detailed data set, provided to us by the Highway and Bridge division of the Massachusetts Department of Transportation. For each auction in our study, we observe the full set of items involved in construction, along with ex-ante estimates and ex-post realizations of the quantity of each item, a 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. Furthermore, our setting is particularly conducive to the study of risk aversion. Bridge maintenance projects are highly standardized, and so heterogeneity across projects is well captured by the characteristics observed in our data. The winner of each auction is determined entirely by the expected cost of the project given the bidder's unit bids. Participating bidders are all pre-qualified by the DOT and neither historical performance, nor external quality considerations play a role in the allocation of contracts. In addition, while there is substantial variation between the ex-ante DOT estimates and the

[6]To highlight the effects of the counterfactual policies themselves, we report the results of all counterfactuals assuming that bidders have correct quantity estimates (but still interpret these estimates with uncertainty) in both the baseline and the counterfactual. Results in the case that bidders use the quantity predictions estimated in our first stage are similar. We report them in the appendix for robustness.

[7]Neither moral hazard nor hold up problems are considered in our model. Moral hazard might make lump sum auctions more attractive as imposing more risk on bidders would induce more thrifty uses of material. However, the extent of moral hazard is limited by the contractors' ability to influence quantities given DOT restrictions and supervision. Hold up problems would strengthen our results.

ex-post realizations of item quantities, all changes to the original project specification must be approved by an on-site DOT project manager or engineer, limiting the scope of moral hazard. Finally, while previous work on highway procurement auctions has discussed the role of ex-post renegotiation of unit-prices and a disincentive for bid skewing due to a possibility of having a winning bid rejected by the DOT, neither of these forces is applicable in our setting. Unit price renegotiation occurs in a negligible number of cases in our data, and MassDOT does not reject the winning bidder as a matter of policy.[8]

Connections and Contributions to the 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).[9] Athey and Levin (2001) first established the theoretical framework demonstrating that bid-skewing arises in equilibrium when bidders are better informed about ex-post quantities than the auctioneer. Using a general modeling framework, Athey and Levin establish a number of empirical predictions and test them in the context of US timber auctions. Notably, they test the hypothesis that bidders have superior information (beyond what is given to them by the auctioneer) by comparing the direction of bid skews: profitable skews are indicative of superior information. They find significant evidence of superior information, as well as evidence that there is little informational differentiation between the top two bidders. We discuss analogous exercises in our reduced form section and find similar results in our setting as well. Furthermore, Athey and Levin note that the absence of total skewing (e.g. penny bidding) in their setting is inconsistent with risk neutral bidders in their model, and suggest risk aversion as a more fitting explanation for what they observe. Using the Athey and Levin framework, we construct a structural model that allows us to quantify the costs – realized and hypothetical – of scaling auctions in practice.

Bajari, Houghton, and Tadelis (2014) ("BHT") studies a setting similar to ours: the auctions used to procure highway construction contracts in California. As in our setting, BHT observe item-level unit bids submitted in a scaling auction in which awards are allocated based on engineers' quantity estimates, but compensated based on realized quantities.[10]

[8] In a handful of cases, MassDOT has withdrawn the auction all together after receiving bids, citing internal mis-estimation in the project specification, and has re-posted the auction anew after making adjustments. The same bidders were eligible to participate in the revised auction.

[9] More recently, De Silva, Dunne, Kosmopoulou, and Lamarche (2016) apply a framework similar to Bajari, Houghton, and Tadelis (2014) to assess the effects of a DOT's commitment to reducing the scope of project changes.

[10] There are several notable differences between the setting in Bajari, Houghton, and Tadelis (2014) and ours. Unlike MassDOT, the California DOT imposes tighter limits on quantity overruns, and does occasionally reject bidders with mathematically unbalanced bids. Furthermore, while the overall level of bid

However, the study's main focus is on adaptation costs—costs incurred from disruptions in work-flow due to inadequate preliminary planning. BHT propose a structural model for bidding in which bidders are risk-neutral and have correct (on average) expectations over what the final quantity of each item used will be. They then use conditions derived from this model in conjunction with data on ex-post negotiated change orders, adjustments to unit prices, extra work-orders, and deductions (due to failures on the part of the contractor) to identify the expected cost of these adjustments that is paid by the California DOT.

Our paper differs from Bajari, Houghton, and Tadelis (2014) in several significant ways. First, because BHT is primarily concerned with measuring adaptation costs, it does not aim to predict bids at the item level. By contrast, our paper is focused on predicting bids for auctions in counterfactual settings. Our approach incorporates risk and risk aversion to rationalize interior bids, allowing us to capture substitution patterns between items with a micro-founded generative model of unit bid setting.[11] Our model characterizes equilibrium bids at the auction-bidder-item level as a function of the item's historical quantity variance, the bidder's private cost type and distribution of opponent types, and the level of risk aversion in the auction. Our identification strategy leverages variation in unit bids across auctions that each bidder participated in, as well as variation across auctions that items identified by the DOT as "highly skewed" appeared in.[12]

Our approach allows us to estimate the distribution of bidder cost types in each auction, as well as the coefficient of bidder risk aversion. These parameters, along with those governing the item quantity distributions, jointly characterize the equilibrium bid distribution in each setting.[13] Using our estimates, we are able to predict the equilibrium bids that would arise in each counterfactual. We can thus assess policy-relevant outcomes: the expected cost to the DOT, as well as the utility to prospective bidders (which may impact entry). To our knowledge, no counterfactual analysis of scaling auctions, nor any assessment of their performance in the context of mechanism design has been done before.

skewing, as evidenced by the relationship between quantity overruns and price overruns (as in figure 7) across all highway and bridge projects in Massachusetts is similar to that in California, this relationship is particularly pronounced among the bridge maintenance projects that our analysis focuses on.

[11] Bajari et al. model bidders as risk neutral, but subject to a heuristic penalty function in bidders' utility that convexly penalizes deviations of unit bids from the DOT's cost estimates for each item. They estimate that the penalty coefficient is small and negative (suggesting bid-skewing is encouraged, rather than penalized), but not statistically significant. As part of preparing our paper, we replicated their methodology on our data set, and found our model substantially better in back-predicting item bid spreads.

[12] By contrast, Bajari et al. use aggregate optimality conditions such that each observation entering their moment condition is at the bidder-auction level. They then estimate a mean cost type across all bidders and auctions, as well as mean coefficients on adaptation costs, etc.

[13] Note that one cannot evaluate counterfactual outcomes by extrapolating from the empirical score distribution. Changes to the auction setting will change the equilibrium score distribution, and so it is necessary to compute the equilibrium from primitives in each counterfactual.

More generally, our paper relates to the literature on multi-dimensional auctions, and scoring auctions in particular (auctions in which bids on different dimensions of interest are aggregated into a single-dimensional score to determine the winner). Che (1993) characterizes the equilibria of auctions that employ a two-dimensional scoring rule (quality and price) with single-dimensional bidder types. Asker and Cantillon (2008) extend this to a more general setting, allowing for multi-dimensional bidder types and general quasi-linear scoring rules, by showing that a mapping of multi-dimensional attributes onto a single dimensional "pseudo-type" is sufficient to characterize equilibria up to payoff equivalence. Both of these papers assume that bidders are risk neutral, and the result on "pseudo-types" does not extend directly to the risk averse case. In our paper, we model single-dimensional bidder types, as this is the most parsimonious way to ensure a unique monotonic equilibrium.[14] However, we plan to extend our approach to a more general type space in future work. Furthermore, while our identification strategy leverages the particular properties of scaling auctions, our work may provide methodological insights for estimation and prediction in more general multi-dimensional auctions as well.

Our paper also relates to a rich literature on the theory and estimation of equilibrium bidding in auctions with risk averse bidders. Maskin and Riley (1984) and Matthews (1987) first characterized the optimal auction in the presence of risk averse bidders with independent private values (IPV). While we do not relate our results to the optimal mechanism in this version of the paper, an evaluation of the DOT cost savings under the optimal mechanism, following Matthews (1987), is under preparation for a future draft. Campo, Guerre, Perrigne, and Vuong (2011) first established semi-parametric identification results for estimating risk aversion parameters in single-dimensional first price auctions in an IPV setting. As in their approach, we exploit the heterogeneity across items being auctioned and a parameterization of the bidders' utility function for identification. However, as our identification leverages the optimal spread of unit bids across items at each bidder's equilibrium score, we do not require any restrictions on the distribution of the bidders' private value distribution for estimation.

Our paper is structured as follows. In Section 2, we give an overview of how the players involved with procurement auctions—contractors and DOT managers—have treated bid skewing in practice. We then present an illustrative example of equilibrium bidding in our setting to demonstrate how uncertainty, risk aversion, and competition influence the inter-

[14] To see why a multi-dimensional bidder type model is substantially more complicated, note that a monotonic equilibrium in our setting requires a single-dimensional ranking of bidder types: a bidder with a better type should have a higher chance of winning (and therefore a lower score). Whereas in the risk neutral case (as in Bajari et al. (2014)), better bidders are those with lower expected costs for completing a project, risk averse bidders are compared by the certainty equivalent of their profits from completing a project. As demonstrated in section 2, the certainty equivalent entails an interaction between bidders' item costs and item bids, making straight-forward comparisons in a general case difficult.

pretation of bids that we may see in practice. In section 3, we discuss our dataset and present reduced form evidence that the bids we observe in our data support our model. In section 4, we present a full theoretical model of equilibrium bidding. In section 5, we present a structural model for estimating the auction primitives that underlie the bids in our data. In section 6, we present our structural estimates. Finally, in section 7, we present our counterfactual predictions and discuss their implications for policy.

2 Bid Skewing and Material Loss to the DOT

2.1 Scaling Auctions in Highway and Bridge Procurement

Like most other states, Massachusetts manages the construction and maintenance for its highways and bridges through its Department of Transportation (DOT). In order to develop a new project, DOT engineers assemble a detailed specification of what the project will entail. This specification includes an itemized list of every task and material (item) that is necessary to complete the project, as well as the engineers' estimates of the quantity with which each item will be needed, and a market unit rate for its cost. The itemized list of quantities is then advertised to prospective bidders.[15]

Any contractor who has been pre-qualified for a given project can submit a bid for the contract to implement it. Pre-qualification entails that the contractor is able to complete the work required, given their staff and equipment. Notably, it does not depend on past performance in any way. In order to submit a bid, a contractor posts a per-unit price for each of the items specified by the DOT. Since April 2011, all bids have been processed through an online platform, Bid Express, which is also used by 36 other state DOTs.[16] All bids are private until the completion of the auction.

Once the 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 the process of construction, it is common for items to be used in quantities that deviate from the DOT engineer's specification. All changes, however, must be approved by an on-site DOT manager. The winning contractor is ultimately paid the sum of her unit price bid multiplied by the *actual* quantity of each item used.

While contractors' ability to influence the item quantities that are ultimately used is limited, bidders may be able to predict which items will over/under-run the DOT's estimates.

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

[16]Scaling auctions using paper-bids were used for over a decade prior to the introduction of Bid Express.

Consequently, DOT officials have expressed concerns that bidders may manipulate unit prices to take advantage of government inaccuracies and extract rents from the taxpayer till.

2.2 Views of Bid Skewing by Contractors and DOT 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".

DOT 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 the DOT'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*.[17]

> *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.*[18]

However, as the definition for material unbalancedness is very broad, FHWA statute requires that a bid be *mathematically* unbalanced as a precondition. A *mathematically unbalanced* bid is defined as one, "structured on the basis of nominal prices for some work and inflated prices for other work."[19] 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 DOT officials to examine the possibility of material unbalancedness. However, in practice, such bids are ubiquitous, and substantial challenges by the DOT 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.[20] However, only 2.5% of projects have seen bidders rejected across all justifications, a

[17] 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)).

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

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

[20] See figures 20a and 20b in the appendix for more details.

handful of which were due to unbalanced bids.[21]

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).[22] 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, there is a theoretical basis to question the relationship between mathematical and material unbalancedness. As we demonstrate, bid skewing plays dual roles in bidders' strategic behavior. On the one hand, bidders extract higher ex-post profits by placing higher bids on items that they predict will overrun 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 overruns are largely competed away in equilibrium, but the reduction in ex-ante risk is passed on to the DOT in the form of cost-savings.

[21] Note that 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.

[22] 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.

2.3 An 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 $\sigma_r^2 = 1$.[23] We assume that the actual quantities are exogenous to the bidding process, and do not depend on who wins the auction in any way. Furthermore, we will assume that the bidders' expectations are identical across both bidders.[24]

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

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

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

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

Bidders are endowed with a standard CARA utility function over their earnings from the

[23] As we discuss in section 4, we assume that the distributions of q_c^a and q_r^a are independent conditional on available information regarding the auction. This assumption, as well as the assumption that the quantity distributions are not truncated at 0 (so that quantities cannot be negative) are made for the purpose of computational traceability in our structural model. Note that if item quantities are correlated, bidders' risk exposure is higher, and so our results can be seen as a conservative estimate of this case.

[24] 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.

project with a common constant coefficient of absolute risk aversion γ:

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

Note that 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 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. Note that 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:[25]

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

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

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

Furthermore, as we discuss in section 4, the optimal selection of bids for each bidder i can be described as the solution to a two-stage problem:

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

[25] See section 4 and the appendix for a detailed derivation.

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

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

To see how this decision process can generate bids that appear mathematically unbalanced, suppose, for example, that the common CARA coefficient is $\gamma = 0.05$, and consider a bidder in this auction who has type $\alpha^i = 1.5$.[26] Suppose, furthermore, that the bidder has decided to submit a total score of \$500. There are a number of ways in which the bidder could construct a score of \$500. For instance, she could bid her cost on concrete, $b_c^i = \$12$, and a dollar mark-up on traffic cones: $b_r^i = (\$500 - \$12 \times 10)/20 = \$19$. Alternatively, she could bid her cost on traffic cones, $b_r^i = \$18$, and a two-dollar mark-up on traffic cones: $b_c^i = (\$500 - \$18 \times 20)/10 = \$14$. Both of these bids would result in the same score, and so give the bidder the same chances of winning the auction. However, they yield very different expected utilities to the bidder. Plugging each set of bids into equation (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

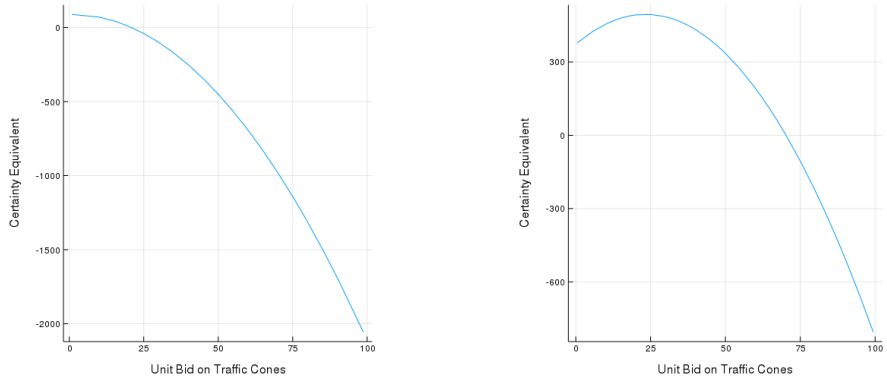
$$12 \times (\$2) + 16 \times (\$0) - \frac{0.05 \times 2}{2} \times (\$2)^2 - \frac{0.05 \times 1}{2} \times (\$0)^2 = \$23.80.$$

In fact, further inspection shows that the optimal bids giving a score of \$500 are $b_c^i = \$47.78$ and $b_r^i = \$1.12$, yielding a certainty equivalent of \$87.98. The intuition for this is precisely that described by Athey and Levin (2001), and the contractors cited by Stark (1974): the bidder predicts that concrete will overrun in quantity – she predicts that 12 tons will be used, whereas the DOT estimated only 10 – and that traffic cones will underrun – she predicts that 16 will be used, rather than the DOT's estimate of 20. When the variance terms aren't too large (relatively), the interpretation is quite simple: every additional dollar bid on concrete is worth approximately 12/10 in expectation, whereas every additional dollar bid on traffic cones is worth only 16/20.

However, the incentive to bid higher on items projected to overrun is dampened when the variance term is relatively large. This can arise when the coefficient of risk aversion is

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

relatively high or when the variance of an item's ex-post quantity distribution is high. More generally, as demonstrated in equation (1), the certainty equivalent of profits is increasing in the expected quantity of each item, $\mathbb{E}[q_c^a]$ and $\mathbb{E}[q_r^a]$, but decreasing in the variance of each item σ_c^2 and σ_r^2.



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

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

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

2.4 Bid Skewing in Equilibrium

As we discuss in section 4, the auction game described above has a unique Bayes Nash Equilibrium. This equilibrium is characterized following the two-stage procedure described on in section 4.2: (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' expecta-

tions 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 overruns 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 laid out on page 11 when bidders are risk averse with CARA coefficient $\gamma = 0.05$, and when bidders are risk neutral (e.g. $\gamma = 0$). To hone in on the effects of risk in particular, and not mis-estimation, we will assume that the bidders' expectations of ex-post quantities are perfectly correct (e.g. the realization of q_c^a is equal to $\mathbb{E}[q_c^a]$, although the bidders do not know this ex-ante, and still assume their estimates are noisy with Gaussian error).

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

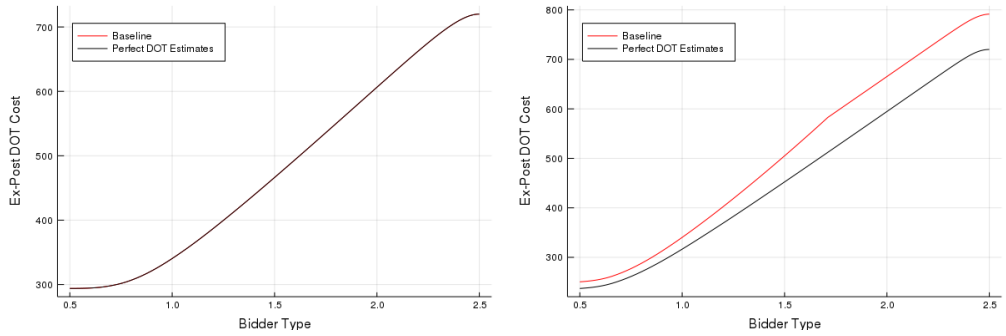
	Risk Neutral Bidders	Risk Averse Bidders
Noisy Quantity Estimates	\$326.76	\$317.32
Perfect Quantity Estimates	\$326.76	\$296.26

Table 1: Comparison of Expected DOT Costs

In Table 1, we present the expected (ex-post) DOT cost in each case. This is the expectation of the amount that the DOT would pay the winning bidder $q_c^a b_c^w + q_r^a b_r^w$ at the equilibrium bidding strategy in each setting, taken with respect to the distribution of the type of the lowest (winning) bidder.[27] When bidders are risk neutral ($\gamma = 0$), the equilib-

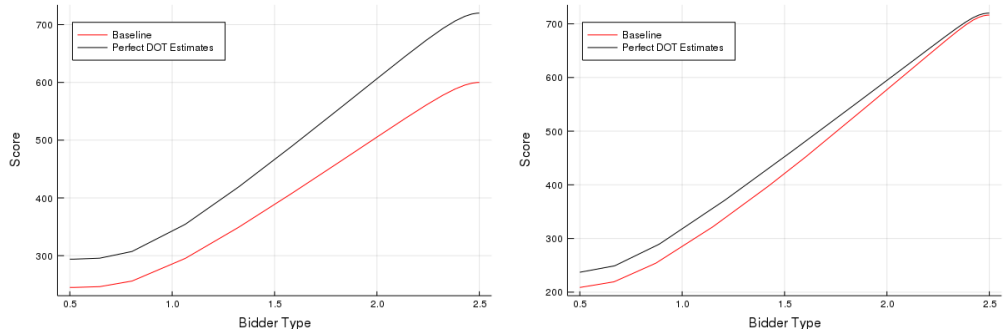
[27]In order to simulate equilibria, we need to assume a distribution of bidder types. For this ex-

rium cost to the DOT does not change when the DOT improves its quantity estimates. The reason for this is that since $\gamma = 0$, the variance term in equation (1) is zero regardless of the level of the noise in quantity predictions. As the bidders' quantity expectations $\mathbb{E}[q_c^a]$ and $\mathbb{E}[q_r^a]$ are unchanged, the expected revenue of the winning bidder (corresponding to the expected cost to the DOT) is unchanged as well.



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

Figure 2: Equilibrium DOT Cost/Bidder Revenue by Bidder Type



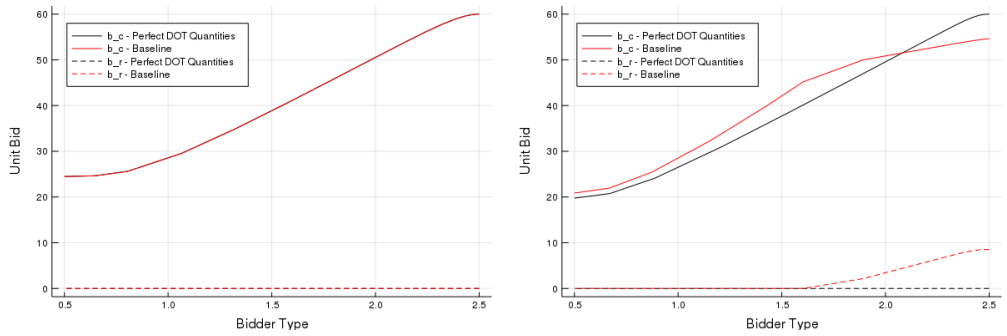
(a) Risk Neutral Bidders (b) Risk Averse Bidders: $\gamma = 0.05$

Figure 3: Equilibrium Score Functions by Bidder Type

In figure 2a, 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. Note that 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

ample, 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.

submit are different. In figure 3a, 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 overrunning 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 4: Equilibrium Unit Bids by Bidder Type

Figure 4a 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 overrun 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 2b, 3b and 4b, 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 overruns, and so there is no meaningful different to over-bid on one item over the other. The analysis of the optimal bid (conditional on a score) here is analogous to that under risk neutrality, and so we defer details to the appendix.

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

Table 2: Comparison of expected DOT costs under different levels of bidder risk aversion

While the general observation that reducing uncertainty may result in meaningful cost savings to the DOT, the degree of these savings depends on the baseline level of uncertainty in each project, as well as the degree of bidders' risk aversion and the level of competition in each auction (constituted by the distribution of cost types and the number of participating bidders). To illustrate this, we repeat the exercise summarized in Table 1 over different degrees of risk aversion and different levels of uncertainty. In Table 2, we present the expected DOT cost under the baseline example and under the counterfactual in which the DOT eliminates quantity risk, as well as the percent difference between the two, for a range of CARA coefficients.[28] The bolded row with a CARA coefficient of 0.05 corresponds to

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

the right hand column of Table 1. We repeat this exercise across different magnitudes of prediction noise in Table 21, in the appendix.

3 Data and Reduced Form Results

3.1 Data

Our data comes from MassDOT and covers 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. If we keep 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 maintenance. 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.[29]

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. And all contractors have had 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 6 MassDOT districts around the state. The Project Manager assigns a Resident Engineer to monitor work on a particular project out in the field and

[29] See table 19 for a breakdown of the number of projects in our data, by year.

to be the first to decide whether to approve or reject underruns, overruns, and Extra Work Orders (EWOs).[30] Underruns and overruns, as the DOT defines them and as we will define them here, apply to the items specified in the initial project design and refer to the difference between 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 underruns and overruns 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 2, the DOT's estimate, which can be thought of 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 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 5 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.

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

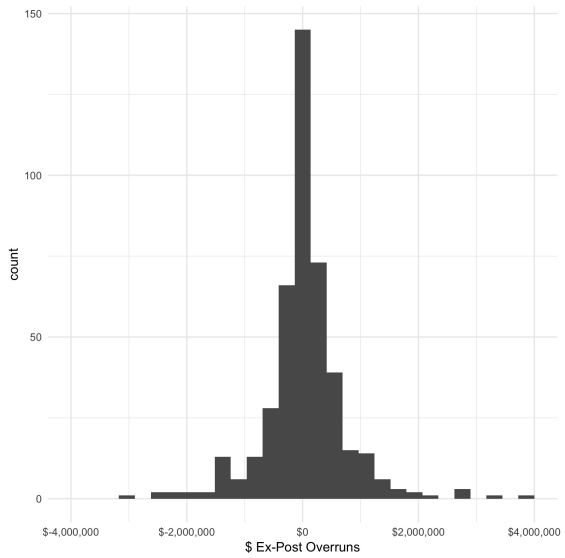


Figure 5: Net Over-Cost (Ex-Post Quantities) Across Bridge Projects

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 less 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.

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 overrun (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 overrun: 21.97% on average.

In addition to the firm's identity, there are a number of factors which may influence its competitiveness in a given auction. One such factor is the firm's distance from the project.

Common Firm	Common Firm	Rare Firm
Number of Firms	24	92
Total Number of Bid Submitted	2263	620
Mean Number of Bid Submitted Per Firm	94.29	6.74
Median Number of Bid Submitted Per Firm	63.0	2.5
Total Number of Wins	351	89
Mean Number of Wins Per Firm	14.62	0.97
Median Number of Wins Per Firm	10	0
Mean Bid Submitted	$2,774,941	$4,535,310
Mean Ex-Post Cost of Bid	$2,608,921	$4,159,949
Mean Ex-Post Overrun of Bid	9.7%	21.97%
Proportion of Bids on Projects in the Same District	28.19	15.95
Proportion 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

Although we do not observe precise locations for each project in our data, we observe which of the 6 geographic districts that MassDOT jurisdiction is broken into each project belongs to. We then geocode the headquarters of each firm by district, and compare districts for each project-bidder pair. 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 measure of competitiveness is specialization—firms with extensive experience bidding on and implementing a certain type of project may find it cheaper to implement an additional project of the same sort. Our data involves three distinct project types, according to the DOT taxonomy: Bridge Reconstruction/Rehabilitation projects, Bridge Replacement projects, and Structures Maintenance projects. 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. The mean specialization of a common firm is 24.44%, while the mean specialization of a rare firm is 2.51%. As projects have varying sizes, we compute a measure of specialization in terms of project revenue as well. We define a revenue-dominant firm (within a project-type) as a firm that has been awarded more than 1% of the total money spent by the DOT across projects of that project type. Among common firms, 51.67% of bids

submitted were by firms that were revenue dominant in the relevant project type; among rare firms, the proportion of bids by revenue dominant firms is 11.8%.

A third factor of competitiveness is each firm's capacity – the maximum number of DOT projects that the firm has ever had open while bidding on another project – and its utilization – the share of the firm's capacity that is filled when she is bidding on the current project.[31] 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.[32]

Description of Quantity Estimates and Uncertainty

As we discuss in section 2, 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 item 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).[33]

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.[34]

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

[32] Note that 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.

[33] A number of different factors may influence the extent of item over/under-runs in a given project: the type of maintenance needed, underlying state of the structure, time since assessment and skill of the project designer, chief among them. While our dataset is insufficient to robustly estimate the causal effects of these features on overruns, we present a brief discussion of the variation observed across DOT designers and project managers in the appendix.

[34] 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.

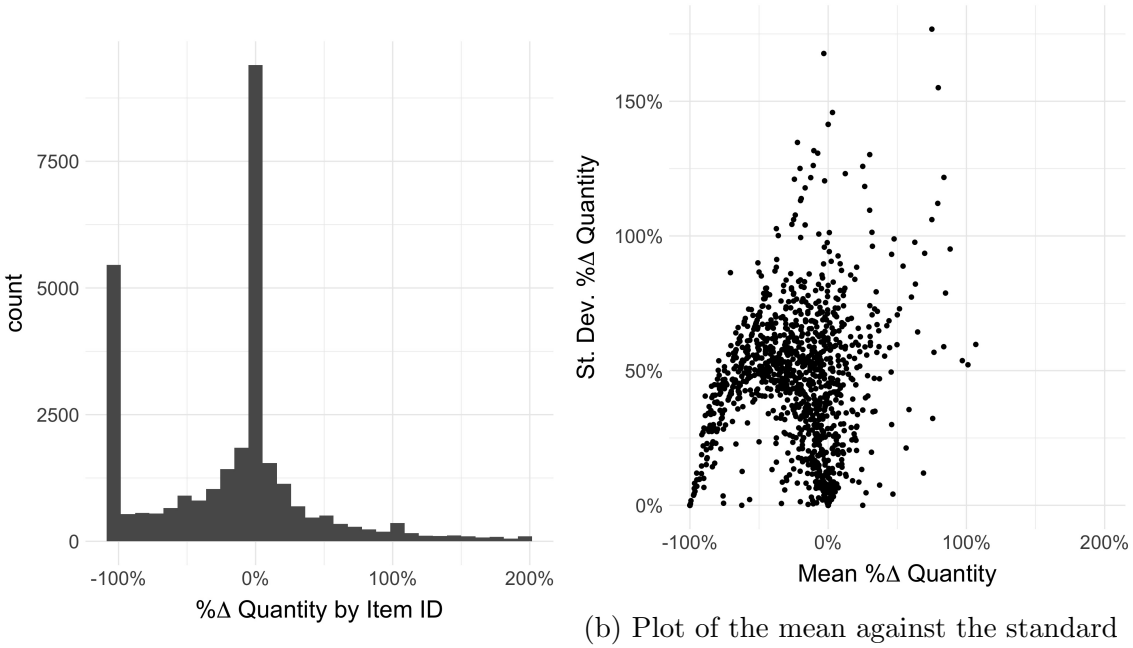
For each item, in every auction, we observe the quantity with which MassDOT predicted it would be used at the time of the auction – q_t^e in our model – the quantity with which the item was ultimately used – q_t^a – and a *blue book* DOT estimate for the market rate for the unit cost of the item. The DOT quantities are typically inaccurate: 76.7% of item observations in our data had ex-post quantities that deviated from the DOT estimates. Figure 6a presents a histogram of the percent quantity overrun across observations of items. The percent quantity overrun 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 overruns are 0%, another 18% involve items that aren't used at all (so that the overrun is equal to -100%). The remaining overruns are distributed, more or less symmetrically, around 0%. Furthermore, quantity overruns vary across observations of the same item in different auctions. Figure 6b plots the mean percent quantity overrun for each unique item with at least 2 observations against its standard deviation . While a few items have standard deviations close to 0, the majority of items have overrun standard deviations that are as large or larger than the absolute value of their means. That is, the percent overrun of the majority of unique items varies substantially across observations.[35] 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.

3.2 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 7, we plot the relationship between quantity overruns and the percent by which each item was overbid above the blue book cost estimate by the winning bidder.[36] 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 overrun 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, initially stated project length in days, number of bidders, and fixed effects for the year in which the project was opened for bidding, project

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

[36] 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 overrun 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 overrun across item-project pairs.

(b) Plot of the mean against the standard deviation of percent quantity overruns within each unique item.

Figure 6

type, resident engineer, project manager, and project designer. Specifications that include an array of additional controls, like item fixed effects, produce a very similar slope.[37] We use a similar procedure for all residualized bin-scatters in this section.

As figure 7 demonstrates, there is a significant positive relationship between percent quantity overruns and percent overbids by the winning bidder. A 1% increase in quantity overruns corresponds to a 0.085% increase in overbids on average.[38] This suggests that the winning bidder is able to correctly predict which items will overrun on average. As in the example in Section 2, items predicted to overrun generally recieve 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 skewing strategically.

Furthermore, the bid skewing relationship is similar across bidders beside the winner. Figure 8a plots the residualized bin-scatter of percent overbids against percent quantity overruns 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 overruns 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

[37]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 include untruncated versions in an online appendix for robustness.

[38]See the appendix for a full regression report.

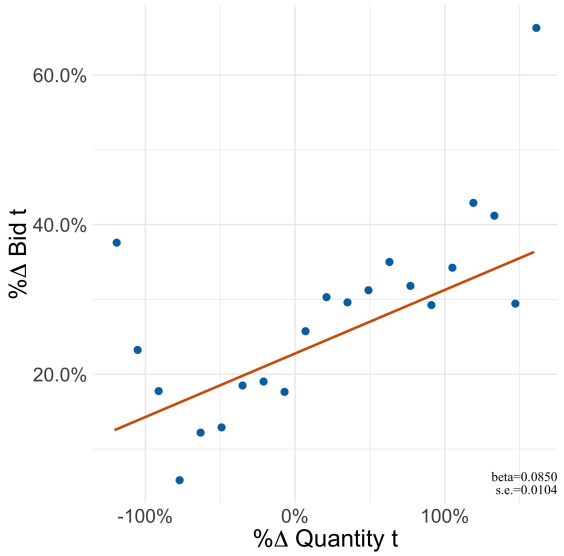


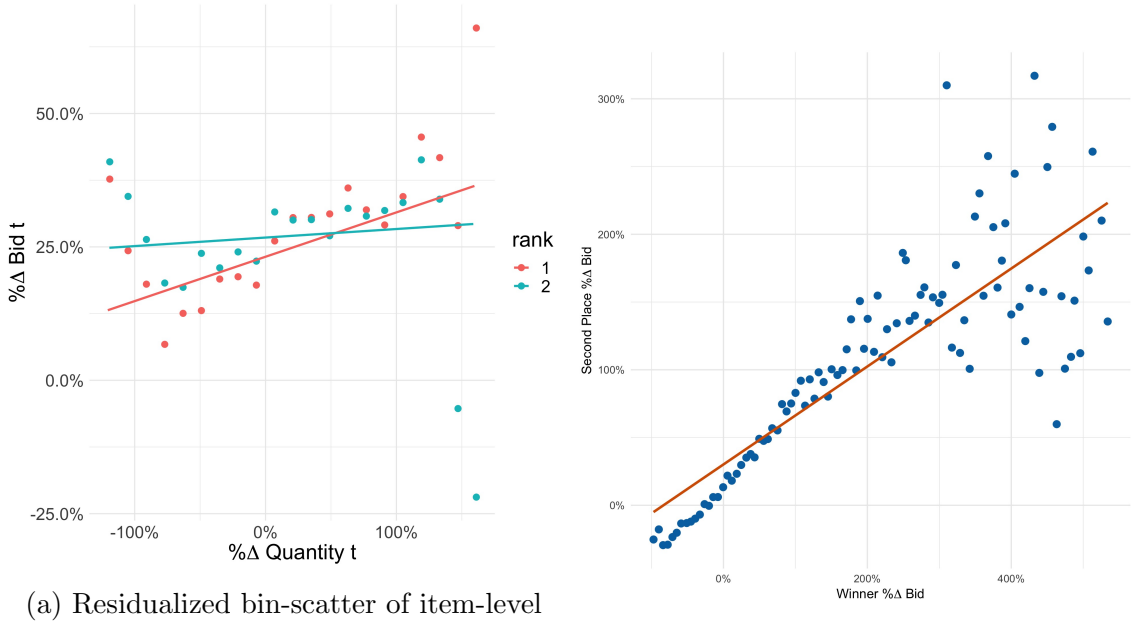
Figure 7: Residualized bin-scatter of item-level percent winner overbid against percent quantity overrun

similar total scores. Figure 8b plots a residualized bin-scatter 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.[39] Together, these figures suggest that bidders have access to the same information regarding quantity overruns.

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.

In addition to interior bids, risk aversion has several testable theoretical implications. First, risk averse bidders are predicted to bid more aggressively on projects that are worth

[39]Note that the percent overbids in figure 8b appear to be substantially larger than those in figure 8a. This is because while large overbids occur in the data, they are relatively rare and so are averaged down in the percent quantity overrun binning of figure 8a.



(a) Residualized bin-scatter of item-level percent overbid by the rank 1 (winning) and rank 2 bidder, against percent quantity overrun

(b) Bin-scatter of item-level percent overbid (not residualized for the moment but not for a meaningful reason)

Figure 8

more. A true reduced form test for this would require a ceteris perebis comparison of bid outcomes on identical auctions that only vary on project size. However, a suggestive proxy for aggressive bidding is the percent net over-cost: the percent by which the total amount paid to the winner in each project exceeds the total project value given by the DOT's blue book unit cost estimates. As shown in figure 9, this relationship is generally negative in our data. Interpreting percent net over-cost as a proxy for markups, this suggests that bidders extract less rents (percentage-wise) in auctions with higher stakes, as risk averse behavior would imply.

Furthermore, as we discuss in section 2, risk averse bidders balance the incentive to bid high on items that are projected to overrun 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 10a and 10b, 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

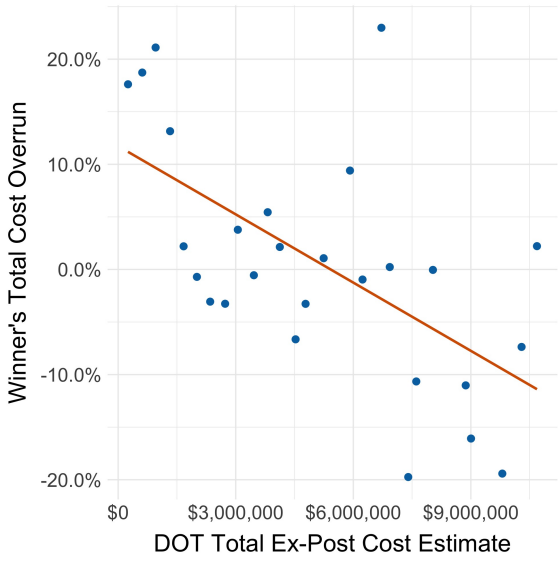
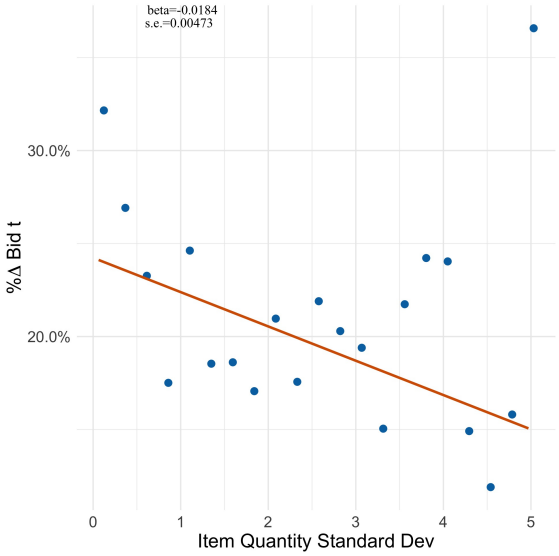


Figure 9: Bin-scatter of the percent net over-cost against the total DOT estimate for the project cost. DOT estimates are calculated with blue book cost estimates and ex-post quantity realizations.

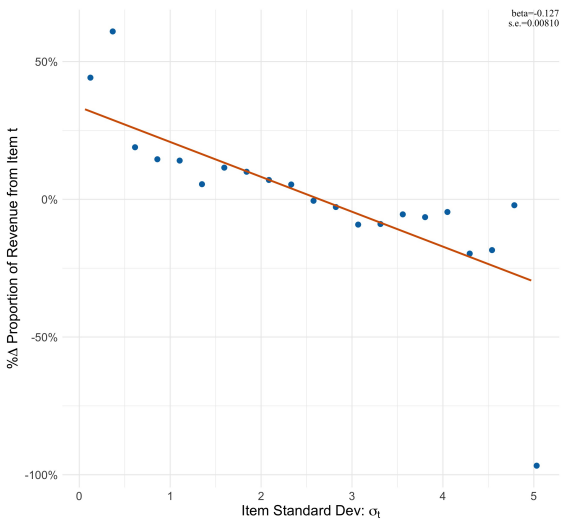
in section 5.[40] For every item, in every auction, our first stage gives us an estimate of the variance of the error on the best prediction of what the ex-post quantity of that item would be, given the information available at the time of bidding.

In figure 10a, we plot a residualized binscatter of the winning bidder's percent overbid on each item against the item's standard deviation – the square root of the estimated prediction variance. The relationship is negative on the whole, suggesting that bidders do bid lower on items with higher variance. Note, however, that this analysis does not consider 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

[40] As we discuss in section 5, 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 just Hamiltonian Monte Carlo 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 bin-scatter of item-level percent overbid against the square root of estimated item quantity variance



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

Figure 10

for bidding high on an item:

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

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. In figure 10b, we plot the residualized bin scatter of the %Δ Cost 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.

4 A Structural Model for Bidding With Risk Aversion

4.1 Setup

A procurement project is characterized by T items, each of which is needed in a different quantity. MassDOT (henceforth, "the buyer" or '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 buyer 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.[41]

The auction is simultaneous with sealed bids, but both the set of $m > 1$ participating bidders and the buyer'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.[42]

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 that market price of the materials—generally things like 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,

[41] 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 overruns 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 harm 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.

[42] 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.

publicly known distribution with a well behaved density $f(\alpha^i)$ over a compact subset $[\underline{\alpha}, \overline{\alpha}]$ of $\mathbb{R}_+$.[43] 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 $\mathbf{b}^i = \{b_1^i, \ldots, b_T^i\}$, 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:

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

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_t^i multiplied by the ex-post quantity q_t^a 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_t^a for each item.[44]

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:[45]

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

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

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

if she wins the auction. Note that 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.

[43] 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.

[44] Note that only the winner of the auction incurs any costs. All losing bidders receive no further cost nor revenue from the project at hand, once the auction is complete.

[45] Note that 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 independently 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

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

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

where 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. 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. Note that the expectation in the first term is with respect to $\mathbf{q}^a$.

Bidders form their expectations 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:

$$\begin{aligned}
&\left(1 - \mathbb{E}_{\varepsilon}\left[\exp\left(-\gamma \sum_{t=1}^{T} (q_t^b - \varepsilon_t) \cdot (b_t^i - \alpha^i c_t)\right)\right]\right) \cdot (\Pr\{s^i < s^j \text{ for all } j \neq i\}) \\
&= \left(1 - \exp\left(-\gamma \sum_{t=1}^{T} q_t^b (b_t^i - \alpha^i c_t) - \frac{\gamma \sigma_t^2}{2}(b_t^i - \alpha^i c_t)^2\right)\right) \cdot (\Pr\{s^i < s^j \text{ for all } j \neq i\}).
\end{aligned}$$

where the first equality is given by rewriting $q_t^a = q_t^b - \varepsilon_t$, so that the expectation operator in the profit term is with respect to the distribution of ε. The second equality follows from the closed form solution to this expectation.[46]

4.2 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 buyer 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.[47] 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

[46] $\mathbb{E}[\exp(-\gamma c\varepsilon)] = \exp(-\gamma \mu_\varepsilon c + \frac{\gamma^2 \sigma_\varepsilon^2}{2} c^2))$ when c is a constant and $\varepsilon \sim \mathcal{N}(\mu_\varepsilon, \sigma_\varepsilon^2)$.

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

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, submitting a score s – that is, the bidder's expected utility from winning the auction using the optimal vector of bids $\mathbf{b}$ that yield s – is determined entirely by her choice of s, and is monotonically increasing in s. Following a 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:

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

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

Note that 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:[48]

$$b^*_{i,t}(s) = \alpha^i c_t + \frac{q^b_t}{\gamma\sigma^2_t} + \frac{q^e_t}{\sigma^2_t \sum_{p=1}^{T}\left[\frac{(q^e_p)^2}{\sigma^2_p}\right]}\left(s - \sum_{p=1}^{T}\left[\alpha^i c_p q^e_p + \frac{q^b_p q^e_p}{\gamma\sigma^2_p}\right]\right). \tag{7}$$

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 expected utility given the equilibrium distribution of opponent scores.

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

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

where 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 is standard in auction theoretic analysis (see Milgrom and Segal (2002), for example), the optimal strategy is described by the first order condition:

$$\gamma\sum_{t=1}^{T}\left(q^b_t - \gamma\sigma^2(b^*_{i,t}(s^*_i) - \alpha^i c_t)\right)\frac{\partial b^*_{i,t}(s^*_i)}{\partial s} = \sum_{k\neq i}\frac{h^j(s^*)}{1 - H^j(s^*)}\left[\exp\left(\gamma\sum_{t=1}^{T} q^b_t(b^*_{i,t}(s^*_i) - \alpha^i c_t) - \frac{\gamma\sigma^2_t}{2}(b^*_{i,t}(s^*_i) - \alpha^i c_t)^2\right) - 1\right], \tag{8}$$

where $h^j(\cdot)$ is the pdf of contractor j's score distribution. Note that the exponent on the RHS is one over the certainty equivalent of the profit from winning - we will denote this as $\exp(\gamma\bar{\pi})$ as shorthand for expositional purposes. Substituting $\frac{\partial b^*_{i,t}(s^*_i)}{\partial s}$ by taking the derivative

[48] Note that 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 6, the bids predicted by our estimated model do quite well at matching the data.

of equation (7) with respect to s and evaluating it at s_i^*, we obtain a global optimality first order condition:

$$\frac{\gamma^2}{\sum_{p=1}^{T}\left[\frac{(q_p^e)^2}{\sigma_p^2}\right]}\left(\sum_{p=1}^{T}\left[\alpha^i c_p q_p^e + \frac{q_p^b q_p^e}{\gamma\sigma_p^2}\right] - s_i^*\right) = \sum_{k\neq i}\frac{h^j(s_i^*)}{1-H^j(s_i^*)}\left[\exp(\gamma\bar{\pi})-1\right]. \tag{9}$$

Note, however, that while equation 9 characterizes the equilibrium score s_i^* for bidder i, the equilibrium vector of bids *conditional* on s_i^* is defined entirely by the optimality of the bids with respect to bidder i's expected utility from winning the auction using s_i^*. That is, conditional on an equilibrium choice of score, the optimal bids for bidder i are given by equation 7, evaluated at the equilibrium score.

5 Econometric Model

We now present a multi-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 for each project type, and each bidder's efficiency type α in 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, as described in equations (7) and (9).

Stage 1a: Estimating the Posterior Distribution of q_t^a

In the model presented in section 4, 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.

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 the signal mean q_t^b and variance σ_t^2 for each item t.

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:

$$q_{t,n}^a = \widehat{q_{t,n}^b} + \eta_{t,n} \text{ where } \eta_{t,n} \sim \mathcal{N}(0, \hat{\sigma}_{t,n}^2) \tag{10}$$

$$\text{such that} \qquad \widehat{q_{t,n}^b} = \beta_{0,q} q_{t,n}^e + \vec{\beta}_q X_{t,n} \text{ and } \hat{\sigma}_{t,n} = \exp(\beta_{0,\sigma} q_{t,n}^e + \vec{\beta}_\sigma X_{t,n}). \tag{11}$$

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, $q_{t,n}^e$, and a matrix of item-project characteristics $X_{t,n}$. We estimate this model with Hamiltonian Monte Carlo as an efficient implementation of a likelihood method optimized for a GLM and use the posterior mode as a point estimate for the second stage of estimation.[49] We demonstrate the goodness of fit in section 6.

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 and 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 data set contains a unit bid for every item, submitted by every participating bidder in every auction that

[49] Note that it is possible to estimate our first and second stage jointly using Hamiltonian Monte Carlo, adding further fidelity to the effect of the first stage estimates on the second stage moments along the entire posterior distribution. However, as we prefer GMM for the second stage for this version, we make do with the posterior mode. We could also simply run the second stage GMM along the posterior distribution and compute a full second stage posterior this way, but this would be very computationally burdensome, and so we do not do so at this time.

we see. In particular, we have three main sources of heterogeneity: (1) bids submitted by different bidders in an auction with the same project characteristics, item, etc.; (2) bids submitted by the same bidders across different items and different auctions with 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 4 predicts that the optimal unit bid for item t for a bidder of type α_n^i in auction n is given by:

$$b^*_{t,i,n}(s^*_{i,n}) = \alpha^i_n c_{t,n} + \frac{q^b_{t,n}}{\gamma\sigma^2_{t,n}} + \frac{q^e_{t,n}}{\sigma^2_{t,n}\sum_{p=1}^{T_n}\left[\frac{(q^e_{p,n})^2}{\sigma^2_{p,n}}\right]}\left(s^*_{i,n} - \sum_{p=1}^{T_n}\left[\alpha^i c_{p,n}q^e_{p,n} + \frac{q^b_{p,n}q^e_{p,n}}{\gamma\sigma^2_{p,n}}\right]\right), \quad (12)$$

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

$$\frac{\gamma^2}{\sum_{p=1}^{T_n}\left[\frac{(q^e_{p,n})^2}{\sigma^2_{p,n}}\right]}\left(\sum_{p=1}^{T_n}\left[\alpha^i_n c_{p,n}q^e_{p,n} + \frac{q^b_{p,n}q^e_{p,n}}{\gamma\sigma^2_{p,n}}\right] - s^*_{i,n}\right) = \sum_{k\neq i}\frac{h_n(s^*_{i,n})}{1-H_n(s^*_{i,n})}\left[\exp(\gamma\bar{\pi}) - 1\right], \quad (13)$$

where $\bar{\pi} = \sum_{t=1}^{T_n} q^b_{t,n}(b^*_{t,i,n}(s^*_{i,n}) - \alpha^i_n c_{t,n}) - \frac{\gamma\sigma^2_{t,n}}{2}(b^*_{t,i,n}(s^*_{i,n}) - \alpha^i c_{t,n})^2$.

As discussed above, we identify $q^b_{t,n}$ and $\sigma^2_{t,n}$ with a statistical model of ex-post quantities conditional on item-project characteristics using the full history of auctions in our data. To reduce the dimensionality of our parameter space, we model the bidder-auction efficiency type α^i_n onto a bidder-specific fixed effect and a regression model of bidder-auction characteristics:

$$\alpha^i_n = \alpha^i + \beta_\alpha X_{i,n}.$$

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 that appears 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 see the need for Assumption 1, note that 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. We will therefore examine relaxations of the independence assumption in an upcoming revision.

Note that Assumption 1 implies that the optimal score $s^*_{i,n}$ is also observed with error:

$$s^*_{i,n} = \sum_{t}^{T_n} b^d_{i,t,n} q^e_{t,n} + \bar{\nu}_{i,n} = s^d_{i,n} + \bar{\nu}_{i,n},$$

where $\bar{\nu}_{i,n} = -\sum_{t=1}^{T_n} \nu_{t,i,n} q^e_{t,n}$ is also mean-zero, conditional on the project characteristics of auction n. Write $\theta_2 = (\gamma, \{\alpha^i\}, \vec{\beta}_\alpha | \beta_{0,q}, \vec{\beta}_q, \beta_{0,\sigma}, \vec{\beta}_\sigma, \vec{\beta}_s, \vec{\sigma}_s)$. By definition, the bidder-item-auction level error on each unit bid is given by:

$$\begin{aligned}
\nu_{t,i,n} = b^d_{t,i,n} - \boldsymbol{\alpha^i_n} & \left(c_{t,n} - \frac{q^e_{t,n}}{\sigma^2_{t,n} \sum_{p \in \mathcal{T}(n)} \left[\frac{(q^e_{p,n})^2}{\sigma^2_{p,n}} \right]} \left[\sum_{p \in \mathcal{T}(n)} c_{p,n} q^e_{p,n} \right] \right) - \\
& \frac{1}{\boldsymbol{\gamma}} \left(\frac{q^b_{t,n}}{\sigma^2_{t,n}} - \frac{q^e_{t,n}}{\sigma^2_{t,n} \sum_{p \in \mathcal{T}(n)} \left[\frac{(q^e_{p,n})^2}{\sigma^2_{p,n}} \right]} \left[\sum_{p \in \mathcal{T}(n)} \frac{q^b_{p,n}}{\sigma^2_{p,n}} q^e_{p,n} \right] \right) - \\
& \frac{q^e_{t,n}}{\sigma^2_{t,n} \sum_{p \in \mathcal{T}(n)} \left[\frac{(q^e_{p,n})^2}{\sigma^2_{p,n}} \right]} \left[s^d_{i,n} + \boldsymbol{\bar{\nu}_{i,n}} \right]. \quad (14)
\end{aligned}$$

where

$$\alpha^i_n = \alpha^i + \beta_\alpha X_{i,n}. \quad (15)$$

Note that $\nu_{i,t,n}$ is linear in α_n^i and $\frac{1}{\gamma}$, as well as in $\bar{\nu}_{i,n}$. Furthermore, under Assumption 1, since $\mathbb{E}[\nu_{t,i,n}] = 0$ and is orthogonal to the matrix of item-auction features $X_{t,n}$ and bidder-auction features $X_{i,n}$, we have that $\mathbb{E}[\bar{\nu}_{i,n}] = 0$ as well.

We therefore define a demeaned bid error

$$\tilde{\nu}_{t,i,n} = \nu_{i,t,n} - \frac{q_{t,n}^e}{\sigma_{t,n}^2 \sum_{p \in \mathcal{T}(n)} \left[\frac{(q_{p,n}^e)^2}{\sigma_{p,n}^2}\right]} \bar{\nu}_{i,n}, \tag{16}$$

and form the following moment conditions, under Assumption 1:

$$\mathbb{E}\left[\tilde{\nu}_{t,i,n} \cdot Z_{t,i,n} | X_{t,n}, X_{i,n}\right] = 0,$$

where Z is each of the following instruments:

- Indicator for unique firm IDs[50]
- Indicator for being a "top skewed item"
- The bidder-auction feature vectors that comprise $X_{i,n}$.

Identification

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 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 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 $X_{i,n}^j$, and asymptotically zero in the number of

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

auctions. There are 14 such moments, one for each column of the feature matrix $X_{i,n}$. Each of these moments can be thought of as informing the identification of the coefficient β_α^j.

For complete details on the moment construction, see section C.2. Note that our moment conditions use only the optimality of bidders' unit price bids (given the scores that are observed in equilibrium).[51]

5.1 Bayesian Sampling with Hamiltonian Monte Carlo

In addition to our main GMM approach, we estimate a (fully) parametric version of our structural model using Hamiltonian Monte Carlo.[52] Bayesian analysis facilitates the modeling of hierarchical relationships in bidders' efficiency types – across auctions for the same bidders, and across bidders in similar auctions. In our GMM approach, we account for these relationships in the form of bidder fixed effects, and a regression function of auction-bidder characteristics. However, a more sophisticated GMM treatment would be difficult, given the high dimensionality of the parameter space and the amount of data available. As such, we consider both approaches in our paper. We present the details of a preliminary Bayesian specification for the second stage of our structural estimation along with results from an HMC fit of the model in section E of the appendix.

6 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 $\widehat{q}_{t,n}^b$ as well as the variance term $\widehat{\sigma}_{t,n}^2$ in table 13 in the appendix. A histogram of the resulting variance terms themselves are plotted in figure 11, 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.[53] In addition, we demonstrate the model fit of our first stage in figure 16 and table 12 in the appendix.

In the second stage, we estimate a common CARA coefficient γ, as well as a bidder-auction specific efficiency type $\alpha_n^i = \alpha^i + \beta_\alpha X_{i,n}$ for every bidder-auction pair in our data

[51] This is conceptually similar to the classical GPV use of the empirical bid distribution to model bidders' beliefs over the distribution of opponent bids that they face.

[52] Hamiltonian Monte Carlo is an efficient algorithm for sampling the posterior distribution of a statistical model. See Betancourt (2017) for an accessible complete explanation.

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

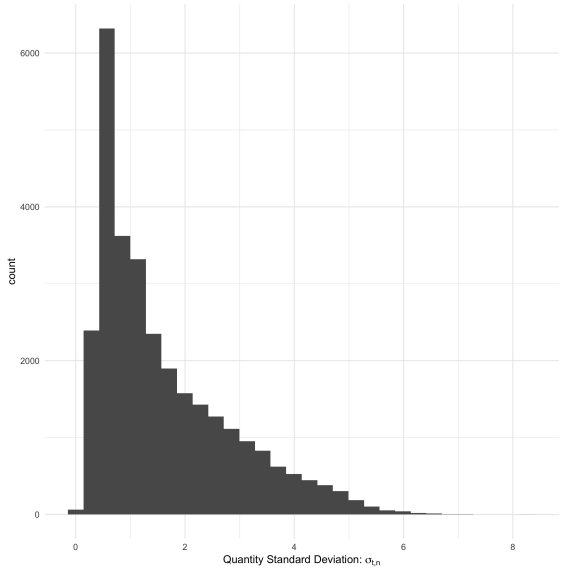


Figure 11: A histogram of standard deviation estimates for each item t in each project n

using the GMM estimator presented in section 5. We summarize the results in tables 5, 6 and 7. The full parameter estimates are presented in table 14 in the appendix. The coefficient of risk aversion γ in our data is estimated to be about 0.046. This corresponds to a certainty equivalent of \$15,000 for a 50-50 lottery for \$100,000 (or \$0 upon a loss).[54] As we report in table 5, the 95% confidence interval around our estimate is $(0.032, 0.264)$. 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.[55]

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

Table 5

In table 6, we present summary statistics of our estimates of bidder-auction efficiency types.[56] We break down the results by project type to highlight the differences between

[54]Note that 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.

[55]At the moment, the bootstrap is only over the second stage, holding the first stage estimates fixed. A full two-stage bootstrap, which requires substantially more computation time, will be presented in a future draft.

[56]There are a few of decisions made by the econometrician in estimation. We considered different thresholds on the number of auctions in which a firm must have participated in order to have a separate firm fixed effect. We also identified several outlying items: items that constituted large fractions of the project cost

		$\widehat{\alpha}_n^i$			
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.

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.[57] 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 – her unit 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 over the bidder's private costs only.

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

and were always estimated and used in unit quantities. These items might better be represented as lump sum items, over which uncertainty is poorly captured in our quantity model. The substance of our results is robust to these considerations, however. We present the results under different thresholds and when large lump items are excluded in the appendix as a robustness check.

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

		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 $\hat{\alpha}_n^i$

projects as well. This may be due, in part, to inaccurate prediction of the ex-post quantities. However, note that 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 18 and 19, and table 15 in the appendix. Figure 18 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 19 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 15 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.

7 Counterfactual

7.1 Perfectly Predicted DOT Quantities

In order to draw conclusions from our results, we return to the discussion in section 2. How much money *would* the DOT save if it were able to perfectly predict the actual quantities that will be 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.[58]

[58]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. As 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

Note that it is not sufficient to simply invert the econometric model of bidding described in section 5 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 2, 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. 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 coefficient of risk aversion, we use the estimate $\hat{\gamma} = 0.046$ from the second 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 $\widehat{\alpha}_n^i$ from our second stage onto auction characteristics.[59] The details of the equilibrium construction are presented in section A of the appendix.

In Figure 12, we plot a histogram of the (a) percentage and (b) dollar savings to the DOT from the perfect quantity prediction counterfactual. To calculate these savings, we compute the equilibrium bids for every efficiency type α twice: first under the baseline setting, and second 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, $\boldsymbol{\sigma}^2 \approx 0$.[60] 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 $\boldsymbol{q}^a$.[61] The dollar gains in figure 12b are computed by taking the difference between the expected DOT cost under the baseline

bidders would eventually realize that the DOT's quantities are accurate.

[59] 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.

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

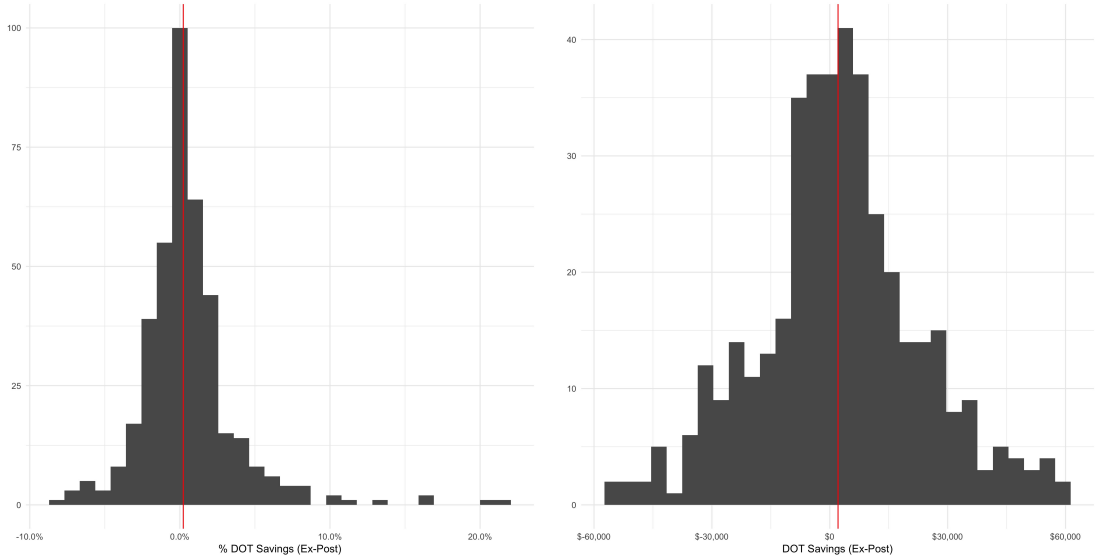
[61] More concretely, let $g(\alpha)$ and $G(\alpha)$by the density and cumulative probability functions of bidders' efficiency types in a given auction. Let $g^1(\alpha) = Ng(\alpha)(1 - G(\alpha))^{N-1}$ be the density of the first order statistic of g—the density of the lowest type bidder, when there 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 $\int_{\underline{\alpha}}^{\overline{\alpha}} g^1(\tilde{\alpha}) \sum_t q_t^a b_t^*(\tilde{\alpha}) d\tilde{\alpha}$.

setting, and under the counterfactual setting for each auction. The percent gains in figure 12a 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 12c. We calculate bidder utility gains by taking 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.[62] We present summary statistics for all three metrics in table 8.

We predict that the median expected saving to the DOT from eliminating uncertainty about ex-post quantities is about $2,203 or 0.23% of the baseline expected project cost. However, the standard deviation of savings is about $24,704 (4.25%) and the 25th and 75th percentiles are -$9,355 (-1.02%) and $13,987 (1.60%) 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 $\boldsymbol{q}^b$, which often differ from the true quantities. That is, whereas in the baseline, bidders optimize unit bids with regards 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 $17.61, and the 25th and 75th percentiles are $3.76 and $43.35, 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.

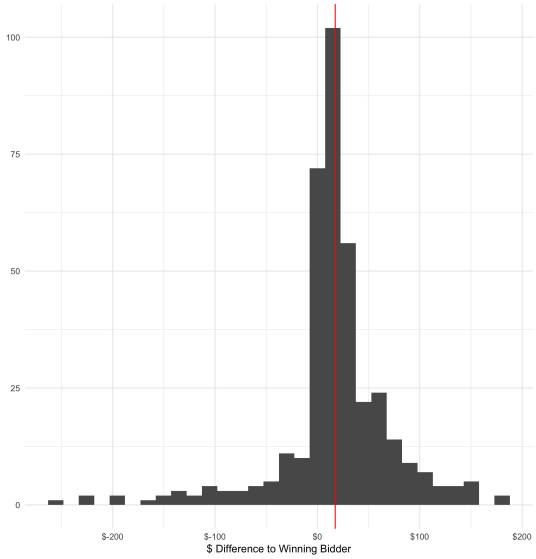
The projected expected DOT savings from eliminating risks detailed in table 8 and figure 12 reflect the two channels by which eliminating uncertainty changes the bidders' problem: (1) it eliminates risk, raising the value of the project and encouraging more aggressive bids; (2) it gives bidders access to the accurate ex-post quantities, allowing bidders to perfectly

[62]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.



(a) DOT % Savings

(b) DOT Dollar Savings



(c) Winning Bidder Utility Gains

Figure 12: Percent and dollar expected DOT ex-post Savings, and bidder utility gains from a counterfactual in which risk is eliminated. The median is highlighted in red in each case.

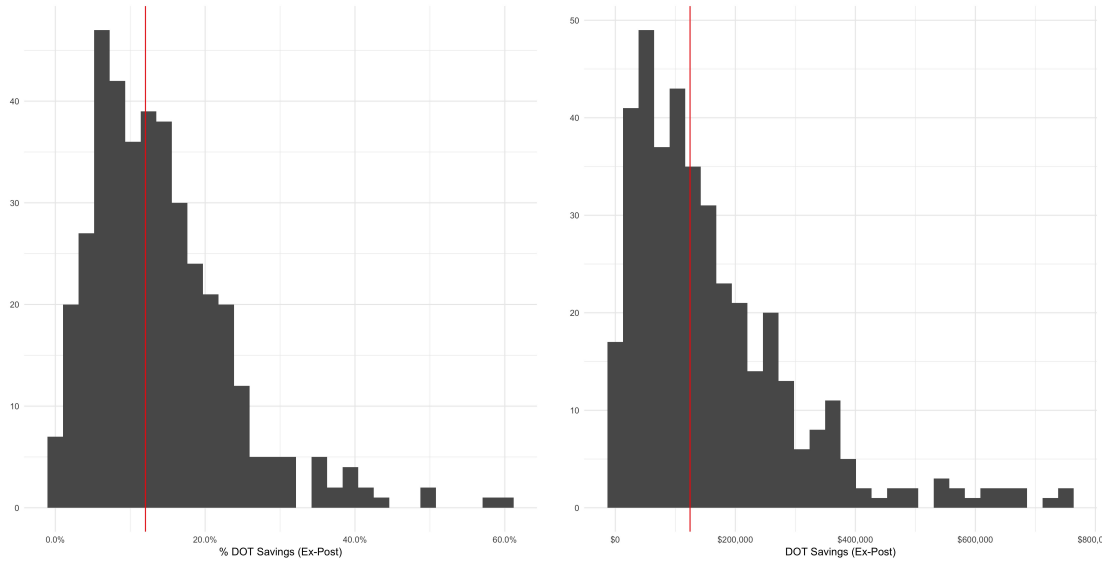
Statistic	Mean	St. Dev.	25%	Median	75%
Net DOT Savings	\$2,145.37	\$24,704.09	− \$9,354.61	\$2,203.49	\$13,987.89
% DOT Savings	0.70%	4.25%	−1.02%	0.23%	1.60%
Bidder Gains	\$6.64	\$145.87	\$3.76	\$17.61	\$43.35

Table 8: 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. Note: Results are truncated at the top and bottom 1% to exclude extreme outliers from the mean/SD calculations.

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 table 9 and figure 13. Absent bidder mis-optimization due to inaccuracies in their quantity projections, the median expected saving to the DOT is \$125,187 or 11.98% 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 \$4.81 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).

Statistic	Mean	St. Dev.	25%	Median	75%
Net DOT Savings	\$172,513.80	\$165,129.50	\$61,569.34	\$125,187.10	\$226,318.90
% DOT Savings	13.74%	9.05%	7.18%	11.98%	18.25%
Bidder Gains	\$19.16	\$124.55	−\$8.48	\$4.81	\$37.64

Table 9: 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. Note: Results are truncated at the top and bottom 1% to exclude extreme outliers from the mean/SD calculations.



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

Figure 13: Percent and dollar expected DOT ex-post Savings, and bidder utility gains 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 in red in each case.

7.2 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 departments of American DOTs, 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.[63] Second, they incentivize the winning bidder to minimize costs (as all costs are privately incurred and

[63] Around 2007, the MBTA – the segment of MassDOT responsible for construction and maintenance of the public transportation system in Massachusetts – switched from scaling auctions to lump sum auctions for the majority of its procurement. We spoke to the chief engineer about the decision for this transition in 2017. Chief among his reasons was the assertion that the scope of MBTA projects is much more difficult to define (and therefore spec out ex-ante) than 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.

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 2, 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 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

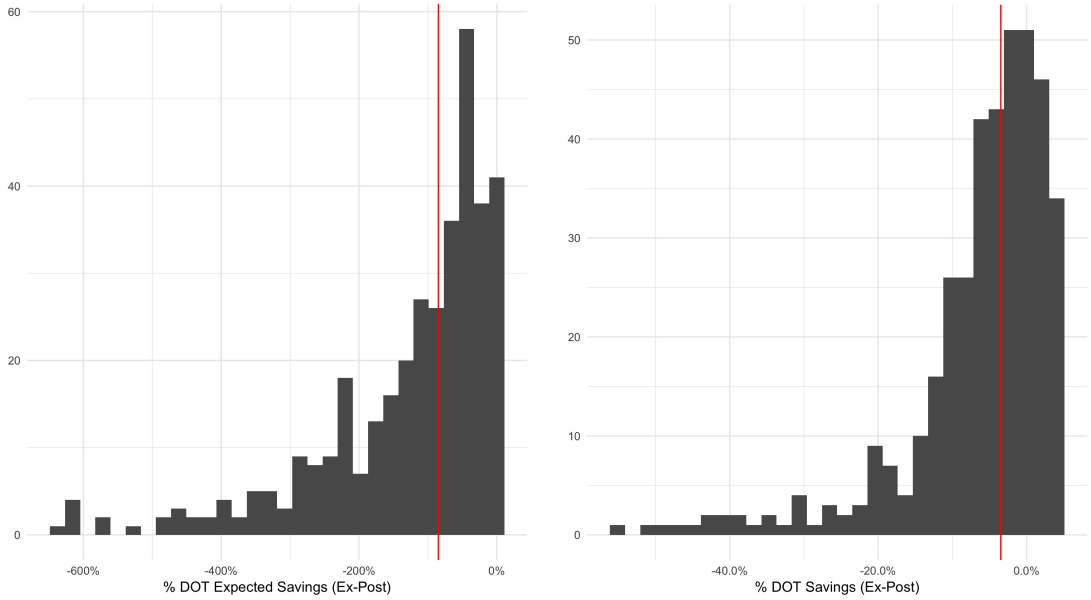
$$\sum_t (\mu q_t^a + (1-\mu) q_t^e) \cdot b_t,$$

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 is characterized as in section 4.2 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 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 14a, and under a $\frac{1}{2}$-sharing auction in figure 14b. The median expected loss from moving to a lump sum auction is 84.84%, while the median expected loss from a $\frac{1}{2}$-sharing auction is 3.47% – both with fat tails. Summary statistics for each case are presented in table 10.[64]



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

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

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

8 Entry

It is well known that an increase in competition benefits an auctioneer. In this section we evaluate the entry of an additional contractor to each auction in our data. First, we estimate the expected amount that the DOT would save if an additional contractor were to enter. We do this by computing the equilibrium bid function in each auction under the baseline (as

[64]Note that 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 overruns (and lower risk) at a competitive score (even if the score itself rises).

DOT % Cost Change	Mean	St. Dev.	25%	Median	75%
Lump Sum	−127.93%	129.70%	−175.11%	−84.84%	−38.08%
$\mu = 1/2$	−6.84%	12.00%	−9.49%	−3.47%	0.39%

Table 10: Summary of expected DOT percent cost change from switching to a μ-sharing auction with $\mu = 0$ (lump sum) and $\mu = 1/2$. Note: Results are truncated at the top and bottom 1% to exclude extreme outliers from the mean/SD calculations.

in the counterfactuals described in figure 12), 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, under the assumption that bidders enter if they anticipate to profit more than the cost of entry and total entry is set in equilibrium.

8.1 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.[65] 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

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

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.[66] We plot the distribution of lower bounds on the cost of entry K in each auction in figure 15a. In figure 15b, we plot the expected savings to the DOT from the entry of an additional bidder. Summary statistics are presented in table 11. The median lower bound on entry costs is \$1,959, while the median DOT savings amount to \$49,335. The distribution of DOT savings is quite fat tailed, however. While the mean lower bound on entry costs is \$2,316, the mean DOT saving is \$82,583. This suggests that there is substantial potential value to encouraging entry with a relatively modest guaranteed bonus payment to the winning bidder.

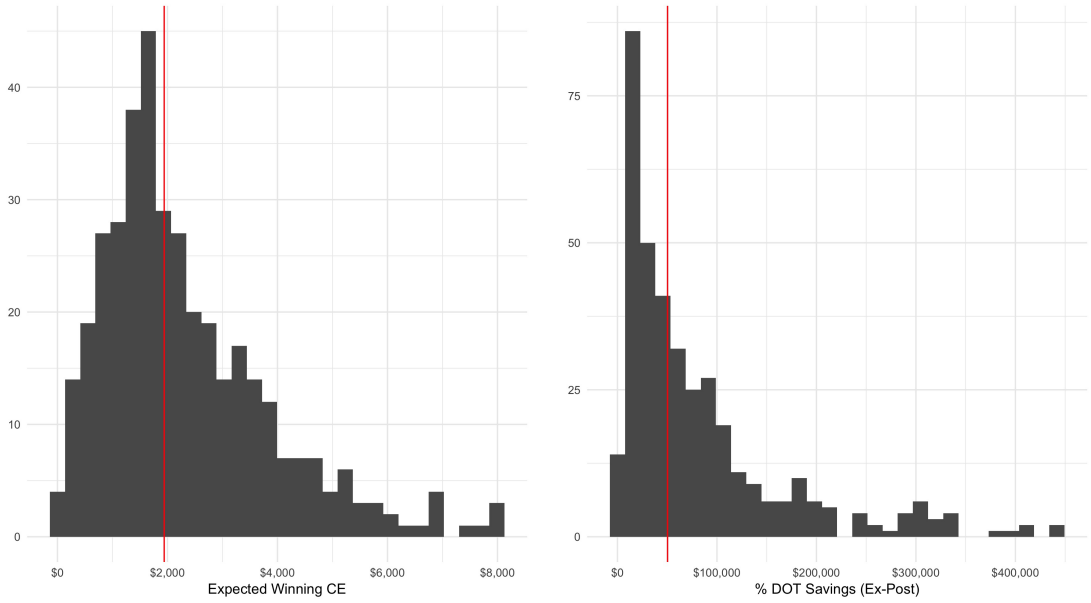
Statistic	Mean	St. Dev.	25%	Median	75%
Net DOT Savings	\$82,583.25	\$87,568.51	\$22,296.89	\$49,335.35	\$103,379.50
% DOT Savings	8.90%	8.45%	2.06%	5.65%	13.47%
Bidder Certainty Equivalent	\$2,315.80	\$1,524.88	\$1,264.95	\$1,959.42	\$3,135.44

Table 11: Summary of the welfare impacts of an additional bidder (not paying an entry cost) to each auction. Note: Results are truncated at the top and bottom 1% to exclude extreme outliers from the mean/SD calculations.

9 Conclusion

This paper studies the bidding behavior of construction firms that participate in scaling procurement auctions hosted by the Massachusetts Department of Transportation. In particular, we analyze the incentives for bidders to strategically *skew* their bids. We show that while

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



(a) Distribution of lower bounds on the cost of entry

(b) Distribution of the expected dollar savings to the DOT from the entry of an additional bidder to each auction

Figure 15: Welfare impacts of an additional entry to each auction

bidders do skew, placing high bids on items they predict will overrun the DOT's quantity estimates and low bids on items they predict will underrun, this is not necessarily indicative of rent extraction. For risk averse bidders, skewing facilitates diversifying bidders' exposure to the risk of items being used in quantities far outside their expectations. In a competitive environment, such as the one in MassDOT's bridge maintenance auctions, skewing generates substantial savings to the DOT. If bidders were compensated entirely based on the DOT's quantity estimates (or equivalently, using a lump sum auction), they would not be able to skew their bids. However, in this case, bidders would be responsible for all unanticipated modifications to the project specification and raise their bids on the whole to account for the added risk. Our estimates suggest that the DOT would subsequently pay nearly 85% more for the median project.

While switching to a lump sum auction would increase DOT expenditures, increasing bidders' exposure to risk a little bit may not be as harmful. A mixed compensation auction in which bidders are paid half on the DOT's estimates and half on the realized quantities only increases the median project's cost to the DOT by 3.5%. For a few projects, this mixed auction may even save the DOT a little. This suggests that policies that limit bidders' ability to fully optimize their bids—such as the minimum bid requirement considered by MassDOT—may be helpful in reducing DOT expenditures. In an upcoming draft of this paper, we examine a counterfactual with a minimum bid requirement to find out.

References

Alesina, A. and G. Tabellini (2007). Bureaucrats or politicians? Part I: A single policy task. *American Economic Review 97*(1), 169–179.

Alesina, A. and G. Tabellini (2008). Bureaucrats or politicians? Part II: Multiple policy tasks. *Journal of Public Economics 92*(3-4), 426–447.

Altshuler, A. and D. Luberoff (2004). *Mega-Projects: The Changing Politics of Urban Public Investment*. Brookings Institution Press.

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

Athey, S. and P. Haile (2006). Empirical models of auctions. pp. 1–48.

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

Bafumi, J. and A. Gelman (2006). Fitting Multilevel Models When Predictors and Group Effects Correlate. *Political Science*.

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

Bajari, P., R. McMillan, and S. Tadelis (2009). Auctions versus negotiations in procurement: An empirical analysis. *Journal of Law, Economics, and Organization 25*(2), 372–399.

Berman, R. (2015, Dec). A major infrastructure bill clears congress. *The Atlantic*.

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

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

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*.

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

Evdokimov, K. (2010). Identification and Estimation of a Nonparametric Panel Data Model with Unobserved Heterogeneity. *Princeton University: Discussion Paper*, 1–62.

Flyvbjerg, B. (2009). Survival of the unfittest: Why the worst infrastructure gets built-and what we can do about it. *Oxford Review of Economic Policy 25*(3), 344–367.

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.

Haile, P. A. and E. Tamer (2003). Inference with an Incomplete Model of English Auctions. *Journal of Political Economy 111*(1), 1–51.

Hinze, J. (2010). *Construction Contracts*. McGraw-Hill Education.

Kahneman D. Lovallo, D. (1993). Timid choices and bold forecasts: A cognitive perspective on risk taking.

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

Krishna, V. (2009). *Auction theory*. Academic press.

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

Lewis, G. and P. Bajari (2011). Procurement contracting with time incentives: Theory and evidence. *Quarterly Journal of Economics 126*(3), 1173–1211.

Lewis, G. and P. Bajari (2014). Moral hazard, incentive contracts, and risk: Evidence from procurement. *Review of Economic Studies 81*(3), 1201–1228.

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

Masscases.com (1984). Department of labor and industries vs. boston water and sewer commission, 18 mass. app. ct. 621.

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.

Nilson, E. (2018, Aug). Democrats still want to make infrastructure week happen. *Vox*.

Pitsis, T. S., M. Kornberger, and S. Clegg (2004). The art of managing relationships in interorganizational collaboration. *Management 7*(3), 47–67.

Porter, R. H. and J. D. Zona (1993). Detection of Bid Rigging in Procurement Auctions. *Journal of Political Economy 101*(5), 255–268.

Puentes, R. (2015, Jan). Why infrastructure matters: Rotten roads, bum economy. *Brookings*.

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.

Shemkus, S. (2015, aug). State's Aging Infrastucture Worries Businesses.

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

Somaini, P. (2015). Identification in auction models with interdependent costs.

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.

Team, S. D. (2018a). Rstan: the r interface to stan. *R package version 2*(17), 3.

Team, S. D. (2018b). Stan modeling language users guide and reference manual. *Version 2*, 18.

US Department of Transportation (1988). Memorandum: Bid analysis and unbalanced bids. Technical report, Washington DC.

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 $\boldsymbol{q}^b$ and 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:

$$\alpha_n^i \sim \text{LogNormal}(\mu_n^\alpha, \sigma_n^{\alpha 2}) \tag{17}$$

where $\mu_n^\alpha = X_n \beta_\alpha$ and σ_n^α is project-type specific. 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 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:

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

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

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

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

$$s'(\alpha)\sum_{t=1}^{T}\left[\left(\gamma q_t^b - \gamma^2\sigma^2(b_t^*(s(\alpha)) - \alpha c_t)\right)\frac{\partial b_t^*(s(\alpha))}{\partial s}\right] = [\exp(\gamma\bar{\pi}(\alpha)) - 1]\sum_{j\neq i}\frac{f(\alpha)}{1-F(\alpha)}, \quad (18)$$

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:

$$\max_{\mathbf{b}(s(\alpha))}\left[1 - \exp\left(-\gamma\sum_{t=1}^{T} q_t^b(b_t(s(\alpha)) - \alpha c_t) - \frac{\gamma\sigma_t^2}{2}(b_t(s(\alpha)) - \alpha c_t)^2\right)\right] \quad (19)$$

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

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

Note that 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 19 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:

$$b_t^*(\cdot) = \max\left\{\begin{array}{l} \alpha c_t + \frac{q_t^b}{\gamma\sigma_t^2} + \frac{q_t^e}{\sigma_t^2\sum_{t:b_t^*(\cdot)>0}\left[\frac{(q_t^e)^2}{\sigma_t^2}\right]}\left(s(\alpha) - \sum_{t:b_t^*(\cdot)>0}\left[\alpha c_t q_t^e + \frac{q_t^b}{\gamma\sigma_t^2}q_t^e\right]\right) \\ 0 \end{array}\right. \quad (20)$$

Note that when all items have interior bids, this is equivalent to equation 7. We solve the ODE in (18) numerically using a state-of-the-art stiff ODE solver using the DifferentialEquations library in Julia.[67] At every evaluation of equation (20) in the ODE solver, we compute the optimal bid vector at every score by numerically solving the program in (19) using the IPOPT optimization suite through JuMP framework. We then compute the partial deriva-

[67] We would like to particularly thank the lead developer of DifferentialEquations.jl for helping us work through numerical issues in getting this to work.

tive $\frac{db_t^*(\cdot)}{s}$ using the (analytical) derivative of equation (20), evaluated at the optimal bids found with the numerical solver.

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

B Entry Cost Proofs

B.1 Proof of Entry Cost Lower Bound

Lemma 1. *The cost of entry K is bounded from below by the certainty equivalent of participating in the auction absent an entry cost.*

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:

$$(1 - \exp(-\gamma\bar{\pi}(\alpha))) \cdot (1 - F(\alpha))^{N^*-1} \geq 1 - \exp(-\gamma K) \tag{21}$$

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:

$$\mathbb{E}[u(\pi(s(\alpha), \alpha))] \geq 0, \tag{22}$$

$$\mathbb{E}[u(\pi(s(\alpha),\alpha))] = \underbrace{(1-\exp(\gamma K))}_{\text{Utility on entering and losing}} \cdot \underbrace{\left[1-\left(\prod_{k\neq i}(1-F(\alpha))^{N^*-1}\right)\right]}_{\text{Prob of losing}} +$$

$$\underbrace{\left(1-\exp\left(\gamma K-\gamma\sum_{t=1}^{T}q_t^b(b_t^*(s(\alpha))-\alpha c_t)-\frac{\gamma\sigma_t^2}{2}(b_t^*(s(\alpha))-\alpha^i c_t)^2\right)\right)}_{\text{Expected utility conditional on entering and winning}} \cdot \underbrace{\left(\prod_{k\neq i}(1-F(\alpha))^{N^*-1}\right)}_{\text{Prob of win w/ } s(\alpha)}.$$

Substituting and rearranging inequality 22, 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 21.

Next, consider a bidder in our setting – who must decide whether or not to invest the entry cost K to discover her type and then enter the auction. The bidder will choose to enter if and only if her ex-ante expected utility from entering (upon paying K) is positive:

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

where $\mathbb{E}[u(\pi(s(\tilde{\alpha}),\tilde{\alpha}))]$ is the bidder's expected utility from entering 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: $\int_{\underline{\alpha}}^{\overline{\alpha}}[\cdot]\, f(\tilde{\alpha})d\tilde{\alpha}$.

$$\mathbb{E}_\alpha[\mathbb{E}[u(\pi(s(\tilde{\alpha}),\tilde{\alpha}))]] =$$

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

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

Rearranging equation (23), we have that if $\mathbb{E}_\alpha[\mathbb{E}[u(\pi(s(\tilde{\alpha}),\tilde{\alpha}))]]\geq 0$, then:

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

Solving inequality (22) for K, we obtain that the certainty equivalent of entering the auction at zero cost provides a lower bound on the entry cost of the N^*th bidder.

□

C Technical Details

C.1 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 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.[68] We then add summary statistics for

[68] 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.

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.[69] The resulting matrix is very high dimensional, and so we project the matrix onto its principle components, and use the first 15.[70] 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 (11).

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.[71] 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

[69] 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)

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

[71] 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.

use $X_{i,n}$ and the unique bidder ideas to model α_n^i in equation 15.

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.

C.2 Second Stage Estimation Details

Sample Moment Functions

We now define the sample moment functions used in the second stage of our estimation procedure. 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 point-estimate of θ_1. Denote $\theta_2 = (\gamma, \{\alpha^i\}_i, \vec{\beta}_\alpha)$ for the vector of second stage parameters, and $\mathcal{T}_s$ for the set of "top skew items".

We make use of the following 3 types of moments:

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

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

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

where

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

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 esimate 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:

$$\hat{\theta}_2 = \min_{\theta_2} \left\{ \frac{1}{|\mathcal{I}^*|} \sum_{i^*} m^1_{i^*}(\theta_2|\hat{\theta}_1)^2 + m^2_s(\theta_2|\hat{\theta}_1)^2 + \frac{1}{J} \sum_{j=1}^{J} m^3_j(\theta_2|\hat{\theta}_1)^2 \right\}$$

where $\mathcal{I}^*$ is the set of unique firm IDs, and J is the number of columns in $X_{i,n}$. This optimization problem is solved subject to the constraint that $\alpha^i_n(\theta_2)$ be non-negative for every i and n.[72]

[72] 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.

D Estimation Results Tables

First Stage Model Fit

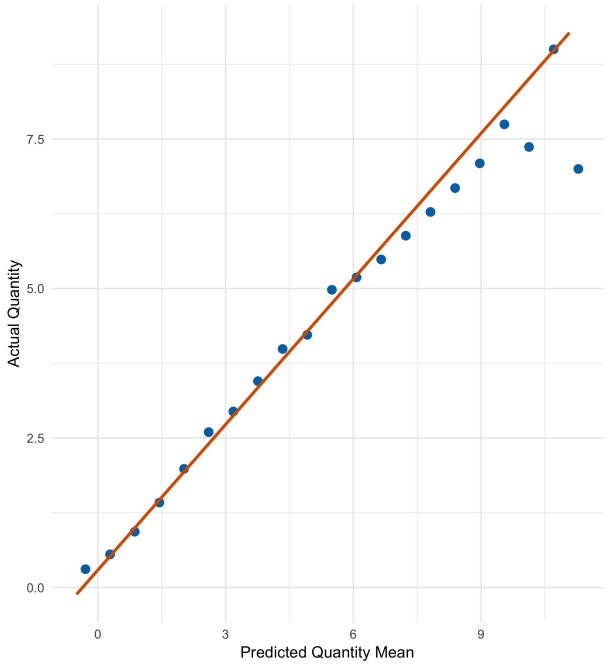


Figure 17: A bin scatter of actual quantities vs model predictions

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

Table 12: Regression report for figure 16

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 13: 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 14: Parameter estimates for the Second Stage GMM estimation

Second Stage Model Fit

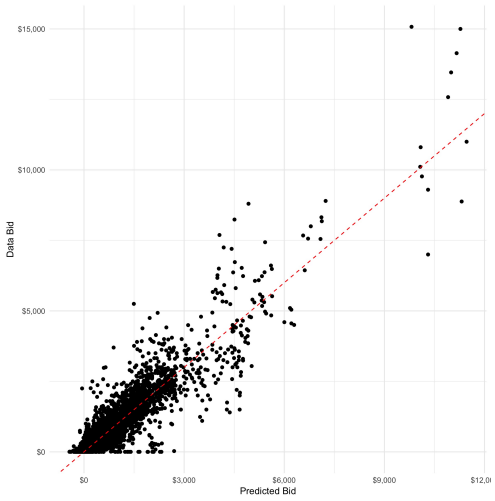


Figure 18: A scatter plot of actual quantities vs model predictions.
Note: Unit bids are scaled so as to standardize quantities so exact dollar values are not representative.

	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

Table 15: Regression report for figure 18

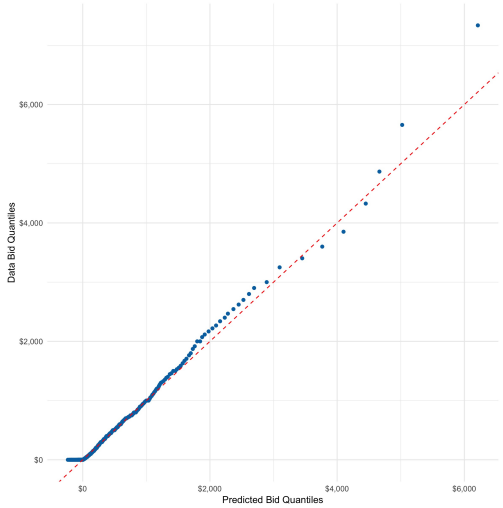


Figure 19: Quantile-Quantile plot of predicted bids against data bids. Quantiles are presented at the 0.0001 level and truncated at the top and bottom 0.01% for clarity. The 45-degree line is dashed in red for reference.
Note: Unit bids are scaled so as to standardize quantities so exact dollar values are not representative.

E Bayesian Sampling

We make the following additional assumptions for a Bayesian approach.

First, we choose priors on the structural parameters. Note that the particular choice of priors has/is being experimented with and results do not appear to be very sensitive to it thus far.[73]

$$\gamma = \frac{1}{\gamma_{raw}} \text{ where } \gamma_{raw} \sim \mathcal{N}_+(5, 1)$$

$$\alpha_n^i \sim \mathcal{N}_+(1, 0.5)$$

The key additional assumption is the modeling of the measurement error on observed bids. For GMM, we assumed only that $b^d_{t,i,n} = b^*_{t,i,n} + \nu_{t,i,n}$ with $\mathbb{E}[\nu_{t,i,n}] = 0$. For the Bayesian approach, we model the distribution of IID draws:

$$\nu_{t,i,n} \sim \mathcal{N}(0, \sigma_b^2),$$

where σ_b is given a prior distribution and estimated.[74]

Note, however, that by the formula for $b^*_{i,t,n}$, the optimal bid (given the auction data and structural parameters), the optimal bid for each item is a function of the optimal total score

[73] We model γ as an exponential transformation to allow for higher flexibility in its level estimate while keeping the raw parameters on a similar scale for computational efficiency.

[74] We use the prior $\sigma_b \sim \mathcal{N}(0, 3)$ at the moment.

$s^*_{i,n}$. We do not observe the optimal score, however - we observe only an "observed" score

$$s^d_{i,n} = \sum_{t \in \mathcal{T}(n)} b^d_{i,t,n} q^e_{t,n} \equiv s^*_{i,n} - \sum_{t \in \mathcal{T}(n)} \nu_{i,t,n} q^e_{t,n}.$$

Note that by construction, the distribution of the error on the observed score is known given the assumptions above:

$$\sum_{t \in \mathcal{T}(n)} \nu_{i,t,n} q^e_{t,n} \sim \mathcal{N}\left(0, \left(\sigma_b^2 \sum_{t \in \mathcal{T}(n)} \left[(q^e_{t,n})^2\right]\right)\right) \text{ given } q^e_{t,n} \text{ and } X_{t,n}$$

Putting these together, we model:

$$s^*_{i,n} \sim \mathcal{N}\left(s^d_{i,n}, \left(\sigma_b^2 \sum_{t \in \mathcal{T}(n)} \left[(q^e_{t,n})^2\right]\right)\right) \text{ and } b^d_{i,t,n} \sim \mathcal{N}\left(b^*_{i,t,n}(s^*_{i,n}), \sigma_b^2\right)$$

Posterior Mode Results for Bridge Auctions The following are summary statistics of the posterior mode of the HMC samples, analogous to those in section 6. Note that while the estimated CARA coefficient here is higher than the GMM estimate, this is in part due to the level of aggregation in the GMM estimate. While we aggregate bidders who appear in less than 10 auctions together for GMM – assigning them the same bidder-specific fixed effect – we treat each bidder-auction pair as an independent draw from the distribution of efficiency types in this estimation procedure. In an upcoming revision, we will present results for an extended Bayesian analysis in which relationships between bidder-auction draws are modeled in a hierarchical fashion, and correlations between bid errors are allowed.

$1/\widehat{\gamma}$	SE
1.66	0.002

Table 16: Estimates for the CARA coefficient. Note that the modal $\widehat{\gamma}$ here is $1/1.77 \approx 0.57$.

Project Type	Mean	Sd	Q1	Median	Q3
Bridge Reconstruction/Rehab	1.157	0.384	0.889	1.095	1.399
Bridge Replacement	1.146	0.338	0.893	1.094	1.333
Structures Maintenance	1.04	0.412	0.819	1.001	1.223

Table 17: Summary statistics of α^i_n estimates by project type.

Project Type	Mean	Sd	Q1	Median	Q3
Bridge Reconstruction/Rehab	9.08	24.93	-6.97	4.55	18.48
Bridge Replacement	3.93	22.21	-9.83	0.54	13.42
Structures Maintenance	61.23	641.64	-3.36	9.3	35.65

Table 18: Summary statistics of estimated bidder markups given $\hat{\alpha}_n^i$.

F Additional Tables and Figures

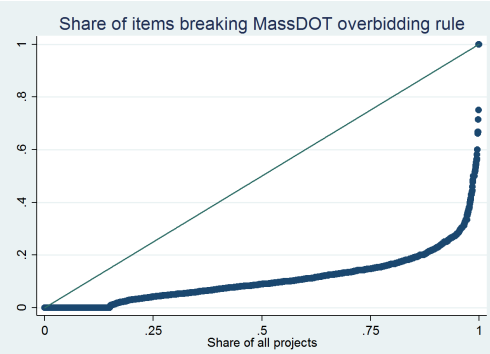
F.1 Distribution of Projects by Year in Our Data

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

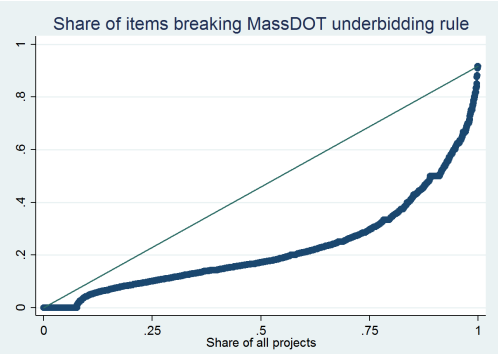
Table 19: Distribution of projects by year in our data

F.2 Shares of Projects with "Unbalanced" Bids

Most projects have a substantial portion of unit bids that should trigger a mathematical unbalancedness flag.



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



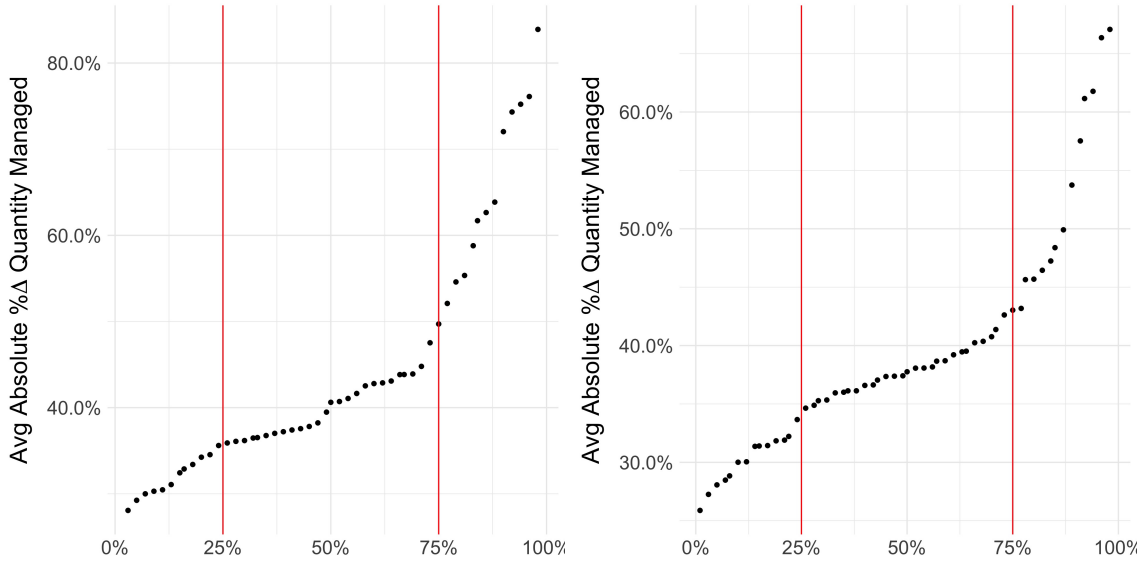
(b) Share of projects (x-axis) that have a particular share of their items breaking the MassDOT underbidding rule (y-axis)

Figure 20

F.3 Discussion of Quantity Uncertainty vis-a-vis Designer and Project Manager Identities

Although many factors could influence the percent overrun for each item, one factor of note is the identity of the designer, resident engineer and project manager in charge. A designer who is less experienced, for example, might be more prone to mis-estimates in the project specification. A project manager who is less experienced might be more prone to making mistakes that necessitate changes. Figures 21a and 21b show the average absolute value of percent quantity overruns across items in projects managed by each project manager or designed by each designer, respectively. There are 53 unique project managers and 57 unique designers. The median project manager worked on 6 projects in our data set, with a mean of 8.4 and a maximum of 38. The median designer worked on 3 projects, with a mean of 7.8 and a maximum of 147 (this is the in-house MassDOT designer team, in contrast to the others, who are consultants). While there is not a clear relationship between absolute overruns and experience, and it is possible that the variation in overruns stems from differences in the projects that each project manager/designer is involved with, the heterogeneity in overruns across project managers and designers suggests that the choice or training of the staff employed by MassDOT could be an avenue for reducing levels of uncertainty.[75]

[75]The full distributions of the number of auctions that each project manager and designer participated in, as well as a plot of average absolute overruns against the number of auctions are included in the appendix, for reference.



(a) Project Managers

(b) Designers

Figure 21: Average absolute value of percent quantity overruns across items managed by each project manager (a) and each designer (b).

G Additional Discussion of the Toy Model

G.1 Savings from Eliminating Risk by Risk and Risk Aversion

In this section, we present additional simulation results for the toy model discussed in section 2. The parameters of the example are described in Table 20 below. In Table 21, we present

	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 20: Auction parameters from the toy model.

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 $\sigma_r^2 = 2 \times 1 = 2$.[76]

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

G.2 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

[76] Note that 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.

unit bids multiplied the DOT's quantity estimates:

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

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

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

The key intuition for bid skewing is as follows. Suppose that the bidders' expectations of the actual 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. Note that 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 charcterized by the optimality of $s(\theta)$ with respect to the expected profits of a bidder with expected total cost θ:

$$E[\pi(s(\theta^i))|\theta^i] = \left(\frac{12}{10} \cdot s(\theta^i) - \theta^i\right) \cdot \text{Prob}\left(s(\theta^i) < s(\theta^j)\right) \tag{25}$$

$$= \left(\frac{12}{10} \cdot s(\theta^i) - \theta^i\right) \cdot \left(1 - F(\theta^i)\right), \tag{26}$$

where the second equality follows from the strict monotonicity of the equilibrium.[77]

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 $s(\theta)$:

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

[77] 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.

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:

$$E[\pi(\theta^i)] = 12 \cdot \frac{s(\theta^i)}{10} - \theta^i \tag{27}$$

$$= \frac{\int_{\theta^i}^{\overline{\theta}}\left[1 - F(\tilde{\theta})\right]d\tilde{\theta}}{1 - F(\theta^i)}. \tag{28}$$

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

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

$$s(\theta) = \min\left\{\frac{q_c^e}{q_c^b}, \frac{q_r^e}{q_r^b}\right\} \cdot \left[\theta + \frac{\int_{\theta}^{\overline{\theta}}\left[1 - F(\tilde{\theta})\right]d\tilde{\theta}}{1 - F(\theta)}\right]. \tag{29}$$

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.

Note that 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.