Online Appendix for:

Scaling Auctions as Insurance: A Case Study in Infrastructure Procurement

(Not for Publication)

Valentin Bolotnyy [*]   Shoshana Vasserman [†]

March 2023

[*]Hoover Institution, Stanford University. Email: vbolotnyy@stanford.edu
[†]Stanford Graduate School of Business and NBER. Email: svass@stanford.edu.

A Additional Tables and Figures

Distribution of Projects by Year in Our Data

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

Table 1: Distribution of projects by year in our data

Estimated Number of Employees for Most Common Firms

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

Table 2: All 24 most common firms in our sample are privately owned, and so there is no publicly available, verifiable information on their revenues or expenses. The numbers of employees presented here were drawn from Manta, an online directory of small businesses, and cross-referenced with LinkedIn, on which a subset of these firms list a range of their employee counts, as of November 2018. Note that there is some ambiguity as to who "counts" as an employee, as such firms often hire additional construction laborers on a project-by-project basis. The "family owned" label is drawn from the firms' self-descriptions on their websites.

Bid Level-Weighted Bins

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

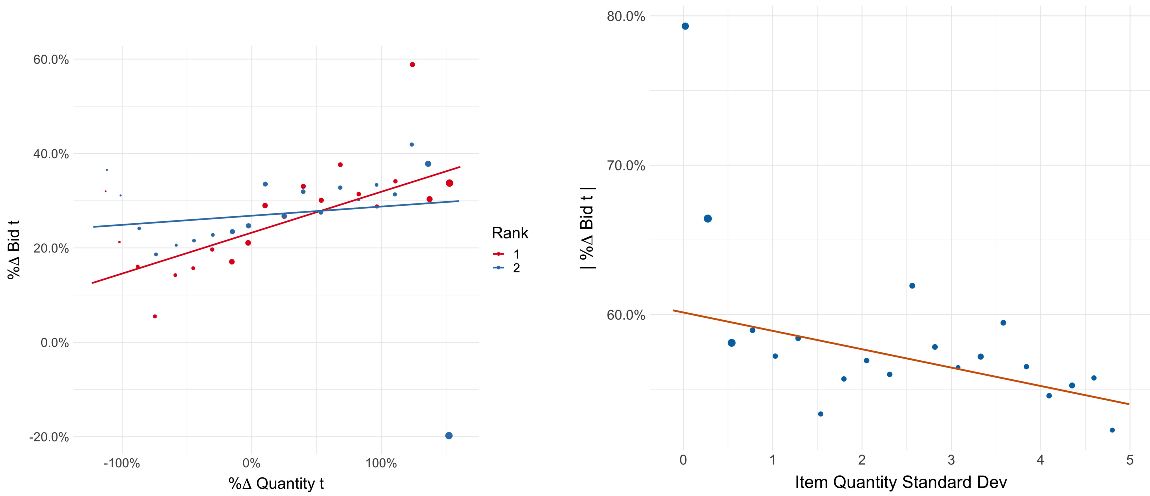

(a) Replication of Figure 3a with weighted bins.

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

Figure 1

Items with the highest and lowest quantity variance To demonstrate the types of items that often have high and low quantity variance, we plot the frequency of items with at least 10 item-auction observations with estimated standard deviations among the highest and lowest 5%.

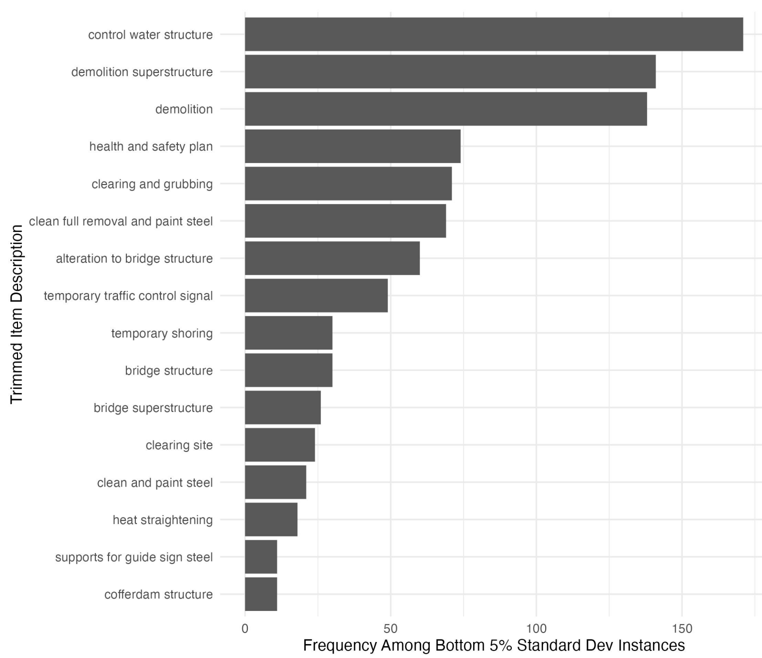

Figure 2: Trimmed description of items with at least 10 instances among the 5% lowest $\widehat{\sigma}_{t,n}$

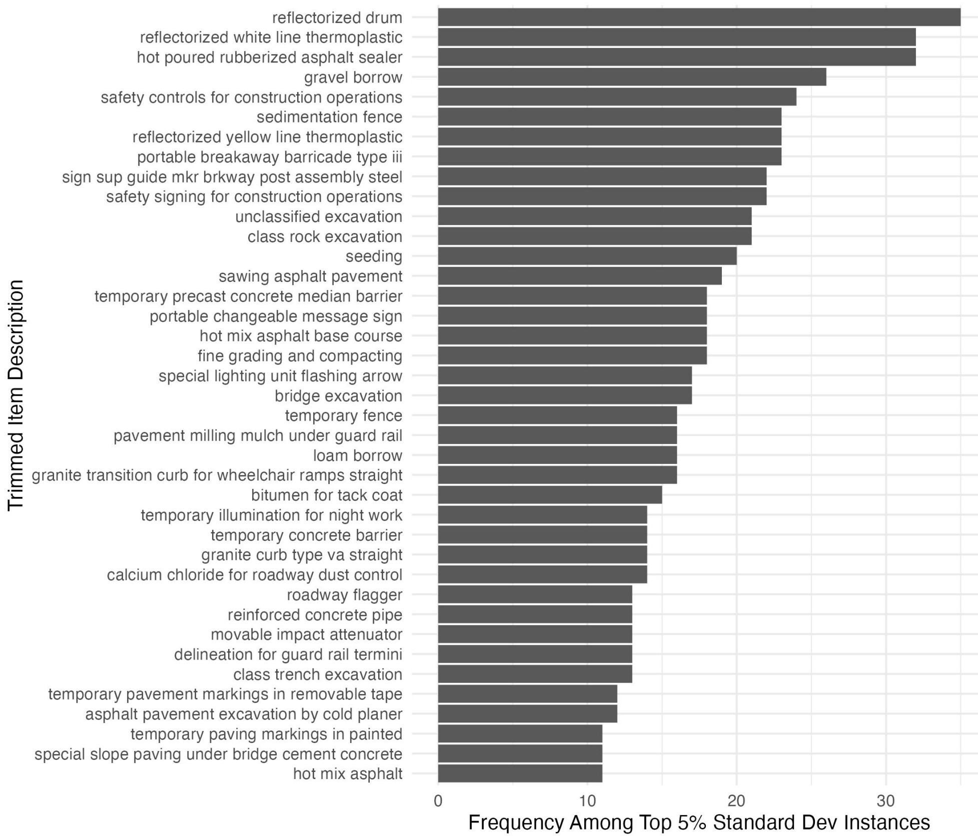

Figure 3: Trimmed description of items with at least 10 instances among the 5% highest $\widehat{\sigma}_{t,n}$

B Illustrative Example

Consider the following simple example of infrastructure procurement bidding. Two bidders compete for a project that requires two types of inputs to complete: concrete and traffic cones. MassDOT ("the DOT" for short) estimates that 10 tons of concrete and 20 traffic cones will be necessary to complete the project. Upon inspection, the bidders determine that the actual quantities of each item that will be used—random variables that we will denote q_c^a and q_r^a for concrete and traffic cones, respectively—are normally distributed with means $\mathbb{E}[q_c^a] = 12$ and $\mathbb{E}[q_r^a] = 16$ and variances $\sigma_c^2 = 2$ and $\sigma_r^2 = 1$.[1] 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.[2]

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

[2]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-

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 describe each bidder's cost as a multiplicative factor α of the market-rate cost estimate for each item: $c_c = \$8$/ton for each ton of Concrete and $c_r = \$12$/pack for each pack of 100 traffic cones. Each bidder i knows her cost 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: $b_{i,c}$ and $b_{i,r}$. 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:

$$s_i = 10b_{i,c} + 20b_{i,r}.$$

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

Bidders are endowed with a standard CARA utility function over their earnings from the project. For simplicity in this example, we assume that the bidders share a common constant coefficient of absolute risk aversion γ:

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

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

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

$$\pi(\mathbf{b}_i, \alpha_i, \mathbf{c}, \mathbf{q}^a) = q_c^a \cdot (b_{i,c} - \alpha_i c_c) + q_r^a \cdot (b_{i,r} - \alpha_i c_r),$$

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

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

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

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

site DOT official, thereby limiting contractors' ability to influence ex-post quantities.

winning:[3]

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

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

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

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

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

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

That is, at every possible score 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 $\alpha_i = 1.5$.[4] 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, $b_{i,c} = \$12$, and a dollar mark-up on traffic cones: $b_{i,r} = (\$500 - \$12 \times 10)/20 = \$19$. Alternatively, she could bid her cost on traffic cones, $b_{i,r} = \$18$, and a two-dollar mark-up on traffic cones: $b_{i,c} = (\$500 - \$18 \times 20)/10 = \$14$. Both of these bids would result in the same score, and so give the bidder the same 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 $b_{i,c} = \$47.78$

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

[4] 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 $\alpha_i \times c_r = 1.5 \times \$12 = \$18$.

and $b_{i,r} = \$1.12$, 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 over-run in quantity – she predicts that 12 tons will be used, whereas the DOT estimated only 10 – and that traffic cones will under-run – she predicts that 16 will be used, rather than the DOT's estimate of 20. When the variance terms aren't too large (relatively), the interpretation is quite simple: every additional dollar bid on concrete is worth approximately 12/10 in expectation, whereas every additional dollar bid on traffic cones is worth only 16/20.

However, the incentive to bid higher on items projected to over-run is dampened when the variance term is relatively large. This can arise when the coefficient of risk aversion is relatively high or when the variance of an item's ex-post quantity distribution is high. More generally, as demonstrated in equation (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.

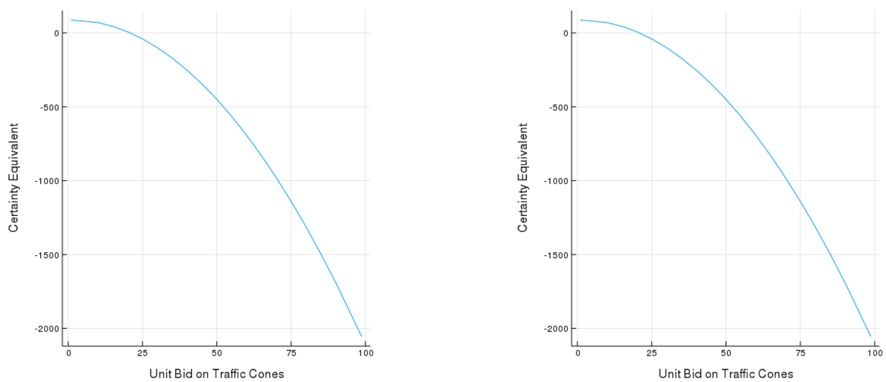

(a) Score = \$500 (b) Score = \$1000

Figure 4: 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 4 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, $b_{i,r} = \$1.12$. In the second case, however, the optimal bid is much higher at $b_{i,r} = \$23.33$. 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.

B.1 Bid Skewing in Equilibrium

As we discuss in Section 5, the auction game described above has a unique Bayes Nash Equilibrium. This equilibrium is characterized following the two-stage procedure described in Appendix A: (1)

given an equilibrium score $s(\alpha)$, each bidder of type α submits the vector of unit bids that maximizes her certainty equivalent conditional on winning, and sums to $s(\alpha)$; (2) The equilibrium score is chosen optimally, such that there does not exist a type α and an alternative score $\tilde{s}$, so that a bidder of type α can attain a higher expected utility with the score $\tilde{s}$ than with $s(\alpha)$.

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

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

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

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

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

Table 3: Comparison of Expected DOT Costs

In Table 3, 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 $q_c^a b_{w,c} + q_r^a b_{w,r}$ at the equilibrium bidding strategy in each setting, taken with respect to the distribution of the type of the lowest (winning) bidder.[5] When bidders are risk neutral ($\gamma = 0$), the equilibrium cost to the DOT does not change when the DOT improves its quantity estimates. The reason for this is that since $\gamma = 0$, the variance term in Equation (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.

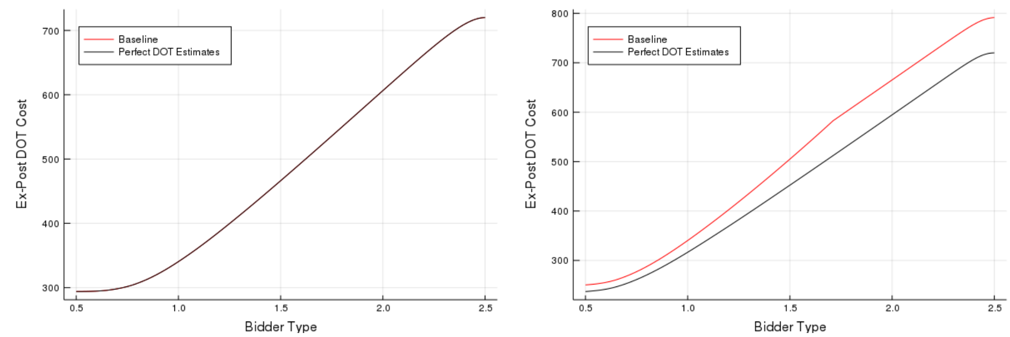

(a) Risk Neutral Bidders (b) Risk Averse Bidders: $\gamma = 0.05$

Figure 5: Equilibrium DOT Cost/Bidder Revenue by Bidder Type

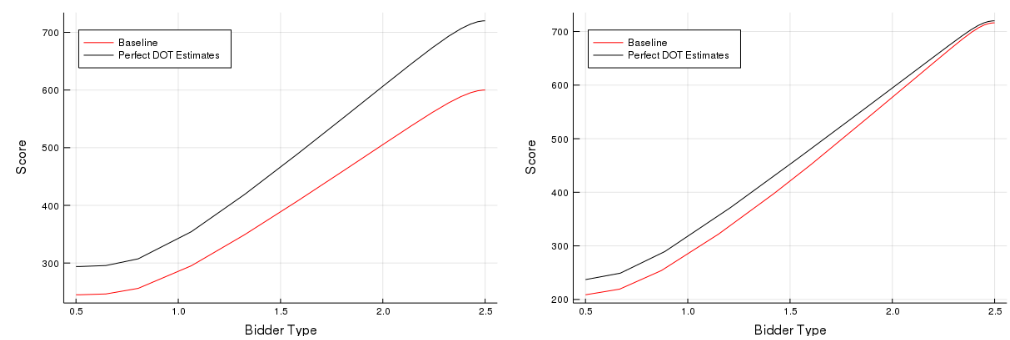

(a) Risk Neutral Bidders (b) Risk Averse Bidders: $\gamma = 0.05$

Figure 6: Equilibrium Score Functions by Bidder Type

In Figure 5a, we plot the revenue that each type of bidder expects to get in equilibrium when bidders are risk neutral. The red line corresponds to the baseline setting, in which the DOT underestimates the ex-post quantity of concrete, and overestimates the ex-post quantity of traffic cones. The black line corresponds to the counterfactual in which both quantities are precisely estimated, and bidders have no residual uncertainty about what the quantities will be. While the ex-post cost to the DOT is the same whether or not the DOT quantity estimates are correct, the unit bids and resulting scores that bidders submit are different. In Figure 6a, we plot the equilibrium score for each bidder type when bidders are risk neutral. The score at every bidder type is smaller

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

under the baseline than under the counterfactual in which the DOT discerns ex-post quantities. This is because the scores in the counterfactual correspond to the bidders' expected revenues, while the scores in the baseline multiply bids that are skewed to up-weight over-running items by their under-estimated DOT quantities. See the Appendix for a full derivation and discussion of the risk neutral case.

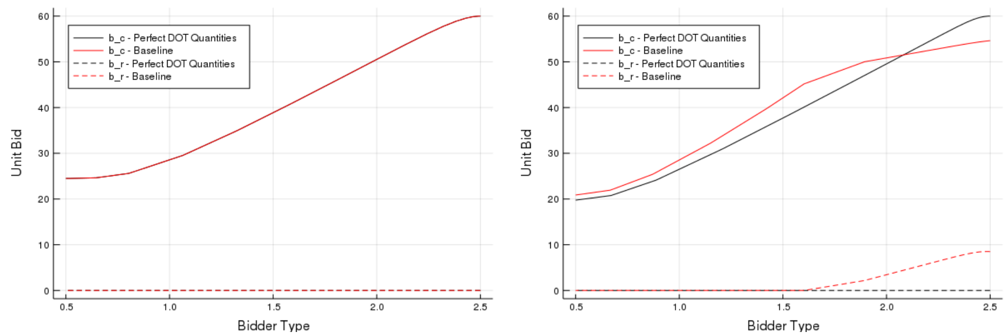

(a) Risk Neutral Bidders (b) Risk Averse Bidders: $\gamma = 0.05$

Figure 7: Equilibrium Unit Bids by Bidder Type

Figure 7a plots the unit bid that each type of bidder submits in equilibrium when bidders are risk neutral. As before, the red lines correspond to the baseline setting in which the DOT mis-estimates quantities, whereas the black lines correspond to the counterfactual setting in which the DOT discerns ex-post quantities perfectly. The solid line in each case corresponds to the unit bid for concrete $b_c(\alpha)$ that each α type of bidder submits in equilibrium. The dashed line corresponds to the equilibrium unit bid for traffic cones $b_r(\alpha)$ for each bidder type. Notably, in every case, the optimal bid for each bidder puts the maximum possible amount (conditional on the bidder's equilibrium score) on the item that is predicted to over-run the most, and \$0 on the other item. This is a direct implication of optimal bidding by risk neutral bidders, absent an external impetus to do otherwise. As noted by Athey and Levin (2001), this suggests that the observations of *interior* or *intermediately-skewed* bids in our data, as well as in Athey and Levin's, are inconsistent with a model of risk neutral bidders. Other work, such as Bajari, Houghton, and Tadelis (2014) have rationalized interior bids by modeling a heuristic penalty for extreme skewing that represents a fear of regulatory rebuke. However, no significant regulatory enforcement against bid skewing has ever been exercised by MassDOT, and discussions of bidding incentives in related papers as well as in Athey and Levin (2001) suggest that risk avoidance is a more likely dominant motive.

In Figures 5b, 6b and 7b, we plot the equilibrium revenue, score and bid for every bidder type, when bidders are risk averse with the CARA coefficient $\gamma = 0.05$. Unlike the risk-neutral case, the DOT's elimination of uncertainty regarding quantities has a tangible impact on DOT costs. When the DOT eliminates quantity risk for the bidders, it substantially increases the value of the project for all of the bidders, causing more competitive bidding behavior. Seen another way, uncertainty regarding ex-post quantities imposes a cost to the bidders, on top of the cost of implementing the project upon winning. In equilibrium, bidders submit bids that allow them to recover all of their costs (plus a mark-up). When uncertainty is eliminated, the cost of the project decreases, and so the total revenue needed to recover each bidder's costs decreases as well. Note, also, that the elimination of uncertainty results in different levels of skewing across the unit bids of different items. Whereas under the baseline, bidders with types $\alpha > 1.6$ place increasing interior bids on traffic cones, when risk is eliminated, this is no longer the case. However, this is subject to a tie breaking rule—when the DOT perfectly predicts ex-post quantities, there are no over-runs, and so there is no meaningful different to overbid on one item over the other. The analysis of the optimal bid (conditional on a score) here is analogous to that under risk neutrality, and so we defer details to the appendix.

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

Table 4: 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 3 over different degrees of risk aversion and different levels of uncertainty. In Table 4, 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.[6] The bolded row with a CARA coefficient of 0.05 corresponds to the right hand column of Table 3.

C 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,

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

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 $c_{i,c}$ for each unit of concrete and $c_{i,r}$ for each traffic cone used. Thus, at the time of bidding, each bidder i expects to incur a total cost

$$\theta_i \equiv \mathbb{E}\left[q_c^a c_{i,c} + q_r^a c_{i,r}\right] = 12c_{i,c} + 10c_{i,r},$$

if she were to win the auction. Each bidder i submits a unit bid for each of the items: $b_{i,c}$ and $b_{i,r}$. The winner of the auction is then chosen on the basis of her *score*: the sum of her unit bids multiplied the DOT's quantity estimates:

$$s_i = 10b_{i,c} + 20b_{i,r}.$$

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:

$$\begin{aligned} E[\pi(b_{i,c}, b_{i,r})|c_{i,c}, c_{i,r}] &= \underbrace{\mathbb{E}\left[(q_c^a b_{i,c} + q_r^a b_{i,r}) - (q_c^a c_{i,c} + q_r^a c_{i,r})\right]}_{\text{Expected profits conditional on winning}} \times \underbrace{\text{Prob}(s_i < s_j)}_{\text{Probability of winning}} \\ &= \left((12b_{i,c} + 10b_{i,r}) - \theta_i\right) \times \text{Prob}\left((10b_{i,c} + 20b_{i,r}) < (10b_{j,c} + 20b_{j,r})\right). \end{aligned}$$

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 $c_{i,c} = \$70$ and $c_{i,r} = \$3$, 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: $b_{i,c} = \$75$ and $b_{i,r} = \$12.50$, yielding a net profit of \$155. However, if instead, she bids $b_{i,c} = \$99.98$ and $b_{i,r} = \$0.01$, 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 $12b_{i,c} + 20b_{i,r} = \$1000$— the exact revenue that she would earn upon winning the auction. She could still bid $b_{i,r} = \$0.01$, for example, but then she would need to bid $b_{i,c} = \$83.33$, resulting in a revenue of \$1000 and a profit of \$130 if she wins the auction. A quick inspection shows that no choice of $b_{i,c}$ and $b_{i,r}$ could improve her expected revenue at the same score.

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

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

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

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

$$\{b_c(\theta_i), b_r(\theta_i)\} = \arg\max_{\{b_c, b_r\}} \Big\{12b_c + 40b_r - \theta_i\Big\} \text{ s.t. } 10b_c + 50b_r = s(\theta_i).$$

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

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

$$E[\pi(s(\theta_i))|\theta_i] = \left(\frac{12}{10} \cdot s(\theta_i) - \theta_i\right) \cdot \text{Prob}\left(s(\theta_i) < s(\theta_j)\right) \tag{2}$$

$$= \left(\frac{12}{10} \cdot s(\theta_i) - \theta_i\right) \cdot \left(1 - F(\theta_i)\right), \tag{3}$$

where the second equality follows from the strict monotonicity of the equilibrium.[7]

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 $s(\theta)$:

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

[7]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 $b_r(\theta_i) = 0$. If bidder i wins the auction, she expects to earn a markup of:

$$E[\pi(\theta_i)] = 12 \cdot \frac{s(\theta_i)}{10} - \theta_i \tag{4}$$

$$= \frac{\int_{\theta_i}^{\overline{\theta}}\left[1 - F(\tilde{\theta})\right]d\tilde{\theta}}{1 - F(\theta_i)}. \tag{5}$$

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:

$$s(\theta) = \min\left\{\frac{q_c^e}{q_c^b}, \frac{q_r^e}{q_r^b}\right\} \cdot \left[\theta + \frac{\int_{\theta}^{\overline{\theta}}\left[1 - F(\tilde{\theta})\right]d\tilde{\theta}}{1 - F(\theta)}\right]. \tag{6}$$

Suppose that $\frac{q_r^e}{q_r^b} \leq \frac{q_c^e}{q_c^b}$. Then bidder i will bid $b_r^*(\theta_i) = \frac{s(\theta_i)}{q_r^e}$ and $b_c^*(\theta_i) = 0$. Consequently, if bidder i wins, she will be paid $q_r^b \cdot b_r^*(\theta_i) = \left[\theta_i + \frac{\int_{\theta_i}^{\overline{\theta}}\left[1 - F(\tilde{\theta})\right]d\tilde{\theta}}{1 - F(\theta_i)}\right]$ as in our example.

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

D Discussion of Policy Inefficacy If Bidders are Risk Neutral

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

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

Thus, bidding is effectively one-dimensional, and all of the properties of standard single-item first price auctions with risk neutral bidders apply. In particular, not only would the policies to reduce uncertainty or switch to a lump sum auction be ineffective (which follows directly from the model as "risk" does not enter into bidder preferences), but so would a policy to subsidize bidders. This result is a consequence of revenue equivalence: the decrease in expected costs from the entry of an additional bidder is equal to the expected profit that this bidder would earn upon entering.[8] As such, the equilibrium number of entrants to a given auction is necessarily efficient: incentivizing an additional bidder would cost the full additional surplus that this bidder would bring. By contrast, revenue equivalence does not apply with risk averse bidders, and the additional bidder's expected utility from entering may be lower than the decrease in expected costs if she enters.

E Views of Bid Skewing by Contractors and MassDOT Managers

Bid Skewing Among Contractors

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

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

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

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

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

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

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

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

MassDOT Challenges to Bid Skewing

Concerns that sophisticated bidding strategies may allow contractors to extract excessively large payments have led to a number of lawsuits about MassDOT's right to reject suspicious bids. The Federal Highway Administration (FHWA) has explicit policies that allow officials to reject bids that are deemed manipulative. However, the legal burden of proof for a manipulative bid is quite high. In order for a bid to be legally rejected, it must be proven to be *materially unbalanced*.[9]

[9] 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)).

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.[10]

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."[11] In other words, it is a bid that appears to be strategically skewed. In order to discourage bid skewing, many regional DOTs use concrete criteria to define mathematically unbalanced bids. In Massachusetts, a bid is considered mathematically unbalanced if it contains any line-item for which the unit bid is (1) over (under) the office cost estimate and (2) over (under) the average unit bid of bidders ranked 2-5 by more than 25%.

In principle, a mathematically unbalanced bid elicits a flag for MassDOT officials to examine the possibility of material unbalancedness. However, in practice, such bids are ubiquitous, and substantial challenges by MassDOT are very rare. In our data, only about 20% of projects do not have a single item breaking MassDOT's overbidding rule, and only about 10% of projects do not have a single item breaking the underbidding rule. Indeed, most projects have a substantial portion of unit bids that should trigger a mathematical unbalancedness flag. However, only 2.5% of projects have seen bidders rejected across all justifications, a handful of which were due to unbalanced bids.[12]

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).[13] In

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

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

[12] 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.

[13] 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.

January 2017, MassDOT attempted to require a minimum bid for every unit price item in a various locations contract due to bid skewing concerns. SPS New England, Inc. protested, arguing that such rules preclude the project from being awarded to the lowest responsible bidder. The Massachusetts Assistant Attorney General ruled in favor of the contractor on August 1, 2017.

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

F Solving Portfolio Problems

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

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

The primal formulation of this problem is

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

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

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

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

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

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

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

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

$$\min_v \left\{\sum_i h_i(x_i^*) - v\,(s)\right\} \tag{7}$$

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

Simplifying this further,

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

Thus, the solution to the original problem is the v_k^* that minimizes Equation (7) with k non-zero components with the form $x_i^* = \frac{a_i+v_k^* q_i}{2b_i}$. Noting that $\frac{a_i+vq_i}{2b_i} > 0 \iff v > \frac{-a_i}{q_i}$, we propose the following algorithm for solving the problem:

1. Rank $\{i\}$ in order of $\frac{-a_i}{q_i}$ (lowest to highest). Note that under this sorting, for any v, if $v \leq \frac{-a_j}{q_j}$ for some j, then $v \leq \frac{-a_k}{q_k}$ for all $k > j$. Consequently if $v \leq \frac{-a_j}{q_j}$ for some j, then $h_k(x_k^*(v)) = 0$ for all $k > j$ as well, so that we do not need to consider the contributions of elements $k > j$ in the objective.
2. For each k, let $\tilde{v}_k$ to be the value of v that minimizes Equation (7) on the interval $(\frac{-a_k}{q_k}, \frac{-a_{k+1}}{q_{k+1}}]$. Iteratively search over indices k, computing $\tilde{v}_k$, and compare them to find the global minimizer among them.

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

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

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

$$\tilde{v}_k^* = \min\left\{\frac{2s - \sum\limits_{i \leq k} \frac{a_i q_i}{b_i}}{\sum\limits_{i \leq k} \frac{1}{b_i} q_i^2}, \frac{-a_{k+1}}{q_{k+1}}\right\}. \tag{8}$$

Thus, we can compute $\tilde{v}_k^*$ for each sequential index k with two operations by considering just the kth contribution to Equation (8). Finally, we compare each $\tilde{v}_k^*$ across indices k to find the global minimizer.

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

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

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

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

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

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

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

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

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

G Model Extensions

G.1 Constant Relative Risk Aversion

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

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

$$u(\pi) = \frac{\pi^{1-\gamma}}{1-\gamma}. \tag{9}$$

As in our main model, we will assume that bidders are differentiated along their cost efficiency type α. For simplicity, in this extension, we will assume that all bidders within the same auction share a common CRRA coefficient γ.

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

$$E[u(x)] = \frac{\exp\left[(1-\gamma)\mu + \frac{1}{2}(1-\gamma)^2\sigma^2\right]}{1-\gamma}. \tag{10}$$

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

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

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

$$\pi(\mathbf{b}) = \sum_t q_t^a \cdot (b_t - \alpha c_t) = \sum_t q_t^e \cdot \left(\frac{q_t^a}{q_t^e}\right) \cdot (b_t - \alpha c_t). \tag{11}$$

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

$$\pi(\mathbf{b}) = \left(s - \alpha \sum_k q_k^e c_k\right) \cdot \left[\sum_t \left(\frac{q_t^a}{q_t^e}\right) \cdot \frac{q_t^e \cdot (b_t - \alpha c_t)}{s - \alpha \sum_k q_k^e c_k}\right] \tag{12}$$

$$= R(s) \cdot \left[\sum_t \left(\frac{q_t^a}{q_t^e}\right) \cdot \frac{q_t^e \cdot (b_t - \alpha c_t)}{R(s)}\right] \tag{13}$$

$$= R(s) \cdot \left[\sum_t w_t \exp(p_t)\right], \tag{14}$$

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

While Equation (14) looks like a portfolio, it is not a lognormal random variable. However, taking logs and a second order Taylor approximation thereof, we achieve:

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

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

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

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

$$\log(E[u(\pi(\mathbf{b}))]) \approx \log\left(\frac{[R(s)]^{1-\gamma}}{1-\gamma}\right) + \frac{1-\gamma}{[R(s)]^2} \cdot \left[\sum_t R(s) \cdot q_t^e \cdot (b_t - \alpha c_t) \cdot \left(\bar{p}_t + \frac{1}{2}\sigma_t^2\right) - \frac{\gamma}{2}\sum_t (q_t^e)^2 \sigma_t^2 \cdot (b_t - \alpha c_t)^2\right]. \quad (15)$$

Computing Optimal Bids As in the CARA case, the expected utility of profits upon winning in Equation (15) yields a portfolio problem. For a given score s, the bidder will choose a vector of bids $b_t(s)$ that maximize Equation (15) subject to the condition that they are all non-negative and sum to s. Given the non-negativity constraints, there is again no closed form solution for optimal unit bids. However, at a given solution, every non-zero optimal unit bid has the following form:

$$b_t^* = \alpha \cdot \left[c_t - \frac{\sum_{k:b_k>0} q_k^e c_k}{q_t^e \sigma_t^2 \sum_{k:b_k>0} \frac{1}{\sigma_k^2}}\right] + \frac{R(s)}{\gamma} \cdot \left[\frac{\frac{\bar{p}_t}{\sigma_t^2} + \frac{1}{2}}{q_t^e} - \frac{\sum_{k:b_k>0} \frac{\bar{p}_k}{\sigma_k^2} + \frac{1}{2}}{q_t^e \sigma_t^2 \sum_{k:b_k>0} \frac{1}{\sigma_k^2}}\right] + \left[\frac{s}{q_t^e \sigma_t^2 \sum_{k:b_k>0} \frac{1}{\sigma_k^2}}\right]. \quad (16)$$

Estimation Our estimation procedure mirrors that of the CARA model, making adjustments as necessary for the alternative set of assumptions. First, we estimate a first stage model of log-quantity-ratio predictions and variances. Denoting an auction by n, we assume that the posterior distribution of each $\log(q_{t,n}^a/q_{t,n}^e)$ is given by a statistical model that conditions on item characteristics (e.g. the item's type classification), observable project characteristics (e.g. the project's location, project manager, designer, etc.), and the history of DOT projects. In particular, we model the realization of the actual quantity of item t in auction n as:

$$\log(q_{t,n}^a/q_{t,n}^e) = p_{t,n}^b + \eta_{t,n}, \text{ where } \eta_{t,n} \sim \mathcal{N}(0, \sigma_{t,n}^2) \quad (17)$$

$$\text{such that} \quad p_{t,n}^b = \beta_p X_{t,n} \text{ and } \hat{\sigma}_{t,n} = \exp(\beta_\sigma X_{t,n}). \quad (18)$$

As in the CARA case, we estimate this model with Hamiltonian Monte Carlo and use the posterior means $\widehat{p}_{t,n}^b$ and $\widehat{\sigma}_{t,n}$ directly in the second stage of our procedure. Plugging the first stage estimates into Equation (16), we obtain a prediction of each bidder i's unit bid for each item t given her observed score $s_{i,n}$ in auction n. We therefore adapt Assumption 1 (Section 6) once more and use the analogous moment conditions to the CARA case to estimate a project-wide CRRA parameter

γ_n for every auction and an efficiency coefficient for every bidder-auction pair $\alpha_{i,n}$. Given the structure of CRRA, bidder wealth is a relevant part of the utility function, that influences the optimal portfolio choice. While this again makes use of several approximations, we account for wealth by augmenting $R(s) \equiv W + s - \alpha \sum_k q_k^e c_k$ throughout our model. We allow wealth W to vary from project to project, and estimate it by projecting onto the matrix of average project-bidder characteristics X_n. We present our estimates in Table 5 below. The resulting estimates of α are higher than in the heterogeneous γ CARA case in the main model of our paper. However, the CRRA coefficients are mostly between 0.6 and 0.8, which is within the region found in related studies. See pg 133 of Campo, Guerre, Perrigne, and Vuong (2011) for a discussion.

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

Table 5: CRRA Estimate Summary Statistics Across All Projects

Computing Equilibrium Outcomes We construct equilibria analogously to the CARA case. For simplicity, we consider only equilibrium bidding conditional on the observed number of bidders in each auction, and do not account for endogenous entry decisions.

To see how this is done, note that the analog of Equation (18) (Appendix A) is given by:

$$\begin{aligned} \mathrm{EU}(\sigma(\alpha_i), \alpha_i) = \exp\left[(1-\gamma) \cdot V(\sigma(\alpha_i), \alpha_i)\right] \cdot (1 - F(\sigma^{-1}(\sigma(\alpha_i))) \\ + W^{1-\gamma} \cdot F(\sigma^{-1}(\sigma(\alpha_i)), \end{aligned} \tag{19}$$

where

$$\begin{aligned} V(\sigma(\alpha_i), \alpha_i) = \log\left(R(\sigma(\alpha_i), \alpha_i)\right) + \\ \frac{1}{R(\sigma(\alpha_i), \alpha_i)} \cdot \left[\sum_t q_t^e \cdot (b_t(\sigma(\alpha_i)) - \alpha_i c_t) \cdot \left(\bar{p}_t + \frac{1}{2}\sigma_t^2\right)\right] \\ - \frac{\gamma/2}{R(\sigma(\alpha_i), \alpha_i)^2} \cdot \left[\sum_t (q_t^e)^2 \sigma_t^2 \cdot (b_t(\sigma(\alpha_i)) - \alpha_i c_t)^2\right], \end{aligned} \tag{20}$$

and $b_t(\sigma(\alpha_i))$ is chosen according to the portfolio problem, with $R = W + \sigma(\alpha_i) - \alpha_i \sum_r q_r^e c_r$. The equilibrium function $\sigma(\alpha)$ is thus given by the solution to the differential equation:

$$\sigma'(\alpha) = \frac{f(\alpha)}{1 - F(\alpha)} \frac{1}{(1-\gamma)\frac{\partial V(\sigma(\alpha), \alpha)}{\partial s}} \left[1 - \frac{W^{1-\gamma}}{\exp\left[(1-\gamma) \cdot V(\sigma(\alpha), \alpha)\right]}\right], \tag{21}$$

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

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

Table 6: Summary of Counterfactual Changes in Expected DOT Spending

subject to the boundary condition that the highest α type is indifferent between participating in the auction or not: $V(\sigma(\overline{\alpha}), \overline{\alpha}) = \log(W)$. We present summary results for the lump sum and no-risk counterfactuals in Table 6.

G.2 Asymmetric Bidder Types

To construct counterfactuals in Section 8, we assume that bidder types (α, γ) are fully characterized by a one-dimensional meta-type τ, which distributed IID, according to a common auction-specific distribution. Under these assumptions, we construct the unique symmetric equilibrium in monotone strategies. In this section, we discuss an extension to a model in which bidder types are characterized by different known CARA coefficients, and separate, uncorrelated IID draws of efficiency types α. In this case, a symmetric equilibrium may not exist and we must consider asymmetric equilibria.

Modeling and Estimation To simplify our analysis for demonstrative purposes, we split bidders into two groups based on the number of auctions that they participated in throughout our data. Of the 25 unique bidder ids in our sample, we label 14 as "frequent" bidders because they participated in 50 or more auctions; we label the remaining 11 "infrequent".[14] Note that while this splits the sample of bidders roughly in half, the distributions of participation are quite different in the two groups: the average "frequent" firm participated in 123 auctions, whereas the average "infrequent" firm participated in 20.

In order to generate counterfactuals with an internally consistent set of type estimates, we re-estimate our second stage model under the constraint that bidder risk aversion types $\gamma_{i,n}$ are homogeneous within their frequency groups. To do this, we modify the second part of Equation

[14]Recall that the 25 unique bidder ids refer to 24 "common" firms, who participated in at least 30 auctions, and 92 "rare" firms who are grouped together because they did not participate in enough auctions to identify an individual firm fixed-effect in the bidder type regressions. Here, we split the "common" firms further in order to divide our sample into *highly* active firms, and the rest.

(28) in Appendix C.3.2 as follows:

$$\frac{1}{\gamma_{i,n}} = \beta_{0,\gamma}^{g(i)} + \beta_{X,\gamma} X_n + \nu_{\gamma,i,n}, \tag{22}$$

where $g(i) \in \{\mathrm{f}, \mathrm{r}\}$ denotes whether the bidder is a frequent bidder and X_n is the vector of average bidder-auction characteristics in auction n. Leaving the rest of the model unchanged, we estimate the modified second stage following the GMM procedure specified in our paper.

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

Table 7: 2-Type Estimate Summary Statistics Across All Projects

Our results, summarized in Table 7, show that frequent bidders are not generally more cost efficient than infrequent bidders. While the median cost multiplier is about the same in both groups, the 25th percentile of infrequent firms is about 8% more efficient than the corresponding percentile of frequent firms, but the 75th percentile of infrequent firms is 3% less efficient than frequent firms. On the other hand, we find that frequent bidders exhibit slightly less risk aversion than less frequent bidders. Neither of these results are surprising: frequent bidders likely specialize in work for the DOT. As such, they participate in auctions where they do not have a particularly strong efficiency advantage. But their risks are also spread across more auctions, and so the weight of uncertainty in any particular project is lower.

Computing Equilibrium Outcomes To compute counterfactuals, we assume that bidder efficiency types α are drawn IID from an auction-wide distribution (characterized similarly to the distribution of τ in our paper) and that the estimated group-specific CARA coefficients are publicly known. Under these assumptions, by Reny and Zamir (2004), there exists an asymmetric equilibrium in monotonic strategies such that each bidder type bids according to the monotone function $\sigma_g : [\underline{\alpha}, \overline{\alpha}] \to \mathbb{R}$.[15] For simplicity, we consider only equilibrium bidding conditional on the observed numbers of bidders in each auction and do not account for endogenous entry decisions.

[15]Uniqueness may not be guaranteed and so our construction procedure may be thought of as imposing an equilibrium selection criterion.

Given strategies σ_g, the probability that i will win the auction under bid s is given by:

$$\prod_{j\neq i}\left[1-H_j\left(s\right)\right]=\prod_{j\neq i}\left[Pr(s<\sigma_j(\alpha_j))\right] \tag{23}$$

$$=\prod_{j\neq i}\left[1-F(\sigma_j^{-1}(s))\right] \tag{24}$$

$$=\prod_{j\neq i}\left[1-F(\varphi_j(s))\right] \tag{25}$$

where the final equality relabels the inverse bid function for convenience. Applying logic analogous to that in Appendix A of our paper, we derive the following set of differential equations:

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

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

$$\text{where}\quad V(\varphi_g(\tilde{s}))=1-\exp(-\gamma_g\cdot CE(\tilde{s},\gamma_g,\varphi_g(\tilde{s}))) \tag{26}$$

$$\text{and}\quad CE(\tilde{s},\gamma_g,\varphi_g(\tilde{s}))=\sum_{t=1}^{T}q_t^b\left(b_t^*(\tilde{s})-\varphi_g(\tilde{s})c_t\right)-\frac{\gamma_g\sigma_t^2}{2}\left(b_t^*(\tilde{s})-\varphi_g(\tilde{s})c_t\right)^2. \tag{27}$$

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

There are three possible cases. The "standard" case[16] requires that the highest and lowest α types of both frequency groups submit the same score, and that the highest α type of the more risk averse group (who has no chance of winning) earn a certainty equivalent of zero. This applies when there is only one bidder of the less risk averse group, or if project portfolio risk is irrelevant as in the no-risk counterfactual. However, as noted in Hubbard and Kirkegaard (2019), the standard conditions do not generally induce an equilibrium in an asymmetric auction with more than 2 bidders when the bidders have different supports for their value distributions. In our case, while both frequency groups have the same support of cost efficiency types α, their different CARA

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

coefficients induce different (overlapping) supports on the certainty equivalents of their portfolios at each score.

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

To compute equilibria for each auction, we solve a boundary value problem with a shooting algorithm based on Hubbard and Paarsch (2014). To account for bifurcation, we allow for two regions in the ODE solution: a homogeneous region in which only one group of bidders compete against each other, and a heterogeneous region that follows the equations described in this section. We find the beginning of the heterogeneous region by iteratively checking for the first score at which both groups of bidders have a type willing to participate in bidding.[17] While shooting methods are known to be highly sensitive to the step size of their integration method, we found that a modified shooting method with high-step-size Euler integration worked most consistently for our problem.

To demonstrate how asymmetric equilibria look in comparison to symmetric equilibria with similar informational assumptions, in Table 8 we present a comparison of the no-risk and lump-sum counterfactuals under a symmetric 1-type (single, publicly known CARA coefficient) model and under the asymmetric 2-type (two groups with different CARA coefficients) model described above for a sample of auctions in our data.[18] The sample accounts for roughly half of the projects in our dataset and is similar in distribution to projects of the Bridge Reconstruction and Rehabilitation category. Among the projects included, the average γ in the 1-type specification is 0.49 (with standard deviation 0.083 and median 0.029). By contrast, the average frequent type γ is 0.040 (s.d., 0.065; median, 0.028) and the average infrequent type γ is 0.048 (s.d., 0.048; median, 0.027).[19]

In general, we find that the results between the 1-type and 2-type cases are within a few

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

[18] To see a full description of the 1-type case, see the 2022 version of this paper.

[19] We chose projects to include in the sample on the basis of computational efficiency. This tends to select for projects with lower overall risk aversion and costs, as these projects exhibit more stable numerical behavior at the tails. The total sample is 248 projects. However as we were not able to get convergence on all projects in all auction formats, our comparison includes 237 projects in the no risk counterfactual and 204 projects in the lump sum counterfactual.

percentage points of each other. For the median auction in this sample, the 2-type model predicts a risk premium that is 0.4 percentage points higher and a lump sum cost that is 0.2 percentage points lower than the 1-type model. A similar pattern holds at the average and every quartile as well. This suggests that the value of lowering risk by using a scaling format or by reducing uncertainty may be a bit lower if bidders are asymmetric, relative to a symmetric model in which all bidders in the same auction have the same level of risk aversion.

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

Table 8: Comparisons of 1-type vs 2-type Counterfactuals for a Sample of Auctions

References

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

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

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.

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

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

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

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

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

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

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

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

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