Self-Control at Work

Supreet Kaur
Michael Kremer
Sendhil Mullainathan[*]

Abstract

Self-control problems change the logic of agency theory: workers not only fail to work as hard as employers would like, they fail to work as hard as they themselves would like. In response, firms can use incentive pay to affect the self-control problem, not just moral hazard. We describe the results of a year-long field experiment on data entry workers designed to test the empirical importance of these ideas. First, we find that workers will choose *dominated* contracts—which pay less for every output level but have a steeper slope—in order to motivate themselves. Second, their effort increases significantly as the (randomly assigned) payday gets closer. Third, these two effects are linked: the demand for dominated contracts (and their benefits) is concentrated amongst those with the highest payday effects. Finally, as workers gain experience, they appear to learn about their self control problems: the correlation between the payday effect and the demand for the dominated contract grows with experience. Both payday and contract effects are quantitatively large in magnitude when benchmarked against the impact of a change in the slope of incentives or of a year of education. These results together suggest that self-control, in this context at least, meaningfully alters the firm's contracting problem.

---

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

We thank Lawrence Katz, David Laibson, John Matsusaka, Ted O'Donoghue and participants at numerous seminars for helpful comments, and Raluca Dragusanu for outstanding research assistance. We thank Dominik Bulla, Karthik R, and the Centre for Innovative Financial Design (IFMR) for operational support and Dan Bjorkegren, Matthew Darling, and Nicolas Papamichael for expert research assistance in 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: workers do not work as hard as firms would like (Holmstrom 1979; Grossman and Hart 1983). Wages do not reflect the full benefits of work—the employer provides some insurance—and so workers are imperfectly incentivized. Introspection suggests another problem at work: self-control.[1] Looking to the future, they would like to work hard. Acting in the moment, though, they would rather relax. This raises a new tension: workers do not work as hard as they themselves would like.

This tension on the worker side changes the logic of contract design. A simple example illustrates how. Suppose a firm earns revenue by entering data and faces a penalty C for not completing by a certain date. It hires a worker to enter this data at some *wage w* and gives her a penalty *c* for not meeting the deadline. The penalty creates risk for the worker: the data could turn out to be complex to enter so that even at her best effort the worker may not be able to meet the deadline. This increased risk requires compensation. Self-control changes this simple logic. A worker with self-control problems may see benefits to an increased penalty in tomorrow's contract (assuming she is sophisticated as in O'Donoghue and Rabin 1999). An increased penalty will motivate her to work harder tomorrow, which she values today. As a result, she may not need compensation. This generates a striking prediction. Sophisticated workers with self-control problems may prefer a *dominated* contract, one that pays less for every output realization. In effect, the incentive scheme can be an implicit commitment device (Laibson 1997; Ashraf et. al. 2006; Gine et. al. 2010).

This observation separates work from several other self-control domains that have been studied (e.g. savings and smoking). An independent agent—the employer—has both the means and motives to reduce worker self control. Sharp monetary incentives or firing for tardiness and inadequate performance can ameliorate self-control problems. This could alter our understanding of work arrangements. Employers can now actually increase productivity beyond

[1] Frederick et al (2002) and DellaVigna 2009 review the self-control literature. Prominent models include Laibson (1996), O'Donoghue and Rabin (1999; 2001), and Fudenberg and Levine (2006). Bernheim, Ray and Yeltekin (2011) and Banerjee and Mullainathan (2009) examine it in the development context. Gul and Pesendorfer (2001; 2004) provide a different account of the demand for commitment.

what workers could achieve as full residual claimants.[2] This argument makes two key assumptions: self-control at work is quantitatively important and workers are sufficiently sophisticated as to value the implicit self-control benefits of stronger incentives.

We performed a field experiment to test these assumptions.[3] In the experiment, data entry workers are paid weekly according to a piece rate that depends on the number of accurate fields entered each day. Workers were randomized into two conditions. Workers in the *Control* condition face the standard piece rate of w. Workers in the *Choice* condition are allowed to choose a target T for the day: if they meet the target they receive the standard piece rate w; if they fail, they receive half the piece rate $w/2$. Like the penalty in the fictional example, the target increases penalty without increasing reward. In one sense, a zero target contract dominates every positive target contract. Workers were allowed to choose their targets either the evening before (for the next day) or that morning; they were randomly assigned into these *Morning* and *Evening* choice conditions. All contract assignments were randomly assigned daily for each worker.[4] To measure the impact of time horizon, workers were also randomized into different payday groups: all were paid weekly but the exact day of payment varied. These randomizations were at the worker level and paydays once assigned were fixed. Our experiment takes place over 13-months in Mysore, India. Workers were typically high school graduates for whom this was their primary source of earnings.

[2] Clark (1994) makes this case for the rise of the factory during the industrial revolution. O'Donoghue and Rabin (1999;2006) formalize how firms use deadlines to motivate a procrastinator. They also produce interesting implications for screening which we examine only briefly here (see Section IV.F.). DellaVigna and Malmendier (2004) also study contract design in a different context. Kaur et al (2010) discuss the work context. Firms could also increase output by solving the free riding problem in team production (Cheung (1969) and Alchian and Demsetz (1972)).

[3] Ariely and Wertenbroch (2002) and Burger et. al. (2008) provide evidence of self control in work. Though quite interesting, both papers involved student populations and smaller stakes. Shearer (2004), Gneezy and List (2006), Bandiera et al (2007), Fehr and Goette (2008), and Hossain and List (2009) use field experiments to study other features of worker psychology.

[4] In addition to the control and two choice conditions above, we also randomize workers into one of three dominated contracts *without choice*. While not useful for understanding self-control, these provide useful calibration benefits. This is reported in greater detail in Section IV. We also randomized workers seats so that every 1 to 3 weeks they moved to different seats. This allows us to estimate peer effects as discussed in Kaur, Kremer and Mullainathan (2010, 2011).

We find three main results. First, workers work harder as the payday gets closer. They earn 8% more on paydays than at the beginning of the weekly pay cycle.[5] This is not concentrated on the payday: production rises smoothly through pay cycle. Second, dominated contracts are chosen on 35% of the workers days where they are offered (among workers that are present). The ability to choose a dominated contract increases productivity: the *Choice* condition, shows a Local Average Treatment Effect of 6%. Third, these two effects are related. We find significant heterogeneity in the payday effect and this heterogeneity predicts the impact of *Choice*. Workers with above median payday effects are 49% more likely to choose the dominated contracts, and show a 20% treatment effect of *Choice* on output. The workers with the biggest self-control problems appear to value the dominated contracts the most. We argue in Section V that while any one of these findings may be explained in some other way (e.g. workers trying to signal ability to employers) only the self control interpretation fits these three facts together.

We calibrate the magnitude of these effects against two benchmarks. We compare the magnitudes to a simple OLS estimate of the returns to education: the payday effect is comparable to a little more than a 1-year increase in education and the ability to choose the dominated contract increases output at the same level roughly of two-thirds of a year of education. We also calibrate a simple model to estimate the implied time inconsistency—the extent to which the discount rate changes. To do this, requires mapping out the cost of effort curve which we do using an exogenous change in the overall piece rate which happened after the experiment. The payday effect suggests a discount rate of roughly 5% per day. The *Choice* treatment suggests that the difference in discounting of benefits between the self that chooses the contract and the one that works to be at least 18%. For the half of the workers with the largest payday effects, the impact of choice rises to roughly 2.5 years of education or a time inconsistency between the chooser and doer of 64%.

Two other findings stand out. First, workers seem to learn about their self-control problems. Early on, many workers experiment with dominated contracts when in the *Choice* condition. As they gain experience, workers diverge: some choose positive targets more while others choose

[5] To avoid selection bias due to selective attendance, we report production results with absent workers coded as zero production. Workers are told the next day's contract assignment before leaving for the day.

zero targets more regularly. The workers who increase demand are also those with the highest payday effects, suggesting accurate learning. The treatment effect of *Choice* on productivity also increases with worker experience: as workers sort better, the benefits of choice increase. In contrast, the effect of paydays neither declines nor increases with experience. These results suggest that workers learn about the extent of their self-control problems but do not necessarily learn away these problems (possibly suggesting the availability of external commitment devices to help with self control at work is low).

Second, inconsistent with the simplest self-control models, we fail to find a difference between *Morning* and *Evening* choice: targets chosen for the next day are the same as for the same day. *Ex post analysis* of the data suggests a possible reason: workers may face uncertainty the evening before that will be realized when they reach work. Variability in computer speed and time they can reach work (e.g. due to uncertain buses) may offset a greater desire for self control the evening before. In fact, targets are higher the evening before when these measures of uncertainty are low and this reverses when uncertainty is high. Self-control models that allowed for such uncertainty that is revealed over time could easily explain these results but we ourselves did not make such a prediction prior to running the experiment.[6]

Finally, we find mixed evidence of heterogeneous treatment effects. Workers with above average productivity were 40% more likely to demand dominated contracts but interestingly benefited less from their provision. This hints at the possibility that there may be variability in the extent of sophistication. At the end of the experiment, we conducted surveys to elicit subjective measures of self-control at work and in other domains (such as smoking).[7] As a whole, these measures, however, showed little predictive power, either because of a lack of statistical power or because these abstract measures are inherently noisy.

II. Model

[6] We were motivated to do this analysis when we noticed a significant drop in demand for the targets during days of slow network speed. This led us to see whether this uncertainty mediated the evening-morning effect.
[7] For example, agreement with statements like, "Some days I don't work as hard as I would like."

We present a simple principal-agent model that incorporates worker self-control. Since our empirical work will focus on the *demand* for contracts, we focus only on worker utility under different contracts: we do not explicitly derive the optimal contract here.

An agent exerts unobservable continuous effort *e* to produce stochastic binary output, y. Output equals 1 with probability p(e) and equals 0 otherwise. Output is perfectly observable, and the agent is paid in period T as a function of it. Write L and H for the pay in the low and H states and Δ = H-L. Agents exert effort in period 1, output is realized and pay is given in period T, and incentive contracts are signed in period 0.

Agents discount a payoff at horizon τ (i.e. τ periods in the future) by d(t) where d(t) decreasing in t, d(t) ≤1 and d(0)=1. For the time consistent case, we write $d(t)=\delta^t$ where δ equals the *daily* discount factor and t is measured in days. For time inconsistent agents, we assume that the impatience for a delay of *s* periods $\frac{d(t+s)}{d(t)}$ is decreasing in t for any fixed *s*.[8] These assumptions generate a discount function that matches what is used to model hyperbolic discounting: people are particularly impatient for receiving a payment today versus tomorrow, less impatient for tomorrow versus day after tomorrow and more generally their impatience falls with the horizon. For the empirical work, we further implicitly assume that d(1) is sufficiently less than 1. The agent is risk averse, with concave utility, u(•), over income and a cost of effort *c(e)* which is convex in *e*.[9]

[8] A hyperbolic discount factor $d(\tau) = (1+\alpha\tau)^{-\gamma/\alpha}$, where α captures deviations from exponential discounting (Lowenstein and Prelec 1992), will satisfy this property. Quasi-hyperbolic discounting (see Laibson 1997) or more broadly, present bias (see Benhabib et al 2007) satisfies these assumptions with one possible caveat. A quasi-hyperbolic function has $d(\tau)=\beta\delta^{\tau}$ so that $\frac{d(1)}{d(0)}=\beta\delta$ and $\frac{d(\tau+1)}{d(\tau)}=\delta$ for $\tau > 0$. So the rate of time preference is strictly decreasing between τ= 0 and τ= 1 and the same between other periods. The conflict is purely about now versus later with no conflict in future periods. For our purposes we are interested in conflict even in future horizons: for example we will look not just at a pay*day* effect but also affects as the payday gets closer.

[9] In writing this utility function we are assuming that the agent only consumes pay when it is given. If agents have access to perfect credit markets and suffer from no other psychological biases (e.g. mental accounting) then consumption utility would not depend on the actually date of pay.

We will write utility as $U_{\{0,1\}}^{\{C,I\}}$ to indicate the utility for either a time consistent (C) or time inconsistent (I) agent at time 0 or 1. Based on our assumptions:

$$U_1^C = \delta^T[p(e)u(H)+(1-p(e))u(L)]-c(e)$$
$$U_0^C = \delta^{T+1}[p(e)u(H)+(1-p(e))u(L)]-\delta c(e)$$
$$U_1^I = d(T)[p(e)u(H)+(1-p(e))u(L)]-c(e)$$
$$U_0^I = d(T+1)[p(e)u(H)+(1-p(e))u(L)]-d(1)c(e)$$

Note that for the time consistent $U_0^C = \delta U_1^C$ but $U_0^I \neq d(1)U_1^I$.

Optimal effort for the time consistent denoted by e_C^* is given by the first order condition:

$$\delta^T[p'(e_C^*)u'(H)+(1-p'(e_C^*))u'(L)]-c'(e_C^*)$$

For the time inconsistent, desired effort differs between time 0 and 1, is denoted by e_0^* and e_1^*. Since time 1 chooses effort the realized effort is given by the first order condition:

$$d(T)[p'(e_1^*)u'(H)+(1-p'(e_1^*))u'(L)]-c'(e_1^*)$$

The difference in these formulas trivially illustrates the payday effect. As T goes down (a closer payday), for a time consistent agent, since $\delta \approx 1$ there should be no noticeable changes in e_C^*.[10] Since we are assuming that *d(1)* < 1 e_1^* will decrease in *T*.

Prediction 1: As the lag between effort and compensation decreases, a time inconsistent agent will supply greater effort. In contrast, there will be no noticeable changes in effort provision if the agent is time consistent.

Period 0's desired effort level is given by the first order condition:

$$d(T+1)[p'(e_0^*)u'(H)+(1-p'(e_0^*))u'(L)]-d(1)c'(e_0^*)$$

Since period 0 thus weighs the benefits of effort relative to the costs more heavily than period 1 $\left(\frac{d(T+1)}{d(T)} > \frac{d(T)}{1}\right)$ period 0 desires more effort than is provided: $e_0^* > e_1^*$. This is the heart of the time inconsistency problem.

[10] When presenting results in Section VII, we return to this assumption and explicitly calibrate the exponential discount rate that is implied by our empirical results.

This inconsistency changes period 0's "demand" for different contracts. Suppose we change the payment in the low state L. For the time consistent worker at time 0:

$$\begin{aligned} \frac{\partial U_0^C}{\partial L} &= \delta^{T+1}[1-p(e_C^*)]u'(L)+\delta\frac{\partial e_c^*}{\partial L}\left(\delta^T p'(e_C^*)u'(H)+(1-p'(e_C^*))u'(L)]-c'(e_C^*)\right) \\ &= \delta^{T+1}[1-p(e_C^*)]u'(L) \\ &> 0 \end{aligned}$$

This is intuitive: more pay in the low state generates an income effect. The disincentive effect can be disregarded because of the envelope theorem—the agent was already equating marginal cost and benefit of effort. Importantly, period 1 who chooses the effort has the same first order condition as period 0 who is valuing the contract. As a result, a decrease in L lowers utility and an increase raises utility.

A time inconsistent worker has a different perspective. As far as she is concerned, the marginal costs and benefits of effort, as she weighs them, have not been equalized by period 1. So the impact of a change in L on incentives must also be considered:

$$\frac{\partial U_0^I}{\partial L}=d(T+1)[1-p(e_1^*)]u'(L)+\frac{\partial e_1^*}{\partial L}\left(d(T+1)p'(e_1^*)u'(H)+(1-p'(e_1^*))u'(L)]-d(1)c'(e_C^*)\right)$$

As before the first term is positive: change L and income changes. The second term—the incentive effect—is negative. This is easiest to see for a reduction in L: lowering pay in the low state increases effort. Since 0 wants more effort and this incentive effect raises utility for 0 and offsets the income effect. Thus, if the agent's self-control problem is sufficiently severe, in period *0* utility may rise when *L falls.* This means that even holding pay constant in the high state (*H*) she may be happier with a contract that pays her less in the low state. Thus a *dominated* contract—one that pays less in some states and no more in all states—can improve her utility. This leads to the next prediction.

Prediction 2: A (sophisticated) time inconsistent agent may prefer a dominated contract that increases the marginal returns to effort. In contrast, a time consistent agent would never prefer such a contract. Providing a (sophisticated) time inconsistent agent the option to select a dominated contract will increase effort, output, and earnings.

Note that a direct implication of our model is that there will be a correlation between our first two sets of predictions. Time consistent workers will not be affected by the timing of compensation, nor will they select dominated contracts. In contrast, sophisticated time inconsistent workers would select a dominated contract[11]:

Prediction 3: *An agent that is affected by the timing of compensation will be more likely to select and benefit from a dominated contract.*

When will time inconsistent sophisticates be most likely to prefer the dominated contract? Trivially in our model period *0* could choose a dominated contract, whereas period *1* would not. If we expand our model to allow for *s+1* periods prior to effort the term in front of the incentive effect $\frac{\partial e_1^*}{\partial L}$ would become $\frac{d(T+s+1)}{d(T+s)} - d(1)$ and would therefore be increasing in *s*. In other words, demand for the dominated contract would increase with the horizon between the choice and effort period. Inconsistent individuals are more likely to choose a dominated contract for next week than for tomorrow. [12,13]

Prediction 4: A time inconsistent agent will be more likely to prefer a dominated contract farther in advance of the effort period.

An easy implication of this observation is that the principal can use the incentive scheme to increase productivity and utility of the period 0 self. Consider the case of a risk netural agent where u(w) is linear. In this case the optimal incentive scheme for a time consistent agent is L=0 and H =1, that is, the agent owns the output.

[11] Note that paydays should lead to production increases among all time inconsistent workers, whereas only time inconsistent sophisticates should demand commitment. In a model that incorporates naivete, the naifs would attenuate the correlation towards zero.

[12] In a quasi-hyperbolic model, this prediction would be more trivial: there will be demand for commitment for tomorrow but no demand for commitment today. The horizon of commitment would not otherwise since the model is exponential except for today versus tomorrow.

[13] Our stylized model predicts that workers will only demand targets in advance of the effort period. However, the continuity of time means that even in the morning of the workday, workers may value targets—for example, because the morning self wants the afternoon self to work hard. Thus, our prediction is simply that there will be *less* demand for the dominated contract in the morning than the evening, not that there will be none.

From before, we know that for a time inconsistent agent, less effort will be realized under this scheme than 0 would like. Following the previous notation let e* be the effort when the agent owns the output. Suppose the firm perturbs this scheme by x so that:

$$L = -\frac{x}{p(e_1^*)}$$

$$H = \frac{x}{(1 - p(e_1^*)}$$

Since the probability of the low state is $p(e_1^*)$ this merely increases incentives without producing an income effect (for constant effort). The incentive effect on period 1 is positive. As a result, we can see that this scheme generates positive net benefits to both firm and the worker at time 0. The worker at time 0 would prefer it and the firm would weakly prefer it (this can be made a strong preference by letting the firm share in some of the increased revenues). This leads to the following implication:

Prediction 5: A time consistent agent has maximum productivity when she owns the output. A time inconsistent agent, however, will produce more with an employer providing a different incentive scheme than the agent simply owning the output.

In traditional agency theory, the firm provides insurance. Here we can see the firm can increase productivity as well.

We do not directly test this implication but we see it indirectly when we examine nonlinear (dominated) incentive contracts and examine impacts on productivity and workers' willingness to choose them. The derivation of this prediction implicitly shows another prediction. Time inconsistency generates incentive schemes that give super-normal incentives: workers will be incentivized at a rate that is larger than the actual impact of effort on output. [14]

III. Experiment Design

III.A. Experimental Context

[14] Both these predictions are seen in O'Donoghue and Rabin (1996). The first is implicit in their Propositions 2 and 3 where the employer is able to increase probability of task completion. The second is a direct consequence of employers in their model willing to give sharper penalties for missing a deadline than project value.

We test these predictions within an Indian data entry firm in the city of Mysore, located in one of the country's major data entry hubs. In this firm, workers use data entry software to type information from scanned images into fields on their computer screen. To control for quality, we measured accuracy using dual entry of data, with manual checks by separate quality control staff when there were discrepancies. These are standard practices in the data entry industry. Workers were paid piece rates for production as a function of how many accurate fields they entered every day. The specific piece rate schedule depended on the contract assignment (see Section IV.C.) but all contracts were functions of accurate fields entered. At all times, the screen showed the worker total fields entered so far that day.

In this context, the *physical* production function itself is completely individualistic. This means that incentive schemes do not need to be concerned about production externalities.[15] Moreover, the incentives to work came primarily from the piece rate contracts. There were no penalties for being late or leaving early. There were also few reputational concerns or potential for promotion to drive effort since workers were hired for a fixed duration job as is common in this industry. Of course there may be some residual career concerns and we discuss how this would affect our findings in Section VIII.

Employees were recruited through the standard procedures used by the firm with which we worked—from the pool of resumes submitted by walk-ins to the firm and solicitations via posters and announcements in surrounding villages. Applicants were required to have completed tenth grade education and be at least eighteen years of age. Employees were hired in order of application. Upon joining the firm, workers received about 2 weeks of training. This included technical instruction on the data entry software, the production task, and other aspects of computer usage. They were also trained on the two types of incentive contracts and the four contract treatments. During the initial part of training, workers were paid a flat stipend of Rs. 100/day while they learned the task. Trainees then worked under assignment to the control contract. At the end of the training period, they were assigned to the dominated contract for two days under the low and medium targets, respectively. This gave them the opportunity to observe

[15] This does not rule out social externalities in production. See Kaur et al. (2010; 2011) for an analysis of those in this context.

their production under both types of incentive schemes before beginning contract randomizations.[16]

III.B. Treatments

To test prediction 1, we randomized employees into three payday groups—Tuesday, Thursday, and Saturday. One-third of workers were assigned to each group at the beginning of the study, and these assignments determined which day of the week each worker received her full weekly pay. For example, on Tuesday evening of each week, employees in the Tuesday payday group were paid for work completed since the previous Wednesday. Randomly assigning paydays removes other reasons that specific days might impact effort. For example, workers might work less hard on Fridays since they'd like to enjoy their Friday evening out with friends. Alternatively they might work harder on Monday after a weekend's rest. In this design, the same day is a payday for some (randomly chosen) workers and not for others. As a result, we can identify the effect of aligning compensation with effort by comparing production on paydays with production on non-paydays.[17]

To test Prediction 2, the demand for dominated incentive schemes, we focus on two types of contracts. The first is a linear "control" contract that paid a piece rate wage of w for each unit of production. The second is a nonlinear "dominated" contract that imposed a production target. Under this latter contract, workers received the piece rate of w if they met the target, but only received $w/2$ for each entered field if they fell short of the target. As shown in Figure 1, the control contract dominates the treatment contract in earnings—for any given production level, earnings are *always* weakly higher under the control contract.

Every day one quarter of employees were placed in *Control* and simply received the linear contract. One quarter were placed in *Target* and were assigned to the dominated contract, with an exogenously chosen production target. The imposed target was selected from three target

[16] The training period for some workers (particularly those that were the first to joined the project) lasted longer than 2 weeks. However, the structure of the training remained the same, regardless of duration.

[17] In a quasi-hyperbolic model, where time horizon is defined by a day, one can make a sharper prediction: production will only be higher on the payday itself. Of course if the time horizon is longer than a day or if discounting is hyperbolic (and not quasi-hyperbolic) this sharp prediction will fail.

levels—low, medium, or high.[18] One quarter of workers were assigned to *Evening choice,* in which they chose their own target the evening before (for the next day). They could always choose a target of 0 (and many did), which is the equivalent of choosing the simple linear contract. One quarter were placed in *Morning choice* and chose their targets in the morning when they arrived to work. To make the information workers have similar across these conditions, all workers were told their treatment assignment for each day the evening before.

We randomized contract treatments daily at the individual level. We used a balanced design, where 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.[19] This ensured that each of the four treatment cells had an equal number of observations, both within each worker and across the entire sample. The vector of assignments was independent across workers. As an example, Appendix Table 1 displays the contract assignments for 5 workers in the sample over a 24-day period. Daily variation enables us to generate a large number of data points within the study period. This will give us sufficient power for a richer set of analyses, such as looking at trends in behavior over time.[20,21]

In addition to the above predictions, our design provides the opportunity to test whether heterogeneity in treatment effects is predicted by correlates of self-control that have been posed elsewhere in the literature. These include preference reversals in estimated discount rates; workers' own assessments of their self-control problems; failed attempts at quitting addictive behaviors; and measures of ability such as productivity, education and IQ. Appendix A provides additional details about the experiment context and protocols.

[18] For about half the randomization period (mostly in the first half of the study), the Target Assignment treatment consisted of assignment to "low" or "medium" targets only. Assignment to the "high" target was added later, as worker production levels increased. These target levels are explained in greater detail in Appendix A.

[19] During the period when the Target Assignment treatment consisted of assignment to only low or medium targets, randomizations were on an 8-day cycle: 2 Control; 2 Target Assignment (1 low and 1 medium target); 2 Evening Choice; and 2 Morning Choice. During the period when the Target Assignment treatment consisted of assignment to low, medium, and high targets, randomizations were on a 12-day cycle: 3 Control; 3 Target Assignment (1 low, 1 medium, and 1 high target); 3 Evening Choice; and 3 Morning Choice. The proportional weight on each of the four treatments therefore remained 25% at all times.

[20] A drawback of daily randomization is the potential for bias from inter-temporal substitution in effort. In section IV.B below, we empirically test for this concern.

[21] Note that daily randomization reduces the likelihood of potential Hawthorne effects. It is unlikely that workers would persistently alter their behavior each day for a year in response to knowledge of their treatment status.

IV. Results

IV.A. Summary Statistics

Panel A of Table 1 displays participants' characteristics. Most workers were males (74%). 63 workers reported their age on resumes or elsewhere in their application. These workers ranged in age from 19 to 38 years, with a mean of 24. We collected information on educational attainment and experience during baseline surveys administered to 101 of the 111 workers.[22] Employees had 13 years of education on average. The majority had taken a computer course and had an email address prior to joining the firm.

The key outcomes of interest are worker output and demand for dominated contracts. As defined above, output is measured as the number of accurate fields entered in a day. We have 2 measures of demand for dominated contracts under the Choice treatments—an indicator for whether a positive target was chosen and the target level chosen.

Table 2 reports summary statistics for measures of these outcomes. Column (1) provides means and standard deviations for the 8-month period during which the contract randomizations were run. This constitutes the main Analysis Sample—when both contract and payday treatments occurred simultaneously—and is comprised of 102 workers and 8,423 observations. Attendance was 88 percent and mean production conditional on attendance was 6,094 accurate fields per day in the Analysis sample. This amounts to mean daily earnings of Rs. 183 (conditional on attendance). Column (2) reports statistics for the entire 11-month period during which the payday randomizations were run. This constitutes the Full Payday Sample and is comprised of 111 workers and 11,744 observations. While mean attendance was the same in the full payday period, mean production was somewhat higher (6,433 fields). This difference stems from the fact that the payday period ran for 3 additional months, and therefore reflects production increases over time by workers. For consistency of analysis, we will use the Analysis Sample throughout the empirical analysis that follows. To demonstrate that restricting analysis to this sample does not impact the payday results, all tables pertaining to the payday treatments also

[22] In this and other information presented in Table 2, some of the employees that were hired in later stages of the project were not surveyed because of clerical oversight.

contain a column showing regression estimates for the Full Payday Sample. Appendix Table 2 shows that the 8,673 observations in the Analysis Sample are spread evenly across the treatment cells due to the balanced randomization design; the minor differences in observations in each cell are caused by worker turnover—vacancy time until worker replacement, variation in start day from first payday, and the random order of contract assignments.

IV.B. Payday Effects on Production (Test 1)

We now turn to our first prediction: productivity should increase as the payday increases. We estimate an OLS model of the form:

$$Y_{i,t} = \alpha_0 + \alpha_1 Payday_{i,t} + \boldsymbol{\alpha_2 W}_i + \boldsymbol{\alpha_3 D}_t + \boldsymbol{\alpha_4 S}_{i,t} + \mu_{i,t} \qquad (1)$$

$Payday_{i,t}$ is an indicator for whether worker *i* had a payday on date *t*. $\mathbf{W}_i$, $\mathbf{D}_t$, and $\mathbf{S}_{i,t}$ are vectors of worker, date, and seating assignment dummies, respectively, and are included to increase the precision of the payday coefficient estimate. $\mu_{i,t}$ is the residual error. $Y_{i,t}$ is the number of accurate fields entered by worker *i* on date *t*. When workers are absent, this variable is coded as 0.[23] Random assignment of paydays ensures that: $E[Payday_{i,t}\mu_{i,t}]=0$. α_1 is therefore interpretable as the causal impact of paydays on production.

Columns (1)-(4) of Table 3 provide estimates of the payday effect on output in the Analysis Sample. Column (1) estimates the specification in regression model (1). It shows that workers produce 215 fields more on average on paydays than non-paydays (significant at the 1% level). There is persistent serial correlation in output, which we control for in column (2), and continue to find a positive and significant effect of the payday. Average output is roughly 5,300 fields. These coefficients (and the rest in the Table) suggest a treatment effect of 2.6% to 10% increased production depending on the specification and whether we are comparing paydays to all days or to the day furthest from the payday.

In columns (3) and (4) we examine in detail the dynamics of the payday cycle. Instead of a single dummy for the payday, we include dummies for each of the days leading up to the payday.

[23] Simply dropping absentee observations from the sample would produce selection and could bias estimates. In this context, the zero assignment has an economic interpretation: it also corresponds to the workers' earnings that day.

$$Y_{i,t} = \alpha_0 + \sum_k \alpha^k_1 \text{Payday}^{-k}_{i,t} + \boldsymbol{\alpha_2 W}_i + \boldsymbol{\alpha_3 D}_t + \boldsymbol{\alpha_4 S}_{i,t} + \alpha_5 Y_{i,t-1} + \alpha_6 Y_{i,t-2} + \mu_{i,t} \tag{2}$$

where $\text{Payday}^{-k}_{i,t}$ are dummies indicating that the payday is *k* days away and $Y_{i,t-1}$ and $Y_{i,t-2}$ are lag production controls.[24] Figure 2 graphs the regression coefficients from column (4) of Table 5. The days that immediately follow a payday, and are therefore furthest from the next payday, are when employees are least productive. Production then rises steadily through the pay cycle.[25] This dynamic is interesting since it suggests that quasi-hyperbolic models (Laibson 1996) do not fit our data well. If the discounting happens only at the level of a day (or a more frequent unit), then we would expect the effects to be concentrated on the day of the payday. Instead we see a steady increase.

Since our output measures includes both attendance and productivity while at work, columns (5) to (8) separately examine the impact on productivity using both the Analysis (columns 5 to 8) and the Full Sample (9 and 10) and varying the inclusion of controls. All are linear probability models except for column 8, which uses a probit estimator. Unsurprisingly, people are more likely to show up to work on the payday by 4.8 percentage points; if nothing else, the benefit of collecting one's pay should increase attendance. In columns (6) and (7), however, we see this "check collection" motive is an incomplete explanation: attendance increases steadily in the days before as well. This is consistent with our model since the benefits of showing up at work will appear sooner when the payday is closer. This effect is magnified by the fact that the workers know they will also earn more because they work harder once at work.

These payday effects might imply that paying people daily will increase total output. But in what we have shown some (or all) of the increased effort on paydays may reflect a substitution of effort from other days. Such substitution will happen whenever the cost of effort is not

[24] We will look at earnings as the primary outcome. In Appendix Table 3 we show how our results decompose into hours and productivity conditional on hours worked.

[25] Figure 2 shows that production steadily increases as the payday approaches. It does not, however, pin down the shape of this increase—one could fit a linear, convex, or concave curve through the confidence intervals. Some time inconsistency models predict convexity. At the extreme, β-δ and dual-self models where time periods are defined as days would predict that all the effect should be concentrated on the payday itself; our results would seem to refute this special case. However, this is not necessarily true under alternate ways of conceptualizing time periods or the horizon of β and δ. More generally, predictions in any model will be sensitive to how the discount function is defined and how time periods are specified. As discussed above, we are interested in testing core predictions that are common across self-control models. We therefore focus here on the qualitative prediction that production should be higher closer to paydays—a result that bears out strongly in the data.

separable across days; for example, workers may be tired from prior days' work effort or may be less motivated if they had high income recently. In Section IV.B below, we test directly for these substitution effects. Our results suggest that in fact effort across days is independent. While this hints that more frequent payments would increase outputs (that the payday effects reflect net new output), one should keep in mind that we have no direct experimental data on payment frequency. Even if this were true, there might be several reasons for not paying workers daily. In our particular case, for example, there are costs of making more frequent payment that are substantive. Cash management poses significant costs to the employer. On a different note, even if more frequent payment solves the work self control problem, it may exacerbate the *consumption* self control problem. Specifically infrequent payment may be an implicit savings commitment device.

Calibration of Effect Sizes

How can we understand these effect sizes? To calibrate these effects, throughout the paper we will compare the production increases to two other values. The first calibration is straightforward. An OLS estimate in our data finds a return to an additional year of education of 501 fields, or 9% of mean production. Calibrated in this way, these results suggest that the timing of pay alone has the equivalent of 29% to 112% of the impact of a year of education.

A more theoretically insightful calibration would be to use our payday coefficients to calibrate the implied discount rate. Let e_P be the effort on payday and e_N be the effort on non-paydays; P refer to the payday time index and N be the index in time for the non-payday. Write $\pi'(e)$ to be the marginal return to the worker of a unit of effort. In the model above $\pi'(e)=p'(e)[u'(H)-u'(L)]$. Since the subjects face a continuous output and a linear pay scheme the original model does not match this context. Since we have a linear pay scheme for the experiment we can model $\pi'(e)$ as $u'()k$ where k is the coefficient on the linear pay scheme.

From the first order condition above we can write:

$$d(P)\pi'(e_P) = c'(e_P)$$
$$d(N)\pi'(e_N) = c'(e_N)$$

Assuming a linear approximation for utility of money, for small changes we can write:

$$[d(P)-d(N)]\pi'(e_P) \approx c'(e_P)-c'(e_N)$$

Our data only gives us the *output* impact of the payday. This formula suggests that converting the output impact into an estimate of the discount rate *(d(P)*) requires knowledge about the marginal cost of effort. Producing this requires another treatment. The easiest way to do this would be to use variation in the linear pay scheme. For this reason, at the end of the study (after contract randomizations were finished), we randomized workers into two piece rate wages: Rs. 3 (their usual wage) and Rs. 4 per 100 accurate fields. Each worker received each piece rate 5 times over a 10-day period in random order, with approximately half the workers in the office assigned to each wage within any given day. If we write these two pay schemes as $\pi_1(e)$ and $\pi_2(e)$ we can write:

$$\begin{aligned}
&d(N)\pi_1'(e_1) = c'(e_1)\\
&d(N)\pi_2'(e_2) = c'(e_2)\\
&\Rightarrow d(N)[\pi_2'(e_2)-\pi_1'(e_1)] = c'(e_2)-c'(e_1)
\end{aligned}$$

We find that the higher piece rate increases productivity by 11%. Thus a 33% increase in wages increases output by 11% or an elasticity of 0.33. By assuming a constant elasticity, we can use the change in pay impact to calibrate the output change due to the payday (or any intervention). Assuming e_2 is chosen in this way we can write:

$$\begin{aligned}
&[d(P)-d(N)]\pi'(e_P) \approx d(N)[\pi_2'(e_2)-\pi_1'(e_1)] \Rightarrow\\
&\frac{[d(P)-d(N)]}{d(N)} \approx \frac{\pi_2'(e_2)-\pi_1'(e_1)}{\pi'(e_1)}
\end{aligned}$$

This tells us that the percentage change in discount rate can be estimated by the percentage change in slope needed to get the same output effect as the payday scheme. In Table 5, the production difference between paydays and the beginning of the pay cycle is 428 fields, or 8%. Thus, the production increase on paydays is roughly comparable to the impact of raising the piece rate by 24% between the beginning and end of the pay week. This suggests that the discount factor changes 24% between non-paydays and paydays. If we use the full estimation in Figure 1, this suggests that relative to day 0 (the day after payday), the discount factor rises by approximately 5% each day over the pay week cycle.

As a whole, these results are consistent with Prediction 1 above. Might they be driven by other factors as well? First, perhaps the results reflect natural impatience as alluded to earlier. Since payments are made weekly, the maximum gap between paydays and non-paydays is about 6 days and the average gap is 3 days. Could a reasonable discount factor produce the increase in production of 8% on paydays? Recall that for a time consistent individual, the marginal cost of effort equals δ^T times the marginal benefit of effort. We know from the calibration above, that the payday increase is equivalent to a 24% increase in the piece rate or the marginal benefit of effort. This suggests that $\delta^3 \approx 1/(1.24) \approx .81$. This implies a *daily* discount rate of 7 percent, which is highly implausible. This calculus also rules out variants of impatience, such as credit constraints. Second, suppose that there are transaction costs to simply showing up for work. Workers may then appear to have higher output on paydays simply because they are more likely to attend. We cannot address this problem by simply examining productivity conditional on attendance (as opposed to total output) on paydays since those who attend may be a selected sample.[26] The results in column (6), however, provide direct evidence that the `show up to collect the paycheck effect' does not drive our results: we find increased productivity in the days *before* the payday. These pre-payday effects could not be explained by a transaction cost of attendance. Third, perhaps the payday effect is not about higher production on the payday but lower production on the day after. Perhaps people take the day after off. Again the dynamics belie this explanation: we find an increase in the days leading up to the payday and the payday effect is neither concentrated on the payday nor the day after. Finally, and perhaps most importantly, suppose that consumption and work effort are complements. People might work harder on paydays because they look forward to consumption that day. A night out with friends might be more enjoyable after a hard day at work. Once again, it is hard to see how this would explain the pre-payday effects; one would require the complementarities with consumption to depend (in a decreasing way) on prior days' work effort. Finally, rather than looking to rule out

[26] In Appendix Table 4, we show regression results where the dependent variable is production conditional on attendance. However, these results are difficult to interpret. Since attendance is higher closer to paydays due to the self-control benefits of paydays, the composition of workers in attendance is correlated with distance from payday. For example, when workers face negative productivity shocks like sickness, they may be more likely to come into work on paydays than non-paydays (i.e. the higher cost of effort is more likely to be justified when benefits of effort are more immediate). Alternately, if low ability workers also have greater self-control problems, the mean ability of the worker pool will be lower on paydays. Such selection problems could lead average production, conditional on attendance to be *lower* on paydays than non-paydays because different groups of workers are being compared across days. This also applies to columns (5)-(8) of Appendix Table 3, where we condition outcomes on attendance.

explanations of the payday effect *in isolation*, one should examine how they fit the full set of facts, a point we return to in Section V.[27]

Finally, testing Prediction 3 involves correlating heterogeneity in the payday effect with the demand for dominated contracts. For this test to have any validity, there must be heterogeneity in payday effects. We test directly for this by interacting a worker fixed effect with the payday dummy:

$$Y_{i,t} = \phi_0 + \phi_1 Payday_{i,t} + \boldsymbol{\phi_2 W_i} + \boldsymbol{\phi_3} Payday_{i,t}*\mathbf{W}_i + \boldsymbol{\phi_4 D_t} + \boldsymbol{\phi_5 S_{i,t}} + \phi_6 Y_{i,t-1} + \phi_7 Y_{i,t-2} + \mu_{i,t} \quad (3)$$

$Payday_{i,t}*\mathbf{W}_i$ is a vector of interactions between the vector of worker dummies and the payday dummy and all other variables are as defined above. The p-value of the F-test of joint significance of the coefficients in $\boldsymbol{\phi_3}$ is 0.000, suggesting significant heterogeneity in the payday effect.

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

For the contract treatments, we first analyze take-up of the dominated contracts. Take-up can be measured in two ways. First, we focus on days when the worker was present both the day before and the day of the Choice assignment. Absent workers would not have a chance to choose for that day in *Morning* Choice or the day after in *Evening* choice. Measured in this way 35% of the Analysis Sample chooses a positive target (see Appendix Table 5). Since this constitutes a selected sample, we also analyze take-up for all 4,193 Choice observations. We define target choice to be 0 if a worker was absent the day of Choice assignment. This is sensible if we think that not showing up to work indicates a preference for a 0 production level. For consistency, we also define target choice to be 0 if a worker was absent the day before Choice assignment, and therefore did not receive notice of assignment as per protocols and could not select a target if assigned to Evening Choice. These conventions provide a lower bound on the level of demand

[27] We find another piece of evidence consistent with our model. In India, festivals involve large expenditures by households (Banerjee and Duflo 2006). Under convex effort costs, time consistent workers should not show large production spikes in the days leading up to festivals (which are perfectly foreseeable); time inconsistent workers, however, would be expected to show such spikes. Indeed, we find that average production increases by 15% in the week prior to major festivals (significant at 1%).

for dominated contracts by workers. Under this definition, the take-up rate across observations is 28%, as is the mean of the workers' take-up rates. Figure 3 plots a histogram and kernel density estimate of worker take-up rates. The figure reveals substantial variation in demand for targets. Some workers (16% of the sample) always chose a target of 0. The bottom quarter of the distribution chose positive targets less than 10% of the time. The top quarter of the distribution chose positive targets at least 60% of the time.

The average target chosen is 974 fields. To interpret this, we compute the probability that a worker would have failed to meet the chosen target *if her output distribution matched the distribution under control*: this counterfactual represents what the period *1* self would produce if not tied to a positive target by the period *0* self. Specifically, 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 of the 1,168 observations in which a worker was assigned to Choice, 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 vector of 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 provides an estimate of the probability that the worker would have missed her chosen target under the control contract.

The mean probability of workers who would have missed their chosen targets (if they had been assigned to the control contract) is 8.3% across the 1,168 observations (Row 1 of Appendix Table 6). The mean of the worker averages for this statistic is 8.8%. Figure 6 presents the distribution of worker averages; it shows considerable variation among workers in the aggressiveness of their selected targets. For about 60% of workers, the mean target miss probability was 5% or less. The top quarter of the distribution selected targets associated with 15% or higher probability of incurring the penalty. In Row 2 of Appendix Table 6, we show the proportion of times workers actually failed to reach their targets under Choice (conditional on choosing a positive target). The mean worker missed her chosen targets 2.5% of the time.

When deciding on a target level, time inconsistent workers would weigh the motivational benefits to their future selves against the probability of incurring these costs. The penalty for missing one's target is substantial: half of one's piece rate earnings for the day. If shocks generate uncertainty in output (see Section IV.D), then choosing overly aggressive targets can be extremely costly—either due to the financial penalty, or from having to achieve the target even on days when the cost of effort turns out to be high. One might think that the model would give guidance as to the magnitude of targets we might expect. But the marginal cost of effort is unobservable, making it hard to make these predictions. For example, if the penalty is large enough *ex post* workers will always stretch to meet the goal. But *ex ante* some of this stretching is inefficient.

Testing whether these targets are binding, therefore, requires directly examining their impact on production. We estimate:

$$Y_{i,t} = \lambda_0 + \lambda_1 \text{Choice}_{i,t} + \lambda_2 \text{Target}_{i,t} + \boldsymbol{\lambda_3 W}_i + \boldsymbol{\lambda_4 D}_t + \boldsymbol{\lambda_5 S}_{i,t} + \lambda_6 Y_{i,t-1} + \lambda_7 Y_{i,t-2} + \mu_{i,t} \qquad (4)$$

$\text{Choice}_{i,t}$ is an indicator for whether worker *i* was assigned to one of the Choice treatments on day *t*; $\text{Target}_{i,t}$ is an indicator for the Target Assignment treatment. As before, $Y_{i,t}$ measures production, $\mathbf{W}_i$, $\mathbf{D}_t$, and $\mathbf{S}_{i,t}$ are vectors of worker, date, and seating assignment dummies, respectively, and $Y_{i,t-1}$ and $Y_{i,t-2}$ are controls for lagged production. Due to random assignment, $E[\text{Choice}_{i,t}\mu_{i,t}] = E[\text{Target}_{i,t}\mu_{i,t}] = 0$. The key coefficient of interest is λ_1. It represents the ITT estimate of giving workers the option to take-up the dominated contract.

In columns (1)-(4) of Table 4, we estimate variants of the above regression model for the Analysis Sample. Column (1) shows *Choice* increased production by 111 fields on average (2% of mean production). This effect is significant at the 10% level. Being assigned to the low target did not significantly increase production in the sample overall. In contrast, assignment to the medium and high targets led to average production increases of 213 fields (4% of mean production) and 335 fields (6% of mean production) respectively. These estimates are significant at the 5% level.

Column (2) separately estimates *Evening* and *Morning* choice: *Evening* choice increases output by 150 fields (3% of mean production, significant at the 5% level); *Morning Choice* increases output by 73 fields, but is not significant. Columns (3) and (4) limit the analysis to those observations in which workers were assigned to the Control or Choice treatments—the Target Assignment observations are excluded—and show similar results.

Columns (5)-(8) examine the impact of contract assignments on attendance using a linear probability model. Unlike the payday treatment, which led to substantial attendance increases, the contract treatments do not appear to impact whether employees show up to work on average.

The above estimates show that the Choice treatments increase production by about 2%, implying a local average treatment effect of approximately 6%. Using the piece rate treatment calibration logic from above, we can lower bound the time inconsistency at 18%—that is, the difference in discounting of benefits between the self that chooses the contract and the one that works to be at least 18%. This is a lower bound since the dominated contract may not achieve optimal effort (from the perspective of today's self).

We can also use the hourly data (where hours are defined as calendar hours (8-9 am, 9-10 am, etc) to paint a qualitative picture of how production changes as workers approach (and surpass) their targets. Of course since targets are set endogenously, these numbers cannot be interpreted causally but are informative nonetheless. In column (1) of Appendix Table 7, we regress hourly production on a dummy that equals 1 for the calendar hour in which the worker hit her target that day and 0 otherwise, along with worker, date, and computer fixed effects.[28] In columns (2) and (3) we we add a series of dummies for the hours immediately before and after the target was hit; the omitted category is 4 or more hours before the target was reached. These coefficients are graphed in Panel A of Figure 4. In all specifications, we see that output increases as workers

[28] Production is defined as 0 in hours when workers did not work. In cases where a worker did not manage to reach her target during the day, the hour when she reached it is coded at 7 pm (which is when the office closed for the day).

reach the target and then falls off. We can reject that production in the hour when the target was reached equals production in the hour after it was reached at the 1% significance level.[29,30]

As discussed in the payday results above, our interpretation of treatment effects estimates relies on the assumption that labor supply is separable across periods. In the case of the contract treatments, if hard work under targets increases the cost of effort on future days, this would change our interpretation of the estimates in Table 4. To test for this concern, we exploit the random ordering of contract treatment assignment. If effort costs are not independent across periods, then today's production should be lower if the worker was assigned to a high effort treatment last period (such as choice or assignment to the high target). In Appendix Table 8, we explore various specifications in line with this approach and find no evidence that there is dependence in effort between periods. For example, in column (1), we regress production on dummies for yesterday's contract assignment. We cannot reject that being assigned to choice or a target (relative to being assigned to control) has no impact on the next day's production. In fact, the coefficient estimates in all specifications are usually positive (though insignificant). In contrast, if there were intertemporal substitution, these coefficients should be negative. The positive (though insignificant) coefficient could itself be of interest. It could reflect some type of habit-formation in which working becomes *easier* if you have been working recently.

Another potential concern arises from the block randomization design of the contract treatments. Since workers are assigned to each treatment a fixed number of times within each 8- or 12-day period, treatment assignment on a given day is correlated with the probability of future treatments within each block. For example, conditional on receiving Choice today, a worker is less likely to receive Choice (and more likely to receive the Control contract)

[29] Note that there are compositional issues in the sample of which the distance from the target dummies are estimated. For example, workers who reach their targets at 5 pm are never observed 3 hours after the target is reached—the value of the 3 hours after target dummy will always be 0 in these cases. This raises selection problems that make it difficult to interpret the coefficients on the distance dummies. The comparison in which we are most interested is between the hour when the target is reached and the hour right after. We observe workers in the hour after the target is reached in 98.7% of cases; so while the selection problem undermines the validity of this comparison, it is unlikely to drive the difference in the coefficients. In Appendix Figure 3, we graph the proportion of worker-days observed in each distance category.

[30] As a benchmark, we repeat this analysis for the Assignment to Target observations in columns (4)-(6) of Appendix Table 7 and Panel B of Figure 5. We find similar patterns in how workers behave around the point at which they hit the targets we exogenously impose on them.

tomorrow. This could induce a mechanical correlation that affects what is being captured by the coefficients on the treatment assignment dummies—a concern that would not arise under independent randomization. We test for this concern in Appendix Table 9. For each observation in the Analysis sample, we compute the probabilities of receiving each contract assignment in that worker-day; these probabilities are determined by the worker's previous assignments in that randomization block. We then directly control for these probabilities in a regression of production on the contract dummies. The results indicate that the assignment probabilities have little predictive power and their inclusion has little impact on the estimated treatment effects. This is confirmed by an F-test of joint significance of the probability controls—the test p-value is 0.45.

Finally, as we did for the payday treatment effects, we check for heterogeneity in treatment effects of Choice. Using only Control and Choice observations, we regress production on: a dummy for assignment to Choice; worker fixed effects; interactions of each worker fixed effect with the Choice dummy; and date fixed effects, computer fixed effects, and lag production controls. The p-value of the F-test of joint significance of the interaction coefficients is 0.003. We interpret this as evidence for heterogeneity in the individual treatment effects of Choice.

IV.C. Correlation Between Payday and Contract Effects (Prediction 3)

The payday and contract results each support the predictions of time inconsistency models. As noted above, we also see substantial heterogeneity among workers in the payday and contract effects. We now explore this heterogeneity by checking whether the two sets of effects are correlated—whether those that are most affected by the payday treatments are also those that select and derive the greatest benefits from the dominated contracts.

We first construct a measure of the payday effect for each worker. We define the payday differential as:

$$\text{Payday Difference}=\frac{\text{(Mean Production on Paydays)-(Mean Production on Nonpaydays)}}{\text{Mean Production in Sample}}$$

We compute this differential for each worker in the Analysis Sample, using only those observations in which the worker was assigned to the Control contract. We then define a worker as having a high payday difference if her differential is above the mean differential in the sample.[31,32]

In Table 5, we test whether workers with high payday differentials are more likely to demand the dominated contract. We use two outcome variables to measure take-up: the target level chosen (which includes targets choices of 0) and a binary indicator that equals 1 if the worker chose a positive target. Following the conventions described above, in this and future regressions, we define both take-up dependent variables as 0 on days that the worker was absent the day before or day of assignment to Choice (see Section IV.B). In each column, we regress the dependent variable on the high payday differential dummy and controls. On average, workers that are more affected by paydays select a target that is 353 fields higher and are 13.8 percentage points more likely to select a positive target. These coefficients correspond to a striking 47% and 49% of the mean target level and take-up rate, respectively, and are both significant at the 1% level.

In Table 6, we explore whether workers with high payday differentials derive more benefit from the contracts. In column (1), we provide the estimates of the average treatment effects of Assignment to Choice and Assignment to a Target for reference. In column (2), we estimate:

$$Y_{i,t} = \theta_0 + \theta_1 \mathrm{Choice}_{i,t} + \theta_2 \mathrm{Choice}_{i,t} * \mathrm{HighDiff}_{i,} + \theta_3 \mathrm{Target}_{i,t} + \theta_4 \mathrm{Target}_{i,t} * \mathrm{HighDiff}_i + \boldsymbol{\theta}_5 \mathbf{W}_i + \boldsymbol{\theta}_6 \mathbf{D}_t + \boldsymbol{\theta}_7 \mathbf{S}_{i,t} + \theta_8 Y_{i,t-1} + \theta_9 Y_{i,t-2} + \mu_{i,t} \quad (5)$$

$\mathrm{HighDiff}_i$ is the indicator for whether worker *i* has an above average payday differential. We are interested in the coefficients on the interactions—θ_2 and θ_4. If the workers that are most affected by paydays are also those that benefit the most from the dominated contracts, then these coefficients will be positive.

[31] In the results presented in Section IV.A., we also see production increases in the days leading up to paydays. We use the payday-nonpayday difference for simplicity.

[32] Note that since we can only compute this statistic for workers that were assigned to the Control contract on both paydays and non-paydays during their employment, it cannot be computed for some workers that were in the sample for shorter periods of time. This reduces our sample size for this analysis from 8,423 to 8,240 observations.

The results indicate that this is indeed the case. The average treatment effects of Choice and Target Assignment for workers with low payday differentials are statistically indistinguishable from 0. In contrast, compared to the production of low differential workers under each of the treatments, high differential workers produce about 480 fields more under Choice and Target Assignment on average. These coefficients are significant at the 1% level and their magnitudes correspond to 9% of mean production. Using our estimate of the return to education, providing high payday difference workers with simply the *option* to select targets leads to production increases comparable to a one-year increase in education. In addition, using our benchmark production elasticity of 0.33 (see Section IV.A), the local average treatment effect for high difference workers is comparable to a 64% increase in the piece rate wage.

In column (3), we explore how these treatment effects vary on paydays versus non-paydays. The payday difference statistic measures the extent to which a worker is affected by aligning the compensation period with the effort period. It constitutes an imperfect proxy for the level of a worker's self-control problem. On non-paydays, when the effort and compensation period are not aligned, time inconsistency is likely to create greater distortions on effort for workers with greater self-control problems. We therefore see that it is these workers that benefit the most from the provision of targets on non-paydays. Specifically, on non-paydays, high payday difference workers produce 735 fields (14% of mean production) more under Choice than Control. At the same time—allowing for heterogeneity in the extent to which both types of treatments help a worker overcome her time inconsistency problem—the high difference workers are the ones that are most helped by paydays. They therefore have less need for dominated contracts to solve the self-control problem on paydays than those workers for whom paydays don't produce large benefits. As a result, we see that on paydays the Choice treatment is relatively more beneficial for workers with a low payday difference. The estimated coefficients on Target assignment in column (3) tell a similar story.

In columns (4)-(6), we repeat this analysis for attendance as the dependent variable. While the average treatment effects on attendance are indistinguishable from zero, we see in column (5) that this masks substantial heterogeneity. The Choice treatment does not affect attendance of low payday difference workers, and increases the attendance of high difference workers by 5.8

percentage points. This effect is significant at 1%. If workers face a self-control problem in not just how hard they work in the office, but also in the decision to show up to work (as implied by the payday results), then this effect is consistent with a model of time inconsistency with sophisticated agents. There is a sizeable fixed cost of attendance—for example, up to a 2-hour commute in each direction. Workers that are sophisticated enough to pick targets are also sophisticated enough to know that in the absence of a target, they will be tempted to exert low effort. Consequently, they're more likely to go in when they can select targets, because they know their earnings on those days will justify paying the attendance fixed cost. This is consistent with the results in column (6), which match the trends in column (3). Choice boosts attendance for high difference workers especially on non-paydays, whereas on paydays the impact of Choice is relatively lower for high difference workers than low difference ones.[33]

In Appendix Table 11, we perform the analogous exercise for payday treatment effects—we examine whether those most impacted by the contract treatments have higher increases in effort on paydays. We find that treatment effects under Choice and Target Assignment are highly predictive of payday spikes. However, take-up of dominated contracts is not predictive of payday effects. [34]

IV.D. Morning and Evening Choice (Prediction 4)

In Table 7, we examine the impact of evening choice on take-up of dominated contracts. The OLS specification is as before:

$$\text{Take-up}_{i,t} = \delta_0 + \delta_1 \text{Eve}_{i,t} + \boldsymbol{\delta_2 W}_i + \boldsymbol{\delta_3 D}_t + \boldsymbol{\delta_4 S}_{i,t} + \mu_{i,t}. \qquad (6)$$

[33] Is the large production effect of Choice on high difference workers driven completely by the impact on attendance? In Appendix Table 10, we estimate the average treatment effects of Choice on production and attendance as 395 fields and 4.4 percentage points, respectively, for high payday difference workers. For these workers, mean production conditional on attendance is 5581. As a simple calibration, 5581*0.044 = 245 < 395. This implies that the entire effect of Choice on Production for these workers is not driven by attendance increases. In column (3), we regress production conditional on attendance on the contract treatment dummies. While the coefficients are positive and significant, they are difficult to interpret (see footnote 32).

[34] A different interaction between dominated contracts and paydays may be interesting: is there greater or lower demand for dominated contracts on paydays? One intuition suggests that self-control problems may be lower on paydays but this need not be the case: this depends on whether the payday as a motivator is a substitute or a complement for the dominated contract as a motivator. In our data we find no difference in demand for dominated contracts on paydays versus non-paydays on average.

where $Eve_{i,t}$ is an indicator that equals 1 if worker *i* was assigned to Evening Choice on date *t* and equals 0 if the worker was assigned to Morning Choice, $Take\text{-}up_{i,t}$ measures take-up of the dominated contract by worker *i* on date *t*. As before, we use two measures of take-up—the target level chosen and a binary indicator for whether a positive target was chosen. Column (1) shows no difference in take up of dominated contracts the evening before.

Why might this prediction fail? Ex post analysis and qualitative work suggest a hypothesis. In choosing the night before, workers may fear uncertainty that may affect productivity the next day, uncertainty which may be realized after showing up to work, making workers *less* likely to commit to a target the day before.

An important source of this volatility is network speed fluctuations. These 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. The wait time between images could range from one second to over five minutes. Some computers in the office were more sensitive to network fluctuations than others. We asked the office management staff to consult workers in identifying the set of computers that were perceived as more sensitive to network slowdowns. Management did not know the list would be used for this purpose.

We then tested whether the computers identified as more uncertain are indeed more sensitive to overall network fluctuations. In Appendix Table 12, we show that these computers are more sensitive and illustrate this in Appendix Figure 5. When workers arrive to the office in the morning, they receive new information on the network speed and can use this to inform their target choice. This information is especially valuable for workers on bad computers, since network shocks will greatly impact their productivity.

To test this, we estimate:

$$\text{Take-up}_{i,t} = \gamma_0 + \gamma_1 \text{Eve}_{i,t} + \gamma_2 \text{BadComputer}_{i,t} + \gamma_3 \text{Eve}_{i,t} * \text{BadComputer}_{i,t} + \boldsymbol{\gamma_4 W}_i + \boldsymbol{\gamma_5 D}_t + \mu_{i,t} \quad (8)$$

in columns (3)-(6) of Table 7. Columns (3)-(4) provide some evidence that workers assigned to bad computers are less likely to demand targets in the sample as a whole. In columns (5)-(6), we estimate the above OLS regression model. Our predictions hold strongly in these results. When assigned to a good computer, selected targets are 168 fields *higher* on average in the evening than the morning. However, when assigned to bad computers, selected targets are 82 fields *lower* on average in the evening than the morning. These are sizeable magnitudes (equivalent to 22% and 11%, respectively, of the mean target levels chosen by workers in the sample overall). The results are similar if we use our binary measure of demand. Workers on good computers are 6.6 percentage points more likely to pick a positive target in the evening than the morning. In contrast, those on bad computers are more likely to demand targets in the morning than the evening.

A second source of uncertainty faced by workers stemmed from external constraints on time. For example, as discussed in Appendix A, workers that lived in more remote areas faced long and uncertain commute times. These impacted morning arrival time and therefore how much the worker could produce in a day. In addition, some workers had duties or binding constraints on time outside the office. This made it more difficult to absorb production shocks. For example, if the network unexpectedly slowed down, it would have been harder for workers with external constraints to stay late in the office to ensure their targets were met. Much of the uncertainty from these sources was resolved by the morning of the workday—by then employees knew their arrival time at the office and would have had a better sense of duties at home for that day.

Thus, we expect workers with greater external constraints to be relatively more likely to demand the dominated contract in the morning. To test this prediction, we estimate:

$$\text{Take-up}_{i,t} = \varphi_0 + \varphi_1\text{Eve}_{i,t} + \varphi_2\text{Constraint}_i + \varphi_3\text{Eve}_{i,t}*\text{Constraint}_i + \boldsymbol{\varphi_4}\mathbf{D}_t + \boldsymbol{\varphi_5}\mathbf{S}_{i,t} + \mu_{i,t} \quad (9)$$

Constraint_i is a measure of the external constraints faced by worker *i*. In columns (7)-(10) of Table 7, we present results for two binary measures of the constraint variable. In columns (7)-(8), the variable measures workers' response to a question in the endline survey that asked them to agree or disagree 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." The high constraint indicator takes a value of 1 if the worker's response was "Agree Strongly" and 0 otherwise. The results indicate that workers with more uncertain commute times select targets more often under Morning Choice than Evening Choice, and the opposite is true for workers with less uncertain commute times. The endline survey asked four questions related to external constraints. For the analysis in Columns (9)-(10), we compute a Constraint Index for each worker by averaging his or her answers to the four questions. The high constraint indicator equals 1 if the worker's constraint index score was above the sample mean score and equals 0 otherwise. The results in columns (9) and (10) are similar to those in columns (7) and (8), respectively.

Collectively, our results suggest that when uncertainty is high, this pattern reverses and take-up is higher in the morning—after a substantive portion of the uncertainty has been resolved.[35]

IV.E. 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? We examine how workers' choices and treatment effects evolve with their experience. We define $\text{experience}_{i,t}$ as the number of workdays worker *i* has been in the Analysis Sample on date *t*.[36]

In Figure 5, we explore how demand for the dominated contracts evolved with worker experience for high and low payday difference workers. For each value of the experience variable, we compute the proportion of high difference workers that choose a positive target under assignment to Choice (and were present the day of and day before Choice assignment); these values are plotted in closed circles. The open circles plot the value of this statistic for the low difference workers. The figure shows little initial difference in mean take-up rates between high and low difference workers. However, as workers gain experience, we see a divergence. Over time, those that have the largest self-control problems (as measured by our payday

[35] Of course, if demand for motivating future selves were sufficiently strong the worker might choose some target (albeit a smaller one) even in the presence of uncertainty.

[36] Recall that days during which workers are in training are not included in the Analysis Sample. As a result, $\text{experience}_{i,t}=1$ on worker *i*'s first day of contract randomizations. Note also that the experience variable suffers from selective attrition.

difference proxy) end up demanding the dominated contract at substantially higher rates than the workers that do not have large payday impacts.[37]

We explore these trends more formally in Panel A of Table 8. In columns (1) and (2), we regress each of our measures of take-up on the log of the experience variable and our standard controls. In both columns, we cannot reject that mean demand for the dominated contract does not change over time in the sample as a whole. In columns (3) and (4), we add the high payday difference indicator and an interaction of high payday difference with log experience. The results are consistent with the trends in Figure 6. As low difference workers gain experience, they decrease their take-up of the targets; a 1% increase in experience is associated with about a 0.066 percentage point decrease in take-up (significant at 5%). In contrast, we cannot reject that the demand among high difference workers stays constant over time. The F-tests (of) whether the log experience coefficient and interaction coefficient sum to zero are insignificant (with p-values of 0.493 and 0.895 in columns (3) and (4), respectively). As a result, at higher values of experience, the high payday difference workers exhibit substantially higher demand on average for the dominated contracts than the low difference workers.

These results are consistent with a story in which workers initially try the dominated contracts. Over time, they continually receive opportunities to observe their production under targets—both through Target assignment and potentially also when on Choice assignment. Those workers for whom the targets do not yield utility benefits stop selecting the dominated contract. In contrast, the workers with large self-control problems see that the targets are helpful, and continue to select them. Consistent with Test 3 of our model, this latter group of workers correlates with the group that is most affected by paydays.

Next, in Panel B of Table 8, we test whether the treatment effects on production persist over time. For reference purposes, in column (1), we first regress production on: log experience; dummies for Choice assignment, Target assignment, and Payday; and our vector of standard controls. As before, we define experience as the number of workdays the employee has been in

[37] The figure also shows that there is variation in the level of day-to-day take-up within each group of workers over time. This is not surprising since the composition of workers assigned to Choice changes each day with the contract randomizations. In addition, day-to-day shocks (such as network speed fluctuations) impact take-up.

the Analysis Sample. Not surprisingly, we see that production increases strongly with experience. The remaining results in column (1) are consistent with those presented in earlier tables.

In column (2), we add interactions of log experience with each of the treatment variables of interest: Choice, Target assignment, and Payday. We are interested in the coefficients on the interactions. If treatment effects diminish over time—for example, once the novelty of the treatments wears off—then these coefficients will be negative. Instead, the results in column (2) reveal positive interaction coefficients. The interaction of log experience with Choice assignment is positive and significant at the 5% level. This is consistent with the findings in Panel A, which indicate that the workers that derive the largest benefits from the dominated contracts are the ones that are most likely to select them over time. In addition, the interaction with Target assignment is also positive and significant at the 10% level. The coefficient on the payday interaction is essentially 0, indicating that the payday effect is constant over time on average.

In column (3), we repeat this exercise using a different measure of experience: a binary indicator for whether the worker has been in the Analysis Sample for more than two calendar months. The coefficients Choice and Target interaction coefficients are positive (and insignificant). The interaction on the payday coefficient is now negative, but insignificant. In column (4), we check for persistence in the payday effect in the full 11-month payday sample and again find no change in effects over time. Together, columns (2) to (4) provide compelling evidence that the treatment effects of the contract treatments and paydays persist over time, and some evidence that the effects of Choice increase over time.

Overall, we see that workers select and derive steady benefits from the dominated contracts throughout the experiment. Similarly, the production increases on paydays persist week after week. Given the long horizon of the study, our results imply that time inconsistency is a perpetual problem in the workplace. They lend credence to our view that many workplace features can plausibly be interpreted as arrangements that seek to solve self-control problems.

IV.F. Heterogeneity of Treatment Effects

How well do observable worker characteristics and survey measures predict self-control? In Table 9, we use external measures to explore heterogeneity in demand for dominated contracts and in treatment effects of the contract and payday treatments. Each column conducts this analysis for a different potential correlate of ability or self-control problems; the correlate analyzed in each column is specified at the top of that column. In Panel A, we show coefficient estimates from a regression of target level chosen on the correlate and controls. In Panel B, we regress the binary take-up indicator on the correlate and controls. In Panel C, we report estimates from a regression of production on: the correlate; dummies for Choice and Payday; interactions of each of these dummies with the correlate; and controls.

In columns (1)-(3), we look for evidence on whether more able workers are differentially affected by self-control. The correlate in column (1) is a measure of worker productivity: whether the worker's mean production under assignment to the control contract is above the sample mean. In column (2), the correlate is years of education. Both these correlates positively predict dominated contract demand—for example, high productivity workers are 40% more likely to choose positive targets under Choice. Interestingly, these measures do not positively predict treatment effects of choice, potentially indicating higher levels of sophistication among higher ability workers. In columns (3), we look at a measure of IQ—the sum of the worker's scores on the Raven's Matrix and Digit Span tests. IQ does not predict take-up or contract treatment effects. In addition, none of the ability correlates predicts treatment effects of paydays.

The literature on psychology and economics has proposed a range of correlates of self-control problems. In the remaining columns, we examine the predictive power of some of these correlates, collected through the endline surveys. In columns (4)-(6), we look at measures of self-control problems based on self-reports by workers. The correlate in column (4) is the Self-control Factor, obtained from a factor analysis on the endline survey data. In column (5), we construct a Self-Control Index from the endline survey responses by averaging each worker's responses to the 9 self-control questions in the endline survey. Both the Self-control Factor and Self-control index values have been de-meaned in the analysis. In column (6), we use self-reports of addictive behaviors by male workers. In this column, the correlate equals 1 if the worker said he had tried to quit drinking, smoking, or chewing tobacco and failed, and equals 0 otherwise. Each of these three columns shows similar results. These three correlates from the endline surveys positively predict demand for the dominated contract, and also positively predict treatment effects of the contracts. However, among these, only the coefficients on the Self-control Factor are generally significant. None of the correlates predicts the payday effect.

In columns (7)-(8), we look at outcomes from the discount rate exercise, in which we asked workers to trade off cash rewards between different time horizons—a standard way of testing for self-control problems in the literature. In column (7), we look at impatience. Our measure of impatience is defined as the proportion of times in the 6 questions the worker chose the smaller immediate reward rather than the larger delayed reward. In column (8), we look at preference reversals. Here, our self-control correlate is defined as the proportion of times a worker chose the larger immediate reward in the short horizon, but then displayed patience when choosing between the same amounts in the long horizon. As in the case of the Self-control Index and Addictive behaviors, these correlates do not appear to predict demand for the dominated contracts—the coefficients in Panels A and B are positive but significant. Also as before, we see in Panel C that workers with greater self-control problems (as measured by these correlates) are less productive on average. We also see evidence that these workers benefit more from the contract treatments. For example, the coefficient on the interaction between proportion of impatient responses and Choice assignment is 706 and is significant at the 5% level. It indicates that workers that show impatience in all 6 questions produce 706 more fields (13% of mean production) under Choice than Control. As before, the correlates do not seem to predict payday behavior.

Overall, we find some support that proxies of self-control posed in the literature correlate with behavior under the contract treatments. In contrast, none of these proxies is correlated with the payday effects. Thus, while these survey-based proxies have some predictive power, we find that the strongest predictor of effects under each set of treatments is workers' behavior under the other set of treatments (see Section VII.C).

V. Alternative Explanations

The results are largely consistent with a self-control agency model. Could they be explained without appealing to self-control? We will argue that while any one result could be explained by other factors, it is hard to fit an alternative theory to the full pattern of results: the production increases on paydays; demand for dominated contracts and treatment effects of Choice; and the correlation between the payday and contract effects.

First, could workers be choosing dominated contracts because they are confused? The experiment was designed to minimize this possibility. During the training period, all workers were subjected to targets. At the end of training, we administered a quiz that tested their comprehension of the contracts. The mean score was 93%, indicating that workers understood the contracts.[38] Even after training, given the length of the experiment, workers receive (in both assignment and choice treatments) a great deal of experience with these contracts. As a result, the qualitative interviews with workers suggest they were well aware that the dominated contracts are dominated.

Second, could workers be choosing dominated contracts to signal ability to employers? It is not clear that demand for the dominated contract actually should serve as a positive signal. Since the employer observes production directly, there is no reason to believe a worker that can achieve high production under the control contract should not appear more impressive than one that needs a dominated contract to increase output. Still to mitigate this, we advertised the job as a one-time employment opportunity. Of course that might not have worked fully. The stumbling block for this—as with all the other explanations—is how it would explain both the payday effect and the correlation of dominated contract choice with the payday effect.

Third, a different psychological explanation could be that targets serve as non-monetary impact of targets. It need not be, as we have modeled them, that targets are merely *monetary* motivators. It is possible that the targets generate intrinsic motivation: the desire to hit the target

[38] In Appendix Table 13, we test whether quiz score is correlated with demand for the dominated contracts. If workers mistakenly chose dominated contracts because they did not understand the contract treatments, then we would expect quiz score to be negatively correlated with take-up. Instead, quiz performance positively predicts take-up, although the coefficients in columns (1)-(4) are insignificant. In addition, as noted above, education strongly predicts take-up of the dominated contract.

alone may motivate workers. With data such as ours, of course, one cannot separate intrinsic from extrinsic motivation generated by the target. Of course, without time inconsistency it is unclear how this would explain the payday findings. As a result, while our data cannot rule out nonmonetary motivations provided by the target, it does suggest that time inconsistency is needed in this case as well.

Finally, while it may not explain choice of dominated contracts, could income targeting explain the payday effects? If workers target a fixed weekly income level, then small amounts of impatience or the realization of shocks could lead workers to backload effort closer to the payday. We can test for such behavior directly in the data. Income targeting implies a sharp decrease in marginal utility for income levels above the weekly target (see Camerer et. al. 1997). Two pieces of evidence suggest this is not happening in our data. First, we saw in our test for intertemporal substitution, exogenous production increases caused by target assignment do not lead to production decreases on subsequent days (see Appendix Table 8). 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. We test for this in Appendix Table 14. Under income targeting, the interactions in columns (2) and (3) should be negative. However, the interactions are positive (though largely insignificant).[39]

VI. Conclusion

These results are interesting because they may help us explain workplace arrangements. Agency theory understands workplace arrangements—the existence of bosses, worker discipline and other pejorative mechanisms (even physical punishment Chwe 1990)—in one of two ways. The

[39] We also test an additional prediction of income targeting. Since the impact of day-to-day shocks is adjusted within the payweek to arrive at the weekly target, the variance in production among payweeks should be less than the variance in production among weeks defined according to some other arbitrary cycle, such as calendar weeks. To check this, we compare production across workers' payweek cycles with production across 4 artificial weekly cycles, created by shifting forward days from the worker's actual pay cycle. For example, for a worker assigned to the Saturday pay group, her true pay week is from each Monday to Saturday. The 4 artificial cycles for this worker would be from Tuesdays to Mondays, Wednesdays to Tuesdays, Thursdays to Wednesdays, and Fridays to Thursdays. For each worker, we then compute the standard deviation of weekly production across her actual payweeks and across each of the 4 associated artificial weekly cycles. On average, the standard deviation of weekly production for actual payweeks is 1838. The mean standard deviations for each of the 4 artificial cycles are *lower* than this, ranging from 1731 to 1809. Overall, all 5 standard deviation estimates are close to each other (within 5% or less from the payweek mean). This provides further support against weekly income targeting.

firm exists to provide insurance. This insurance creates moral hazard. These workplace arrangements exist to mitigate that moral hazard. The second view 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.” Joint production necessitates the need for monitoring (Alchian and Demsetz 1972).

Our results suggest a different way to understand a diverse host of workplace arrangements. Might certain contract features (such as nonlinear contracts) be thought of as partly reflecting self-control benefits? Might discipline at the workplace or workplace rules be thought of as demand for features to help workers avoid the temptation to shirk? Might the organization of production itself, such as the presence of a boss and task division with deadlines serve to mitigate self-control problems? Clark (1994) advances this interpretation of the industrial revolution for example. One of the major changes in the organization of production in economic history has been the transition from the putting out system (under which workers were paid piece rates according to work performed and could choose production levels and work hours themselves) to the more rigid workplace system that is the norm today (with features like assembly lines, production minimums, rigid work hours, and hefty punishments for even momentary lapses in behavior). One interpretation is that increases in capital since the industrial revolution place a premium on increasing labor productivity (Clark 1994). Finding ways of reducing worker self-control problems could be one response to this problem. Under this view, the rise of the factory and its associated disciplinary infrastructure was in part an attempt to solve self-control problems.

Indeed, the process of development—which entails movements from agriculture to manufacturing, or from cottage industry to factory work—may increase labor productivity not just through technological innovation, but also because the work arrangements associated with these advances mitigate workers’ self-control problems. (See Clark 1994 and Kaur et al. 2010 for further exposition of this view).

We may even need to enrich how we conceptualize the production function. Take the basic prediction that the self-control problem increases as the returns to work are further in the future. This mechanism suggests a new look at a variety of naturally occurring production function differences. In agriculture in developing countries, should we view productivity in long-horizon crops differently from productivity in short-horizon crops? Might farmers choose shorter horizon crops because effort distortions are smaller when effort and compensation are more closely aligned? Might the move from farm work to formal sector work with regular pay have self-control productivity benefits?

These arguments 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.

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): 635-672.

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 Esther Duflo. 2006. "The Economic Lives of the Poor." MIT Working Paper No. 06-29.

Banerjee, Abhijit, and Sendhil Mullainathan. 2009. "The Shape of Temptation: Implications for the Economic Lives of the Poor." Mimeo, Harvard University.

Benhabib, Jess, Alberto Bisin, and Andrew Schotter. 2007. "Present Bias, Quasi-Hyperbolic Discounting, and Fixed Costs."

Bernheim, B. Douglas, Debraj Ray and Sevin Yeltekin, 2011, "Poverty and Self-Control", working paper.

Burger, Nicholas, Gary Charness, and John Lynham. 2008. "Three Field Experiments on Procrastination and Willpower." Available: www.bec.ucla.edu/Charness_Procrastination.pdf

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.

Cheung, S.N.S., 1969. *The Theory of Share Tenancy*. University of Chicago Press, Chicago, pp: 188.

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.

Chwe, Michael S.1990. Why Were Workers Whipped? Pain in a Principal-Agent Model, *Economic Journal* 100, pps. 1109-1121.

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.

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.

Giné, Xavier & Goldberg, Jessica & Yang, Dean, 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.

Holmstron, 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. 2009. “The Behavioralist Visits the Factory: Increasing Productivity Using Simple Framing Manipulations.” NBER Working Paper No. 15623.

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

Kaur, Supreet, Michael Kremer, and Sendhil Mullainthan. 2011. “Works Well with Others: Peer Effects and Social Motivation.” Mimeo, Harvard University.

Laibson, David I. 1996. "Hyperbolic Discount Functions, Undersaving, and Savings Policy," NBER Working Paper 5635, National Bureau of Economic Research, Inc

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

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

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. 1999. "Doing It Now or Later," *American Economic Review*, 89(1): 103-124.

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

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

Figure 1
Incentive Contracts

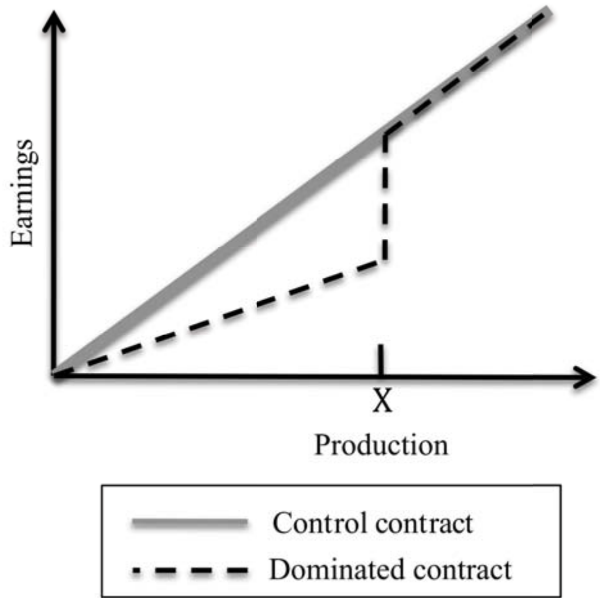

Notes: This figure displays the two types of incentive contracts offered to workers. The linear "control" contract paid a piece rate wage of *w* for each accurate field entered. The nonlinear "dominated" contract imposed a production target, X; workers were paid *w* for each accurate field if they met the target, but only received *w*/2 for each field if they fell short of the target. Thus, earnings are equivalent under both contracts for output levels above X. However, if a worker fails to achieve X, earnings are substantially less under the dominated contract.

Figure 2
Production over the Pay Cycle

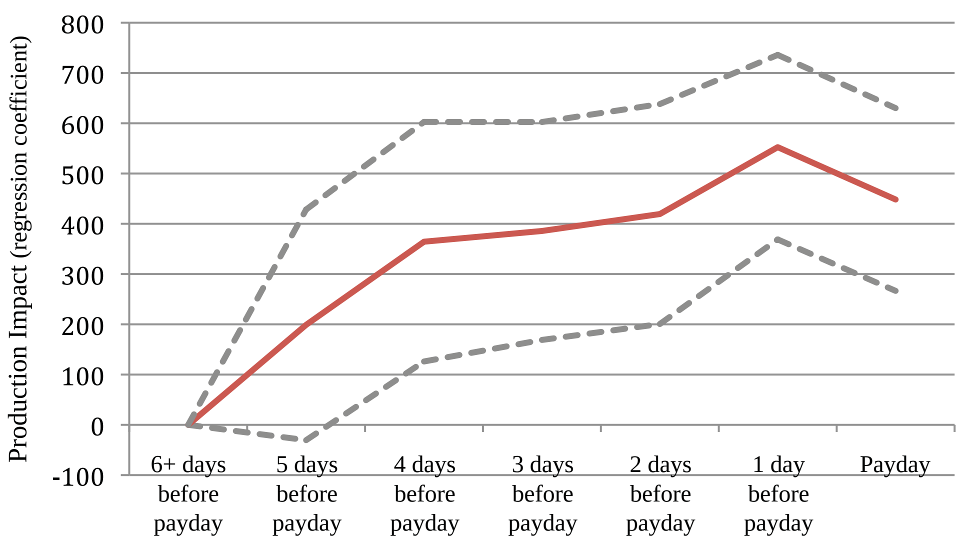

Notes: This figure graphs the coefficients and confidence intervals from a regression of production on 6 binary indicators that capture distance from a worker's next payday (payday, 1 day before payday, 2 days before payday, etc). The regression includes controls for lagged production as well worker, date, and seat assignment fixed effects. Confidence intervals are based on robust standard errors. Note these coefficients correspond to those shown in column (4) of Table 5.

Figure 3
Take-up of Dominated Contracts: Distribution of Worker Means

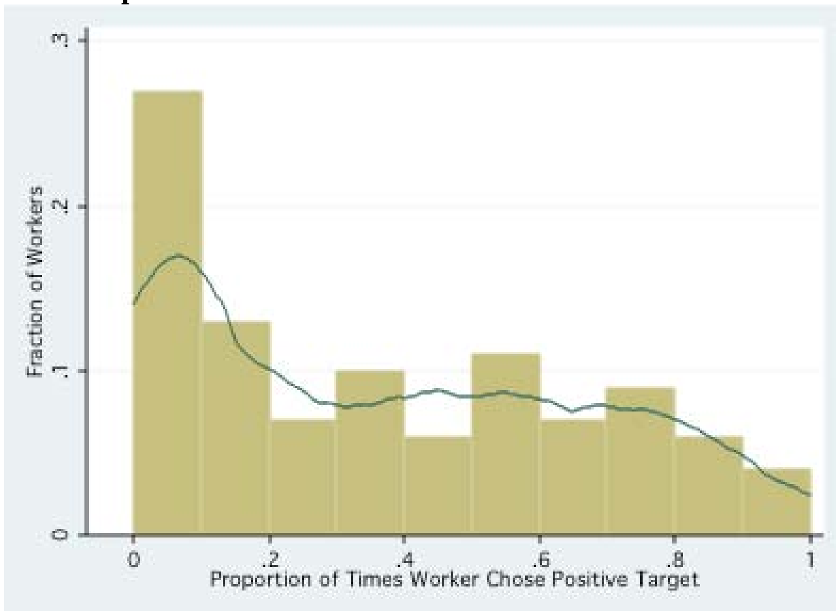

Notes: This figure shows the distribution of take-up rates of the dominated contract by workers. A worker's take-up rate is defined the proportion of times the worker selected a positive target when assigned to Choice (and was present the day before and day of Choice assignment). The distribution is shown for the 101 workers in the Analysis Sample that were assigned to Choice at least once.

Figure 4
Production Behavior Around Target Thresholds

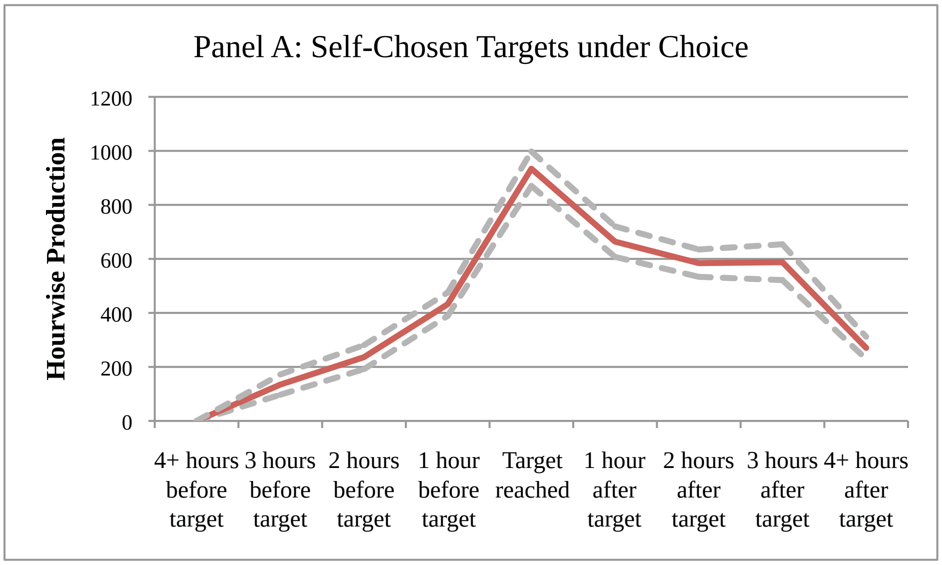

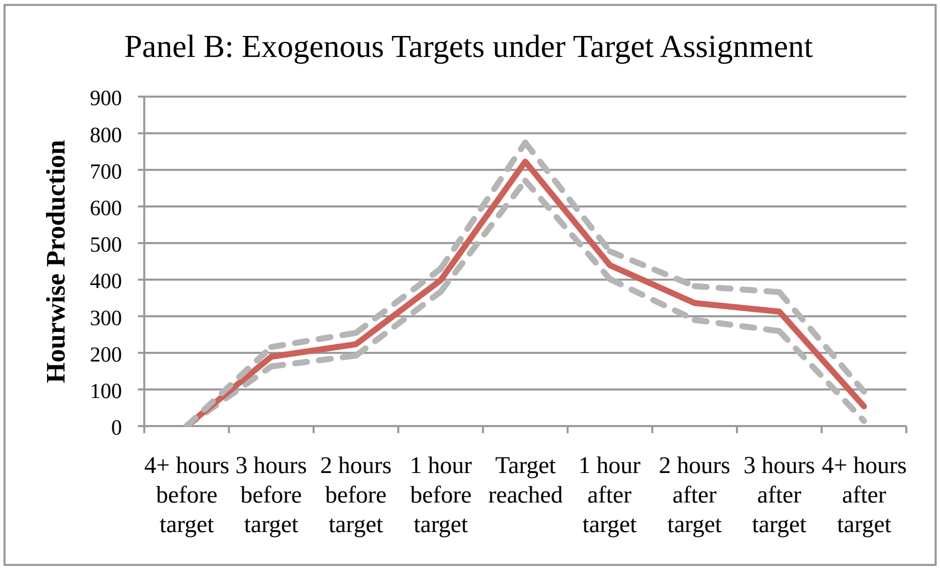

Notes: This figure graphs the coefficients and associated 95% confidence intervals from regressions of hourwise production (production by worker i on date t in calendar hour h) on binary indicators of distance from when a worker achieved her target within a day and controls for worker, date, and seat assignment fixed effects. Standard errors are corrected to allow for clustering by worker-day. Figure 1 reports estimates from days when workers selected positive targets under Assignment to Choice and were present. Figure 2 reports estimates from days when workers were assigned to exogenous targets under Target Assignment and were present. The coefficients plotted in Panels A and B correspond to those shown in columns (2) and (5), respectively, of Table 8.

Figure 5
How the Demand for Dominated Contracts Changes with Experience

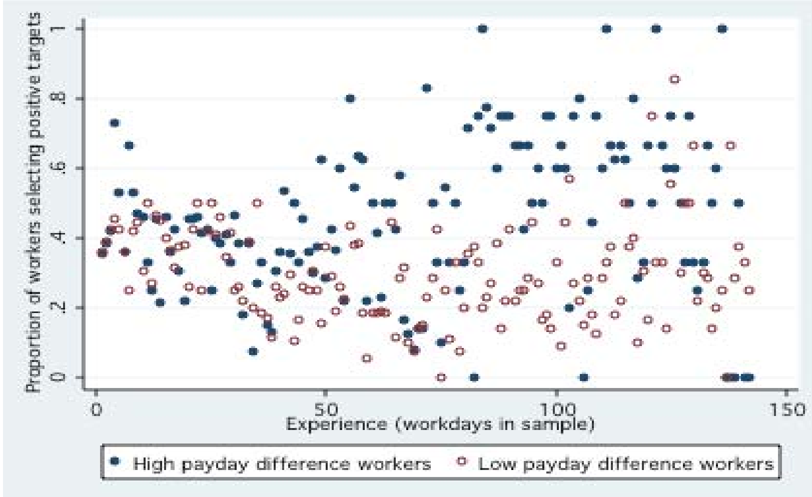

Notes: This figure shows how demand for the dominated contract evolves with worker experience. The x-axis measures worker experience, defined as the number of workdays the worker has been in the sample. The y-axis measures the proportion of workers that chose a positive target under assignment to Choice for each value of the experience variable. The closed and open circles show the value of this statistic separately for high payday difference workers and low payday difference workers, respectively. High difference workers are those for whom the mean payday production difference—the difference in production on paydays and non-paydays under control, divided by overall mean production under control—is higher than the sample average payday difference. The proportion of times positive targets were chosen is computed using Choice observations in the Analysis Sample in which the worker was present the day before and day of Choice assignment.

Table 1
Worker Characteristics and Survey Responses

	Mean	Standard Deviation	10th pctile	90th pctile	Obs
	(1)	(2)	(3)	(4)	(5)
A. Worker Characteristics					
Proportion female	0.26	0.44	--	--	111
Age	24	4	20	29	63
Years of education	13	2	11	15	101
Completed high school	0.84	0.37	--	--	101
Used computer prior to joining firm	0.67	0.47	--	--	101
Taken computer training course	0.70	0.46	--	--	101
Had email address prior to joining firm	0.60	0.49	--	--	101
B. Performance on Tests Administered During Experiment					
Contracts comprehension quiz: percentage score	93	13	80	100	79
Raven's Matrix score	35	12	17	49	107
Digit Span composite score	27	6	20	36	109
IQ composite score (Raven's Matrix plus Digit Span)	62	15	40	84	106
C. Discount Rate Measurement					
Proportion of times worker chose smaller immediate reward	0.31	0.28	0	0.67	58
Proportion of times worker displayed preference reversal	0.17	0.23	0	0.67	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
"At the end of the day, 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	3.11	4.17	70
Worker has tried to quit an addictive behavior and failed	0.12	0.33	--	--	51
Factor analysis: self-control factor	0.00	0.86	-1.37	0.72	70
E. Endline Survey: Self-Reported Measures of External Constraints					
Worker agreed or agreed strongly with the statement:					
"If I miss one bus or train, the next one I can take is much later."	0.61	0.49	--	--	70
"I do not have much flexibility in how late I can stay in the office because I have to leave at a certain time."	0.56	0.50	--	--	70
Constraints index: mean of responses to all 4 constraints questions (1=disagree strongly; 5=agree strongly)	3.61	0.85	2.25	4.50	70

Notes: This table presents summary statistics for the 111 workers that participated in the study. Panel A presents statistics on worker characteristics, gathered from a baseline survey. Panel B provides information on tests that were administered to workers: a quiz that tested their comprehension of the contract treatments; their performance on the Raven’s Matrix IQ test; their total score on the Digit Span test (administered forwards and backwards in each English and the local language); and the sum of their Raven’s Matrix and Digit Span score. Panel C describes worker behavior during a discount rate exercise in which they traded 3 sets of off 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 but showed patience by choosing the larger reward in the long horizon). Panel D summarizes responses by workers to questions during the endline survey that asked them to agree or disagree with statements relating to their self-control behavior. It also reports summary statistics for the Self-Control Factor, which is determined using a Factor Analysis on all the endline survey questions. Panel E provides details of responses to 4 endline survey questions that asked about external constraints.

Table 2
Summary Statistics of Outcome Measures

	Analysis Sample	*Full Payday Sample*
	(1)	(2)
Attendance	0.88	0.88
	(0.33)	(0.32)
Production	5337	5665
	(3404)	(3651)
Production conditional on attendance	6094	6433
	(2935)	(3193)
Indicator for whether positive target was selected under Choice treatment	0.35	--
	(0.48)	
Target level selected under Choice treatment	974	--
	(1502)	
Number of workers in sample	102	111
Number of observations in sample	8,423	11,744

Notes: This table shows summary statistics for four outcome measures: attendance rates, production conditional on attendance (where production is measured as the number of accurate fields entered by a worker in a day), take-up rates for the commitment contract (where take-up is defined as whether a worker selected a positive target when assigned to Choice), and the target levels selected by workers when assigned to Choice (this includes observations where the worker selected a target of 0). These statistics are summarized for each of 2 samples—the 11-month period during which payday treatments were run; the 8-month period during which both the contract and payday treatments were run. The table presents means for each measure and standard deviations are shown in parentheses.

Table 3
Treatment Effect of Paydays on Worker Production

	Analysis Sample								*Full Payday Sample*	
	Production	*Production*	*Production*	*Production*	*Attendance*	*Attendance*	*Attendance*	*Attendance*	*Production*	*Attendance*
	OLS	OLS	OLS	OLS	OLS Linear Probability Model	OLS Linear Probability Model	OLS Linear Probability Model	Probit Marginal Effects	OLS	OLS Linear Probability Model
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Payday	215 (70)***	140 (63)**	225 (68)***	428 (94)***	0.048 (0.009)***	0.055 (0.010)***	0.077 (0.013)***	0.065 (0.009)***	414 (85)***	0.072 (0.011)***
1 day before payday			353 (76)***	539 (95)***		0.035 (0.011)***	0.053 (0.013)***	0.043 (0.010)***	469 (85)***	0.047 (0.011)***
2 days before payday			207 (94)**	417 (113)***		0.016 (0.014)	0.037 (0.016)**	0.028 (0.013)*	394 (103)***	0.041 (0.014)***
3 days before payday				374 (112)***			0.026 (0.017)	0.019 (0.012)	412 (101)***	0.036 (0.014)***
4 days before payday				332 (123)***			0.047 (0.017)***	0.039 (0.012)***	254 (105)**	0.036 (0.014)***
5 days before payday				176 (119)			0.023 (0.017)	0.019 (0.013)	115 (107)	0.017 (0.014)
Production on previous workday		0.355 (0.016)***	0.353 (0.016)***	0.355 (0.016)***					0.349 (0.013)***	
Production from two workdays ago		0.135 (0.015)***	0.136 (0.015)***	0.137 (0.015)***					0.139 (0.012)***	
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	8423	8423	8423	8423	8423	8423	8423	7920	11744	11744
R2	0.50	0.59	0.59	0.59	0.11	0.11	0.11	0.14	0.60	0.11
Dependent variable mean	5337	5337	5337	5337	0.88	0.88	0.88	0.87	5665	0.88

Notes: This table reports the effects of paydays and distance from paydays on worker production and attendance. Production is defined as the number of accurate fields completed by a worker in a day, and equals zero on days workers are absent. Attendance is an indicator variable that takes the value of 1 if a worker is present and 0 otherwise. Payday is an indicator variable for whether the day was the worker's assigned payday. 1 day before payday to 5 days before payday are indicators for whether the current day is 1 day to 5 days away, respectively, from her assigned payday. Production on previous workday and Production from two workdays ago are lag production controls that measure worker i's production on workdates t-1 and t-2, respectively. Columns (1)-(8) show results from the Analysis Sample—the 8-month period during which both payday and contract randomizations were run. Columns (9)-(10) report results from the Payday Sample—the full 11-month period during which payday randomizations were run. All columns except column (8) show estimates from OLS regressions. Column (8) gives results from a probit regression where the dependent variable is attendance. Estimated discrete changes for each dummy variable are reported. All regressions include fixed effects for each date in the sample, each worker in the sample, and each computer seating assignment. Robust standard errors are reported in parentheses.

Table 4
Treatment Effects of Contract Assignment on Worker Production

Observations	*Dependent variable: Production*				*Dependent variable: Attendance*			
	All obs	All obs	Control & Choice	Control & Choice	All obs	All obs	Control & Choice	Control & Choice
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Assignment to choice	111		120		0.007		0.007	
	(59)*		(59)**		(0.009)		(0.009)	
Assignment to evening choice		150		156		0.01		0.014
		(69)**		(69)**		(0.01)		(0.010)
Assignment to morning choice		73		84		-0.00		0.001
		(69)		(69)		(0.01)		(0.010)
Assignment to low target	3	3			-0.002	-0.00		
	(90)	(90)			(0.013)	(0.01)		
Assignment to medium target	213	213			-0.006	-0.01		
	(91)**	(91)**			(0.013)	(0.01)		
Assignment to high target	335	334			-0.005	-0.01		
	(150)**	(150)**			(0.019)	(0.02)		
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lag production controls	Yes	Yes	Yes	Yes	No	No	No	No
Observations	8423	8423	6310	6310	8423	8423	6310	6310
R2	0.59	0.59	0.60	0.60	0.15	0.15	0.11	0.11
Dependent variable mean	5337	5337	5311	5311	0.88	0.88	0.88	0.88
Proportion choosing a positive target (conditional on attendance)	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35
Proportion choosing a positive target (target=0 when absent)	0.28	0.28	0.28	0.28	0.28	0.28	0.28	0.28

Notes: This table estimates the treatments effects of the contracts on production. The dependent variable in columns (1)-(4) is production. Production is defined as the number of accurate fields completed by the worker in a day, and equals zero on days workers are absent. The dependent variable in columns (5)-(8) is attendance. Attendance is a binary variable that takes the value of 1 when workers are present and 0 otherwise. The regressions in columns (1)-(2) and (5)-(6) include all observations in the Analysis sample. The regressions in columns (3)-(4) and (7)-(8) include observations in which workers were assigned to the Control contract or to one of the Choice treatments (Evening choice or Morning choice). Each column shows the estimates from an OLS regression of the dependent variable on indicators for contract treatment. Assignment to choice is an indicator that equals 1 if the worker was assigned to one of the Choice treatments (Evening Choice or Morning Choice) and equals 0 otherwise. All regressions include worker, date, and seat fixed effects. In addition, columns (1)-(4) also include controls for lagged production. Robust standard errors are reported in parentheses.

The bottom 2 rows present the proportion of times workers chose a positive target when assigned to a Choice treatment. The first row limits the estimate to observations in which a worker was present the day of and day before Choice assignment. The second row includes all Choice observations and codes target choice as 0 if the worker was absent the day of or day before Choice assignment. Standard deviations are reported in parentheses.

Table 5
Heterogeneity in Take-up of Dominated Contracts: Correlation with Payday Impact

Dependent variable	*Target level chosen*	*Positive target indicator*
	(1)	(2)
High payday production difference	353	0.138
	(129)***	(0.044)***
Seat fixed effects	Yes	Yes
Date fixed effects	Yes	Yes
Lag production controls	Yes	Yes
Observations	4098	4098
R2	0.22	0.20
Dependent variable mean	759	0.28

Notes: This table tests whether workers that are most affected by paydays are also more likely to demand the dominated contract. The dependent variables are two measures for take-up of the dominated contract. The first variable is the continuous measure of the target level selected by workers when assigned to Choice. The second variable is a binary indicator for whether the worker selected a target above zero when assigned to a Choice treatment. Both dependent variables are defined as 0 if a worker was absent the day before or day of Choice assignment. Each column shows results from an OLS regression of the dependent variable on an indicator for whether a worker had a high payday differential, computer fixed effects, and date fixed effects. High payday production difference is an indicator that equals 1 if the worker's payday differential—defined as the difference between mean production on paydays and non-paydays under Control divided by mean production under assignment to Control—is above the sample average. Standard errors allow for clustering by worker and are reported in parentheses.

Table 6
Heterogeneity in Contract Treatment Effects: Correlation with Payday Impact

Dependent variable	*Production*	*Production*	*Production*	*Attendance*	*Attendance*	*Attendance*
	(1)	(2)	(3)	(4)	(5)	(6)
Assignment to choice	118	-69	-146	0.007	-0.016	-0.028
	(60)*	(74)	(84)*	(0.009)	(0.010)	(0.011)**
Assignment to choice *		482	735		0.058	0.091
High payday production difference		(126)***	(144)***		(0.019)***	(0.022)***
Assignment to choice *			401			0.064
Payday			(179)**			(0.024)***
Assignment to choice * Payday *			-1314			-0.178
High payday production difference			(288)***			(0.041)***
Assignment to a target	153	-35	-48	-0.003	-0.019	-0.024
	(71)**	(86)	(96)	(0.010)	(0.012)*	(0.013)*
Assignment to a target *		483	673		0.042	0.066
High payday production difference		(148)***	(168)***		(0.022)*	(0.025)***
Assignment to a target *			68			0.026
Payday			(219)			(0.029)
Assignment to target * Payday *			-972			-0.120
High payday production difference			(348)***			(0.049)***
Payday			-183			-0.009
			(153)			(0.021)
High payday difference *			1178			0.164
Payday			(234)***			(0.032)***
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Lag production controls	Yes	Yes	Yes	No	No	No
Observations	8240	8240	8240	8240	8240	8240
R2	0.60	0.59	0.59	0.11	0.11	0.11
Dependent variable mean	5355	5355	5355	0.875	0.875	0.875

Notes: This table reports estimates of how heterogeneity in treatment effects of the contracts correlates with effects of paydays. Columns (1)-(3) report results from OLS regressions in which the dependent variable is production. Production is defined as the number of accurate fields entered in a day and equals zero on days workers are absent. Columns (4)-(6) report results from OLS Linear Probability Models in which the dependent variable is a binary indicator for attendance. Columns (1) and (4) show regressions of the dependent variable on indicators for contract treatment. Columns (2) and (4) add interactions of each of the contract treatment indicators with an indicator for High payday production difference. High payday production difference equals 1 if the worker's payday difference—defined as the difference between mean production on paydays and non-paydays under Control divided by mean production under assignment to Control—is above the sample average. For each of the contract treatments, columns (3) and (6) add a triple interactions of the contract treatment indicator with the High payday difference indicator and an indicator for whether the current day was the worker's assigned payday (along with the pair wise double interactions between these variables). All regressions include worker, date, and seat fixed effects. In addition, columns (1)-(3) include controls for lagged production. Robust standard errors are reported in parentheses.

Table 7
Demand for the Dominated Contract: Impact of Uncertainty

			Definition of High Uncertainty Indicator							
			Worker is assigned to high uncertainty computer				Worker's morning arrival time is sensitive to bus/train schedules		Worker has above average score on Constraints Index	
Dependent Variable	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Evening choice	-18	-0.002	-26	-0.003	168	0.066	87	0.027	126	0.041
	(37)	(0.012)	(32)	(0.010)	(72)**	(0.022)***	(46)*	(0.016)	(65)*	(0.021)*
High uncertainty indicator			-134	-0.013	-20	0.027	282	0.104	329	0.099
			(63)**	(0.016)	(82)	(0.022)	(206)	(0.062)	(193)*	(0.061)
Evening choice * High uncertainty indicator					-230	-0.082	-253	-0.070	-233	-0.070
					(78)***	(0.024)***	(97)**	(0.029)**	(89)**	(0.029)**
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No
Seat fixed effects	Yes	Yes	No	No	No	No	Yes	Yes	Yes	Yes
Observations	4193	4193	4193	4193	4193	4193	3106	3106	3106	3106
R2	0.34	0.33	0.33	0.32	0.33	0.32	0.13	0.12	0.13	0.12
Dependent variable mean	767	0.28	767	0.28	767	0.28	803	0.3	803	0.3

Notes: This table tests whether the horizon of choice and uncertainty impact demand for the dominated contract. The observations used are the Assignment to Choice treatment observations in the Analysis Sample. Two outcome measures of demand are used. Target level chosen is the continuous measure of the target level selected by workers when assigned to Choice. Positive target indicator is a binary indicator for whether the worker selected a target above zero when assigned to a Choice treatment. Both dependent variables are defined as 0 if a worker was absent the day before or day of Choice assignment.

Evening choice is a dummy that equals 1 if the worker was assigned to Evening Choice and 0 if the worker was assigned to Morning Choice. 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 constraint indicator measures workers' response to a question in the endline survey that asked them to agree or disagree 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," and takes a value of 1 if the worker's response was "Agree Strongly" and 0 otherwise. In columns (9)-(10), the high constraint indicator equals 1 if a workers constraint index score—computed by averaging her responses to the 4 questions on constraints in the endline survey—is above the mean score in the sample. Standard errors are reported in parentheses. 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)-(10).

Table 8: Trends Over Time

Panel A: Trends in Demand for the Dominated Contract

Dependent variable	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*
	(1)	(2)	(3)	(4)
Log experience	-102	-0.025	-128	-0.066
	(96)	(0.032)	(88)	(0.030)**
High payday production difference			-337	-0.087
			(277)	(0.100)
High payday production difference *			189	0.062
Log experience			(76)**	(0.026)**
Observations	4098	4098	4098	4098
Dependent variable mean	759	0.28	759	0.28
F-test p-value: (Log experience) +			0.493	0.895
(High payday difference * Log experience) = 0				

Panel B: Persistence in Treatment Effects Over Time

Dependent variable	*Production*	*Production*	*Production*	*Production*
Sample	Analysis Sample	Analysis Sample	Analysis Sample	Payday Sample
	(1)	(2)	(3)	(4)
Log experience	257	176		314
	(66)***	(73)**		(54)***
More than two months in sample			459	
			(148)***	
Assignment to choice	109	-269	27	
	(59)*	(174)	(76)	
Assignment to choice *		104		
Log experience		(51)**		
Assignment to choice *			144	
More than two months in sample			(114)	
Assignment to a target	145	-264	42	
	(70)**	(217)	(93)	
Assignment to a target *		113		
Log experience		(62)*		
Assignment to a target *			174	
More than two months in sample			(135)	
Payday	140	131	225	110
	(63)**	(179)	(82)***	(169)
Payday *		3		11
Log experience		(51)		(45)
Payday *			-142	
More than two months in sample			(116)	
Observations	8423	8423	8423	11744
Dependent variable mean	5337	5337	5337	5665

Notes: Panel A examines trends in demand for the dominated contract. All columns use observations from the Analysis Sample in which the worker was assigned to a Choice treatment. The dependent variable in columns (1) and (3) is the target level selected by workers when assigned to Choice. The dependent variable in columns (2) and (4) is a binary indicator for whether the worker selected a target above zero when assigned to Choice. Both variables are defined as 0 if a worker was absent the day before or day of Choice assignment. The log of worker experience is defined as the number of workdays an employee has been in the Analysis sample. The high payday difference indicator equals 1 if the worker's mean payday production difference—the difference in production on paydays and non-paydays under control, divided by overall mean production under control—is higher than the sample average. All regressions control for seat and date fixed effects and lagged production. Regressions (1)-(2) also include worker fixed effects. Standard errors are corrected to allow for clustering by worker.

Panel B tests whether treatments effects on production persist over time. All columns show results from OLS regressions in which the dependent variable is production—defined as the number of accurate fields completed by a worker in a day, and equals zero if the worker is absent. Regressions (1)-(3) use observations from the Analysis Sample; Regression (4) uses observations from the Full Payday Sample. Log experience is as defined above. More than 2 months in sample is a binary indicator for whether the worker has been in the sample for more than 2 calendar months. All regressions include lagged production controls and worker, date, and seat fixed effects. Robust standard errors are reported in parentheses.

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

	Correlate of Ability			*Correlate of Self-control Problems*				
	High productivity worker	*Years of education*	*IQ test index score*	*Factor analysis: Self-control factor*	*Self-control index*	*Addictive behaviors dummy (males only)*	*Discount rate: Proportion of impatient responses*	*Discount rate: Proportion of preference reversals*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Dependent Variable: Target Level Chosen (Choice observations)								
Correlate	472	120	-1	124	148	352	230	533
	(159)***	(43)***	(5)	(74)*	(140)	(239)	(342)	(573)
Observations	4187	4056	4089	3106	3106	2245	2454	2454
R2	0.12	0.23	0.22	0.25	0.25	0.29	0.27	0.28
Panel B: Dependent Variable: Positive Target Indicator (Choice observations)								
Correlate	0.116	0.029	-0	0.056	0.057	0.139	0.070	0.143
	(0.049)**	(0.015)*	(0)	(0.025)**	(0.046)	(0.082)	(0.115)	(0.189)
Observations	4187	4056	4089	3106	3106	2245	2454	2470
R2	0.11	0.19	0.19	0.22	0.21	0.26	0.23	0.23
Panel C: Dependent Variable: Production (Choice and Control observations)								
Correlate	237	127	13	-255	-209	-247	-1254	-826
	(93)**	(65)*	(6)**	(103)**	(148)	(395)	(318)***	(461)*
Assignment to choice	1428	147	150	91	91	38	113	116
	(147)***	(73)**	(73)**	(82)	(83)	(114)	(83)*	(83)*
Assignment to choice *	-258	51	-1	167	215	429	706	765
Correlate	(154)*	(45)	(5)	(92)**	(146)	(263)	(305)**	(447)*
Payday	241	153	181	156	156	115	135	131
	(74)***	(74)**	(71)**	(87)*	(87)*	(119)	(106)	(107)
Payday *	-148	27	-0	31	77	58	-53	-234
Correlate	(121)	(40)	(4)	(87)	(138)	(246)	(308)	(544)
Observations	6304	6101	6149	4674	4674	3376	3701	3701
R2	0.57	0.56	0.55	0.57	0.57	0.55	0.59	0.58

Notes: This table tests whether correlates of ability and self-control explain heterogeneity in results. Columns (1)-(3) use measures of ability. The correlate in column (1) is an indicator for whether the worker's mean production under assignment to the control contract is above the sample mean. The correlates in columns (2) and (3) are, respectively, years of education and composite IQ score, defined as the sum of the worker's score on the Raven's Matrix and Digit Span tests. Both education and IQ have been de-meaned. The correlates in columns (4)-(8) are various proxies for self-control. The column (4) correlate is the Self-control Factor, obtained from a principal factors analysis on the endline survey data. The column (5) correlate is a Self-Control Index, obtained by averaging each worker’s responses to the 9 self-control questions in the endline survey. Both the Self-control Factor and Self-control index have been de-meaned. The correlate in column (6) 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 (7)-(8) are computed from the discount rate exercise, in which workers traded off cash rewards between different time horizons. The column (7) correlate measures the proportion of times the worker chose the smaller immediate reward instead of the larger delayed reward. The column (8) 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 dependent variable in Panel A is the mean target level chosen by workers under Choice. The dependent variable in Panel B is a binary indicator for whether the worker selected a positive target under Choice. Both these dependent variables are defined as 0 if a worker was absent the day before or day of Choice assignment. Both Panels A and B report estimates from an OLS regression of the dependent variable on the specified correlate and date and computer fixed effects; columns (2)-(8) also contain lag production controls. The dependent variable in Panel C is production—defined as the number of accurate fields entered in a day, and equal to zero on days workers are absent. Panel C reports estimates from a regression of production on: the specified correlate; Choice and Payday dummies; interactions of each of these dummies with the correlate; date and seat fixed effects; and controls for lagged production. In each regression, standard errors are corrected to allow for clustering by worker. 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.

APPENDIX A: Study details

Production Task

Workers entered information from scanned images into fields on their screen (see Appendix Figure 2). Once a worker finished entering data from an image, the software automatically sent the data to a central server and fetched the next image. This meant workers could not select the images on which they worked. Output was measured as the number of accurate fields entered. The data entry software displayed both the total and accurate number of fields entered so far that day (with about a 15 minute delay), so employees always had real time information on their own output.

Workers faced some uncertainty in production due to shocks. Two types of shocks are particularly relevant in our context. First, the office experienced network speed fluctuations that impacted productivity. Some computers were more sensitive to these fluctuations than others. As a result, workers were randomly assigned to seats in the office and these assignments changed every 1-3 weeks. Second, many employees commuted from surrounding villages using buses and trains, with some traveling up to two hours in each direction. Those from more remote locations faced increased uncertainty in morning arrival times and therefore production.

Paydays

Workers received their wages in cash on their assigned payday. Once they finished work for the day, they reported to the office manager, who computed and paid out their earnings for the previous week (including that day). If employees were absent on their payday, they could collect their owed earnings when they returned to work at no penalty.[40]

Contract Treatments

Workers were paid piece rates based on output. The control contract paid Rs. 0.03 for each accurate field entered, regardless of production amount. The dominated contract paid Rs. 0.03 per accurate field if the worker met the day's production target, and Rs. 0.015 per accurate field otherwise.[41]

Under the Assignment to a Target treatment, workers were assigned to low, medium, and high targets. These 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 worker production under the control contract.[42] During the last month of

[40] The office was closed every alternate Saturday; in those weeks, the Saturday pay group was paid on Friday. More generally, if the office was closed on a scheduled payday due to a holiday, the payment day for members of that group was moved to another day (almost always the day immediately before or after). Such adjustments to the payday schedule were announced in advance and also marked on the posted monthly office calendar.

[41] Under both contracts, workers were also paid a small flat daily wage of Rs. 15 conditional on attendance. Since the base wage constituted about 8 percent of mean daily earnings, the overwhelming majority of worker compensation was tied explicitly to output. Data entry firms in the region commonly use an incentive structure that combines a base wage with payment tied to production. Earnings of workers in the experiment were at par with or slightly higher than those of workers in other data entry firms in the region.

[42] As noted in footnote 24, we did not assign employees to the high target for most of the first half of the study given its level of difficulty.

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.

Before leaving work each day, employees were required to report to an office staff member in a separate area of the office. At that time, they were told their contract assignment for the next day. For example, employees were informed of Wednesday's assignment on Tuesday evening. If the assignment was Evening Choice, they also selected their target for Wednesday at that time. If the assignment was Morning Choice, then they selected their target upon arriving in the office Wednesday morning. This exchange was confidential and took place away from other workers.

Office Structure and Timeline

The office was open each day from 8:45 am to 6:30 pm, five to six days per week except holidays. Employees could choose when they worked, except for two 15-minute periods each day when work activity was halted for "server maintenance". In accordance with the norms of the firm with which we worked, employees were given tea in an outside area at 11 am and 3:30 pm each day. Workers could select the length of their tea breaks and lunch breaks. They were also free to check email, play computer games, or leave the office at any time.

The project ran for 15 months. During the first 2 months, the management staff established protocols, recruited subjects, and trained the new hires. After this, the contract and payday randomizations ran for 4 months. There was then a 2-month break while the office underwent changes to the data entry software and task. During this time, workers were generally not paid the standard piece rates and there were no contract randomizations. The contract and payday randomizations then resumed for another 4 months. In the final 3 months of the project, we ran endline activities and surveys. We did not randomize workers into the four contract treatments during this time, but we continued to adhere to the payday assignments. Thus, the contract treatments ran for an approximate total of 8 months and the payday treatments for 11 months.[43]

Sample construction

The office held 64 data entry operators at a time. Due to employee turnover, 111 workers participated in the experiment.[44] When an employee quit, the management staff hired a replacement from a database of persons that had submitted applications for the job. As in the initial recruitment, workers were hired in order of application date. The gender composition of employees was kept fixed—if a female quit, the worker hired to replace her was a female. Each new worker "inherited" all the assignments of his or her predecessor—payday group, vector of contract assignments, and seating assignment. The payday and seat assignments took effect immediately. New hires began their scheduled contract assignments after completing training.

[43] There were small stretches of time when operations were interrupted due to problems like breakdown of the electricity generator—a reality of business operations in developing countries. These periods never lasted longer than 2 weeks, and often only lasted 1-3 days. Randomizations were suspended during these periods. There were also 3 instances where we paused randomizations while workers were retrained before a change in the data entry task.

[44] 138 workers were hired over the course project. However, 27 of these workers quit before completing training or were only employed outside the 11-month randomization period. As a result, the sample consists of 111 workers.

APPENDIX B:
Appendix Figures and Tables

Appendix Figure 1
Treatment Design

		Contract Assignment (Assignment changes daily)			
		Control Contract (0.25)	*Assigned to Target* (0.25)	*Evening Choice* (0.25)	*Morning Choice* (0.25)
Payday Assignment (Assigned once in beginning of study)	*Tuesday Payday* (0.33)				
	Thursday Payday (0.33)				
	Saturday Payday (0.33)				

Notes: This chart provides an overview of the treatment design. One-third of workers were assigned to each of the three payday groups. This assignment was done once for each worker, when the worker joined the firm, and remained fixed for the duration of the project. Workers were orthogonally assigned to each of the four contract treatments exactly 25% of the time. The assignments changed daily.

Appendix Figure 2
Workers' Data Entry Screen

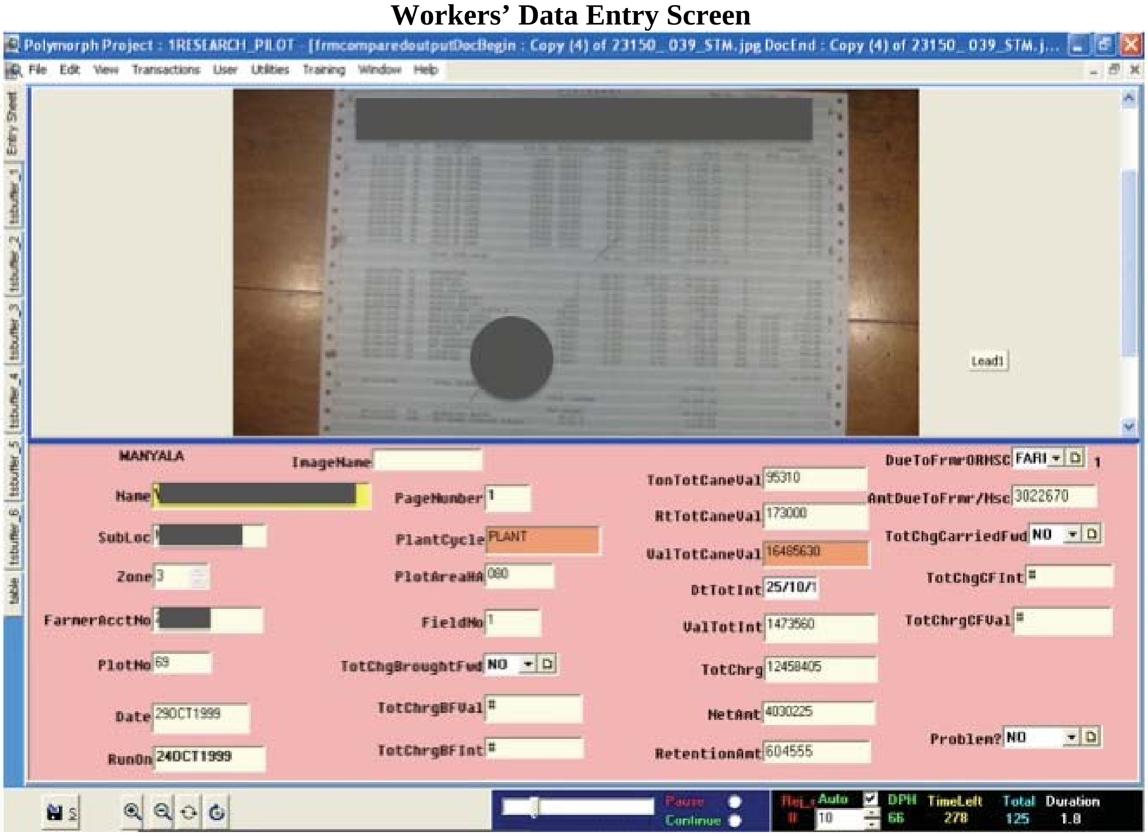

Notes: The figure displays a screen shot of a typical data entry screen. Workers viewed scanned images of records in the top half of their screen and entered information from these images into the appropriate fields at the bottom half of the screen. Identifying information from the records has been covered for confidentiality.

Appendix Figure 3
Proportion of Observations around Target Thresholds

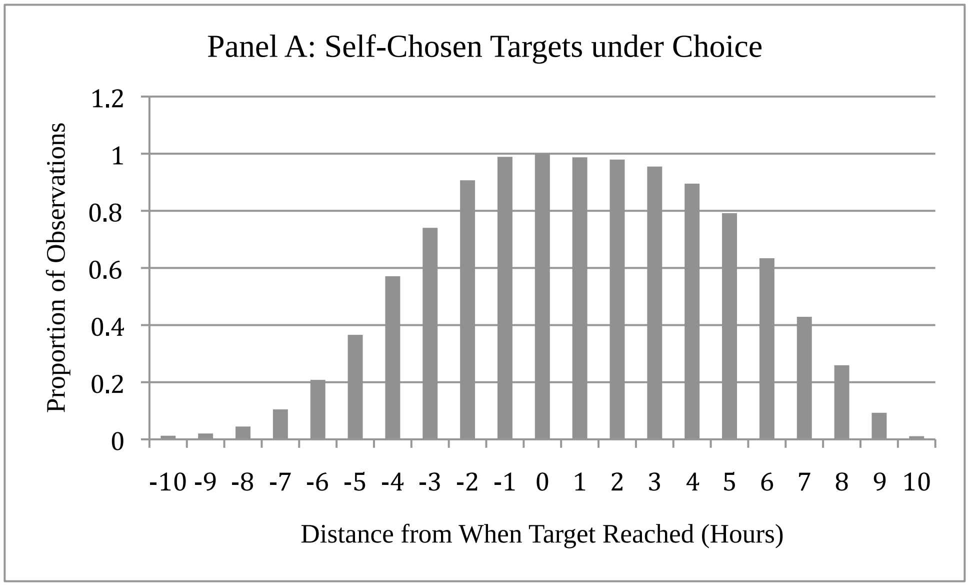

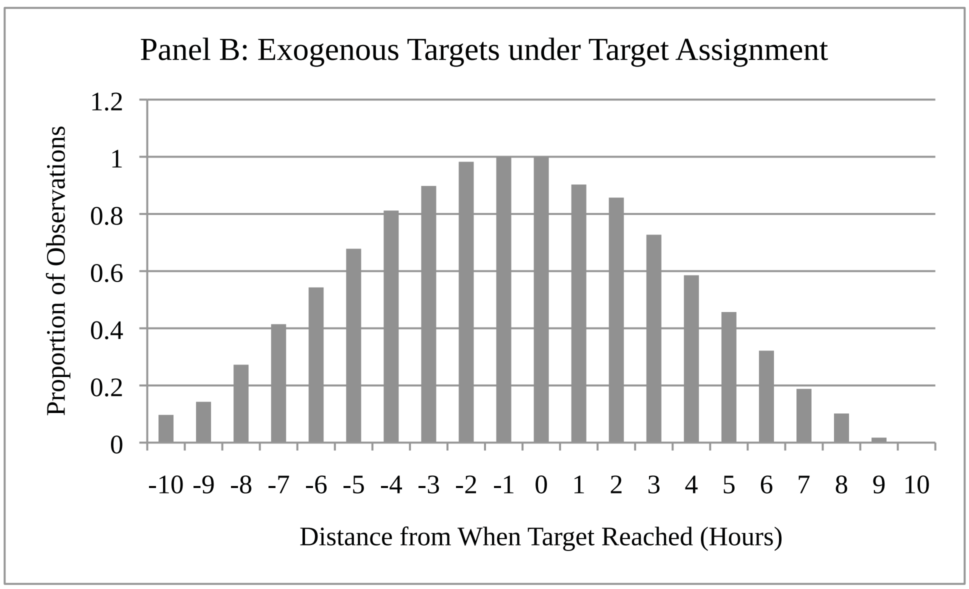

Notes: This figure shows in what proportion of worker-days we observe workers at a given distance from when they reach their targets. The x-axis measures distance in hours from when a worker reached her target for the day. The y-axis measures the proportion of observations for which that distance falls between 8 am – 6 pm (the hours of operation of the office). Panel A computes these statistics for observations in which workers were assigned to Choice, chose positive targets, and were present. Panel B computes these statistics for observations in which workers were assigned to a Target and were present.

Appendix Figure 4
Production Volatility

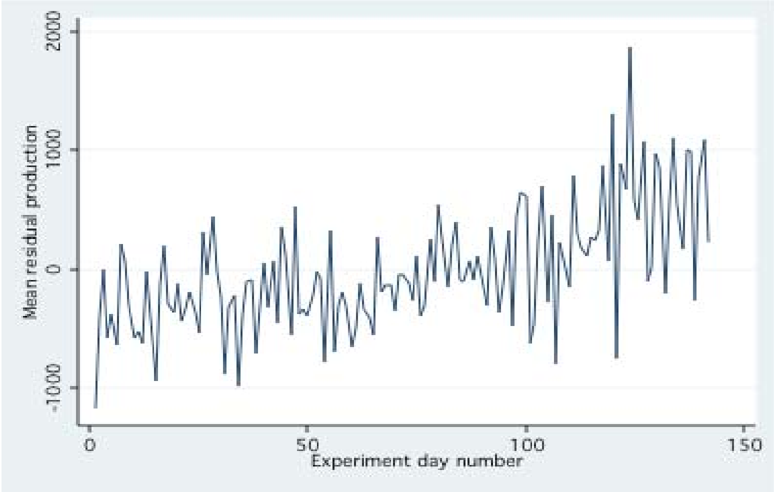

Notes: This figure shows how production varies across days in the Analysis Sample. Residual production is defined as the residual from a regression of production on a vector of worker dummies and lagged production controls. The x-axis measures the day number of the experiment and the y-axis measures the mean of the production residuals for that day.

Appendix Figure 5
Correlation of Computer Production Volatility with Office-level Shocks

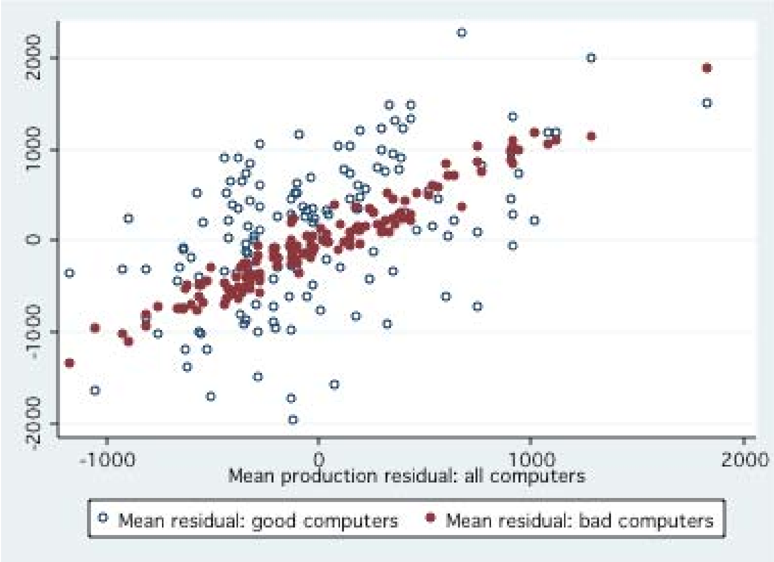

This figure shows the relative production volatility of good and bad computers. Residual production is defined as the residual from a regression of production on a vector of worker dummies and lagged production controls. The x-axis measures the mean of the production residuals across all observations within a day. The open circles plot the mean production residual of workers assigned to good computers within a day. The closed circles plot the mean production residual of workers assigned to bad computers.

Appendix Table 1
Randomization Example

	Worker 1	*Worker 2*	*Worker 3*	*Worker 4*	*Worker 5*
Day 1	Evening Choice	Evening Choice	Control	Target	Evening Choice
Day 2	Morning Choice	Control	Evening Choice	Morning Choice	Control
Day 3	Control	Evening Choice	Morning Choice	Control	Target
Day 4	Morning Choice	Morning Choice	Control	Morning Choice	Control
Day 5	Target	Target	Morning Choice	Target	Morning Choice
Day 6	Control	Control	Evening Choice	Control	Evening Choice
Day 7	Target	Control	Target	Target	Morning Choice
Day 8	Morning Choice	Evening Choice	Target	Control	Target
Day 9	Control	Target	Control	Morning Choice	Control
Day 10	Target	Target	Evening Choice	Evening Choice	Morning Choice
Day 11	Evening Choice	Morning Choice	Morning Choice	Evening Choice	Evening Choice
Day 12	Evening Choice	Morning Choice	Target	Evening Choice	Target
Day 13	Evening Choice	Control	Evening Choice	Evening Choice	Control
Day 14	Morning Choice	Target	Target	Morning Choice	Evening Choice
Day 15	Target	Evening Choice	Morning Choice	Morning Choice	Evening Choice
Day 16	Target	Control	Control	Target	Control
Day 17	Evening Choice	Morning Choice	Target	Evening Choice	Target
Day 18	Morning Choice	Morning Choice	Evening Choice	Evening Choice	Target
Day 19	Target	Control	Target	Target	Morning Choice
Day 20	Control	Target	Control	Control	Target
Day 21	Control	Morning Choice	Morning Choice	Morning Choice	Evening Choice
Day 22	Evening Choice	Evening Choice	Morning Choice	Target	Morning Choice
Day 23	Morning Choice	Evening Choice	Control	Control	Control
Day 24	Control	Target	Evening Choice	Control	Morning Choice

Notes: This table provides an example of the daily contract randomizations. The four contract treatments were: Assignment to the Control contract; Assignment to a Target (at either the low, medium, or high target level), Evening Choice, and Morning Choice. The table shows the contract treatment assignments for five workers over a 24-day period of the study. Workers were assigned to each of the four treatments exactly 3 times over each 12-day period. The order of the assignments was random and changed every 12 days. The vectors of treatment assignments were independent across workers.

Appendix Table 2
Randomization Balance

	Payday Treatments			Contract Treatments				Total
	Tuesday Payday	Thursday Payday	Saturday Payday	Control Contract	Target Assignmt	Evening Choice	Morning Choice	
Proportion of observations	0.33	0.33	0.34	0.25	0.25	0.25	0.25	1.00
Number of observations	2,788	2,809	2,826	2,117	2,113	2,088	2,105	8,423

Notes: This table shows how the 8,423 observations in the Analysis Sample were spread over each of the treatment cells. The Payday Treatments involved randomly assigning workers into one of three payday groups—Tuesday, Thursday, and Saturday—which determined on which day of the week employees were paid their weekly earnings. Workers were randomized into four contract treatments: Assignment to the linear control contract, Assignment to the dominated contract with an exogenously set Target, Evening choice (in which workers chose their preferred contract the evening before the workday), and Morning choice (in which workers chose their preferred contract the morning of the workday). Workers were randomly assigned to each of the 4 contract treatments 3 times over every 12 workdays.

Appendix Table 3
Decomposition of Treatment Effects

Dependent variable	*Workday Length*	*Workday Length*	*Production*	*Production*	*Workday Length | Attend=1*	*Workday Length | Attend=1*	*Production | Attend=1*	*Production | Attend=1*
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
				Panel A: Payday Treatments				
Payday	20	40	-60	26	2	10	-69	-11
	(4)***	(7)***	(38)	(55)	(3)	(4)**	(38)*	(53)
1 day before payday		39		119		17		102
		(7)***		(54)**		(4)***		(53)*
2 days before payday		30		135		16		104
		(8)***		(68)**		(5)***		(67)
3 days before payday		20		169		9		137
		(8)**		(65)**		(5)*		(65)**
4 days before payday		22		87		-0		35
		(9)***		(71)		(5)		(69)
5 days before payday		14		30		4		-2
		(9)		(69)		(5)		(67)
Minutes logged into software			17	17			12	12
			(0.39)***	(0.39)***			(1)***	(1)***
(Minutes logged into software)^2			-0.01	-0.01			-0.003	-0.003
			(0.00)***	(0.00)***			(0.001)**	(0.001)**
Observations	8423	8423	8423	8423	7376	7376	7376	7376
R2	0.27	0.28	0.86	0.86	0.30	0.31	0.84	0.84
Dependent variable mean	404	404	5337	5337	461	461	6094	6094
				Panel B: Contract Treatments				
Assignment to choice	9	-1	48	-13	6.17	5.77	35	-8
	(4)**	(5)	(35)	(46)	(2.69)**	(3.33)*	(34)	(43)
Assignment to a target	4	-6	84	4	6.37	2.47	84	121
	(5)	(6)	(41)**	(53)	(3.14)**	(4.02)	(40)**	(70)*
Assignment to a choice * High payday difference		27		158		1.37		4
		(9)***		(71)**		(5.68)		(51)
Assignment to a target * High payday difference		27		205		11.47		217
		(11)**		(84)**		(6.43)*		(83)***
Minutes logged into software			17	17			12	12
			(0.39)***	(0.39)***			(1)***	(1)***
(Minutes logged into software)^2			-0.01	-0.01			-0.003	-0.003
			(0.00)***	(0.00)***			(0.001)**	(0.001)**
Observations	8240	8240	8240	8240	7213	7213	7213	7213
R2	0.27	0.28	0.86	0.86	0.30	0.30	0.84	0.84
Dependent variable mean	404	404	5355	5355	461	461	6118	6118

Notes: This table decomposes treatment effects on production into minutes worked and production conditional on minutes worked. The dependent variable in columns (1)-(2) is workday length—the number of minutes between a worker's first login and last logout in the data entry software within a day. The dependent variable in columns (3)-(4) is production—the number of accurate fields entered in a day. Both variables equal zero on days workers are absent. The dependent variables in columns (5)-(6) and (7)-(8) are workday length conditional on being present and production conditional on being present, respectively. Panel A investigates impacts of paydays. Panel B investigates behavior under the contract treatments. High payday differenceis a binary indicator that equals 1 if a worker's mean payday production difference—the difference in production on paydays and non-paydays under control, divided by overall mean production under control—is higher than the sample average. All regressions include controls for lag production and worker, date, and seat fixed effects. Robust standard errors are reported in parentheses.

Appendix Table 4
Production Conditional on Attendance over the Pay Cycle

Dependent variable:	Production | Attend=1	Production | Attend=1
Sample	*Analysis Sample*	*Payday Sample*
	(1)	(2)
Payday	35	44
	(69)	(65)
1 day before payday	268	242
	(68)***	(64)***
2 days before payday	218	169
	(85)**	(81)**
3 days before payday	205	193
	(81)**	(78)**
4 days before payday	25	38
	(90)	(80)
5 days before payday	37	63
	(85)	(80)
Lag production controls	Yes	Yes
Worker fixed effects	Yes	Yes
Date fixed effects	Yes	Yes
Seat fixed effects	Yes	Yes
Observations	7376	10341
R2	0.74	0.73
Dependent variable mean	6093	6432

Notes: This table reports the effects of paydays and distance from paydays on worker production conditional on attendance, defined as the number of accurate fields completed by the worker in a day (conditional on the worker being present). Each column shows results from an OLS regression of the dependent variable on indicators of distance from payday, lag production controls, and worker, date, and seat fixed effects. Column (1) reports results from observations in the Analysis sample in which workers were present. Column (2) reports results from observations in the Payday sample in which workers were present. Robust standard errors are reported in parentheses.

Appendix Table 5

Take-up of Dominated Contracts: Summary Statistics

Statistic	*Worker-day means*	*Worker means*
	(1)	(2)
Proportion choosing a positive target (conditional on attendance)	0.35	0.36
	(0.48)	(0.31)
Proportion choosing a positive target (target=0 when absent)	0.28	0.28
	(0.45)	(0.26)
Number of observations	4193	101
	worker-days	workers

Notes: This table reports the proportion of times workers selected positive targets when assigned to a Choice treatment. The first row of the table summarizes choice behavior when the worker was present both the day before and the day of Choice assignment. The second row of the table includes absentee observations, and defines target choice to be 0 if a worker was absent the day before or day of Choice assignment. Column (1) presents means for the Analysis sample of 4,193 choice observations as a whole. Column (2) summarizes the worker means for each statistic, computed over the 101 workers that were assigned to Choice at least once during their employment. Standard deviations for each statistic are reported in parentheses.

Appendix Table 6

Magnitude of Chosen Targets: Summary Statistics

Statistic	*Worker-day means*	*Worker means*
	(1)	(2)
Probability of missing chosen target if assigned to control contract	0.083	0.088
	(0.158)	(0.118)
Proportion of times chosen target was actually missed	0.012	0.025
	(0.109)	(0.118)
Number of observations	1168	84
	worker-days	workers

Notes: This table reports means and standard deviations of statistics that describe the aggressiveness of targets chosen by workers. Row 1 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 Choice, 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 vector of 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 us the probability that the worker would have missed her chosen target under the control contract. Row 2 reports the actual proportion of times workers' production fell below their chosen targets. Column (1) presents means of the test statistic in each row for the 1,168 choice observations in which positive targets was chosen. Column (2) summarizes worker means, computed over the 84 workers that chose a positive target at least once.

Appendix Table 7
Hourly Production around Target Thresholds

Observations	Self-Chosen Targets (Assignment to Choice)			Exogenous Targets (Assignment to a Target)		
Dependent variable	*Hourwise production*	*Hourwise production*	*Hourwise production*	*Hourwise production*	*Hourwise production*	*Hourwise production*
	(1)	(2)	(3)	(4)	(5)	(6)
3 hours before target reached		135	-13		189	48
		(19)***	(15)		(14)***	(12)***
2 hours before target reached		236	14		224	80
		(23)***	(87)		(16)***	(14)***
1 hour before target reached		432	88		398	176
		(22)***	(20)***		(16)***	(16)***
Hour in which target was reached	597	934	498	543	722	510
	(51)***	(33)***	(30)***	(25)***	(26)***	(25)***
1 hour after target reached		664	223		439	247
		(29)***	(28)***		(19)***	(19)***
2 hours after target reached		584	170		336	206
		(26)***	(24)***		(43)***	(22)***
3 hours after target reached		588	198		313	226
		(34)***	(31)***		(27)***	(25)***
4+ hours after target reached		270	133		54	164
		(21)***	(18)***		(20)***	(20)***
Worker's mean production under control in current hour			0.90			0.85
			(0.02)***			(0.02)***
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	13,948	13,948	13,948	20,174	20,174	20,174
R2	0.26	0.29	0.42	0.24	0.27	0.41
Dependent variable mean	847	847	847	730	730	730
F-test p-value: (Hour in which target reached) = (1 hour after target reached)		0.00	0.00		0.00	0.00

Notes: This table describes production levels around the hours in which workers achieved their targets.

The dependent variable in each regression is hourwise production, defined as production by worker i on day t in calendar hour h. For each date in the sample, there are 11 calendar hours, which correspond to the times the office was open and it was possible for employees to work: the first hour is 8-9 am and last is 6-7 pm. Production is coded as 0 if a worker did not work in a certain hour.

The table displays 8 binary variables that capture distance from when the target was reached. For example, the "Hour in which target was reached" is an indicator variable that equals 1 in the calendar hour in which a worker achieved her target for the day, and equals 0 otherwise. In cases where workers failed to achieve their target, the hour in which they reached their target is coded as the hour after the office closed (7-8 pm). "Worker's mean production under control in current hour" is a covariate that equals the sample mean of worker i's production in hour h on days in the Analysis Sample when worker i was assigned to the Control contract and was present. The last row of the table displays p-values from an F-test of whether the coefficient on Hour in which target was reached equals the coefficient on 1 after target reached in the regression shown in that column. Standard errors are corrected to allow for clustering by worker-day.

Columns (1)-(3) report estimates from days when workers selected positive targets under Assignment to Choice and were present. Columns (4)-(6) report estimates from days when workers were assigned to exogenous targets under Target Assignment and were present.

Appendix Table 8
Test for Inter-temporal Substitution in Effort across Days

Dependent variable	Production	Production	Production	Production	Production
Observations: Today's assignment	*All contracts*	*Control & Choice*	*Control only*	*Control only*	*Control only*
	(1)	(2)	(3)	(4)	(5)
Assigned to choice yesterday	1	7	193		
	(67)	(78)	(163)		
Assigned to a target yesterday	19	12	243		
	(77)	(87)	(179)		
Assigned to a high target yesterday				372	
				(267)	
Assigned to choice the past 2 days in a row					54
					(144)
Assigned to a target the past 2 days in a row					-74
					(274)
Worker fixed effects	Yes	Yes	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	8423	6310	2117	2117	2117
R2	0.50	0.51	0.55	0.55	0.55
Dependent variable mean	5337	5311	5283	5283	5283

Notes: This table tests for evidence on whether there is inter-temporal substitution in worker effort across days. The dependent variable in each column is production, which is defined as the number of accurate fields completed by worker i on date t, and equals 0 if the worker was absent.

The observations in column (1) are all the observations in the Analysis sample. Column (2) restricts analysis to those observations in which a worker was assigned to the Control contract or to Choice on day t. Columns (3)-(5) further restrict analysis to only those observations in which a worker was assigned to the Control contract on date t.

Columns (1)-(3) show results from an OLS regression of production on an indicator for whether worker i was assigned to Choice on date t-1, and an indicator for whether worker i was assigned to a Target on date t-1. Column (4) shows an OLS regression of production on an indicator for whether worker i was assigned to the High Target on date t-1. Column (5) shows an OLS regression of production on an indicator for whether worker i was assigned to Choice on dates t-1 and t-2, and an indicator for whether worker i was assigned to a Target on date t-1 and t-2. All regressions include controls for worker, date, and seat fixed effects. Robust standard errors are reported in parentheses.

Appendix Table 9

Predictive Power of Contract Assignment Probability

Dependent variable	*Production*	*Production*	*Production*
	(1)	(2)	(3)
Assignment to evening choice	150		161
	(69)**		(78)**
Assignment to morning choice	73		131
	(69)		(76)*
Assignment to low target	3		20
	(90)		(100)
Assignment to medium target	213		207
	(91)**		(102)**
Assignment to high target	334		405
	(150)**		(159)**
Pr(evening choice)		106	-65
		(161)	(182)
Pr(morning choice)		-155	-296
		(158)	(176)*
Pr(low target)		-63	-93
		(203)	(225)
Pr(medium target)		260	43
		(207)	(234)
Pr(high target)		-88	-487
		(373)	(396)
Worker fixed effects	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes
Lag production controls	Yes	Yes	Yes
Observations	8423	8423	8423
R2	0.59	0.59	0.59
Dependent variable mean	5337	5337	5337
F-test of joint significance of probability controls (p-value):		0.39	0.45

Notes: This table tests whether probability of assignment to a contract treatment predicts output. The dependent variable in each regression is production, which is defined as the number of accurate fields completed by worker i on date t, and equals 0 when a worker is absent. Column (1) reports results from an OLS regression of the dependent variable on dummies for each contract assignment treatment. Column (2) shows a regression of production on the probabilities of worker i receiving each contract treatment on date t. The probabilities are computed using the worker's previous assignments within the randomization block. Column (3) shows a regression of production on all the treatment assignment dummies and probability controls. All regressions include controls for lagged production and worker, date, and seat fixed effects. Robust standard errors are reported in parentheses. The bottom row of the table reports the p-value of an F-test of joint significance of the 5 probability controls in columns (2) and (3).

Appendix Table 10
Contract Effects: High Payday Difference Workers

Dependent variable	*Production*	*Attendance*	*Production | Attend=1*
	(1)	(2)	(3)
Assignment to choice	395	0.044	144
	(116)***	(0.016)***	(73)**
Assignment to a target	452	0.023	376
	(121)***	(0.019)	(83)***
Worker fixed effects	Yes	Yes	Yes
Computer fixed effects	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes
Lag production controls	Yes	No	Yes
Observations	3216	3216	2706
R2	0.53	0.14	0.72
Dependent variable mean	4696	0.84	5581

Notes: This table estimates the treatments effects of the contracts on production for high payday difference workers. The table uses observations from only those workers in the Analysis sample whose payday production difference was above the sample average. The payday production difference is computed as the difference between the worker's mean production on paydays and non-paydays under Control divided by mean production under assignment to Control. The dependent variable in column (1) is production. Production is defined as the number of accurate fields completed by the worker in a day, and equals zero on days workers are absent. The dependent variable in columns (2) is attendance—a binary indicator that takes the value of 1 when workers are present and 0 otherwise. The dependent variable in column (3) is production conditional on attendance—it equals production when a worker is present and is missing otherwise. Each column shows the estimates from an OLS regression of the dependent variable on indicators for Choice and Target assignment. All regressions include fixed effects for each date in the sample, each worker in the sample, and each seating assignment. In addition, columns (1) and (3) also include controls for lagged production. Robust standard errors are reported in parentheses.

Appendix Table 11

Heterogeneity in Payday Treatment Effects: Correlation with Contract Effects

Dependent variable	*Production*	*Attendance*	*Production*	*Attendance*	*Production*	*Attendance*
	(1)	(2)	(3)	(4)	(5)	(6)
Payday	156	0.044	-207	-0.007	-247	-0.013
	(166)	(0.025)*	(173)	(0.023)	(182)	(0.023)
High dominated contract take-up * Payday	-2	-0.001				
	(249)	(0.032)				
High choice treatment effects * Payday			802	0.107		
			(251)***	(0.035)***		
High target treatment effects * Payday					884	0.118
					(248)***	(0.033)***
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Lag production controls	Yes	No	Yes	No	Yes	No
Observations	2117	2117	2117	2117	2115	2115
R2	0.65	0.24	0.65	0.24	0.65	0.24
Dependent variable mean	5283	0.87	5283	0.87	5186	0.87

Notes: This table tests whether behavior under the contract treatments predicts payday effects. All regressions report estimates from observations from the Analysis Sample in which the worker was assigned to the Control contract. Columns (1), (3), and (5) report estimates from an OLS model where the dependent variable is production. Production is defined as the number of accurate fields entered in a day and equals 0 when the worker is absent. Columns (2), (4), and (6) report estimates from an OLS Linear Probability Model in which the dependent variable is a binary indicator for attendance.

High dominated contract take-up rate is a dummy that equals 1 if the worker's take-up rate of the dominated contract—the proportion of times the worker chose a positive target under assignment to choice on nonpaydays—is above the mean take-up rate in the sample. Columns (1) and (2) show regressions of the dependent variable on an indicator for whether the day was the worker's assigned payday and an interaction of the payday dummy with the high take-up dummy.

High choice treatment effects is a dummy that equals 1 if the worker's choice differential—defined as the difference between mean production under assignment to Choice versus Control on nonpaydays—is above the sample average. Columns (3) and (4) show regressions of the dependent variable on the payday dummy and an interaction of the payday dummy with the high choice treatment effects dummy.

High target treatment effects is a dummy that equals 1 if the worker's choice differential—defined as the difference between mean production under assignment to a Target versus Control on nonpaydays—is above the sample average. Columns (5) and (6) show regressions of the dependent variable on the payday dummy and an interaction of the payday dummy with the high target treatment effects dummy.

All regressions include controls for worker, date, and seat fixed effects. Note that the number of observations changes between columns because contract treatments were assigned randomly and workers were in the sample for varying lengths of time. Each measure requires that we have observations for a worker on non-paydays under a certain set of contract assignments—assignment to Choice, assignment to Choice and Control, and assignment to a Target and Control, respectively. For some workers, these measures cannot be computed and they are therefore not included in the analysis. Robust standard errors are reported in parentheses.

Appendix Table 12
Sensitivity of Bad Computers to Production Shocks

Dependent variable	*Production residual*	*Production residual*
	(1)	(2)
Bad computer indicator	-139	-130
	(72)*	(79)
Mean of day's residual production	0.454	
	(0.097)***	
Bad computer indicator * Mean of day's residual production	0.261	
	(0.116)**	
Mean of day's residual production for bad computers		0.373
		(0.098)***
Bad computer indicator * Mean of day's residual production for bad computers		0.263
		(0.117)**
Mean of day's residual production for good computers		0.088
		(0.087)
Bad computer indicator * Mean of day's residual production for good computers		-0.020
		(0.099)
Observations	8423	8423
R2	0.024	0.024
Dependent variable mean	0.000	0.000

Notes: This table reports estimates on whether workers assigned to bad computers are more sensitive to productivity shocks. The table shows results from OLS regressions in which the dependent variable is the residual from a regression of production on a vector of worker dummies and lagged production controls. Column (1) shows a regression of the production residual on: an indicator for whether the worker was assigned to a computer identified by workers and management as being likely to become slow during network fluctuations; the mean residual production for that day (by averaging the dependent variable across the observations for that day, excluding the worker's own observation); and an interaction of the 2 covariates. Column (2) shows a regression of the production residual on the bad computer indicator; the mean of residual production that day, computed separately for the good and bad computers (excluding the worker's own observation); and interactions of these latter two covariates with the bad computer indicator. Standard errors are clustered by computer (seat) assignment.

Appendix Table 13
Correlation between Contract Quiz Score and Demand for Dominated Contracts

Dependent Variable:	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*	*Target level chosen*	*Positive target indicator*
	(1)	(2)	(3)	(4)	(5)	(6)
Scored 100% on quiz	175	0.001				
	(182)	(0.058)				
Quiz score (percentage)			280	-0.11		
			(682)	(0.21)		
Education (years)					233	0.059
					(49)***	(0.015)***
Date fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3582	3582	3582	3582	3582	3582
R2	0.10	0.11	0.10	0.11	0.15	0.14
Dependent variable mean	804	0.29	804	0.29	804	0.29

Notes: This table tests whether a worker's score on the contracts quiz, which tested comprehension of the contracts treatments and was administered before workers began contract randomizations, is correlated with demand for the dominated contract. Columns (1), (3), and (5) show estimates of OLS models in which the dependent variable is the target level selected by a worker when assigned to Choice. Columns (2), (4), and (6) show estimates of OLS Linear Probability Models in which the dependent variable is a binary indicator for whether the worker selected a target above zero when assigned to a Choice treatment. Both dependent variables are defined as 0 if a worker was absent the day before or day of Choice assignment. Columns (1)-(2) show regressions of the dependent variable on an indicator that equals 1 if the worker scored 100% on the contracts quiz, and equals 0 if the worker scored below 100%. Columns (3)-(4) show regressions of the dependent variable on the continuous measure of the worker's percentage score in the quiz. Columsn (5)-(6) show regressions of the dependent variable on years of eduction. All regressions include date and seat fixed effects. Standard errors are clustered by worker. In 3 instances, workers who initially scored below 80% on the quiz were re-trained and retook the quiz before the start of contract randomizations; their quiz score in the data equals their score on their second quiz. Due to a clerical error, the quiz was not administered to all workers. Observations are for those workers in the Analysis sample that took the contracts quiz.

Appendix Table 14
Test for Income Targeting: Production Increases Before Paydays

Dependent variable	*Prodn*	*Prodn*	*Prodn*
	(1)	(2)	(3)
Payday	200	110	105
	(90)**	(145)	(129)
1 workday before payday	331	181	268
	(99)***	(147)	(132)**
Assigned to Target yesterday		-99	-76
		(103)	(92)
Payday *		259	179
Assigned to Target yesterday		(191)	(169)
1 workday before payday *		341	122
Assigned to Target yesterday		(205)*	(185)
Lag production controls	Yes	No	Yes
Worker fixed effects	Yes	Yes	Yes
Date fixed effects	Yes	Yes	Yes
Seat fixed effects	Yes	Yes	Yes
Observations	4344	4344	4344
R2	0.61	0.52	0.61
Dependent variable mean	5352	5352	5352

Notes: This table tests for income targeting by workers within their payweeks. All regressions report estimates from observations from the Analysis Sample in which the worker was assigned to the Control contract or to a Target on the previous day. All columns report estimates from an OLS model where the dependent variable is production. Production is defined as the number of accurate fields entered in a day and equals 0 when the worker is absent. Payday is an indicator variable for whether the current day is the worker's assigned payday. 1 workday before payday is an indicator for whether the current day is 1 workday before the worker's assigned payday. Assigned to Target yesterday is an indicator that equals 1 if the worker was assigned to a target yesterday, and equals 0 if the worker was assigned to the control contract yesterday. All regressions include controls for worker, date, and seat fixed effects. Columns (1) and (3) also include controls for lag production. Robust standard errors are reported in parentheses.