Procurement Contracting with Time Incentives: Theory and Evidence *

Patrick Bajari
University of Minnesota and NBER

Gregory Lewis
Harvard University

March 23, 2009

Abstract

In public sector procurement, social welfare often depends on the time taken to complete the contract. A leading example is highway construction, where slow completion times inflict a negative externality on commuters. Recently, highway departments have introduced innovative contracting methods that give contractors explicit time incentives. We characterize equilibrium bidding and efficient design of these contracts. We then gather a unique data set of highway repair projects awarded by the Minnesota Department of Transportation that includes both innovative and standard contracts. Descriptive analysis shows that for both contract types, contractors respond to the incentives as the theory predicts, both at the bidding stage and after the contract is awarded. Next we build a structural econometric model that endogenizes project completion times, and perform counterfactual policy analysis. Our estimates suggest that switching from standard contracts to designs with socially efficient time incentives would increase welfare by over 19% of the contract value; or in terms of the 2009 Mn/DOT budget, $290 million. We conclude that large improvements in social welfare are possible through the use of improved contract design.

*We are grateful to Mn/DOT for data, and to Rabinder Bains, Tom Ravn and Gus Wagner for their help. We would also like to thank John Asker, Susan Athey, Matt Gentzkow, Ariel Pakes, Chad Syverson and participants at the IIOC, Stony Brook, NBER IO and WBEC conferences for helpful comments and suggestions. Jason Kriss, Tina Marsh, Maryam Saeedi and Connan Snider provided excellent research assistance

1 Introduction

It is estimated that approximately 15 percent of world output is accounted for by public sector procurement. Designing efficient mechanisms for procurement is therefore essential for guaranteeing the efficient allocation of goods and services. In the United States, auctions are typically used to award procurement contracts to the lowest qualified bidder. In many contexts, however, social welfare depends upon the time to complete the contract. Unfortunately, standard procurement mechanisms do not weight expected completion time in selecting the winning bidder. This suggests that it may be possible to increase social welfare by including project completion time in the auction design.

We take as a case study the design of time incentives in award procedures for highway procurement. Highway repair generates significant negative externalities for commuters through increased gridlock and commuting times. For example, Interstate 35W is a main commuting route in Minneapolis carrying over 175,000 commuters per day. If a highway construction project results in a 30 minute delay each way for commuters on this route, the daily social cost imposed by the construction would be 175,000 hours. If we value time at $10 an hour, this is a social cost of $1.75 million per day. But in standard highway contracts, contractors have poor incentives to internalize this externality. For example, highway contractors in Minnesota pay damages of between $150 and $3500 dollars per day late, depending on the contract size. Given these weak incentives, it is likely that the observed completion times will be inefficiently slow.

Recently, state highway departments in the US have started to experiment with innovative auction designs that provide explicit time incentives. The most sophisticated is called A+B bidding. Here contractors submit a dollar bid for labor and materials, the "A" part, and a total number of days to complete the project, the "B" part. The bids are scored using both the A and the B bid and the project is awarded to the contractor with the lowest score. The winning contractor may also receive incentive payments (disincentives) for completing the project earlier (later) than the days bid. Standard highway contracts are "A-only" contracts because they do not weight project completion time in selecting the contractor.

In this paper, we compare these contracts designs both theoretically and empirically. We start by building a theoretical model of A+B contracts that includes standard A-only contracts and other commonly used contract designs as special cases. We characterize equilibrium bids and project completion times. The theory shows that for many A+B designs,

contractors will bid completion times that differ substantially from their internal estimates of how long it will take. However, by optimally selecting the scoring rule for an A+B auction, we demonstrate that even with this strategic bidding the observed outcome will be both ex-ante and ex-post efficient. Our results imply that the more commonly used A-only contracts will generally result in an inefficiently slow project completion time.

In our empirical analysis we build a unique dataset of bids and project completion times for 248 A-only and 29 A+B contracts awarded by the Minnesota Department of Transportation between 1997-2007. Our data includes detailed information on the progress of the projects, damages charged for late project completion on A-only contracts, and the scoring rules used in A + B contracts. We find that the observed bidding patterns are broadly consistent with the predictions of our theory. In particular we show that in both A-only and A+B contracts, bidders respond to the time incentives, adjusting their bids to manipulate the incentive structure they will face if they win the auction. We also show that time incentives do work: A+B contracts are completed significantly faster, at a rate increasing in the weight given to time in the scoring rule. That said, we find that for some forms of A+B contracts, bidders do not behave quite as the theory predicts, suggesting that some designs should in practice be preferred over others.

Next we propose a method for structurally estimating the contractor's time-related costs. Our theoretical model implies that the observed project completion time equates the marginal benefit and marginal cost of delay. In A-only contracts, the marginal benefit of delay is exogenously specified by state regulations regarding daily damages for late completion, which vary with contract size and year. So given this exogenous variation in marginal benefits, we can use the first order condition to back out the marginal costs. Then having estimated the structural time cost parameters, we can infer the counterfactual completion time under different incentive structures.

In the final section of the paper, we estimate the welfare gain to using the efficient auction mechanism. Our estimates suggest that the welfare gain is large, over 19% of the contract value in lost surplus in an average contract. Moreover, since budget constraints may be important in practice, we show that a constrained efficient policy could achieve much of this gain (8.5% of contract value, or around 145M a year) at relatively low cost to (0.87% of contract value, or around 14.7M). This motivates our main conclusion, which is that including stronger time incentives in highway procurement through better contract design would result in large social welfare improvements.

This paper is related to four main literatures. There is a literature in engineering on the role of time incentives in highway procurement (see for example Arditi, Khisty and Yasamis (1997) and Herbsman, Chen and Epstein (1995)). These papers are primarily descriptive, and do not provide the counterfactual welfare analysis that we obtain from our richer approach.

The second is the large theoretical literature on optimal procurement (see for example Laffont and Tirole (1987), Manelli and Vincent (1995), Branco (1997)). In our analysis of the A+B auctions, we follow the existing literature on scoring auctions starting with Che (1993), allowing for multidimensional type as in Asker and Cantillon (2008b)). We focus on welfare-maximizing, rather than cost-minimizing contract design, thereby avoiding complex multidimensional screening issues — see Asker and Cantillon (2008a) for an analysis of optimal scoring auctions. We also emphasize the importance of ex-post incentives in procurement as in Tirole (1986) and Bajari and Tadelis (2001), although the contractual completeness of the A+B form limits the renegotiation issues emphasized in those papers.

Third, there is a growing empirical literature on auctions with multidimensional attributes. Krasnokutskaya and Seim (2005) and Marion (2007) consider outcomes from mechanisms where the contract is not awarded solely based on price. Athey and Levin (2001) and Bajari, Houghton and Tadelis (2007) analyze multidimensional bidding in timber auctions and highway procurement respectively, emphasizing how the bids determine ex-post behavior. Finally, our paper is related to earlier work on analysis of bidding for highway contracts (see Porter and Zona (1993), Hong and Shum (2002), Bajari and Ye (2003), Jofre-Benet and Pesendorfer (2003), Krasnokutskaya (2004),Li and Zheng (2006), Silva, Dunne, Kankanamge and Kosmopoulou (2008), Einav and Esponda (2008) and Gil and Marion (2009)).

Section 2 presents an overview of the highway procurement process. Sections 3, 4 and 5 respectively contain the theoretical, empirical and counterfactual policy analysis. Section 6 concludes. All proofs and tables are in the appendix.

2 The Highway Procurement Process

Building or repairing a highway is a complex activity. Here we emphasize key features of the process that inform our later modeling choices. Figure 1 gives a simple timeline. First, the state highway department (the "owner") announces the project. Interested bidders may need to pre-qualify in order to bid, either by submitting a "bidding bond" that will be forfeited if

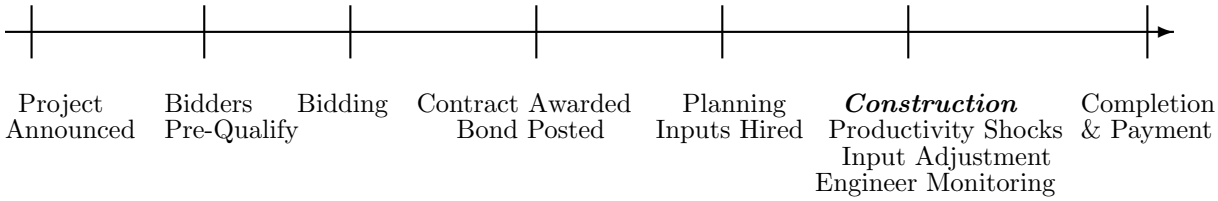

Figure 1: **Timing of Events in Highway Procurement Process**

they win the contract and then don't accept it; or by submitting some form of credentials. In Minnesota, a bidding bond equal to 5% of the bid suffices. Next, the contractors bid on the contract, according to the bidding rules laid out in the announcement. The contract is awarded by law to the lowest scoring bidder, so that contractor reputation plays little role in the process. The winner posts a "contract bond" guaranteeing the completion of the contract according to specifications. This bond is typically secured through a third party whose will take on the contractor's obligations in the event of default.

Once the contract is awarded, the contractor must plan how to structure the various distinct activities, such as excavation or grading, that make up the construction project. To do this, they work out how long each activity will take for a standard crew size, and then use sophisticated software to work out the optimal sequence to complete the activities in using the "critical path method" (Clough, Sears and Sears 2005). The key feature of this technique is that some activities are designated as critical, and must be completed on time to avoid delay, while others are off the critical path and have some time slack. After doing this, the contractor compares the time taken to the time allowed under the contract: if the plan would allow timely completion with a little contingency time on the side, and if there are no time incentives, then he is done. On the other hand, if there are incentives for quicker completion, the contractor must undertake cost-benefit analysis to determine whether it is worth hiring additional capital or labor to speed up the process.

Next, the construction begins. During the process, the project engineer conducts random checks on the quality of the materials and monitors whether everything is completed according to the plan specifications. Productivity shocks may affect the rate at which any activity is completed, and the contractor must continually check progress against the planned time path. If necessary, additional inputs may need to be hired, especially when there is a delay on a critical path activity. At the end of the process, the contractor is paid the amount bid less any damages assessed for late completion.

3 Theory

With the above process in mind, we describe a general framework for the analysis of highway contracts with time incentives. In the model, contractors have multidimensional private costs that depend both on the cost of materials, and on the time to completion. This reflects the increased labor, rental and subcontracting costs of quick completion. These private costs are uncertain due to productivity shocks, so that winning contractors may have to adapt their construction plans after the close of the auction. Bidders compete on both contract price and on the time to completion, and the contract is awarded on the basis of a scoring rule.

Auction Format: n risk-neutral contractors bid on a highway procurement contract. A bid is a pair (b, d^b) indicating the base payment b that the winning contractor will receive, and the number of days $d^b \in [0, d^E]$ that the contractor commits to complete the contract in (the "B-part"). The upper bound d^E is the project engineer's estimate of the maximum time the project should take to complete. The bids are ranked according to the scoring rule $s = s(b, d^b) = b + c_U d^b$, and the contract is awarded to the contractor with the lowest score. The constant $c_U > 0$ in the scoring rule is known as the *user cost*. The contract also specifies ex-post incentives: a per day incentive payment $c_I \geq 0$ to be paid when the winning contractor completes the contract in advance of d^b, and per day disincentive $c_D > 0$ reducing the winning contractor's base payment when the contractor completes the contract later than d^b. We restrict to the case with $c_I \leq c_D$, as is true in all of the contracts we examine. The three parameters (c_U, c_I, c_D) define the incentive structure.

Signals and Payoffs: Losing bidders receive a payoff normalized to zero. The winning contractor has a payoff given by:

$$\pi(b, d^b, d^a; \theta) = b + 1(d^b > d^a)(d^b - d^a)c_I - 1(d^b < d^a)(d^a - d^b)c_D - c(d^a; \theta) \tag{1}$$

where d^a is the actual days taken to complete the contract and $c(d^a; \theta)$ are the input costs incurred. So his profit is just his bid plus incentive payments (if any), less incentive payments (if any), less input costs. The cost function $c(d; \theta)$ depends both on how long the contractor takes and on the realization of a cost parameter $\theta \in \mathbb{R}^k$, which may depend on productivity shocks. It is assumed to be twice continuously differentiable and strictly convex in d for all θ, with $\lim_{d \to 0} c'(d; \theta) = -\infty$ and $\lim_{d \to \infty} c'(d; \theta) > 0$. This captures both the idea that slower completion may lower costs since activities on the critical path need not be accelerated; and

also that delay eventually becomes costly, since there are daily costs associated with having labor and capital employed on site.

For example, suppose that the total project cost is the sum of materials, equipment and labor costs. Equipment is rented at daily rate r. The hourly wage rate $w(h)$ is strictly increasing and convex in the number of hours h worked per day, due to the cost of overtime. Then if θ_M and θ_L are the total materials cost and the total labor-hours required on the project, the contractor will optimally complete the contract at a uniform rate, and we have

$$c(d;\theta) = \theta_M + \theta_L \, w\left(\frac{\theta_L}{d}\right) + rd \tag{2}$$

so that total costs are strictly convex in d, and under mild assumptions on $w(\cdot)$, marginal costs are initially falling in d but then rising.

At the time of bidding, the contractor may be uncertain as to the realization of his private costs conditional on winning the contract. We assume that each contractor i does not observe θ_i, but instead receives only a private signal $x_i \in \mathbb{R}^k$ affiliated with θ_i at the time of bidding. We assume also that the pairs (x_i, θ_i) are *independently* drawn from some continuous joint distribution F_i, which has common compact support $\mathbb{X} \times \Theta$ for all bidders. Thus the auction framework fits into the independent private values (IPV) framework, but with the added complications of a scoring mechanism, ex-post incentives and noisy ex-ante signals.[1]

Equilibrium: A (Bayes-Nash) equilibrium of the game comprises a set of bidding strategies $(\beta_1(x), \cdots, \beta_n(x))$ of the form $\beta_i(x) = (b_i(x), d_i^b(x))$, that are mutual best-responses; and a profit maximizing ex-post completion time strategy $d^a(d^b;\theta)$ that depends on the days bid d^b and the cost realization θ.

Social Welfare and Efficiency: Social welfare is given by $W(d^a;\theta) = V - c(d^a;\theta) - d^a c_S$. It reflects the total social value of the highway project V, less the contractor's private costs, less the social costs imposed on motorists by the construction. The social costs are assumed to be linear in the days taken, with the daily social cost equal to a constant c_S.

We say that a contract design is *ex-post efficient* if the incentive structure is such that the contractor chooses d^a to maximize welfare W for any realization of θ. A contract design is *ex-ante efficient* if the winning bidder is always the bidder who generates the highest

[1] The IPV assumption is important only for the bidding characterization in Proposition 2; all the other results will hold for more general information structures.

expected social welfare $E[W(d^a;\theta)|x_i]$ in equilibrium. These correspond intuitively to productive efficiency (any contractor to whom the job is given maximizes social welfare) and allocative efficiency (the contract is allocated to the contractor who maximizes social welfare in expectation).

Examples: Many specific contract designs have been used by local and state transportation authorities to provide contractors with time incentives. The three most popular such designs are lane rental, incentive/disincentive contracts, and A+B contracting (Herbsman et al. 1995). In addition, most standard highway contracts provide limited time incentives by specifying damages that will be charged if the contract finished late. All of these are special cases of our model.

In *lane rental contracts*, the winning contractor pays the procurer a daily "lane rental" for each day taken on the job. The contract is awarded to the lowest bidder. This corresponds in our model to the case where the bidders are constrained to bid $d^b = 0$, and the disincentive c_D is equal to the lane rental.

In *incentive/disincentive (I/D) contracts*, the contract is awarded solely on the bid amount, and the target number of days d^E is specified by the project engineer, who uses the project plans to come up with a reasonable completion time. Time incentives are provided by daily incentive and disincentive payments for finishing early and late, respectively. This is a special case of our model where the bidders are constrained to bid $d^b = d^E$.

In *standard contracts*, as in I/D contracts, the contract is awarded solely on the bid amount, and the project engineer sets d^E. But here there are no positive incentives ($c_I = 0$), and the disincentives c_D are not generally not project specific, being set out in statewide specifications. We have been told that they are set to cover the costs of having the project engineer out on the site for longer — in other words, they take no account of commuter costs. Again, this corresponds to our model with the constraint that bidders bid $d^b = d^E$. Finally, in *A+B contracts*, contractors bid both an amount of money and a number of days — this is exactly the model we set out above.

Discussion: The model includes a number of non-standard features that we believe are important in thinking about time incentives in procurement. First, it allows for contractors to have multidimensional private costs, reflecting their relative advantages in cheap versus quick project completion. Second, it allows for ex-ante cost uncertainty, which is extremely realistic given that contractors routinely have to adjust their plans and thus their cost expectations

during the contract. In comparing different contract designs, it will be important to think about their efficiency in the face of ex-post contractor adaptation.

Nonetheless, we have not modeled a few features that play a role in real-world contracting. We assume that the project engineer can successfully monitor the construction and ensure that the finished project meets the contract specifications. In this he is aided by the mandatory bonding requirements, which limit incentives to shirk on quality. We also abstract away from the measurement of the number of days taken by the contractor, d^a. In practice, the project engineer determines d^a from the number of work days charged, a measure that takes into account reasonable delays due to unforeseen weather circumstances, necessary work stoppages, and change orders. From talking to project engineers, we know that d^a is to some extent the outcome of negotiation between the contractor and the project engineer as to what is fair and reasonable. We assume that contractors are risk neutral, but most results extend easily to the case of risk aversion. Finally, we do not model the input choices of the contractor at the beginning of the contract, prior to the resolution of uncertainty. Adding this would not be hard, but adds additional complexity without changing any of the results. Later we discuss how this affects the interpretation of our counterfactual results.

Analysis of the Timing Game: Our analysis proceeds by backward induction. First we analyze the optimal timing choice of a contractor who has won the auction and has learned his true cost parameter θ. Next, we consider how contractors will bid in the auction, given that they will proceed optimally in the timing subgame that follows. Our goal is to assess the efficiency of the differing contract designs.

In Figure 2 we depict the "moving parts" of the timing problem. The left panel depicts the incentives under a lane rental contract. Each extra day taken costs the contractor c_D in lane rentals, so his marginal cost of delay is just equal to the rental rate. Overtime and other input costs imply he faces a declining marginal benefit of delay, depicted as the curve $-c'(d;\theta)$ in the figure. Profit maximization requires that he equate the marginal benefits and costs of delay, and so he completes at d^a.

In the right panel, we show a standard contract. There the project engineer has specified the target time d^E, and each day late incurs damages c_D. On the other hand, finishing early gives no bonus, so the marginal cost of delay jumps discontinuously from $c_I = 0$ to $c_D > 0$ at d^E, as depicted in the picture. Again, the optimal completion time d^a occurs where $-c'(d;\theta)$ cuts the step function describing the marginal costs of delay. To formalize this analysis for

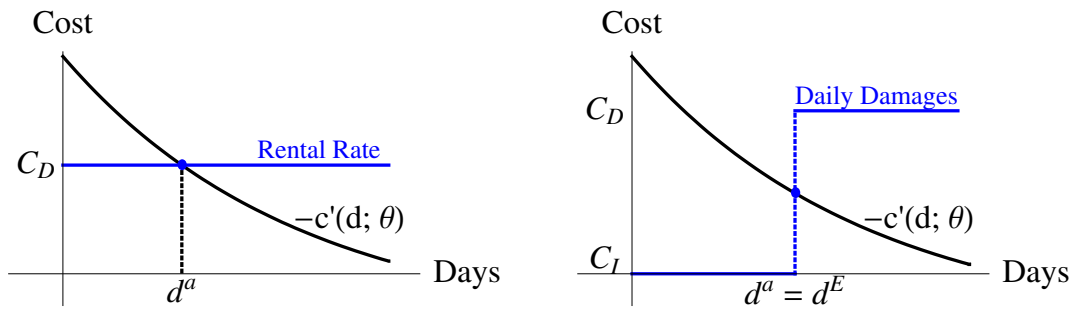

Figure 2: **Completion Time in Lane Rentals and Standard Contracts.** The figure depicts the marginal benefit to delay $-c'(d;\theta)$ curve and the incentive structure for lane rentals and standard contracts. In the left panel, lane rental imposes a constant cost of delay c_D, so the contractor optimally completes at d^a, equating marginal benefit and cost of delay. The right panel depicts a standard contract in which there are no positive incentives, but damages are charged after the specified completion time, d^E. The optimal completion time is $d^a = d^E$.

the general case, it will be useful to define an *incentive relation* $I(d, d^b)$ that specifies the marginal costs of delay to the contractor:

$$I(d, d^b) = \begin{cases} c_I & , \ d < d^b \\ [c_I, c_D] & , \ d = d^B \\ c_D & , \ d > d^b \end{cases} \tag{3}$$

Then the optimal completion time will just be where the marginal benefit of delay curve cuts the incentive relation. Also, as the number of days bid d^b increases, the discontinuous jump from c_I to c_D (if any) shifts right, and so the optimal completion time will increase for some realizations of θ. We formalize this in a lemma:

Lemma 1 (Optimal Completion Time) *The optimal completion time $d^a(d^b;\theta)$ is a function, satisfying $-c'(d^a;\theta) \in I(d^a, d^b)$ $\forall d^b$ $\forall \theta$. It is (weakly) increasing in d^b if $c_I < c_D$ and constant in d^b if $c_I = c_D$, for any realization of θ.*

Bidding on Time: Next, we consider how the contractor should choose the number of days to bid d^b in an A+B contract.[2] For any realization of θ, her incentive payments $IP\left(d^a(d^b;\theta), d^b, \theta\right)$ depend entirely on how many days she bid.[3] For example, if she bids a

[2]The contractor has no choice in the other contract designs.

[3]The incentive payments are $IP\left(d^a(d^b;\theta), d^b, \theta\right) = 1(d^a > d^b)(d^a - d^b)c_D - 1(d^a < d^b)(d^b - d^a)c_I$.

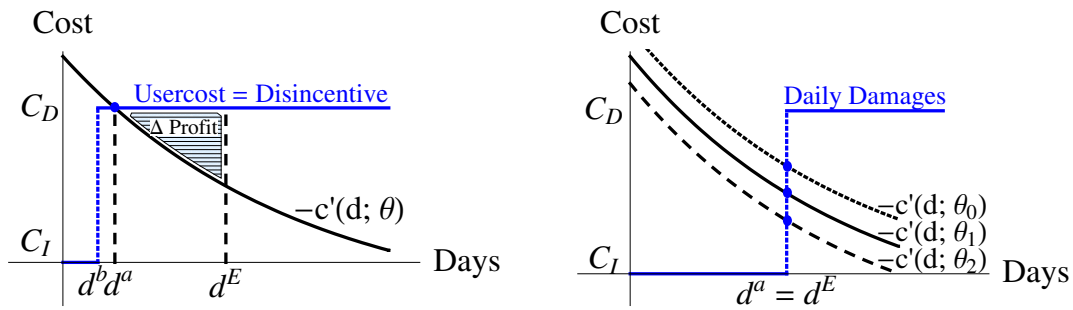

Figure 3: **Bidding and Stickiness in Completion Time** In the left panel, we depict the gains to bidding $d^b < d^E$ in a A + B contract where $c_I = 0 < c_U = c_D$, for a particular realization of θ. The additional profits from this strategy are shown in the shaded region. In the right panel, we compare the optimal completion times in standard contracts for different realizations of θ. In all cases, $d^a = d^E$, indicating that completion times will be "sticky" at d^E in standard contracts.

very low number of days, she will often optimally finish late and pay out damages at rate c_D per day, with incentive payments of $-(d^a(d^b;\theta) - d^b)c_D$. The upside is that she is more competitive during the bidding stage. To analyze this trade-off, fix a desired total score $s = b + c_U d^b$. Substituting out for b in the payoff function (1), it follows that the optimal days bid is:

$$d_i^b(x) \in \arg\min_d \; E\left[c(d^a;\theta) - IP\left(d^a(d;\theta), d, \theta\right) + c_U d | x\right] \tag{4}$$

To see these trade-offs graphically, look at the left panel of Figure 3. In this example, the contractor faces zero positive incentives ($c_I = 0$) and equal user costs and disincentives ($c_U = c_D$). Then it is better for her to bid a low number of days $d^b < d^E$ for any fixed s, even if this results in damages. This allows her to increase the A-part of her bid by $(d^E - d^b)c_U$, holding s fixed. On the other hand, she pays out $(d^a - d^b)c_D = (d^a - d^b)c_U$ in damages, and also incurs additional private costs $\int_{d^a}^{d^E} -c'(s;\theta)ds = c(d^a;\theta) - c(d^E;\theta)$. This leaves the shaded region as increased profits.

For certain incentive structures, it *always* pays to bid a "very low" or "very high" number of days. To formalize this, it will be helpful to define a lower and upper bound on the completion times. Define $\underline{d} = \inf\{d : \exists\theta : -c'(d;\theta) = c_D\}$ and similarly $\overline{d} = \sup\{d : \exists\theta : -c'(d;\theta) = c_I\}$.[4] A contractor will never want to complete before $\underline{d}$ because his marginal private costs always exceed $c_D > c_I$ before that point. Similar logic shows he will never complete after $\overline{d}$. Then we have the following result:

[4] If $c_I = 0$, the upper bound may not exist, but then the case in proposition 2 below won't arise.

Lemma 2 (Bidding on Time) *In equilibrium, if $c_I < c_D \leq c_U$, the contractor bids $d^b \leq \underline{d}$. If $c_U \leq c_I < c_D$, the contractor bids $d^b \geq \bar{d}$. If $c_I = c_U = c_D$ the contractor is indifferent among all choices of d^b. Finally, if $c_I < c_U < c_D$, the days bid is $d_b = d_i^b(x) \in (\underline{d}, \overline{d})$, decreasing in c_U, increasing in c_I and c_D.*

This is intuitive. If $c_I = c_U = c_D$, incentives are completely flat regardless of the days bid, so the contractor is indifferent. On the other hand if positive incentives are low, as in the left panel of Figure 3, it is profitable to bid as few days as you could reasonably expect to complete the contract in, to "lock-in" a bonus for finishing early. Finally, if the usercost is in-between the positive and negative incentives, the choice of days to bid is no longer straightforward, as no one choice is best for all cost realizations θ. Instead the contractor minimizes the expected cost as in (4), including the "opportunity cost" $c_U d^b$ of a forgone higher A-part bid in this calculation.

Ex-Post Efficiency: Now that we know how contractors bid their days, and how they complete given the incentive structure induced by that choice, we are in a position to characterize the conditions required for ex-post efficiency of the various contract designs. The simplest case is a lane rental. The social cost of delay is c_S, so if the procurer sets $c_D = c_S$, the contractor will internalize the social costs and complete in the socially efficient time.

Only slightly harder is the case of the standard contract. Consider the right panel of Figure 3, which shows the marginal benefit of delay for different realizations of θ. As you can see, given the discontinuity in the incentive relation at d^E, the contractor will end up completing the contract in d^E days in all cases. But this immediately proves that the standard contract is not ex-post efficient, since efficiency would require different completion times for the different cost realizations. In fact, standard contracts will in general be inefficient unless the project engineer imposes such an unrealistic time d^E that contractors will always complete early or late. To rule out this case, we make an assumption:

Assumption 1 (Efficiency may require adaptation) *There exist $\theta_L, \theta_H \in \Theta$ such that $-c'(d^E; \theta_L) < c_S$ and $-c'(d^E; \theta_H) > c_S$.*

All this says is that there are some favorable cost realizations for which social efficiency requires completion before d^E, and some bad realizations which for which the socially optimal completion time is late. Under this assumption, we can provide a general characterization of ex-post efficient contract design:

Proposition 1 (Ex-Post Efficiency) *If assumption 1 holds, in equilibrium:*

(a) Lane rentals are ex-post efficient iff $c_D = c_S$.

(b) I/D contracts are ex-post efficient iff $c_I = c_D = c_S$.

(c) Standard contracts are not ex-post efficient.

(d) A+B contracts are ex-post efficient iff either $c_I = c_U = c_S \leq c_D$ or $c_I \leq c_U = c_S = c_D$.

The first three results follow the logic above. For A+B contracts, the contract designer has the luxury of picking one of the incentives "incorrectly", either giving low positive incentives, or overly high negative incentives. The reason for this is that the contractor can correct the incentive structure herself by bidding the right number of days, and indeed this is what happens in equilibrium. This may be of practical importance, since highway construction officials are often reluctant to offer positive incentives for budgetary reasons.[5] Ex-post efficient outcomes can be achieved without positive incentives by either lane rentals or A+B bidding, if appropriately designed.

Analysis of Bidding: To complete the analysis, we consider how a contractor should bid in the auction. Define the *pseudo-cost* $P_i(x)$ of a contractor with private signal x as their expectation of the sum of their private costs, their incentive payments and the B-part of the score, given that they complete in the optimal time as in proposition 1 and bid the number of days optimally as in proposition 2. Formally, it is:

$$P_i(x) = E\left[c(d^a;\theta) + IP(d^a, d_i^b(x);\theta) + c_U d_i^b(x)|x\right] \tag{5}$$

This is analogous to the pseudo-value of Asker and Cantillon (2008b), and indeed from Theorem 1 of their paper, all equilibrium outcomes of the full scoring auction can be obtained by looking at equilibrium outcomes of the standard lowest-bid auction where contractors have costs equal to their pseudo-costs. With that insight, the following result comes easily:

Proposition 2 (Bidding and Ex-Ante Efficiency) *In any equilibrium:*

(a) The strategies take the form $\beta_i(x) = (s_i(x) - c_U d_i^b(x), d_i^b(x))$, where $s_i(x)$ satisfies the first-order condition

$$s_i(x) = P_i(x) + \frac{1}{\sum_{j\neq i} h_j(s)}$$

[5] This thinking is in some sense wrongheaded, since the procurer will have to pay for any time incentives offered, regardless of how they are structured; but it may come out of a different budget.

where $h_j(s)$ is the hazard function of the distribution of scores submitted by bidder j, and $P_i(x)$ is bidder i's pseudo-cost.

(b) If the equilibrium is symmetric and the contract is ex-post efficient, it is also ex-ante efficient. Standard contracts will be ex-ante efficient in symmetric equilibria if $c_D = c_S$.

The intuition for the first part is that the single-dimensional pseudo-cost completely summarizes the total expected costs of a winning contractor, and that this is just a standard first price auction with the pseudo-cost replacing the cost. Simple algebra then yields the first order condition. For the ex-ante efficiency results, notice that if the equilibrium is symmetric, the bidder with the lowest pseudo-cost will win the auction. If the contract design is ex-post efficient, this is the welfare maximizing allocation, since the pseudo-costs correspond closely to the welfare function. On the other hand, standard contracts are not ex-post efficient, but if the daily damages are set equal to the daily social cost, the scoring rule at least weights time appropriately, so that the contract is allocated to the right bidder.

4 Empirical Analysis

The theory outlined above indicates how contracts should be designed in order to maximize social welfare. In the remainder of the paper, we analyze data from contracts let by the Minnesota Department of Transportation (Mn/DOT). Our dataset is unusually rich, as it includes detailed data on both standard contracts and the newer A+B contracts. This enables us to examine two interesting questions. First, we present some descriptive evidence on the bidding and outcomes in both standard and A+B contracts, showing that the predictions of the theory are largely confirmed, although there are some interesting deviations that may motivate a particular design. Second, we do some policy analysis based on counterfactual welfare simulations, showing that large welfare gains can be achieved for a reasonable increase in the Mn/DOT budget.

4.1 The Data

We have data from Mn/DOT on all the highway procurement contracts completed in Minnesota during the period 1997-2007. To get this dataset, we merged data from three main

sources. The first is the publicly available bidding abstracts, which detail who is letting the contract (which level of government), what the contract is for, who bid and what amount they bid, as well as the project engineer's initial cost estimate. This is formed by the Mn/DOT project engineer based on the specified quantities for the project materials and "blue book" prices for the various contract items. The cost estimate explains a large fraction of the variation in the observed bids.[6]

The second source of data is data on the contract completion time. For a subset of completed standard contracts (248 of them), Mn/DOT gave us the actual diary data used by Mn/DOT project engineers to record contract progress. Using their own software, a program called FieldOps, we measured both the total days actually taken on the contract and the days the project engineer had allowed — thus deducing whether the contract was completed early or late. Most contracts are "working day" contracts, in which it is the responsibility of the project engineer to count the number of "working days" used by the contractor, to assess whether he is late and decide whether to charge (liquidated) damages. This allows the project engineer to make allowances for unforeseeable construction delays. Given this, the contract days in the bidding abstract don't match observed outcomes all that well, and the diary data is crucial for the analysis.

Our final data source is A+B contract data, which we got from Mn/DOT's office of construction and innovative contracting. This contract design has typically only been used on highly time-sensitive projects, and so there are only 29 such contracts, with 123 bids placed in the auctions. Here we we observe all the bids (both the A-part and the B-part), the identity of the bidders, the scoring rule and the incentives c_I and disincentives c_D. In addition there are some other features not captured in the formal model, such as a minimum number of days to bid in some contracts; and capped positive incentives in others. For 23 of the A+B contracts we also observe the final completion time (the remaining contracts are still outstanding).

We present summary statistics on the standard contracts in Table 1, focusing on the contracts for which we have diary data, since these are extensively used in the analysis. A typical contract has value just over \$1 million. The winning bid is 94.4%, on average of the engineer's estimate, though the average bid is higher at 106.7% of the engineer's estimate. It is clear that the diary subsample is not representative of the full set of Mn/DOT contracts, which are on average twice as big at \$2 million. We will try correct for this in the policy analysis.

[6] A linear regression of all of the bids in our data on the cost estimate alone has an R^2 of 0.97.

These contracts are of relatively short duration, on average 34 days. Contracts are generally completed on time, although in the event that they are completed late, damages are assessed in only 29% of cases. This is because the project engineer has discretion over when to assess penalties. Damages range from nothing to as high as 1.05% of contract value. Time penalties or damages are specified in the standard Minnesota contract specifications, as detailed in Table 2. For example, in a typical \$1M contract, the penalty for being a day late is \$1000. Also in Table 2, we present detailed summary statistics on completion time. Smaller contracts are more likely to finish early or on time than larger contracts. In fact, none of the contracts of size less than \$50 000 finish late, perhaps reflecting the fact that the penalties in these contracts are a larger fraction of the total contract value.

Some summary statistics on A+B auctions are presented in Table 3. These contracts are much larger, of average value \$11 million and estimated duration 151 days. Notice that these contracts do indeed get completed earlier than specified: on average the winner bids only 123 days, and typically finishes early, earning average incentive payments of \$41600 per contract. There are four basic kinds of incentive structure that are observed in the A+B data: (i) those with no positive incentives, and equal user cost and disincentive ($0 = c_I < c_U = c_D$); (ii) those with small positive incentives ($0 < c_I < c_U = c_D$); (iii) those with entirely equal incentives and disincentives ($0 < c_I < c_U = c_D$); and (iv) those with user cost higher than the other incentives ($c_U > c_I = c_D > 0$). We will see how these different incentives affect bidder behavior in what follows.

4.2 Bidding Behavior

To start, we look at the days bid in A+B contracts, as shown in the left panel of Figure 4. The theory predicts that bidders should bid fewer days when positive incentives are lower than disincentives, so as to "lock-in" the gains to finishing early. In fact it makes the strong prediction that with $c_I < c_U = c_D$, the optimal bid is $d^b \leq \underline{d}$, the lowest days they could possible want to complete in. In the data, we see that weaker incentives are correlated with lower days bid, as predicted. Also, in the one contract in which usercosts were set above the disincentive, many contractors actually phoned Mn/DOT and enquired if they could bid zero days as is optimal. They were told that they could not, as their bids could be rejected as "irregular", but we see that they bid a low number of days (look at the triangles).

To confirm that there is a statistically significant difference in bidding behavior, we run a

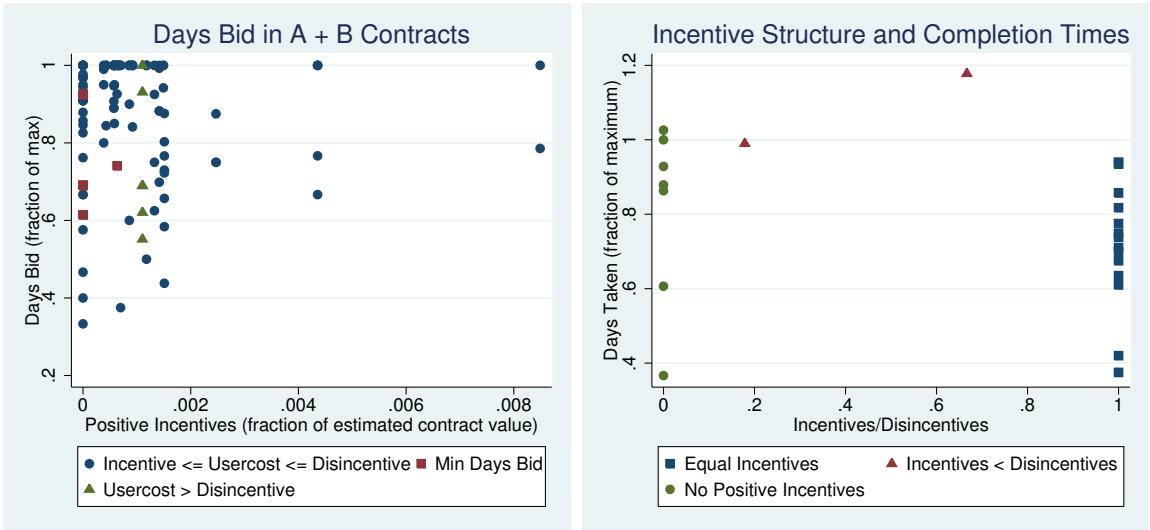

Figure 4: **Bidding and Competion in A+B Contracts.** This left panel shows how the number of days bid varies with the incentive structure. The right panel depicts how the final days taken relates to the days bid, for varying incentive structures.

tobit regression of days bid (as a percentage of the maximum allowed) on dummies for the different incentive structures.[7] A tobit is used to account for the censoring of days bid both below (at the minimum days) and above (at the maximum days). The outcome is shown in the first column of Table 4. We find that relative to the case with equal incentives (the omitted group), the days bid is around 15% lower with zero incentives or a higher usercost. The dummy for small positive incentives has no statistically significant effect.

Next, we look the days taken relative to days bid (right panel of Figure 4). The theory predicts that with zero or small positive incentives, the contractors should almost always finish late (since they bid $d^b \leq \underline{d}$). In fact with small incentives, the actual completion time tends to be within 10% of the days bid, and often early. This would make sense if the incentive structure were given ex-ante, as in the standard contracts, since the wedge between incentives and disincentives should produce "stickiness" around d^b. But in A+B contracts, the contractor should have bid such a low d^b such that finishing late is almost inevitable. This may indicate that the bidders are not bidding days as the theory predicts. By contrast, the theory has nothing to say about the case with equal incentives, but the data shows that contractors bid up the days and finish early with certainty, thus earning incentive payments. The fully theoretically equivalent strategy of bidding low days, increasing the A-bid, finishing late and paying out damages is not chosen.

[7] We treat the incentive structure as exogenous, given the limited data.

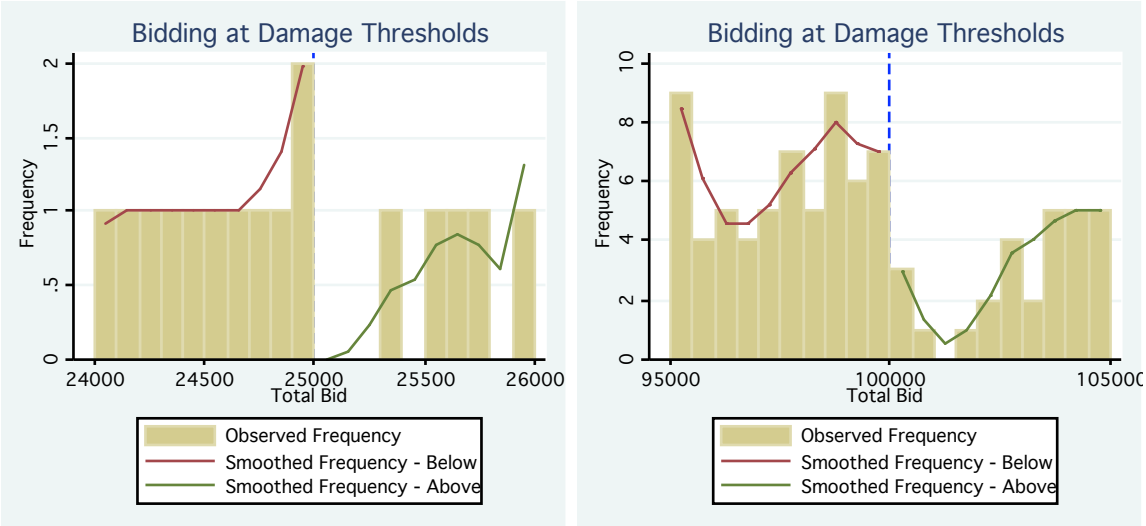

Figure 5: **Bid Manipulation in Standard Contracts.** The figure shows the relative frequency of bids around the thresholds at which liquidated damages increase in standard contracts. In the left panel we show bids around $25000; in the right panel, we have bids around $100 000.

Time incentives in standard contracts are weaker, but they still appear to play a role in bidding. The damage specifications depend on the final value of the contract, given by the winning bid. Damages increase across pre-defined thresholds. For example, one threshold is at $25000 so a contract worth $24900 attracts damages of $150 per day; while a contract worth $25100 attracts damages of $300 a day. In theory this gives incentives to contractors to bid just below these thresholds, so that in the event they win, their potential damages are lower. Figure 5 shows does actually occur. Both panels show histograms of the frequency of bids on either side of a threshold ($25 000 in the left panel, $100 000 in the right), as well as local linear regressions smoothing the frequencies on either side. It is clear that bids immediately to the left of the threshold are far more frequent than those to the right. The gap is particularly clear for low contract values; as contract values increase, the gap tends to get smaller. The joint hypothesis that the density is continuous across the thresholds is formally rejected, which suggests that contractors do pay attention to the time incentives.[8]

To conclude the bidding analysis, we look at the actual dollar bid across standard and A+B

[8]The formal test is given in McCrary (2008), and has a test statistic based on the difference between the smoothed frequencies at the threshold itself. As is typical of nonparametric tests, the result is sensitive to bandwidth specification, and so we focus more on graphs themselves.

contracts. In column 2 of Table 4 we show the results of an OLS regression of markup over the engineer's estimate (i.e. bid / engineer's estimate) on the log engineer estimate, disincentives and incentives (both as a fraction of estimate) plus dummies for the main task on the contract and the district it was in. The results show that markups over engineer's estimate significantly increase with the disincentives, and decrease with the incentives. This is exactly what the theory would predict.

4.3 Completion Times and Incentives

Now we consider how the incentive structure impacts the final completion time. Recall that our theory predicts that contract completion time will be "sticky" around d^E in standard contracts, so that many contracts will be completed exactly on time. Look at the left panel of Figure 6, which is a histogram of the days late across contracts. In the figure, a contract exactly on time has been added to the bin to the left of 0, and so you can see that over 20% of the contracts were completed either just on time or a day early. This is *exactly* what the theory predicts: at exactly the time when the penalties kick in, contractors choose to finish their work. This is a strong validation of the theory.

We end with a basic question: do high powered time incentives yield quicker contract completion? The right panel of Figure 6 sheds light on this. A+B contracts are generally completed much faster than the engineer's estimate, and where both large positive incentives and disincentives are given, the contractor typically earns the maximum possible bonus ("capped incentives"). There does not appear to be an obvious difference between the contracts with no positive incentives and those where positive incentives were given. By contrast, standard contracts are often late.

To look at this more formally, we ran an OLS regression of days taken (as a fraction of the maximum days) on the relative disincentives and positive incentives. Recall that for A+B contracts, disincentives are almost always equal to user costs, so the disincentives are a measure of user costs for those contracts.[9] The results are reported in column 3 of Table 4. We see that the stronger the time incentives, in the form of high usercosts/disincentives, the quicker the contract is completed. By contrast, the presence of positive incentives has no statistically significant effect on completion time, which is as the theory predicts — the actual

[9]There is only one case in which the two are not equal, and that contract is removed for the purposes of this regression.

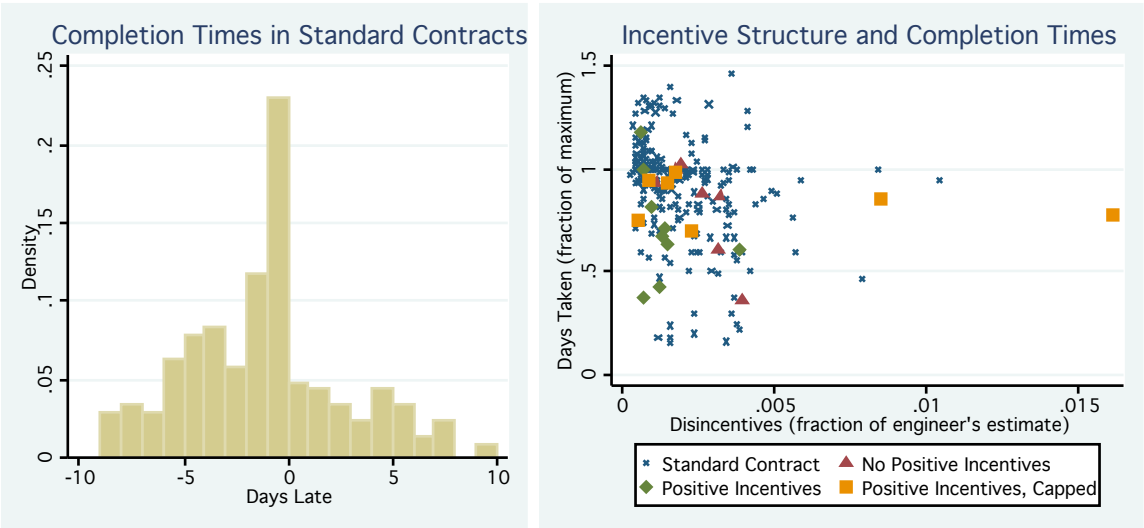

Figure 6: **Completion Times and Incentives.** The histogram in the left panel is of the difference between actual and contractually specified completion time in standard contracts. The huge spike just before 0 indicates many contracts are completed exactly on time, as the theory predicts. The right panel shows how the completion time varies across standard and A+B contracts.

days taken should be constant in expectation across incentive structures with $c_I \leq c_U = c_D$. This regression is meant to be primarily descriptive, as the OLS approach does not account for the clustering of contracts at 0 seen in Figure 6. We will take account of these econometric issues in the structural model below. But it does suggest that higher powered time incentives do speed up contract completion.

Overall, the predictions of the theory stand up well. An important deviation from the theory is that bidders on A+B contracts do not bid an unreasonably low number of days when there are no positive incentives. This is possibly because they do not understand the game, or more likely because they are concerned that their bid will be rejected as irregular. This implies that the incentive structure with zero positive incentives will generally not deliver ex-post efficiency, even if the user cost is set equal to true social costs, because contractors will not respond to positive cost shocks by completing early. On the other hand, with equal incentives, the mechanism is "strategy-proof" in an intuitive way: the number of days bid has no important implications for ex-post behavior. This suggests that designing A+B contracts to provide equal positive and negative incentives contracts may be desirable in practice.

5 Policy Analysis

The central contention of the paper is that building time incentives into highway procurement contracts could yield large social welfare gains. To assess this, we need to evaluate the counterfactual welfare of both producers and commuters when contracts are re-designed to incorporate explicit time incentives. Our strategy for doing this is as follows. First, we estimate the contractor's private costs by looking at how their behavior changes as damages vary across contracts. To do this, we use first order conditions motivated directly by our theoretical model. Second, we build measures of commuter user costs, using a methodology very similar to that used by the engineers at Mn/DOT. With these in hand, we consider simple counterfactual policy changes in we introduce lane rentals, which as we saw earlier have a relatively simple efficient design. Of particular interest is a realistic case in which the lane rental is a constant fraction of the user-cost, which turns out to be constrained efficient when the owner (in this case Mn/DOT) faces budget constraints.

5.1 Estimating Contractor Costs

We estimate the contractor's private costs by looking at their behavior as damages vary in standard contracts. We focus on standard contracts because damages are exogenously specified at the state level, whereas in A+B contracts there is a concern that the incentive structure is endogenous. Since our diary data subsample consists of only 248 contracts, it is not possible for us to use a completely flexible, nonparametric approach when analyzing the impact of penalties on project completion times. We will use a parametric model in our estimation.

The idea behind our estimation approach is to use the first order conditions from the theory model developed in the earlier sections. This will allow us to account explicitly for the discontinuity in incentives as the contractor moves from "early" to "late", which produce a mass of on time completions, as seen in Figure 6.[10] The first order conditions must satisfy:

$$-c'(d;\theta) = 0 \text{ if } d < d^E \tag{6}$$

$$-c'(d;\theta) = c_D \text{ if } d > d^E \tag{7}$$

$$0 < -c'(d;\theta) < c_D \text{ if } d = d^E \tag{8}$$

[10] A similar econometric approach was taken in Reiss and White (2005).

These are easily interpreted as saying that firms only complete early when their marginal benefits to delay reach zero; only complete late when their marginal benefits to delay equal the cost of delay c_D, and otherwise complete on time.[11]

In specifying the marginal benefit of delay function for any given contract, we need to account for contract heterogeneity. A natural way to homogenize the contracts is to assume that the marginal benefits of delay scale with the project engineer's estimate of contract duration. This seems reasonable, as the project engineer's estimate is based on extensive analysis of the project specifications, and explains over 90% of the variation in days taken in our data. Formally, suppose that the contractor's cost function is homogenous of degree 1 (i.e. a contract of twice the size, completed in twice the days, costs twice as much). It immediately follows that the marginal benefit of delay will be homogenous of degree 0.[12] Then if the cost shocks scale with contract duration, we have for an arbitrary contract:

$$-c'(d;\theta) = -c'(\frac{d}{d^E};\frac{\theta}{d^E}) = -c'(\widetilde{d};\widetilde{\theta}) = \alpha_0 + \alpha_1\widetilde{d} + \widetilde{\theta} \tag{9}$$

where $\widetilde{d} = d/d^E$; and $\widetilde{\theta} = \theta/d^E$ is assumed to have normal distribution with mean zero and standard deviation σ^2. This a very simple and parsimonious specification — linear marginal benefit and normal errors — but it turns out to fit the data quite well, and is appropriate given the limited number of observations in our sample. We have experimented with other specifications, but they have not contributed to fit.

We proceed by maximum likelihood. The log likelihood for each observation, $\ell_i(\widetilde{\theta})$ is:

$$\ell_i(\widetilde{\theta}) = \begin{cases} \log\left(\phi\left(\frac{-\alpha_0-\alpha_1\widetilde{d}}{\sigma}\right)\right) - \log(\sigma) & , \widetilde{d} < 1 \\ \log\left(\Phi\left(\frac{(c_D/d^E)-\alpha_0-\alpha_1\widetilde{d}}{\sigma}\right) - \Phi\left(\frac{-\alpha_0-\alpha_1\widetilde{d}}{\sigma}\right)\right) & , \widetilde{d} = 1 \\ \log\left(\phi\left(\frac{(c_D/d^E)-\alpha_0-\alpha_1\widetilde{d}}{\sigma}\right)\right) - \log(\sigma) & , \widetilde{d} > 1 \end{cases} \tag{10}$$

where ϕ and Φ denote the standard normal pdf and cdf respectively. The likelihood function is very similar to that of a censored tobit, except that instead of having a mass point on one side at the point of censoring, we have a mass point at on-time completion, $\widetilde{d} = 1$.

We allow the error standard deviation σ to vary with the contract size, specifying that σ is

[11] We assume that firms believe the damages will be charged with certainty. If in fact they believe they will only be imposed some fraction of the time, this methodology will overestimate the private costs, and is thus conservative in assessing welfare gains to improved contract design.

[12] See for example Mas-Colell, Whinston and Green (1995), Theorem M.B.1.

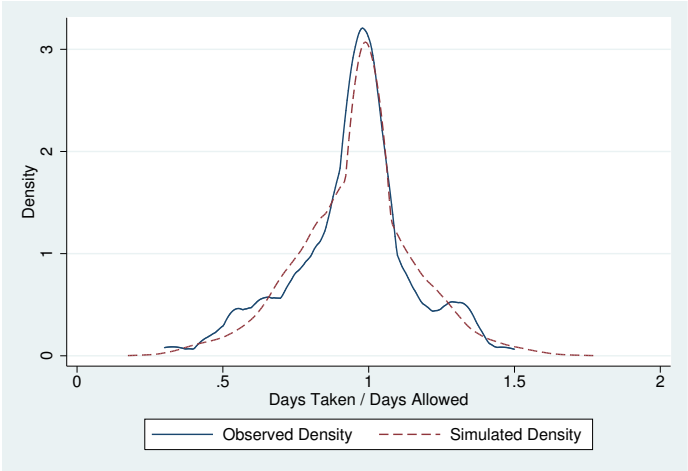

Figure 7: **Fit of the Timing Model** The figure shows kernel density estimates of the normalized completion times $d/\bar{d}$ for both the actual and data simulated from the structural model.

a linear function of the log contract size. This is motivated by the observation that in the data, even after normalizing through by the engineer's days estimate, smaller contracts have more variable outcomes. We have also experimented with more complicated specifications, including additional covariates, but they are generally insignificant. We trim the top and bottom 5% of observations, which amounts to contracts where the contract was completed in less than 30% of the allotted time, or more than 153% of the allotted time. This allows for a much better fit. To obtain standard errors, we bootstrap the ML procedure.

Our results are in Table 5. For an average contract, which lasts 37.4 days, we estimate the marginal benefit of delay to be $18895 - 510d$, a line that reaches zero just before the allotted completion time.[13] In other words, we predict absent any negative shocks, contractors will complete slightly early. This is both consistent with the data and the fact that contractors typically build a little "contingency" time into their project plans. We find that the marginal benefit of delay is initially positive, but decreasing in days, as assumed in the theory. We find also that larger contracts have smaller cost shocks, relative to the number of days assigned to the contract. To assess fit, it is useful to compare some sample and predicted moments. In the data, 29.5% of the contracts finish late, and conditional on being late, they finish 15.7% over the days allotted. The model matches this data well, predicting that 34% of contracts will finish late, and that they will finish 15.7% late, as in the data. This good fit is also evident from comparing the kernel density plots of the normalized completion times against those simulated from the structural model, as in Figure 7.

[13]Due to the trimming, the average contract size is slightly higher than that reported in Table 1.

5.2 Estimating User Costs

We next construct measures of road user costs for the roads under construction in 167 of the contracts.[14] This is necessary because Mn/DOT generally only calculates user costs for A+B projects, or where they have been asked to evaluate the costs of a delay in construction. We develop an estimate of the daily user cost for those contracts as follows:

$$User\ Cost_t = Delay_t \times Time\ Value_t \times Traffic_t \tag{11}$$

The daily user cost is estimated as the per user delay (in hours), times the time value for the average commuter (in $/hour), times the average daily traffic on that road. In this we follow closely the actual Mn/DOT methodology for computing user costs. Their methodology accounts also for additional wear and tear in the case where commuters are re-routed, and for the possibility of cars with multiple occupancy. By neglecting these factors, we hope to get conservative user cost estimates.

The first element that we need to calculate is the average delay due to the construction project. In practice, this depends on whether the project engineer decides to close down lanes but still leave the road open, or to close the road and detour commuters. Whenever the road is left open, commuters must slow down, and Mn/DOT generally assumes that their commuting speed over that section is cut in half. On the other hand, if rerouted, the detour is generally longer than the original route, and that also causes delays. The decision of whether to reroute or close lanes is spelled out in the project plans, and so Mn/DOT uses this information when calculating user costs. Since we do not have access to the information, we take a conservative stance, computing the minimum of these two alternatives.

To do this, for each contract where the location data allows us to pinpoint it on a map, we use google maps to outline the construction zone, and its length (see the left panel of Figure 8). We then use google to calculate a travel time for that route. If the route is left open, we assume that the travel time will double, which means the delay will be equal to the original travel time. We also do our best to construct a likely detour around that section of the road, as in the right panel of Figure 8. We then use google to estimate the time required to drive the detour, getting a delay estimate as the difference between the travel times. In

[14] We have gone to considerable effort to get this data: for the remaining contracts, Mn/DOT themselves no longer appear to have easily available copies of the original plans indicating the location of the construction, which is necessary to compute the user costs.

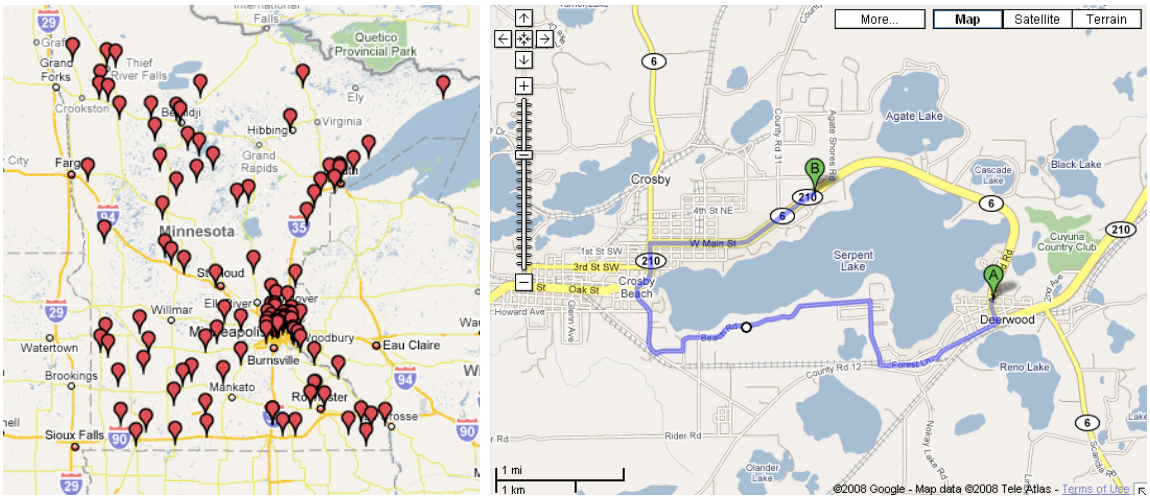

Figure 8: **Delay Calculations:** The left panel shows the locations of the highway construction projects used in the counterfactual. The right panel shows an example of a detour calculation around a section of route 6.

practice, Mn/DOT does a less-virtual version of this exercise, actually sending personnel to drive the detour and original route at different times of day and record the data. Overall then, delays are calculated as:

$$Delay_t = \text{Min}\{Travel\ time_t, Detour\ Time_t - Travel\ time_t\}$$

To get an estimate of the daily traffic, we use traffic volume measures from Mn/DOT at each location in the dataset.[15] For a measure of the time value, we use the rate of \$12/hour, which is the rate used in the Mn/DOT calculations. We throw out one outlier with a usercost estimate of over \$600 000 per day. We summarize all these measures, including the usercost we compute, in Table 6. The estimated average delay is 9.11 minutes, which with around 18550 daily commuters on a typical road results in an average daily usercost of \$15626.

5.3 Counterfactuals

Now that we have contractor costs and social benefits, we consider three counterfactual changes. First, a change in which Mn/DOT charges a lane rental equal to the current damages specified in standard contracts (see Table 2). This affects all contracts of a given size equally. Next, we consider two policies that are more contract specific: one in which

[15] See http://www.dot.state.mn.us/traffic/data/html/volumes.html

the Mn/DOT assesses a lane rental equal to 10% of the estimated usercost, and another in which the lane rental is equal to the estimated usercost. The former policy is constrained efficient, in that it is the best policy subject to the requirement that the required increase in the Mn/DOT budget is less than a given amount (in this case \$17573 a contract or at 370 contracts a year, \$6.5 million a year).[16] The latter policy is the efficient contract design.[17]

The procedure for obtaining the counterfactual estimates is as follows. First, we calculate the mean counterfactual completion times for each contract by simulating draws from the cost shock distribution, and then computing the mean optimal completion time under the new incentives. Then we calculate the additional private costs incurred by the contractor — since he will now complete earlier — under the estimated parameters of the time cost model. This by itself is enough for welfare analysis, since we can compare the social gains from the average reduction in user costs to the private cost increase faced by the contractor.

But it will also be interesting to estimate the cost to Mn/DOT of implementing this policy. To get a sense of how costly actually implementing this policy would be, we compute a net cost to Mn/DOT as:

$$\text{Cost Mn/DOT}_t = (\Delta\text{Private Costs}_t + \Delta\text{Damages}_t) \times 1.2 - \Delta\text{Damages}_t \tag{12}$$

where the multiplier 1.2 is a bidding markup, chosen to be reasonably conservative (compare it, for example, with the markup of the average bidder over engineer's cost of 1.07). This formula accounts for the fact that bidders will mark up both their expected damages and higher private costs, so that Mn/DOT will be a net loser on the lane rentals it collects.

One concern is that the subsample of contracts for which we have user cost data is not representative of the contracts undertaken in the state of Minnesota as a whole. To address this, we use propensity score re-weighting. First, we estimate the probability that a contract is in the subsample given key observables like contract size, district, government authority, type of contract (e.g. bridge-building) and primary material used (e.g. asphalt, concrete).

[16] Maximizing the social welfare function subject to a budget constraint — payments to contractors less rentals less than some B — yields a first order condition under which the ratio of usercosts across two projects must be equal to the ratio of the expected marginal benefits of contractor delay. Under our specifications, contractor optimization leads to the the RHS of this expression being equal to the ratio of the per day rental rates, implying that the constrained efficient policy is to set the rental rates as a constant fraction of the usercosts across all contracts, where the fraction is determined by the budget.

[17] For bigger changes, we are leaning more heavily on our parametric specification to recover estimates of the contractor's private costs, and so the results from this last case should be viewed with some caution.

This is done using a probit on the full sample of contracts in Minnesota (3702 contracts). Second, we weight the relevant counterfactual statistics for each observation in the subsample by the inverse of the estimated probability of inclusion. If the unobservables have the same distribution across the full sample and subsample, this will correct for selection.

The results are in Tables 7 and 8. In the first table, we see that introducing lane rentals at current damages reduces the average number of days taken from 33.9 to 32.5; while a lane rental of 10% of usercost reduces the mean completion time to 31.8 days. The efficient design has a strong effect, creating incentives for contractors to complete in only 18.9 days. The estimated private time reduction costs under the status quo are small, on average $499 with an average contract size of around $1 million, suggesting the current time incentives are weak. Under the first counterfactual policy, these are estimated to rise to $1712.8, and all the way to $138543.9 under the efficient lane rentals. In all counterfactuals, the damages increase dramatically, because the contractor must start paying damages from the moment the contract starts. The re-weighted results are remarkably similar, which is re-assuring since it suggests that the counterfactual subsample is reasonably similar to the full sample.

What are the changes in producer costs and user welfare? The magnitudes are striking! From the re-weighted results, the smaller lane rental policy improves user welfare by $26676 per contract, at a cost to producers of only $1217, for a net welfare gain of $25459 per contract. The more aggressive policies have even larger effects, since they tie incentives directly to the cost to commuters. The 10% usercost lane rental improves net welfare by $163808 per contract, and the efficient policy yields net gains of $388653 per contract. This is around 19.4% of contract value, or in terms of the Mn/DOT budget around $310M a year.[18]

The problem is of course that the fully efficient policy would be extremely costly to Mn/DOT, raising the issue of whether taxpayers would be willing to foot the bill. We estimate the cost at $199000 per contract, or about 10% more. Instead, the constrained efficient policy of setting lane rentals at 10% of usercost may be a good practical policy proposal. The expected funding increase there is only $17573 per contract, or less than 1%. Compare this with the expected gains to users of around 8.5% per contract. The huge gap between the costs and benefits comes from matching incentives to usercosts. Providing high-powered incentives on high usercost projects really speeds them up, sometimes at little cost to the contractor, as in the case when he draws favorable cost shocks. The current incentives in standard contracts

[18]The Mn/DOT budget for 2009 is $1.69B. Source: Minnesota Transportation Appropriations for 2008-9 (http://www.house.leg.state.mn.us/hrd/as/85/as143.html)

are both low-powered and "one-size fits all", which does little to aid commuters [19]

Despite our best efforts, there are missing elements in the above counterfactual analysis. On one hand, we ignore the fact that under the new contract regime, contractors may be selected for on the basis of their ability to complete quickly, and therefore the winning bidders may actually have lower costs than those we estimated earlier. Moreover, they may make different decisions with respect to the hiring of labor and rental of capital than they currently do, enabling them to complete quicker without incurring the high costs we project when we hold labor and capital decisions fixed. These suggests we underestimate welfare gains.

On the other hand, with accelerated contracts and high-powered incentives comes increased need for monitoring the contractors. This is costly, since Mn/DOT would have to employ additional personnel on site. Similarly, if contractors are risk averse, time incentives increase contract risk and they will bid higher. In that sense we underestimate the costs to Mn/DOT. One more caveat is that we do not account for general equilibrium effects due to the bidding up of input prices — if one of these policy changes were implemented throughout Minnesota construction, one might expect that construction labor costs would rise as the premium to quick completion rose. Nor do we account for the deadweight loss of additional taxation needed to finance these time incentives. Nonetheless, the projected welfare gains are so large even under conservative assumptions and sub-optimal policy changes that it seems unlikely that our conclusions will be reversed by accounting for these effects.

6 Conclusion

This paper has shown that building time incentives into highway procurement contracts is important from the perspective of social welfare. From a theoretical perspective, we have shown that the A + B scoring auctions used by Mn/DOT and other agencies can be used to achieve efficient ex-ante allocations and ex-post outcomes, provided the incentives are correctly set. We have provided a characterization of the optimal incentive structure. Taking this theory to a rich dataset, we have found evidence that suggests that setting equal incentive and disincentive ex-post payments may achieve the best outcomes in practice.

We have also shown that the standard highway procurement contracts are generally ex-post

[19] That said, the few A+B contracts that have been used are a step in the right direction, as they do exactly what we suggest — provide high powered incentives on high usercost contracts.

inefficient, and thus that contract completion time might change if stronger time incentives were provided. Results from our counterfactual model demonstrate that there are large social gains to increased time incentives. We conclude that increasing the time incentives provided to contractors through careful auction and contract design would improve social welfare.

References

Arditi, David, Jotin Khisty, and Firuzan Yasamis, "Incentive/Disincentive Provisions in Highway Contracts," *Journal of Construction Engineering and Management*, September 1997, *123* (3), 302—207.

Asker, John and Estelle Cantillon, "Procurement when Price and Quality Matter," 2008. unpublished.

– **and** – , "Properties of Scoring Auctions," 2008. forthcoming in RAND Journal of Economics.

Athey, Susan and Jonathan Levin, "Information and Competition in U.S. Forest Service Timber Auctions," *The Journal of Political Economy*, apr 2001, *109* (2), 375–417.

Bajari, Patrick and Lixin Ye, "Deciding Between Competition and Collusion," *Review of Economics and Statistics*, 2003, *85* (4), 971–989.

– **and Steven Tadelis**, "Incentives Versus Transaction Costs: A Theory of Procurement Contracts," *The RAND Journal of Economics*, 2001, *32* (3).

– **, Stephanie Houghton, and Steven Tadelis**, "Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs," January 2007. unpublished.

Branco, Fernando, "The Design of Multidimensional Auctions," *Rand Journal of Economics*, 1997, *28* (1), 63—81.

Che, Yeon-Koo, "Design Competition Through Multidimensional Auctions," *The RAND Journal of Economics*, 1993, *24* (4).

Clough, Richard H., Glenn A. Sears, and S. Keoki Sears, *Construction Contracting*, John Wiley & Sons, 2005.

Einav, Liran and Ignacio Esponda, "Endogenous Participation and Local Market Power in Highway Procurement," 2008. Working paper, NYU Stern and Stanford.

Gil, Ricard and Justin Marion, "The Role of Repeated Interactions, Self-Enforcing Agreements and Relational [Sub]Contracting: Evidence from California Highway Procurement Auctions," 2009. Working paper, University of Santa Cruz.

Herbsman, Zohar, Wei Tong Chen, and William Epstein, "Time is Money: Innovative Contracting Methods in Highway Construction," *Journal of Construction Engineering and Management*, 1995, *121*, 273—281.

Hong, Han and Matthew Shum, "Increasing Competition and the Winner's Curse: Evidence from Procurement," *The Review of Economic Studies*, oct 2002, *69* (4), 871–898.

Jofre-Benet, Mireira and Martin Pesendorfer, "Estimation of a Dynamic Auction Game," *Econometrica*, September 2003, *71* (5), 1443–1489.

Krasnokutskaya, Elena, "Identification and Estimation in Highway Procurement Auctions under Unobserved Auction Heterogeneity," 2004. PIER Working Paper 05-006.

– **and Katja Seim**, "Bid Preference Programs and Participation in Highway Procurement Auctions," 2005. Working Paper, University of Pennsylvania.

Laffont, Jean-Jaques and Jean Tirole, "Auctioning Incentive Contracts," *Journal of Political Economy*, 1987, *95* (5), 921—937.

Li, Tong and Xiaoyong Zheng, "Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions," July 2006. CEMMAP Working Paper No. CWP13/06.

Manelli, Alejandro and Daniel Vincent, "Optimal Procurement Mechanisms," *Econometrica*, 1995, *63* (3), 591—620.

Marion, Justin, "Are Bid Preferences Benign? The Effect of Small Business Subsidies in Highway Procurement Auctions," *Journal of Public Economics*, 2007, *91*, 1591—1624.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green, *Microeconomic Theory*, Oxford University Press, 1995.

McCrary, Justin, "Manipulation of the Running Variable in a Regression Discontinuity Design: A Density Test," *Journal of Econometrics*, 2008, *142* (2).

Porter, Robert and Douglas Zona, "Detection of Bid Rigging in Procurement Auctions," *Journal of Political Economy*, 1993, *101* (3), 518—538.

Reiss, Peter and Matthew White, "Household Electricity Demand, Revisited," *Review of Economic Studies*, 2005, *72*, 853–883.

Silva, Dakshina De, Timothy Dunne, Anuruddha Kankanamge, and Georgia Kosmopoulou, "The impact of public information on bidding in highway procurement auctions," *European Economic Review*, 2008, *52* (1), 150—181.

Tirole, Jean, "Procurement and Renegotiation," *The Journal of Political Economy*, 1986, *94* (2), 235–259.

7 Appendix

Proof of Lemma 1: Taking the FOC of the profit function $\pi(b, d^a, d^b; \theta)$ with respect to d^a yields that $-c'(d^a; \theta) = I(d^a, d^b)$. Then since $-c(d; \theta)$ is strictly decreasing and $I(d^a, d^b)$ is increasing, there is a unique solution for given $(d^b; \theta)$, implying $d^a(d^b; \theta)$ is a function. The comparative static follows from the FOC since $I(d^a; d^b)$ is decreasing in d^b if $c_I < c_D$, and constant otherwise. ∎

Proof of Lemma 2: In the case where $c_I = c_U = c_D$, d^a does not depend on d^b. Simplifying (4), we get $E\left[c(d^a; \theta) - c_u d^a\right]$, independent of d^b, so indifference follows. If $c_I < c_U = c_D$ it suffices to show $d_i^b(x) \leq \underline{d}$ weakly optimal for all realizations of θ and strongly optimal sometimes. Fix θ and compare the payoffs to $d_i^b(x) = 0$ versus any deviation d'. There are two possibilities. Either $d' \leq d^a(d'; \theta)$ in which case the contractor will complete early late both on-path and in the deviation, with zero payoff difference. Or $d' > d^a(d'; \theta)$ in which case the contractor completes early in the deviation and late on-path, with payoff difference $c_U d' - c_I(d' - d^a(d'; \theta)) - c_D d^a(d^b; \theta) > 0$, so bidding d^b is strictly better. Finally, by definition of $\underline{d}$, it will sometimes be the case that $d' > d^a(d'; \theta)$ if $d' > \underline{d}$. The case with $c_I = c_U < c_D$ is similar. Finally, the comparative statics results follow directly from (4). ∎

Proof of Proposition 1: Ex-post efficiency requires that $d^a = d^*$, where d^* solves $-c'(d; \theta) = c_S$. For lane rentals, we have $-c'(d^a; \theta) = c_D$, so $c_D = c_S$ is both necessary and sufficient for the result. In I/D contracts with $c_I = c_D$, the optimal completion time sets $-c'(d^a; \theta) = c_D = c_I$, so $c_S = c_I = c_D$ suffices. It is also necessary, since by assumption 1, there exists θ_L with $-c'(d^E; \theta) < c_S$, so for efficient completion $c_I = c_S$. Similarly the existence of θ_H implies we must have $c_D = c_S$, and the necessity follows. Standard contracts are a special case of the I/D contract with $c_I = 0 < c_D$, and so are immediately ex-post inefficient. Finally, for A+B contracts, if $c_I < c_U = c_S = c_D$, by proposition 2, the bidder will bid $d^b \leq \underline{d}$ and by definition of $\underline{d}$ complete on time or late, so $-c'(d^a; \theta) = c_D = c_S$ is sufficient for efficiency. Similar logic gives sufficiency for $c_I = c_U = c_S \leq c_D$. For necessity notice that certainly either c_D or c_I must equal c_S, or the completion time will not be efficient. Moreover, if this holds and $c_I < c_U < c_D$, then by proposition 2 the contractor bids $d^b \in (\underline{d}, \overline{d})$. By definition of $(\underline{d}, \overline{d})$, there are realizations of θ for which the contractor wants to complete early, and some late, and then since c_S cannot equal both c_I and c_D, the completion time is inefficient. ∎

Proof of Proposition 2: The objective function faced by a bidder with signal x_i is to

maximize $E\left[\pi(b, d^b, d^a; \theta)|x_i\right] \Pi_{j\neq i}\left(1 - G_j(s)\right)$, where $G_j(s)$ is the equilibrium distribution of rival j's scores. Substituting in the optimal number of days and simplifying, we get the standard objective function in an FPA $(s - P_i(x))\,\Pi_{j\neq i}\left(1 - G_j(s)\right)$ with the score s replacing the bid b and the pseudo-cost $P_i(x)$ taking the place of the cost. The FOC follows directly from taking a derivative in s in the objective function. For (b), if the contract is ex-post efficient, the pseudo-cost simplifies to $E\left[c(d^a; \theta) + d^a c_S\right]$ plus a constant that is equal across bidders, and then by inspecting the welfare function it is clear that the bidder with the lowest pseudo-cost maximizes expected social welfare. If the bidding strategies are symmetric, the winner has the lowest pseudo-cost, and so the result follows. For standard contracts, if $c_D = c_S$ then again the bidder with the lowest pseudo-cost maximizes expected social welfare. ∎

Table 1: Summary Statistics: Bidding and Outcomes for Standard Contracts

	Diary Data Available				No Diary Data
	Mean	Standard Deviation	Min	Max	Mean
Engineer's Estimate	1085825	1574707	8869	11388853	2000276
Bid	1131726	1649215	7954.34	12979986	2094701
Markup	1.067	0.224	0.394	1.993	1.071
Winning Bid	1025031	1489726	7954	9549736	1938585
Winning Markup	0.944	0.175	0.394	1.462	0.964
Days Allowed	34.101	25.493	3	141	—
Days Taken	33.415	27.018	0.5	136.3	—
Prob. Enforced	0.290	0.458	0	1	—
Total Damages	714.18	3061.76	0	27000	—
Total Damages (% of EngEst)	0.174	0.135	0	1.05	—
N	762	762	762	762	6142

Means and standard deviations on firm-level variables are re-weighted to avoid oversampling contracts with many bidders.

Table 2: Damage Specifications and Time Outcomes for Standard Contracts

	Damages per Day ($)		Time Outcomes		
Contract Value ($)	1995–2004	2005-	# Obs	Avg. Days Late	% Late
Below 25K	75	150	7	-6.4	0
25K-50K	125	300	5	-2.6	0
50K-100K	250	300	19	-3.7	15.8
100K-500K	500	600	64	-2.7	20.3
500K-1M	750	1000	30	2.7	50.0
1M-2M	1250	1500	35	-1.37	37.1
2M-5M	1750	2000	15	3.37	53.3
5M-10M	2500	3000	7	3.4	57.1
Over 10M	3000	3500	0	—	—

The contract value is equal to the winning bid in the auction.

Table 3: Summary Statistics: Bidding and Outcomes for A+B Contracts

	Mean	Standard Deviation	Min	Max
Engineer's Estimate	11546182	18847279	618947	99154104
Maximum Days	154.37	242.81	15	1067
Usercost	10858.20	7342.00	3000	28000
Daily Incentives	5093.50	4117.06	0	10000
Daily Disincentives	10858.20	7726.10	3000	30000
No positive incentives (%)	34.5	—	—	—
Small positive incentives (%)	10.3	—	—	—
Equal incentives (%)	51.7	—	—	—
Higher usercost (%)	3.5	—	—	—
A-Bid ($)	13237849	22639349	601335	137423952
B-Bid (Days)	135.70	217.33	5	987
Markup (Bid / EngEst)	1.099	0.190	0.799	1.656
A-Bid of Winner	11861851	19844433	601335	102843344
B-Bid of Winner	122.34	193.14	6	987
Winning Markup	1.002	0.136	0.799	1.386
Total Incentive Payment	40926.09	63244.39	-30000	250000
N	123	123	123	123

Means and standard deviations on firm-level variables are re-weighted to avoid oversampling contracts with many bidders. "No positive incentives" means $0 = c_I < c_U = c_D$, "Small positive incentives" means $0 < c_I < c_U = c_D$, "Equal Incentives" means $c_I = c_U = c_D$, and "Higher usercost" means $c_I = c_D < c_U$.

Table 4: Descriptive Regressions

	Days Bid (fraction of Max)	$ Bid (fraction of estimate)	Days Taken (fraction of Max)
No positive incentives	-0.1523		
	(0.0745)		
Small positive incentives	0.0930		
	(0.0615)		
Higher usercost	-0.1963		
	(0.0516)		
Log Engineer Estimate		0.0166	
		(0.0051)	
Disincentives / EngEst		49.3047	-38.2252
		(6.8859)	(18.8496)
Incentives / EngEst		-34.8935	15.7458
		(7.0540)	(20.6784)
Bridge Dummy		0.0241	0.0239
		(0.0108)	(0.0630)
Grading Dummy		0.0320	-0.0158
		(0.0121)	(0.0740)
Milling Dummy		-0.0025	0.0613
		(0.0105)	(0.0659)
Traffic Control Dummy		-0.0331	-0.0902
		(0.0140)	(0.0733)
District FE	no	yes	yes
R^2	—	0.0848	0.0845
N	123	6635	270

Robust standard errors are in parentheses, clustered by contract. The first column is of a modified tobit, used because of censoring at the minimum and at the maximum days bid. The excluded group is equal incentives and disincentives; and the data is only from A+B auctions (where there is bidding on days). In the second and third columns regression is by OLS, and fixed effects by district and the primary task in the contract are included. All bids are used in column 2. In column 3, only the diary subsample plus the A+B contracts are used, since we only have timing data for these contracts.

Table 5: Completion Time

Estimated Cost Function for Average Contract ($\approx$ 37 days)		
Variable	Coefficient	Standard Error
α_0	18895.2	3445.5
α_1	-510.12	93.3
σ	4317.9	943.4
Marginal Effect: Log Contract Size on σ	-572.7	266.9

Estimates of the marginal time cost function — the intercept α_0 and slope α_1 — and the shock standard deviation σ, for an average contract. Standard errors obtained by bootstrapping.

Table 6: Summary Statistics: User Costs of Construction

	Mean	Standard Deviation	Min	Max
Delay (in minutes)	9.11	14.2	1	114
Traffic (daily commuters)	18551.27	28254.01	300	160000
Estimated User Cost ($ per day)	15625.74	29425.39	56	274773.3

Table 7: Summary Statistics for Counterfactual Data

	Mean	Std Deviation	Mean (re-weighted)	Std Deviation (re-weighted)
No change				
Days Taken	33.90	28.15	36.86	30.98
Private Time Reduction Costs	503.53	643.73	545.63	703.92
Damages Paid	2139.88	2785.60	2432.39	3142.15
Lane Rental = Current Damages				
Days Taken	32.48	27.49	35.44	30.25
Private Time Reduction Costs	1720.84	2205.28	1833.03	2416.06
Damages Paid	50649.37	75307.33	57887.33	85427.21
Lane Rental = 10% of Usercost				
Days Taken	31.81	28.20	34.71	31.29
Private Time Reduction Costs	7042.28	25348.10	7389.20	25053.41
Damages Paid	46470.12	91166.06	49237.40	91361.21
Lane Rental = Usercost (efficient policy)				
Days Taken	18.92	26.56	21.55	29.66
Private Time Reduction Costs	136986.71	278635.19	145158.22	296037.69
Damages Paid	109651.47	260161.57	129759.57	286165.38

Counterfactual estimates of statistics of interest, for three counterfactual lane rental policies. Estimates are from simulations based on a sample of 167 contracts where we have detailed traffic and construction data. The last two columns are constructed by re-weighting the outcomes by the inverse of their probability of being included in the sample used for the counterfactual.

Table 8: Counterfactual Welfare Measures

	Mean	Std Deviation	Mean (re-weighted)	Std Deviation (re-weighted)
Lane Rental = Current Damages				
Δ User Welfare	26676.01	61441.43	27842.03	62116.99
Δ Private Costs	1217.32	1581.65	1287.40	1729.95
Δ Cost to Mn/DOT	11162.68	16233.00	12635.87	18370.01
Lane Rental = 10% of Usercost				
Δ User Welfare	146275.86	613484.89	170651.13	708984.20
Δ Private Costs	6538.75	25204.26	6843.56	24930.21
Δ Cost to Mn/DOT	16712.55	44688.17	17573.28	43338.79
Lane Rental = Usercost (efficient policy)				
Δ User Welfare	501325.60	1316048.13	533265.56	1326104.65
Δ Private Costs	136483.18	278263.23	144612.59	295651.56
Δ Cost to Mn/DOT	185282.14	344444.12	199000.55	366089.41

Counterfactual estimates of welfare measures, for three counterfactual lane rental policies. Estimates are from simulations based on a sample of 167 contracts where we have detailed traffic and construction data. The last two columns are constructed by re-weighting the outcomes by the inverse of their probability of being included in the sample used for the counterfactual.