

Self-Control at Work

Supreet Kaur

Columbia University and National Bureau of Economic Research

Michael Kremer

Harvard University and National Bureau of Economic Research

Sendhil Mullainathan

Harvard University and National Bureau of Economic Research

Self-control problems change the logic of agency theory by partly aligning the interests of the firm and worker: both now value contracts that elicit future effort. Findings from a year-long field experiment with full-time data entry workers support this idea. First, workers increase output by voluntarily choosing dominated contracts (which penalize low output but give no additional rewards for high output). Second, effort increases closer to (randomly assigned) paydays. Third, the contract and payday effects are strongly correlated within workers, and this correlation grows with experience. We suggest that workplace features such as high-powered incentives or effort monitoring may provide self-control benefits.

I. Introduction

Agency theory emphasizes a tension between workers and firms. Because employers provide insurance, workers do not benefit fully from their effort. This creates moral hazard: workers do not work as hard as the

We thank Lawrence Katz, David Laibson, John Matsusaka, Derek Neal, Ted O'Donoghue, three anonymous referees, and participants at numerous seminars for helpful comments. Raluca Dragusanu provided outstanding research assistance. We thank Dominik Bulla, Karthik R, and the Centre for Innovative Financial Design for operational

Electronically published October 30, 2015
[*Journal of Political Economy*, 2015, vol. 123, no. 6]
© 2015 by The University of Chicago. All rights reserved. 0022-3808/2015/12306-0001\$10.00

employer would like (Holmstrom 1979; Grossman and Hart 1983). Introspection suggests another tension at work. Self-control problems mean that workers often do not work as hard as they themselves would like.¹ Looking to the future, they would like to work hard; when the future arrives, they may end up slacking.²

Self-control at work thus differs from self-control in other contexts, such as savings or smoking (Laibson 1997; Ashraf, Karlan, and Yin 2006; Giné, Karlan, and Zinman 2010). In many of these other contexts, the market provides commitment only if the consumer demands it sufficiently so as to create a new institution. Worker self-control problems, however, hurt employer profits directly. As a result, both the firm and the employee have self-interest in curbing them. The workplace exists to organize effort provision by its workers. The same features that mitigate moral hazard— incentive contracts and job design features such as fixed hours of work— can also mitigate self-control problems. In other words, the employer has both the means and motives to (implicitly) provide commitment devices.³

We build a simple model to create testable predictions from these ideas.⁴ The model provides one stark prediction. In agency models, workers must be compensated for a sharpening of incentives: they face extra risk. With self-control, no compensation may be needed. Workers who are aware of their self-control problem (sophisticated in the sense of O'Donoghue and Rabin [1999]) will value sharper incentives as a way to motivate future selves. As a result, sophisticated workers may voluntarily choose a dominated contract, one that pays less for a low-output realization and the same for high-output realizations.⁵

support and Abdulla Al-Sabah, Dan Bjorkegren, Matthew Darling, Nicolas Papamichael, Frank Schilbach, and Zhaoning Wang for excellent assistance with specific parts of the project. We gratefully acknowledge funding from Harvard University, the International Finance Corporation, and the National Science Foundation. All errors are our own. Data are provided as supplementary material online.

¹ Frederick, Loewenstein, and O'Donoghue (2002) and DellaVigna (2009) review the self-control literature. Prominent models include Laibson (1996, 1997), O'Donoghue and Rabin (1999, 2001), and Fudenberg and Levine (2006). Banerjee and Mullainathan (2010) and Bernheim, Ray, and Yeltekin (2013) examine self control in the development context. Gul and Pesendorfer (2001, 2004) provide a different account of the demand for commitment.

² Baicker, Mullainathan, and Schwartzstein (2012) refer to this as *behavioral hazard*. In moral hazard, inefficiencies arise because people face the wrong "price." In behavioral hazard, psychology generates inefficiencies even when facing the right price.

³ This suggests a potential additional rationale for organizing production in firms, in addition to providing workers with incentives or addressing free riding in team production (Cheung 1969; Alchian and Demsetz 1972).

⁴ O'Donoghue and Rabin (1999, 2006) formalize how firms use deadlines to motivate procrastinators and produce interesting implications for screening. Kaur, Kremer, and Mullainathan (2010) provide a general discussion of how self-control may affect work arrangements. DellaVigna and Malmendier (2004), Eliaz and Spiegler (2006), and Spiegler (2011) study contract design in other contexts.

⁵ This prediction (and our model) presumes that there is a limited availability of other external devices to help workers with their self-control problems.

In addition to predicting demand for dominated contracts, our model also suggests that the timing of pay affects effort. As the payday gets closer, the source of the self-control problem diminishes: the rewards of work and the cost of work are closer together in time. As a result, output should increase. The model also suggests an important role for heterogeneity. Workers with greater self-control problems should show both larger payday effects and a greater desire for dominated contracts.

The primary contribution of this paper is a 13-month field experiment to test these predictions.⁶ A growing literature emphasizes the importance of natural environments, realistically high stakes, and sufficiently long durations in experimental tests of theory (Levitt and List 2007).⁷ In the experiment, full-time data entry workers in India are paid on the basis of the number of accurate fields entered each day. First, to test the demand for dominated contracts, on random days workers were offered the option to choose a target for the day: if they met the target, they received the standard piece rate; if they fell short of the target, they received only half the piece rate for their output. On average, workers selected positive targets—which correspond to choosing dominated contracts—36 percent of the time. The option to choose a dominated contract increases production and earnings, with a treatment on the treated effect of 6 percent for those workers who accept the dominated contract. A production increase of this size corresponds to that induced by an 18 percent increase in the piece rate wage (computed using exogenous wage changes). We show that this is a lower bound on the extent of time inconsistency. It implies that workers value the net benefits of future effort by at least 18 percent more at the time of contract choice than at the time they actually exert that effort.

Second, to test the impact of paydays, workers were randomized into different payday groups: all were paid weekly but the exact day of pay-

⁶ Ariely and Wertenbroch (2002) test for procrastination in effort by hiring university students to proofread text over 3 weeks with a maximum total payout of \$10. They find that allowing students to self-impose intermediate deadlines raises performance. Burger, Charness, and Lynham (2009) pay students to study and examine the impact of externally imposed deadlines and the relationship with willpower depletion; the deadlines actually lower performance. However, since students are not given the option to self-impose penalties for procrastination, the results are more difficult to interpret. More generally, there is scant empirical evidence on self-control in realistic workplace settings involving nonstudent populations, high stakes (i.e., full-time earnings), and long durations. Two recent papers document results consistent with time inconsistency among bank workers (Cadena et al. 2011) and bicycle taxi drivers (Dupas and Robinson 2013).

⁷ In pioneering work on this point, List (2006) shows that while experienced sports card traders exhibit gift exchange in the lab, this effect is strongly attenuated in an actual market environment. In a field experiment on gift exchange, Gneezy and List (2006) find that treatment effects on worker effort wear off after a few hours, suggesting that short-run responses can provide misleading estimates. A growing literature uses field experiments to test features of worker effort other than self-control (e.g., Shearer 2004; Bandiera, Barankay, and Rasul 2007; Fehr and Goette 2007; Hossain and List 2012). For excellent reviews of this literature, see Levitt and List (2009) and List and Rasul (2011).

ment varied. Worker output is 8 percent higher on payday than at the beginning of the weekly pay cycle. An effect of this magnitude corresponds to a 24 percent increase in the piece rate or about one additional year of education in our sample. A calibration of our model suggests that the pay cycle effect cannot be explained in an exponential discounting framework; it requires an exponential discount rate of 4 percent per day or 1.65×10^6 percent per year.

Third, we find substantial heterogeneity in the extent of the payday and contract effects. This heterogeneity is in fact predictive: workers with above-mean payday effects are 49 percent more likely to choose dominated contracts. Providing these workers the option to choose a dominated contract increases their output by 9 percent, implying a treatment on the treated effect on output of 28 percent for those workers who select the contract. This implies that for these workers, dominated contracts have production impacts comparable to an 85 percent increase in the piece rate wage.⁸

Fourth, the option to choose dominated contracts has bigger treatment effects when the payday is further away. This is consistent with the fact that the self-control problem is less severe closer to the payday, and the dominated contract therefore has less scope to affect effort.

While these results broadly support self-control models, one finding does not. Workers are no more likely to select dominated contracts for the more distant future than they are for the nearer future: take-up on the morning of the workday and the evening before is similar. Ex post analysis suggests a possible reason: workers face output uncertainty—for example, from network speed fluctuations or uncertain commute times—that is (partly) resolved when they arrive to the office. When such uncertainty is low, workers are indeed more likely to demand targets the evening before work than the morning of work.

We also find evidence of learning. While payday effects do not change with experience, the demand for dominated contracts does. Early on, many workers experiment with these contracts when offered the option. As they gain experience, the correlation between payday effects and choice of dominated contracts increases. After 2 months of experience, workers with high payday effects are 20 percentage points (73 percent) more likely to select dominated contracts than workers with low payday effects.

Note that the dominated contract is merely a construct that precisely isolates the demand for self-control in an experiment. We are not arguing that giving employees such choices is necessarily the optimal contract for

⁸ Although we find a strong correlation between the payday and contract choice effects, neither effect is well predicted by conventional “lab experiment” measures of time inconsistency, such as subjects’ choices among cash payments at different times. These results line up well with the interesting lab results of Augenblick, Niederle, and Sprenger (2014), who find present bias in effort tasks but fail to find it for cash discounting tasks among student subjects.

time-inconsistent workers. Our results indicate that workers will demand incentives that help them overcome self-control problems, but such incentives could take a variety of forms. In our model, the dominated contract helps solve the self-control problem by creating high-powered incentives around a discrete threshold; this general feature is a common ingredient in contracts in a wide variety of settings. Many firms provide discrete bonuses for meeting targets, such as in sales in the United States (Oyer 2000; Larkin 2014) or data entry in India. In addition, most jobs have production minimums—such as a 40-hour week or an output requirement—below which the penalty is not a commensurate percentage loss in earnings but rather the threat of being fired altogether. In some instances, employers remove the worker's ability to choose certain dimensions of effort, for example, through rigid hours or, more extremely, assembly lines that make it impossible to slow one's pace. More traditional explanations, such as fixed costs or team production, are of course important in understanding these arrangements. Our results suggest that self-control considerations may also potentially be relevant.

While this paper tests for self-control problems, a companion paper (Kaur, Kremer, and Mullainathan 2014) examines how self-control problems affect equilibrium labor market contracts and job design. Present bias among workers generates higher-powered equilibrium incentives in a simple agency model with free entry of firms, and sufficient present bias reverses the standard result in agency theory that firms insure workers against risk at the expense of weaker incentives and hence lower output. Instead, the distribution of output may second-order stochastically dominate the distribution of wages, and hence present-biased workers may exert more effort and produce more output working in firms than as self-employed owner-operators. To the extent that workers are risk averse and hence dislike high-powered incentives on output, self-control problems make firms more likely to adopt costly technologies to monitor effort (such as banning telecommuting) and contractually obligate workers to put in a prespecified level of effort. If workers have heterogeneous time preferences, firms will face an adverse selection problem; in equilibrium, even workers who are not present-biased may have to accept contracts with higher-powered incentives or costly effort monitoring. In contrast to other models of equilibrium interaction between present-biased and time-consistent agents—in which present-biased agents are naive and hence can be exploited by sophisticated time-consistent agents (Della-Vigna and Malmendier 2004; Eliaz and Spiegel 2006; Gabaix and Laibson 2006)—the presence of present-biased agents makes time-consistent agents worse off.

The rest of the paper is organized as follows. Section II lays out a simple model of effort choice and demand for contracts by time-consistent and time-inconsistent workers. Section III explains the experimental design

and our context. Section IV presents results. Section V discusses possible alternative explanations. Section VI presents conclusions.

II. Choice of Contracts and Effort by Time-Consistent and Time-Inconsistent Workers

In this section, we use a simple model to derive empirically testable predictions to distinguish the behavior of time-consistent and time-inconsistent workers. In our experiment, workers receive piece rates based on their output plus a small show-up payment. Each day, some workers are offered a choice between two types of incentive contracts. The first is a linear piece rate contract. The second is a dominated contract, which pays less than the linear piece rate for low output levels but pays the same as the linear piece rate for high output levels. While contract choice is made daily, workers are paid weekly: on one day per week, they receive their cumulative earnings from the preceding 7 days. In what follows, we generate predictions on how our experimental setup distinguishes between time-consistent and time-inconsistent workers and enables us to calibrate the extent of time inconsistency.

Assume that worker i has the per-period utility function $y_t - \alpha^i c(e_t)$, where y_t is income received in period t , e_t is effort in period t , $c(\cdot)$ is the cost of effort, and $\alpha^i > 0$ reflects individual variation in effort costs. We will focus below on the case in which $c(e_t) = e_t^\theta$, with $\theta > 1$. However, as we discuss in the proofs, our propositions will hold to a first-order approximation under a more general $c(\cdot)$ that is increasing, convex, and twice differentiable, with $\lim_{e_t \rightarrow \infty} c'(e_t) = \infty$. For simplicity, we will also focus on the case in which $\alpha^i = 1$ for all workers.

We write $D^i(t)$ to denote worker i 's discount factor, where $D^i(t) \in \{D^C(t), D^I(t)\}$. Time-consistent workers discount the future using an exponential discount factor: $D^C(t) = \delta^t$. Time-inconsistent workers have a hyperbolic discount factor, $D^I(t)$: for any delay s , $D^I(t + s)/D^I(t)$ is strictly increasing in t .⁹ The time-inconsistent worker is at least as impatient as the time-consistent one: $D^I(t + 1)/D^I(t) \leq \delta$ for all t . We additionally assume that all workers are sophisticated: they know $D^i(t)$ and accurately predict their own future actions.

Timing and production.—There are T periods. In each period, the worker chooses effort e_t , which determines output that period. Each period also has its own distinct contract, which depends on output in that period. The contract for period t is signed k periods in advance in period $t - k$.

⁹ The hyperbolic discount function $D(t) = (1 + \mu t)^{-\nu/\mu}$ (Lowenstein and Prelec 1992) satisfies this property. Note that for a quasi-hyperbolic function (see Laibson 1997), $D'(t) = \beta\delta^t$, $D'(1)/D'(0) = \beta\delta$, and $D'(t + 1)/D'(t) = \delta$ for $t > 0$, so it satisfies this property for $t = 0$. In what follows, we model time-inconsistent agents with general hyperbolic preferences but also briefly discuss the case of quasi-hyperbolic preferences.

Output is a deterministic function of effort, and we choose units so that output equals effort. We begin by considering a simple case in which wages are an affine function of output: $w(e_t) = a + be_t$, where a is a base wage and b is the piece rate. (In proposition 2 below, we will also consider a more complicated contract.)

In period T , output is realized and workers are paid for the total output from their effort in periods 1 to T . Thus, income is 0 in all periods except for period T : $y_t = 0$ for all $t \neq T$ and $y_t = \sum_{s=1}^T w(e_s)$ for $t = T$.

The timing in the model and the experiment match to some degree. In the experiment, a period is a day and T is 6 days (a week of work). Then in both the model and the experiment workers (i) have a distinct contract each day and (ii) are paid on the basis of the sum of these contracts at the end of the pay week on day T .

But the timing in the model and the experiment do differ in two subtle but important ways. First, in the model, workers choose a single effort e_t for each period. In reality, workers constantly choose effort throughout the day. Second, contract choice in the model happens k periods in advance. In the experiment, it happens either in the morning before work or in the previous evening (after the previous day's work). To reconcile this with the model, it is convenient to think of each day as having three distinct time periods: before work, during work, and after work. Then, in the experiment, effort is exerted during work, and contract choice is made either $k = 1$ periods in advance (i.e., the morning before work) or $k = 2$ periods in advance (i.e., the previous evening). These differences, though notable, should not change the qualitative predictions of the model: what ultimately matters for our predictions is that pay (the rewards of effort) happens after effort is exerted and that contract choice is made before effort is exerted.

Optimal effort.—Note that under quasi-linear utility, optimal effort in each period is separable from effort choice in other periods. We therefore focus on the worker's choice for a particular period, t . In period $t - k$, when contemplating optimal effort in period t , worker i discounts the cost of period t 's effort by $D^i(k)$. Since payment occurs in period T , she discounts the payoff from that effort by $D^i(T - t + k)$. So in period $t - k$ the worker's preferred period t effort maximizes

$$\max_e \{D^i(T - t + k)w(e) - D^i(k)c(e)\}. \tag{1}$$

In contrast, the period t self will choose effort for that period according to

$$\max_e \{D^i(T - t)w(e) - D^i(0)c(e)\}. \tag{2}$$

Let $e_{t,s}^i$ denote the optimal effort in period t from the perspective of worker i 's period s self. From the perspective of period $t - k$, the optimal

effort in period t , $e_{i|t-k}^i$, is pinned down by the first-order condition (FOC) from expression (1):¹⁰

$$c'(e_{i|t-k}^i) = \frac{D^i(T-t+k)}{D^i(k)} b, \quad (3)$$

where b is the piece rate in the affine contract $w(e)$. The effort level actually exerted by the worker in period t , $e_{i|t}^i$, is given by substituting t for $t-k$ in FOC (3):

$$c'(e_{i|t}^i) = \frac{D^i(T-t)}{D^i(0)} b = D^i(T-t)b. \quad (4)$$

Note that since $D^i(T-t-1) > D^i(T-t)$, equation (4) implies that $e_{i|t+1}^i > e_{i|t}^i$ for all workers. Thus, output increases as the payment period approaches.

LEMMA 1 (Equivalence of discounting and piece rate changes). Let $e_{i|t+1}^i$ denote the effort level chosen in period $t+1$ under piece rate b . The piece rate b' needed to generate effort $e_{i|t+1}^i$ in period t is given by

$$b' = \frac{D^i(T-t-1)}{D^i(T-t)} b. \quad (5)$$

Proof. Define b' as the piece rate that the worker must be paid in period t to elicit effort equivalent to $\tilde{e} \equiv e_{i|t+1}^i$, which is the effort level exerted by the worker in period $t+1$. The FOC for worker i at period $t+1$ under the original piece rate is $D^i(T-t-1)b = c'(\tilde{e})$. The FOC for worker i at period t under the alternative piece rate b' is $D^i(T-t)b' = c'(\tilde{e})$. This implies

$$\frac{D^i(T-t-1)}{D^i(T-t)} = \frac{b'}{b}.$$

QED

Intuitively, $D^i(T-t-1)/D^i(T-t)$ is how much more the period T wage payment is valued by the worker in period $t+1$ relative to period t . This is exactly how much more the worker would need to be paid for her

¹⁰ The regularity conditions and the properties of $c'(\cdot)$ guarantee single-peakedness of the maximand. Hence there will be a unique maximum either at the interior or at the corner where $e = 0$. It is straightforward to verify that the maximum will not be zero if the derivative of the maximand with respect to e is positive at $e = 0$. That will be the case if $c'(0) < D^i(T)b$ for all i . Workers for whom this condition is not satisfied would not participate in the program described below; hence we assume that this condition is satisfied for all workers. Under this condition, the FOCs will be both necessary and sufficient for a global maximum.

period t self to decide to exert effort level $e^i_{t+1|t+1}$ in period t . Because they both change the perceived returns to effort in a period, there is an equivalence between how the discount factor and piece rate affect output.

PROPOSITION 1 (Pay cycle effect). The proportional increase in output from period t to $t + 1$ is given by

$$\frac{e^i_{t+1|t+1} - e^i_{t|t}}{e^i_{t|t}} = \varepsilon \left[\frac{D^i(T - t - 1) - D^i(T - t)}{D^i(T - t)} \right], \tag{6}$$

where ε is the elasticity of output with respect to the piece rate. The term in brackets reduces to $(1/\delta) - 1$ if workers are exponential discounters and is greater than $(1/\delta) - 1$ if they are time inconsistent. This implies that output increases over the pay cycle will be larger for time-inconsistent workers than for time-consistent ones.

Proof. Since $c(e) = e^\theta$, $\varepsilon = 1/(\theta - 1)$ and is constant over effort. Thus,

$$\varepsilon = \frac{1}{\theta - 1} = \left(\frac{e^i_{t+1|t+1} - e^i_{t|t}}{e^i_{t|t}} \right) / \left(\frac{b' - b}{b} \right).$$

Substituting in for b'/b from lemma 1 gives the relationship between the output increase over the pay cycle and the change in the discount factor:

$$\frac{D^i(T - t - 1) - D^i(T - t)}{D^i(T - t)} = \frac{1}{\varepsilon} \left(\frac{e^i_{t+1|t+1} - e^i_{t|t}}{e^i_{t|t}} \right).$$

Since

$$\frac{D^i(T - t - 1)}{D^i(T - t)} > \frac{D^C(T - t - 1)}{D^C(T - t)} = \frac{1}{\delta},$$

expression (6) implies that the output increase from period t to $t + 1$ will be larger for time-inconsistent workers.

Note that for a general $c(\cdot)$ function, the elasticity of output to the piece rate will change from $e = e^i_{t|t}$ to $e = e^i_{t+1|t+1}$. In this more general case,

$$\frac{\partial e^i_{t|t}}{\partial b} \frac{b}{e^i_{t|t}} \approx \left(\frac{e^i_{t+1|t+1} - e^i_{t|t}}{e^i_{t|t}} \right) / \left(\frac{b' - b}{b} \right),$$

and expression (6) will hold as a first-order approximation. To compare the magnitude of the pay cycle increase across the two types of workers, consider a time-consistent and time-inconsistent worker (where we allow the two workers to have different values of α^i), both of whom exert the same effort level in period t . Then, both workers will have the same elas-

tivity in period t , implying that the output increase from period t to $t + 1$ will be larger for time-inconsistent workers. QED

In the experiment, workers are assigned to a weekly payday. We will show in Section IV that for an exponential discounter—for any reasonable value of δ —output should not noticeably change over a weekly pay cycle. Proposition 1 will enable us to calibrate the level of discounting over the pay week, using an experimentally obtained elasticity measure to estimate θ .

Contract choice.—For time-consistent workers, since

$$\frac{D^c(T - t + k)}{D^c(k)} = \frac{D^c(T - t)}{D^c(0)} = \delta^{T-t},$$

the FOCs for optimal effort (3) and (4) are exactly the same. Hence the effort level chosen by period t is also optimal from the perspective of period $t - k$: $e_{t|t-k}^i = e_{t|t}^i$ for time-consistent workers. The reason is that utility from the standpoint of the period t self is simply a multiple of utility from the perspective of the period $t - k$ self. Both selves weigh the benefits of income at the payday relative to the costs of effort at time t exactly the same.

In contrast, from the perspective of a time-inconsistent worker in period $t - k$, the period t self will supply too little effort. Specifically, since $D^i(T - t + k)/D^i(k) > D^i(T - t)/D^i(0)$, $e_{t|t-k}^i > e_{t|t}^i$ for time-inconsistent workers. Because the period $t - k$ self weighs the benefits of effort relative to the costs more heavily than the period t self, the period $t - k$ self desires more effort than the period t self. This is the essence of the time inconsistency problem.

In proposition 2, we show that this will lead time-inconsistent workers to demand dominated contracts, which punish workers by paying less than $w(\cdot)$ if effort is below a threshold and pay the same as $w(\cdot)$ for effort above the threshold. See figure 1 for an example of such a dominated contract. Note that we do not make any claims about optimal contracting in this setting; rather, the dominated contract is a convenient device that enables us to test for time inconsistency.

PROPOSITION 2 (Demand for dominated contracts, bounds on time inconsistency).

- a. Suppose that in period $t - k$ workers are offered the following dominated wage schedule, which allows them to choose a target output level, $X_{t|t-k}$, for period t :

$$v_{t|t-k}(e) = \begin{cases} a + b^{\theta}e & e < X_{t|t-k} \\ a + be & e \geq X_{t|t-k} \end{cases}$$




FIG. 1.—Incentive contracts. This figure displays the two types of incentive contracts offered to workers. The linear control contract paid a piece rate wage of b for each accurate field entered. The nonlinear dominated contract imposed a production target, X ; workers were paid b for each accurate field if they met the target but received only $b/2$ for each field if they fell short of the target.

where $0 \leq b^p < b$. Time-inconsistent workers' period $t - k$ selves will strictly prefer $v_{t|t-k}(\cdot)$ over $w(\cdot)$ and will choose $X_{t|t-k} \in (e_{t|t}^i, e_{t|t-k}^i]$. In contrast, time-consistent workers will never strictly prefer $v_{t|t-k}(\cdot)$.

b. Define $x_{t|t-k}^i$ as the proportional increase in output under $v_{t|t-k}(\cdot)$ relative to $w(\cdot)$ at time t . Then $x_{t|t-k}^i$ is bounded above by the elasticity times the level of time inconsistency between periods $t - k$ and t :

$$x_{t|t-k}^i \leq \varepsilon \left[\frac{D^i(T - t + k)/D^i(k)}{D^i(T - t)/D^i(0)} - 1 \right]. \tag{7}$$

Proof. The period $t - k$ self chooses $X_{t|t-k}$ to maximize its utility, subject to the constraint that the period t self will choose the level of effort that maximizes its utility given the dominated wage schedule with target $X_{t|t-k}$.

For time-consistent workers, $e_{t|t-k}^i = e_{t|t}^i$, and so $v_{t|t-k}(\cdot)$ has no benefits and will never be strictly preferred. For such workers,

$$\frac{D^c(T - t + k)/D^c(k)}{D^c(T - t)/D^c(0)} = 1,$$

and thus $x_{t|t-k}^i = 0$.

In contrast, for time-inconsistent workers, $e_{t|t-k}^i > e_{t|t}^i$, so the period $t - k$ self will potentially want to influence the period t self's choice of effort. We restrict attention to targets in the range $X_{t|t-k} > e_{t|t}^i$. The reason is that for any target $X_{t|t-k} \leq e_{t|t}^i$, the period t self would choose effort $e_{t|t}^i$, and so from the perspective of period $t - k$, $v_{t|t-k}(\cdot)$ would elicit the same effort as $w(\cdot)$.

To characterize the optimal $X_{t|t-k}$, we solve backward, starting with the period t self's problem given $X_{t|t-k}$. Let $e_t^*(X_{t|t-k})$ denote optimal effort from the perspective of period t given contract $v_{t|t-k}(\cdot)$ with target $X_{t|t-k}$. Let $e_t^{i(p)}$ denote optimal effort from the perspective of period t under the penalty piece rate b^p ; $e_t^{i(p)}$ is implicitly defined by the FOC

$$c'(e_t^{i(p)}) = D^i(T - t)b^p.$$

Note that $e_t^{i(p)} < e_{t|t}^i$ since $b^p < b$.

Since $X_{t|t-k} > e_{t|t}^i$, by the convexity of $c(\cdot)$, the period t self will never choose effort greater than $X_{t|t-k}$. Thus, in the range of effort $[X_{t|t-k}, \infty)$, period t will choose $X_{t|t-k}$ and earn piece rate b ; in the range of effort $(0, X_{t|t-k})$, period t will choose $e_t^{i(p)}$ and earn piece rate b^p . Then, $e_t^*(X_{t|t-k})$ is given by whichever of these provides higher utility to period t :

$$e_t^*(X_{t|t-k}) = \begin{cases} X_{t|t-k} & D^i(T - t)(bX_{t|t-k}) - c(X_{t|t-k}) \\ & \geq D^i(T - t)(b^p e_t^{i(p)}) - c(e_t^{i(p)}) \\ e_t^{i(p)} & \text{otherwise.} \end{cases}$$

There is always some $X_{t|t-k} > e_{t|t}^i$ such that $e_t^*(X_{t|t-k}) = X_{t|t-k}$. To see this, write $X_{t|t-k} = e_{t|t}^i + \mu$. As $\mu \rightarrow 0$, the time t self will be arbitrarily close to indifferent between exerting $X_{t|t-k}$ under $v_{t|t-k}(\cdot)$ and exerting $e_{t|t}^i$ under $w(\cdot)$. Hence for any $b^p < b$, there will be some μ small enough such that the worker's period t self will prefer $X_{t|t-k}$ over $e_t^{i(p)}$.

We now characterize the period $t - k$ self's optimal choice of $X_{t|t-k}$. Note that by FOC (3), period $t - k$ will prefer to induce $e_{t|t-k}^i$ over any effort level higher than this. In addition, since, by the convexity of $c(\cdot)$, the period t self's utility is decreasing in e for $e > e_{t|t-k}^i$, the period t self is weakly more likely to choose $e_t^*(X_{t|t-k}) = X_{t|t-k}$ for $X_{t|t-k} = e_{t|t-k}^i$ than for some target greater than $e_{t|t-k}^i$. Therefore, $X_{t|t-k} \in (e_{t|t}^i, e_{t|t-k}^i]$. In addition, since the period $t - k$ self's FOC is satisfied with equality at $e_{t|t-k}^i$, $bD^i(T - t + k)/D^i(k) - c'(e_{t|t-k}^i)$ will be positive and decreasing in e for $e < e_{t|t-k}^i$. Thus, the period $t - k$ self will prefer to induce as high an effort level as possible in the range $e \leq e_{t|t-k}^i$.

The maximum level of effort, $e_t^{i(\max)}$, that the period t self can be induced to supply under the dominated contract is implicitly given by

$$D^i(T - t)[be_t^{i(\max)} - b^p e_t^{i(p)}] - [c(e_t^{i(\max)}) - c(e_t^{i(p)})] = 0.$$

Then $e_{t|t-k}^i$ can be induced in period t if $e_{t|t-k}^i \leq e_t^{i(\max)}$. Thus, we have

$$X_{t|t-k} = \min\{e_{t|t-k}^i, e_t^{i(\max)}\}.$$

Finally, following the logic in the proposition 1 proof, we define b'' as the alternative piece rate wage that would induce period t to choose $X_{t|t-k}$ under the wage schedule $w(e)'' = a + b''e$. Thus, following the FOC in equation (4), b'' is defined as $D^i(T - t)b'' - c'(X_{t|t-k}) = 0$. The FOC in equation (3) implies $D^i(T - t + k)b - D^i(k)c'(X_{t|t-k}) \geq 0$ since $X_{t|t-k} \leq e_{t|t-k}^i$. Together, these conditions imply

$$\frac{b''}{b} \leq \frac{D^i(T - t + k)}{D^i(T - t)D^i(k)}.$$

Plugging this into the elasticity formula

$$\varepsilon = \frac{1}{\theta - 1} = \left(\frac{e_{t|t-k}^i - e_{t|t}^i}{e_{t|t}^i} \right) / \left(\frac{b'' - b}{b} \right)$$

gives expression (7):

$$\frac{D^i(T - t + k)/D^i(k)}{D^i(T - t)} - 1 = \frac{D^i(T - t + k)/D^i(k)}{D^i(T - t)/D^i(0)} - 1 \geq \frac{x_{t|t-k}^i}{\varepsilon}.$$

For a more general $c(\cdot)$,

$$\frac{\partial e_{t|t}^i}{\partial b} \frac{b}{e_{t|t}^i} \approx \left(\frac{e_{t|t-k}^i - e_{t|t}^i}{e_{t|t}^i} \right) / \left(\frac{b'' - b}{b} \right),$$

and expression (7) will hold as a first-order approximation. QED

In the experiment, workers can choose between a linear piece rate contract and a dominated contract in which they self-impose a target. For time-consistent workers, $e_{t|t-k}^i = e_{t|t}^i$, and so the dominated contract $v_{t|t-k}(\cdot)$ has no benefits and will never be strictly preferred. In contrast, while time-inconsistent workers' period t selves will prefer $e_{t|t}^i$, their period $t - k$ selves will prefer $e_{t|t-k}^i$ and will want to induce as high an effort level as possible in the range $e \leq e_{t|t-k}^i$ (see the proof). Consequently, time-inconsistent workers will strictly prefer $v_{t|t-k}(\cdot)$ to $w(\cdot)$ because they can use the dominated contract to induce their future self to work harder.

The ratio on the right-hand side of expression (7),

$$\frac{D^i(T-t+k)/D^i(k)}{D^i(T-t)/D^i(0)},$$

reflects the convexity of $D^i(\cdot)$. The numerator captures how the period $t-k$ self discounts the costs of effort at time t relative to the benefits of pay at time T ; the denominator captures this ratio for the period t self. If the period $t-k$ and t selves valued the benefits relative to the costs exactly the same, then this ratio would be 1. Indeed, for exponential discounters, $(\delta^{T-t+k}/\delta^k)/(\delta^{T-t}/\delta^0) = 1$. This underscores that time-consistent workers' output would not increase if these workers were offered dominated contracts. If workers had quasi-hyperbolic preferences, this ratio would equal $1/\beta$, which captures the level of time inconsistency between current and future periods. Note that the observed output increase, $x_{|t-k}^i$, is a lower bound on the deviation from time consistency because the period $t-k$ self may not be able to induce its preferred effort level of $e_{|t-k}^i$: $X_{|t-k}^i \leq e_{|t-k}^i$.

PROPOSITION 3 (Correlation of pay cycle and dominated contract effects). Suppose that some workers are time-consistent exponential discounters with discount function $D^C(\cdot)$ and the others are time inconsistent with discount function $D^I(\cdot)$. Then, the magnitude of the output increase over the pay cycle, $e_{t+1|t+1}^i/e_{|t}^i$, will be positively correlated with demand for dominated contracts and the extent to which their provision increases output.

Proof. This follows from propositions 1 and 2. Time-inconsistent workers will exhibit larger output increases over the pay cycle. Moreover, only time-inconsistent workers can be expected to demand dominated contracts and increase output in response to being offered the choice of such contracts. QED

If the population includes exponential discounters with different discount rates and hyperbolic discounters with different $D(\cdot)$ functions, then the correlation will not be one. The pay cycle increase reflects impatience: how much wages are discounted when the payment period is further away. In contrast, the output increase under dominated contracts reflects time inconsistency: the extent to which the period $t-k$ and t selves differ in weighing costs versus benefits of effort in period t . In practice, this correlation will also be weakened if some time-inconsistent workers are naive; in this case, all time-inconsistent workers would exhibit pay cycle effects but only sophisticates would choose (and be affected by) dominated contracts.

When workers are time inconsistent, the temporal distance between the moment of contract choice, moment of effort, and moment of compensation is what generates scope for the dominated contract to affect effort

provision. In propositions 4 and 5 below, we describe how changes in the extent of these distances will lead to changes in dominated contract effects.

PROPOSITION 4 (Decrease in dominated contract effects on output over the pay cycle). For time-inconsistent workers, if $e_{t|t-k}^i$ can be induced in period t using the dominated contract, then $x_{t|t-k}^i > x_{t+1|t+1-k}^i$. In contrast, for time-consistent workers, $x_{t|t-k}^i = x_{t+1|t+1-k}^i = 0$.

Proof. For time-inconsistent workers,

$$\begin{aligned} x_{t|t-k}^i &= \varepsilon \left[\frac{D^i(T-t+k)/D^i(k)}{D^i(T-t)/D^i(0)} - 1 \right] \\ &> \varepsilon \left[\frac{D^i(T-t-1+k)/D^i(k)}{D^i(T-t-1)/D^i(0)} - 1 \right] \\ &\geq x_{t+1|t+1-k}^i, \end{aligned}$$

where the equality comes from the assumption that $e_{t|t-k}^i$ can be induced in period t (see the proposition 2 proof); the strict inequality is due to the hyperbolicity assumption that $D^i(t+s)/D^i(t)$ is increasing in t ; and the final weak inequality comes from the definition of $x_{t+1|t+1-k}^i$ using proposition 2.

In contrast, for time-consistent workers, $x_{s|s-k}^i = 0$ in all periods, so trivially, there will be no change in dominated contract effects over the pay cycle. QED

Proposition 4 holds when the period t self can be induced to exert $e_{t|t-k}^i$. In lemma A1 in online appendix A, we show that this will always be the case for θ sufficiently large. Proposition 4 states that the impact of dominated contracts on output will be larger further away from the pay period. The reason is that $e_{t|t-k}^i$ becomes closer to $e_{t|t}^i$ as T approaches. As a result, the level of the time inconsistency problem gets smaller closer to the pay period, and there is therefore less scope for the dominated contract to increase effort.

PROPOSITION 5 (Horizon of choice). When selecting a dominated contract for period t , time-inconsistent workers will choose to induce a weakly higher effort level when contract choice is made further in advance of period t .

Proof. Consider two possible values of k , k_1 and k_2 , where $k_1 < k_2$. Using the definition of hyperbolicity,

$$D^i(T-t)/D^i(0) < D^i(T-t+k_1)/D^i(k_1) < D^i(T-t+k_2)/D^i(k_2).$$

Using the first-order conditions, this implies $e_{t|t}^i < e_{t|t-k_1}^i < e_{t|t-k_2}^i$. Consequently, following the logic in proposition 2, the period $t-k_2$ self will prefer to induce an effort level greater than the period $t-k_1$ self. However, if an effort level greater than $X_{t|t-k_1}$ is not inducible, then the period

$t - k_2$ self will choose $X_{t-k_2} = X_{t-k_1}$ and induce the same effort level as the period $t - k_1$ self. QED

In the experiment, we vary whether contract choice occurs in the evening before the workday or the morning of the workday. As discussed above, one way to map this to the model is to assume that there are three subperiods within each 24-hour period: (i) the evening before the workday, (ii) the morning of the workday, and (iii) the workday, during which time effort is exerted. Then, we can think of period i as $t - k_2$, period ii as $t - k_1$, and period iii as t .

The model is deterministic and abstracts from the possibility of shocks to output or to the cost of effort. Suppose instead that there were some probability $p > 0$ that instead of output equaling e , output equaled $\min\{0, e + j\}$, where j is a mean zero normal error term with variance σ_j^2 . Then the dominated contract derived above would be less attractive for both time-consistent and time-inconsistent workers, because there would be some states in which workers would face the penalty even if $e = X_{t-k}$. Thus time-consistent workers would strictly prefer $w(\cdot)$ to $v_{t-k}(\cdot)$. In contrast, time-inconsistent workers could still prefer a dominated wage schedule: for small enough p , they would prefer $v_{t-k}(\cdot)$ to $w(\cdot)$ by continuity. Since they would incur penalties with positive probability, they might choose less aggressive target effort levels. Thus, in the stochastic output case, $1 + \alpha(1 - \theta)$ is a less tight lower bound on the extent to which the worker deviates from exponential discounting.

Shocks to the cost of effort, for example, from illness or a family emergency, would make dominated contracts less attractive not only because workers might miss the target, losing $b - b^p$ per unit of output, but also because even under smaller shocks, exerting the effort to reach the target might yield little surplus to the workers. These factors could lead workers in time 0 to reject even dominated contracts with very small X_{t-k} .

In addition, as discussed further below, the risk of shocks to output or the cost of effort might vary over time. Under such time-varying stochastic shocks, time-inconsistent workers might select dominated contracts some times but not other times. Finally, workers will have more information about the shocks as they get closer to the moment of work; this could dampen the prediction that because of time inconsistency, dominated contracts will be more appealing further in advance of the time of effort. The fact that uncertainty may vary over the pay cycle or get resolved closer to the effort period may weaken the predictions in propositions 4 and 5.

In this section we have considered worker choice of effort and contracts in response to exogenously determined menus of contracts.¹¹ In

¹¹ Another issue is that we take the period between payments as exogenous. Endogenizing the length of time between pay periods is an interesting issue in its own right. Even if more frequent payment mitigates the work self-control problem, it may carry transaction costs and

the next sections we first use the predictions to test for time inconsistency in worker effort. We then calibrate the extent to which workers depart from standard exponential discounting using the propositions above.

III. Experiment Design

A. *Experimental Context*

To assess the empirical relevance and magnitude of time inconsistency, we worked with an Indian data entry firm in Mysore—a region that is a major data entry hub. Using the firm's infrastructure—office space, entry software, and operational protocols—we designed and managed a field experiment over 13 months.

Workers used data entry software to type information from scanned images into fields on their computer screen (see app. fig. 1). The software provided them with information on their own output with about a 15-minute delay. Following standard practice in the data entry industry, workers were paid piece rates based on the number of accurate fields entered. Accuracy was measured using dual entry of data, with manual checks of discrepancy by separate quality control staff. Workers were paid a piece rate of Rs. 0.03 (rupees) for each accurate field entered (see below) plus a small flat daily show-up fee of Rs. 15 that constituted about 8 percent of their compensation. They earned zero on days they were absent. Thus in the language of the model in Section II, $w(e) = \text{Rs. } 15 + \text{Rs. } 0.03e$. Pay levels were at par with or slightly higher than those paid by other data entry firms in the region.

Employees were recruited through the standard procedures used by the firm: the pool of resumes submitted by walk-ins and solicitations via posters and announcements in surrounding villages. Applicants were required to have completed tenth grade and be at least 18 years old. Workers were told they were being hired for a one-time contract and were not provided reference letters upon completion of the job.

Roughly three-quarters of workers were male. Among those who reported age on their resumes, average age was 24 years (panel A, table 1). Workers averaged 13 years of education; most had taken a computer course and had an e-mail address prior to joining the firm.¹² Many employees commuted from surrounding villages using buses and trains, with some traveling up to 2 hours in each direction.

exacerbate the consumption self-control problem, since infrequent payment may be an implicit savings commitment device. Indeed, there were several instances over the course of the project in which workers asked management to withhold their earnings for weeks at a time because this would help them save for lump-sum expenditures.

¹² In this and other information presented in table 1, some employees hired in later stages of the project were not surveyed because of clerical oversight.

TABLE 1
SUMMARY STATISTICS

	Mean (1)	Standard Deviation (2)	Observations (3)
A. Worker Characteristics			
Proportion female	.26	.44	111
Age	24	4	63
Years of education	13	2	101
Completed high school	.84	.37	101
Used computer prior to joining firm	.67	.47	101
Had e-mail address prior to joining firm	.60	.49	101
B. Performance on Tests Administered during Training			
Contracts comprehension quiz: percentage score	93	13	79
IQ composite score (Raven's Matrix plus Digit Span)	62	15	106
C. End Line Survey: Discount Rate Measurement			
Proportion of times worker chose smaller immediate reward	.31	.28	58
Proportion of times worker displayed preference reversal	.17	.23	58
D. End Line Survey: Self-Reported Measures of Self-Control Problems			
Worker agreed or agreed strongly with the statement:			
"Some days I don't work as hard as I would like to"	.76	.43	70
"I get tempted to leave work earlier than I would like"	.40	.49	70
"I wish I had better attendance at work"	.86	.35	70
"It would be good if there were rules against being absent because it would help me come to work more often"	.73	.45	70
Self-control index ^a	3.43	.55	70
Worker has tried to quit an addictive behavior and failed (males only)	.12	.33	51
Factor analysis: self-control factor	.00	.86	70

NOTE.—This table presents summary statistics for the 111 workers who participated in the full study (contract and payday treatments). In the discount rate exercise (panel C), workers traded off three sets of cash awards (Rs. 20 vs. Rs. 24, Rs. 50 vs. Rs. 57, and Rs. 100 vs. Rs. 110) under two different horizons: short horizon (the smaller amount today vs. the larger amount in 3 days) and long horizon (the smaller amount in 14 days vs. the larger amount in 17 days). Panel C reports statistics on the proportion of times the worker chose the smaller immediate reward out of the six questions and the number of times the worker showed preference reversal (chose the smaller immediate reward in the short horizon and chose the larger reward in the long horizon). Panel D summarizes responses to questions that asked workers to agree or disagree with statements about self-control behavior. The self-control factor (panel D) was determined using a factor analysis on the full set of end line survey questions.

^a Mean of responses to all nine self-control questions (1 = disagree strongly, 5 = agree strongly).

New recruits received about 2 weeks of training before contract randomizations began (see below). During the first 4 days, they were paid a flat stipend while receiving instruction on the data entry software and production task. During the next 4 days, they worked under assignment to the control contract with wage schedule $w(e)$. They also received training on the contract treatments (described below) during this time. After this, they were assigned to the dominated contract for 2 days under the low and medium targets, respectively. This gave them the opportunity to observe their production under both types of incentive schemes before beginning contract randomizations. The mean score on a quiz that workers took to verify they understood the contracts was 93 percent. Throughout the experiment, workers were randomly assigned to seats in the office, and these assignments changed every 1–3 weeks, since some computers were slower and more sensitive to network speed fluctuations than others.

B. Treatments

To test proposition 1, employees were randomized into three payday groups, which were paid in the evenings of Tuesday, Thursday, and Saturday, respectively, for work completed over the previous 7 days. For example, the Thursday payday group workers were paid when leaving the office on Thursdays throughout the experiment; payment was for work completed from the previous Friday to that Thursday. Workers were instructed to stop working at least 20 minutes early on paydays to allow sufficient time for their output to be computed and earnings disbursed before they left to catch their bus or train. The payday randomizations allow us to control for other day-of-the-week factors that might affect effort, such as a post- or preweekend effect.

To test proposition 2, we used two contracts. The linear “control” contract paid a piece rate wage of b (Rs. 0.03) for each field entered accurately. The nonlinear “dominated” contract paid a piece rate of b if workers met a target, but only $b/2$ for each entered field if they fell short of the target. As shown in figure 1, for any given production level, earnings are always weakly higher under the control contract than under the dominated contract.

Each day, each worker was independently randomized into one of four contract treatments. In the first, workers were assigned to the control contract. In the second, they were assigned to the dominated contract, with an exogenous target imposed; the target was selected from three target levels: level 1, level 2, and level 3.¹³ Of course, imposing targets

¹³ The three target levels were set at 3,000, 4,000, and 5,000 accurate fields, respectively. In the first month of randomizations, these corresponded to the 30th, 50th, and 70th per-

could increase output regardless of whether a worker is time inconsistent. To test our model, we rely on the remaining treatments, which gave workers the *option to choose a dominated contract*, in which they chose their own target. They could always choose a target of zero (and many did), which is the equivalent of choosing the simple linear control contract. In the third treatment, *morning option to choose a dominated contract*, workers chose their targets in the morning when they arrived to work.¹⁴ Finally, in the fourth treatment, *evening option to choose a dominated contract*, workers chose their targets the evening before the workday.

To make workers' information similar across these conditions, all workers were told their treatment assignment for each day the evening before. Every worker received each of the four contract treatments in random order exactly 25 percent of the time over every 8-day or 12-day work period. As an example, appendix table 1 displays the contract assignments for five workers in the sample over a 24-day period.

Figure 2 provides an overview of the experiment time line and design. Our sample of 102 workers and 8,423 observations covers the 8-month period when both contract and payday treatments occurred simultaneously.¹⁵ Appendix table 2 verifies that treatment assignments were balanced across the sample.

IV. Results

A. Pay Cycle Effects on Production (Test 1)

Workers produce 215 fields more, on average, on paydays than on non-paydays on a base of roughly 5,300 fields (table 2, col. 1). Effects persist controlling for serial correlation in output (col. 2).

To examine dynamics over the weekly pay cycle more fully, we estimate a model with a full set of indicators for each day in the pay week (with 6 or more days from the next payday as the omitted category). The coefficients from this regression are displayed in table 2, column 3. Employees are least productive on the days furthest from their next payday. Production then rises through the pay cycle.¹⁶ Earnings follow a similar

centiles, respectively, of production under the control contract. Initially, the target assignment was only to level 1 and level 2 targets. Assignment to the level 3 target was added later, as production levels increased. During the last month of contract randomizations, we changed these levels to 4,000, 5,000, and 6,000 accurate fields to correspond to increases in worker production over time.

¹⁴ Note that if the length of time periods is a day, then we would expect this last treatment to have no effect.

¹⁵ The payday treatments were run for 3 additional months (during end line activities). All payday effects reported below are similar in a sample covering this longer period.

¹⁶ We lack the power to pin down the exact shape of the increase in output over the pay week; one could fit a convex, linear, or concave curve through the confidence intervals in fig. 2.




FIG. 2.—Experiment time line

pattern, as shown in column 4 of table 2 and plotted in figure 3. The average change in output and earnings from the beginning to the end of the pay week is 8 percent. This magnitude corresponds to approximately 1 additional year of education in our sample.¹⁷

Note that output and earnings dip slightly (and insignificantly) from the day before the payday to the payday itself, likely because workers were required to stop work at least 20 minutes early to collect their pay. Specifically, average piece rate earnings per hour worked are Rs. 27 on both the payday and the day before payday. If we assume that workers would have worked 20 minutes longer on paydays, this implies they would have earned about Rs. 9 more on paydays on average, which would more than compensate for the observed payday dip of Rs. 3.¹⁸

Attendance also increases steadily over the pay cycle (col. 5), consistent with increased effort closer to paydays.¹⁹ In general, the payday cycle affects both the extensive margin—attendance and workday length—and the intensive margin (app. table 3, panel A).

The pay cycle dynamics suggest that quasi-hyperbolic models (Laibson 1997) do not fit our data well. These models would predict that the effects arise only on the payday itself. Instead, we see a steady increase. These dynamics also rule out an explanation based solely on workers showing up to work on the payday to collect their checks or an explanation focused on workers taking the day after payday off.

¹⁷ We also estimated a specification with a linear control for number of days before the next payday, a payday dummy, and standard controls. The coefficient on the linear control indicates that earnings increase, on average, by Rs. 3—or 2 percent—per day leading up to the payday (significant at the 1 percent level).

¹⁸ Moreover, workers may have wanted to leave work early on paydays to make purchases, e.g., if they were credit constrained, had time inconsistency in consumption, or worried about demands from relatives if they hung on to cash.

¹⁹ Results are similar if a probit estimator is used instead of a linear probability model.

TABLE 2
PAY CYCLE TREATMENT EFFECTS

	DEPENDENT VARIABLE					
	Production (1)	Production (2)	Production (3)	Earnings (4)	Attendance (5)	Production (6)
Payday	215 (70)***	140 (63)**	428 (94)***	14.09 (2.99)***	.077 (.013)***	
1 day before payday			539 (95)***	17.19 (3.02)***	.053 (.013)***	
2 days before payday			417 (113)***	13.54 (3.60)***	.037 (.016)**	
3 days before payday			374 (112)***	11.82 (3.57)***	.026 (.017)	
4 days before payday			332 (123)***	10.15 (3.91)***	.047 (.017)***	
5 days before payday			176 (119)	5.91 (3.79)	.023 (.017)	1,071 (239)***
Piece rate increase						Yes 550
Lag dependent variable controls	No	Yes	Yes	Yes	No	Yes
Observations	8,423	8,423	8,423	8,423	8,423	8,423
R ²	.4961	.5889	.5909	.57	.11	.76
Dependent variable mean	5,337	5,337	5,337	172	.88	9,361

NOTE.—The sample in cols. 1–5 is the experiment sample. The dependent variables equal zero if a worker was absent. Payday is a binary indicator for whether that day was the worker's assigned payday. In cols. 3–5, "X days before payday" are binary indicators for whether the current day is X calendar days away from the worker's assigned payday; the omitted category in these columns is 6 or more days away from the payday. In col. 6, the sample is observations after the end of the experiment in which workers' wages were randomized. Piece rate increase is a binary indicator that equals one if the worker's piece rate was Rs. 0.04 per accurate field that day and equals zero if the worker's piece rate was Rs. 0.03 per accurate field. All regressions include fixed effects for each date in the sample, each worker in the sample, and each computer seating assignment. Regressions 2, 3, and 6 also include controls for lagged production (production on the previous workday and 2 workdays ago); similarly, regression 4 includes controls for lagged earnings. Robust standard errors are reported in parentheses.

* Significant at the 10 percent level.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

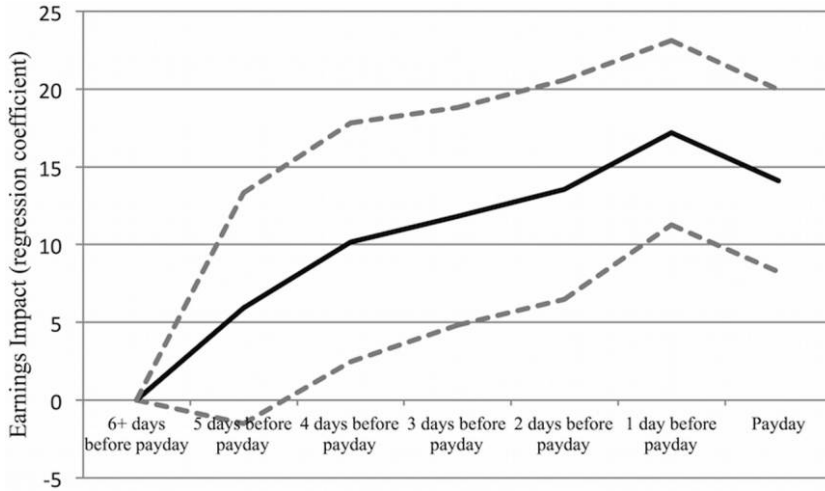


FIG. 3.—Earnings over the pay cycle. This figure graphs the coefficients and 95 percent confidence intervals from a regression of earnings on six binary indicators that capture distance from a worker’s next payday (payday, 1 day before payday, 2 days before payday, etc.). The omitted category is 6 or more days before payday. Note that these coefficients correspond to those shown in column 4 of table 2.

Further evidence consistent with present bias is provided by festivals, which involve large, perfectly foreseeable expenditures. Under convex effort costs, time-consistent workers should smooth production; in contrast, time-inconsistent workers’ output would spike before festivals. Indeed, production increases by 8 percent in the week prior to major festivals (app. table 4).

Calibration of the implied discount rate.—We can use proposition 1 to calibrate the discount rate implied by the weekly production cyclicity. In order to estimate the elasticity of output to the piece rate, after contract randomizations were finished, workers were randomly offered one of two piece rate wages: Rs. 0.03 (their usual piece rate) and Rs. 0.04 per accurate field. Each worker received each piece rate five times over a 10-day period in random order. This 33 percent increase in wages increased output by 11 percent, for an elasticity of 0.33 (table 2, col. 6). Note that under the assumption that $c(e) = e^\theta$, with elasticity equal to $1/(\theta - 1)$, this implies $\theta = 4$.

Since the average output increase over the 6 days from the beginning to the end of the payweek is 8 percent (table 2, col. 3), on average,

$$\frac{e_{t+1|t+1}^i - e_{t|t}^i}{e_{t|t}^i} = \frac{0.08}{6} = 0.013.$$

Proposition 1 allows us to back out the implied change in the discount factor:

$$\frac{d^i(T-t-1) - d^i(T-t)}{d^i(T-t)} = \frac{0.013}{0.33} = 0.04.$$

Thus on average, the daily increase in discounting is 4 percent. If workers were time-consistent exponential discounters, this would require an annual discount rate of 1.65×10^6 percent. Standard estimates for the exponential discount rate in the literature are about 5 percent per year (e.g., Engen, Gale, and Scholz 1994; Hubbard, Skinner, and Zeldes 1994; Engen, Gale, and Uccello 1999)—far lower than those we estimate. Of course, the exact discount rate implied by the calibration above should be taken with a grain of salt since the model does not perfectly correspond to reality. As discussed above, output and effort costs may be stochastic. Some workers might be able to smooth intertemporally using savings or credit: we would expect such workers to show more modest pay cycle increases, deepening the puzzle of our finding such large effects. While we model utility as quasi-linear, there may indeed be some income effects; however, these would generate substitutability between effort in different periods—behavior that we empirically rule out below in Section IV.B. In addition, the effort elasticity may not be exactly 0.33 and may not be constant everywhere. While all these factors suggest caution regarding the precision of the calibrated discount rate of 4 percent per day, it seems hard to imagine that one could fail to reject the hypothesis of exponential discounting at rates of about 5 percent per year.

While standard estimates of implied daily discount rates under exponential discounting do not match our data, estimates that allow for hyperbolic discounting are much more consistent. For example, fitting laboratory data to a hyperbolic model, Kirby and Marakovic (1995) estimate discounting of 1–3 percent per day over short horizons. Calibrations from field data also produce such large estimates (e.g., Shui and Ausubel 2005; Paserman 2008; Fang and Silverman 2009).

Pay cycle effects are not unique to our setting. For example, in large US firms, salespeople increase the frequency of their sales over the fiscal year, with a spike in the last quarter when bonuses are computed and paid; a trend that remarkably resembles our figure 2 (Oyer 1998, fig. 1).²⁰ Factory workers in pre-Industrial Revolution England exhibited similar dynamics over their pay cycle (Clark 1994).²¹

²⁰ We thank an anonymous referee for pointing us to this paper.

²¹ Such “pay cycle” behavior has been documented in other domains as well. For example, Shapiro (2005) documents time inconsistency in consumption by showing that the caloric intake of food stamp recipients in the United States declines by 10–15 percent over the food stamp month.

While our simple model predicts that more frequent pay will increase output, this may not be true more generally and could conflict with other objectives. Time inconsistency in consumption means that workers may value more infrequent payments to help save for lumpy expenditures (e.g., Ashraf et al. 2006). Consistent with this, some workers in our experiment asked us to withhold their earnings to help them save. Similarly, the nine members of the experiment's managerial staff—who were paid fixed salaries—chose to receive their earnings monthly rather than weekly. Of course, more frequent pay may also be undesirable because of transaction costs or long output horizons, like lengthy sales cycles.

B. Demand for and Treatment Effects of Dominated Contracts (Test 2)

On average, when given the option to choose a dominated contract, workers take up the dominated contract by selecting a positive target 36 percent of the time when present (table 3). This is based on the sample of workers who were present both the day before and the day of the treatment assignment (and thus were informed of their treatment the evening before as per protocol and were able to select targets). As a conservative estimate, if we code workers who are absent the day before or the day of assignment as choosing zero targets, the mean take-up rate across workers is 28 percent.²²

Figure 4 plots the distribution of worker take-up rates; 16 percent of workers always chose a target of zero. The bottom quarter of the distribution chose positive targets less than 10 percent of the time. The top quarter chose positive targets at least 60 percent of the time. As discussed in Section II, in a deterministic model in which workers had a fixed type, hyperbolic workers would always choose a positive target and exponential discounters would never choose a positive target. However, time-varying stochasticity in output or effort costs could create within-worker variation in choice of contract. Network slowdowns, assignment to a slow computer, changes in difficulty of batches of data, sickness, or family emergencies could make the risk of shocks to output or the cost of effort greater on certain days (see Sec. IV.E). In addition, workers might go through periods of present-biasedness, for example, because of variation in family circumstances, seasonal variation in other income sources, or shocks that increase or decrease exposure to habit-forming goods such as alcohol, tobacco, or caffeine. Access to other motivational

²² Under the null hypothesis of time consistency, take-up of dominated contracts should be zero. We view 36 percent as a sizable take-up rate and interpret it as a useful and important summary statistic in support of our model. However, ultimately our test of proposition 2 is whether the contract treatments affect production and earnings.

TABLE 3
 CONTRACT TREATMENTS: TAKE-UP OF DOMINATED CONTRACTS (Summary Statistics)

	Mean	Standard Deviation
Dominated contract chosen: conditional on attendance	.36	.31
Dominated contract chosen: target = 0 if absent	.28	.26

NOTE.—The sample is observations in which workers were assigned the option to choose a dominated contract. The first row limits analysis to observations in which a worker was present the day before or the day of treatment assignment. The second row codes target choice as 0 if absent the day before or the day of assignment.

devices that reduce the need for dominated contracts could also vary across days (see, e.g., Kaur et al. 2010).

Table 4 presents treatment effects of giving workers the option to choose targets. Relative to assignment to the control contract, assignment to the option to choose a dominated contract treatments increased production by 120 fields or 2 percent (significant at the 5 percent level; col. 1). Looking within these treatments, the evening option to choose a dominated contract increased output by 3 percent (significant at the

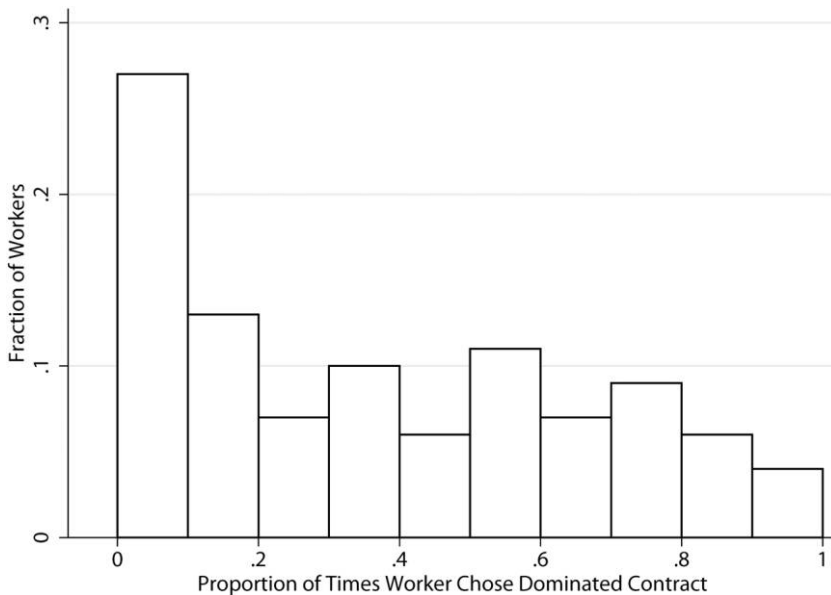


FIG. 4.—Take-up of dominated contracts: distribution of worker means. A worker's take-up rate is the proportion of times the worker chose a dominated contract (i.e., selected a positive target) when given the option (conditional on being present the day before and the day of assignment to the option to choose a dominated contract treatment). The distribution is shown for the 101 workers in the sample who were assigned the option to choose at least once.

TABLE 4
 CONTRACT TREATMENTS: TREATMENT EFFECTS OF CONTRACTS

	DEPENDENT VARIABLE				
	Production				
	Control and Option Sample (1)	Control and Option Sample (2)	Full Sample (3)	Earnings: Full Sample (4)	Attendance: Full Sample (5)
Option to choose dominated contract	120 (59)***				
Evening option to choose dominated contract		156 (69)**	150 (69)**	4.60 (2.17)**	.01 (.01)
Morning option to choose dominated contract		84 (69)	73 (69)	2.32 (2.17)	-.00 (.01)
Target imposed:					
Level 1 target			3 (90)	-1.55 (2.88)	-.00 (.01)
Level 2 target			213 (91)**	3.13 (2.89)	-.01 (.01)
Level 3 target			334 (150)**	5.01 (4.80)	-.01 (.02)
Observations: worker-days	6,310	6,310	8,423	8,423	8,423
R^2	.60	.60	.59	.57	.15
Dependent variable mean	5,311	5,311	5,337	172	.88

NOTE.—The dependent variables equal zero if the worker is absent. Columns 1 and 2 limit analysis to observations in which workers were assigned to the control contract or given the option to choose a dominated contract. Columns 3–5 include the full sample. All regressions include worker, date, and computer seating assignment fixed effects. Regressions 1–3 also include lagged production controls, and regression 4 includes lagged earnings controls. Results from ordinary least squares regressions are shown; robust standard errors are reported in parentheses.

* Significant at the 10 percent level.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

5 percent level) and the morning option to choose a dominated contract insignificantly increased output, though we cannot reject that these two coefficients are equal (col. 2). Not surprisingly, exogenously imposing a target on workers increased output—with larger effects for higher targets (col. 3).

Earnings also increased when workers were given the option to choose their own targets (col. 4). As with output, the evening option to choose a dominated contract increased average earnings by 3 percent and the morning option had a positive but insignificant impact on earnings. In contrast, the exogenously imposed targets did not increase earnings significantly because of how often workers missed them, causing workers to lose half their piece rate earnings for the day (see below). The contract treatments have no effects on attendance on average (col. 5). They also do not alter the quality of output; we see no effect on accuracy (app. table 3, panel B).

The implied treatment on the treated effect of choosing a positive target on output is approximately 6 percent. Given the estimated elasticity of 0.33, the magnitude of this effect corresponds to an 18 percent increase in the piece rate. Using proposition 2, we can back out the implied bound on the departure from time consistency. The treatment on the treated effect of 6 percent implies that across workers on average,

$$\frac{d^i(T) - d^i(t)d^i(T-t)}{d^i(t)d^i(T-t)} \geq \frac{0.06}{0.33} = 0.18.$$

On average, workers value the benefits of wages on payday, relative to the costs of effort on the workday, by at least 18 percent more at the time of contract choice than in the moment of effort—a major departure from time consistency.

We estimate that the chosen targets are aggressive enough that workers would have missed them 9.1 percent of the time if they had been assigned to the control contract that day (table 5, panel A, col. 1). Recall that in the model, workers choose targets in such a way that their future selves never miss them. In the actual experiment, workers missed their chosen targets under the option to choose a dominated contract treatment (conditional on choosing a positive target and being present) 2.6 percent of the time (table 5, panel B, col. 1). When they missed their chosen targets, the earnings loss corresponded to almost half their mean daily earnings. The percentage of times workers missed the exogenously imposed targets—which were set at the 30th, 50th, and 70th percentiles of the control contract output distribution—was larger, at 8.6–14.1 percent (table 5, panel B, col. 1). This helps explain why the target-imposed treatments did not significantly raise earnings (table 4).

TABLE 5
CHANCE OF MISSING TARGETS: SUMMARY STATISTICS

	All Workers (1)	High Payday Effect Workers (2)	Low Payday Effect Workers (3)
A. Predicted Probability of Missing Target (Using Control Contract Output Distribution)			
Self-chosen target	.091 (.120)	.118 (.146)	.073 (.096)
B. Proportion of Times Target Was Actually Missed			
Self-chosen target	.026 (.122)	.052 (.187)	.008 (.036)
Target imposed:			
Level 1 target	.086 (.155)	.120 (.198)	.063 (.115)
Level 2 target	.113 (.197)	.117 (.240)	.110 (.167)
Level 3 target	.141 (.265)	.177 (.315)	.115 (.224)

NOTE.—Panel A reports the probability that workers would have missed their chosen targets if they had been assigned to the control contract that day. This is computed as follows. For observations in which workers were in attendance, we estimate a regression of production on worker, date, and computer fixed effects; lag production controls; payday distance dummies; contract assignment dummies; and log experience. For each observation in which a worker was assigned to the option to choose a dominated contract, selected a positive target, and was present, we predict the worker’s production under the control contract on that day using the estimates from the above regression. To this predicted value, we add the worker’s residuals from the above regression to arrive at a vector of potential production values, which we fit to a lognormal distribution. Evaluating the cumulative distribution function of this distribution at the chosen target level gives an estimate of the probability the worker would have missed her target under the control contract. Panel B, row 1, reports the proportion of times production was below workers’ chosen targets. Rows 2–4 of panel B report this statistic when targets were exogenously imposed on workers. High payday effect workers are those whose payday effect—the difference between mean production on paydays and nonpaydays under assignment to the control contract divided by mean production under the control contract—is above the sample average. Column 1 presents these statistics for the workers who chose a positive target at least once and for whom the payday effect can be computed (8,240 worker-days and 90 workers); cols. 2 and 3 report these statistics separately for high and low payday effect workers (5,024 worker-days and 54 workers, and 3,216 worker-days and 36 workers, respectively).

These results are consistent with the existence of stochastic shocks to output or the cost of effort. For example, at any given time, there was a risk that network slowdowns could severely impede productivity for the remainder of the day. Note that shocks create additional costs for workers beyond the financial penalty, such as having to stay in the office late to meet their target on days when there are negative shocks to output or being unable to leave early when there are unexpectedly high effort costs.

In the presence of shocks, even modest target levels could prompt risk-averse workers to work hard to ensure they clear their target before a shock arrives. Once the target is achieved, the return to effort is not zero; workers would continue working until the marginal cost of effort equals its marginal return—discounted from the view of the self that is exerting the effort. In this way, targets set below mean output levels can generate output and earnings increases among risk-averse agents. Overall, the results in tables 4 and 5 support the view that when workers chose their own targets, they did so sensibly—balancing the motivational benefits with the risk of lost earnings—leading them to choose target levels that increased their average earnings.

Note that in the model, labor supply is separable across periods. Consistent with this, we find no evidence that higher effort on one day increases the cost of effort on subsequent days (app. table 5). Specifically, assignment to the option to choose a dominated contract or target imposed (relative to the control contract) does not reduce production the next day.²³

C. Heterogeneity in Preferences: Correlation between Payday and Contract Effects (Tests 3 and 4)

The payday and contract choice results each suggest that at least some workers are time inconsistent. For example, 41 percent of workers have a 10 percent or larger pay cycle effect. Similarly, after 2 months of experience, 49 percent of workers select dominated contracts 25 percent of the time or more. More formally, we strongly reject that workers are homogeneous in these effects. To test for heterogeneity in payday effects, we regress production on a payday dummy, worker fixed effects, interactions of each worker fixed effect with the payday dummy, and standard controls. The p -value of the F -test of joint significance of the interaction coefficients is .000. Similarly, to test for heterogeneity in treatment effects of contracts, we limit the sample to control and option to choose a dominated contract observations and regress production on worker fixed effects, an option to choose a dominated contract assignment dummy, interactions of each worker fixed effect with this dummy,

²³ Since workers are assigned to each treatment a fixed number of times in each 12-day period, assignment on a given day is correlated with the probability of future treatments in each block. This mechanical correlation could affect the estimates in table 4. In app. table 6, we control for the probabilities of the worker receiving each contract assignment for that observation given the worker's previous assignments in that randomization block. An F -test of joint significance of the probability covariates has a p -value of .45, indicating that the assignment probabilities have little predictive power. Their inclusion also has little impact on the estimated treatment effects.

and standard controls. The p -value of the F -test of joint significance of the interaction coefficients is .003.

Proposition 3 predicts that the payday and contract effects will be positively correlated. To test this, we define the payday effect for each worker as

$$\begin{aligned} \text{Payday effect} = & (\text{Mean production on paydays} \\ & - \text{Mean production on nonpaydays}) \\ & \div \text{Mean production in sample.} \end{aligned}$$

This measure is computed using only observations in which workers were assigned to the control contract treatment.²⁴ Note that we chose this as our summary measure of a worker's payday effect at the start of the empirical analysis because this measure does not take a strong *ex ante* stance on the nature of time inconsistency. The prediction that is common to both hyperbolic and quasi-hyperbolic models is that of output increases on paydays. Even in hyperbolic models, increases closer to payday are expected to be most pronounced. In addition, calibrating a hyperbolic parameter for each worker using the increase over the full workweek would necessarily require (arbitrary) functional form assumptions. We therefore use the simple proportional difference in means between paydays and nonpaydays. In *ex post* analysis, we have confirmed that the results are robust to other measures that capture the pay cycle effect.

On average, workers with an above-average payday effect are 13.8 percentage points more likely to select a positive target and select targets that are 351 fields higher (table 6, panel A). These coefficients correspond to a striking 47 percent and 49 percent of the mean take-up rate and target level, respectively, and are both significant at the 1 percent level. The payday effects heterogeneity is not driven by other potentially correlated interpersonal differences among workers, such as productivity (table 6, panel A, col. 3). Workers with large payday effects also increase production more in response to dominated contracts. In table 6, panel B, column 1, the interaction between the option to choose a dominated contract and the high payday effect dummies is 482 fields: 9 percent of mean production (significant at the 1 percent level), implying a treatment on the treated effect of 28 percent. The effect on earnings has the same magnitude (table 6, panel B, col. 2) and is shown graphically in figure 5.

²⁴ We can compute this statistic only for workers who were assigned to the control contract on both paydays and nonpaydays during employment. This reduces our sample size for this analysis from 8,423 to 8,240 observations.

TABLE 6
PAY CYCLE EFFECTS: CORRELATION WITH DOMINATED CONTRACT EFFECTS

	A. TAKE-UP OF DOMINATED CONTRACTS		
	Dependent Variable		
	Dominated Contract Chosen (1)	Target Level Chosen (2)	Target Level Chosen (3)
High payday effect worker	.138 (.044)***	351 (129)***	338 (126)***
High-productivity worker			-105 (134)
Observations: worker-days	4,098	4,098	4,098
R^2	.2	.22	.22
Dependent variable mean	.28	759	759
	B. TREATMENT EFFECTS OF CONTRACTS		
	Dependent Variable		
	Production (1)	Earnings (2)	Attendance (3)
Option to choose dominated contract	-69 (74)	-2.24 (2.34)	-.016 (.010)
Option to choose dominated contract \times high payday effect worker	482 (126)***	15.15 (3.99)***	.058 (.019)***
Target imposed	-35 (86)	-3.82 (2.74)	-.019 (.012)*
Target imposed \times high payday effect worker	483 (148)***	14.31 (4.71)***	.042 (.022)*
Observations: worker-days	8,240	8,240	8,240
R^2	.59	.57	.11
Dependent variable mean	5,355	173	.875

NOTE.—The table present ordinary least squares regressions; standard errors are reported in parentheses. The panel A sample is observations in which workers were given the option to choose a dominated contract. If a worker was absent the day before or the day of treatment assignment, both dependent variables are coded as zero. High payday effect worker is a binary indicator for whether the worker's mean payday effect (the difference in production on paydays and nonpaydays under assignment to the control contract, divided by mean production under the control contract) is above the sample average. High productivity worker is a binary indicator for whether the worker's mean production is above the sample average. Regressions include computer and date fixed effects and lagged earnings controls. Standard errors are clustered by worker. The panel B sample comprises all observations. The dependent variables equal zero if a worker was absent. Each regression includes worker, date, and computer seat assignment fixed effects. The col. 1 and 2 regressions also include lagged production controls and lagged earnings controls, respectively. Robust standard errors are reported.

* Significant at the 10 percent level.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

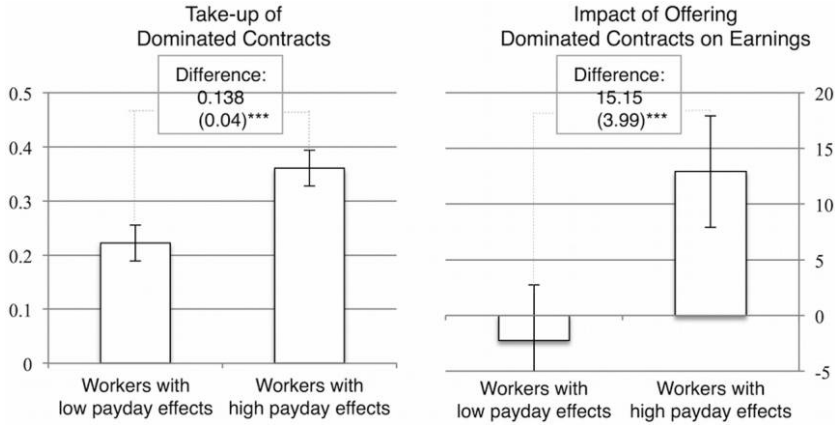


FIG. 5.—Pay cycle effects: correlation with contract choice and earnings impact. These figures show differences in take-up rates and treatment effects of dominated contracts. Workers with low (high) payday effects are those whose payday effect—the difference in production on paydays and nonpaydays under assignment to the control contract, divided by mean production under the control contract—is below (above) the sample average. The top of each chart displays point estimates and standard errors corresponding to regressions shown in table 6. Each bar corresponds to the estimated mean for each group, along with 95 percent confidence intervals.

High payday effect workers are also more likely to show up to work when assigned the option to choose a dominated contract or target imposed (table 6, panel B, col. 3).²⁵ This provides additional evidence that high payday effect workers demand dominated contracts.

Given our labor supply elasticity estimate of 0.33, the 9 percent intent to treat effect implies that providing high payday effect workers with simply the option to select targets leads to production increases comparable to a 27 percent increase in the piece rate wage. This magnitude corresponds to a 1-year increase in education. Using proposition 2, the treatment on the treated effects allow us to bound the level of time inconsistency of the workers that, on the basis of their pay cycle behavior, appear most time inconsistent. When these workers choose dominated contracts,

$$\frac{d^l(T)/d^l(t)}{d^l(T-t)/d^l(0)} - 1 \geq \frac{0.28}{0.33} = .084;$$

²⁵ The large production effect of the option to choose a dominated contract on high payday impact workers does not seem to be driven completely by the impact on attendance. For high payday impact workers, the average treatment effects on production and attendance are 395 fields and 4.4 percentage points, respectively. For these workers, mean production conditional on attendance is 5,581 fields. As a simple calibration, $5,581 \times 0.044 = 245 < 395$. Moreover, regressing production conditional on attendance on the contract treatment dummies yields positive and significant coefficients (app. table 3), although these are difficult to interpret since attendance is endogenous.

this implies that the relative value of the wage benefits to effort costs is 84 percent higher at the time of contract choice than at the time of effort.²⁶

The workers most affected by payday effects also select more aggressive targets (table 5, cols. 2 and 3). We estimate that high payday effect workers would have missed their selected targets 11.8 percent of the time had they been under the control contract and actually miss them 5.2 percent of the time. In contrast, these statistics are 7.3 percent and 0.8 percent, respectively, for low payday effect workers.²⁷

We do not see changes in take-up of dominated contracts over the pay cycle (table 7, panel A).²⁸ Consistent with proposition 4, when high payday effect workers are closer to their payday (and the self-control problem is therefore smaller), the treatment effect of the option to choose a dominated contract is smaller. For these workers, the earnings impact of being able to choose targets is Rs. 33, or 21 percent, lower on paydays than on nonpaydays (table 7, panel B, col. 1). On average, the earnings impact of the option to choose treatment declines by Rs. 3 (or 2 percent) per day as the payday approaches (panel B, col. 2). The effect of being assigned to an exogenous target follows a similar pattern for these workers. In contrast, low payday effect workers—who are not affected by the dominated contract treatments—exhibit no detectable trends over the pay cycle (panel B, cols. 3 and 4).

In the above analysis, we regress contract choice on payday effects, rather than the other way around, because contract choice—in particular the acceptance of dominated contracts—will depend not only on whether workers are time consistent but also on their degree of sophistication. A regression of payday effects on contract choice will thus be subject to an errors-in-variables problem. Nonetheless, we show regressions of this type in appendix table 9 and plot a corresponding figure in appendix figure 2. Specifically, using our standard specification (with earnings as the dependent variable), there is some evidence that workers with above-average take-up rates of dominated contracts have steeper earnings increases over the pay cycle (app. fig. 2, panel A). Since workers with different take-up rates also differ in their productivity levels, repeating this analysis using log production strengthens this result: the two

²⁶ There are no significant differences in the elasticity of output with respect to wages between workers with above- and below-average payday effects (app. table 7, cols. 2 and 3).

²⁷ When low payday effect workers miss their targets (whether they are self-chosen or exogenously imposed), this decreases their future probability of taking up the dominated contract by 12 percentage points and also decreases the target levels they select. In contrast, missing targets appears to have no impact on the future take-up behavior of high payday effect workers on average (app. table 8).

²⁸ Note that table 7 regressions use observations after 1 month of experience to examine trends after workers have learned about the contracts (see Sec. IV.D). Using all observations does not qualitatively change the results but slightly decreases their precision.

TABLE 7
CORRELATION IN PAYDAY AND DOMINATED CONTRACT EFFECTS:
TRENDS OVER THE PAY CYCLE

A. DEMAND FOR DOMINATED CONTRACTS (Dependent Variable: Dominated Contract Chosen)				
	High Payday Effect Workers		Low Payday Effect Workers	
	(1)	(2)	(3)	(4)
Payday dummy	-.013 (.030)		-.002 (.024)	
Day in pay cycle (linear control)		.006 (.005)		-.002 (.004)
Observations	1,252	1,252	1,967	1,967
Dependent variable mean	.319	.319	.23	.23
B. TREATMENT EFFECTS OF CONTRACTS (Dependent Variable: Earnings)				
	Payday Dummy (1)	Day in Pay Cycle (Linear Control) (2)	Payday Dummy (3)	Day in Pay Cycle (Linear Control) (4)
Pay cycle measure	32.88 (7.59)***	6.32 (1.49)***	-3.40 (6.00)	.69 (1.09)
Option to choose dominated contract	21.58 (4.55)***	25.49 (6.81)***	-4.41 (3.14)	-9.33 (4.93)*
Option to choose dominated contract × pay cycle measure	-33.33 (9.16)***	-3.33 (1.70)**	11.93 (6.90)*	2.39 (1.28)*
Target imposed	18.90 (5.22)***	26.17 (7.96)***	-.67 (3.63)	-1.73 (5.75)
Target imposed × pay cycle measure	-32.71 (10.70)***	-4.24 (2.04)**	-1.28 (8.30)	.278 (1.51)
Observations	2,502	2,502	3,947	3,947
Dependent variable mean	159	159	196	196

NOTE.—Columns 1 and 2 report results for high payday effect workers: those whose mean payday effect (the difference in production on paydays and nonpaydays under assignment to the control contract, divided by mean production under the control contract) is above the sample average. Columns 3 and 4 report results for workers with a below-average payday effect. Ordinary least squares regressions are presented; robust standard errors are reported in parentheses. Analysis is limited to observations after about 1 month (20 workdays) of experience to reflect learning (see Sec. IV.D). The panel A sample is observations in which workers were given the option to choose a dominated contract. If a worker was absent the day before or the day of treatment assignment, the dependent variable is coded as zero. Regressions include computer and date fixed effects and lagged earnings controls. The panel B sample comprises observations from all contract treatments. Each regression includes worker, date, and computer seat assignment fixed effects. The col. 1 regression also includes lagged earnings controls.

* Significant at the 10 percent level.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

groups have significantly different pay cycle trends. Specifically, workers with above-average take-up of dominated contracts have greater output increases 2 days before their payday (16 log points; significant at 1 percent), 1 day before their payday (13 log points; significant at 5 percent), and on their payday (10 log points; significant at 5 percent) (app. table 9, col. 1). On average, the slope of the output increase over the pay cycle for workers with high dominated contract demand is more than twice as large as those with low demand (app. table 9, col. 2).

D. Learning over Time

As workers gain experience, do they learn about the value of the dominated contracts or perhaps find other ways around their self-control problems? Averaging across all workers, we do not find significant trends in take-up of dominated contracts (table 8, col. 1).

However, this masks substantial heterogeneity. Figure 6 plots experience (number of workdays in the experiment) against the proportion of workers choosing positive targets (i.e., dominated contracts). High payday effect workers are shown in closed circles and low payday effect workers are shown in open circles. Mean take-up rates of dominated contracts among the two groups are initially similar. As they gain experience, there is a divergence: low payday effect workers decrease take-up of dominated contracts while high payday effect workers increase take-up (albeit insignificantly). After 2 months of experience, high payday effect workers are 20.6 percentage points, or 73 percent, more likely to select positive targets than low payday effect workers (p -value of .000; table 8, panel A, col. 3).

The impact of paydays on output does not change with experience, suggesting that underlying self-control problems do not change over time (table 8, panel B). However, the treatment effect of giving workers the option to choose a dominated contract grows with experience. This is consistent with the trends in table 8, panel A, which indicate that the group of workers that benefits most from the dominated contracts is more likely to select them over time. Given the long horizon of the study, the results in table 8 imply that time inconsistency is a persistent problem in the workplace.

Appendix table 10 examines these learning trends as a function of a worker's initial take-up rate (the proportion of times the worker selected a positive target under the option to choose a dominated contract in the first 10 workdays of experience). Among high payday effect workers, 25 percent have an initial take-up rate of zero. After 2 months of experience, such workers increase their take-up by 15.4 percentage points (or 48 percent) on average, while there is no detectable change for those with initially higher take-up rates (cols. 1 and 2). While only suggestive,

TABLE 8
CHANGES IN OUTCOMES WITH WORKER EXPERIENCE

	EXPERIENCE MEASURE		
	Log Number of Days Worked (1)	Log Number of Days Worked (2)	More than 2 Months Worked (3)
A. Take-Up of Dominated Contract			
Experience measure	-.040 (.027)	-.066 (.030)**	-.096 (.057)*
High payday effect worker	.139 (.043)***	-.087 (.100)	.074 (.050)
Experience measure × high payday effect worker		.062 (.026)**	.132 (.054)**
Dependent variable mean	.28	.28	.28
B. Treatment Effects on Earnings			
Experience measure	7.45 (2.11)***	4.67 (2.34)**	-3.82 (4.93)
Option to choose dominated contract	3.38 (1.87)*	-8.26 (5.54)	-.87 (2.28)
Experience measure × option to choose dominated contract		3.20 (1.63)**	8.93 (3.78)**
Payday	4.84 (1.99)**	4.65 (5.72)	4.36 (2.48)*
Experience measure × payday		.05 (1.62)	.96 (3.79)
Dependent variable mean	172	172	172

NOTE.—Ordinary least squares regressions; standard errors are reported in parentheses. In cols. 1 and 2, the measure of experience is the log of the number of workdays a worker has been in the sample. In col. 3, the measure is a binary indicator for more than 50 workdays (~2 months) of experience. The panel A sample comprises 4,098 observations in which workers were given the option to choose a dominated contract. The dependent variable is defined as zero if a worker was absent the day before or the day of treatment assignment. High payday effect worker is a binary indicator for whether the worker's mean payday effect (the difference in production on paydays and nonpaydays under assignment to the control contract, divided by mean production under the control contract) is above the sample average. All regressions control for date and computer seat assignment fixed effects and lagged earnings. Standard errors are clustered by worker. In panel B, the sample comprises all 8,423 observations. Production is defined as zero when a worker is absent. The covariates in each regression are dummies for option to choose a dominated contract, target imposed (not shown), payday, and interactions of the experience measure with each indicator. All regressions also include worker, date, and computer seat fixed effects and lagged earnings controls. Robust standard errors are reported.

* Significant at the 10 percent level.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

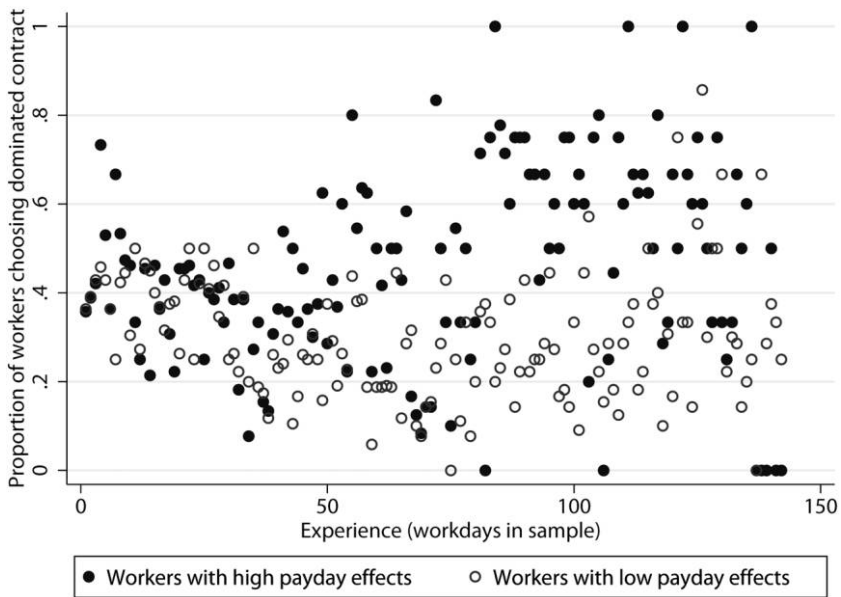


FIG. 6.—How the demand for dominated contracts changes with experience. Worker experience is the number of workdays the worker has been in the sample. The proportion of times positive targets were chosen is computed for each value of the experience variable using observations in which the worker was given the option to choose a dominated contract (conditional on being present both the day before and the day of treatment assignment). High (low) payday effect workers are those whose mean payday effect—the difference in production on paydays and nonpaydays under assignment to the control contract, divided by mean production under the control contract—is above (below) the sample average.

this is consistent with a story in which workers with high payday effects but zero initial take-up are naive and become sophisticated as they experience the dominated contracts through the contract randomizations over time. In contrast, among low payday effect workers, those whose initial take-up rate is above 50 percent sharply decrease take-up with experience, while those with initially lower take-up rates remain stable over time (cols. 3 and 4). Again, this is consistent with a story of learning among low payday effect workers, whose output does not increase, on average, from the option to choose.

E. Morning and Evening Choice (Test 5)

Contrary to our initial expectations, on average across the whole sample, workers did not select higher targets in the evening before work than in the morning of work (table 9, cols. 1 and 2). Note that positive take-up of dominated contracts in the morning of the workday implies that time

TABLE 9
DEMAND FOR THE DOMINATED CONTRACT: IMPACT OF UNCERTAINTY AND TIMING OF CHOICE

	DEFINITION OF HIGH UNCERTAINTY INDICATOR							
	Worker's Assigned Computer Is Sensitive to Network Fluctuations				Worker's Morning Arrival Time Is Sensitive to Bus/Train Schedules			
	Dominated Contract Chosen (1)	Target Level Chosen (2)	Dominated Contract Chosen (3)	Target Level Chosen (4)	Dominated Contract Chosen (5)	Target Level Chosen (6)	Dominated Contract Chosen (7)	Target Level Chosen (8)
Evening option to choose dominated contract	-.002 (.012)	-18 (37)	-.003 (.010)	-26 (32)	.066 (.022)***	168 (72)**	.027 (.016)	87 (46)*
High uncertainty indicator			-.013 (.016)	-134 (63)**	.027 (.022)	-20 (82)	.104 (.062)	282 (206)
Evening option to choose dominated contract × high uncertainty indicator					-.082 (.024)***	-230 (78)***	-.070 (.029)**	-253 (97)**
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Seat fixed effects	Yes	Yes	No	No	No	No	Yes	Yes
Observations	4,193	4,193	4,193	4,193	4,193	4,193	3,106	3,106
R ²	.33	.34	.32	.33	.32	.33	.12	.13
Dependent variable mean	.28	767	.28	767	.28	767	.3	803

NOTE.—This table tests whether the horizon of contract choice and uncertainty affect demand for the dominated contract. The sample comprises worker-day observations in which workers were given the option to choose a dominated contract (in the evening or morning). Dominated contract chosen is a binary indicator for whether the worker selected a target above zero. Target level chosen is the continuous target level selected. Both dependent variables are defined as zero if a worker was absent the day before or the day when given the option to choose a contract. Evening option to choose dominated contract is a dummy that equals one if the worker was assigned to choose the evening before the workday and equals zero if the worker was assigned to choose the morning of the workday. In cols. 3–6, the high uncertainty indicator equals one if the worker was assigned to a computer that was highly sensitive to office network speed and equals zero otherwise. In cols. 7 and 8, the high uncertainty indicator equals one if during the end line survey the worker “agreed strongly” with the statement “The bus/train schedules really impact whether I can get to work on time because if I miss one bus or train, the next one I can take is much later.” All regressions include date fixed effects. Robust standard errors are reported in cols. 1 and 2. Standard errors are corrected to allow for clustering by computer in cols. 3–6 and by worker in cols. 7 and 8.

* Significant at the 10 percent level.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

inconsistency operates at time periods shorter than a day. Consistent with this, 40 percent of the workers in the end line survey agreed with the statement, “Some days I get tempted to leave work earlier than I would like” (table 1).

Why might proposition 5 fail? As discussed in Section II, dominated contracts are less attractive when agents face exogenous risks. Ex post analysis and qualitative work suggest that in the evening before the workday, workers faced two types of uncertainty that were partially resolved by the morning of the workday; agents thus sometimes faced greater costs of choosing targets the evening before work than the morning of work.

First, network speed fluctuations affected the rate at which workers could send data entered from an image to the central server and retrieve the next image for entry. This wait time ranged from 1 second to over 5 minutes. When workers arrived to the office in the morning, they received new information on network speed and could use this to inform their target choice. This information was especially valuable for workers on “bad” computers, since network shocks greatly affected productivity for these computers.

To test whether network uncertainty deterred workers from choosing targets the evening before work, we asked the office management staff to consult workers to identify which computers were more sensitive to network slowdowns. Management did not know that the list would be used for this purpose. The computers identified as more uncertain are indeed more sensitive to overall network fluctuations (see app. table 11 and app. fig. 3). Assignment to the more uncertain computers decreases mean output by 313 fields or 6 percent (app. table 12)—a magnitude that corresponds to an 18 percent reduction in the piece rate based on our elasticity estimate. Workers respond by picking targets that are 134 fields lower, on average. This suggests that the uncertainty that causes workers to shy away from dominated contracts is significant. Consistent with the results in tables 3, 4, and 5, workers appear to trade off income risk against the self-control benefits of dominated contracts.

When workers are assigned to a good computer (i.e., a computer that is not as sensitive to network fluctuations), they are 6.6 percentage points more likely to choose a dominated contract when given the choice the evening before production than the morning of production. However, when assigned to a bad computer, they are 1.6 percentage points less likely to choose a positive target in the evening than in the morning (table 9, col. 5).

Second, many workers also faced uncertainty regarding commute time and thus arrival time, which was resolved by the time they arrived to work in the morning. In the end line survey, those who “agree strongly” with the statement “The bus/train schedules really impact whether I can get to work on time because if I miss one bus or train, the next one I can take is

much later” select targets more often the morning of production than the evening before production. The opposite is true for workers with less uncertain commute times (table 9, col. 7).

These results are consistent with the hypothesis that, all else equal, a greater gap between the period of contract choice and the period of effort increases target levels; however, there are greater expected costs of choosing targets before the uncertainty of network speed and arrival time are resolved. In our data, when uncertainty is lower and similar between the evening before and the morning of work, workers are more likely to choose a dominated contract in the evening—further from the moment of temptation. However, when uncertainty is high the evening before production but is reduced by the next morning, take-up is higher in the morning. These findings indicate that contract demand can interact strongly with uncertainty.

E. Correlates of Take-Up and Treatment Effects

While payday effects strongly predict demand for dominated contracts, we see much less predictive power from a range of self-control correlates commonly used in the literature on psychology and economics. In columns 1–3 of table 10, we look at measures of self-control problems based on self-reports by workers during the end line survey. The correlate in column 1 is the demeaned self-control factor, obtained from a factor analysis on the end line data. In column 2, we construct a demeaned self-control index by averaging each worker’s responses to nine self-control questions. In column 3, we use a binary indicator for whether male workers said they had tried to quit drinking, smoking, or chewing tobacco and failed. Each of these three columns shows similar results. Workers with higher values of the correlates are less productive on average. Each correlate positively predicts demand for the dominated contract and positively predicts treatment effects of the contracts. However, among these, only the coefficients on the self-control factor are generally significant. None of these correlates predicts the payday effect.

Laboratory measures of time preference—computed by asking workers to make binary choices between monetary rewards at different time horizons (see the note to table 1)—also have limited predictive power. In column 4, the correlate is impatience: the proportion of times the worker chose a smaller immediate reward rather than a larger delayed reward. The column 5 correlate is preference reversals: the proportion of times a worker chose the smaller immediate reward in the short horizon but then displayed patience when choosing between the same amounts in the long horizon. These measures positively (but insignificantly) predict demand for dominated contracts. As before, workers with greater values of these measures are less productive, on average, but have larger contract treatment effects.

TABLE 10
 HETEROGENEITY IN TREATMENT EFFECTS: LAB AND SURVEY CORRELATES OF SELF-CONTROL

CORRELATE OF SELF-CONTROL						
Self-Control Factor (1)	Self-Control Index (2)	Addictive Behaviors Dummy (3)	Discount Rate: Impatient Responses (4)	Discount Rate: Preference Reversals (5)	Years of Education (6)	IQ Test Index Score (7)
A. Dependent Variable: Dominated Contract Chosen						
Correlate	.057 (.046)	.140 (.082)*	.068 (.115)	.143 (.189)	.029 (.015)**	-.001 (.002)
Observations	3,106	2,245	2,454	2,470	4,056	4,089
R ²	.22	.26	.23	.24	.19	.19
B. Dependent Variable: Target Level Chosen						
Correlate	147 (140)	354 (238)	220 (342)	533 (572)	122 (43)***	-1 (5)
Observations	3,106	2,245	2,454	2,454	4,056	4,089
R ²	.25	.29	.27	.28	.23	.22
C. Dependent Variable: Earnings						
Correlate	-8.01 (3.26)**	-8.69 (12.59)	-39.73 (10.05)***	-25.54 (14.56)*	4.02 (2.03)*	.39 (.19)**
Option to choose dominated contract	2.84 (2.60)	1.27 (3.62)	3.60 (2.63)	3.71 (2.61)	4.59 (2.31)**	4.72 (2.33)**

Option to choose dominated contract × correlate	5.19 (2.91)*	6.81 (4.65)	13.77 (8.37)	22.03 (9.65)**	23.75 (14.17)**	1.71 (1.42)	-0.02 (.16)
Payday	5.48 (2.72)**	5.50 (2.74)**	4.22 (3.68)	4.78 (3.29)	4.65 (3.33)	5.40 (2.30)**	6.30 (2.20)***
Payday × correlate	.83 (2.77)	2.34 (4.30)	1.75 (7.82)	-1.39 (9.70)	-7.71 (17.32)	.60 (1.25)	.00 (.12)
Observations	4,674	4,674	3,376	3,701	3,701	6,101	6,149
R ²	.55	.55	.53	.57	.56	.54	.53

NOTE.—The self-control correlate in col. 1 is the self-control factor, obtained from a principal factors analysis on the end line survey. The col. 2 correlate is a self-control index, obtained by averaging each worker's responses to the nine self-control questions in the end line. The col. 3 correlate equals one if the worker said he has tried to quit drinking, smoking, or chewing tobacco and failed, and equals zero otherwise (male workers only). The col. 4 and 5 correlates are from the discount rate exercise (see the note to table 1). The col. 4 correlate is the proportion of times the worker chose the smaller immediate reward instead of the larger delayed reward. The col. 5 correlate measures preference reversals: the proportion of times a worker chose the smaller immediate reward in the short horizon but the larger delayed reward in the long horizon. The col. 6 and 7 correlates are, respectively, years of education and composite IQ score (the sum of the Raven's Matrix and Digit-Span test scores). The self-control correlate in each column has been demeaned. In panels A and B, the sample is observations in which workers were assigned the option to choose a dominated contract. The dependent variables in panels A and B are defined as zero if a worker was absent the day before or the day of treatment assignment. The panel C sample is observations in which workers were assigned to the control contract or the option to choose a dominated contract. All regressions include date and computer seat assignment fixed effects and lagged earnings controls. Observations change between columns because not all workers provided education information or took the IQ tests and because the end line survey and discount rate exercise were administered only at the end of the project. Ordinary least squares regressions are shown. Standard errors are clustered by worker.

* Significant at the 10 percent level.

** Significant at the 5 percent level.

*** Significant at the 1 percent level.

Education (col. 6) positively predicts take-up of dominated contracts but does not predict treatment effects. IQ (col. 7)—the sum of the worker's scores on the Raven's Matrix and Digit Span tests—does not predict any of the effects.²⁹

The strong correlation between the payday and contract effects (documented in Sec. IV.C) indicates that there are stable interpersonal differences across field behaviors—evidence for which has been limited in the literature. However, the findings in table 10 are consistent with those of other studies: laboratory and survey measures of self-control predict field behavior, but often to a limited extent (see, e.g., Chabris et al. 2008), perhaps because of the various measurement issues with laboratory measures (see table 1 in Chabris, Laibson, and Schuldt 2008; Augenblick et al. 2014). Or this suggests that self-control is context dependent: predicting it in the workplace requires measures specific to that context.

V. Alternative Explanations

The results are largely consistent with a self-control agency model. Could they be explained within the context of a standard exponential discounting model? We argue that while other models could explain any one result, self-control problems are required—at least to some degree—to fit the full pattern of results: the production increases on paydays, sustained demand for dominated contracts and treatment effects of contract choice, and the correlation between the payday effects and demand for dominated contracts.

First, could workers be choosing dominated contracts because they are confused? Recall that during the training period, we assigned workers to the various contracts and also tested their comprehension using a contract quiz—the mean score on which was 93 percent. Take-up is not being driven by those who have worse understanding of the contracts: quiz performance is positively (although insignificantly) correlated with take-up, and education strongly predicts take-up. Moreover, demand for dominated contracts persists over the long horizon of the study.

Second, could workers be choosing dominated contracts to signal ability to employers? Since the employer observes production directly, there is no reason to believe that a worker who is productive under the control contract is not more impressive than one who needs to rely on a dominated contract. Moreover, it is unclear why workers with larger payday effects should be more likely to signal ability, or why workers would be

²⁹ There is also some heterogeneity in effects by gender. The option to choose has larger treatment effects on output and earnings for men than for women. However, there is no difference in the pay cycle effects by gender.

more likely to signal ability the evening before the workday than the morning when assigned to good computers.

Could weekly income targeting explain the payday effects? Income targeting implies a sharp decrease in marginal utility for income levels above the target (see Camerer et al. 1997; Dupas and Robinson 2013). Two pieces of evidence suggest that this is not happening in our data. First, as we saw in our test for intertemporal substitution, exogenous production increases do not decrease production on subsequent days (see app. table 5). Second, a targeting model delivers an even finer testable prediction: an unexpected production increase today will lead to a larger reduction in tomorrow's effort if the worker is closer to her payday, because there are fewer subsequent days over which the adjustment needs to be made. In appendix table 13, we examine the impact of being assigned to a target (and thus increasing production) the previous day. We see no evidence that this reduces production, especially around the payday. Finally, since the impact of day-to-day shocks is adjusted within the pay week to arrive at the weekly target, under income targeting the variance in production among pay weeks should be less than the variance in production among weeks defined according to some other arbitrary cycle, such as calendar weeks. We see no evidence of this (results available on request).

Finally, a different psychological explanation could be that the targets are not merely monetary motivators. Targets may also generate intrinsic motivation: the desire to hit the target may increase effort (Amabile and Kramer 2011). With data such as ours, of course, one cannot separate intrinsic from extrinsic motivation generated by the target. However, without time inconsistency, it is unclear how this would explain the payday findings, the correlation between the payday and contract effects, or the higher take-up of targets in the evening versus the morning when uncertainty is low. As a result, while our data cannot rule out nonmonetary motivations, they do suggest that time inconsistency is needed in this case as well.

VI. Conclusion

We find that many workers are present biased, that this substantially affects their effort, and that they are sophisticated enough about this present bias to choose dominated contracts in an experiment. Output increases over the pay cycle imply a daily discount rate of 4 percent. Workers with above-average payday effects choose dominated contracts 43 percent of the time, and being offered the option to choose dominated contracts increases earnings by 9 percent. For other workers, payday effects and demand for dominated contracts are smaller, pointing to the importance of heterogeneity in self-control problems.

In a companion paper (Kaur et al. 2014), we derive results on equilibrium contracts and job design when at least some workers are subject to self-control problems and there is free entry of firms.³⁰ We consider a standard agency model in which observable output is a stochastic function of unobservable worker effort and workers are risk averse. The standard result in such agency models is that equilibrium labor market contracts partially insure workers at the expense of some reduction in the steepness of worker incentives, effort, and output relative to self employment. We show that present bias among workers will lead firms to offer higher-powered incentives and that sufficiently strong present bias reverses the standard partial insurance result.³¹ With sufficiently strong present bias, the distribution of output second-order stochastically dominates the distribution of wages, and hence present-biased workers exert more effort and produce more output as employees than as self-employed owner-operators.

We also find implications for the organization of production and job design. The steep incentives needed to motivate unobservable effort by present-biased workers will impose risk on workers. As a result, to the extent that workers are risk averse, present bias will make firms and workers more willing to expend resources on adopting technologies and designing jobs so as to make effort observable; this enables contracting on effort rather than stochastic production. (For example, firms may require employees to work fixed hours in a factory or office rather than allowing them to telecommute or choose their work hours.)

If workers are heterogeneous in their (unobservable) time preferences, then firms will face an adverse selection problem. In Kaur et al. (2014), we show that, in general, time-consistent workers will be made worse off by the presence of present-biased workers. Depending on parameter values, they may have to accept contracts with either (1) higher-powered incentives and thus more risk, (2) costly effort monitoring, or (3) pooling with present-biased workers and thus lower expected wages than in the absence of such workers

In contrast to other models of equilibrium interaction between present-biased and time-consistent agents—in which present-biased agents are naive and hence are exploited in equilibrium by others who can better predict their behavior (DellaVigna and Malmendier 2004; Eliaz and Spiegel 2006; Gabaix and Laibson 2006)—the presence of present-biased agents makes time-consistent agents worse off.

³⁰ Much of this was in Sec. VI of an earlier draft of this paper.

³¹ Free entry of firms implies that under all contracts, firms will make zero expected profits, with expected wages equal to expected production. When firms are competing to hire workers, workers will wind up being compensated for earning less when output realizations are low by being better compensated when output is high.

In our end line survey, many employees expressed a desire for rules to help them work harder under pure piece rates: 78 percent of workers agreed with “Some days I don’t work as hard as I would like to” and 87 percent agreed with “I wish I had better attendance at work” (see table 1). Some—but not all—expressed demand for workplace rules to increase effort. For example, 70 percent agreed with “It would be good if there were rules against being absent because it would help me come to work more often” while 24 percent disagreed.

Our finding of self-control problems among at least a subset of workers, together with the results on job design and effort monitoring above, may help shed light on the role of factory discipline in the Industrial Revolution and on some contemporary debates about human resource practices. Prior to the Industrial Revolution, textiles were often produced in a cottage industry system in which self-employed producers worked in their homes. This evolved into the “putting-out” system, under which workers rented space on factory floors, were free to choose their output and work hours, and sold their output. During the Industrial Revolution, this system was replaced by factory discipline, in which even piece rate workers were subject to dismissal or heavy fines for minor deviations such as stepping away from their machine, eating, talking, whistling, or looking out the window (Clark 1994).

Many historians and some economists (e.g., Thompson 1967; Marglin 1974, 2008) have argued that the introduction of the new management technology of factory discipline was as important to the Industrial Revolution as any purely technological innovation. They tend to see factory discipline as imposed on workers by capitalists, and perhaps as made possible only by the dispossession of farmers by enclosure.

Clark (1994) turns this interpretation on its head, with a much more benign view of the role of factory discipline. He notes that under the putting-out system, workers “frequently kept irregular hours, often taking off Monday (‘St. Monday’) and even Tuesday and working long hours on Thursday and Friday” (151). Clark posits that workers valued the constraints imposed on their behavior by factory discipline because this helped mitigate their self-control problems. To underscore this point, he highlights that even workers paid piece rates were often under factory discipline. For example, among 32 linen mills in Belfast in the 1890s, 29 imposed fines for minor tardiness and 21 locked out pieceworkers who were a few minutes late, causing not just workers but also the firm to lose output for the entire day. In this view, the emergence of factory discipline could be seen as reflecting optimal contracts among workers and employers, and there would be no need for labor market regulations imposing limits on the workweek, for example. Similarly, under this view, monitoring of workers—a hotly debated contemporary issue,

as evidenced by recent debates about telecommuting and keystroke monitoring technologies—may have both welfare and output benefits.

Our results suggest a way to encompass the sharply divergent perspectives on these issues. We find strong heterogeneity in the extent of time inconsistency and in the earnings benefits of imposing dominated contracts. As discussed in Kaur et al. (2014), this can create an adverse selection problem, and there is no presumption that equilibrium contracts or job design will be socially efficient. Regulations on the workweek or on the acceptable level of effort monitoring could potentially be welfare improving.

Agency theory traditionally understands workplace arrangements—the existence of bosses and worker discipline—in one of two ways. The first view is that the firm exists to provide insurance. This insurance creates moral hazard. Workplace arrangements exist to mitigate that moral hazard. The second view—that joint production necessitates the need for monitoring (Alchian and Demsetz 1972)—is summarized in a story by Steven Cheung (1983, 8): “On a boat trip up China’s Yangtze River in the 19th Century, a titled English woman complained to her host of the cruelty to the oarsmen. One burly coolie stood over the rowers with a whip, making sure there were no laggards. Her host explained that the boat was jointly owned by the oarsmen, and that they hired the man responsible for flogging.”

The discussion above suggests a potentially different way to understand a diverse host of workplace arrangements. Discipline at the workplace—such as the coolie in Cheung’s story—may reflect demand for arrangements to help avoid the temptation to shirk. Do job features such as assembly lines, production minimums, rigid work hours, and hefty punishments for even small lapses in behavior such as tardiness have self-control benefits? Might this help explain why the movement from farm to factory work has typically been accompanied by increases in labor productivity (see Kaur et al. 2010)? Are workplace incentive contracts—which often embody high-powered incentives in some form—at least partially structured to provide self-control benefits? Could the organization of production itself serve to mitigate self-control problems?

Time inconsistency also has implications for how we conceptualize the production function. For example, in subsistence agriculture, the motivation problem may be larger for crops with longer planting cycles. Indeed, the move from agriculture to formal sector work with regular pay—a key component of the historical development process—could have productivity benefits partly due to the effect on self-control (see Kaur et al. 2010).

These possibilities are, of course, speculative. However, given that we find strong evidence that self-control problems distort worker effort at economically meaningful magnitudes, a closer exploration of these possibilities is warranted in future research.

References

- Alchian, Armen A., and Harold Demsetz. 1972. "Production, Information Costs, and Economic Organization." *A.E.R.* 62 (December): 777–95.
- Amabile, Teresa, and Steven Kramer. 2011. "The Power of Small Wins." *Harvard Bus. Rev.* 89 (May): 70–80.
- Ariely, Dan, and Klaus Wertenbroch. 2002. "Procrastination, Deadlines, and Performance: Self-Control by Precommitment." *Psychological Sci.* 13 (3): 219–24.
- Ashraf, Nava, Dean Karlan, and Wesley Yin. 2006. "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines." *Q.J.E.* 121 (2): 635–72.
- Augenblick, Ned, Muriel Niederle, and Charles Sprenger. 2014. "Working over Time: Dynamic Inconsistency in Real Effort Tasks." Working Paper no. 18734, NBER, Cambridge, MA.
- Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein. 2012. "Behavioral Hazard in Health Insurance." Working Paper no. 18468, NBER, Cambridge, MA.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2007. "Incentives for Managers and Inequality among Workers: Evidence from a Firm Level Experiment." *Q.J.E.* 122:729–75.
- Banerjee, Abhijit, and Sendhil Mullainathan. 2010. "The Shape of Temptation: Implications for the Economic Lives of the Poor." Working Paper no. 15973, NBER, Cambridge, MA.
- Bernheim, B. Douglas, Debraj Ray, and Sevin Yeltekin. 2013. "Poverty and Self-Control." Working Paper no. 18742, NBER, Cambridge, MA.
- Burger, Nicholas, Gary Charness, and John Lynham. 2009. "Three Field Experiments on Procrastination and Willpower." Manuscript, Univ. California, Los Angeles. http://www.bec.ucla.edu/Charness_Procrastination.pdf.
- Cadena, Ximena, Antionette Schoar, Alexandra Cristea, and Héber M. Delgado-Medrano. 2012. "Fighting Procrastination in the Workplace: An Experiment." Working Paper no. 16944, NBER, Cambridge, MA.
- Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. 1997. "Labor Supply of New York City Cabdrivers: One Day at a Time." *Q.J.E.* 112: 407–41.
- Chabris, Christopher, David Laibson, Carrie Morris, Jonathan Schuldt, and Dmitry Taubinsky. 2008. "Individual Laboratory-Measured Discount Rates Predict Field Behavior." *J. Risk and Uncertainty* 37 (2–3): 237–69.
- Chabris, Christopher, David Laibson, and Jonathan Schuldt. 2008. "Intertemporal Choice." In *The New Palgrave Dictionary of Economics*, 2nd ed., edited by Steven N. Durlauf and Lawrence E. Blume. New York: Macmillan.
- Cheung, Steven N. S. 1969. "Transaction Cost, Risk Aversion, and the Choice of Contractual Arrangements." *J. Law and Econ.* 12 (1): 23–42.
- . 1983. "The Contractual Nature of the Firm." *J. Law and Econ.* 26 (April): 1–21.
- Clark, Gregory. 1994. "Factory Discipline." *J. Econ. Hist.* 54:128–63.
- DellaVigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *J. Econ. Literature* 47:315–72.
- DellaVigna, Stefano, and Ulrike Malmendier. 2004. "Contract Design and Self-Control: Theory and Evidence." *Q.J.E.* 119:353–402.
- Dupas, Pascaline, and Jonathan Robinson. 2013. "Daily Needs, Income Targets and Labor Supply: Evidence from Kenya." Working Paper no. 19264, NBER, Cambridge, MA.

- Eliaz, Kfir, and Ran Spiegler. 2006. "Contracting with Diversely Naive Agents." *Rev. Econ. Studies* 73:689–714.
- Engen, Eric M., William G. Gale, and John Karl Scholz. 1994. "Do Saving Incentives Work?" *Brookings Papers Econ. Activity* (1): 85–180.
- Engen, Eric M., William G. Gale, and Cori Uccello. 1999. "The Adequacy of Household Saving." *Brookings Papers Econ. Activity* (2): 65–187.
- Fang, Hanming, and Dan Silverman. 2009. "Time Inconsistency and Welfare Program Participation: Evidence from the NLSY." *Internat. Econ. Rev.* 50 (4): 1043–77.
- Fehr, Ernst, and Lorenz Goette. 2007. "Do Workers Work More If Wages Are High? Evidence from a Randomized Field Experiment." *A.E.R.* 97 (1): 298–317.
- Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." *J. Econ. Literature* 40 (2): 351–401.
- Fudenberg, Drew, and David Levine. 2006. "A Dual Self Model of Impulse Control." *A.E.R.* 96:1449–76.
- Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Q.J.E.* 121 (2): 505–40.
- Giné, Xavier, Dean Karlan, and Jonathan Zinman. 2010. "Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation." *American Econ. J.: Appl. Econ.* 2 (4): 213–35.
- Gneezy, Uri, and John List. 2006. "Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets Using Field Experiments." *Econometrica* 74 (5): 1365–84.
- Grossman, Sanford J., and Oliver D. Hart. 1983. "An Analysis of the Principal-Agent Problem." *Econometrica* 51 (January): 7–45.
- Gul, Faruk, and Wolfgang Pesendorfer. 2001. "Temptation and Self-Control." *Econometrica* 69 (6): 1403–35.
- . 2004. "Self Control, Revealed Preferences and Consumption Choice." *Rev. Econ. Dynamics* 7 (2): 243–64.
- Holmstrom, Bengt. 1979. "Moral Hazard and Observability." *Bell J. Econ.* 10 (Spring): 74–91.
- Hossain, Tanjim, and John List. 2012. "The Behavioralist Visits the Factory: Increasing Productivity Using Simple Framing Manipulations." *Management Sci.* 58 (12): 2151–67.
- Hubbard, Glenn, Jonathan Skinner, and Stephen Zeldes. 1994. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." *Carnegie-Rochester Conf. Ser. Public Policy* 40:59–125.
- Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan. 2010. "Self-Control and the Development of Work Arrangements." *A.E.R.* 100 (2): 624–28.
- . 2014. "Equilibrium Contracts and Job Design When Workers Are Present-Biased." Manuscript, Harvard Univ. <http://scholar.harvard.edu/kremer/>.
- Kirby, Kris, and Nino Marakovic. 1995. "Modeling Myopic Decisions: Evidence for Hyperbolic Delay-Discounting with Subjects and Amounts." *Org. Behavior and Human Decision Processes* 64:22–30.
- Laibson, David I. 1996. "Hyperbolic Discount Functions, Undersaving, and Savings Policy." Working Paper no. 5635, NBER, Cambridge, MA.
- . 1997. "Golden Eggs and Hyperbolic Discounting." *Q.J.E.* 112 (2): 443–77.
- Larkin, Ian. 2014. "The Cost of High-Powered Incentives: Employee Gaming in Enterprise Software Sales." *J. Labor Econ.* 32 (2): 199–227.

- Levitt, Steven, and John List. 2007. "What Do Laboratory Experiments Measuring Social Preferences Reveal about the Real World?" *J. Econ. Perspectives* 21: 153–74.
- . 2009. "Field Experiments in Economics: The Past, the Present, and the Future." *European Econ. Rev.* 53 (1): 1–18.
- List, John. 2006. "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions." *J.P.E.* 114 (1): 1–37.
- List, John, and Imran Rasul. 2011. "Field Experiments in Labor Economics." In *Handbook of Labor Economics*, vol. 4A, edited by Orley Ashenfelter and David Card, 103–28. Amsterdam: Elsevier.
- Loewenstein, George, and Drazen Prelec. 1992. "Anomalies in Intertemporal Choice: Evidence and an Interpretation." *Q.J.E.* 107 (May): 573–97.
- Marglin, Stephen Alan. 1974. "What Do Bosses Do? The Origins and Functions of Hierarchy in Capitalist Production." *Rev. Radical Polit. Econ.* 6 (2): 60–112.
- . 2008. *The Dismal Science: How Thinking Like an Economist Undermines Community*. Cambridge, MA: Harvard Univ. Press.
- O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing It Now or Later." *A.E.R.* 89 (1): 103–24.
- . 2001. "Choice and Procrastination." *Q.J.E.* 116 (1): 121–60.
- . 2006. "Incentives and Self-Control." In *Advances in Economics and Econometrics: Theory and Applications; Ninth World Congress*, vol. 2. Econometric Society Monographs, vol. 42. New York: Cambridge Univ. Press.
- Oyer, Paul. 1998. "Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality." *Q.J.E.* 113 (1): 149–85.
- . 2000. "Theory of Sales Quotas with Limited Liability and Rent Sharing." *J. Labor Econ.* 18 (3): 405–26.
- Paserman, Daniele. 2008. "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation." *Econ. J.* 118 (531): 1418–52.
- Shapiro, Jesse. 2005. "Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle." *J. Public Econ.* 89:303–25.
- Shearer, Bruce. 2004. "Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment." *Rev. Econ. Studies* 71 (2): 513–34.
- Shui, Haiyan, and Lawrence M. Ausubel. 2005. "Time Inconsistency in the Credit Card Market." Working paper, Univ. Maryland. <http://www.ausubel.com/creditcard-papers/time-inconsistency-credit-card-market.pdf>.
- Spiegler, Ran. 2011. *Bounded Rationality and Industrial Organization*. Oxford: Oxford Univ. Press.
- Thompson, E. P. 1967. "Time, Work-Discipline, and Industrial Capitalism." *Past and Present*, no. 38 (December 1): 56–97.