

NBER WORKING PAPER SERIES

PROCUREMENT CONTRACTING WITH TIME INCENTIVES:
THEORY AND EVIDENCE

Patrick Bajari
Gregory Lewis

Working Paper 14855
<http://www.nber.org/papers/w14855>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2009

Bajari would like to thank the National Science Foundation, the University of Minnesota Grant in Aid program and the Center for Urban and Regional Affairs at the Humphrey Institute at the University of Minnesota for generous research support. 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 IOOC, 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. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2009 by Patrick Bajari and Gregory Lewis. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Procurement Contracting with Time Incentives: Theory and Evidence
Patrick Bajari and Gregory Lewis
NBER Working Paper No. 14855
April 2009, Revised September 2010
JEL No. D02,D21,D44,H57,L0,L74,L78

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 based on scoring auctions that give contractors explicit time incentives. We characterize equilibrium bidding and efficient design of these contracts. We then gather an extensive data set of highway repair projects awarded by the California Department of Transportation between 2003 and 2008 that includes both innovative and standard contracts. Comparing similar contracts in which the innovative design was and was not used, we show that the welfare gains to commuters from quicker completion substantially exceeded the increase in the winning bid. Having argued that the current policy is effective, we then develop a structural econometric model that endogenizes participation and bidding to examine counterfactual policies. Our estimates suggest that while the current policy raised commuter surplus relative to the contractor's costs by \$359M (6.8% of the total contract value), the optimal policy would raise it by \$1.52B (29%).

Patrick Bajari
Professor of Economics
University of Minnesota
4-101 Hanson Hall
1925 4th Street South
Minneapolis, MN 55455
and NBER
bajari@econ.umn.edu

Gregory Lewis
Department of Economics
Harvard University
125 Littauer Center
1805 Cambridge Street
Cambridge, MA 02138
and NBER
glewis@fas.harvard.edu

1 Introduction

Public sector procurement accounts for approximately 15 percent of world output. Designing efficient procurement mechanisms is therefore essential for guaranteeing the efficient allocation of many 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 makes for a good case study, as it generates significant negative externalities for commuters through increased gridlock and commuting times. For example, US 101 is an important highway through Silicon Valley, 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 implies 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 California are given relatively generous deadlines, and even then are only penalized with damages of up to \$40,000 per day late. 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 contract 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 winning contractor.

In this paper, we evaluate this new contract design both theoretically and empirically. We start by building a simple 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. Intuition suggests that the right policy is to “tax” contractors

for each day they take, thereby forcing them to internalize the externality. We show that the time weight in the A+B scoring rule acts like a tax, and so this contract design can achieve efficient outcomes. By contrast, the standard contract design amounts to a “soft quota” policy, and will generally result in an inefficiently slow project completion time.

Armed with these insights, we move on to the heart of the paper: an analysis of over 1300 contracts awarded by the California Department of Transportation between 2003 and 2008. Our data includes detailed information on the contract provisions and award procedures, summaries of the bids and bidders, time taken during construction, and damages charged for late completion. Examining the A+B bid data, we see that contractors on average bid to complete the work much quicker than the plans require, around 60% of the engineer’s days estimate. From the payment data, we see that they almost always follow through and complete on time, whether the target is one they bid on themselves (A+B), or one given by the design engineer (standard). So by multiplying a project-specific estimate of the daily social cost to commuters by the days saved by accelerating the target date, we get an average reduction in the negative externality to commuters of \$6.4M per contract. Using a reasonably detailed understanding of the A+B assignment rule, we then construct a control group of standard contracts for the A+B contracts, and compare their winning bids using both regression and matching estimators. The winning bid is on average \$1.5M higher in A+B contracts, so the policy raises commuter welfare by substantially more than the change in procurement costs.

Next, we estimate the additional costs accrued by contractors when accelerating construction. The theory implies that bidders equate the weight on days in the scoring rule with their marginal costs of acceleration. By manipulating this first order condition and making a parametric assumption about the form of the acceleration costs, we can back out the shape of the cost curve. Taking advantage of policy variation in the way the weights are assigned across districts, we estimate this curve and find that it is sharply convex. This implies that small reductions in completion time are reasonably cheap, but large reductions are expensive and require far bigger incentives. So a relatively cost effective policy for Caltrans would be to award more contracts by A+B auction than is current practice, but give weaker time incentives.

To test this theory, we estimate auxiliary models of contractor participation and win probability, and use these models together to examine counterfactual scenarios. We find that the fully efficient policy raises commuter welfare relative to contractor costs by a large amount,

nearly 29% of the total contract value (for these contracts, \$1.5B). Moreover, since acceleration costs to contractors are passed through to Caltrans, which may face budget constraints, we show that a policy with smaller time incentives could achieve most of these gains (\$1.2B) without higher total acceleration costs than under the current policy. This motivates our main conclusion, which is that including stronger time incentives in highway procurement through more sophisticated contract design would substantially raise commuter surplus relative to contractor costs.

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)), as well as a recent report from the National Cooperative Highway Research Program (Fick, Cackler, Trost and Vanzler 2010). These papers take more of a descriptive approach than we do here.

The second is the large theoretical literature on regulation and optimal procurement (see for example Weitzman (1974), Laffont and Tirole (1987), Manelli and Vincent (1995), Branco (1997)). In analyzing A+B auctions, we follow the existing literature on scoring auctions starting with Che (1993) and extended 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).

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 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), Marion (2007), De Silva, Dunne, Kankanamge and Kosmopoulou (2008), Einav and Esponda (2008), Gil and Marion (2009) and Lewis and Bajari (2010)).

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 Highway Procurement in California

Highway procurement in California takes place in three phases, as illustrated by Figure 1. Once a need for construction has been identified, Caltrans designs the project: a set of plans embodies this design. The design engineer will also develop an estimate of the project cost (“the engineer’s estimate”) and a target number of working days for project completion (“the engineer’s days estimate”). Based on advice from the traffic operations unit, the design will specify a maximum number of lanes that can be closed at each phase of the project, and during which hours of the day closures may occur. Last, the engineer will make a recommendation as to whether a standard or A+B contract design should be used, usually based on the size of the project and the projected negative externality (given the lane closure schedule). This recommendation goes for approval to headquarters.

Once a decision has been made, the terms of the contract are summarized in a set of special provisions. In a standard design, the provisions specify that the contract will be awarded to the lowest responsive bidder, and the winning firm will have to complete the contract within the engineer’s days estimate. In the A+B design, the contract is awarded based on a scoring rule, and the winning firm must complete the contract within the number of days they bid.

At this point the project is advertised, and interested parties can obtain copies of the plans and special provisions. In California, bidders are pre-qualified by having a contractor license for that kind of work, and by submitting a bond equal to 10% of their bid. This bond will be forfeited if they win the contract and then don’t complete it (alternately the bond issuer, typically a third party, may pay someone else to complete the work). Next, the contractors bid on the contract, according to the bidding rules laid out in the special provisions. In the case of A+B contracts, if the days bid seems unrealistic, the winning contractor may be required to provide a convincing project timeline for on-time completion, otherwise Caltrans may seek to disqualify their bid.¹

Once the contract is awarded, the construction phase starts. The contractor must plan how to structure the various distinct activities, such as excavation or grading, that make up the construction project, operating within the constraints of the lane closure schedule. They determine how long each activity will take for a standard crew size, and then apply the “critical path method” (Clough, Sears and Sears 2005) to devise the optimal sequence to complete the activities in. The key feature of this technique is that some activities are

¹Generally speaking, Caltrans avoids disqualification where possible, as it can lead to costly lawsuits.




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

designated as critical, and must be completed on time to avoid delay, while others are off the critical path and have some time slack. If the derived plan allows for timely completion plus some contingency time (usually 5%) then it is used; otherwise the contractor will have to change the planned inputs and iterate. For A+B contracts, the construction plan may have to be worked out before bidding.

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. He also designates days on which environmental conditions make work difficult as “weather days”, and as “other days” those on which the contractor cannot work through no fault of his own (e.g. a general strike). These days do not count towards the contract deadline. Still, productivity shocks may affect the rate at which any activity is completed, and contractors check progress against the planned time path and amend plans if necessary. At the end of the process, the contractor is paid the amount bid less any damages assessed for late completion. The daily damages are equal to the usercost in A+B contracts, and equal to the liquidated damage rate in standard contracts:²

$$\text{Liquidated Damages} = \left(\text{LD}\% \times \frac{\text{Engineer's Estimate}}{\text{Engineer's Days Estimate}} \right) + \text{RE Office Expenses} \quad (1)$$

The LD% depends on the type and size of contract, and ranges from 10% to 20%. The home office expenses are estimates of office rental and related costs for the resident engineer.

3 Theory

Before turning to the data, it will be useful to specify a simple model to frame the empirical analysis. Essentially, the problem faced by Caltrans is twofold: On the one hand, they

²The engineer’s days term excludes days on which planting is done on the highway (“plant establishment days”). A new formula was adopted in 2009, outside of our sample period.

want to award highway contracts to the most cost effective or efficient bidder. For this, an auction mechanism is the preferred approach. On the other hand, they want to regulate an externality — the delay to commuters caused by highway construction — when they don’t know how expensive it is for contractors to accelerate construction. If the available policy tools were either a quota or a tax on the time to completion, this would correspond to the problem analyzed in Weitzman (1974).

In practice, these two problems are solved jointly by first specifying a set of time incentives as part of the contract, and then awarding that contract, typically by auction. We refer to the combination of time incentives plus award mechanism as a contract design. To analyze the theoretical performance of the A+B and standard contract designs, we now set up a model of an A+B auction that subsumes the standard design as a special case:

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 received by the winning contractor, and the contract days $d^B \in [0, d^E]$. 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 $c_U \geq 0$ in the scoring rule is known as the *user cost*. The contract also specifies ex-post time incentives: a per day incentive $c_I \geq 0$ and disincentive $c_D > 0$ that are applied when the winning contractor completes the job before or after the contract days. We restrict our analysis to the case $c_I \leq c_D$, which holds for all of the contracts we examine. The three parameters (c_U, c_I, c_D) define the incentive structure.

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

$$\pi(b, d^B, d^T; \theta) = b + 1(d^B > d^T)(d^B - d^T)c_I - 1(d^B < d^T)(d^T - d^B)c_D - c(d^T; \theta) \quad (2)$$

where d^T is the actual days taken to complete the contract and $c(d^T; \theta)$ are the costs incurred. Profit is bid plus incentive payments (possibly negative) less costs.

The cost function $c(d^T; \theta)$ is “long-run”, in the sense that it represents the cost of completing the contract in d^T days, for a contractor of type θ , *given optimal input choices*.³ We assume a textbook long-run cost curve, U-shaped in d for all θ , with a minimum at the efficient scale

³Once construction starts and inputs are hired, the contractor will face a “short-run” cost function $c(d; \theta, K)$ for fixed inputs $K = K(d; \theta)$, but this is not the relevant curve at the bidding stage.

of construction (which is presumably close to the engineer’s estimate). Contract acceleration is costly because it requires working at an inefficiently large scale. Formally, the function is assumed twice continuously differentiable and strictly convex in d for all θ . The type θ reflects contractor-specific cost parameters, such as their expertise with working on a tight schedule, their relationships with subcontractors and input suppliers, and their current managerial capacity.⁴ We assume that each contractor i draws their cost parameter θ_i from some contractor-specific distribution F_i with common compact support $\Theta \subset \mathbb{R}^k$ for all i . This is an (asymmetric) independent private values (IPV) framework.

Strategies and Equilibrium: A (Bayes-Nash) equilibrium of the game comprises a set of bidding strategies $(\beta_1(\theta), \dots, \beta_n(\theta))$ of the form $\beta_i(\theta) = (b_i(\theta), d_i^B(\theta))$, that are mutual best-responses; and a completion time $d^T(d^B; \theta)$.

Social Welfare and Efficiency: Social welfare is given by $W(d^T; \theta) = V - c(d^T; \theta) - d^T c_S$. 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 completion time $d^T(d^B; \theta)$ is welfare maximizing for all types θ . We say that a contract design is *ex-ante efficient* if the winning bidder is always the bidder who generates the highest social welfare $W(d^T; \theta)$ in equilibrium. These notions decouple regulating the winning bidder (ex-post efficiency) from choosing that bidder (ex-ante efficiency).

Standard Contracts: The standard contract design is a special case.⁵ The contract is awarded solely on the bid amount, and the design engineer sets d^E . There are no positive incentives ($c_I = 0$), and the disincentives c_D are called liquidated damages, determined by the formula (1). This corresponds to our model with the constraint that bidders bid $d^B = d^E$.

We proceed in two parts. First we analyze how effective the two designs are in regulating the externality, and then we look at whether the auction mechanism is ex-ante efficient.

Efficient Regulation: The standard contract design is ex-post inefficient. To see this,

⁴For example, the construction company C.C. Myers repaired an important ramp between I80 and I580 near San Francisco after an explosion, doing so in 17 days as opposed to the 50 estimated. In an interview, the owner mentioned as important factors for the speedy completion a collaboration with a steel fabricator to get girders made quickly, an ambitious plan that allowed for work to be done while still waiting for some of the inputs, and his crews working all day in 12-hour shifts (Pogash 2007).

⁵We omit analysis of two other important contract designs with time incentives: lane rental and incentive/disincentive contracts. They are also special cases of our general model - details available on request.




Figure 2: **Completion Time in Standard and A+B Contracts.** Both panels show the marginal private cost of acceleration curve(s), $-c'(d; \theta)$. The left panel shows a standard contract, with damages charged after the specified completion time d^E . The optimal completion times vary with the type θ , but all types choose to complete at d^E , which is inefficient. In the right panel, the contractor bids to complete where the usercost intersects his acceleration cost curve, and completes on time because of the incentive structure. If the usercost equals the social cost, this is efficient.

consider the left panel of Figure 2. The downward-sloping curves are three possible marginal private cost of acceleration curves, corresponding to three different types for the winning contractor. As depicted, all types have positive costs of acceleration at the engineer's days estimate, and all would prefer to operate at a lower scale and complete more slowly. But they are penalized at a rate of c_D when late, giving them an incentive for on-time completion. By contrast, there is no bonus for being early, which implies the time incentives are discontinuous at d^E . This design effectively amounts to a quota of d^E days, where the penalty for exceeding the quota is c_D . Facing these incentives, all three types will complete exactly on time. But since efficiency requires that different types complete at different times to equate their private marginal costs and benefits of acceleration, this design is inefficient.

Instead, the contractor should be forced to internalize the externality imposed on commuters by construction. The simplest way to do this is a tax of c_S per day. As it turns out, the A+B design looks like a quota, but behaves like a tax. Look at the right panel of Figure 2. There, $c_I < c_D$ so that the contractor faces discontinuous incentives at d^B , like a quota. But the contractor can choose d^B in the auction. The scoring rule specifies the trade-off: each extra day bid forces the contractor to lower their bid by c_U to maintain the same score. This essentially amounts to a tax on each day of c_U . Not surprisingly then, if the usercost is set equal to the social cost c_S , this creates the right incentives for ex-post efficiency.

The only complication is that the contractor is not bound by the days bid, except insofar

as the contract enforces compliance. Whenever $c_D < c_U$, the *effective* tax is only c_D , and so the contractor will bid zero days but complete where the marginal cost of acceleration equals c_D .⁶ This design detail is important in practice: in some of the first A+B contracts they let, Caltrans set $c_D < c_U$, and the contractors bid just one day. As a consequence, in all the recent contracts they have issued, Caltrans has set $c_I < c_U = c_D$. Overall, for ex-post efficiency in the A+B design, it is sufficient that the effective cost of delay is c_S , which is best accomplished by setting $c_U = c_S$ and $c_I \leq c_U \leq c_D$.

Ex-Ante Efficiency: The next question is whether the auction mechanism awards the contract to the contractor who will maximize social welfare. As Asker and Cantillon (2008b) show, in a scoring auction the equilibrium strategies of bidders are functions of their pseudo-costs, defined here as:

$$P_i(\theta) = c(d^T, \theta_i) + IP(d^T, d^B) + c_U d^B \quad (3)$$

where d^B and d^T are optimally chosen in the manner described above, and $IP(d^T, d^B)$ denotes any incentive payments that will be paid or charged. The pseudo-cost has three parts: contractor costs given the construction plan; the incentive payments received under that plan; and the B-score $c_U d^B$. This one-dimensional pseudo-cost is a sufficient statistic for the bidder's type θ .

Having transformed the types into pseudo-costs, we can then apply standard results from the auction literature. Provided the bidders are symmetric, the auction mechanism will award the contract to the bidder with the lowest pseudo cost. Now, if the contract is ex-post efficient, the pseudo-cost is equal to the social cost of the project, and so awarding it on the basis of pseudo-cost yields ex-ante efficiency. More generally, even when the contract design is not ex-post efficient, it may still be best to award based on pseudo-cost. This is true of standard contracts if incentives are appropriately set. Suppose, for example, that all bidders have positive benefits of delay up until the engineer's target date. Then since there are no bonuses for being early, all contractors will complete either on-time or late. By setting $c_D = c_S$, the procurer can force contractors to internalize the social costs of late completion as part of their pseudo-cost. Then awarding the contract on the basis of pseudo-cost selects the contractor who maximizes welfare subject to the constraint that they cannot finish early — a second best result — and this is ex-ante efficient.

In summary, the problem of regulating the externality and awarding the contract can be

⁶Bidding to complete one day earlier than actually planned earns c_U from the scoring rule, and costs c_D in damages, so is always desirable with those incentives, which implies a corner solution of zero days.

largely decoupled. By taxing the externality through a weight in the scoring rule, the contract design is made ex-post efficient. By awarding the contract by auction, we achieve the usual gains from competition among bidders.

Cost of time incentives: Later in the paper, we consider how much the use of A+B rather than standard contracts costs Caltrans in practice. It would be nice to have a theoretical prediction to take to the data. Our intuition is that it should be quite costly, as high powered time incentives provide incentives for contract acceleration, and the acceleration costs should be passed through in the form of higher bids. But without knowing more about the specifics of the pseudo-cost distribution, the pass through rate could be anything from below zero to well over one. To see this, note that the first order condition for bidding is

$$b_i(\theta) = \underbrace{c(d^T, \theta_i)}_{\text{costs}} + \underbrace{IP(d^T, d^B)}_{\text{anticipated incentives}} + \underbrace{\frac{1}{\sum_{j \neq i} h_j(s)}}_{\text{markup}} \quad (4)$$

where $h_j(s)$ is the hazard function of the distribution of scores submitted by bidder j . Under the assumption that $c_I < c_U < c_D$, contractors will finish on time, and $IP(d^T, d^B) = 0$. So consider a switch from a standard to an A+B contract design. If type θ optimally bids and completes at d^T given the usercost c_U , they will experience a cost increase of $c(d^T; \theta) - c(d^E, \theta)$, known to be positive. But the markup term in (4) depends on the full distribution of rival cost functions via their scores, and may go up or down. For example, suppose there were two companies: one that used cheap unskilled labor that could complete the job at a slow rate, but could not cheaply accelerate; and another using more expensive skilled labor, that was good at acceleration. The first company might be able to maintain a high markup in auctions for standard contracts, due to the cost asymmetry, but in the A+B contracts that cost advantage would diminish and its markup would fall. On the other hand if the same company were both cheap and quick, the A+B design would magnify their competitive advantage, and their markup would rise. The conclusion is that the relationship between costs and bids is complex, and it would be unwise to deduce cost changes from bid changes.

Discussion: The model is simple and gives clean intuition, but it abstracts away from some real-world features. First, we assume that contractors choose their target completion time at the time of bidding, and then everything follows according to plan (so that target and actual dates coincide). We also assume that contractors are risk neutral. But events may arise during construction which require adaptation of the construction plans, or cause

unavoidable delays.⁷ This creates risk.

Contractor risk aversion would explain why Caltrans makes provision for delays that are not the contractor’s fault by counting them as weather or other days. It would also explain why contractors build contingency time into their construction plans. Relative to the above model, contractors may choose to operate at a slightly higher scale to “insure” against being late due to an unexpected event. Risk aversion also matters for the bidding. Bigger time incentives lead to greater variation in final payments for the same target date, and though this risk can be partially offset by “insuring” through bigger scale, the bids should pass through not only the increased acceleration costs and markup but also a risk premium. But the main intuition about contract design in terms of taxes and quotas will survive in the more complex model: though the contractor problem is more complex, as long as he is made to internalize the externality, the design will be efficient.

Second, 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 do not model the monitoring costs as part of the welfare function, as these costs are typically very small. Finally, we do not model the restrictions on lane closure that are typically part of the contract design: these are taken as exogenous. The social cost c_S should therefore be interpreted as the daily externality when these restrictions are enforced.

4 Empirical Analysis

The theory suggests how contracts should be designed in order to maximize social welfare. In the remainder of the paper, we examine the A+B contract design in practice, using data from contracts let by Caltrans during the period 2003-2008. Our dataset is rich, including detailed data on the contract provisions, bids and outcomes for both standard and A+B contracts. This enables us to answer two practical questions: Do A+B contracts work well in practice? And could Caltrans adopt better policies? We answer the first using the language and methodology of the policy evaluation literature. For the second we build, estimate and simulate counterfactual outcomes using a structural model of contractor behavior.

⁷In Lewis and Bajari (2010) we consider a richer model which includes a construction phase after the bidding in which ex-ante unknown state variables are realized.

4.1 Data

To construct our dataset, we downloaded publicly available files from Caltrans. There are three main kinds of data: special provisions, bid summary and final payments. Special provisions data were collected for all non-emergency contracts let in the sample period.⁸ The special provisions detail the amendments to the standard construction contract terms, and are usually a pdf file of approximately 200 pages. From these files we extracted the provisions relating to time and liquidated damages by text parsing. These allowed us to classify the contracts as either A+B contracts, standard contracts, or some other type. A different section described the lane closure restrictions. Since these are depicted in graphical form, we had to hand-code this data, which we did for all A+B-eligible contracts (see below for the definition of “eligible”), and for a stratified random sample of 10% of the remaining standard contracts. We imputed the data for the remaining contracts.⁹ From this data we constructed two variables. The first is “lane closure ratio” which is the ratio of the number of lanes the contractor may close (averaged across hours of the day) to the total lanes of the highway. The second is “reopening penalty”, a dummy for whether an explicit monetary penalty is specified in the contract for late lane reopening.¹⁰

Bid summary data includes the location and description of each project, the design engineer’s estimate of the project cost, the required project duration, and bidder names, addresses and bids.¹¹ It also includes the funding source, so we can classify a project as federal or not. We use this to construct a measure of the straight-line distance between the contractor’s address and the project location. Last, we collected final payment data, which includes all payments made by Caltrans to the winning contractor during the life of a project.¹² This data only encompasses projects completed during the sample period. The final payment files also give a breakdown of the working days, weather days, other days and contract change order days over the course of the project, as well as the contract start and end date.¹³

We augmented this data with data on traffic volumes and the percentage of trucks at each

⁸Data available here: http://www.dot.ca.gov/hq/esc/oe/project_ads_addenda/

⁹The imputed data is used only in the counterfactual simulations.

¹⁰For example, if all lanes should be open by 8am, and the contractor is still working at 8:15, the contractor will be charged a significant penalty, on average \$4,000 per 15 minutes. In contracts without such a penalty, the PE can suspend the contract instead.

¹¹Data available at: <http://www.dot.ca.gov/hq/esc/oe/planholders/bidsum.php>

¹²Data available at: <http://www.dot.ca.gov/hq/asc/oap/payments/public/ctnums.htm>

¹³Contract change order days are additional days allowed when the project design is amended during the course of the project by means of a change order.

of the locations.¹⁴ This enabled us to construct a simple measure of the negative externality to commuters during project construction. Our formula for the social cost is:

$$\text{Social Cost}_t = \text{Delay}_t \times \text{Time Value}_t \times \text{Traffic}_t \quad (5)$$

For time value, we use a weighted average of a value for cars (\$12 per vehicle) and for trucks (\$28 per vehicle). These are the values used by Caltrans in their own calculations in 2008. To estimate the delay, we assume that traffic will be slowed from 55MPH (the speed limit on a typical California freeway) to 35MPH over the full length of the construction zone, or five miles, whichever is shorter. In most construction zones, a slowdown of at least this magnitude is to be expected either due to actual lane closures or to traffic control. These assumptions are chosen to be conservative, and imply that we will never project a delay of more than 3:11 minutes per commuter on any project.¹⁵ The social cost is then arrived at by multiplying the traffic, delay and average time value.

The Caltrans traffic operations divisions in each district also calculate social costs, for the purposes of deciding on A+B assignment and to set the user cost in the A+B scoring rule. As far as we can tell, the methodology varies across districts and individual employees, and consequently takes into account different possible sources of delay.¹⁶ The most sophisticated calculations go beyond the direct delays we consider, also allowing for increased vehicle operating costs (due to delays) and for queuing. Indeed, queues can often lead to massive social costs and so the lane closure restrictions are often designed precisely to ensure that lane closures occur only at times of day when they are not likely to result in a queue.

From the raw data we removed contracts where either the special provisions or bid summaries were missing; contracts which were neither standard nor A+B contracts; and bids with highly implausible ratios of bid to engineer's estimate (less than 20% or over 300%). We also chose to focus on a subset of contracts which seemed *a priori* to potentially require lane or shoulder closure: barrier construction, bridge repair or resurfacing, new lane or ramp construction, road rehabilitation and widening/realignment.¹⁷ The A+B design is very rarely used outside

¹⁴We used the 2008 measures at: <http://traffic-counts.dot.ca.gov/index.htm>

¹⁵We also used Google Maps to calculate a detour route for a typical commuter around the construction zone, but in all cases the detour caused a delay of more than 3.11 minutes and so this data was discarded.

¹⁶We have seen three quite different calculations: one similar to ours, initially developed by headquarters in 2004; the more sophisticated spreadsheet described in the main text, developed by district 10 in 2007; and an actual calculation for a contract in district 12, which is based on assessing the delay due to a temporary reduction in roadway capacity.

¹⁷Omitted categories: cleaning, drainage, plant establishment, retrofitting, sealing, signs and slope recon-

of these categories. Likewise, we restricted the analysis to those Caltrans districts that use the A+B design with some frequency (they are SF Bay Area (4), Fresno (6), Riverside/San Bernadino (8), San Diego (11) and Orange County (12)). Of these, district 4 is by far the most frequent user of the format, and plays an important role in our analysis.

4.2 Descriptive Evidence

Table 1 summarizes many of the contract characteristics. Data from our chosen districts comprises 603 standard and 79 A+B contracts, which we compare in the first two columns. A+B contracts are typically much bigger, in terms of both cost and time required, have higher traffic counts and higher estimated social costs. We also see that A+B contracts are much more likely that to include an explicit monetary penalty for late lane re-opening. The standard contracts in the omitted districts are typically smaller in both size and potential welfare loss, and less than 1% of them are A+B.¹⁸

Finally, our social cost estimates are on average much higher than the usercosts for the A+B contracts. This is a little surprising, as Caltrans policy is to set the usercost equal to their calculated social cost estimate, and our social cost estimates should be more conservative than theirs. We are not sure how to account for the difference: it could be that Caltrans assumes no negative externality during hours when lane closures are not permitted (though we have been told this is not the case), or that sometimes they set the usercost below their social cost estimate to avoid having the B-score be “too large” relative to the contract size.¹⁹ Given how conservative our calculations have been, however, we will proceed under the assumption that they are a lower bound on the true social costs.

The differences in characteristics between standard and A+B contracts suggests non-random assignment. To learn about the assignment mechanism, we asked the officials in a number of districts how they decided on the A+B contracts. All of the districts except district 4 said they followed a rule mandated by headquarters, under which a project should be assigned A+B status if the engineer’s estimate was over \$5M and it had an estimated daily social cost of over \$5,000. District 4, on the other hand, indicated they used the A+B design

struction/protection.

¹⁸The LA district is unusual in that contracts are big and lane closures cause large delays, but the A+B design is not used. Instead incentive/disincentive designs are used on time sensitive segments of projects.

¹⁹Some Caltrans engineers expressed concern that large time incentives would push the contractors to genuinely accelerate construction, rather than just trim excess fat from their construction schedules.

whenever the engineer’s days estimated exceeded 100 days. In all districts, there was room for exceptions, for example to use the design when there were other political / economic issues that made timely completion important, and to disqualify the design when the contract had the potential for third party problems.²⁰

To test this in the data, we first constructed a set of size eligible contracts, consisting of those that were either over \$5M (outside of district 4) or over 100 days (in district 4). We then estimated a probit for A+B assignment, using a small set of RHS covariates. The results are shown in Table 2, split out for district 4, other districts and combined. Looking at the combined column, we see that the average marginal effect on traffic is highly significant and large: an additional 10,000 traffic increases the probability of A+B assignment by approximately 1.4%. District 4 is also far more likely to assign the A+B design for a given contract size (consistent with the fact that it has no social cost threshold), and similarly bridge resurfacing contracts are more likely to be A+B, though this is only significant at 10%.²¹ Using a simple rule that predicts A+B assignment only for eligible contracts with predicted A+B probabilities over 50%, we can correctly predict nearly 92% of the assignments. The fact that we can do so well with a very simple model and five covariates suggests that the districts acted largely in accordance with their policies.

Similarly, we wanted to understand how the usercosts were chosen for the A+B contracts. Again, most districts set the usercost equal to the estimated social cost, as suggested by headquarters. But district 4 says they typically set the usercost using the liquidated damage formula of equation (1), rather than basing it on a social cost calculation. Table 3 shows the results of a regression of log usercost on covariates, where the log specification is chosen because of the ratio form of the liquidated damage formula. For district 4, we can explain 80% of the variation, again suggesting that their description is accurate. We are less successful for the other districts (R^2 of 0.42), but this is probably not surprising given the limited amount of data and the variation in methodology for calculating social costs across districts.

²⁰Third party problems arise, for example, when a utility company needs to evacuate a piece of land during some phase of the construction. Then if the company is late in fulfilling its obligations, the contractor cannot proceed and can claim compensation from Caltrans for the delay. This is presumably more problematic when the contractor is on a tight schedule.

²¹We have tried additional controls for other kinds of work, but none of them are significant.

4.3 Policy Evaluation

We now turn to the question of whether the A+B design works well in practice. To best take advantage of our data, we proceed in two parts. First, we ask whether the A+B design accelerates project delivery at reasonable cost *assuming that bidders fulfill the contract terms*. This part of the analysis uses all the data. Then we look to see whether in fact ex-post behavior by contractors differs across A+B and standard contracts in terms of actual delivery, contract violations etc. This second part uses only the data from completed contracts.

Does the A+B design incentivize bidders to commit to finish the project faster? Recall that under the standard design, the project deadline is set by the design engineer; whereas in A+B the bidders may commit to any time faster than that. So a sensible measure of acceleration is the ratio of the winning days bid to the engineer’s days estimate. In the A+B contracts the average ratio is 59.5%!²² To put this into dollar terms, we calculate a welfare gain to commuters as the days accelerated (engineer’s estimate less winning bid) times the daily social cost estimate. This comes out to an average of \$6.46M per contract.²³

We know already that the A+B contracts are a selected sample of the full set of contracts, and so it would be inappropriate to make inferences for the whole sample based on these outcomes. Viewed through the lens of the treatment effects literature, what we see is an average treatment effect on the treated (ATT). The A+B “treatment” caused an acceleration of around 40% in the target completion time for these contracts.²⁴ What we would like to know is what this costs Caltrans.

To answer this question, we need a consistent estimate of how much the winning bid on each A+B contract would have been had it been a standard contract. We construct a control group consisting of all eligible standard contracts. Since we don’t have good control contracts for bridge resurfacing (almost all of them are A+B), and sparse data above the 95th percentile of contract size (\$90M +) we omit these from the subsequent analysis. In Table 4 we show that this control group is well balanced relative to the A+B group on most observables (including types of work). The statistically significant differences relate to

²²The median is very similar, 60.0%, so this is not driven by a few outliers.

²³We assume throughout that commuters are not delayed on days where no construction occurs, such as weekends, weather and other days. If delays also occur on these non-working days, our estimates of the welfare gain will be too low by a factor of between 1.6 and 1.8, since each working day attracts between 0.6 and 0.8 non-working days (depending on contract size).

²⁴The causal language is appropriate here since we are talking about commitments, not outcomes: in the absence of the treatment, the commitment would have been to complete in the engineer’s days.

traffic, lane closures and re-opening penalties, which we will control for in what follows.

We can now look at whether these A+B contracts cost more than they would have had they been standard contracts. We estimate this counterfactual cost in two ways. First, we run an OLS regression of the log winning bid on a variety of covariates, including A+B status, for the subsample consisting only of the treatment and control groups.²⁵ A log-linear specification seems appropriate since one would expect errors to scale with contract size. We control for contract size, traffic, lane closures and re-opening penalties, type of work, year and district. All these control variables are pre-determined at the time of A+B assignment.

The regression results are shown in Table 5. In all columns, we see a significant positive coefficient on the A+B dummy, implying that these contracts cost 7-11% more than if the design had been standard.²⁶ The coefficient increases as we add fixed effects and additional controls. Notice also that traffic and lane closures do not enter significantly, which is reassuring as it suggests that the factors important for assignment may be orthogonal to the contractor's costs. The fit is very good, as seen from the high R^2 terms. This is because the engineer's estimate is a superb control for contract costs.

Our second approach is a non-parametric matching estimator. We match each A+B contract to the standard contract that is closest on the observables, and compare the winning bids.²⁷ We try to force exact matches on the categorical variables district, type of work and year where possible and then within those cells to match on all the covariates from columns (1) and (2) of Table 5. We are able to exactly match 66% of the time. Averaging the differences in bids, we get a sample estimate of the average effect on the treated (ATT).

To facilitate direct comparison, we also convert the percentage costs from the OLS regressions into an average dollar effect, and summarize all the results in Table 6. The OLS estimates indicate an increase in the winning bid of between \$1.23M and \$1.77M (our preferred specification). The matching estimator yields an increase of between \$1.63M and \$2.13M, albeit with much higher standard errors. Combined with the commuter welfare estimate from earlier, we estimate that the difference between the gain to commuters and the higher payments by Caltrans is between \$4.8M and \$5.8M per contract. This compares favorably with the

²⁵We don't use the full sample because of the well-known problems with enforcing linearity over parts of the covariate space in which there is no overlap of treatment and control — running the full sample regression leads to the odd result that A+B contracts do not cost more.

²⁶Strong, Tometich and Raadt (2005) find a similar cost increase of 7.5% for A+B contracts in Minnesota.

²⁷We also perform the linear bias-correction of Abadie, Drukker, Herr and Imbens (2001) since even small differences in average contract size across the treatment and control might otherwise bias the results.

average A+B contract size of \$23M. Although this is not a social welfare measure — since additional contractor costs are unobserved, and bid changes may be a poor proxy for reasons discussed earlier — it still suggests that the policy provided considerable value to taxpayers.

Discussion: We conclude this subsection with some comments regarding the interpretation of the results. First, this is an estimated effect on the treated: these contracts were selected by Caltrans on the basis of being time sensitive projects, and we wouldn't expect a \$5M welfare gain for a random highway contract (as we will see in the structural results). Second, these results could be biased by selection on unobservables. This seems less of a concern than it might be in other settings since we have a good understanding of what drives A+B assignment and can argue that the omitted factors have limited impact. For example, third party issues are often cited as a reason that an eligible contract was not assigned A+B status, and we do not control for them in the evaluation. But as long as the third party issues don't impact contractor costs *when operating on a standard schedule* — and they shouldn't, by design — they will not bias the results. In addition, we have an excellent control variable in the design engineer's estimate. For these reasons, we may be relatively confident of our estimate of the cost increase and welfare gain.

4.4 Comparing Ex-Post Performance

Until now, we have assumed that A+B and standard contracts perform similarly after the contract is signed. But one might be concerned that contractors are less likely to meet their obligations in A+B contracts, for example by shirking on quality, completing late, or by convincing project engineers to award them unnecessary extra weather days. Table 7 summarizes the ex-post performance of treatment and control contracts. A+B contracts are typically shorter in duration (because days are bid down), finish in about the time bid, get slightly more weather days, CCO days and other days, finish late slightly more often, and get slightly higher deductions. Deductions are charged whenever there is some form of contract violation, such as late completion or failure to meet specifications. They thus shed some light on the question of whether contractors are more likely to violate contract terms on A+B contracts by performing sub-standard work, or hiring illegal labor etc.²⁸

²⁸This approach was also taken in Fick et al. (2010), where they looked at a sample of 455 standard and 22 contracts with time incentives from a particular state highway agency in 2007, and compared quality deductions and quality adjustments. They found no statistically significant differences.

Are any of these differences significant after conditioning on observables? To test this, we regress the various outcome measures on covariates and A+B status for the set of treatment and control contracts. The results in Table 8 show that A+B contracts actually take *fewer* working days, weather days and total days conditional on observables, though not significantly so. And though deductions are higher, this too is not significant, nor is the estimated coefficient particularly large relative to the average contract size. The data is therefore consistent with the hypothesis that A+B and standard contracts have comparable ex-post performance. This is not surprising: once the contract days have been specified (either by bidding or by the design engineer), the contracts themselves are virtually identical.²⁹ The only substantive difference is that the daily penalty for being late is typically higher in A+B contracts (average of approximately \$12K per day vs \$7K for standard). This higher penalty seems to be sufficient to enforce compliance with the accelerated construction schedule.

5 Counterfactual Policy Analysis

Our analysis thus far suggests that the A+B contract design, as used by Caltrans, has worked well. We now want to examine the trickier question of whether they could do better. To do this, we will need to make some parametric assumptions and estimate a particular structural model. The payoff is that we will be able to examine counterfactual scenarios. In particular, we will compare the outcomes from the observed policy with two policies based on using the A+B design on all contracts (the “tax” approach), and two policies based on eliminating the A+B program and instead just reducing the engineer’s days (the “quota” approach). Our approach will be to focus on modeling the bidding behavior of contractors, since we have seen that ex-post, contractors complete almost exactly in the time they committed to. In what follows we will assume on-time completion.

We systematically proceed through the auction process. First, we estimate a simple entry model in which contractors decide which auctions to participate in on the basis of both their own and the auction characteristics. Next, we model days bids in the A+B auctions, specifying a parametric form for the contractor’s acceleration costs and then exploiting a first order condition from the theory to estimate this curve. Last, we estimate the probability

²⁹In a standard contract, the relevant clause in the special provisions is that “the work shall be diligently prosecuted to completion before the expiration of X days”; whereas the A+B contract demands “diligent completion before the expiration of the number of working days bid.”

of winning the auction conditional on bidder characteristics, residuals from the previous regressions and (in the case of A+B auctions) the B-part of their score. By proceeding in this way, we avoid the selection problems that would result if we only modeled part of the process. Last, we simulate counterfactual outcomes for the observed set of firms and contracts under alternate policies.

5.1 Entry

We use a simple static logit model of entry to simulate the set of bidders who will enter each auction. We differentiate firms on three dimensions: their distance from the contract site, their size, and whether they are in-state or out-of-state. To define firm size, we first compute the total volume of contracts held by the firm over the sample period. We then form a coarse binary measure, “Big Firm”, by classifying firms that held over \$50M worth of work as big. There are 21 such firms, which together account for around 25% of the total volume. We treat the bidder and auction characteristics as exogenous throughout, although in several of the specifications we compare A+B with contracts in the control group (rather than the full sample) to limit concerns of selection bias.

The results are shown in Table 9. Columns (1)-(4) are for only A+B and control group contracts, while column (5) is for the full sample. In all specifications, contract size and distance enter negatively and highly significantly, while big firm and engineer’s days enter positively. Fewer firms participate on bigger contracts, although the interaction between big firm and size is positive. This presumably reflects differential capacity constraints. Firms farther away from the site are also less likely to enter. Big firms participate in more auctions.

There seems to be little evidence that A+B contracts attract more or less participation relative to the control group, with no joint significance in the A+B coefficients in columns (2)-(4). This is the case when we interact the A+B dummy with big firm, distance and in-state in column (3), implying that A+B contracts do not attract relatively more big firms, or closer firms, conditional on the other observables. When we include firm fixed effects to control for firm-specific unobservables, the significance of the A+B coefficients is even smaller. The broad picture that emerges is that matching is essentially based on distance and size, rather than A+B status, and for that reason we use the full sample coefficients in column (5) for the rest of the structural estimation. Of course this does not rule out matching on time-varying unobservables, such as a situation in which firms with excess

capacity (unobserved) are differentially more likely to participate in A+B auctions. We will try to control for this potential selection effect in later regressions, by including the residuals from the regression in (5) as controls.

5.2 Estimating Acceleration Costs

Next, we use the bidding behavior of contractors in A+B auctions to back out acceleration costs. Recall the contractor's cost curve is $c(d; \theta)$. Define the number of days accelerated as $\tilde{d} = d^B - d^E$, and the base cost $\bar{c}(\theta^M) \equiv c(d^E; \theta^M)$ as the cost of completing the project on the design engineer's schedule, where θ^M is the part of the type that determines materials and other base costs. Letting θ^T be the part of the type determining acceleration costs, we define $c_A(\tilde{d}; \theta^T) \equiv c(d^E - \tilde{d}; \theta) - \bar{c}(\theta^M)$ as the cost of completing the project at a faster rate.

This cost curve is latent. But in an A+B contract, the theory tells us that the contractor will bid to equate the marginal cost of acceleration with the usercost. Let bidders be indexed by i , let contracts be indexed by j and let x_{ij} be a vector of bidder and contract covariates. Adopting a polynomial form for marginal costs, we get the following first order condition:

$$c_{U,j} = c'_A(\tilde{d}_{ij}; x_{ij}, \theta_{ij}^T) = \tilde{d}_{ij}^\alpha e^{x_{ij}\beta + \xi_j + \theta_{ij}^T} \quad (6)$$

where ξ_j is a contract-specific unobservable. This specification is simple but flexible. It allows for linear marginal costs as a special case ($\alpha = 1$) but admits both concave and convex marginal costs. It also implies total acceleration costs of $\frac{1}{1+\alpha} c_{U,j} \tilde{d}_{ij}$, which have the attractive property that they are equal to zero if $\tilde{d}_{ij} = 0$ (no acceleration). Taking logs and re-arranging terms, we obtain our main estimating equation:

$$\log \tilde{d}_{ij} = \frac{1}{\alpha} (\log c_{U,j} - x_{ij}\beta - \xi_j - \theta_{ij}^T) \quad (7)$$

We will estimate this by OLS.

Identification: Before estimation it is useful to think through the identification of the model. The argument is sketched in Figure 3. To start, suppose that there were no type variation at all, so that there was a single marginal cost of acceleration $c'_A(\tilde{d})$ to identify. Then since contractors equate their marginal costs with the usercost, variation in usercosts would induce contractors to bid different days. By matching the usercost with the days bid, as shown in the left panel of the figure, we can trace out the marginal cost curve.




Figure 3: **Identification.** The left panel shows the identification argument for a fixed type. Exogenous variation in the usercosts from C_1 to C_3 causes the firm to bid different days, d_1^B to d_3^B . The marked points can be plotted, and with sufficient variation, the entire marginal acceleration cost curve can be traced out. In the right panel, we show how differences in the days bid for a fixed usercost allows us to measure how much heterogeneity there is across firms.

Now, let us allow for heterogeneity in types. In this case, for any given usercost we observe a distribution of days bid, as in the right panel of the Figure. The variance of this distribution is informative as to the degree of heterogeneity. We make two assumptions: first, that marginal costs are everywhere increasing in type, and second that type and usercost are independent of each other. The first condition is purely technical and rules out multiple intersections between the marginal cost curves of different types (they all intersect once at d^E). The second is important, and we return to it below. Then to get the marginal cost curve for the γ -th quantile of the type distribution, we proceed analogously to before: matching usercosts with the γ -th quantile of days bid. All in all, the cost curve is non-parametrically identified over the support of the usercost distribution. As is usual, this argument extends to the case of non-identical contracts and covariates by conditioning.

The key non-parametric identification assumption is that the types and usercosts are independent. Returning to our linear estimating equation, the practical analogue is that the error $\xi_j + \theta_{ij}$ must be uncorrelated with the log usercost and other regressors x . This gives rise to two important endogeneity concerns: that Caltrans sets the usercost based on some contract-level unobservable ξ_j that is also correlated with contractor costs (endogenous policy), and that the set of types θ_{ij} that participate depends on the usercost (endogenous entry). As we have already seen, there does not appear to be any average difference in entry behavior across standard and A+B contracts, which limits fears of endogenous entry. Nonetheless, we will employ a control function approach to address this concern.

The endogenous policy concern is harder to evaluate. From the regressions in Table 3, we are fairly confident that District 4 sets its usercosts based on the liquidated damage formula used statewide. From the formula in equation (1), we see that conditional on the type and size of contract, the liquidated damages co-vary with the design engineer’s estimate of home office expenses, which are likely orthogonal to the contractor’s acceleration costs (especially since it is the *project* engineer who will supervise the work). But although this is plausibly exogenous variation, it is not good variation, since these office expenses are a small part of the liquidated damages.

Our best source of variation is inter-district variation, coming from the differences in policy for setting usercost in district 4 versus in other districts. From Table 3, we know that usercosts are set considerably lower in district 4, conditional on the observables. So we can identify the elasticity of days bid with respect to usercost by comparing the days bid on similar contracts in district 4 versus other districts, accounting for any mean differences in acceleration costs with district fixed effects. This does not eliminate endogeneity concerns, of course. Although we have a good idea of how usercosts are computed in other districts from having seen the relevant spreadsheets for some contracts — they are based on estimates of social costs derived from the construction plans, detailed traffic patterns and lane closures — we cannot be sure that none of the elements entering these calculations is correlated with acceleration costs. But since acceleration costs are basically a function of contractor expertise and available capacity, it seems reasonable to believe that they are orthogonal to social costs, and therefore to the usercosts in the other districts.

Estimation Results: The results of the regressions are reported in Table 10. In the first column we include neither fixed effects nor any firm characteristics, and find a non-significant effect of usercost on days bid. As we proceed through the regressions, successively adding controls for firm characteristics and fixed effects, this coefficient grows to approximately 0.35 in columns (3) and (4) and becomes statistically significant. In column (4), we have included an entry residual from the participation model in column 5 of Table 9, defined as $(1 - \hat{p})$, where \hat{p} is the predicted probability of entry. This does not enter significantly, and has minimal impact on the estimated coefficient on log usercost, which suggests that the endogenous entry concern is not important here.³⁰

The other coefficients have the expected signs. Contract size is significant in every specifica-

³⁰We have experimented with a more flexible control function, implemented as a piecewise linear spline with knots at the quartiles, but the coefficients are not jointly coefficient (p-value 0.7107).

tion, with longer contracts attracting more acceleration and higher estimated costs contracts getting less. This is consistent with a scenario in which contractors are constrained in how much work they can do per day, and so when there’s more work per day (higher cost or fewer days) they offer to accelerate less. Big firms and firms closer to the contract site offer speedier construction, as one would expect. Finally, firms offer more acceleration when they are less constrained by lane closure requirements (higher lane closure ratio), and when they don’t face explicit penalties for re-opening a lane late on any construction day.

Our preferred specification is column (3), with an estimated coefficient on log usercost of 0.348. From equation (7) this coefficient has a structural interpretation as $1/\alpha$, implying α roughly equal to 3, and thus sharply convex marginal costs of acceleration. This is consistent with what we know about the construction process: small reductions in time can be achieved at the same scale by cutting out contingency time, but larger reductions require genuine acceleration and operation at an inefficient scale.

The Type Distribution: So far, we have been able to remain agnostic as to the distributions of ξ_j and θ_{ij} . But for simulation purposes, we need to be able to draw contract shocks and individual-specific types. This requires deconvolution of the error term $\xi_j + \theta_{ij}$. Given our limited data, it does not make sense to take the fully non-parametric approach of Li and Vuong (1998). Instead, we assume that both ξ_j and θ_{ij} are normally distributed, iid across contracts and bidders, and independent of each other, with $\xi_j \sim N(0, \sigma_\xi^2)$ and $\theta_{ij} \sim N(0, \sigma_\theta^2)$. The deconvolution process is relatively straightforward: we estimate σ_θ^2 from the variance of differences in residuals within contracts, and then back out σ_ξ^2 as the difference between the overall variance in residuals and our estimate of σ_θ^2 . We estimate $\sigma_\xi^2 = 0.13$, implying a standard deviation of 36% in the costs of acceleration across contracts. Bidder heterogeneity accounts for much more of the variance, with σ_θ^2 estimated as 0.23, and thus a standard deviation of nearly 50% across bidders for the same contract.

The normal-normal model fits the observed density fairly well, as we will see when we simulate moments from the full structural model.³¹ The independence assumptions are obviously strong, though necessary. In particular, the iid assumption across periods for θ_{ij} implies no persistence in a given firm’s type across contracts. Our bidder observables — big firm, in-state firm and log distance — will allow for some persistence, but ultimately there’s no denying the fact that this is not a dynamic model. We return to this point in the discussion of the counterfactual results.

³¹We have also tried conditional deconvolution, $\sigma_\xi^2 = \sigma_\xi^2(x_j)$, but this doesn’t appear to aid model fit.

5.3 Who wins the auction?

Last, we need a model of who wins the auction conditional on bidder covariates and type, days bid and the weight given to time in the scoring rule. We want to avoid modeling the full auction, because doing this convincingly would take us far afield. Recall that in a scoring auction, bidders compete on the basis of pseudo-costs, and if bidders are symmetric, the bidder with a lowest pseudo-cost wins. Writing the pseudo-cost out in a way that separates the acceleration and base costs, we get:

$$\begin{aligned} P_{ij} &= c_A(d_{ij}^B; \theta_{ij}^T, \xi_j, x_{ij}) + \bar{c}(\theta_{ij}^M, \xi_j, x_{ij}) + c_{U,j}d_{ij}^B \\ &= \gamma_j + \frac{\alpha}{1+\alpha}c_{U,j}d_{ij}^B + x_{ij}\delta + \theta_{ij}^M \end{aligned} \quad (8)$$

where in the second line we use the fact that $c_A(d_{ij}^B; \theta_{ij}^T, \xi_j, x_{ij}) = \frac{1}{1+\alpha}c_{U,j}(d_j^E - d_{ij}^B)$ and sweep all contract-level terms including $\frac{1}{1+\alpha}c_{U,j}d_j^E$ into a fixed effect γ_j . This index form motivates a conditional logit model for the probability of winning, where contract fixed effects capture the term γ_j , and $c_{U,j}$ is set to zero in standard contracts.

There are a few problems with this. First, since bidders are asymmetric, it may not be the case that the bidder with the lowest pseudo-cost wins — but this is a reasonable approximation to the truth. Second, the distribution of θ_{ij}^M almost certainly exhibits heteroskedasticity, with bigger contracts having more variation in base costs across bidders. To deal with this, we normalize the B-score $c_U d_{ij}^B$ by the engineer’s cost estimate. Finally, the base type θ_{ij}^M may be correlated with either the days bid d^B (through correlation with θ_{ij}^T) or the covariates x_{ij} . To address this, we include residuals from both the bidding and entry models as controls in the conditional logit.³²

The results are shown in the first column of Table 11. The first column includes only our bidder observables, and shows that big firms, in-state firms and closer firms are more likely to win, though only significantly so for big firms. In the second column, we include the normalized B-score as a control, and this enters highly significantly and negatively. This is unsurprising since we know that lower scores win the auction. The most interesting results are in the third column. Both residuals enter significantly as controls. The negative coefficient on the entry residual — significant at 10% — implies that firms that were predicted to be unlikely entrants (i.e. high residual) are unlikely to be winners at the end of the day. More

³²Since the days bid regression was estimated using the subsample of A+B contracts, we have no bid residuals for standard contracts, and normalize it to zero.

important for us, the positive coefficient on the bid residual implies a positive correlation between base and acceleration costs. This may be because there is some fixed input that drives both, like managerial capacity. Had we omitted this control, we would have incorrectly concluded that the positive selection of bidders with low acceleration costs was entirely due to the A+B mechanism, rather than it being a combination of the mechanism and this underlying correlation (see the coefficients on normalized B-score in columns (2) and (3)).

From a big picture perspective, positive selection makes A+B auctions more attractive. One might have thought that there was a sharp trade-off between fast and cheap contract completion, with some firms specializing in fast completion and others in cheap completion. In that case the use of A+B contracts would sharply raise base costs relative to the standard design, by favoring the fast but expensive contractors, and dealing with the commuter externality may not actually yield much of a welfare gain. What we have learned from these regressions is that it is the same kinds of firms that participate in both kinds of contract — contrary to a specialization story — and that “fast” and “cheap” are positively correlated. This is good for social welfare. A downside of this is that the best firms — those with spare capacity, say — have increasing market power as the usercost increases, and their markups should reflect this. So while using A+B contracts appears to be particularly good for total welfare given our observations thus far, it may raise procurement costs substantially.

5.4 Counterfactuals

We are now able to model counterfactual changes and simulate their outcomes using our dataset. The important maintained assumption is that the distribution of types and contract shocks estimated above remains constant as we change policies. This requires our earlier assumption of random A+B assignment conditional on the observables, since otherwise we cannot expect the distribution of ξ_j for standard contracts to be the same as we estimated for A+B contracts. We are also assuming random entry conditional on the observables.

We consider four policies. The first two are A+B policies. In the first, the usercost is set equal to the estimated social cost (the efficient policy) and all contracts are assigned A+B status.³³ In the second, the usercost is assigned to be only 10% of the estimated social cost (the “budget” policy). Policies of this type — where the usercosts are a constant

³³We are poorly identified outside of the support of the usercosts. So we bound the usercost above at the engineer’s estimate / engineer’s days, so the B-score can be no bigger than the engineer’s estimate.

fraction of the social cost — are constrained efficient when the constraint bounds the total increase in contractor acceleration costs over all contracts. This second policy is motivated by the observation that since acceleration costs are so convex, Caltrans can get more “bang for buck” without increasing overall contractor costs — and therefore hopefully without increasing Caltrans costs — by using A+B contracts frequently but with small incentives.³⁴

The last two policies are “quota” policies, where Caltrans elects to reduce the engineer’s days by 10% uniformly to accelerate contract completion, rather than using the A+B design. The theory suggests this should be inefficient and perform worse, but this policy is much simpler and so it is interesting to know how much worse it would do.³⁵ We consider a “hard” quota, where the penalty for violating it is essentially infinite, and so contractors will certainly complete on time; and a “soft” quota, where the current liquidated damages rates continue to apply for being late.

The simulation procedure is as follows. Fix a policy. Then for each contract, randomly draw a set of participants using the estimated logit model, and independently draw contract shocks ξ_j and types θ_{ij} for each participant from the estimated normal distributions.³⁶ Then using the bidding model, determine the level of acceleration that equates marginal benefits and costs for each participant, and work out the associated total acceleration costs.³⁷ Finally, work out the win probability for each contractor using the conditional logit model, and average across participants to get expected acceleration and costs.³⁸ We repeat this simulation process 1000 times, and report the average outcomes across all simulations.³⁹ To get confidence intervals for the outcomes, we estimate the bidding, type distributions and win probability regressions exactly as before using a sequence of 100 bootstrap samples and then run the simulation as before using the bootstrapped coefficients.⁴⁰ We choose not to

³⁴The “small incentives” approach is also attractive because it leaves less room for a firm with an advantage on both base and acceleration costs to exert market power.

³⁵We thank Justin Marion for making this suggestion.

³⁶Whenever the entry model randomly draws zero entrants, we force the draw of a single bidder, where each firm has chance proportional to their probability of participation.

³⁷We constrain the days bid to be no less than 20% of the engineer’s days — essentially assuming infinite marginal costs at that point — since we never observe bids below this threshold in the data.

³⁸In quota contracts, we must account for differences in acceleration costs across bidders when calculating win probabilities. From equation (8) our estimated coefficient on the normalized B-score in the conditional logit on winning has a structural interpretation as $\frac{\alpha}{1+\alpha}$ × the coefficient on the sum of the normalized B-score and acceleration costs. We therefore inflate the estimated reduced-form coefficient by $\frac{1+\alpha}{\alpha}$ and multiply the result by the acceleration cost, including this term in the index used to calculate win probabilities.

³⁹We have experimented with more or less repetitions, but there is little simulation error at this size.

⁴⁰We discard bootstrap samples where the coefficient on log usercost is negative, as the simulation results become meaningless. This occurs in approximately 4% of cases.

bootstrap the participation model since it is quite precisely estimated off a very large sample and estimation is computationally demanding. The bootstrap samples are clustered by contract and stratified by year, so that the fixed effects are estimable in all samples. We report 95% confidence intervals from the 2.5% and 97.5% quantiles of the bootstrap simulations.

Model Fit: A number of sample moments and their simulated counterparts are shown in Table 12. Fitting these moments is a fairly strong test of the structural model as a whole, since our estimation approach was to sequentially estimate a series of individual models for entry, bidding and winning, rather than simultaneously matching all the moments as in a simulated GMM procedure. The advantage is that the estimation procedure is considerably more transparent, but specification errors at each step may accumulate.

Consider the first two columns, which compare a set of important sample and simulated moments under the observed A+B assignment. The participation model fits well, although it slightly over-predicts the distance of participating bidders, and under-predicts the fraction of big firms on contracts over \$10M. Moving to the bidding model, the estimated acceleration is within the 95% confidence interval for all moments, and extremely close for the key statistic, the mean acceleration of winning bidders. Finally, the win probability model also does pretty well, matching the probability of winning conditional on having the lowest score fairly closely, although doing a worse job on the win probability in contracts with large B-components. Overall the fit is good, especially since virtually none of these moments was being directly matched during the estimation procedure.

Results: Turning now to the counterfactual A+B policies in columns (3) and (4), we see that average contract acceleration is lower than in the data. This is essentially because of selection: the Caltrans policy assigned A+B status to big contracts, and in the counterfactual policies A+B status is also assigned to small contracts, so mean acceleration falls. Comparing the optimal and budget policies, we see that although the usercosts in the optimal policy are on average over 5 times bigger, there is less than twice the acceleration, due to the convexity of the acceleration curve.⁴¹ The hard quota policy achieves exactly 10% acceleration, as one would expect, but the soft quota does slightly worse because it doesn't always bind. In addition to promoting more acceleration from all bidders, the time incentives also induce positive selection in that the winning bidder is more likely to be the contractor offering to do the project fastest — 65% of the time under the optimal policy.

⁴¹The usercosts are not ten times bigger because of the cap we enforce on the size of the usercosts.

Table 13 compares the simulated results in welfare terms. The optimal policy would reduce completion time to 64% of the engineer’s estimate, the budget policy would still achieve 76% completion time, while the hard quota would get 90% (by design) and the soft quota would get 90.4%. Under the existing policy, commuters gain nearly \$5M per contract. Gains in the counterfactual policies are smaller from \$2.5M for the optimal policy through to \$0.5M for the quota policies, once again because these policies are not selective. Acceleration costs vary sharply across policies. The optimal policy is expensive, at \$372K a contract (due to the big incentives). By contrast, the budget policy with its small incentives, induces quite good acceleration at an average cost of only \$44000 per contract. The hard and soft quotas are even cheaper, at \$10000 and \$5000 a contract. The big relative difference in the quota acceleration costs arises from the rare occasions when all the participants have high type and would prefer to finish on time, but the hard quota forces the winner to finish fast.⁴²

Moving onto the totals, the bottom line is that the counterfactual A+B policies perform substantially better. In particular, the “budget” policy raises commuter welfare relative to contractor costs by \$1.2B (compared with the current policy at \$359M) over the sample period, while reducing the total acceleration costs of contractors. To put the magnitude of these gains into perspective, the total value of these contracts (as assessed by the engineer’s estimate) is \$5.24B, so the \$1.2B gain from the budget policy is 23%. Interestingly, even the soft quota performs better than the current policy. So if A+B assignment for all contracts is not feasible for reasons not modeled here, a hybrid of more aggressive scheduling for most contracts, along with A+B assignment for the most time sensitive might still be a substantial improvement. Except in the case of the hard quota, the simulated outcomes are estimated reasonably precisely, and the confidence intervals rarely overlap.

Sensitivity Analysis: It would be reassuring if our results were not heavily dependent on the details of our structural model. In Table 14 we show outcomes for the counterfactuals under three alternate scenarios. The first is that A+B assignment is restricted to the 161 size-eligible contracts. We consider this scenario to address the concern that we are using the functional form of the structural model to make inferences for contracts for which we have never observed A+B assignment. The broad welfare conclusion is similar: the A+B policies perform better than the quotas. The hard quota now performs much better than the soft quota, because the liquidated damage rates are insufficient to deter late completion in

⁴²With steeper cost curves, the hard quota sometimes induces very high costs. Accordingly, our bootstrapped confidence intervals are sometimes virtually unbounded, and we just put a “-” in the table to indicate a very large number.

big contracts. Notice also that for the optimal policies, the welfare gain is \$500M less than before. This implies that assigning A+B status to small contracts accounts for a substantial share of the potential welfare gain.

In the second scenario we examine the sensitivity of the results to our social cost estimates. Recall that although we had been quite conservative, our social costs were much higher than the Caltrans usercosts. So here we divide them by three (making them roughly comparable to Caltrans), and rerun the simulations. Obviously the welfare gains are much smaller (by a factor of more than 3), but they are still substantive — \$311M in the case of the budget policy. Finally, in the last scenario we change α , the parameter governing the shape of the acceleration curve. To implement this, we rerun our bidding model under the constraint that the coefficient on log usercost is 0.175 (i.e. half of what it was estimated to be), and proceed as before. This implies $\alpha \approx 5.7$, a much steeper cost curve. To rationalize the observed data with this new curve, marginal costs are very low at first, though rising quicker than before.

The results of this experiment are interesting. The conclusions are if anything reinforced: even if the cost curve is much steeper than we had estimated, both the efficient and welfare policies still deliver welfare gains between \$1.3B and \$1.5B. But the case for the budget policy becomes even stronger: total acceleration costs are now only \$16M total, and the welfare gain is higher than before. The hard quota now performs far worse than the other policies, which suggests it is not a particularly robust policy choice. Overall, across all of these scenarios, the welfare gain from the budget policy at least \$311M, which is still large relative to the total size of the contracts (5.9%).

Discussion: Strong time incentives will lead contractors to accelerate, and their increased costs will be passed through to Caltrans. In fact, while we estimate that the acceleration cost for the A+B contracts is approximately \$474,000 per contract, in the reduced form analysis we estimated an increase in the winning bid of around \$1.5M. As we argued earlier, the positive correlation between base and acceleration costs we saw in the data can generate large markups. Risk aversion may also explain why there is such a gap between acceleration costs and the bids. This raises the issue of how costly this policy will be to Caltrans: if costs are passed through to bids with this kind of multiplier, it may create budget problems. It is this concern that motivates the “budget” policy, where small time incentives are used. Instead of modeling the pass-through rate explicitly, we show that the budget policy incurs lower acceleration costs than the current A+B policy, and so for any constant pass-through rate, this policy will be cheaper. This makes it a pragmatic policy choice.

Our analysis ignores dynamics and general equilibrium effects. In terms of dynamics, if types were perfectly persistent and capacity constrained, then an expansion of the A+B program might fare worse than we project, as the best types would be unable to participate on all of these contracts. But it seems plausible that type is determined by something like capacity, which evolves, so as some firms become busy, others will be able to step in. The general equilibrium concern is that accelerating all contracts would cause input prices to rise. This is not necessarily right: if overall demand for highway construction contracts remains constant, long-run input demand will also stay constant. The question is whether contractors operating at bigger scale over shorter periods of time will lead to bottlenecks in input supply, and this seems far from clear. Given the large and robust benefits to providing even small time incentives, it certainly seems worth trying.

6 Conclusion

The A+B program appears to be a policy success, with the gains to commuters substantially higher than the costs to Caltrans. Yet our analysis shows that there are still substantial inefficiencies. Expanding the use of the A+B contract design and tightening time requirements on standard contracts would significantly increase commuter welfare relative to contractor costs. The current policies reserve the A+B design for big contracts, and assign commensurately big time incentives. But as we have shown, this sacrifices about a third of the potential welfare gain (by ruling out small contracts) and is costly because of the convex acceleration costs. Small across-the-board time incentives may increase commuter welfare relative to contractor costs, while still lowering procurement costs.

References

- Abadie, Alberto, David Drukker, Jane Herr, and Guido Imbens**, “Implementing Matching Estimators for Average Treatment Effects in Stata,” *The Stata Journal*, 2001, 1 (1), 1–18.
- 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.
- , **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.
- De Silva, Dakshina, 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.
- Einav, Liran and Ignacio Esponda**, “Endogenous Participation and Local Market Power in Highway Procurement,” 2008. Working paper, NYU Stern and Stanford.
- Fick, Gary, Ells Tom Cackler, Steve Trost, and Lee Vanzler**, “Time-Related Incentive and Discincentive Provisions in Highway Construction Contracts,” Technical Report 652, National Cooperative Highway Research Program, Washington, D.C. 2010.
- 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.
- Lewis, Gregory and Patrick Bajari**, “Incentives and Adaptation: Evidence from Highway Procurement in Minnesota,” 2010. Working Paper, Harvard University.
- Li, Tong and Quang Vuong**, “Nonparametric Estimation of the Measurement Error Model Using Multiple Indicators,” *Journal of Multivariate Analysis*, 1998, 65, 135–165.
- **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.
- Pogash, Carol**, “A Miracle-Worker Highway Man Rides the Bonus Train,” *New York Times*, June 2007.
- Porter, Robert and Douglas Zona**, “Detection of Bid Rigging in Procurement Auctions,” *Journal of Political Economy*, 1993, 101 (3), 518–538.
- Strong, K, C. Tometich, and N. Raadt**, “Cost Effectiveness of Design-Build, Lane Rental, and A + B Contracting Techniques,” in “Mid-Continent Transportation Research Symposium” Ames, IA 2005.
- Weitzman, Martin L.**, “Prices vs. Quantities,” *Review of Economic Studies*, 1974, 41 (4), 477–491.

Table 1: Summary Statistics

	Standard (sample districts)	A+B	Standard (other districts)	A+B
Engineer's Estimate (\$M)	5.282 (16.46)	22.31 (25.84)	4.203 (12.27)	22.88 (23.98)
Winning Bid (\$M)	4.656 (15.43)	21.37 (31.18)	3.866 (11.21)	17.11 (14.99)
Winning Bid / Engineer's Estimate	0.906 (0.257)	0.925 (0.217)	0.903 (0.222)	0.842 (0.142)
Number of Bidders	4.949 (2.573)	5.405 (2.594)	4.687 (2.527)	4 (2.191)
Engineer's Days Estimate	125.3 (191.5)	321.6 (215.9)	117.9 (158.8)	223.3 (77.11)
Traffic	76208.3 (84053.6)	120324.6 (72717.5)	33426.4 (51914.0)	84354.2 (111990.9)
Lane Closure Ratio	0.467 (0.116)	0.433 (0.0950)	0.490 (0.111)	0.444 (0.0742)
Reopening Penalty	0.418	0.835	0.214	0.667
Estimated daily welfare loss	33977.8 (46946.3)	51130.9 (45913.6)	15258.0 (28903.3)	53835.1 (76767.1)
Usercost (if A+B)		15242 (13889.6)		27189.7 (12762.7)
N	603	79	689	6

Summary statistics for the full dataset, split into four groups. Sample districts are 4 (SF Bay), 6 (Fresno), 8 (Riverside/San Bernardino), 11 (San Diego) & 12 (Orange County). The other seven districts assign few A+B contracts and are not included in the subsequent analysis. SE for binary variables are not reported.

Table 2: Probability of A+B Assignment

	A+B Assignment		
	District 4	Other Districts	All
Engineer's Estimate (\$M)	0.0089*** (0.0031)	-0.0017 (0.0020)	0.0016 (0.0015)
Engineer's Days Estimate	-0.0005 (0.0003)	-0.0003 (0.0002)	-0.0004** (0.0002)
Traffic (thousands)	0.0017** (0.0007)	0.0009* (0.0005)	0.0014*** (0.0004)
Bridge Resurfacing	0.2172 (0.1973)	0.3049 (0.2273)	0.3020* (0.1666)
District 4			0.2556*** (0.0605)
N	84	92	176
% Correctly Predicted (eligible only)	65.88	71.28	71.51
% Correctly Predicted (all)	82.61	94.38	91.94

Coefficients are average marginal effects from a probit of AB assignment on the covariates using the sample of eligible contracts only. "Eligible Contracts" refers to contracts with engineer's days > 100 in District 4, and engineers estimate > \$5M plus usercost > \$5K elsewhere. Reported standard errors are robust. Significance is denoted by asterisks: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. The % correctly predicted is the fraction of A+B assignment decisions predicted by a rule that only assigns AB status to eligible contracts with $\hat{p} > 0.5$.

Table 3: Usercost Regressions

	Log Usercost		
	District 4	Other Districts	All
Log Engineer's Estimate	0.898*** (0.063)	0.468** (0.167)	0.759*** (0.082)
Log Engineer's Days	-0.895*** (0.095)	0.021 (0.219)	-0.637*** (0.120)
Log Traffic	0.005 (0.051)	0.011 (0.052)	0.023 (0.045)
Plant Establishment	0.375** (0.168)	0.256 (0.200)	0.256** (0.125)
Log Liquidated Damage %	0.550** (0.242)	-0.603 (0.544)	0.215 (0.250)
District 4			-0.713*** (0.103)
N	51	27	78
R^2	0.80	0.42	0.71

Robust standard errors reported. Significance is denoted by asterisks: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Summary Statistics: Treatment and Control Groups

	Standard (Control)	A+B (Treatment)	Difference
Engineer's Estimate (\$M)	19.03	19.22	-0.19
Winning Bid (\$M)	16.30	17.33	-1.04
Winning Bid / Engineer's Estimate	0.87	0.92	-0.05
Number of Bidders	5.96	5.56	0.40
Engineer's Days Estimate	361.93	315.63	46.30
Traffic	92821.25	122398.36	-29577.11**
Lane Closure Ratio	0.47	0.42	0.05***
Reopening Penalty	0.52	0.83	-0.32***
Estimated daily welfare loss	43373.70	52903.67	-9529.97
Types of Work			
Bridge Repairs	0.11	0.08	0.03
Construction	0.13	0.15	-0.03
Rehabilitation	0.39	0.41	-0.02
Widen and Realign	0.27	0.27	0.00
N	95	66	161

Summary statistics for a subsample of contracts. The treatment group is a subset of all A+B contracts, excluding contracts over \$90M, bridge resurfacing contracts, and contracts that would not ordinarily be eligible for A+B assignment. The control group is the corresponding subset of eligible standard contracts. A t-test for differences in means is conducted for each variable, and the outcome is reported in the third column. Significance is denoted by asterisks: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Policy Evaluation: OLS

	Log Winning Bid		
	(1)	(2)	(3)
A+B	0.074** (0.037)	0.106*** (0.036)	0.107*** (0.038)
Log Engineer's Estimate	0.944*** (0.021)	0.941*** (0.021)	0.940*** (0.021)
Log Engineer's Days	0.063* (0.036)	0.081** (0.041)	0.072* (0.043)
Log Traffic			0.006 (0.017)
Lane Closure Ratio			0.138 (0.129)
Reopening Penalty			0.025 (0.045)
District/Year/Work Type FE	No	Yes	Yes
R^2	0.9534	0.9714	0.9718
N	161	161	161

Coefficients are from OLS regression of winning bid on covariates. District, year and type of work fixed effects are included where indicated. Reported standard errors are robust. Significance is denoted by asterisks: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Policy Evaluation

	OLS			Nearest Neighbor	
Δ Avg. Winning Bid (\$M)	1.230** (0.597)	1.743*** (0.558)	1.767*** (0.587)	1.626* (0.986)	2.135** (0.969)
Δ Avg. Commuter Welfare (\$M)	6.981	6.981	6.981	6.981	6.981
Δ (C. Welfare - WinBid) (\$M)	5.751	5.238	5.215	5.355	4.847

The average treatment on the treated (ATT) is estimated. The first three columns correspond to the OLS regression above, with robust standard errors calculated via the delta method. The last two columns are the ATT estimated via nearest neighbor matching, with exact matching on district, year and type of work, and bias correction using the variables in OLS regressions (2) and (3) respectively. Exactly one match is used. Exact matches are achieved in 66% of cases. Significance is denoted by asterisks: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Summary Statistics: Ex-Post Data

	Standard	A+B
Contract days	193.9 (177.0)	157.6 (106.1)
Working days used	177.9 (144.3)	157.9 (108.2)
Weather days	69.21 (55.74)	73.50 (71.41)
Contract change order days	15.30 (26.82)	19.23 (25.86)
Other days	7.170 (25.47)	11.15 (22.60)
Fraction Late	0	0.0250
Daily penalty if late	6694.7 (3443.6)	11907.0 (8626.7)
Deductions	4635.2 (12885.5)	35307.9 (133301.6)
N	47	40

Ex-post summary statistics for a subsample of contracts. The A+B (treatment) group is a subset of all A+B contracts, excluding contracts over \$90M, bridge resurfacing contracts, and contracts that would not ordinarily be eligible for A+B assignment. The standard (control) group is the corresponding subset of eligible standard contracts. SE for binary variables are not reported.

Table 8: Ex-post analysis

	Working Days	Weather Days	Other Days	Total Days	Deductions
A+B	-8.942 (8.865)	-21.424 (19.677)	6.689 (6.589)	-23.677 (23.185)	19258.663 (21716.804)
Contract days	0.782*** (0.077)	-0.006 (0.079)	-0.021 (0.032)	0.754*** (0.120)	-112.945 (113.203)
Engineer's Estimate (\$M)	1.448** (0.676)	1.015 (1.113)	0.033 (0.288)	2.497** (0.976)	3099.237 (2479.001)
Traffic (thousands)	0.113 (0.073)	-0.087 (0.135)	0.108* (0.061)	0.134 (0.151)	-322.334 (247.879)
Lane Closure Ratio	-42.136 (35.718)	-144.448** (58.724)	46.664 (44.223)	-139.920 (93.651)	-2.21e+04 (98615.740)
Reopening Penalty	-9.177 (10.359)	38.053* (19.845)	-15.711 (9.738)	13.165 (24.720)	6566.810 (22013.903)
District/Year/Work Type FE	Yes	Yes	Yes	Yes	Yes
R^2	0.9669	0.2958	0.3602	0.8264	0.2608
N	87	87	87	87	87

Coefficients from OLS regressions of ex-post outcomes on covariates. District, year and type of work fixed effects are included. Reported standard errors are robust. Asterisks denote significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Participation

	Dummy for Participation				All Contracts
	A+B Eligible Contracts				
	(1)	(2)	(3)	(4)	(5)
Log Distance (miles)	-0.930*** (0.035)	-0.930*** (0.035)	-0.922*** (0.048)	-1.162*** (0.095)	-0.884*** (0.018)
Big Firm \times Engest (\$M)	0.015*** (0.004)	0.015*** (0.004)	0.015*** (0.004)	0.022*** (0.007)	0.028*** (0.003)
Big Firm \times Federal Contract	0.382* (0.217)	0.383* (0.216)	0.340 (0.217)	0.383** (0.191)	0.364*** (0.107)
Instate Contractor	-0.807*** (0.236)	-0.808*** (0.236)	-0.858*** (0.277)		-1.098*** (0.111)
Big Firm	1.108*** (0.214)	1.108*** (0.214)	0.999*** (0.222)		1.246*** (0.064)
Federal Contract	-0.213 (0.140)	-0.190 (0.142)	-0.173 (0.143)	-0.185 (0.152)	0.015 (0.057)
Engineer's Estimate (\$M)	-0.011*** (0.003)	-0.011*** (0.003)	-0.011*** (0.003)	-0.014*** (0.005)	-0.022*** (0.003)
Engineer's Days Estimate	0.001** (0.000)	0.000* (0.000)	0.000* (0.000)	0.001 (0.000)	0.001*** (0.000)
Log Traffic	-0.186*** (0.046)	-0.176*** (0.047)	-0.176*** (0.047)	-0.208*** (0.052)	-0.068*** (0.017)
Has PE work	-0.037 (0.170)	-0.003 (0.173)	0.003 (0.175)	0.041 (0.223)	0.019 (0.124)
A+B Contract		-0.099 (0.093)	-0.339 (0.628)	-0.483 (0.644)	
A+B \times Big Firm			0.334* (0.181)	0.160 (0.192)	
A+B \times Log Distance			-0.018 (0.066)	0.010 (0.069)	
A+B \times Instate			0.167 (0.517)	0.236 (0.487)	
District/Year/Work Type FE	yes	yes	yes	yes	yes
Firm FE	no	no	no	yes	no
N	18396	18396	18396	18396	194060
Wald test of A+B coefs (p-value)	-	0.291	0.213	0.663	-

Raw coefficients from logistic regressions of participation on covariates. A data point is a bidder-contract pair. Columns (1)-(4) are estimated only on contracts that are typically eligible for A+B assignment; column (5) is estimated using all contracts. District, year and type of work fixed effects are included; and firm fixed effects are included in column (4). Standard errors are robust and clustered by contract. Asterisks denote significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. The p-value reported in the last row is from a Wald test of the joint significance of the A+B coefficients.

Table 10: Incentives and Acceleration in A+B contracts

	Log days accelerated			
	(1)	(2)	(3)	(4)
Log Usercost	0.167 (0.112)	0.229* (0.120)	0.348** (0.139)	0.350** (0.138)
Log Engineer's Days Estimate	1.269*** (0.178)	1.349*** (0.182)	1.604*** (0.201)	1.610*** (0.201)
Log Engineer's Estimate	-0.200 (0.133)	-0.278** (0.139)	-0.430*** (0.138)	-0.429*** (0.138)
Log Traffic	0.026 (0.073)	0.010 (0.077)	-0.064 (0.064)	-0.064 (0.064)
Has PE work	-0.353 (0.227)	-0.368 (0.233)	-0.261 (0.212)	-0.263 (0.212)
Lane Closure Ratio	0.912 (0.783)	0.986 (0.789)	0.687 (0.673)	0.703 (0.665)
Reopening Penalty	-0.066 (0.190)	-0.075 (0.191)	-0.297* (0.173)	-0.300* (0.171)
Big Firm		0.238** (0.103)	0.257** (0.103)	0.297* (0.156)
Instate Contractor		-0.172 (0.233)	-0.324 (0.231)	-0.343 (0.241)
Log Distance (miles)		-0.074 (0.051)	-0.053 (0.041)	-0.071 (0.072)
Entry Residual				0.244 (0.623)
District/Year/Work Type FE	no	no	yes	yes
R^2	0.310	0.322	0.396	0.396
N	421	421	421	421

OLS Regressions of log days accelerated on covariates, where days accelerated is engineer's days less days bid. "Entry Residual" is the residual from the participation regression of column (5) of Table 9. District, year and type of work fixed effects are included where indicated. Standard errors are robust and clustered by contract. Asterisks denote significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: Winning and Time Incentives

	Probability of Winning		
Big Firm	0.100*** (0.029)	0.080*** (0.027)	0.042 (0.036)
Instate Contractor	0.046 (0.075)	0.056 (0.069)	0.062 (0.065)
Log Distance (miles)	-0.019 (0.012)	-0.014 (0.011)	0.002 (0.012)
B-Score / Engest		-5.295*** (1.074)	-2.726** (1.416)
Bid Residual			0.130* (0.090)
Entry Residual			-0.272* (0.098)
Contract FE	yes	yes	yes
N	2871	2871	2871

Average marginal effects are reported from a conditional logistic regression of the probability of winning on covariates. “B-score” is defined as days bid times usercost, equal to zero for standard contracts. “Bid residual” is the residual from the regression in column (3) of Table 10 above, and is included as a control function (equal to zero for standard contracts). “Entry Residual” is the residual from the participation regression of column (5) of Table 9. Contract FE are included and the model is estimated as a conditional logit. Standard errors are robust and clustered by contract. Significance levels are denoted by asterisks: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 12: Sample and Simulated Moments

	Data	Current Policy	Efficient Policy (100% Welfare)	Budget (10% Welfare)	10% Quota (hard)	10% Quota (soft)
# AB Contracts	78.00	78.00	620.00	620.00	0	
Avg. Usercost	14.80	14.80	15.66	2.97	-	-
# Bidders	4.79	4.82	4.82	4.82	4.82	4.82
# Bidders if AB	5.40	5.31	4.82	4.82	-	-
Avg. bidder distance	104.10	123.08	123.08	123.08	123.08	123.08
Fraction big firms	0.27	0.27	0.27	0.27	0.27	0.27
Fraction big firms if > \$10M	0.55	0.48	0.48	0.48	0.48	0.48
Engdays - days bid	101.46	98.50	54.20	32.36	14.15	13.60
(if AB)	-	(90.03,108.20)	(47.61,62.78)	(21.26,52.43)	(14.15,14.15)	(11.99,13.89)
Engdays - days bid	56.78	49.86	27.79	17.81	7.33	7.02
(if AB & engdays <250)	-	(43.69,58.25)	(21.16,34.76)	(12.07,29.79)	(7.33,7.33)	(6.13,7.24)
Engdays - days bid	141.79	142.40	194.86	109.84	50.49	48.64
(if AB & engdays >250)	-	(126.59,159.52)	(164.06,239.69)	(69.49,174.00)	(50.49,50.49)	(43.43,49.67)
Engdays - win days bid	126.39	127.71	72.03	40.41	14.15	13.77
(if AB)	-	(112.75,141.73)	(64.08,79.99)	(27.27,62.92)	(14.15,14.15)	(12.43,14.01)
Engdays - win days bid	69.65	63.73	35.75	22.14	7.33	7.10
(if AB & engdays <250)	-	(54.28,74.51)	(28.80,41.93)	(15.43,35.55)	(7.33,7.33)	(6.32,7.27)
Engdays - win days bid	177.60	185.44	265.26	137.70	50.49	49.28
(if AB & engdays >250)	-	(159.58,211.21)	(220.86,305.56)	(88.11,211.26)	(50.49,50.49)	(44.78,49.96)
Pr(Win) if low B-score	0.49	0.52	0.65	0.49	-	-
(if AB)	-	(0.46,0.60)	(0.55,0.71)	(0.45,0.57)		
Pr(Win) if low B-score	0.51	0.58	0.72	0.66	-	-
(if AB & b-score/engest>0.2)	-	(0.50,0.67)	(0.59,0.80)	(0.51,0.74)		

Simulation results. The first column is the observed data; the second column is simulated moments under the observed A+B assignment; the third through sixth columns are counterfactuals. The third and fourth columns are policies of assigning all contracts to be A+B, with usercost set either at the consumer welfare or some fraction of it. The fifth and sixth columns show policies reducing the number of engineer's days by 10% in all contracts, with an infinite ("hard") penalty or a soft penalty equal to the liquidated damages that would usually be assessed if late. 95% confidence intervals are given in parentheses, and are generated by bootstrapping the regressions in Tables 10 and 11, and taking the 2.5th and 97.5th percentiles of the simulated results based on the bootstrapped coefficients. Confidence intervals are not reported for participation moments, since the participation model was not bootstrapped.

Table 13: Counterfactual Welfare

	Current Policy	Efficient Policy (100% Welfare)	Budget (10% Welfare)	10% Quota (hard)	10% Quota (soft)
B-Days/Engdays (%)	69.13 (66.16,71.60)	63.73 (52.66,72.82)	77.59 (59.87,85.59)	90.00 (90.00,90.00)	90.44 (90.11,91.71)
Mean Commuter Gain (\$M)	5.08 (4.57,5.72)	2.82 (2.35,3.39)	1.98 (1.59,2.62)	0.63 (0.63,0.63)	0.61 (0.54,0.62)
Mean Acc. Cost(\$M)	0.474 (0.045,0.834)	0.372 (0.033,0.636)	0.044 (0.005,0.065)	0.011 (0.007,0.054)	0.005 (0.000,0.018)
Mean Net Gain (\$M)	4.60 (3.93,5.36)	2.45 (2.12,2.91)	1.93 (1.54,2.58)	0.62 (0.57,0.62)	0.60 (0.52,0.62)
Total Acc. Cost (\$M)	36.95 (3.52,65.07)	230.52 (20.38,394.22)	27.13 (2.91,40.32)	6.73 (4.36,33.69)	3.14 (0.16,11.37)
Total Net Gain (\$M)	359.06 (306.33,417.80)	1518.15 (1317.33,1802.10)	1197.87 (954.15,1598.89)	383.20 (356.23,385.56)	373.22 (324.89,383.40)

Counterfactual welfare results under different policies. The first column is simulated outcomes under the observed A+B policy, averaged over A+B contracts. The second and third columns are policies of assigning all contracts to be A+B, with usercost set either at the consumer welfare or some fraction of it. The fourth and fifth columns show a “quota” policy of reducing the number of engineer’s days by 10% in all contracts, with an infinite (“hard”) penalty or a soft penalty equal to the liquidated damages that would usually be assessed if late. 95% confidence intervals are given in parentheses, and are generated by bootstrapping the regressions in Tables 10 and 11, and taking the 2.5th and 97.5th percentiles of the simulated results based on the bootstrapped coefficients. For the hard quota, these percentiles are sometimes unreasonably large: we put a “-” in those cases.

Table 14: Sensitivity Analysis

		Current Policy	Efficient Policy (100% Welfare)	Budget (10% Welfare)	10% Quota (hard)	10% Quota (soft)
Scenario 1	Total Cost (\$M)	37.81	215.39	20.16	9.23	2.70
(eligible only)	Total Welfare (\$M)	358.21	970.11	705.05	380.69	255.40
Scenario 2	Total Cost (\$M)	37.81	99.95	7.92	9.23	3.17
(social cost / 3)	Total Welfare (\$M)	94.20	411.10	311.29	120.74	122.28
Scenario 3	Total Cost (\$M)	20.53	129.08	16.54	114.27	0.98
($\alpha \times 2$)	Total Welfare (\$M)	376.61	1468.07	1308.95	275.65	380.33

Counterfactual welfare results under different policies and scenarios. The columns correspond to the policies in Table 13. The first scenario restricts A+B assignment to only size-eligible contracts. The second scenario redoes the calculations with social costs one third of those we estimated. The final scenario assumes $\alpha = 5.7$, twice as big as estimated, implying a much steeper cost curve.