

Procurement Contracting with Time Incentives: Theory and Evidence

Patrick Bajari*

University of Minnesota and NBER

Gregory Lewis†

Harvard University

February 1, 2009

Abstract

In public sector procurement, social welfare often depends on how long it takes to complete the contract. A leading example is highway construction, where slow completion times inflict a negative externality on commuters. In standard highway contracts, contractors pay only a small penalty for late project completion. As a result, contractors do not fully internalize the negative externality of delays, resulting in welfare losses compared to an efficient allocation. Recently, state highway departments introduced an innovative contracting method called A+B bidding to reduce this inefficiency, where contractors bid on project completion time in addition to total project cost. In this paper we compare these two market designs theoretically and empirically. We characterize equilibrium bidding and efficient design, showing that A+B contracts can achieve the social optimum. We then gather a unique data set of highway repair projects awarded by the Minnesota Department of Transportation which include both contract forms. Our descriptive empirical analysis suggests that observed behavior is broadly consistent with the predictions of our theoretical model. Next we build a structural econometric model which endogenizes project completion times. Our estimates suggest that switching from standard contracts to designs with socially efficient time incentives could increase welfare by over 32% of the contract value on average. We conclude that large improvements in social welfare are possible through the use of improved auction design.

*Email: bajari@umn.edu

†Email: glewis@fas.harvard.edu

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 contracts, however, social welfare depends upon the time to complete the contract. Unfortunately, standard procurement mechanisms do not use a measure of 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 would generate 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 an innovative auction design 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 equilibrium behavior and outcomes in A-only versus A+B contracts. We start by building a theoretical model of A+B contracts, that includes

A-only contracts and other commonly used contract designs as special cases. We characterize equilibrium bids and project completion times. Our results demonstrate that contractors may strategically manipulate their bids by bidding completion times that differ substantially from their internal target project completion times. However, by optimally selecting the scoring rule for an A+B auction, we demonstrate that even with such manipulation 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 277 highway contracts awarded by the Minnesota Department of Transportation between 1997-2007. Our data includes ex-post penalties 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 contracts are rarely completed ahead of schedule when there are no bonuses for early contract completion. By contrast, when there are positive incentives, contractors strategically inflate their days bid and then complete projects early in order to earn incentive payments.

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. Since we observe the incentive structure, we know the marginal benefits and can use this first order condition to back out the marginal costs. The problem is not trivial, since contractors often face discontinuous incentives at the scheduled completion time. Economic theory predicts that the dependent variable — actual completion time — should “bunch” at that point. 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 32% of the contract value in lost surplus in an average contract. 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 Arditì, Khisty and Yasamis (1997) and Herbsman, Chen and Epstein (1995)). These papers are primarily descriptive, and cannot provide the counterfactual welfare analysis that we obtain from our richer model and empirical 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 that touches 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) and Einav and Esponda (2008)).

Section 2 presents the theoretical analysis and Section 3 contains the empirical analysis. The counterfactual is in section 4, and section 5 concludes. All proofs and tables are in the appendix.

2 Theory

In this section, 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, 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 can choose the actual completion time d^a . His payoff is given by:

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

where $c(d; \theta)$ is his private cost function. This is just his bid plus incentive payments (if any), less incentive payments (if any), less private costs. The cost function is parameterized by $\theta \in \mathbb{R}^k$, with k finite. It is assumed to be twice continuously differentiable, strictly decreasing and strictly convex in d for all θ so that completing the contract more slowly lowers costs, but at a decreasing rate. We assume also $c'(d_0; \theta) = 0$ for some finite d_0 , so infinite delay is not profitable.

For example, suppose that the total project cost is the sum of materials and labor costs. Suppose also that 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

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

so that total costs are strictly decreasing and strictly convex in the days taken d .

At the time of bidding, the contractor may be uncertain as to the realization of his private costs, conditional on winning the contract. To formalize this, we assume that each contractor i does not observe θ_i , but instead receives only a vector of private signals $x_i \in \mathbb{R}^k$ affiliated with θ_i at the time of bidding. We assume also that each contractor i draws the pair (x_i, θ_i) *independently* from some 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.

Equilibrium: A (Bayes-Nash) equilibrium of the game comprises a set of bidding strategies $(\beta_1(x), \dots, \beta_n(x))$ 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 $W(b, d^a; \theta) = V - c(d^a; \theta) - d^a c_S$. 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 expected social welfare $E[W(b, d^a; \theta) | x_i]$ 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. They are typically very small. 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 key role in real-world contracting. We assume that the project engineer can successfully monitor the con-

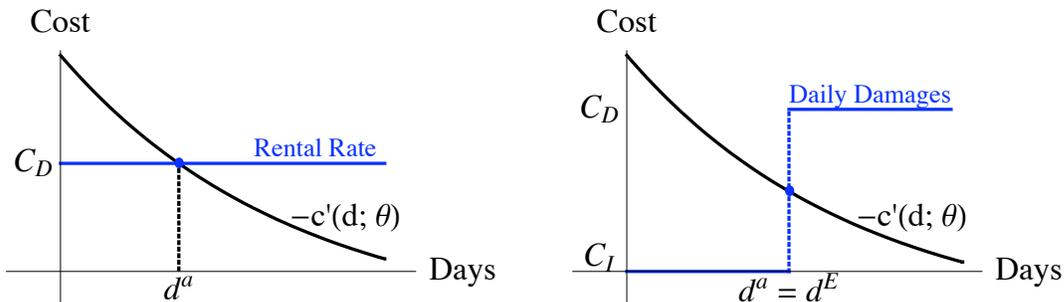


Figure 1: **Completion Time in Lane Rentals and Standard Contracts.** The figure depicts the marginal benefit to delay $-c'(d; \theta)$ curve and the incentive structure for lane rentals and standard contracts. In the left panel, lane rental imposes a constant cost of delay c_D , so the contractor optimally 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 $d^a = d^E$.

struction and ensure that the finished project meets the contract specifications.¹ 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.

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 1 we depict the “moving parts” of the timing problem. Consider first the left panel, which 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. On the other hand since his private costs are concave and decreasing

¹Contractors are limited in their ability to shirk on quality by performance-based specifications, and bonding requirements, which are mandatory on public contracts.

in the number of days taken, he faces a declining marginal benefit of delay². This is 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 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:

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

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 proposition:

Proposition 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.³ For any realization of θ , her incentive payments $IP(d^a(d^b; \theta), d^b, \theta)$ depend entirely on how many days she bid. For example, if she bids a 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

²This is exactly analogous to the marginal cost of abatement curve in environmental economics

³The contractor has no choice in the other contract designs.

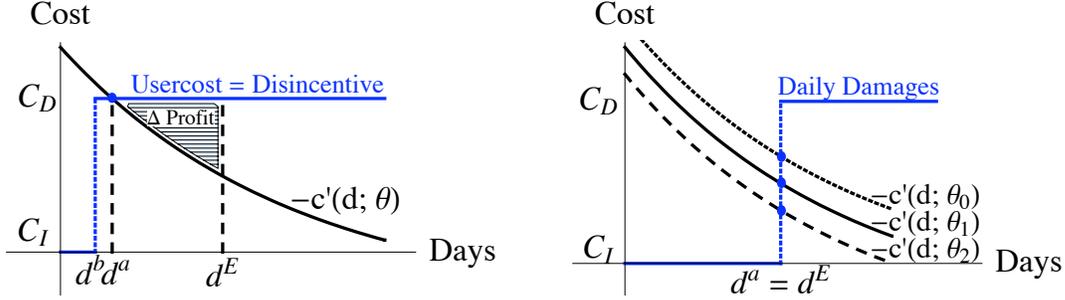


Figure 2: **Bidding and Stickiness in Completion Time** In the left panel, we depict the gains to bidding $d^b < d^E$ in a A + B contract where $c_I = 0 < c_U = c_D$, for a particular realization of θ . The additional profits from this strategy are shown in the shaded region. In the right panel, we compare 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.

function (1), it follows that the optimal days bid is:

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

To see these trade-offs graphically, look at the left panel of Figure 2. 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 $\bar{d} = \sup\{d : \exists \theta : -c'(d; \theta) = c_I\}$.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 \bar{d} . Then we have the following result:

Proposition 2 (Bidding on Time) *In equilibrium, if $c_I < c_D \leq c_U$, the contractor*

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

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 $d_i^b(x) \in (\underline{d}, \bar{d})$, 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 2, 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 2, 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 requires 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 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 3 (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:

$$P_i(x) = E [c(d^a; \theta) - 1(d_i^b(x) > d^a)(d_i^b(x) - d^a)c_I + 1(d_i^b(x) < d^a)(d^a - d_i^b(x)c_D + c_U d_i^b(x)|x)] \quad (5)$$

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:

⁵This 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.

Proposition 4 (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*

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

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 FPA with the pseudo-cost replacing the cost. Simple algebra then yields 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.

3 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 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 estimate the private cost function of contractors in standard contracts, using a first order condition based on the theory above. This estimation exercise allows us to run counterfactual simulations later in the paper to quantify how much welfare could be improved through better design.

3.1 The Data

We have data from Mn/DOT, on all the highway procurement contracts completed in Minnesota during the period 1997-2007 (a total of 3702 contracts). 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. A linear regression of all of the bids in our data — including A-only contracts — on the cost estimate alone has an R^2 of 0.97.

The second source of data is data on the contract completion time. For a subset of completed standard contracts (248 of them), Mn/DOT we have actual diary data used by Mn/DOT project engineers to record the contract progress.⁶ Using their own software, a program called FieldOps, we exported this data and 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. One reason why this data is usually not so easy to obtain is that 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 to decide whether or not to assess damages for late completion. This allows the project engineer to make allowances for bad weather, or other unforeseen construction delays. It also means though that the contract days in the bidding abstract don’t match the observed data very well, and so it is extremely useful to have the diary data for the completion time analysis.

⁶For the remaining contracts, the diary file was either missing or corrupt.

Our final data source was the data on A+B contracts, which we got from Mn/DOT’s office of construction and innovative contracting. This contract design has typically only been used on extremely time-sensitive projects, and so there are only 29 such contracts, with 123 bids placed in the auctions. Here 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). In the next section we present some descriptive evidence on these contracts, before returning to the larger dataset on standard contracts.

3.2 A+B Auctions

Some summary statistics on A+B auctions are presented in Table 1. These contracts are large, of average value \$11 million and estimated duration 151 days. By comparison, our standard contracts are around \$1 million, and last about 34 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. In the rest of this section, we explore this data much more carefully, and in particular see if outcomes vary with the incentive structure in the precise way the theory predicts.

There are four basic kinds of incentive structure that are observed in the data: (a) those with no positive incentives, and equal user cost and disincentive ($0 = c_I < c_U = c_D$); (b) those with small positive incentives ($0 < c_I < c_U = c_D$); (c) those with entirely equal incentives and disincentives ($0 < c_I < c_U = c_D$); and (d) those with user cost higher than the other incentives ($c_U > c_I = c_D > 0$). The breakdown of each type is shown in Table 2. Our analysis will be primarily descriptive, because we lack enough data to do too much more than this. We will treat the incentive structure as exogenous throughout the analysis.⁷

The first set of tests we run is on the bidding behavior. The theory predicts that

⁷We will control for some observable contract characteristics, so in practice we require only the weaker assumption of exogenous contract assignment conditional on those observables.

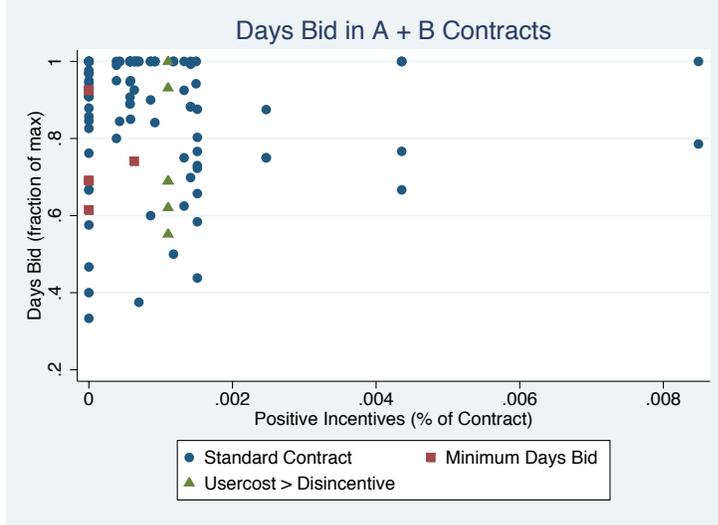


Figure 3: **Days Bid in A+B Contracts.** This figure shows how the number of days bid varies with the incentive structure.

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 $c_U = c_D$, they should bid zero days (see Lemma 2). Likewise, when with high user cost $c_U < c_D$, the optimal bid is zero days. By contrast, in the case of equal incentives ($c_U = c_I = c_D$) the theory makes no predictions on the optimal number of days to bid, as it has no equilibrium impact on payoffs. The observed days bid are shown in Figure 3.

Notice that there appears to be a pattern: the lower the incentives, the lower the days bid. Of the four cases in which the minimum number of days was bid, 3 of these occur with no positive incentives. Similarly, in the one contract where the usercost was set above the disincentive, most of the days bids were well below the maximum bids. All of this goes towards supporting the theory. Yet the theory predicts something even stronger: bids of zero days. In practice, contractors don’t bid zero days for fear that Mn/DOT will classify their bid as “irregular” and disqualify it from the auction. This may explain why we don’t see this theoretical prediction borne out in the data.

To confirm that there is a statistically significant difference in bidding behavior as the incentive structure changes, we run a tobit regression of days bid (as a percentage of the maximum allowed) on dummies for the different incentive structures. A tobit is used to account for the censoring of days bid both below (at the minimum days) and

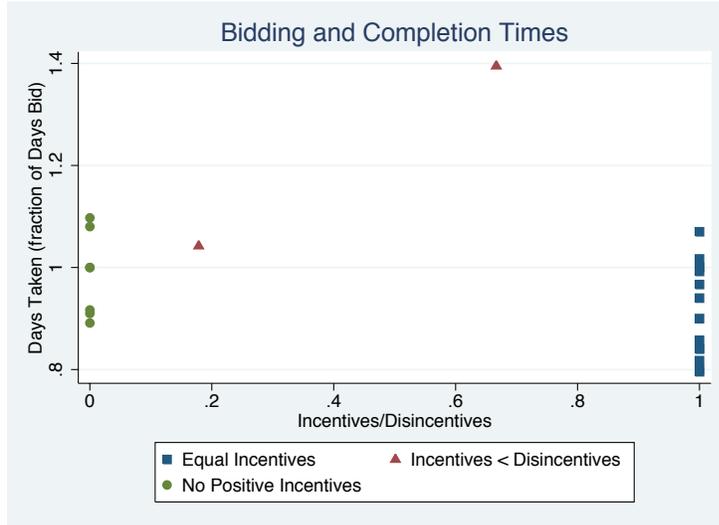


Figure 4: **Bidding and Completion Times.** This shows how the final days taken relates to the number of days bids, for varying incentive structures.

above (at the maximum days). The outcome is shown in the first column of Table 3. 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.

We turn next to the actual dollar bid in these contracts. Theory predicts that the A-bid should be relatively higher in contracts with zero or small positive incentives, since the winning contractor will receive no bonus payment for finishing early. In fact, they should optimally bid zero days and build in all of their anticipated damages into their bid. The same is true for contracts with usercost exceeding disincentives.

In column 2 of Table 3 we show the outcome of an OLS regression of their dollar bid (as a percentage of the engineer’s estimate) on the incentive structure. We add a dummy for whether there are more than two bidders as a control for the competitiveness of the auction; its coefficient is negative and significant, as one would expect. We find that the bids are higher in contracts with no positive incentives and higher usercost, though only significantly so in the latter case. This is in line with the theory, although it is perhaps surprising that the case with no positive incentives is not generating much higher bids — more on this later.

Finally, we consider how the incentive structure impacts the final completion time. Figures 4 and 5 show what we observe in the data. Figures 4 plots the final completion

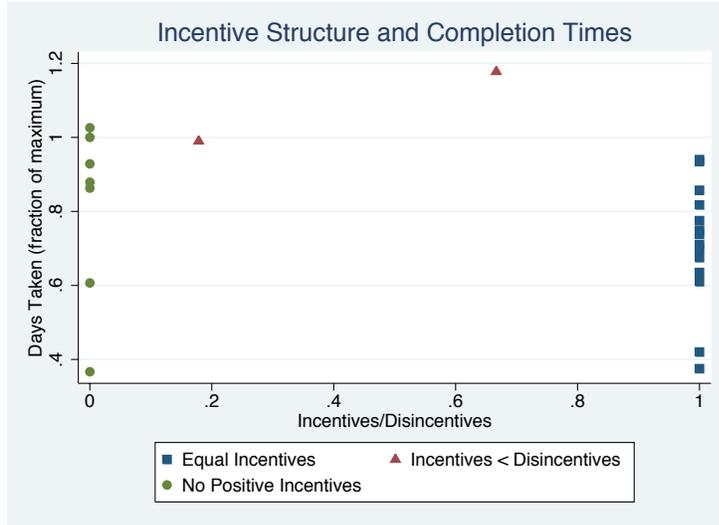


Figure 5: **Incentive Structure and Completion Times.** This shows how the completion time (as a percentage of the maximum time allowed) varies with the incentive structure.

time as a percentage of days bid, against incentives.

The theory predicts that when $c_I < c_D$ the actual completion time should be “sticky” around d^b , since there is a wedge between the benefits to finishing earlier (which are low) and the costs of finishing late (which are high). By contrast, the theory has nothing to say in the case with equal incentives, since the number of days bid has no equilibrium impact on the optimal completion time, nor on profits.

In Figure 4 we do observe some of the “stickiness” predicted by the theory. When there are zero positive incentives, the actual completion time appears to be within 10% of the number of days bid. On the other hand with equal incentives, they tend to finish the job well ahead of schedule, cashing in on the incentive payments. This seems to indicate that contractors prefer to bid high days and then finish early, earning bonuses; instead of using the theoretically equivalent strategy of bidding low days, pushing up A-part of the initial bid, and then paying out damages.

Look now at Figure 5, which shows the relationship between incentive structure and final completion time (as a percentage of the maximum days). Theory predicts that for a given contract all of the observed incentive structures should yield the same final completion time; none should systematically do better than the others. In the figure, we see that almost all incentives structures lead to a reasonably quick completion

time (on average 80% of the maximum with no positive incentives; 70% with equal incentives). It is hard to say whether one incentive structure does better than the other with so little data, but an OLS regression in the third column of Table 3 shows that the difference in days taken is not significant, as the theory would predict.

Overall, the predictions of the theory stand up well, at least in this primarily descriptive analysis. An important deviation from the theory is that bidders do not bid zero or even 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.

3.3 Standard Procurement Auctions

Having shown that A + B contracts do seem to achieve their objective of promoting timely project completion, we turn now to standard contracts, where the time incentives are weak. We summarize this dataset in Table 4, splitting up the contracts for which we have diary data and those where we do not. Examining Table 4 we see that 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 subsample for which we have timing data is not representative of the full set of Mn/DOT contracts, which are on average twice as big, at \$2 million. Since we will argue later that the penalties for delays in large contracts are far too small, this selection bias will if anything strengthen the welfare conclusions we reach towards the end of the paper.

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 (in most cases) to as high as 1.05% of contract value, in a particularly bad case. We show how the damages penalties are specified in Table 5. These are standard contract specifications: in every standard highway contract issued in Minnesota, the project engineer can assess damages per day late according to the contract value, in accordance with the specifications. So for example, in a typical \$1M contract, the penalty for being a day late is \$1000.

Also in that table, 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. In this section, we estimate the contractor's private costs by looking at their behavior as damages vary. As noted earlier, our diary data subsample consists of only 248 contracts. As a result, 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. The first order conditions must satisfy:

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

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

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

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.

We will let θ be unidimensional, and use a simple linear form for the marginal benefit of delay function:

$$-c'(\theta; d) = \alpha_0 + \alpha_1 d + \theta \quad (9)$$

where $\alpha_0 > 0$, $\alpha_1 < 0$ in line with the theory. We will also assume that $\theta \sim N(0, \sigma^2)$, a parametric choice that turns out to fit the data quite well. Our specification is

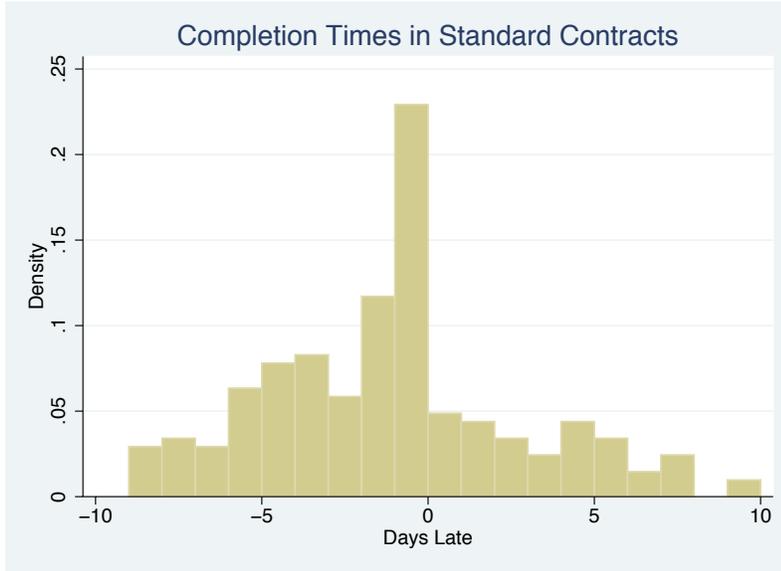


Figure 6: **Completion Times Histogram** The histogram is of the difference between the actual and contractually specified completion time, where 0 is exactly on time. Notice the huge spike just before 0 — many contracts are completed exactly on time, as the theory predicts.

very parsimonious, which is appropriate given the limited number of observations in our sample. We have experimented with other specifications, but they have not contributed to fit.

To see the importance of taking the theory seriously here, examine Figure 6. This is a histogram of the contract completion time relative to the days allowed, so that 0 is exactly on time. Recall that our theory predicts that contract completion time will be “sticky” around d^E in standard contracts (see the right panel of Figure 2). 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.

It also implies that a naive OLS regression of days on incentives would be badly misspecified since there is no way for the errors to account for the clustering of contracts at 0, much in the same way that OLS is inappropriate with censored data. But by using the first order conditions directly, we can avoid these problems. Our approach is to estimate the model by maximum likelihood. An important concern is that the

contracts are not identical, so the marginal benefit of delay will not be equal across contracts. To get some intuition for how the time costs for different contracts should be related, consider again equation 2. In that case, one can show that the marginal benefit of delay function satisfies:

$$-c'(\alpha d; \alpha \theta) = (\alpha \theta) w' \left(\frac{\alpha \theta}{\alpha d} \right) \frac{1}{(\alpha \theta)^2} = \frac{-c'(d; \theta)}{\alpha}$$

because when scaling both the number of days taken and the total amount of work θ up by a constant, the marginal benefit of a single day's delay is smaller than in the smaller contract, where it constitutes a greater fraction of the total work. With that in mind, we assume that the marginal benefit of delay function is homogenous of degree -1. In that case, if the total amount of work on contract i , θ_i , is proportional to the engineer's day estimate d^E — which seems reasonable — we can divide through by d^E to get the cost function for contract i :

$$-c'(d; \theta_i) = d^E c' \left(\frac{d}{d^E}, \theta \right) = d^E \left(\alpha_0 - \alpha_1 \frac{d}{d^E} + \theta \right)$$

where $\theta \sim N(0, \sigma^2)$ as before. This allows us to normalize through by d^E to homogenize the contracts for the estimation, and then rescale by d^E afterwards to get the cost function for any specific contract.

So let $\tilde{d} = d/d^E$. Then we can write down a log likelihood for each observation, $\ell_i(\theta)$:

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

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, $\tilde{d} = 1$.

In implementing this procedure, we allow the standard deviation σ to vary with the contract size, by specifying that σ is a linear function of the log contract size. This is motivated by the observation that in the data, even after normalizing through by

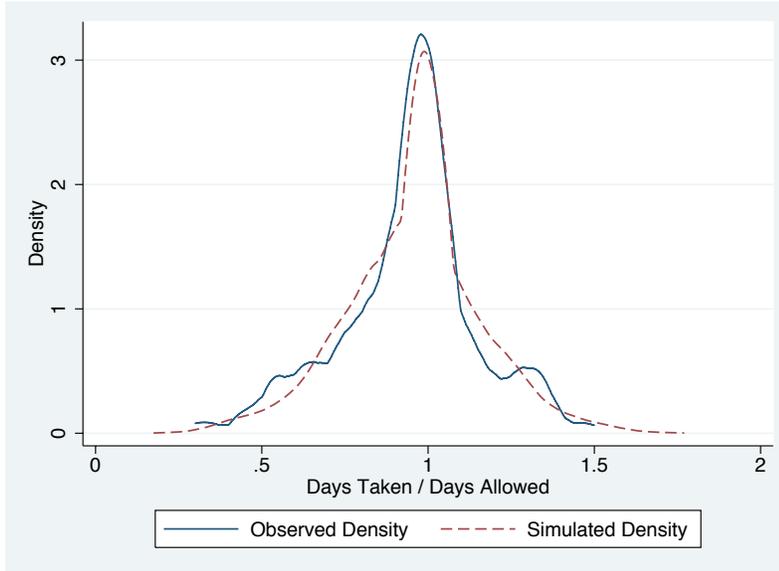


Figure 7: **Fit of the Timing Model** The figure shows kernel density estimates of the normalized completion times d/\bar{d} for both the actual and data simulated from the structural model.

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 6. For an average contract, which lasts 37 days, we estimate the marginal time cost function to be $18895 - 510d$, a line that reaches zero just before the allotted completion time. In other words, we predict that on average, contracts will complete slightly early, consistent with the data. The signs of the coefficients are as expected, and all are significant. In particular, we find 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.

4 Counterfactual

The key contention of the paper is that building time incentives into highway procurement contracts could yield huge social welfare gains. To assess this, we consider a simple counterfactual policy change in which we introduce lane rentals, which as we saw earlier have a relatively simple efficient design. We already have a model of the contractor’s private costs, developed above. We augment this by constructing measures of road user costs for the roads under construction in 147 of the contracts.⁸ This is necessary because Mn/DOT generally only calculates user costs for A+B projects, or projects 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:

$$User\ Cost_t = Delay_t \times Time\ Value_t \times Traffic_t \quad (11)$$

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.

⁸We 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.

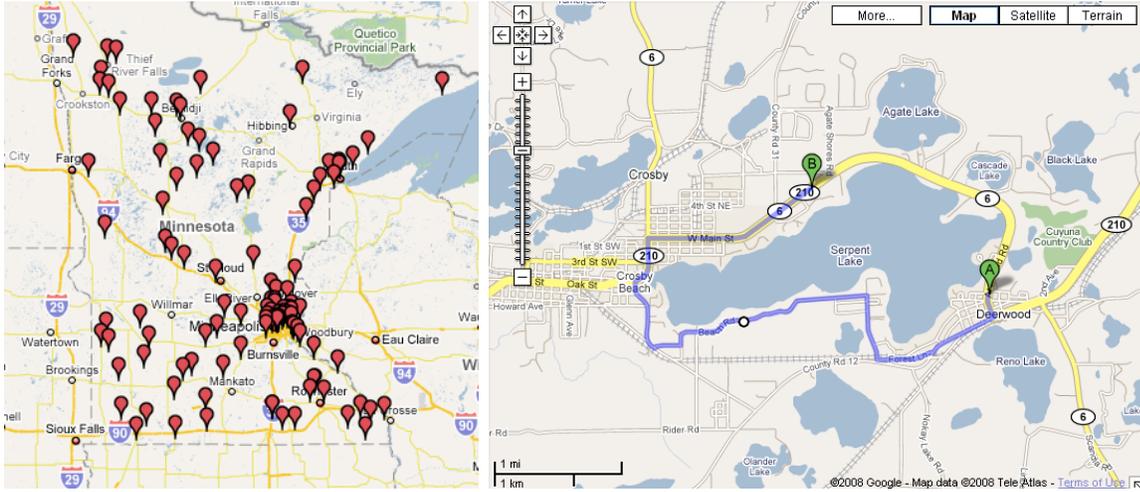


Figure 8: **Delay Calculations:** The left panel shows the locations of the highway construction projects used in the counterfactual. The right panel shows an example of a detour calculation around a section of route 6.

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 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 recording 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.⁹ 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 summarize all these measures, including the usercost we compute, in Table 7. The estimated average delay is 10 minutes, which with around 20 000 daily commuters on a typical

⁹See <http://www.dot.state.mn.us/traffic/data/html/volumes.html>

road results in an average usercost of \$15348.45. We should note that since we cannot locate all contract sites, these contracts may be selected to be in metro areas, and therefore have higher traffic than average. Our counterfactual results should thus be interpreted with this qualification in mind.

We consider two relatively small counterfactual changes: one in which Mn/DOT charges a lane rental equal to the current damages specified in standard contracts (see Table 5), and another in which Mn/DOT charges a percentage lane rental equal to 0.25% of the contract value per day, or \$2500 per day on a \$1 million contract. We also consider a switch to the efficient contract design, where the lane rental is equal to the per day usercost. 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.

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 time 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 now compare an estimate of the social gains from the average reduction in user costs with the private cost increase faced by the contractor.

But it will also be interesting to calculate 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:

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

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, paving) and primary material used (e.g. asphalt, concrete). 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 the selection problem.

The results are in Tables 8 and 9. 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 0.25% per day reduces the mean completion time to 28.6 days. The efficient design has a striking effect, creating incentives for contractors to complete in only 18.8 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, basically 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 then are the changes in producer costs and user welfare? The magnitudes are striking! Concentrating on the re-weighted results, the smaller lane rental policy improves user welfare by \$32300 per contract, at a cost to producers of only \$1279. The more aggressive policies have even larger effects. The 0.25% lane rental improves welfare overall by around \$104 000 per contract, and the efficient policy yields gains of \$639 000 per contract. For the full Minnesota sample, this is around 32% of contract value. This is a substantial amount of money. The problem is of course that this would cost Mn/DOT a lot of money to implement. We estimate that under the first policy, the net contract costs would increase by \$12551 or 0.625%; in the second case by \$73030 or 3.65%; and in the efficient case by 201893, or 10.1%. In all cases though these costs are well below the user gain, so raising taxes to pay for these time incentives would make some sense.

There are two implicit assumptions in this analysis, which bias the results in opposite

directions. On the 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. On the other hand, we ignore 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. 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.

5 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 unique and rich dataset, we have found evidence that suggests that setting equal incentive and disincentive payments ex-post may achieve the best outcomes in practice.

We have also shown that the standard highway procurement contracts are generally ex-post inefficient, and thus considered 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).
- Einav, Liran and Ignacio Esponda**, “Endogenous Participation and Local Market Power in Highway Procurement,” 2008. Working paper, NYU Stern and Stanford.
- 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*, August 2007, 91, 1591—1624.
- Porter, Robert and Douglas Zona**, “Detection of Bid Rigging in Procurement Auctions,” *Journal of Political Economy*, 1993, 101 (3), 518—538.
- 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.

6 Appendix

6.1 Proof of Proposition 1

Examining the FOC of the profit function $\pi(b, d^a, d^b; \theta)$ with respect to d^a yields immediately that $-c'(d^a; \theta) = I(d^a, d^b)$. Moreover 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 immediately from the FOC on noting that $I(d^a; d^b)$ is decreasing in d^b if $c_I < c_D$, and constant otherwise.

6.2 Proof of Proposition 2

In the case where $c_I = c_U = c_D$, d^a does not depend on d^b . Simplifying (4), we get $E[c(d^a; \theta) - c_u d^a]$, independent of d^b , so indifference follows. If $c_I < c_U = c_D$, it will suffice to show that $d_i^b(x) \leq \underline{d}$ is 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).

6.3 Proof of Proposition 3

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, by the above logic. 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 $-c'(d^a; \theta) = c_D = c_S$ as required 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 is true and $c_I < c_U < c_D$, then by proposition 2 the contractor bids $d^b \in (\underline{d}, \bar{d})$. By definition of these bounds then, 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.

6.4 Proof of Proposition 4

The objective function faced by a bidder with signal x_i is to maximize $E [\pi(b, d^b, d^a; \theta) | x_i] \prod_{j \neq i} G_j(s)$, 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 $(s - P_i(x)) \prod_{j \neq i} G_j(s)$ 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 $E [c(d^a; \theta) + d^a c_S]$ 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.

6.5 Tables

Table 1: Summary Statistics: Bidding and Outcomes for A+B Contracts

	Mean	Standard Deviation	Min	Max
Engineer's Estimate	11437330	18943750	618947	99154104
Maximum Days	151.03	238.68	15	1067
Usercost	10756.46	7250.72	3000	28000
Daily Incentives	5168.94	4050.0	0	10000
Daily Disincentives	10756.46	7645.25	3000	30000
A-Bid (\$)	13102066	22743980	601335	137423952
B-Bid (Days)	132.97	214.39	5	987
Markup	1.099	0.190	0.799	1.656
A-Bid of Winner	11861851	19844433	601335	102843344
B-Bid of Winner	123.87	199.28	6	987
Winning Markup	1.003	0.140	6	987
Incentive Payment	41599.39	64886.69	-30000	250000
N	123	123	123	123

Table 2: Observed A + B Incentive Structures

	# Auctions	% of Auctions
No positive incentives	10	34.5
Small positive incentives	3	10.3
Equal incentives	15	51.7
Higher usercost	1	3.5
Total	29	100

“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 3: A+B Regressions

	Days Bid (% of Max)	\$ Bid (% of Est)	Days Taken (% of Max)
No positive incentives	-0.1523 (0.0745)	0.1157 (0.0752)	0.1024 (0.1049)
Small positive incentives	0.0930 (0.0615)	-0.0240 (0.0385)	—
Higher usercost	-0.1963 (0.0516)	0.0634 (0.0328)	—
More than two bidders		-0.1062 (0.0661)	—
N	123	123	23

The excluded group is equal incentives. Robust standard errors are in parentheses, clustered by contract in the first two columns. In the first column, a modified tobit is used to account for the censoring of days bid at the minimum and the maximum days bid. In the second and third columns, the regression technique is OLS. The last two groups are omitted in the final regression because there are too few observations.

Table 4: 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	—
Damages	714.18	3061.76	0	27000	—
Damages (% of Total)	0.174	0.135	0	1.05	—
N	762	762	762	762	6142

Table 5: Damage Specifications and Time Outcomes for Standard Contracts

Contract Value (\$)	Damages per Day (\$)		# Obs	Time Outcomes	
	1995–2004	2005–		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	—	—	—

The over 10M category is excluded since no contracts of that size are observed in the subsample. % late indicates the percentage of contracts that were completed late.

Table 6: Completion Time

Estimated Cost Function for Average Contract (≈ 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 parameters 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 the maximum likelihood procedure.

Table 7: Summary Statistics: User Costs of Construction

	Mean	Standard Deviation	Min	Max
Delay (in minutes)	9.20	14.7	1	114
Traffic (daily commuters)	19225.99	29472.12	300	160000
Estimated User Cost (\$ per day)	14744.71	22388.81	56	156240
N	147	147	147	147

Table 8: Summary Statistics for Counterfactual Data

	Mean	Std Deviation	Mean (re-weighted)	Std Deviation (re-weighted)
No change				
Days Taken	33.91	28.05	36.87	30.81
Private Time Reduction Costs	499.14	639.79	538.99	698.71
Damages Paid	2121.45	2756.54	2407.09	3103.14
Lane Rental = Current Damages				
Days Taken	32.50	27.40	35.45	30.08
Private Time Reduction Costs	1712.76	2201.10	1817.84	2408.07
Damages Paid	50454.10	75072.11	57493.30	85012.95
0.25% Lane Rental				
Days Taken	28.56	22.77	31.06	24.84
Private Time Reduction Costs	26961.39	80690.40	32701.85	92386.72
Damages Paid	145500.25	314574.00	174584.30	362532.33
Lane Rental = Usercost (efficient policy)				
Days Taken	18.79	26.51	21.31	29.57
Private Time Reduction Costs	138543.93	278461.99	147808.92	295571.03
Damages Paid	108857.74	259307.58	128256.48	284764.73

Counterfactual estimates of statistics of interest, for three counterfactual lane rental policies. Estimates are from simulations based on a sample of 147 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 9: Counterfactual Welfare Measures

	Mean	Std Deviation	Mean (re-weighted)	Std Deviation (re-weighted)
Lane Rental = Current Damages				
Δ User Welfare	29280.95	69640.31	32300.72	75347.82
Δ Private Costs	1213.62	1581.23	1278.85	1726.98
Δ Cost to Mn/DOT	11122.87	16191.19	12551.86	18291.67
0.25% Lane Rental				
Δ User Welfare	126498.74	397138.19	136042.89	390826.24
Δ Private Costs	26462.25	80202.81	32162.87	91847.05
Δ Cost to Mn/DOT	60430.45	153463.34	73030.88	175979.08
Lane Rental = Usercost (efficient policy)				
Δ User Welfare	647897.87	2305764.44	786913.31	2808385.11
Δ Private Costs	138044.79	278096.01	147269.94	295193.96
Δ Cost to Mn/DOT	187001.00	344046.94	201893.80	365198.37

Counterfactual estimates of welfare measures, for three counterfactual lane rental policies. Estimates are from simulations based on a sample of 147 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.