Self-Control at Work

Supreet Kaur

Michael Kremer

Sendhil Mullainathan*

Abstract

Workers with self-control problems do not work as hard as they would like. This changes the logic of agency theory by partly aligning the interests of the firm and worker: both now value contracts that elicit more effort in the future. Three findings from a year-long field experiment with data entry workers suggest the quantitative importance of self control at work. First, workers choose dominated contracts—which penalize low output but provide no greater reward for high output—36% of the time to motivate their future selves; use of these contracts increases output by the same amount as an 18% increase in the piece-rate. Second, effort increases as the (randomly assigned) payday gets closer: output rises 8% over the pay week; calibrations show that justifying this would require a 4% daily exponential discount rate. Third, for both findings there is significant and correlated heterogeneity: workers with larger payday effects are both more likely to choose dominated contracts and show greater output increases under them. This correlation grows with experience, consistent with the hypothesis that workers learn about their self-control problems over time. Self-control problems among workers could potentially lead firms to either adopt high-powered incentives or impose work rules to allow monitoring of worker effort.

---

* Supreet Kaur: Department of Economics, Columbia University (kaur@columbia.edu); Michael Kremer: Department of Economics, Harvard University and National Bureau of Economic Research (mkremer@fas.harvard.edu); Sendhil Mullainathan: Department of Economics and National Bureau of Economic Research, Harvard University (mullain@fas.harvard.edu).

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 (IFMR) for operational 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.

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 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 than self-control in other contexts, such as savings or smoking (Laibson 1997, Ashraf et al. 2006, Giné et al. 2010). In many of these other contexts, the market provides commitment only if the consumer demands it sufficiently 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

---

[1] Frederick et al (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). Bernheim, Ray, and Yeltekin (2011) and Banerjee and Mullainathan (2010) 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

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]

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 based on the

---

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.

[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 et al. (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 non-student 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. 2010) and bicycle taxi drivers (Dupas and Robinson 2014).

[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 et al. 2007, Fehr and Goette 2008, and Hossain and List 2012). For excellent reviews of this literature, see Levitt and List (2009) and List and Rasul (2011).

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% of the time. The option to choose a dominated contract increases production and earnings, with a Treatment on the Treated Effect of 6% for those workers who accept the dominated contract. A production increase of this size corresponds to that induced by an 18% increase in the piece rate wage (computed using exogenous wage changes). We show 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% 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 payment varied. Worker output is 8% higher on paydays than at the beginning of the weekly pay cycle. An effect of this magnitude corresponds to a 24% 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% per day or 1.65×10^6 % 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% more likely to choose dominated contracts. Providing these workers the *option* to choose a dominated contract increases their output by 9%, implying a Treatment on the Treated Effect on output of

28% for those workers who select the contract. This implies that for these workers, dominated contracts have production impacts comparable to an 85% 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 are similar. *Ex post* analysis suggests a possible reason: workers face output uncertainty—e.g. 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%) 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 employers giving employees such choices is necessarily the optimal contract for time inconsistent workers. Our results indicate that

[8] Although we find 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 (2013), who find present bias in effort tasks but fail to find it for cash discounting tasks among student subjects.

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 US (Oyer 2000, Larkin 2013) or data entry in India. In addition, most jobs have production minimums—such as a forty-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, like 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.

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 concludes.

While this paper tests for self-control problems, a companion paper, Kaur et al. (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 agents may exert more effort and produce more output working in firms than as self-employed owner-operators. To the extent 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 pre-specified 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 naïve and hence can be exploited by sophisticated time-consistent agents (DellaVigna and Malmendier 2004, Gabaix and Laibson 2006, Elias and Spiegler 2006)—the presence of present-biased agents makes time-consistent agents worse off.

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 seven days. In what follows, we generate predictions on how our experimental set-up distinguishes between time consistent and time-inconsistent workers, and enables us to calibrate the extent of time inconsistency.

Assume 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 where $c(e_t) = {e_t}^\theta$, 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, twice differentiable, with $\lim_{e_t \to \infty} c'(e_t) = \infty$. For simplicity, we will also focus on the case where $\alpha^i = 1$ 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*, $\frac{D^I(t+s)}{D^I(t)}$ is strictly increasing in t.[9] The time-inconsistent worker is at least as impatient as the time-consistent one: $\frac{D^I(t+1)}{D^I(t)} \leq \delta$ for all t. We additionally assume 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 $D(t) = (1 + \mu t)^{-\gamma/\mu}$ (Lowenstein and Prelec 1992) satisfies this property. Note that for a quasi-hyperbolic function (see Laibson 1997), $D^i(t) = \beta\delta^t$, $\frac{D^i(1)}{D^i(0)} = \beta\delta$ and $\frac{D^i(t+1)}{D^i(t)} = \delta$ 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 where 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 both in the model and the experiment workers (i) have a distinct contract each day, and (ii) are paid based on the sum of these contracts at the end of the payweek 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 $e^i_{t|s}$ 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, $e^i_{t|t-k}$ is pinned down by the FOC from expression (1):[10]

$$c'\left(e^i_{t|t-k}\right) = \frac{D^i(T-t+k)}{D^i(k)} b, \tag{3}$$

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

$$c'\left(e^i_{t|t}\right) = \frac{D^i(T-t)}{D^i(0)} b = D^i(T-t)b. \tag{4}$$

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

[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 this condition is satisfied for all workers. Under this condition, the first order conditions will be both necessary and sufficient for a global maximum.

Lemma 1 (Equivalence of Discounting and Piece Rate Changes)*:*
Let $e^i_{t+1|t+1}$ *denote the effort level chosen in period* $t+1$ *under piece rate* b*. The piece rate* b' *needed to generate effort* $e^i_{t+1|t+1}$ *in period* t *is given by:*

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

Proof: Define b' as the piece rate that the worker must be paid in period t to elicit effort equivalent to $\tilde{e} \equiv e^i_{t+1|t+1}$, 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 alternate piece rate b' is: $D^i(T-t)b' = c'(\tilde{e})$. This implies $\frac{D^i(T-t-1)}{D^i(T-t)} = \frac{b'}{b}$. ■

Intuitively, $\frac{D^i(T-t-1)}{D^i(T-t)}$ 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 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 $\frac{1}{\delta} - 1$ *if workers are exponential discounters, and is greater than* $\frac{1}{\delta} - 1$ *if they are time inconsistent. This implies that output increases over the pay cycle will be larger for time inconsistent workers than time consistent ones.*

Proof: Since $c(e) = e^{\theta}$, ε equals $\frac{1}{\theta-1}$ and is constant over effort. Thus, $\varepsilon = \frac{1}{\theta-1} = \left(\frac{e^i_{t+1|t+1} - e^i_{t|t}}{e^i_{t|t}}\right) \Big/ \left(\frac{b'-b}{b}\right)$. Substituting in for $\frac{b'}{b}$ from Lemma 1 gives the relationship between the output increase over the paycycle 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, $\frac{\partial e^i_{t|t}}{\partial b}\frac{b}{e^i_{t|t}} \approx \left(\frac{e^i_{t+1|t+1}-e^i_{t|t}}{e^i_{t|t}}\right) \Big/ \left(\frac{b'-b}{b}\right)$, 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 elasticity in period t, implying that the output increase from period t to $t+1$ will be larger for time inconsistent workers. ■

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 payweek, 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 first order conditions 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 $t-k$: $e^i_{t|t-k} = e^i_{t|t}$ for time consistent workers. This is because 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 $\frac{D^I(T-t+k)}{D^I(k)} > \frac{D^I(T-t)}{D^I(0)}$, $e^i_{t|t-k} > e^i_{t|t}$ 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:*

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

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 $X_{t|t-k} \in (e^i_{t|t}, e^i_{t|t-k}]$. In contrast, time consistent workers will never strictly prefer $v_{t|t-k}(\cdot)$.

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

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

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, $e^i_{t|t-k} = e^i_{t|t}$ 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 $x^i_{t|t-k} = 0$.

In contrast, time inconsistent workers' period $t-k$ selves will prefer $e^i_{t|t-k}$ while their period t selves will prefer $e^i_{t|t}$, so there is potentially scope for the period $t-k$ self to influence the period t self's choice of effort. To characterize the optimal $X_{t|t-k}$, we solve backwards, starting with the period t self's problem given the $X_{t|t-k}$ chosen by the period $t-k$ self. The period t self will solve for the utility maximizing level of effort in the range $[X_{t|t-k}, \infty)$ and the utility maximizing effort in the range $(0, X_{t|t-k})$ and choose whichever yields greater utility. By the convexity of $c(\cdot)$, the period $t-k$ self will never choose $X_{t|t-k} < e^i_{t|t}$; given such an $X_{t|t-k}$, the period t self will never choose effort greater than $X_{t|t-k}$.

Denote as $e_t^{i(p)}$ the optimal level of effort from the perspective of period t when the piece rate is b^p: $e_t^{i(p)} \equiv \text{argmax}_e\{D^i(T-t)(a+b^p e) - c(e)\}$. Note that $e_t^{i(p)} < e^i_{t|t}$ since $b^p < b$. In general, the period t self will choose $X_{t|t-k}$ over $e_t^{i(p)}$ if and only if $D^i(T-t)\big(a + bX_{t|t-k}\big) - c\big(X_{t|t-k}\big) \geq D^i(T-t)\Big(a + b^p e_t^{i(p)}\Big) - c\Big(e_t^{i(p)}\Big)$.

There is always some $X_{t|t-k}$ that will make the time $t-k$ self strictly prefer $v_{t|t-k}(\cdot)$ to the $w(\cdot)$ schedule. To see this, note that since the period $t-k$ self's FOC is satisfied with equality at $e^i_{t|t-k}$, $bD^i(T-t+k)/D^i(k) - c'\big(e^i_{t|t-k}\big)$ will be positive and decreasing in e for $e < e^i_{t|t-k}$. Thus, by the convexity of $c(\cdot)$, higher levels of effort are preferred for $e < e^i_{t|t-k}$. Consequently, the period $t-k$ self will prefer to induce as high an effort level as possible in the range $e \leq e^i_{t|t-k}$.

There is always some $X_{t|t-k} > e^i_{t|t}$ that the period t self can be induced to supply under $v_{t|t-k}(\cdot)$. To see this, write $X_{t|t-k} = e^i_{t|t} + \mu$. As $\mu \to 0$, the time t self will be arbitrarily close to indifferent between exerting $e^i_{t|t} + \mu$ under $X_{t|t-k}(\cdot)$ and exerting $e^i_{t|t}$ 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 any other value of e.

Note that increasing $X_{t|t-k}$ beyond $e^i_{t|t-k}$ will reduce utility for period $t-k$ and by the convexity of $c(\cdot)$ will reduce utility from the standpoint of period t, making it harder to induce to induce $X_{t|t-k}$. Thus for time inconsistent workers, some $X_{t|t-k}$ in the range $e^i_{t|t} < X_{t|t-k} \leq e^i_{t|t-k}$ will maximize period $t-k$'s utility, and $v_{t|t-k}(\cdot)$ will be strictly preferred to $w(\cdot)$.

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: $D^I(T-t)\Big[be_t^{i(\max)} - b^p e_t^{i(p)}\Big] - \Big[c\Big(e_t^{i(\max)}\Big) - c\Big(e_t^{i(p)}\Big)\Big] = 0$. Then, $e^i_{t|t-k}$ can be induced in period t if $e^i_{t|t-k} \leq e_t^{i(\max)}$. Thus, we have $X_{t|t-k} = \min\Big\{e^i_{t|t-k}, e_t^{i(max)}\Big\}$.

Following the logic in the Proposition 1 proof, we define b'' as the alternate 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: $D^I(T-t)b'' - c'\big(X_{t|t-k}\big) = 0$. The FOC in equation (3) implies: $D^I(T-t+k)b - D^I(k)c'\big(X_{t|t-k}\big) \geq 0$ since $X_{t|t-k} \leq e^i_{t|t-k}$. 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 $\varepsilon = \frac{1}{\theta-1} = \left(\frac{e^i_{t|t-k} - e^i_{t|t}}{e^i_{t|t}}\right)\Big/\left(\frac{b''-b}{b}\right)$ gives expression (7): $\frac{D^I(T-t+k)/D^I(k)}{D^I(T-t)} - 1 = \frac{D^I(T-t+k)/D^I(k)}{D^I(T-t)/D^I(0)} - 1 \geq \frac{x^i_{t|t-k}}{\varepsilon}$. For a more general $c(\cdot)$, $\frac{\partial e^i_{t|t}}{\partial b}\frac{b}{e^i_{t|t}} \approx \left(\frac{e^i_{t|t-k} - e^i_{t|t}}{e^i_{t|t}}\right)\Big/\left(\frac{b''-b}{b}\right)$and expression (7) will hold as a first order approximation. ■

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, $e^i_{t|t-k} = e^i_{t|t}$ 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 $e^i_{t|t}$, their period $t-k$ selves will prefer $e^i_{t|t-k}$ and will want to induce as high an effort level as possible in the range $e \leq e^i_{t|t-k}$ (see 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: $\frac{\delta^{T-t+k}/\delta^k}{\delta^{T-t}/\delta^0} = 1.$ 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, $x^i_{t|t-k}$, 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 $e^i_{t|t-k}$: $X_{t|t-k} \leq e^i_{t|t-k}$.

Proposition 3 (Correlation of Pay Cycle and Dominated Contract Effects):
Suppose 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, $e^i_{t+1|t+1}/e^i_{t|t}$, 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. ■

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 vs. benefits of effort in period t. In practice, this correlation will also be weakened if some time inconsistent workers are naïve—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 4-5 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 Paycycle):
For time inconsistent workers, if $e^i_{t|t-k}$ can be induced in period t using the dominated contract, then $x^i_{t|t-k} > x^i_{t+1|t+1-k}$. In contrast, for time consistent workers, $x^i_{t|t-k} = x^i_{t+1|t+1-k} = 0$.

Proof: For time inconsistent workers, $x^i_{t|t-k} = \varepsilon\left[\frac{D^i(T-t+k)/D^i(k)}{D^i(T-t)/D^i(0)} - 1\right] > \varepsilon\left[\frac{D^i(T-t-1+k)/D^i(k)}{D^i(T-t-1)/D^i(0)} - 1\right] \geq x^i_{t+1|t+1-k}$, where: the equality comes from the assumption that $e^i_{t|t-k}$ can be induced in period t (see Proposition 2 proof); the strict inequality is due to the hyperbolicity assumption that $\frac{D^I(t+s)}{D^I(t)}$ is increasing in t; and the final weak inequality comes from the definition of $x^i_{t+1|t+1-k}$ using Proposition 2.

In contrast, for time consistent workers $x^i_{s|s-k} = 0$ in all periods, so trivially, there will be no change in dominated contract effects over the paycycle. ■

Proposition 4 holds when the period t self can be induced to exert $e^i_{t|t-k}$. 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. This is because $e^i_{t|t-k}$ becomes closer to $e^i_{t|t}$ 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 $e^i_{t|t} < e^i_{t|t-k_1} < e^i_{t|t-k_2}$. 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 $X_{t|t-k_2} = X_{t|t-k_1}$ and induce the same effort level as the period $t-k_1$ self. ■

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 if we assume there are 3 sub-periods 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 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 $e = X_{t|t-k}$. Thus time-consistent workers would *strictly* prefer $w(\cdot)$ to $v_{t|t-k}(\cdot)$. In contrast, time-inconsistent workers could still prefer a dominated wage schedule—for small enough p, they would prefer $v_{t|t-k}(\cdot)$ 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, $1 + x(1 - \theta)$ is a less tight lower bound on the extent to which the worker deviates from exponential discounting.

Shocks to the cost of effort, e.g. 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 zero to reject even dominated contracts with very small $X_{t|t-k}$.

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 due to time inconsistency, dominated contracts will be more appealing further in advance of the time of effort. The fact that uncertainty may vary over the paycycle 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 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

III.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 Appendix Figure 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 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

[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 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.

Section II, $w(e) = Rs. 15 + Rs. 0.03e$. 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 eighteen 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 email address prior to joining the firm.[12] Many employees commuted from surrounding villages using buses and trains, with some traveling up to two hours in each direction.

New recruits received about two weeks of training before contract randomizations began (see below). During the first four days, they were paid a flat stipend while receiving instruction on the data entry software and production task. During the next four 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 two 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%. Throughout the experiment, workers were randomly assigned

[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.

to seats in the office and these assignments changed every one to three weeks, since some computers were slower and more sensitive to network speed fluctuations than others.

III.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 seven 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 pre-weekend 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 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 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 could increase output

[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 percentiles, 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

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 5 workers in the sample over a 24-day period.

Figure 2 provides an overview of the experiment timeline and design. Our sample of 102 workers and 8,423 observations covers the eight-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

IV.A. Pay Cycle Effects on Production (Test 1)

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

---

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.

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, Col. 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 pattern, as shown in Col. 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%. This magnitude corresponds to approximately one 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 day before the payday. If we assume 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 (Appendix 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 only arise on the payday itself. Instead,

---

[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 Figure 2.

[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%—per day leading up to the payday (significant at the 1% level).

[18] Moreover, workers may have wanted to leave work early on paydays to make purchases—for example, 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.

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 the payday off.

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% in the week prior to major festivals (Appendix Table 4).

Calibration of Implied Discount Rate

We can use Proposition 1 to calibrate the discount rate implied by the weekly production cyclicality. 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 5 times over a 10-day period in random order. This 33% increase in wages increased output by 11%, for an elasticity of 0.33 (Table 2, Col. 6). Note that under the assumption that $c(e) = e^{\theta}$, with elasticity equal to $\frac{1}{\theta-1}$, this implies $\theta = 4$.

Since the average output increase over the six days from the beginning to the end of the payweek is 8% (Table 2, Col. 3), on average, $\frac{e^i_{t+1|t+1} - e^i_{t|t}}{e^i_{t|t}} = \frac{0.08}{6} = 0.013$. 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%. If workers were time consistent exponential discounters, this would require an annual discount rate of 1.65x10^6%. Standard estimates for the exponential discount rate in the literature are about 5% per year (e.g. Engen et al. 1994, 1999;

Hubbard et al. 1994)—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 inter-temporally 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 which 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% per day, it seems hard to imagine that one could fail to reject the hypothesis of exponential discounting at rates of about 5% 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% per day over short horizons. Calibrations from field data also produce such large estimates (e.g. Paserman 2004; Fang and Silverman 2004; Shui and Ausubel 2004).

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,

Figure 1).[20] Factory workers in pre-Industrial Revolution England exhibited similar dynamics over their pay cycle (Clark 1994).[21]

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.

IV.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% of the time when present (Table 3, Panel A). 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 able to select targets). As a conservative estimate, if we code workers who are absent the day before or day of assignment as choosing zero targets, the mean take-up rate across workers is 28%.[22]

[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 U.S. declines by 10-15% over the food stamp month.
[22] Under the null hypothesis of time consistency, take-up of dominated contracts should be zero. We view 36% as a sizeable 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 impact production and earnings.

Figure 4 plots the distribution of worker take-up rates. 16% of workers always chose a target of zero. The bottom quarter of the distribution choses positive targets less than 10% of the time. The top quarter choses positive targets at least 60% 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 Section IV.E). In addition, workers might go through periods of present-biasedness, for example due to 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 devices that reduce the need for dominated contracts could also vary across days (e.g., see Kaur, Kremer, and Mullainathan 2010).

Panel B of Table 3 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% (significant at the 5% level; Col. 1). Looking within these treatments, *evening option to choose a dominated contract* increased output by 3% (significant at the 5% level) and *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% 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 (Appendix Table 3, Panel B).

The implied Treatment on the Treated Effect of choosing a positive target on output is approximately 6%. Given the estimated elasticity of 0.33, the magnitude of this effect corresponds to an 18% increase in the piece rate. Using Proposition 2, we can back out the implied bound on the departure from time consistency. The TOT effect of 6% 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 the payday, relative to the costs of effort on the workday, by at least 18% 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% of the time if they had been assigned to the *control* contract that day (Table 4, 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% of the time (Table 4, 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% (Table 4, Panel B, Col. 1). This helps explain why the *target imposed* treatments did not significantly raise earnings (Table 3).

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 3 and 4 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 in one day increases the cost of effort in subsequent days (Appendix Table 5). Specifically, assignment to *option to choose a dominated contract* or *target imposed* (relative to the *control* contract) does not reduce production the next day.[23]

[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 3. In Appendix 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-

IV.C. Heterogeneity in Preferences: Correlation Between Payday & Contract Effects (Tests 3-4)

The payday and contract choice results each suggest that at least some workers are time inconsistent. For example, 41% of workers have a 10% or larger pay cycle effect. Similarly, after 2 months of experience, 49% of workers select dominated contracts 25% 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 0.000. 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, and standard controls. The p-value of the F-test of joint significance of the interaction coefficients is 0.003.

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:

$$Payday\ effect = \frac{(Mean\ production\ on\ paydays) - (Mean\ production\ on\ nonpaydays)}{Mean\ production\ in\ sample}.$$

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

test of joint significance of the probability covariates has a p-value of 0.45—indicating that the assignment probabilities have little predictive power. Their inclusion also has little impact on the estimated treatment effects.

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

quasi-hyperbolic models is that of output increases on paydays. Even in hyperbolic models, increases closer to the 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 non-paydays. In ex post analysis, we have confirmed 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 5, Panel A). These coefficients correspond to a striking 47% and 49% of the mean take-up rate and target level, respectively, and are both significant at the 1% level. The payday effects heterogeneity is not driven by other potentially correlated interpersonal differences among workers, such as productivity (Table 5, Panel A, Col. 3). Workers with large payday effects also increase production more in response to dominated contracts. In Table 5, Panel B, Col. 1, the interaction between the *option to choose a dominated contract* and the high payday effect dummies is 482 fields—9% of mean production (significant at the 1% level), implying a Treatment on the Treated Effect of 28%. The effect on earnings is of the same magnitude (Table 5, Panel B, Col. 2) and is shown graphically in Figure 5. High payday effect workers are also more likely to show up to work when assigned to *option to choose a dominated contract* or *target imposed* (Table 5, Panel B, Col. 3).[25] This provides additional evidence that high payday effect workers demand dominated contracts.

[25] The large production effect of *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 5581 fields. As a simple calibration, 5581*0.044 = 245 < 395. Moreover, regressing production conditional on attendance on the contract treatment dummies yields positive and significant coefficients (Appendix Table 3), although these are difficult to interpret since attendance is endogenous.

Given our labor supply elasticity estimate of 0.33, the 9% 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% increase in the piece rate wage. This magnitude corresponds to a one-year increase in education. Using Proposition 2, the TOT effects allows us to bound the level of time inconsistency of the workers which, based on their paycycle behavior, appear most time inconsistent. When these workers choose dominated contracts, $\frac{d^I(T)/d^I(t)}{d^I(T-t)/d^I(0)} - 1 \geq \frac{0.28}{0.33} = 0.84$; this implies that the relative value of the wage benefits to effort costs is 84% higher at the time of contract choice than at the time of effort.[26]

The workers most affected by paydays also select more aggressive targets (Table 4, Cols. 2-3). We estimate that high payday effect workers would have missed their selected targets 11.8% of the time had they been under the *control* contract, and actually miss them 5.2% of the time. In contrast, these statistics are 7.3% and 0.8%, respectively, for low payday effect workers.[27]

We do not see changes in take-up of dominated contracts over the paycycle (Table 6, 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%, lower on paydays than non-paydays (Table 6, Panel B, Col. 1). On average, the earnings impact of the *option to choose* treatment declines by Rs. 3 (or 2%) per day

---

[26] There are no significant differences in the elasticity of output with respect to wages between workers with above and below average payday effects (Appendix Table 7, Cols. 2-3).

[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 (Appendix Table 8).

[28] Note that Table 6 regressions use observations after 1 month of experience, to examine trends after workers have learned about the contracts (see Section IV.D). Using all observations does not qualitatively change the results, but slightly decreases their precision.

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 paycycle (Panel B, Cols. 3-4).

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 (Appendix Figure 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 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%), 1 day before their payday (13 log points; significant at 5%), and on their payday (10 log points; significant at 5%) (Appendix 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 (Appendix Table 9, Col. 2).

IV.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 7, 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%, more likely to select positive targets than low payday effect workers (p-value of 0.000; Table 7, 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 7, 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 7 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 7 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% 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%) on average, while there is no detectable change for those with initially higher take-up rates (Cols. 1-2). While only suggestive, this is consistent with a story in which workers with high payday effects but zero initial take-up are naïve, 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% sharply decrease take-up with experience, while those with initially lower take-up rates remain stable over time (Cols. 3-4). 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*.

IV.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 8, Cols. 1-2). Note that positive take-up of dominated contracts in the morning of the workday implies that time inconsistency operates at time periods shorter than a day. Consistent with this, 40% 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 suggests 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 one second to over five 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 the list would be used for this purpose. The computers identified as more uncertain are indeed more sensitive to overall network fluctuations (See Appendix Table 11 and Appendix Figure 3). Assignment to the more uncertain computers decreases mean output by 313 fields or 6% (Appendix Table 12)—a magnitude that corresponds to an 18% 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 which causes workers to shy away from dominated contracts is significant. Consistent with the results in Tables 3-4, 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 the morning (Table 8, 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 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 8, Col. 7).

These results are consistent with the hypothesis that, all else equal, a greater gap between the period of contract choice and 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 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.

IV.F. 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 9, we look at measures of self-control problems based on self-reports by workers during the endline survey. The correlate in Column (1) is the demeaned Self-Control Factor, obtained from a factor analysis on the endline data. In Column (2), we construct a demeaned Self-Control Index by averaging each worker's responses to 9 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 Table 1 notes)—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.

Education (Column 6) positively predicts take-up of dominated contracts, but does not predict treatment effects. IQ (Column 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 Section 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 9 are consistent with those of other studies: laboratory and survey measures of self-control predict field behavior, but often to a limited extent (e.g., see Chabris et al. 2008). This may be because of the various

[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 women. However, there is no difference in the paycyle effects by gender.

measurement issues with laboratory measures (see Table 1 in Chabris, Laibson, and Schuldt 2008; Augenblick, Niederle, and Sprenger 2013). 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%. 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 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 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 this is not happening in our data. First, as we saw in our test for inter-temporal substitution, exogenous production increases do not decrease production on subsequent days (see Appendix 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, it does 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%. Workers with above average payday effects choose dominated contracts 43% of the time, and being offered the option to choose dominated contracts increases earnings by 9%. 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, Kremer, and Mullainathan (2014)[30]—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. 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.

[30] Much of this was in Section 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.

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 naïve and hence are exploited in equilibrium by others who can better predict their behavior (DellaVigna and Malmendier 2004, Gabaix and Laibson 2006, Elias and Spiegler 2006)—the presence of present-biased agents makes time-consistent agents worse off.

In our endline survey, many employees expressed a desire for rules to help them work harder under pure piece rates. 78% of workers agreed with "Some days I don't work as hard as I would like to" and 87% 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%

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% 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 like 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 only as made possible 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". 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 unpunctuality 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 work week, 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, Kremer, and Mullainathan (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 work week 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 of Steven Cheung (1983): "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 like 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, Kremer, and Mullainathan 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, Kremer, and Mullainathan 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 and Demsetz. 1972. "Production, Information Costs, and Economic Organization." *American Economic Review,* 62 (December): 777-795.

Ali, S. Nageeb. 2011. "Learning Self-Control." *Quarterly Journal of Economics*, 126: 857–893.

Allen, Robert. 1992. "Enclosure and the Yeoman: The Agricultural Development of the South Midlands, 1450-1850.

Amabile, Teresa and Steven Kramer. 2011. "The Power of Small Wins." *Harvard Business Review*.

Ariely, Dan and Klaus Wertenbroch. 2002. "Procrastination, Deadlines, and Performance: Self-control by Precommitment." *Psychological Science*, 13(3): 219-224.

Ashraf, Nava, Dean Karlan, and Wesley Yin. 2006. "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines." *Quarterly Journal of Economics*, 121(2).

Augenblick, Ned, Muriel Niederle, and Charles Sprenger. 2013."Working Over Time: Dynamic Inconsistency in Real Effort Tasks." NBER Working Paper 18734.

Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein. 2012. "Behavioral Hazard in Health Insurance", NBER Working Paper 18468.

Bandiera, Oriana, Iwan Barankay and Imran Rasul. 2007. "Incentives for Managers and Inequality Among Workers: Evidence From a Firm Level Experiment." *Quarterly Journal of Economics*, 122: 729-75.

Banerjee, Abhijit, and Sendhil Mullainathan. 2010. "The Shape of Temptation: Implications for the Economic Lives of the Poor." NBER Working Paper No. 15973.

Benhabib, Jess, Alberto Bisin, and Andrew Schotter. 2007. "Present Bias, Quasi-Hyperbolic Discounting, and Fixed Costs." *Games and Economic Behavior*, 69(2): 205–223.

Bernheim, B. Douglas, Debraj Ray and Sevin Yeltekin, 2012, "Poverty and Self-Control", Working paper.

Burger, Nicholas, Gary Charness, and John Lynham. 2009. "Three Field Experiments on Procrastination and Willpower." Available: www.bec.ucla.edu/Charness_Procrastination.pdf

Cadena, Ximena, Alexandra Cristea, Héber M. Delgado-Medrano, and Antoinnette Schoar. 2011. "Fighting Procrastination in the Workplace: An Experiment." Working paper.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. 1997. "Labor Supply of New York City Cabdrivers: One Day at a Time." *Quarterly Journal of Economics*, 111: 408-441.

Chabris, Christopher, David Laibson, Carrie Morris, Jonathon Schuldt, and Dmitry Taubinsky. 2008. "Individual Laboratory-Measured Discount Rates Predict Field Behavior." *Journal of Risk and Uncertainty,* 37(2-3): 237-269.

Chabris, Christopher, David Laibson, and Jonathan Schuldt. 2008. “Intertemporal Choice.” *Palgrave Dictionary of Economics*.

Cheung, Steven. 1969. “Transaction Cost, Risk Aversion and Choice of Contractual Arrangement”, *Journal of Law and Economics*. 12 (1969): 23.

Cheung, Steven N. S. 1983, “The Contractual Nature of the Firm,” *Journal of Law and Economics*, Vol. 26, No. 1. April, pp. 1-21.

Clark, Gregory. 1994. “Factory Discipline.” *Journal of Economic History*, 54: 128-163.

DellaVigna, Stefano. 2009. “Psychology and Economics: Evidence from the Field.” *Journal of Economic Literature*, 47: 315-372.

DellaVigna, Stefano and Ulrike Malmendier. 2004. “Contract Design and Self-Control: Theory and Evidence.” *Quarterly Journal of Economics*, 119: 353-402.

Duflo, Esther, Michael Kremer, and Jonathan Robinson. 2011. “Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya.” *American Economic Review*, 101 (6): 2350-2390.

Dupas, Pascaline and Jonathan Robinson. 2014. “Daily Needs, Income Targets and Labor Supply: Evidence from Kenya.” NBER Working Paper No. 19264.

Eliaz, Kfir and Ran Spiegler. 2006. “Contracting with Diversely Naïve Agents.” Review of Economic Studies, 73: 689-714.

Engen, Eric M., William G. Gale, and John Karl Scholz. 1994. "Do Saving Incentives Work?" *Brookings Papers on Economic Activity*

Engen, Eric, William Gale, and Cori Uccello. 1999. “The Adequacy of Household Saving.” *Brookings Papers on Economic Activity*

Fang, Hanming and Dan Silverman. 2004. “Time Inconsistency and Welfare Program Participation: Evidence from the NLSY,” Cowles Foundation Discussion Paper 1465, Yale University

Fehr, Ernst and Lorenz Goette. 2008. "Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment." *American Economic Review*, 97(1): 298-317.

Frederick, Shane, George Loewenstein, and Ted O’Donoghue. 2002. “Time Discounting and Time Preference: A Critical Review.” *Journal of Economic Literature*, 40(2): 351- 401.

Fudenberg, Drew and David Levine. 2006. “A Dual Self Model of Impulse Control.” *American Economic Review*, 96: 1449-1476.

Gabaix, Xavier and David Laibson. 2006. “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets.” *Quarterly Journal of Economics*, 121(2): 505-540

Giné, Xavier, Jessica Goldberg, and Dean Yang. 2010. "Identification Strategy: A Field Experiment on Dynamic Incentives in Rural Credit Markets," Policy Research Working Paper Series 5438, The World Bank.

Gneezy, Uri, and John List. 2006. "Putting Behavioral Economics To Work: Testing For Gift Exchange In Labor Markets Using Field Experiments", *Econometrica*, 74(5): 1365-1384.

Grossman, Sanford J., and Oliver D Hart 1983. "An Analysis of the Principal-Agent Problem", *Econometrica*, Vol. 51, No. 1 (Jan., 1983), pp. 7-45.

Gul, Faruk and Wolfgang Pesendorfer. 2001. "Temptation and Self-Control," *Econometrica*, 69(6): 1403-1435.

Gul, Faruk and Pesendorfer, Wolfgang, 2004. "Self Control, Revealed Preferences and Consumption Choice", *Review of Economic Dynamics*, 7(2), 243-264.

Holmstrom, Bengt, 1979. "Moral Hazard and Observability." *The Bell Journal of Economics*, Vol. 10, No. 1, (Spring, 1979), pp. 74-91

Hossain, Tanjim and John List. 2012. "The Behavioralist Visits the Factory: Increasing Productivity Using Simple Framing Manipulations." *Management Science*. 58(7): 1120-1544.

Hubbard, Glenn, Skinner, Jonathan and Zeldes, Stephen. 1994. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving," *Carnegie-Rochester Conference Series of Public Policy*

Kaur, Supreet, Michael Kremer, and Sendhil Mullainthan. 2010. "Self-Control and the Development of Work Arrangements." *American Economic Review.* 100(2): 624-628.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainthan. 2015. "Equilibrium Contracts and Job Design When Workers are Present-Biased"

Kirby, Kris and Nino Marakovic. 1995. "Modeling Myopic Decisions: Evidence for Hyperbolic Delay-Discounting with Subjects and Amounts," *Organizational Behavior and Human Decision Processes*, 64: 22-30.

Laibson, David I. 1996. "Hyperbolic Discount Functions, Undersaving, and Savings Policy." NBER Working Paper No. 5635.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*, 112(2): 443-477.

Larkin, Ian. 2013. "The Cost of High-Powered Incentives: Employee Gaming in Enterprise Software Sales." *Journal of Labor Economics*, forthcoming.

Levitt, Steven and John List. 2007. "What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World?" *Journal of Economic Perspectives*. 21: 153-174.

Levitt, Steven and John List. 2009. "Field Experiments in Economics: The Past, The Present, and The Future." *European Economic Review*. 53(1): 1-18.

List, John. 2006. "The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions." *Journal of Political Economy*. 114(1): 1–37.

List, John and Imran Rasul. 2011. "Field Experiments in Labor Economics." *Handbook of Labor Economics*. Vol. 4(A): 103-228.

Loewenstein, George & Prelec, Drazen, 1992. "Anomalies in Intertemporal Choice: Evidence and an Interpretation," *The Quarterly Journal of Economics*, MIT Press, vol. 107(2), pages 573-97, May

Mahajan, Aprajit and Alessandro Tarozzi. 2011. "Time Inconsistency, Expectations and Technology Adoption: The case of Insecticide Treated Nets." Economic Research Initiatives at Duke (ERID) Working Paper No. 105.

Marglin, Stephen Alan. 1974. "What Do Bosses Do? The Origins and Functions of Hierarchy in Capitalist Production." Review of Radical Political Economics. 6(2).

Marglin, Stephen Alan. 2008. "The Dismal Science: How Thinking Like an Economist Undermines Community. Cambridge, MA: Harvard University Press. 2008.

Oyer, Paul. 1998. "Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality." *Quarterly Journal of Economics*. 113(1):149-185.

Oyer, Paul. 2000. "Theory of Sales Quotas with Limited Liability and Rent Sharing." *Journal of Labor Economics*. 18(3):405-426.

O'Donoghue, Ted and Matthew Rabin. 2006. "Incentives and Self-control." *Econometric Society Monographs*, 42: 215-245.

O'Donoghue, Ted and Matthew Rabin. 2001. "Choice and Procrastination," *Quarterly Journal of Economics*, 116(1): 121-160.

O'Donoghue, Ted and Matthew Rabin. 1999a. "Doing It Now or Later," *American Economic Review*, 89(1): 103-124.

O'Donoghue, Ted and Matthew Rabin. 1999b. "Incentives for Procrastinators", *Quarterly Journal of Economics*, 114(3), 769-816.

Paserman, Daniele. 2004. "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation," *CEPR Discussion Paper 4396*.

Shapiro, Jesse. 2005. "Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle." *Journal of Public Economics*, 89: 303-325.

Shearer, Bruce. 2004. "Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment," *Review of Economic Studies*, 71(2): 513-34.

Shui, Haiyan and Lawrence M. Ausubel. 2004. "Time Inconsistency in the Credit Card Market," Econometric Society North America Summer Meetings

Spiegler, Ran. 2011. *Bounded Rationality and Industrial Organization*. Oxford University Press.

Thompson, E. P. “Time, Work-Discipline, and Industrial Capitalism,” Past & Present, no. 38 (December 1, 1967): 56-97.

Figure 1: Incentive Contracts

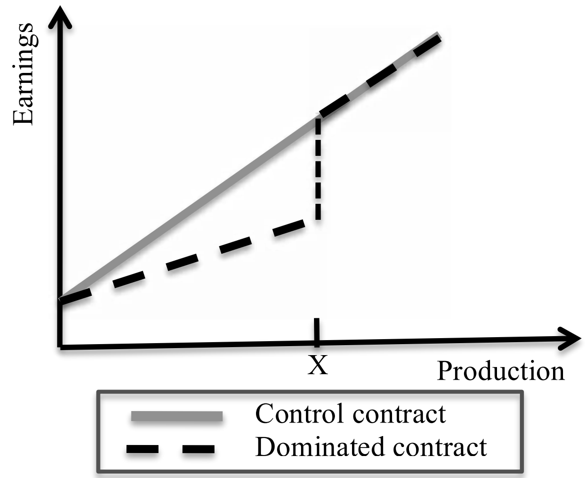

Notes: This figure displays the two types of incentive contracts offered to workers. The linear control contract paid a piece rate wage of *b* for each accurate field entered. The nonlinear dominated contract imposed a production target, X; workers were paid *b* for each accurate field if they met the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

Figure 2: Experiment Timeline

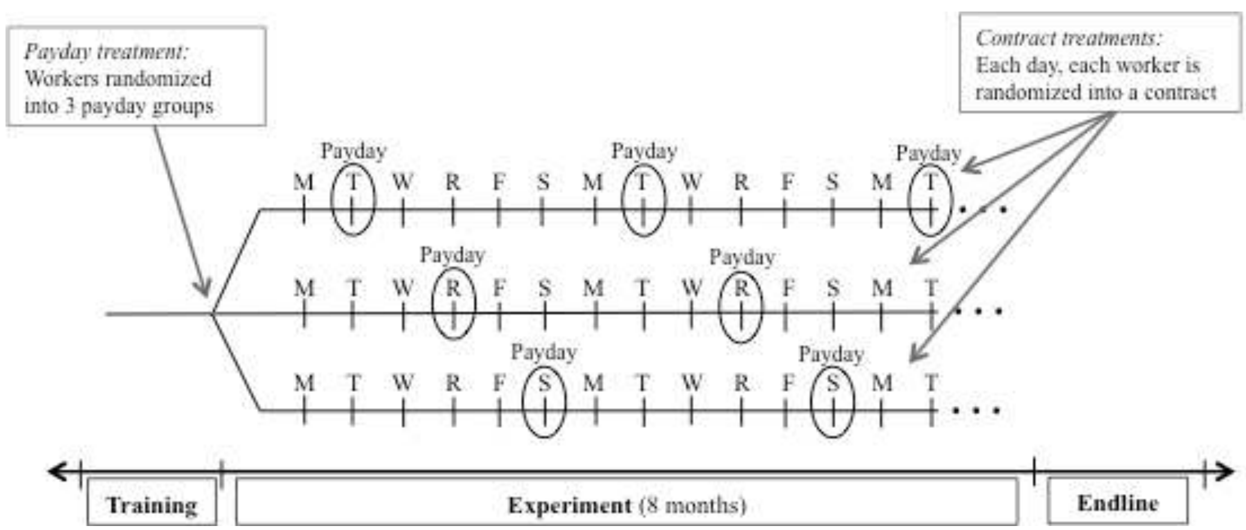

Notes: This figure provides an overview of the experiment timeline.

Figure 3: Earnings over the Pay Cycle

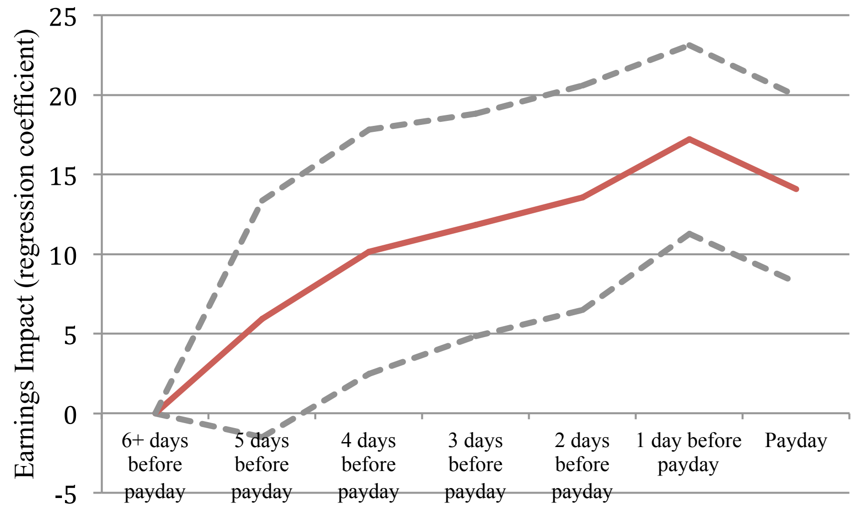

Notes: This figure graphs the coefficients and 95% confidence intervals from a regression of earnings on 6 binary indicators that capture distance from a worker's next payday (payday, 1 day before payday, 2 days before payday, etc). The omitted category is 6 or more days before the payday. Note these coefficients correspond to those shown in Column (4) of Table 2.

Figure 4: Take-up of Dominated Contracts – Distribution of Worker Means

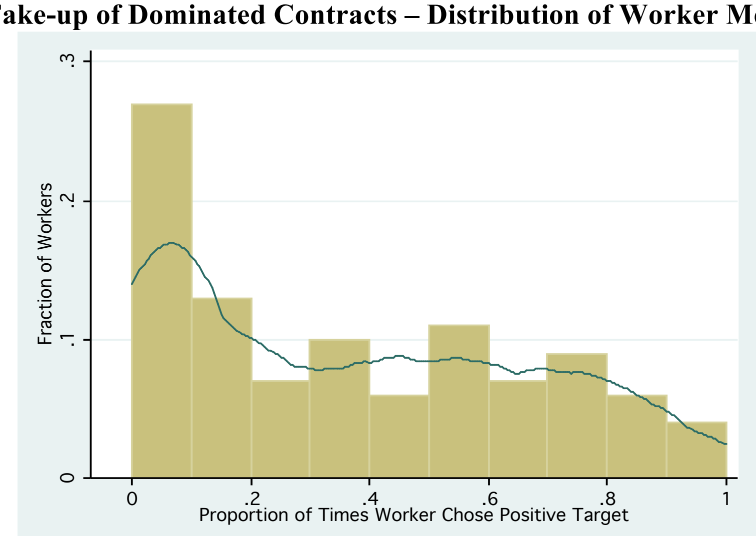

Notes: This figure shows the distribution of take-up rates of the dominated contract by workers. A worker's take-up rate is the proportion of times the worker chose a dominated contract (i.e. selected a positive target) when given the option (conditional on being present the day before and day of assignment to the *option to choose a dominated contract* treatment). The distribution is shown for the 101 workers in the sample that were assigned the *option to choose* at least once.

Figure 5: Pay Cycle Effects - Correlation with Contract Choice and Earnings Impact

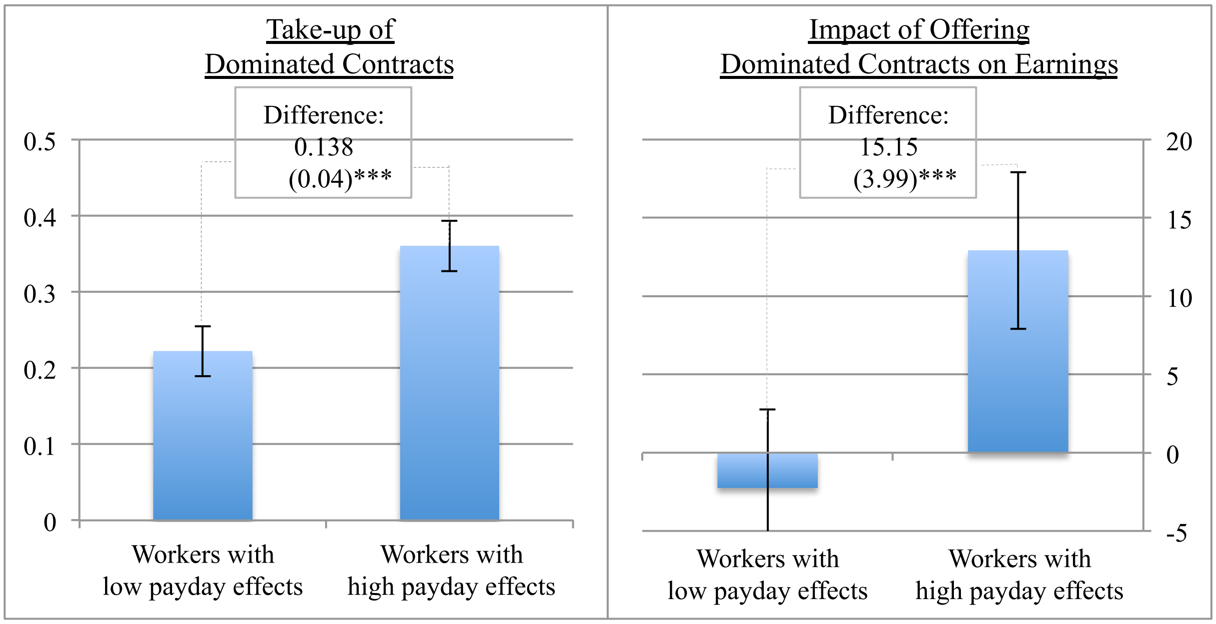

Notes: These figures show differences in take up rates and treatment effects of dominated contracts. Workers with low (high) payday effects are those whose payday effect—the difference in production on paydays and non-paydays under assignment to the *control* contract, divided by mean production under 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 5. Each bar corresponds to the estimated mean for each group, along with 95% confidence intervals.

Figure 6: **How the Demand for Dominated Contracts Changes with Experience**

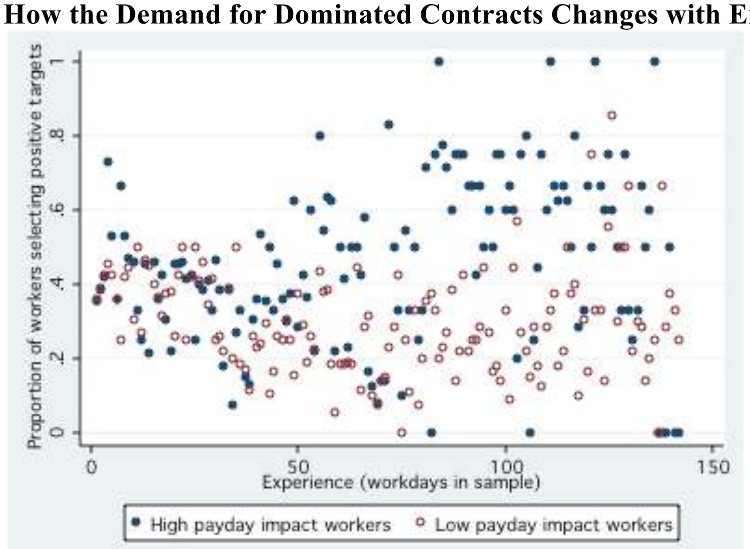

Notes: Worker experience is the number of workdays the worker has been in the sample. The proportion of times positive targets were chosen is computed for each value of the experience variable using observations in which the worker was given the *option to choose a dominated contract* (conditional on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

Table 1: Summary Statistics

	Mean	Std Dev	Obs
	(1)	(2)	(5)
A. Worker Characteristics			
Proportion female	0.26	0.44	111
Age	24	4	63
Years of education	13	2	101
Completed high school	0.84	0.37	101
Used computer prior to joining firm	0.67	0.47	101
Had email address prior to joining firm	0.60	0.49	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. Endline Survey: Discount Rate Measurement			
Proportion of times worker chose smaller immediate reward	0.31	0.28	58
Proportion of times worker displayed preference reversal	0.17	0.23	58
D. Endline 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."	0.76	0.43	70
"I get tempted to leave work earlier than I would like."	0.40	0.49	70
"I wish I had better attendance at work."	0.86	0.35	70
"It would be good if there were rules against being absent because it would help me come to work more often."	0.73	0.45	70
Self-control index: mean of responses to all 9 self-control questions (1=disagree strongly; 5=agree strongly)	3.43	0.55	70
Worker has tried to quit an addictive behavior and failed (males only)	0.12	0.33	51
Factor analysis: self-control factor	0.00	0.86	70

Notes: This table presents summary statistics for the 111 workers that participated in the full study (contract and payday treatments). In the discount rate exercise (Panel C), workers traded off 3 sets of cash awards (Rs. 20 vs. Rs. 24; Rs. 50 vs. Rs. 57; and Rs. 100 vs. Rs. 110) under 2 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 choose the smaller immediate reward out of the 6 questions, and the number of times the worker showed preference reversal (chose the smaller immediate reward in the short horizon, and choosing 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 endline survey questions.

Table 2
Paycycle Treatment Effects

Dependent variable	*Production* (1)	*Production* (2)	*Production* (3)	*Earnings* (4)	*Attendance* (5)	*Production* (6)
Payday	215	140	428	14.09	0.077	
	(70)***	(63)**	(94)***	(2.99)***	(0.013)***	
1 day before payday			539	17.19	0.053	
			(95)***	(3.02)***	(0.013)***	
2 days before payday			417	13.54	0.037	
			(113)***	(3.60)***	(0.016)**	
3 days before payday			374	11.82	0.026	
			(112)***	(3.57)***	(0.017)	
4 days before payday			332	10.15	0.047	
			(123)***	(3.91)***	(0.017)***	
5 days before payday			176	5.91	0.023	
			(119)	(3.79)	(0.017)	
Piece rate increase						1071
						(239)***
Lag dependent variable controls	No	Yes	Yes	Yes	No	Yes
Observations	8423	8423	8423	8423	8423	550
R2	0.50	0.59	0.59	0.57	0.11	0.76
Dependent variable mean	5337	5337	5337	172	0.88	9361

Notes: The sample in columns (1)-(5) is the experiment sample. The dependent variables equal 0 if a worker was absent. Payday is a binary indicator for whether that day was the worker's assigned payday. In columns (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 column (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 1 if the worker's piece rate was Rs. 0.04 per accurate field that day, and equals 0 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 two workdays ago), and similarly regression (4) includes controls for lagged earnings. Robust standard errors are reported in parentheses.

Table 3
Contract Treatments

Panel A: Take-up of Dominated Contracts (Summary Statistics)

Dominated contract chosen: conditional on attendance	0.36
	(0.31)
Dominated contract chosen: target=0 if absent	0.28
	(0.26)

Panel B: Treatment Effects of Contracts

	Dependent variable				
	Production			*Earnings*	*Attendance*
Sample	Control & Option	Control & Option	Full Sample	Full Sample	Full Sample
	(1)	(2)	(3)	(4)	(5)
Option to choose dominated contract	120				
	(59)**				
Evening option to choose dominated contract		156	150	4.60	0.01
		(69)**	(69)**	(2.17)**	(0.01)
Morning option to choose dominated contract		84	73	2.32	-0.00
		(69)	(69)	(2.17)	(0.01)
Target imposed: Level 1			3	-1.55	-0.00
			(90)	(2.88)	(0.01)
Target imposed: Level 2			213	3.13	-0.01
			(91)**	(2.89)	(0.01)
Target imposed: Level 3			334	5.01	-0.01
			(150)**	(4.80)	(0.02)
Observations: worker-days	6310	6310	8423	8423	8423
R2	0.60	0.60	0.59	0.57	0.15
Dependent variable mean	5311	5311	5337	172	0.88

Notes: The Panel A 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 day of treatment assignment. The second row codes target choice as 0 if absent the day before or day of assignment. Means and standard deviations are presented for each row.

In Panel B, the dependent variables equal 0 if the worker is absent. Cols (1)-(2) limit analysis to observations where workers were assigned to the *control* contract or given the *option to choose a dominated contract*. Cols (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 OLS regressions are shown; robust standard errors are reported in parentheses.

Table 4
Chance of Missing Targets: Summary Statistics

	All workers (1)	*High payday effect workers* (2)	*Low payday effect workers* (3)
Panel A: Predicted probability of missing target (using control contract output distribution)			
Self-chosen target	0.091	0.118	0.073
	(0.120)	(0.146)	(0.096)
Panel B: Proportion of times target was actually missed			
Self-chosen target	0.026	0.052	0.008
	(0.122)	(0.187)	(0.036)
Target imposed: Level 1	0.086	0.120	0.063
	(0.155)	(0.198)	(0.115)
Target imposed: Level 2	0.113	0.117	0.110
	(0.197)	(0.240)	(0.167)
Target imposed: Level 3	0.141	0.177	0.115
	(0.265)	(0.315)	(0.224)

Notes: 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 where 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 *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 CDF 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. Panel B, rows 2-4 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 non-paydays under assignment to the *control* contract divided by mean production under the *control* contract—is above the sample average. Col (1) presents these statistics for the workers that 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)-(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).

Table 5
Paycycle Effects:
Correlation with Dominated Contract Effects

	(1)	(2)	(3)
Panel A: Take-up of Dominated Contracts			
Dependent variable	Dominated contract chosen	Target level chosen	Target level chosen
High payday effect worker	0.138	351	338
	(0.044)***	(129)***	(126)***
High productivity worker			-105
			(134)
Observations: worker-days	4098	4098	4098
R2	0.20	0.22	0.22
Dependent variable mean	0.28	759	759
Panel B: Treatment Effects of Contracts			
Dependent variable	Production	Earnings	Attendance
Option to choose dominated contract	-69	-2.24	-0.016
	(74)	(2.34)	(0.010)
Option to choose dominated contract * High payday effect worker	482	15.15	0.058
	(126)***	(3.99)***	(0.019)***
Target imposed	-35	-3.82	-0.019
	(86)	(2.74)	(0.012)*
Target imposed * High payday effect worker	483	14.31	0.042
	(148)***	(4.71)***	(0.022)*
Observations: worker-days	8240	8240	8240
R2	0.59	0.57	0.11
Dependent variable mean	5355	173	0.875

Notes: The table present OLS 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 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 non-paydays 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 is comprised of all observations. The dependent variables equal 0 if a worker was absent. Each regression includes worker, date, and computer seat assignment fixed effects. The Co (1) and (2) regressions also include lagged production controls and lagged earnings controls, respectively. Robust standard errors are reported.

Table 6

Correlation in Payday and Dominated Contract Effects: Trends over the Paycycle

	Sample: High payday effect workers		*Sample: Low payday effect workers*	
	(1)	(2)	(3)	(4)
Panel A: Demand for Dominated Contracts (Dependent variable: Dominated contract chosen)				
Payday dummy	-0.013		-0.002	
	(0.030)		(0.024)	
Day in paycycle (linear control)		0.006		-0.002
		(0.005)		(0.004)
Observations	1252	1252	1967	1967
Dependent variable mean	0.319	0.319	0.230	0.230
Panel B: Treatment Effects of Contracts (Dependent variable: Earnings)				
Paycycle measure	Payday dummy	Day in paycycle (linear control)	Payday dummy	Day in paycycle (linear control)
Paycycle measure	32.88	6.32	-3.40	0.69
	(7.59)***	(1.49)***	(6.00)	(1.09)
Option to choose dominated contract	21.58	25.49	-4.41	-9.33
	(4.55)***	(6.81)***	(3.14)	(4.93)*
Option to choose dominated contract * Paycycle measure	-33.33	-3.33	11.93	2.39
	(9.16)***	(1.70)**	(6.90)*	(1.28)*
Target imposed	18.90	26.17	-0.67	-1.73
	(5.22)***	(7.96)***	(3.63)	(5.75)
Target imposed * Paycycle measure	-32.71	-4.24	-1.28	0.278
	(10.70)***	(2.04)**	(8.30)	(1.51)
Observations	2502	2502	3947	3947
Dependent variable mean	159	159	196	196

Notes: Cols (1)-(2) report results for high payday effect workers--those whose mean payday effect (the difference in production on paydays and non-paydays under assignment to the *control* contract, divided by mean production under the *control* contract) is above the sample average. Cols (3)-(4) report results for workers with a below average payday effect. OLS 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 Section 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 day of treatment assignment, the dependent variables is coded as zero. Regressions include computer and date fixed effects and lagged earnings controls.

The Panel B sample is comprised of 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.

Table 7
Changes in Outcomes with Worker Experience

	Experience Measure		
	Log Number of Days Worked	Log Number of Days Worked	More than 2 Months Worked
	(1)	(2)	(3)
Panel A: Take-up of Dominated Contract			
Experience measure	-0.040	-0.066	-0.096
	(0.027)	(0.030)**	(0.057)*
High payday effect worker	0.139	-0.087	0.074
	(0.043)***	(0.100)	(0.050)
Experience measure *		0.062	0.132
High payday impact		(0.026)**	(0.054)**
Observations	4098	4098	4098
Dependent variable mean	0.28	0.28	0.28
Panel B: Treatment Effects on Earnings			
Experience measure	7.45	4.67	-3.82
	(2.11)***	(2.34)**	(4.93)
Option to choose dominated contract	3.38	-8.26	-0.87
	(1.87)*	(5.54)	(2.28)
Experience measure*		3.20	8.93
Option to choose dominated contract		(1.63)**	(3.78)**
Payday	4.84	4.65	4.36
	(1.99)**	(5.72)	(2.48)*
Experience measure *		0.05	0.96
Payday		(1.62)	(3.79)
Observations	8423	8423	8423
Dependent variable mean	172	172	172

Notes: OLS regressions are shown; standard errors are reported in parentheses. In Cols (1)-(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 is comprised of observations in which workers were given the *option to choose a dominated contract*. The dependent variables is defined as 0 if a worker was absent the day before or day of treatment assignment. High payday impact is a binary indicator for whether the worker's mean payday impact (the difference in production on paydays and non-paydays 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 is comprised of all observations. Production is defined as 0 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.

Table 8

Demand for the Dominated Contract: Impact of Stochasticity 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	
Dependent Variable	*Dominated contract chosen*	*Target level chosen*	*Dominated contract chosen*	*Target level chosen*	*Dominated contract chosen*	*Target level chosen*	*Dominated contract chosen*	*Target level chosen*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Evening option to choose dominated contract	-0.002	-18	-0.003	-26	0.066	168	0.027	87
	(0.012)	(37)	(0.010)	(32)	(0.022)***	(72)**	(0.016)	(46)*
High uncertainty indicator			-0.013	-134	0.027	-20	0.104	282
			(0.016)	(63)**	(0.022)	(82)	(0.062)	(206)
Evening option to choose dominated contract *					-0.082	-230	-0.070	-253
High uncertainty indicator					(0.024)***	(78)***	(0.029)**	(97)**
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Seat fixed effects	Yes	Yes	No	No	No	No	Yes	Yes
Observations	4193	4193	4193	4193	4193	4193	3106	3106
R2	0.33	0.34	0.32	0.33	0.32	0.33	0.12	0.13
Dependent variable mean	0.28	767	0.28	767	0.28	767	0.3	803

Notes: The sample is comprised of worker-day observations in which workers were given the *option to choose a dominated contract* (in the evening or morning). Both dependent variables are defined as 0 if a worker was absent the day before or day of treatment assignment.

Evening option to choose dominated contract is a dummy that equals 1 if the worker was assigned to choose the evening before the workday, and equals 0 if the worker was assigned to choose the morning of the workday. In columns (3)-(6), the high uncertainty indicator equals 1 if the worker was assigned to a computer that was highly sensitive to office network speed, and equals 0 otherwise. In columns (7)-(8), the high uncertainty indicator equals 1 if 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,” during the endline survey. All regressions include date fixed effects. Robust standard errors are reported in columns (1)-(2). Standard errors are corrected to allow for clustering by computer in columns (3)-(6) and by worker in columns (7)-(8).

Table 9: Heterogeneity in Treatment Effects - Correlates of Self-Control

Correlate of self-control	*Factor analysis: Self-control factor* (1)	*Self-control index* (2)	*Addictive behaviors dummy (males only)* (3)	*Discount rate: Proportion impatient responses* (4)	*Discount rate: Proportion preference reversals* (5)	*Years of education* (6)	*IQ test index score* (7)
Panel A: Dependent variable: Dominated contract chosen							
Correlate	0.055	0.057	0.140	0.068	0.143	0.029	-0.001
	(0.025)**	(0.046)	(0.082)*	(0.115)	(0.189)	(0.015)**	(0.002)
Observations	3106	3106	2245	2454	2470	4056	4089
R2	0.22	0.21	0.26	0.23	0.24	0.19	0.19
Panel B: Dependent variable: Target level chosen							
Correlate	122	147	354	220	533	122	-1
	(74)	(140)	(238)	(342)	(572)	(43)***	(5)
Observations	3106	3106	2245	2454	2454	4056	4089
R2	0.25	0.25	0.29	0.27	0.28	0.23	0.22
Panel C: Dependent Variable: Earnings							
Correlate	-8.01	-6.68	-8.69	-39.73	-25.54	4.02	0.39
	(3.26)**	(4.68)	(12.59)	(10.05)***	(14.56)*	(2.03)*	(0.19)**
Option to choose dominated	2.84	2.81	1.27	3.60	3.71	4.59	4.72
contract	(2.60)	(2.64)	(3.62)	(2.63)	(2.61)	(2.31)**	(2.33)**
Option to choose dominated	5.19	6.81	13.77	22.03	23.75	1.71	-0.02
contract * Correlate	(2.91)*	(4.65)	(8.37)	(9.65)**	(14.17)*	(1.42)	(0.16)
Payday	5.48	5.50	4.22	4.78	4.65	5.40	6.30
	(2.72)**	(2.74)**	(3.68)	(3.29)	(3.33)	(2.30)**	(2.20)***
Payday *	0.83	2.34	1.75	-1.39	-7.71	0.60	0.00
Correlate	(2.77)	(4.30)	(7.82)	(9.70)	(17.32)	(1.25)	(0.12)
Observations	4674	4674	3376	3701	3701	6101	6149
R2	0.55	0.55	0.53	0.57	0.56	0.54	0.53

Notes: The self-control correlate in column (1) is the demeaned Self-control Factor, obtained from a principal factors analysis on the endline survey data. The column (2) correlate is a demeaned Self-Control Index, obtained by averaging each worker's responses to the 9 self-control questions in the endline survey. The correlate in column (3) is computed for male workers; it equals 1 if the worker said he has tried to quit drinking, smoking, or chewing tobacco and failed, and equals 0 otherwise. The correlates in columns (4)-(5) are computed from the discount rate exercise, in which workers traded off cash rewards between different time horizons. The column (4) correlate measures the proportion of times the worker chose the smaller immediate reward instead of the larger delayed reward. The column (5) correlate measures preference reversals—the proportion of times a worker chose the larger immediate reward in the short horizon, but then chose the smaller delayed reward when choosing among the same amounts in the long horizon. The correlates in columns (6) and (7) are, respectively, years of education and composite IQ score (the sum of the worker's score on the Raven's Matrix and Digit Span tests). The self-control correlate in each column has been demeaned.

In Panels A-B, the sample is observations where workers were assigned the *option to choose a dominated contract*. The dependent variables in Panels A-B are defined as 0 if a worker was absent the day before or day of treatment assignment. The Panel C sample is observations where workers were assigned to the *control* contract or *option to choose a dominated contract*. All regressions include date and computer seat assignment fixed effects, and lagged earnings controls. Note that observations change between columns because not all workers provided education information or took the IQ tests, and because the endline survey and discount rate exercise were administered only at the end of the project.

OLS regressions are shown. Standard errors are corrected to allow for clustering by worker.