NBER WORKING PAPER SERIES

INCENTIVES AND ADAPTATION:
EVIDENCE FROM HIGHWAY PROCUREMENT IN MINNESOTA

Gregory Lewis
Patrick Bajari

Working Paper 17647
http://www.nber.org/papers/w17647

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2011

We are grateful to Saeid Asgari, Bassem Barsoum, Matthew Cugini, Perry Mayer, Mark Samuelson, Raymond Tritt and Steven Whipple of Caltrans; Rabinder Bains, Tom Ravn and Gus Wagner of Mn/DOT and David Kent of NYSDOT for their help with this paper and related projects. We would also like to thank John Asker, Susan Athey, Raj Chetty, Matt Gentzkow, Ken Hendricks, Jon Levin, Justin Marion, Ariel Pakes, Chad Syverson; seminar participants at Harvard, LSE, MIT, Toronto, UC Davis and Wisconsin; and participants at the AEA, CAPCP, IIOC, Stony Brook, UBC IO, WBEC and the NBER IO / Market Design / PE conferences for useful comments and suggestions. Lou Argentieri, Jorge Alvarez, Minjung Kim, Jason Kriss, Zhenyu Lai, Tom Longwell, Tina Marsh, Maryam Saeedi, Connan Snider and Hao Teng provided excellent research assistance. The authors gratefully acknowledge support from the NSF (grant no. SES-0924371). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Incentives and Adaptation: Evidence from Highway Procurement in Minnesota
Gregory Lewis and Patrick Bajari
NBER Working Paper No. 17647
December 2011
JEL No. D86,H57,L92

ABSTRACT

Procurement projects often encounter unanticipated problems. Deadlines and penalties are one important instrument used to incentivize contractors to adapt their plans. We develop a theory of highway procurement in which contractors must modify their construction rate following a productivity shock. We model how time incentives affect the work rate and time taken, characterizing the efficient contract design. Using new micro-level data from Minnesota that includes day-by-day information on work plans, actual outcomes and delays, we find strong evidence supporting the theory. As an application, we build an econometric model that endogenizes adaptation, and simulate how different incentive structures affect outcomes and the variance of contractor payments. Accounting for the traffic delays caused by construction, switching to a more efficient design would substantially increase welfare without substantially increasing risk.

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

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

An online appendix is available at:
http://www.nber.org/data-appendix/w17647

1 Introduction

Public procurement is big business. The European Commission estimated in 2002 that it amounted to 16.3% of the European Union's GDP (European Commission 2004). This fraction is typical of many developing and developed countries: for example, the World Bank assessed the share in South Africa to be 13% of GDP in 2003 (World Bank 2003). In the United States, highway construction in particular has recently received increased funding through a 6-month, $19.9 billion extension of the existing surface transportation bill.[1]

For many procurement activities, cost and outcome uncertainty is unavoidable. In highway construction, bad weather, unexpected site conditions, poor planning or materials delays can quickly derail a project. When such shocks arise, the question is how contractors will adapt their plans. This depends on the incentives provided by the procurement contract. In cost-plus contracts, where the government pays the contractor's costs plus a fixed fee, it may be easy for the parties to agree on an ex-post efficient plan. By contrast, in a fixed-price contract where there is no cost sharing, the contractor will only adapt if the contract gives him private incentives to do so, or if the contract is renegotiated. Bajari and Tadelis (2001) have argued that the cost-plus design is therefore preferable when adaptation is likely to be necessary, renegotiation is costly, and the cost of specifying a complete contract to cover all possible contingencies is high.

An important distinction lies between shocks that require a whole new project design (i.e. for both parties to adapt), and those that require only the contractor to adapt his plans for executing the old design. In the latter case of one-sided adaptation, project deadlines can in principle effectively substitute for more detailed contracts, by giving the contractor a private incentive to re-optimize their plans whatever state of nature arises.

In this paper, we examine whether time incentives are effective for motivating adaptation in practice. We consider data from state highway construction projects in Minnesota. Several features make this an attractive case study. First, many highway construction activities are relatively routine, with well-specified designs that rarely need to be adapted. Second, our dataset is unusually large and rich, containing day-by-day reports by the project engineers on weather conditions, delays, and planned and actual work hours. This allows us to identify unanticipated shocks by looking at how much work was actually required to complete the project, relative to the initial plans. Finally, highway construction is an area where adap-

[1] Source: American Highway Users Alliance.

tation really matters. When contractors overrun their deadline by 10 days, commuters are inconvenienced by 10 more days of traffic delays. With traffic on a typical state highway in Minnesota exceeding 10,000 vehicles per day, a 5 minute delay to each motorist over a 10 day period amounts to a welfare loss of $100,000, when we value their time at $12/hr. It may be more socially efficient for the contractor to adopt 16-hour workdays rather than continue at the old pace and finish late.

We reach two main conclusions. First, we find that contractors do indeed adapt their work rates in response to construction shocks in order to meet project deadlines. Consistent with the model we develop, we show that such adaptation is limited: for shocks that would require an increase in work rate of an hour per day to meet the deadline, firms tend to increase their work rate by less than half an hour per day. This is not surprising given the weak penalties in Minnesota.

Second, we show that higher penalties would in this case be more effective in improving joint welfare than tighter deadlines. To reach this conclusion, we structurally estimate the contractor's time-related costs. Our approach uses the first order condition from our model that equates the marginal benefit and marginal cost of a change in the work rate. The marginal benefits of speedier construction depend on whether the contract would finish late if the work rate was left unchanged, and on the daily penalties for late completion. These penalties are specified by state regulations, and vary with contract size and year. Given this exogenous variation in marginal benefits, we can use the first order condition to back out the marginal costs. We can then infer the counterfactual completion time under different incentive structures. Combining this with estimates of the externality from traffic delays, we can compare different policies on welfare grounds.

We find that for an average $1.3 million contract, the change in joint welfare from setting a 10% tighter construction deadline is less than $7,000. By contrast, switching to penalties equal to 10% of the traffic delay cost would raise joint welfare by almost $40,000 per contract. We also examine the fully efficient policy, which is a "lane rental" in which the contractor must pay the full traffic delay cost for each day of construction. This has far bigger effects — welfare increases by $400,000 per contract — but is also rather out of sample, and thus hard to evaluate.

Our work quantifies the natural complementarity between deadlines and penalties: without high penalties, contractors respond little to tighter deadlines. Because the social costs of construction vary so much across highway projects, we also argue for tailored penalties, rather than the "one-size-fits-all" approach of statewide contract specifications.

To the best of our knowledge, this is the first empirical paper on the relationship between ex-ante incentives and adaptation in procurement. Early work largely took the form of theory papers emphasizing the role of asymmetric information and moral hazard (McAfee and McMillan (1986), Laffont and Tirole (1986), Laffont and Tirole (1987)). More recent papers focused on analyzing the incentive properties of commonly observed contract forms and arrangements (see for example Bajari and Tadelis (2001), Levin and Tadelis (2010), Bajari, McMillan and Tadelis (2009), Martimort and Pouyet (2008) and Maskin and Tirole (2008)). Issues of adaptation, contractual completeness and renegotiation often arise.

There have been a number of empirical papers analyzing highway procurement in particular (see for example Porter and Zona (1993), Hong and Shum (2002), Bajari and Ye (2003), Jofre-Bonet and Pesendorfer (2003), Krasnokutskaya (2009), Krasnokutskaya and Seim (2010), Li and Zheng (2009), Marion (2007), De Silva, Dunne, Kankanamge and Kosmopoulou (2008), Gil and Marion (2009)). Some have focused on the bidding process, emphasizing features such as the potential for collusion, and preferential treatment of local and small contractors. Others have been concerned with measurement, developing procedures for non-parametric identification of valuations under dynamics or unobserved heterogeneity. One closely related empirical paper is Bajari, Houghton and Tadelis (2007), which analyzes how private information on ex-post materials requirements can lead to ex-ante bidding distortions under unit bidding. Another is Lewis and Bajari (2011), where we examined the use of scoring auctions to award contracts based on both time and price.

More broadly, this paper forms part of the empirical literature on high-powered incentives and their effects on output, which has mainly focused on labor contracts within the firm (see Chevalier and Ellison (1997), Prendergast (1999), Lazear (2000) and Bandiera, Barankay and Rasul (2005)). Finally, a recent literature in public economics has focused on "bunching" at kink points in the tax structure (Saez 2010, Chetty, Friedman, Olsen and Pistaferri 2011). The discontinuity in penalties at the project deadline generates similar "bunching" in our analysis — see Figure 4 — though our econometric approach differs.

The paper proceeds as follows. Section 2 presents an overview of the highway procurement process. Sections 3, 4 and 5 respectively contain the theoretical, descriptive and policy analysis. Section 6 concludes. All tables are to be found in the appendix.

2 The Highway Construction Process

Building or repairing a highway is a complex activity. Here, we emphasize key features of the process in Minnesota that inform our later modeling choices. Figure 1 gives a simple timeline, starting from when the contract is awarded. At that time, the winning contractor must post a “contract bond” guaranteeing the completion of the contract according to the design specification. This bond is typically secured through a third party who will take on the contractor’s obligations in the event of default.

Once the contract is awarded, the contractor must plan how to structure the various distinct activities, such as excavation or grading, that make up the construction project. To do this, they work out how long each activity will take for a standard crew size, and then use sophisticated software to work out the optimal sequence to complete the activities in by using the “critical path method” (Clough, Sears and Sears 2005). The key feature of this technique is that some activities are designated as critical, and must be completed on time to avoid delay, while others are off the critical path and have some time slack. The critical activities are called the “project controlling operations” (PCOs).

The contractor presents his plan to the project engineer in the pre-construction meeting. It is considered good practice to choose a plan that allows some contingency time on the side (around 5% of the time allowed). But a busy contractor may select a plan that allows little or no margin for error, or alternatively plan to finish early and move onto another project. This will obviously be affected by the time incentives that are offered. In Minnesota the incentives are usually simple. The design engineer initially specifies a number of “working days” that the contractor is allowed to take to complete the contract. A “working day” is a day on which the contractor could reasonably be expected to work. Usually this means weekdays (excluding public holidays) with amenable weather conditions. When the contractor works, a working day is charged. When the contractor could have worked, but didn’t, a working day is charged and a note is made of the hours of “avoidable delay”. Finally, when work was not possible, a working day is not charged, and “unavoidable delays” are noted.

Each additional day beyond the number of target working days is charged as a day late. The penalties for being late are specified in the standard contract specifications, which we reproduce in Table 1. Each day late incurs a constant penalty which depends only on the size of the contract. The penalties were last increased in 2005. So the contractor has no external incentive to finish early, but some incentive to avoid finishing late.

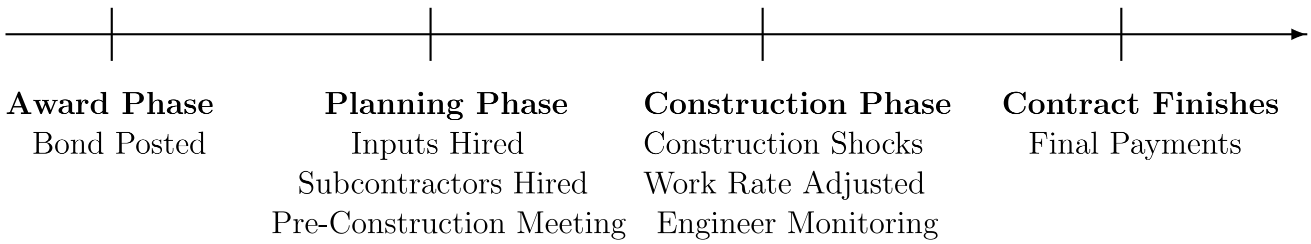

Figure 1: Construction Process

Once the planning is complete, 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 planned specifications. Productivity shocks, materials delays or unexpected site conditions may affect the rate at which any activity is completed, and the contractor must continually check progress against the planned time path. If necessary, the work rate may need to be amended, especially when there is delay on a critical path activity. At the end of the process, the contractor is paid the amount bid less any damages assessed for late completion.

3 Model

With this process in mind, we outline a model of adaptation under time incentives. Contractors have private costs that depend on the amount of capital and the size of the work crew they employ, as well as on the number of hours per day they choose to work. Road construction inflicts a negative externality on commuters, so the efficient outcome requires accelerated construction relative to the private optimum. Faster construction in turn requires either increased scale (more capital and labor) or an increased work rate. We assume that the scale is determined at the start of the project, so that as productivity shocks occur, the firm can only adapt to those shocks by changing their work rate. We want to know how time incentives affect such adaptation, and what implications this has for the socially efficient contract design.

We begin by modeling the production technology using a two-period model. In the first period, the contractor chooses a level of capital K, representing all the factors of production that will be fixed over the length of the project (hired equipment, project manager etc). They also fix the labor L. Following this, a shock θ is realized. This shock is anything that was unanticipated about the amount of work needed to complete the project. For simplicity, we will refer to it as a productivity shock below. Given K, L and θ, the project takes $H(K, L, \theta)$ man-hours to complete.

In the second period, the firm chooses a uniform work rate h (in hours per day).[2] This in turn determines the number of days $d = H(K, L, \theta)/hL$ the project will take, since the number of man-hours of work completed each day is just hL. We impose some economically motivated restrictions on the total hours $H(K, L, \theta)$. Capital substitutes for labor, decreasing the number of man hours required ($H_K < 0$). Labor has declining marginal product, so that adding additional labor increases the number of man hours required ($H_L > 0$), though decreasing the number of days taken for a fixed work rate ($\frac{\partial}{\partial L} d(K, L, \theta, h) < 0$). Last, a good productivity shock corresponds to a low θ ($H_\theta > 0$).

The work-rate decision will be influenced by the time incentives laid out in the contract. We consider time incentives that take the following form: a target completion date d^T and a penalty c_D for each day late. These form of incentives are widely used in highway procurement; other forms of time incentives are called "innovative". One innovative design is the "lane rental" contract. In this design, the contractor pays a rental rate for each day of construction that closes a lane. For construction jobs that require continuous lane closure, this is a special case with a deadline of $d^T = 0$ and lane rental rate c_D.

The contractor is risk-neutral, and pays daily rental rates of r per unit of capital, and an hourly wage $w(h)$ to each worker. For algebraic simplicity, the wage function $w(h)$ is assumed to take the linear form $w(h) = \underline{w} + bh$, a base wage plus an increment that depends on the work rate, reflecting overtime, bonuses for night-time work etc. Overall the contractor's ex-post private costs for a given set of time incentives are:

$$\begin{aligned} C(h, K, L, \theta) &= \underbrace{H(K, L, \theta)w(h)}_{\text{labor costs}} + \underbrace{rdK}_{\text{capital costs}} + \underbrace{\max\{d - d^T, 0\}c_D}_{\text{time penalties}} \\ &= H(K, L, \theta)w(h) + rK\frac{H(K, L, \theta)}{hL} + \max\left\{\frac{H(K, L, \theta)}{hL} - d^T, 0\right\} c_D \end{aligned} \tag{1}$$

Traffic delay costs are assumed to be linear in the days taken, with the daily cost equal to a constant c_T. Calculating this externality as the product of the delay, the amount of traffic and the average time value of commuters, linearity is a good approximation if traffic and delays are constant over time. Realistically, traffic delay costs should vary with traffic

[2] Provided the wage rate $w(h)$ given below is weakly convex, a uniform work rate will be weakly preferable to a more complex construction schedule.

patterns both within and across days, and also with the extent to which construction impedes commuters (e.g. are commuters still delayed when construction ceases for the day?). Here we simplify by assuming that the traffic delay costs are independent of the work rate h.

Discussion: We have assumed that the productivity shock is realized before the work rate decision is made, so that the contractor can choose when the contract will be completed. An alternative would be to make the contractor decide on his work rate before the productivity shock is realized, so that the completion time is stochastic. Both of these are imperfect approximations to a more complex dynamic process. The latter timing assumption is closer to the standard principal-agent models in the theory literature, where the agent chooses an effort level that induces a distribution of (contractible) outcomes. The advantage for us of modeling completion times as deterministic is that we are able to explain why a large fraction of contracts in our data finish exactly on time — because the contractor chooses to finish on time! By contrast, the other model would require that contractors plan to finish on time and that there is a "zero shock" atom, and as we will show, this is not the case.

We have also assumed the contractor is risk-neutral. This assumption is innocuous for most of the analysis, since we work with the ex-post profit function to see how incentives affect adaptation. But for the welfare analysis at the end of this section, the ex-ante joint welfare will vary with the contractor's risk preferences. We discuss the implications of alternative models at that time.

Finally, we assume that the quality of the construction is unaffected by the time incentives. One may worry that the contractor may shirk on quality to save time. This is a version of the famous multi-tasking problem of Holmstrom and Milgrom (1991). In highway procurement, this is less of a concern, as the government employs a project engineer to monitor the construction and ensure that the finished project meets the contract specifications. Low quality construction is penalized by additional penalties laid out in the contract terms, and these penalties can be enforced against the contract bond.[3]

Adaptation: We start the analysis by looking at how the work rate h is chosen, given the realization of the productivity shock. Taking a first-order condition in h, and dealing with

[3] Lewis and Bajari (2011) found that there was almost no difference in the quality penalties charged between highway construction contracts auctioned using scoring auctions (which emphasize cost and time) and standard auctions (which emphasize only cost).

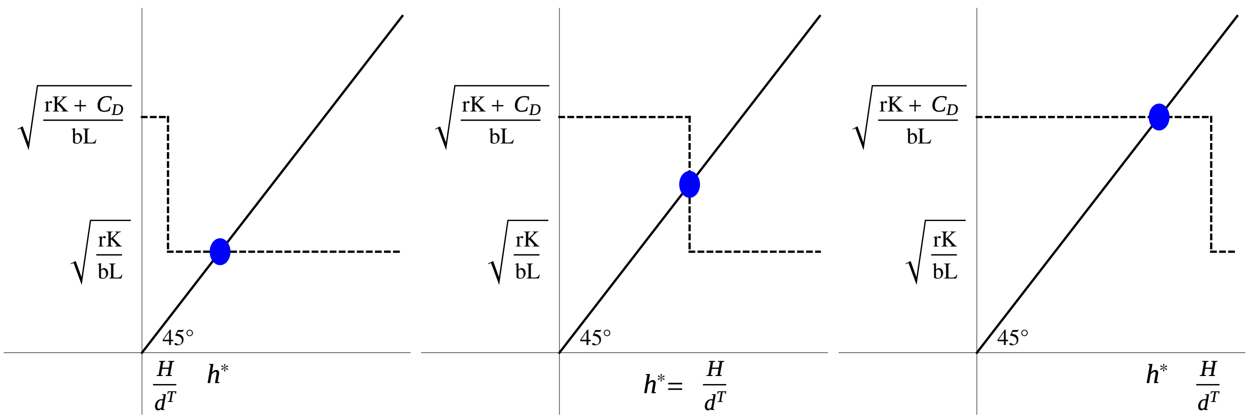

Figure 2: **Optimal Work Rates.** The figure depicts how work rates change with the number of hours of work required to complete the project. In the left panel, a favorable productivity shock means that a slow work rate would suffice for on-time completion, but the contractor works faster to economize on capital costs and will finish early. The middle panel shows a contract where no productivity shock has occurred, and the contractor works at a rate that leads to on-time completion. In the right panel, a negative productivity shock implies a fast work rate is necessary for on-time completion, but the contractor optimally chooses to work slower and will finish late.

various boundary issues resulting from the presence of the max operator in the objective function, we get the following expression for the optimal h^*

$$h^* = \begin{cases} \sqrt{\frac{rK}{bL}} & \text{if } \frac{H}{d^T} < \frac{rK}{bL} \\ \frac{H}{d^T} & \text{if } \frac{H}{d^T} \in \left[\frac{rK}{bL}, \frac{rK+c_D}{bL}\right] \\ \sqrt{\frac{rK+c_D}{bL}} & \text{if } \frac{H}{d^T} > \frac{rK+c_D}{bL} \end{cases} \tag{2}$$

It is easiest to see what's going on here in a picture. There are three cases, corresponding to contracts in which the required work rate for on-time completion $H(\theta)/d^T$ is low (good productivity draws), those where it is intermediate (average productivity draws), and those where is it high (poor productivity draws). These cases are depicted in Figure 2 as the left, middle and right panels, respectively. When the contract is unexpectedly easy to complete, the contractor could work slowly and still complete on time, avoiding the wage premium for accelerated work. The countervailing incentive is that this ties up capital over a longer period, which is costly. Balancing these incentives, the contractor chooses $h^* = \sqrt{\frac{rK}{bL}}$, which is increasing in the rental rate and the capital-labor ratio, and decreasing in the slope of the wage premium. On the other hand, given a middling draw, the contractor chooses a work rate that allows on-time completion ($h^* = H/d^T$), since accelerating to be early is too costly, and slowing down to be late incurs time penalties. Finally, when facing a poor productivity shock, the contractor chooses a work rate of $h^* = \sqrt{\frac{rK+c_D}{bL}}$ and finishes late. Higher time penalties imply faster work rates in this case.

So the contractor adapts his work rate in line with the productivity shocks, although the range of adaptation is bounded. Specifically, for fixed capital and labor, there are a range of shocks θ for which $\frac{H}{d^T} \in \left[\frac{rK}{bL}, \frac{rK+c_D}{bL}\right]$, and in that range, the contractor will accelerate or decelerate work as needed to keep production on time. But no positive productivity shock could induce a slower work rate than $\frac{rK}{bL}$, nor could a sufficiently negative one induce him to work faster than $\frac{rK+c_D}{bL}$. To reduce construction time even after bad shocks, one needs high penalties c_D. These increase the maximum work rate of the contractor.

One important caveat to this analysis is that it is *short-run*, in that we are holding the capital and labor inputs fixed. If the procurer were to consistently offer more aggressive time incentives, contractors would learn to use more capital or labor than is standard, in order to get the jobs done faster. Indeed the evidence presented in Lewis and Bajari (2011) suggests that when high-powered time incentives are offered, the contractors may be willing to adopt entirely non-standard work schedules, signing costly rush orders for inputs, using big work crews, and working 24 hours a day. These long-run changes can only reduce their costs below the short-run levels. Since we never see such high-powered incentives in the data, our entire analysis will focus on the short-run.

Welfare Analysis: We now look at how standard contracts and lane rentals compare in terms of welfare outcomes. For this, it is useful to separate the time incentives from the contractor's other costs (capital rental and wages), and look at how much money the contractor can save by slowing down and completing the contract one day later. Writing $c(d;\theta) = H(K,L,\theta)\left(\underline{w} + b\frac{H(K,L,\theta)}{dL}\right) + rKd$ for the cost of completing in d days for a given K, L, and taking a first order condition, we get:

$$-c'(d;\theta) = \frac{bH(K,L,\theta)^2}{d^2L} - rK \tag{3}$$

We refer to this as the marginal benefit of delay. It is strictly decreasing in d. The interpretation is that while an extra day of construction is useful to a contractor facing a tight work schedule (low d), it is not worthwhile when the pace of construction is already rather slow.

In Figure 3 we depict how the time incentives affect the contractor's choice of completion

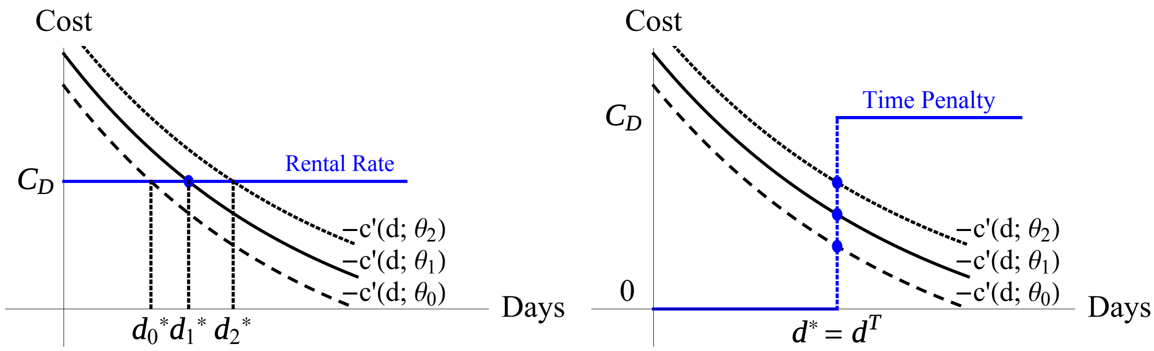

Figure 3: **Completion Time in Lane Rentals and Standard Contracts.** Both panels depict the marginal benefit to delay curve $-c'(d;\theta)$, drawn for three different productivity shocks. In the left panel, a lane rental contract imposes a constant cost of delay c_D, so the contractor optimally completes at d_0^*, d_1^* and d_2^* respectively, in each case equating marginal benefit and cost of delay. In the right panel, the incentive structure is standard, with damages charged after the target completion time d^T. In all cases the contractor will optimally complete exactly on time.

time. The left panel shows the marginal benefit of delay curves for three different productivity shocks under a lane rental contract. Here the contractor faces a daily penalty of c_D right from day one, and thus the incentive structure is flat. Profit maximization implies that he equate the marginal costs and benefits of delay, and so for each shock θ_i he completes at d_i^*.

If the rental rate is set equal to the daily traffic delay costs c_T, the contractor internalizes the negative externality inflicted by the construction, and the social planner's problem is identical to that faced by the contractor. Accordingly, the contractor will hire capital and labor to minimize expected social costs (the sum of private and traffic delay costs), and efficiently choose the work rate given the realization of the productivity shock. Ex-post, the input choices may be sub-optimal, as they are not perfectly adapted to the productivity shock, but this is unavoidable given the timing.

The right panel of Figure 3 shows the same benefit curves under the standard incentive structure. Before the target date d^T, the contractor has no marginal cost of an extra day, since this is not penalized at all. But after the target, each additional day taken attracts c_D in time penalties, and therefore the marginal costs of delay jump discontinuously from zero to c_D at d^T. In the figure this implies that for all three different productivity shocks the contractor will complete exactly on time. There is no incentive to complete early, as delay remains valuable; but also no reason to be late, as delay is not sufficiently worthwhile to offset the time penalties. This implies that completion times should be "sticky" at the target date: we should see many contracts finishing exactly on time.

In contrast to the simple lane rental design, the standard contract design will almost certainly lead to inefficient outcomes. The contractor should efficiently adapt to different productivity shocks by choosing different completion times, but the wedge in incentives makes this privately sub-optimal. In addition, there's little incentive to hire additional capital or labor at the planning stage to increase the probability of quick completion, since finishing early is not rewarded. This makes it difficult for the procurer to design a good incentive structure. On the one hand, setting $c_D = c_T$ at least ensures efficient adaptation for bad productivity draws (it sets the right penalty). But it may be preferable to distort short-run incentives with $c_D > c_T$, setting unreasonably high penalties. This increases the ex-ante incentive to hire additional capital, and thereby gives the contractor ex-post incentives to finish quickly. This is a second-best solution, creating a short-run distortion to offset a long-run distortion.

This analysis is similar to Weitzman (1974) on regulating a firm with unknown costs of compliance. The twist in this two-period model is that the contractor is also ex-ante uncertain, and only learns his costs after the incentive structure has been chosen. The lane rental is essentially a Pigouvian tax, and remains efficient in an ex-ante sense. The standard design is like a quota, and has the usual problem that the regulator has to set it without knowing the underlying costs of the contractor. This leads to inefficiency.

Risk Aversion: An alternative way to look at this problem is to use a standard principal-agent model (Holmstrom and Milgrom 1987). In that model, the contractor is risk-averse, his "effort" is his work rate, the output is the number of days taken, and the productivity shock is the source of output uncertainty. The work rate is non-contractible and the shock unobserved, so that the procurer can contract only on completion time as usual. As we know from that literature, it is no longer optimal to transfer all the risk to the contractor by using the efficient lane rental: giving such high-powered incentives in the presence of productivity shocks increases the variance of the contractor's payments, lowering their expected utility. Weaker incentives are to be preferred here.

It is hard to assess how important risk aversion is in describing contractor's preferences, although papers on skew bidding suggest that they are at least partially risk averse (Athey and Levin 2001, Bajari et al. 2007). Since we cannot analyze this case without making specific functional form assumptions on the contractor's utility, we assume risk neutrality throughout the empirical analysis. We also report how different policies affect the standard deviation of the contractor's payments, so that the reader can assess the impact on risk.

4 Descriptive Analysis

The theory outlined above indicates how contractors should adapt to productivity shocks, and how such adaptation is mediated by the contract design. In the remainder of the paper, we analyze data from contracts let by the Minnesota Department of Transportation (Mn/DOT). Our dataset is unusually detailed, as it includes daily reports by the project engineer on how construction is progressing. This enables us to test the theory, seeing if contractors adapt their work rate in response to productivity shocks, and exhibit the "stickiness" in completion times predicted by the model. Having shown that the theory is largely confirmed, we estimate the contractor's short-run cost curves and use these to run some counterfactual simulations of alternative policies, such as lane rentals.

4.1 The Data

We obtained data from Mn/DOT on all the highway procurement contracts completed in Minnesota during the period 1996-2006. The main data source is the project engineer's own reports on the progress of each project, which are recorded in a program called FieldOps. We refer to this as the "diary data". From these files, we learn how many hours the contractor worked each day, how many hours of work were initially planned, how many hours of avoidable or unavoidable delays were recorded, what the weather conditions were like, and what the current project controlling operation was. We also see how working days were charged, and therefore can deduce whether the project finished early or late. A secondary data source is the bidding abstracts, which detail what the contract is and where it is located, who bid and what amount they bid, as well as the project engineer's initial cost estimate.

The full sample includes contracts with missing data or non-standard time incentives, and projects that do not require lane closure, and therefore are unlikely to significantly impact commuters. We drop these observations, to end up with a sample of 490 contracts. One variable that is not in the data and thus needs to be constructed for the welfare analysis is an estimate of the daily traffic delay cost. We calculated a contract-specific measure by multiplying the average daily traffic around the construction location by an estimate of the time value of commuters ($12/hour), and a conservative estimate of the delay that construction will cause them. Because estimating the delay required detailed manual work on Google Maps, we constructed these estimates only for a subsample of 99 contracts (the delay subsample). More details on both sample selection and the estimation of the traffic delay costs are available in the data appendix.

We present summary statistics on the contracts in Table 2. A typical contract has value just over \$1.3 million, and is of relatively short duration, on average 40 days. During the contract, contractors work for 432 hours, an average work rate of 10.15 hours a day. A substantial number of both avoidable and unavoidable delays are recorded. Contracts are generally completed on time, although in the event that they are completed late, damages are assessed in only 24% of cases. This is because the project engineer has discretion over when to assess penalties. Damages, when assessed, range from \$500 to as high as \$29,000. In Table 1, we present detailed summary statistics on completion time. Smaller contracts are more likely to finish early or on time than larger contracts. In fact, none of the contracts of size less than \$100,000 finish late, perhaps reflecting the fact that the penalties in these contracts are a larger fraction of the total contract value.

Most of the contracts in our sample involve some resurfacing work (nearly 90%), which will require lane closure. The next biggest category is the construction of turn lanes and other work on road shoulders. From this we project an average daily traffic delay cost of around \$15,000, arising from a 10 minute delay to 12,000 commuters.

4.2 Shocks and Adaptation

We use this data to see how contractors adapt to productivity shocks. We construct a measure of the shock θ by comparing the total hours needed to complete the project H with the number of hours initially planned H^P. As bigger contracts will naturally attract bigger shocks, we normalize across contracts by dividing through by the number of days allowed:

$$\widetilde{\theta} = \frac{H - H^P}{d^T}$$

A histogram of the normalized shocks $\widetilde{\theta}$ is shown in the left panel of Figure 4. Shocks are almost symmetric around zero, so that positive and negative shocks are equally likely. Moreover, the shocks can be quite big: the standard deviation is 4.6 hours of work per day, compared with an average workday of 9.6 hours. To see how these shocks affect the work rate, we regress the work rate on the planned work rate, the normalized shock and normalized penalty, and other covariates. One of these is a dummy for whether the contractor is a big firm, defined by having won more than \$10 million in contracts over the sample period; another is a dummy for the contractor being located in Minnesota.

$$h = \alpha_0 h^P + \alpha_1 \widetilde{\theta} + \alpha_2 \frac{c_D}{d^T} + x\beta + \varepsilon \tag{4}$$

The first two columns of Table 3 present the results of these OLS regressions. We find a positive and significant relationship: when a contract requires an extra hour of work per day to remain on-time, the contractor only works around half an hour more each day. This implies that the contract will finish late. This is consistent with the theory shown in Figure 2: there is only partial adjustment for particularly big or small shocks. The coefficient on time penalties is positive, though statistically insignificant. This may be because the time incentives only matter for contracts which experience negative productivity shocks. As we'll see below, the time incentives do appear to bind on these contracts.

There may be a reverse causality problem here, in that when contractors work at a faster rate the contract may actually take more hours to complete (diminishing returns). In that case the shock $\widetilde{\theta}$ and error ε may be correlated. We address this with instrumental variables (IV) regression. Specifically, we instrument for the shock with the number of hours of rain and snow each day. Bad weather can affect the number of hours needed to complete a project, by adversely impacting site conditions. But it has no effect on the optimal work rate except through the hours required, since no working days are charged when it is raining or snowing — in other words, it has no impact on the time incentives. The results reported in the last two columns of Table 3 are qualitatively similar.

4.3 Time Incentives and Completion Times

Having shown that contractors do adapt to productivity shocks, we now consider how the incentive structure affects this adaptation. Recall that our theory predicts that the contract completion time will be "sticky" around the deadline, so that many contracts will be completed exactly on time. Look at the right panel of Figure 4, which is a histogram of the days late across contracts. In the figure, a contract exactly on time has been added to the bin to the left of 0, and so you can see that over 20% of the contracts were completed either just on time or a day early. This is *exactly* what the theory predicts: for moderate shocks, contractors adapt their construction speed to finish on time.

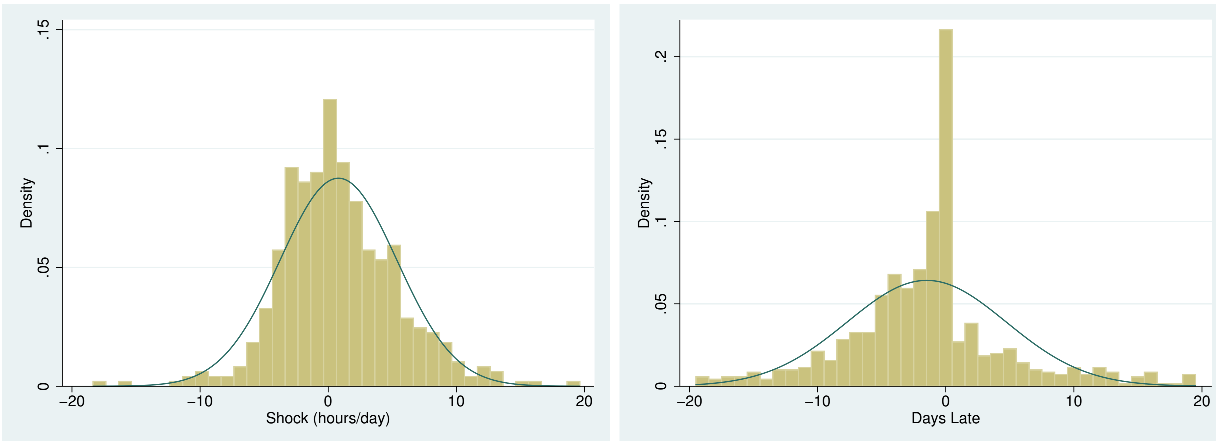

Figure 4: **Shocks and Completion Times.** The left panel is a histogram of the "productivity shock": the difference between actual hours taken and hours planned, divided by the days allowed. The histogram in the right panel is of the difference between actual and contractually specified completion time in standard contracts. The huge spike just before 0 indicates many contracts are completed exactly on time, as the theory predicts. Both histograms have a normal density superimposed.

Comparing the left and right panels of the figure, we note that both are well approximated by a normal density to the left of the deadline. To the right, the completion times histogram has "missing" mass. This too is an implication of the theory: for sufficiently positive shocks, the contractor finishes early, but given the symmetric negative shock, the penalties may prompt them to finish on time, thereby removing mass from after the target date.

To examine this formally requires something more sophisticated than linear regression, as the mass point at zero implies that the conditional mean days taken is not a linear function of the covariates. We set up such a model in the policy analysis below. But we can still try answer a basic question: when incentives are bigger, are contractors late less often?

Table 4 shows the result from a linear regression of an indicator for the contract being late on time penalties and other covariates (the results from a probit are similar). We find that contracts with higher penalty rates, relative to days allowed, are significantly less likely to be late. The effects are quite big: they imply that for an average contract, doubling the penalties would reduce the probability that the contract is late by 17%. High-value contracts are more likely to finish late. Rain and snow are significant predictors of late completion, which makes sense because of their impact on site conditions.

These results together are suggestive of a causal effect of disincentives on adaptation. But the absence of a clean quasi-experiment means that we cannot rule out other explanations.

For example, perhaps contractors complete on time because there are non-pecuniary costs of late completion, such as acquiring a poor reputation. Our understanding of the public procurement process is that reputation doesn't play much of a role, because federal procurement regulations give little discretion for selecting contractors on any basis other than cost or quality. Nonetheless there could be some other explanation for the mass point.

What we would really like to do is use the discontinuity in incentives around the thresholds in Table 1 to show that the probability of being late falls discontinuously around those points. Unfortunately, we have far too few observations to make this practical. So in what follows, we attribute all of the adaptation that we see in the data to the incentive structure (this is implicit in the first order conditions we develop below). This may lead us to overestimate the responsiveness of contractors to incentives. We try to get a sense of how reasonable our results are later in the paper.

5 Policy Analysis

Having found evidence that contractors adjust their work rate in order to meet project deadlines, we would now like to assess different policy proposals for alleviating the negative externality caused by construction. Our strategy for doing this is as follows. We estimate the contractor's short-run private costs of acceleration by looking at how their behavior changes as damages vary, using first order conditions motivated directly by our theoretical model. With these in hand, we consider simple counterfactual policy changes, including accelerated targets and higher penalties. Of particular interest is a realistic case in which the lane rental is a constant fraction of the traffic delay cost, which is constrained efficient when the procurer faces budget constraints.

One important caveat is that our analysis is entirely short-run. What we see in the data is how contractors adapted to different shocks under different incentive structures, given whatever capital and labor choices they had already made. This doesn't tell us how their input choices might change under a counterfactual policy, and so we are forced to hold everything constant. Fortunately, the bias can be signed: since long-run costs should be no more than short-run costs in expectation (since they are solving the same optimization problem with fewer constraints), we will tend to overestimate contractor acceleration costs, and therefore underestimate counterfactual welfare gains.

5.1 Estimation of Contractor Costs

Recall from the theory that the contractor acts to equate the marginal benefit of delay $-c'(d;\theta)$ with the marginal costs of delay, which are determined by the time incentives. We observe both the number of days taken d, and a measure of the shock θ. We also know the exact form of the time incentives. What we don't know is the exact form of the marginal benefit of delay function. So we specify a simple parametric form for it:

$$-c'(d;\theta) = d^T(x\beta + \varepsilon) + \alpha d + \delta\theta = d^T(x\beta + \alpha\widetilde{d} + \delta\widetilde{\theta} + \varepsilon) \tag{5}$$

where $\alpha < 0$, $\delta > 0$, $\widetilde{d} = d/d^T$, $\widetilde{\theta} = \theta/d^T$, and $\varepsilon \sim N(0, \sigma^2)$.[4]

This specification is important for what follows and worth discussing. We assume that the marginal benefits of delay are linear in the days taken d and the shock θ, and that the intercept is proportional to the contract target date d^T. What we have in mind is that contracts with a longer duration are typically more labor intensive, and so the benefit of delay in terms of reduced wages scales with the contract length. We also assume that unobservable benefits are normally distributed. These scale, linearity and normality assumptions, although restrictive, have the advantage of being simple. We will show that we can fit the data quite well despite these restrictions. Nonetheless we will have to be careful about out-of-sample extrapolation that relies solely on these parametric choices.

The idea behind our estimation approach is to use the first order conditions from the theory model developed in the earlier sections.[5] This will allow us to account explicitly for the discontinuity in incentives as the contractor moves from "early" to "late", which produces the mass of on time completions seen in Figure 4. The first order conditions must satisfy:

$$\begin{aligned} -c'(d;\theta) &= 0 \qquad \text{if } d < d^T \\ -c'(d;\theta) &= c_D \qquad \text{if } d > d^T \\ 0 < -c'(d;\theta) &< c_D \qquad \text{if } d = d^T \end{aligned} \tag{6}$$

These say that firms only complete early when their marginal benefits to delay reach zero;

[4]This normal error specification has the disadvantage that for sufficiently negative draws of $\widetilde{\theta}$, the marginal benefits of delay may be everywhere negative, which is implausible. Fortunately, given the parameter estimates below, this problem occurs only infrequently in the counterfactual simulations.

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

only complete late when their marginal benefits to delay equal the cost of delay c_D, and otherwise complete on time. This assumes that firms believe the time penalties will be charged with certainty. But we see in the data that these penalties are imposed only 25% of the time. For this reason, we also consider a version of the model where we use a probit to predict the probability that damages will be charged based on contract characteristics (type of work, size and duration), and then suppose that contractors believe that damages will be charged with this probability when making their completion time choice.[6]

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

$$\ell_i(\widetilde{d}|x,\widetilde{\theta}) = \begin{cases} \log\left(\phi\left(\frac{-x\beta-\alpha\widetilde{d}-\delta\widetilde{\theta}}{\sigma}\right)\right) - \log(\sigma) + \log(\alpha) & , \widetilde{d} < 1 \\ \log\left(\Phi\left(\frac{(c_D/d^T)-x\beta-\alpha\widetilde{d}-\delta\widetilde{\theta}}{\sigma}\right) - \Phi\left(\frac{-x\beta-\alpha\widetilde{d}-\delta\widetilde{\theta}}{\sigma}\right)\right) & , \widetilde{d} = 1 \\ \log\left(\phi\left(\frac{(c_D/d^T)-x\beta-\alpha\widetilde{d}-\delta\widetilde{\theta}}{\sigma}\right)\right) - \log(\sigma) + \log(\alpha) & , \widetilde{d} > 1 \end{cases} \tag{7}$$

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

Identification: It is worth thinking through how the model is identified. After making the assumption that marginal benefits of delay scale with contract length, we have a considerable amount of variation in the normalized time penalty c_D/d^T. This variation pins down the slope of the marginal benefit of delay curve, α. Consider for example two similar contracts that have experienced a similar negative productivity shock, but have different normalized time penalties. Then the extent to which the contractor speeds up in the case with more costly time penalties tells us how expensive acceleration is, or conversely how valuable delay is on the margin. Once the slope is pinned down, the intercept of the delay curve $x\beta + \delta\widetilde{\theta}$ is identified by the distribution of time outcomes for a given set of covariates and shocks.

One concern is that the slope is identified solely off outcomes for on-time and late contracts: for example, we never see how much acceleration there would be if a bonus was awarded for being early. This implies that the marginal benefit of delay curve is locally identified

[6] One problem with this specification is that small contracts are never late in the data, and so in estimating the enforcement probability for these contracts, we are relying entirely on functional form.

around the targeted days, but to get marginal benefits for $\widetilde{d} < 1$, we are relying on the linear functional form. With this in mind, most of our counterfactual simulations require only this local identification.

Results and Model Fit: Our results are in Table 5. We find that marginal benefits of delay have significant and negative slope, and that they are increasing in the size of the shock. The latter is consistent with the theory in equation (3): when there is more work to be done, completing on the same schedule would require a higher work rate, and thus higher wage costs. It follows that delay is more valuable. We also see that contracts with a higher planned work rate have significantly higher marginal benefits of delay, which makes sense.

In the third and fourth columns we allow the error standard deviation σ to vary with the contract size, specifying that σ is a linear function of the log contract value. This is because in the data smaller contracts have more variable outcomes. For an average contract, which lasts 40.6 days, our preferred estimate of the marginal benefit of delay — given in the third column of the table — is $10191 - 9009d$, implying that the average marginal benefit to delay with on time completion is \$1,182. This compares with an average cost of delay of \$1,004, so that for many contracts the time penalties will actually bite.

The fourth column assumes random rather than full enforcement. As we said earlier, we predict the probability of enforcement using a probit, and then compare the expected time penalties to the marginal benefits of delay in the first order condition. None of the observables is a statistically significant predictor of enforcement, although there is some evidence that in bigger contracts enforcement is more lax. When we account for these differential enforcement probabilities, we estimate marginal benefits to delay that are about a third as large as before, because with weak enforcement the only reason to finish on time is that the benefit of delay is small.

We also report how a one standard deviation productivity shock affects the marginal benefit of delay. The estimated impact is large: a negative shock of that magnitude raises the delay benefit by \$3,233 in our preferred specification. This is larger than the standard deviation of the error term, implying that the ex-post shocks explain more of the cost differences than unobserved factors do.

The model fits quite well, although there are some problems caused by assuming a normal distribution for the error term. We examine fit in a number of ways, using our preferred

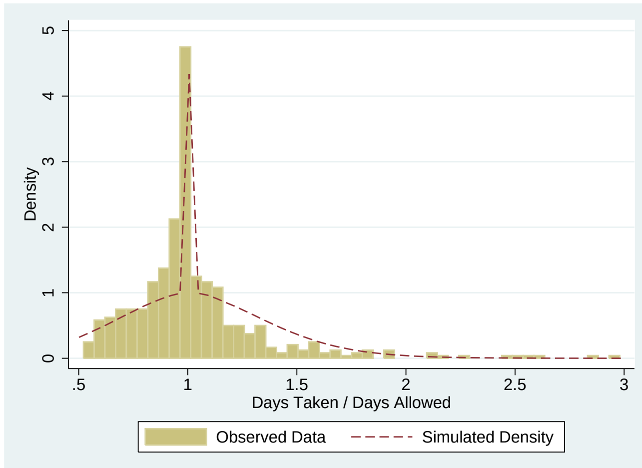

Figure 5: **Model Fit** The figure shows a histogram of the normalized completion times $\widetilde{d}$ in the data, along with a kernel density estimate of completion times simulated from the structural model.

specification, the full enforcement model of column 3. In Figure 5, we show a histogram of the normalized completion time against a kernel density plot of the completion times simulated from the structural model. This is intended as an informal sanity check on the shape of the distribution.[7] The model does a good job of capturing the incentives to be on time, but slightly under-predicts the probability that a contract will finish just early or just late, as it needs a large variance term to explain the contracts that finish very late.

A more rigorous test is to compare some sample and predicted moments. We do this in Table 6. The strengths and weaknesses of the model are quite clear. We do a good job on predicting how long the contract will take, even conditional on contract size. But the model over-predicts the fraction of contracts that will be late, and under-predicts how late they will be conditional on being late. This is probably a consequence of the thin tails of the normal error distribution. The model does a better job of fitting outcomes for big contracts, which is important for the counterfactuals because these contracts are typically the ones with the largest time delay costs to commuters.

Finally, we construct a simulated R^2 measure for the number of days taken d, by constructing the total sum of squares $TSS = \sum_{i=1}^{n}(d_i - \hat{E}[d_i])^2$, residual sum of squares $RSS = \sum_{i=1}^{n}(d_i - \hat{E}[d_i|x, d^T, c_D, \theta])^2$ and defining $R^2 = 1 - \frac{RSS}{TSS}$ as usual. The model does very well on this measure, explaining 87% of the variance in completion times. For comparison, a simple OLS regression with the same set of covariates (x, d^T, θ, c_D) only achieves an R^2 of 66%. This is

[7]This is informal because we are comparing multiple simulated outcomes for each contract to a single outcome for each contract in the data.

encouraging, as it shows that the structural model actually fits better than a simple reduced form, which need not be the case.

5.2 Counterfactuals

Now that we have contractor costs, we consider four counterfactual policy changes. First, we consider tightening the deadline to 90% of the current target, without changing the time penalties. This would be a relatively simple policy to implement. The second policy is more contract specific: changing the penalties to 10% of the traffic delay cost of the contract. This policy is almost neutral with respect to average penalties: the existing average daily penalty on a contract in the sub-sample is \$1,230, while under this counterfactual policy it would be \$1,520. The third policy is a "lane rental contract", where the contractor pays a penalty each day from the beginning of the contract. This penalty is set equal to 10% of the traffic delay cost. Finally, the last policy is the (risk-neutral) welfare maximizing policy of a lane rental with penalty equal to the full traffic delay costs.

The last two policies belong to a class of constrained efficient policies that maximize welfare subject to the constraint that the total costs to the contractors not exceed a certain amount. All members of this class use lane rentals with penalties equal to a fraction of the traffic delay cost, where the fraction depends on the cost constraint. This class is of interest because Mn/DOT faces a budget constraint of its own, and any costs it imposes on contractors through stronger time incentives will be passed-through, possibly at a rate higher than one.

For the simulations, we use the parameter estimates from our preferred specification, given in column 3 of Table 5. This assumes that contractors expect full enforcement. If we instead used the partial enforcement estimates for our simulations, the estimated marginal benefit of delay would be lower, and therefore in our simulations contractors would be more responsive to more aggressive policies. Assuming full enforcement thus leads to conservative estimates of the welfare gains from such policies.

We simulate both productivity shocks and new error draws in our counterfactuals. We do this to see how these policies fare when each contract is exposed to a full range of shocks. The procedure for obtaining the counterfactual estimates is as follows. First, we calculate the mean counterfactual completion times for each contract by simulating the above draws, and then compute the mean optimal completion time under the new incentives. This allows us to calculate the commuter gain from shorter construction, the additional private costs incurred by the contractor in accelerating the contract, and the penalties paid by the contractor. In all cases, these are measured relative to a baseline with no penalties for late completion. We also calculate a welfare gain, defined as the difference between the gain to commuters and the increased costs incurred by contractors. We then average over 10,000 simulations and across contracts.

The results are in Table 7. We find that neither of the small tweaks to existing policies has a big impact on completion times. For example, a tighter deadline implies that the contracts will on average be completed less than half a day sooner. And for many of the contracts 10% of the traffic delay cost is smaller than the current penalties. As a result, the days taken actually increases marginally with the new penalties, although it falls with the introduction of a tighter deadline or a lane rental.

Nonetheless, these more tailored penalties are effective: we see a welfare increase relative to the current policy of around \$51,000 per contract, or about 2.7% of the contract value. This is due to the large heterogeneity in traffic delay costs, which imply that a "one-size-fits-all" penalty regime does a poor job of aligning social and private costs. It is therefore not surprising that in the final column we see large gains to even the 10% lane rental policy, which aligns incentives not only when the contract is running late, but also when things are going well. The gains relative to the current policy are nearly \$200,000 per contract, or about 11% of the contract value. The projected gains from the fully efficient lane rental are substantially bigger.

One concern is that the subsample of contracts for which we have estimated traffic delay costs is not representative of our full sample. To address this, we use propensity score re-weighting. First, we use a probit to estimate the probability that a contract is in the subsample given contract size, duration and type of contract. Second, we weight the relevant counterfactual statistics for each observation in the subsample by the inverse of the estimated probability of inclusion. If the unobservables have the same distribution across the full sample and subsample, this will correct for selection. Looking at the bottom panel of Table 7 we see that the results are similar in percentage terms, though smaller in magnitude. For example, the welfare gain from increasing the penalty is predicted to be \$39,000 per contract, or 2.9%.

We also report the standard deviation of the penalties paid by contractors, across productivity and error draws. While the mean penalty is just a transfer from the contractor to the procurer and doesn't enter welfare calculations even under risk aversion, the standard deviation is a measure of the riskiness of the contract, and matters. We find that these are in general pretty comparable across the first four policies, although they sharply increase in the case of the fully efficient lane rental. To put the magnitude into context, the $47,000 standard deviation in payments from the full lane rental means that if the contractor has a markup of 10% on the average $1.3 million contract, nearly a third of their profit could be wiped out by a one standard deviation negative shock. This might be a barrier to entry for small firms.

Discussion: How reasonable are our estimates? To get a sense of this, we compare our results with those in Lewis and Bajari (2011). That paper examines what happened when the California Department of Transportation offered stronger time incentives through the use of scoring auctions. When the time incentives offered were equivalent to a lane rental of roughly $14,000 per day on much bigger contracts (nearly $22M on average), contractors completed in approximately 60% of the allotted time. Our most ambitious counterfactual — the efficient lane rental — predicts that lane rentals of similar magnitude (around $15,000 per day) applied to much smaller contracts $1.8M in our delay subsample) would have bigger effects (43% completion time). This doesn't seem entirely unreasonable, which is re-assuring.

Our first conclusion is that using a standardized set of time penalties, as Mn/DOT does, is far from optimal. As we showed, by setting contract-specific penalties —without changing deadlines or the average penalty — there could be a welfare gain of 1.9% of contract value. For our full sample, this amounts to $15.2 million. So our first conclusion is that tailored incentives would have a substantial quantitative impact on total welfare.

There is also broad heterogeneity in the ex-post marginal benefits of delay across contractors, both because of ex-ante asymmetries in their production functions, and ex-post productivity shocks. Because of this, contractors would sometimes be willing to finish early if incentivized to do so. The efficient lane rentals give such an incentive, and therefore produce much higher welfare gains — for our sample, a projected gain of $199.5 million. Our second conclusion then is that time incentives really should be structured to provide a constant incentive for accelerated construction.

There are some important limitations of the above analysis. In some cases we are relying on the functional form of the parametric model. This is particularly true of the lane rental policy, where the marginal benefits of delay prior to the contract deadline are inferred based on the slope α of that function, which in turn was identified from late contracts. If the slope is not constant (e.g. marginal benefits are decreasing) we may underestimate the benefits to delay, and hence overestimate the effectiveness of the policy. To address these concerns we have mainly focused on small counterfactual policy changes, so that if the function is locally linear around $\widetilde{d} = 1$, our estimates will be approximately correct.

We also ignore the fact that under the new contract regime, contractors may be selected for on the basis of their ability to complete quickly, and therefore the winning bidders may actually have lower marginal costs than those we estimated earlier (although they may have higher fixed costs). Moreover, they may make different decisions with respect to the hiring of labor and rental of capital than they currently do, enabling them to complete quicker without incurring the high costs we project when we hold labor and capital decisions fixed. This suggests we may underestimate the welfare gains from more high-powered time incentives.

But with these incentives comes an increased need for monitoring the contractors. This may be costly, since Mn/DOT would have to employ additional personnel on site. At least for the small changes considered here, this seems unlikely to be necessary. Finally, we do not account for general equilibrium effects due to the bidding up of input prices, or for the potential deadweight loss incurred in raising the funds for this accelerated procurement.

6 Conclusion

Blunt instruments can sometimes be effective tools. Time incentives are a relatively blunt instrument: instead of contracting on all the events that may require a contractor to adapt his plans, the contract simply specifies a deadline and penalties. What this paper has shown is that these incentives are sufficient to motivate adaptation, at least when the design needn't be amended.

Nonetheless, even blunt instruments need to be designed with care. As the paper has shown, thoughtful deadlines are meaningless without appropriate penalties. In regulating the externality from highway construction, penalties that are proportional to the traffic delay cost perform much better than a "one-size-fits-all" statewide policy.

This paper has focused on highway contracts, because they are cases where the design is close to fully specified and adaptation is generally one-sided. This leaves room for interesting empirical research on cases where the design itself may need adaptation in the face of a shock, such as in military procurement or less routine construction.

References

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

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

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

– **, Robert McMillan, and Steven Tadelis**, "Auctions versus Negotiations in Procurement: An Empirical Analysis," *Journal of Law, Economics and Organization*, 2009, *25* (2), 372–399.

– **, Stephanie Houghton, and Steven Tadelis**, "Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs," 2007. Working Paper, University of Minnesota.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul, "Social Preferences and the Response to Incentives: Evidence from Personnel Data," *Quarterly Journal of Economics*, August 2005, *120* (3), 917–962.

Chetty, Raj, John Friedman, Tore Olsen, and Luigi Pistaferri, "Adjustment Costs, Firm Responses and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records," *Quarterly Journal of Economics*, 2011, *126* (2), 749–804.

Chevalier, Judy and Glenn Ellison, "Risk Taking by Mutual Funds as a Response to Incentives," *Journal of Political Economy*, 1997, *105* (6), 1167—1200.

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

De Silva, Dakshina G., 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.

European Commission, "A report of the functioning of public procurement markets in the EU: benefits from the application of EU directives and challenges for the future," Technical Report, European Commission 2004.

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.

Holmstrom, Bengt and Paul Milgrom, "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 1987, *55* (2), 303–328.

– **and** – , "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership and Job Design," *Journal of Law, Economics and Organization*, 1991, *7*, 24–52.

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

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

Krasnokutskaya, Elena, "Identification and Estimation in Highway Procurement Auctions under Unobserved Auction Heterogeneity," *forthcoming in The Review of Economic Studies*, 2009.

– **and Katja Seim**, "Bid Preference Programs and Participation in Highway Procurement Auctions," *forthcoming in American Economic Review*, 2010.

Laffont, Jean-Jacques and Jean Tirole, "Using cost observation to regulate firms," *Journal of Political Economy*, 1986, *94*, 614—641.

– **and** – , "Auctioning Incentive Contracts," *The Journal of Political Economy*, 1987, *95* (5), 921–937.

– **and** – , *A Theory of Incentives in Procurement and Regulation*, The MIT Press, 1993.

Lazear, Edward P., "Performance Pay and Productivity," *American Economic Review*, December 2000, *90* (5), 1346–1361.

Levin, Jonathan and Steven Tadelis, "Contracting for Government Services: Theory and Evidence from U.S. Cities," *Journal of Industrial Economics*, 2010, *58* (3), 507–541.

Lewis, Gregory and Patrick Bajari, "Procurement Contracting with Time Incentives: Theory and Evidence," *Quarterly Journal of Economics*, August 2011, *126* (3).

Li, Tong and Xiaoyong Zheng, "Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions.," *Review of Economic Studies*, 2009, *76* (4), 1397 – 1429.

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

Martimort, David and Jerome Pouyet, "To build or not to build: Normative and positive theories of public–private partnerships," *International Journal of Industrial Organization*, 2008, *26* (2), 393—411.

Maskin, Eric and Jean Tirole, "Public-Private Partnerships and Government Spending Limits," *International Journal of Industrial Organization*, 2008, *26* (412—420).

McAfee, R. Preston and John McMillan, "Bidding for contracts: a principal-agent analysis," *RAND Journal of Economics*, 1986, *17* (3), 326–338.

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

Prendergast, Canice, "The Provision of Incentives in Firms," *Journal of Economic Literature*, 1999.

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

Saez, Emmanuel, "Do Taxpayers Bunch at Kink Points?," *American Economic Journal: Economic Policy*, 2010, *2*, 180–212.

Weitzman, Martin L., "Prices vs. Quantities," *Review of Economic Studies*, 1974, *41* (4), 477–491.

World Bank, "Country Procurement Assessment Report for South Africa," Technical Report, World Bank 2003.

7 Appendix

Table 1: Damage Specifications and Time Outcomes

	Damages per day ($)		Time Outcomes	
Contract Value ($)	1995-2004	2005-	# Obs	# Late
Below 25K	75	150	1	0
25K-50K	125	300	6	0
50K-100K	250	300	22	0
100K-500K	500	600	145	36
500K-1M	750	1000	102	38
1M-2M	1250	1500	118	50
2M-5M	1750	2000	73	34
5M-10M	2500	3000	23	15

Damage specifications are taken from the Mn/DOT standard specifications for construction contracts. These were last amended in 2005. Time outcomes are based on the diary data.

Table 2: Summary Statistics: Mn/DOT Highway Construction Contracts, 1996-2006

	Mean	Std. Dev.	Min	Max
Contract Value (thousands)	1336.1	1551.7	15.84	9549.7
Days Allowed	40.61	30.05	10	198
Days Taken	43.48	37.40	5.900	271
Planned Hours	391.7	313.1	80	1946.0
Hours Worked	462.5	453.6	29	3290
Avoidable Delays (hours)	70.03	100.2	0	808
Unavoidable Delays (hours)	99.10	142.9	0	1099
Work Rate (hours/day)	10.18	3.470	1.840	24
Contract late?	0.353	0.478	0	1
Penalty applied if late?	0.243	0.430	0	1
Damages (if applied)	6097.8	6036.6	500	27000
Types of Work (not mutually exclusive)				
Beautification	0.0306	0.172	0	1
Bridge Repair	0.104	0.306	0	1
Construction	0.0776	0.268	0	1
Resurfacing	0.898	0.303	0	1
Shoulders / Turn Lanes	0.192	0.394	0	1
Traffic Signals	0.137	0.344	0	1
Delay Subsample				
Daily Traffic	11688.1	15960.0	380	100000
Projected delay (mins)	9.727	7.115	0	20
Daily Traffic Delay Cost	15164.7	20492.6	0	124000
Number of contracts	490			

Summary statistics for selected highway construction contracts let by the Minnesota Department of Transportation between 1996 and 2006. The sample includes all working day contracts for which we have diary data (20 percent of all contracts). The omitted type of work category is "other". "Unavoidable" delays are those that were outside of the contractor's control; "avoidable" delays are those that were preventable in the judgement of the project engineer. "Beautification" work includes landscaping, painting and minor repairs. The delay subsample consists of 99 contracts for which we collected more detailed data on detour options, and projected a daily cost to commuters from traffic delays.

Table 3: Shocks and Work Rate Adaptation

	Work Rate			
	OLS		IV	
	(1)	(2)	(3)	(4)
Shock	0.560***	0.533***	0.404***	0.271***
	(0.049)	(0.057)	(0.103)	(0.102)
Planned Work Rate	0.777***	0.777***	0.647***	0.545***
	(0.066)	(0.077)	(0.104)	(0.114)
Time Penalty / Days Allowed	0.012*	0.018*	0.017**	0.007
	(0.007)	(0.010)	(0.008)	(0.011)
Contract Value / (Days Allowed * 1000)		0.010		0.040***
		(0.009)		(0.013)
Big Firm		0.239		-0.010
		(0.214)		(0.250)
In-State Contractor		0.159		0.073
		(0.328)		(0.323)
N	490	490	490	490
First-stage F-statistic	-	-	29.48	20.94
District/Work/Year FE	no	yes	no	yes

The dependent variable is the work rate in hours per day. The shock variable is the difference between actual and planned hours, divided by the days allowed. Big Firm is a dummy variable indicating that the contracting firm won more than $10 million of contracts over the sample period. The first two columns report results from an OLS regression, while the last two report results from an IV regression where the contract shock is instrumented for by the hours of rain and hours of snow per day. Columns (2) and (4) include year, type of work and district fixed effects. Standard errors are robust. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$). The F-statistic reported for the IV regressions is from an F-test of the joint significance of the instruments.

Table 4: Determinants of Late Contract Completion

	Indicator for Contract Late			
	(1)	(2)	(3)	(4)
Time Penalty / Days Allowed	-0.004**	-0.007***	-0.007***	-0.006***
	(0.002)	(0.002)	(0.002)	(0.002)
Contract Value / (Days Allowed * 1000)	0.006***	0.006***	0.006***	0.005***
	(0.001)	(0.001)	(0.001)	(0.001)
Big Firm			-0.054	-0.026
			(0.045)	(0.045)
In-State Contractor			0.100	0.080
			(0.061)	(0.058)
Hours of rain per day				0.553***
				(0.148)
Hours of snow per day				1.509*
				(0.772)
N	490	490	490	490
R^2	0.04	0.16	0.16	0.21
District/Work/Year FE	no	yes	yes	yes

The dependent variable is the probability that the contract is completed late. All columns are OLS regressions, and standard errors are robust. Significance levels are denoted by asterisks (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$). Big Firm is a dummy variable indicating a firm that won more than $10M in total contract value in the sample. Marginal effects from a probit specification are very similar.

Table 5: Estimates of Marginal Benefit of Delay

	Days Taken / Days Allowed			
	Full Enforcement			Random Enforcement
	Raw Coefficients			
Log Contract Size (β_1)	-4.49	-2.59	-2.08	-0.81
	(3.85)	(3.50)	(4.55)	(1.41)
Planned Work Rate (β_2)	8.29***	6.87***	6.76***	2.08***
	(1.95)	(1.76)	(1.83)	(0.58)
Big Firm (β_3)	2.23	1.56	1.03	-0.73
	(9.55)	(6.79)	(7.61)	(2.96)
In-State Contractor (β_4)	0.26	-2.90	-5.06	-3.32
	(6.08)	(6.24)	(6.62)	(2.01)
Shock (δ)	13.00***	12.14***	12.15***	3.65***
	(1.85)	(1.69)	(1.70)	(0.54)
Slope (α)	-238.75***	-221.89***	-221.84***	-68.69***
	(42.02)	(38.95)	(39.26)	(11.89)
Std. Deviation of Error (σ)	78.57***	70.23***	70.18***	21.63***
	(10.23)	(8.80)	(8.81)	(2.66)
	Estimated Benefit Function for an Average Contract			
Intercept	11037.23	10184.37	10191.13	3122.82
	(1827.26)	(1707.88)	(1736.87)	(529.21)
Slope	-9695.28	-9010.46	-9008.54	-2789.39
	(1706.28)	(1581.77)	(1594.24)	(483.03)
Std. Deviation of Productivity Shock	3458.23	3230.15	3232.69	971.72
	(492.17)	(449.11)	(452.71)	(142.86)
Std. Deviation of Error	3190.40	2851.97	2859.30	881.02
	(415.28)	(357.39)	(422.18)	(124.39)
District FEs	no	yes	yes	yes

Maximum likelihood estimates of the marginal benefit of delay function. The top panel has the raw coefficients; the bottom panel shows the implied function for an average contract. An average contract has value \$1.3M, should be completed in 40 days and has a penalty of \$1K per day. In the third and fourth columns, the standard deviation of the error term is specified as a linear function of the log contract size. The first three columns are from a model that assumes full enforcement of the time penalties; the last column assumes random enforcement, with the probability of enforcement estimated as a function of the contract characteristics. Big Firm is a dummy variable indicating a firm that won more than \$10M in total contract value in the sample. The standard deviation of the productivity shock indicates the effect of a 1 standard deviation increase in the required hours of work per day on the marginal benefit of delay for an average contract. District fixed effects are included where indicated. Standard errors are bootstrapped. Asterisks in the top panel denote significance levels (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

Table 6: Sample and Simulated Moments

	Data	Simulations			
	Mean	Mean	Std. Error	95% L.B.	95% U.B.
Days Taken / Days Allowed	1.04	1.03	0.01	1.01	1.05
Days Taken / Days Allowed if value >$1M	1.10	1.11	0.02	1.07	1.14
Days Taken / Days Allowed if value <$1M	1.00	0.97	0.01	0.94	1.00
Fraction of Contracts Late	0.37	0.45	0.02	0.42	0.48
Fraction Late if value >$1M	0.50	0.54	0.02	0.50	0.58
Fraction Late if value <$1M	0.28	0.39	0.02	0.34	0.43
Days Taken / Days Allowed if late	1.33	1.26	0.02	1.22	1.30
Days Taken / Days Allowed if late & value >$1M	1.30	1.30	0.03	1.25	1.35
Days Taken / Days Allowed if late & value <$1M	1.36	1.23	0.03	1.18	1.29
	Structural Model			OLS	
Simulated R^2 for days taken	0.87			0.66	

Comparison of observed and simulated moments. The first column contains moments observed in the data; the second column contains moments simulated using the maximum likelihood estimates of the model parameters. Columns 3—5 respectively contain the bootstrapped standard errors, and lower and upper bounds of a 95% confidence interval for these moments. The bottom part of the table shows a simulated R^2 for days taken, computed from the total sum of squares and residual sum of squares, where the residuals are the difference between the observed and mean simulated outcome. The OLS comparison uses OLS to predict the days taken, using all the covariates from the structural model (including a constant).

Table 7: Counterfactual Welfare Estimates under Alternative Policies

	Current Policy	10% Tighter Deadline	New Penalties (10% Delay Cost)	Lane Rental (10% Delay Cost)	Lane Rental (100% Delay Cost)
Subsample with Delay Data					
Days Taken	46.84	46.33	46.91	43.57	20.02
	(45.52,48.26)	(45.03,47.80)	(45.63,48.28)	(42.04,45.42)	(17.48,23.63)
Commuter Gain ($K)	64.73	73.92	118.57	274.16	790.88
	(45.09,85.00)	(52.17,97.01)	(88.19,148.23)	(195.84,347.80)	(726.19,844.13)
Acc. Cost ($K)	2.58	2.96	4.71	13.23	184.31
	(1.81,3.38)	(2.11,3.87)	(3.61,5.78)	(9.68,16.34)	(168.08,207.04)
Penalties Paid ($K)	13.06	17.19	8.95	62.82	111.49
	(10.54,16.06)	(14.56,20.34)	(7.03,11.74)	(56.77,69.76)	(85.59,157.57)
Std. Dev. Penalties ($K)	15.95	17.93	12.72	26.50	67.69
	(13.01,18.68)	(14.99,20.61)	(10.19,15.82)	(22.33,29.93)	(55.25,90.63)
Net Gain ($K)	62.16	70.96	113.85	260.93	606.57
	(43.26,81.62)	(50.08,93.12)	(84.44,142.62)	(186.14,331.48)	(525.23,666.29)
Reweighted to match full sample					
Days Taken / Allowed	41.51	41.07	41.58	38.67	19.56
	(40.57,42.62)	(40.13,42.25)	(40.64,42.66)	(37.52,40.09)	(17.70,22.18)
Commuter Gain ($K)	45.60	52.29	86.08	201.11	583.68
	(31.62,60.57)	(36.84,68.60)	(63.47,107.13)	(142.66,256.50)	(538.10,618.51)
Acc. Cost ($K)	1.85	2.14	3.44	9.76	132.55
	(1.29,2.44)	(1.52,2.81)	(2.68,4.19)	(7.08,12.15)	(119.83,151.84)
Penalties Paid ($K)	9.64	12.75	6.22	46.30	80.48
	(7.76,11.64)	(10.90,14.87)	(4.94,7.96)	(42.06,51.43)	(64.69,110.23)
Std. Dev. Penalties ($K)	12.05	13.62	9.21	19.96	47.05
	(9.94,13.96)	(11.46,15.51)	(7.42,11.20)	(16.71,22.34)	(38.83,62.56)
Net Gain ($K)	43.75	50.15	82.63	191.35	451.13
	(30.34,58.14)	(35.32,65.78)	(60.77,102.83)	(135.58,244.35)	(392.99,494.23)

Counterfactual welfare results under different policies. The top panel of results is estimated using the data subsample for which we have social cost measures; the second panel imputes counterfactual results for the full sample by reweighting the counterfactual moments for each contract by the inverse of the probability of appearing in the subsample. This probability is estimated by a probit, conditioning on the observables used elsewhere in the paper. The first column is simulated outcomes under the current policy. In the second column, the contract deadline is accelerated by 10%. In the third, the daily penalties for late completion are set equal to 10% of the traffic delay cost. In the fourth and fifth columns, a lane rental policy is simulated, with rental rates equal to 10% and 100% of the traffic delay cost respectively. All statistics are calculated relative to a scenario where no time incentives are used. For example, "Commuter Gain" for any contract is the average number of days saved by using time incentives times the traffic delay cost. "Acceleration Cost" is the estimated additional cost to the winning contractor of accelerating construction relative to a no time incentives scenario. "Penalties" refers to any penalties the contractor pays, "Std. Dev Penalties" is the standard deviation of those penalties across simulated draws of productivity and cost shocks (averaged over contracts), while "Net Gain" is the difference between commuter gain and contractor acceleration costs. In all cases, mean results are averaged across simulations and contracts. 95% confidence intervals are given in parentheses, and are generated by bootstrapping the regressions and taking the 2.5th and 97.5th percentiles of the simulated results based on the bootstrapped coefficients.