

Incentivizing Behavioral Change: The Role of Time Preferences

Shilpa Aggarwal Rebecca Dizon-Ross Ariel Zucker*
Indian School of Business University of Chicago UC Berkeley

September 28, 2018

Preliminary and incomplete

Abstract

How should the design of incentives vary with the time preferences of agents? We formulate predictions for two incentive contract variations that should increase efficacy for impatient agents relative to patient ones: increasing the frequency of incentive payments, and making the contract “dynamically non-separable” by only rewarding compliance in a given period if the agent complies in a minimum number of other periods. We test the efficacy of these variations, and their interactions with time preferences, using a randomized evaluation of an incentives program for exercise among 3,200 diabetics in India. On average, providing incentives increases daily walking by 1,300 steps or roughly 13 minutes of brisk walking, and decreases the health risk factors for diabetes. Increasing the frequency of payment does not increase effectiveness, suggesting limited impatience over payments. However, making the payment function dynamically non-separable increases cost-effectiveness. Consistent with our theoretical predictions, agent impatience over walking appears to play a role in non-separability’s efficacy: both heterogeneity analysis based on measured impatience and a calibrated model suggest that the non-separable contract works better for the impatient.

*Aggarwal: Indian School of Business shilpa_aggarwal@isb.edu. Dizon-Ross: Chicago Booth School of Business rdr@chicagobooth.edu. Zucker: University of California, Berkeley ariel.d.zucker@gmail.com. This study was funded by the Government of Tamil Nadu, the Initiative for Global Markets, J-PAL USI Initiative, the Chicago Booth School of Business, the Tata Center for Development, the Chicago India Trust, and the Indian School of Business. The study protocols received approval from the IRBs of MIT, Chicago, and IFMR. The experiment was registered on the AEA RCT Registry. We thank Ishani Chatterjee, Rupasree Srikumar, and Sahithya Venkatesan for their great contributions in leading the fieldwork, and Christine Cai, Yashna Nandan, and Emily Zhang for outstanding research assistance. We are grateful to Abhijit Banerjee, Marianne Bertrand, Esther Duflo, Pascaline Dupas, Rick Hornbeck, Seema Jayachandran, Supreet Kaur, Rohini Pande, Devin Pope, Canice Prendergast, Gautam Rao, and Frank Schilbach for helpful conversations and feedback, and to numerous seminar and conference participants for insightful discussions. All errors are our own.

1 Introduction

Incentive design is of core economic interest. While most classic contracting models assume that agents are relatively patient, there is growing evidence that many people are impatient. This raises an important question: What are the implications of agent impatience for the design of incentives? In this paper, we develop predictions for how to adjust incentives for impatience in a contracting setting where principals pay agents to repeatedly engage in tasks over time. Using a randomized controlled trial (RCT) of an incentive scheme for exercise among 3,200 diabetics in India, we then test our predictions by implementing two incentive contract variations designed to increase efficacy for impatient agents relative to patient ones.

Our predictions distinguish between discounting of consumption and of financial payments. As emphasized in the recent time-preference literature, when agents consider future financial payments, the effective discount rate they apply depends on their available borrowing and lending opportunities, and is distinct from the true “primitive” or structural discount rate they apply to future consumption or effort (Augenblick et al., 2015). As a result, we model agents as having separate discounting parameters in the domains of consumption and financial payments. We then identify two contract variations, one per domain, each with increasing efficacy in agents’ impatience in that domain. Both contract variations are predicted to work well for agents with a variety of types of impatience, including those with high discount rates that are time-consistent or constant over time; those with high discount rates that are “time-inconsistent” or decline with delay; and, among time-inconsistent, both those who are “sophisticated” and aware of their own time inconsistency and those who are “naive” and unaware. We use the term “impatience” to capture all of these variants.

The first prediction we test is about the interaction of “dynamic non-separability” and impatience over effort. By dynamically non-separable, we mean that the incentive paid for action in a given period depends on the actions in other periods. For example, the contract might pay people for taking an action on a day if and only if they take that action on at least 5 days in the week. Dynamic non-separability has several advantages and disadvantages that have been discussed before. Our new theoretical insight is that dynamic non-separability interacts with impatience. Relative to dynamically separable contracts, certain dynamically non-separable contracts should increase compliance for those who are impatient over future effort (i.e., who discount future effort heavily) compared to for those who are patient.¹ This is because payments in dynamically non-separable contracts are a function of behavior on multiple days, thus linking the agent’s decisions about exerting effort over time. In a dynamically separable contract, where the agent is paid separately for his behavior in each

¹Section 2 discusses which types of dynamically non-separable contracts our predictions apply to.

period, the agent always compares the financial incentive to the cost of effort *this period* (which is not discounted), and so discount rates over future effort do not matter. In contrast, with dynamically non-separable contracts, the agent compares the incentive to the present *discounted* cost of effort in multiple periods – which will be lower for those who discount future effort more heavily. Intuitively, this type of contract takes advantage of the fact that impatient people discount their future effort.²

Second, we test the prediction that, if agents are impatient over payments, providing more frequent payment increases efficacy. This is the most intuitive prediction of impatience for incentive design and, to our knowledge, the main prediction discussed previously in the literature (besides pre-commitment) for tailoring incentives to impatience,³ with for example Cutler and Everett (2010) proclaiming “the more frequent the reward, the better.” However, there are no empirical evaluations of the effect of changing the payment frequency, and theoretical reasons to question both whether and what types of increases in payment frequency would matter. First, whether payment frequency matters depends on whether discount rates over money are small or large. If credit markets are perfect, discount rates over payment should likely be small (equal to the market interest rate) and increasing payment frequency may have limited impact. If instead credit markets are imperfect, as may be especially true in developing countries, discount rates over payment could approach those over consumption and adjusting payment frequency could have a larger impact. Second, even if discount rates are high, there is still an open question regarding what types of frequency increases would matter. The answer depends on the shape of discount rates over time: if discount rates decline very quickly with lag (as with for example the “quasi-hyperbolic” or “beta-delta” models used in the literature), the gains to increasing frequency are limited unless payments can be made very frequently (e.g., every day), whereas in other models there could be large gains between, say, every month and every week. We test the efficacy of three payment frequencies – monthly, weekly, and daily – to assess whether, and what type of, increases in payment frequency improve compliance. The three frequencies also allow us to explore which model of payment discounting best fits the data.

We test these predictions using an RCT evaluating incentives for behavioral change, which are increasingly prevalent in areas such as health (Duflo et al., 2010; Martins et al., 2009; Morris et al., 2004; Thornton, 2008), education (Fryer, 2011) and the environment

²This logic holds both for sophisticates and naive time-inconsistent (as well as impatient time-consistent), although the logic plays out somewhat differently by type. For sophisticates, non-separability creates a commitment motive: agents exercise today to induce their future selves to exercise. For time-inconsistent naives, who are overoptimistic about their future desire to exercise, it creates an “option value” motive: they exercise today to give their future selves the opportunity to follow-through.

³O’Donoghue and Rabin (1999a) examine how to optimally penalize time-inconsistent procrastinators for delays in a setting where delay is costly to the principal but task costs vary over time.

(Davis et al., 2014; Jayachandran et al., 2017). Tailoring incentives for impatience may be particularly important in the behavioral-change domain since the express purpose of many behavioral-change incentive programs is to address a specific form of impatience: present bias.⁴ Present-biased agents may fail to undertake behaviors with short-run costs but only long-run benefits (e.g., eating right), even if those behaviors are in their own long-run self-interest. Behavioral-change incentive programs thus deliver small, short-run incentives in an attempt to better align agents’ behavior with their long-run interests. However, even in this domain where present bias is a key rationale for using incentives, there is very little evidence on how to adjust incentives for present-biased (or otherwise impatient) agents.

The specific behavior we target in our experiment is exercise among diabetics in India. “Lifestyle diseases” like diabetes and hypertension are exploding policy problems in both developing and developed countries, with the current worldwide cost of diabetes estimated at \$1.3 trillion or 1.8% of global GDP (Bommer et al., 2017) and the burden in India estimated at 2% of GDP (Tharkar et al., 2010). It is widely agreed that one key to decreasing the global burden is to get diabetics to comply with their disease management guidelines: exercise, good diet, and medication adherence. Since these behaviors involve short-run costs and long-run benefits, incentives are a promising approach to improve compliance. We deliver incentives to diabetics and prediabetics for walking, a key component of diabetes management (Qiu et al., 2014; Zanuso et al., 2009). We monitor participants’ walking using pedometers and, if they achieve a daily step target of 10,000 steps, provide them with small financial incentives in the form of mobile recharges (i.e. cell-phone credits). Within the incentives program, we randomly vary (i) whether payment is a linear function of the number of days the agent meets the step target or is “dynamically non-separable,” only rewarding step-target compliance on a given day if the agent meets the step target on a minimum number of other days that week (we use two minimum compliance levels: 4 days and 5 days), and (ii) the frequency with which incentives are paid. We also randomly assign some participants to a pure control group, and some to a “monitoring group” which receives pedometers but no incentives, allowing us to test for the overall effects of incentives on exercise and health.

Our analysis proceeds as follows. We first establish that our incentives program is highly effective at inducing exercise. Providing just 20 INR (0.33 USD) per day of compliance with a daily step target increases compliance by 20 percentage points (pp) off of a base of 30 percent. Average daily steps increase by 1300 or roughly 13 minutes of brisk walking.

We then use our experiment to explore the implications of time preferences for incentive design, presenting three main results. We begin by exploring dynamic non-separability. Our first main finding is that, consistent with our theoretical predictions, moving from a

⁴In contrast, other types of incentives (e.g., incentives for workers) often aim primarily to solve moral hazard issues instead of an “internality” like time inconsistency.

dynamically separable contract to a dynamically non-separable one increases relative efficacy for those who are impatient over effort relative to those who are not. To establish this finding, we first perform heterogeneity analysis based on a baseline measure of discount rates over exercise.⁵ We find that, relative to linear contracts, non-separable contracts increase compliance 4pp more for those whose impatience is above-median relative to those who are below, and 10pp more for those above the 75th percentile. These magnitudes are large relative to the sample-average effect of either contract (20 pp). We also calibrate a model using experimental estimates of the distribution of daily walking costs; the results there also suggest that dynamically non-separable contracts work considerably better for the impatient, with projected compliance in the dynamically non-separable contract relative to the linear increasing by 5pp for each 10pp decrease in the discount factor.

To complete our analysis of dynamic non-separability, we explore its average efficacy, presenting to our knowledge the first empirical comparison of a dynamically non-separable contract with a dynamically separable one. We find that, on average, making the contract dynamically non-separable improves cost-effectiveness: the percent of days on which people hit their step target does not change, but if agents do not meet the step target on at least 4 or 5 days in the week, they do not receive incentives for every day the step target is reached like they do with the linear contract. As a result, the 4-day and 5-day non-separable contracts cost roughly 10 and 15% less while generating the same amount of exercise. Dynamic non-separability has a potential downside, however: it generates more extreme outcomes, working better for some but worse for others. This variation in efficacy makes it important to determine for whom dynamically non-separable contracts work well, highlighting the significance of our finding that the non-separable contracts work better for the impatient.

We next turn to the other dimension of contracts we varied – payment frequency. Our third main result is that increasing the frequency of incentive delivery has limited impact. Incentives delivered at daily, weekly, and monthly frequencies all have equally large impacts on walking, indicating that the model that best fits our sample is one of patience over financial payments. We find additional evidence in support of this conclusion: step-target compliance does not increase as the date of payment delivery approaches. This null finding suggests that, in contrast with the conventional wisdom, increasing incentive frequency is not an effective way to adjust incentives for impatience in our setting. This result is consistent with Augenblick et al. (2015) who find limited impatience in monetary choices but is perhaps still surprising given the limited access to borrowing in our setting.

We conclude the paper with a program evaluation of our incentives program. Our sample has high rates of diabetes and hypertension; regular exercise can prevent complications from

⁵We use Andreoni and Sprenger (2012a)’s convex time budget method; people divide steps over time.

both. We show that the large increases in walking induced by incentives cause moderate improvements in physical health and emotional wellbeing. Incentives improve an index of overall health risk, including measures of blood sugar and body mass index, by a moderately-sized amount. Incentives also improve mental health, suggesting a causal link between exercise and mental health. The health effects are important for policy, suggesting that incentives may be a cost-effective way to decrease the burden of diabetes in India and beyond.

This paper makes several contributions to the literature. First, it contributes to the literature on motivating time-inconsistent or impatient agents and on contract design with impatience. To date, the primary way that researchers have attempted to motivate time-inconsistent agents is to provide commitment devices or contracts that restrict the possible actions of their future selves (e.g., Ariely and Wertenbroch (2002); Kaur et al. (2015); Ashraf et al. (2006); Giné et al. (2010); Duflo et al. (2011); Schilbach (2017); Royer et al. (2015)). Although pre-commitment can be a very useful tool, it is not a panacea: take-up of commitment contracts is generally modest, as discussed in Laibson (2015). Indeed, commitment contracts are only effective for sophisticated time-inconsistent agents, but evidence suggests that a large share of the population is at least partially naive and that commitment can in fact be harmful for partially naive agents (Augenblick and Rabin, 2017; Bai et al., 2017). In contrast, the predictions we test do not require sophistication: they work for multiple types of impatience, including naive time-inconsistency. Beyond the literature on pre-commitment, there is limited work, theoretical or empirical, on how to optimize incentive design for impatience; we discuss the few exceptions in Section 1.1.

Second, we build on the literature examining dynamic contracting (see for example Lazear (1979) and Prendergast (1999)). Many theoretical dynamic contracting papers yield the prediction that it is optimal to defer some component of current pay until the future, making the contract dynamically-nonseparable. To our knowledge, however, there are no empirical comparisons of contracts with and without deferred compensation; since our dynamically non-separable contract entails deferred compensation, our experiment represents the first such comparison.⁶ Our paper is also the first to theoretically show that deferred compensation can work better for impatient agents.

We build on a third body of literature that measures the shape of time preferences. The majority of the recent work has focused on distinguishing whether consumption and payment discount rates are time-consistent or time-inconsistent (Andreoni et al., 2016; Andreoni and Sprenger, 2012a; Augenblick et al., 2015). However, within those classes, there is large variation in the feasible shape and size of discount rates with important policy implications;

⁶Carrera et al. (2017) and Bachireddy et al. (2018) evaluate *separable* contracts which vary payments over time.

to our knowledge, our paper is the first to test the policy implications of that variation.⁷ Finally, we contribute to the growing literature on incentives for health, such as incentives for weight loss (Kullgren et al., 2013; Volpp et al., 2008) and disease monitoring (Labhardt et al., 2011). Although there are several financial incentives trials for exercise for non-diabetics (e.g., Charness and Gneezy (2009); Finkelstein et al. (2008)), as well as one incentivizing 3-month blood sugar control, there is a lack of interventions that incentivize important daily habits among diabetics. Ours represents the first evaluation of incentives to diabetics for daily disease management, and the first trial of incentives for exercise in a developing country.

The paper proceeds as follows. Section 1.1 discusses the literature on contract design and impatience. Section 2 presents the predictions motivating the experiment. Sections 3 and 4 discuss the study setting, design, and data. In Section 5, we present results on incentive design and impatience. Section 6 presents impacts on health outcomes. Section 7 concludes.

1.1 Related literature: Contract design and impatience

We now review the brief theoretical and empirical literature studying how (besides pre-commitment) one can improve incentive design for impatient agents.⁸ O’Donoghue and Rabin (1999a) outline the theoretical implications of time-inconsistent procrastination for the design of “temporal incentive schemes,” which reward agents based on when they complete tasks. Their focus is on avoiding delay; they find that optimal incentives for procrastinators typically involve an increasing punishment for delay as time passes. Carrera et al. (2017) work in a setting where there are one-time “startup costs” for compliance; they show that, in such a setting, offering larger incentives at the beginning of an incentive contract does not empirically decrease procrastination. These studies both differ from ours in their theoretical goals and environments: our work focuses on maximizing the average level of compliance over time rather than on avoiding delay or overcoming startup costs. More similar in spirit to our work, several papers show that worker performance improves toward the end of pay cycles, and attribute this effect to impatience (Clark, 1994; Kaur et al., 2015; Oyer, 1998).⁹ While these papers suggest a potential role for high-frequency payment to improve average performance, with Kaur et al. (2015) suggesting this explicitly, ours is the first to directly test this suggestion.

⁷Two previous papers have explored the shape of discounting using lab-experimental preference measures in the monetary domain. Benhabib et al. (2008) test between hyperbolic, quasi-hyperbolic, and exponential models, but do not have the power to distinguish between them. Tanaka et al. (2010) reject that preferences are purely hyperbolic, quasi-hyperbolic, or exponential.

⁸Dellavigna and Malmendier (2004) study how the firm’s profit-maximizing contract varies with consumer time preferences. Opp and Zhu (2015) study the implications of agent impatience for dynamically self-sustaining agreements when agents can renege on agreements, e.g., settings with upfront payment to workers.

⁹Clark (1994) finds anecdotal evidence that 19th century factory workers often shirked at the beginning of pay cycles, Oyer (1998) finds sales spike among US salespeople near the end of the year when bonuses are paid, and Kaur et al. (2015) show that piece-rate workers increase output as the weekly pay-day approaches.

2 Theoretical predictions

We now present a simple model of incentives to derive predictions for how two incentive contract features – dynamic-nonseparability and payment frequency – interact with time preferences. We consider a model with daylong periods. In each period, individuals experience a utility cost if they walk 10,000 steps (which might be negative) and receive utility from their other consumption in that period, which in our experiment will be consumption of mobile recharges:

$$U = \sum_{t=0}^{\infty} d_c(t) \left(c_t - e_t^{1(w_t=1)} \right)$$

The term e_t is the utility cost from walking 10,000 steps, i.e, the cost of complying with the program exercise target; w_t is an indicator for walking 10,000 steps; c_t are the mobile recharges consumed on day t ; and individuals discount the cost of walking and consumption k days in advance by $d_c(k)$. Because the consumption amounts are small, we model utility as linear in recharges c_t for simplicity, but the model’s qualitative predictions are the same if we relax this assumption. We assume that walking costs e_t are independently and identically distributed (i.i.d.) with cumulative distribution function $F(\cdot)$, are separable from other consumption, and are known in advance. The individual’s problem is to choose c_t and w_t to maximize utility subject to a budget constraint.

Individuals are part of an incentive program where they earn payments for complying with the step target; their budget constraints thus depends on the incentive contract mapping exercise compliance to income. We consider first a separable, linear incentive contract; this contract specifies that walking on each day t will be rewarded with a financial incentive of size m in k_t days. Denote the total value of recharges received in period t as m_t . The form of the budget constraint depends on the availability of borrowing/savings technology. We consider two main cases:

1. **Can borrow and save at an interest rate r .** The lifetime budget constraint becomes $\sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t c_t = \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t m_t$. Thus, on day t , the value of receiving m in rewards k_t days in the future is $\left(\frac{1}{1+r}\right)^{k_t} m$, and so the individual chooses to walk as long as $e_t \leq \left(\frac{1}{1+r}\right)^{k_t} m$.
2. **No savings, borrowing, or storage.** In this case, consumption in a given period is equal to the total amount of recharges received in the period: $c_t = m_t$ in all periods. That would imply that in any given period t , an individual would choose to walk as long as the walking costs are less than the discounted value of consuming the mobile recharge reward k_t days in the future: $e_t \leq d_c(k_t)m$.

More broadly, one can accommodate both of these (and other)¹⁰ cases in a reduced form way by defining a reduced form discount factor parameter representing the amount by which individuals discount rewards received k periods in the future, which encompasses both their “primitive” discount rate and any financial frictions. We denote this discount factor as $d_m(k)$; in case 1, $d_m(k) = \left(\frac{1}{1+r}\right)^k$, whereas in case 2, $d_m(k) = d_c(k)$. Using this new notation, individuals choose to walk on day t as long as walking costs are less than the discounted value of the mobile recharges received for the walk: $e_t \leq d_m(k_t)m$. The probability of compliance with the step target on day t for a reward in k_t days is thus:

$$Pr(w_t = 1|\text{Linear}) = F(d_m(k_t)m). \quad (1)$$

Thus, in a separable, linear contract on which payments are received every X days for the previous X days’ worth of walking (e.g., a weekly contract where payments are received on the last day of the week), the expected days of compliance per payment cycle is:

$$E \left[\sum_{t=1}^X (w_t) | \text{Linear} \right] = \sum_{t=1}^X F(d_m(X-t)m) \quad (2)$$

Predictions regarding dynamic non-separability We now explore the effect of making the contract dynamically non-separable. We focus on a specific form of dynamic non-separability: a “dynamic threshold” wherein payment is a function of the number of periods of compliance in a given time range and there is a minimum threshold level of compliance below which no incentive is received. If agents achieve that threshold, they receive m per day of walking that week. We focus on dynamic thresholds because they are a simple, implementable form of dynamic-non-separability, but the prediction we demonstrate about the interaction between dynamic non-separability and time preferences holds for a broader set of non-separable contracts that display a “dynamic complementarity” (i.e., a period in which the payment for effort is increasing in future effort).¹¹ Note that the daily behavior incentivized in all contracts is to comply with a “static threshold” which asks participants to walk at least 10,000 steps in a given day; what differentiates the dynamic threshold contract is a cross-day or “dynamic” threshold dictating the minimum number of days on which the

¹⁰For example, this approach also nests the case where there is no storage of recharges and time preferences over consumption are domain-specific, so $U = \sum_{t=0}^{\infty} d_m(t)m_t - d_c(t)e_t^{1(w_t=1)}$.

¹¹We conjecture that having a dynamic complementarity is a necessary condition for the prediction to hold. The prediction would thus not hold for contracts that only contain “dynamic substitutabilities” (e.g., paid for at most one day of walking in a week.) Note that dynamic complementarity might not be a sufficient condition.

10,000 step target must be met in the week.¹²

Dynamic non-separability makes an individual’s decision to walk considerably more complicated, as the reward for compliance – and hence the decision to walk – depends on compliance on multiple of the days in the payment period. For simplicity, we illustrate how dynamic non-separability interacts with compliance and time preferences using a shorter payment cycle than used in our experiment: a two-day payment period. For the two-day threshold contract, if the individual complies with the step target on both days, she receives $2m$ on the second day; if she complies on only one day, she receives nothing.

Intuitively, the key difference between the dynamic threshold and linear contracts is that in linear contracts, in each period the individual compares the reward only with her (undiscounted) cost of effort today. Thus, conditional on discount factors over money $d_m(t)$, discount factors over consumption $d_c(t)$ do not affect the decisions to comply;¹³ she complies in period 1 if $e_1 < d_m(1)m$ and in period 2 if $e_2 < m$. The expected total days of compliance w_t over the payment period is thus:

$$E \left[\sum_{t=1}^2 (w_t) | \text{Linear} \right] = F(d_m(1)m) + F(m) \quad (3)$$

In the dynamic threshold contract, in contrast, discount factors over consumption and effort matter. In particular, the individual in period 1 makes a joint decision about whether it is worth it to walk in both periods in order to get paid on the second day, and so compares the present discounted value of effort across both periods with the rewards. On the first day, she complies if the present discounted cost of walking on both days, $e_1 + d_c(1)e_2$, is less than the discounted value of the reward $d_m(1)2m$ and she knows she will follow through on the second day. Restricting to the case where costs are positive (this does not affect the results but simplifies notation), she thus complies if both (i) $e_1 + d_c(1)e_2 < d_m(1)2m$, and (ii) $e_2 < 2m$.¹⁴ Importantly, condition (i) is more likely to be satisfied if agents discount future effort more. On the second day, the agent complies with the step target if she has already walked on the first day, and condition (ii) above holds. The agent’s expected total

¹²While dynamic thresholds interact with time preferences, this is not necessarily true of static thresholds. For example, any step target, such as 10,000 steps, involves a minimum “static threshold” required for payment. Since all steps have to be completed in a single period, however, the performance of static thresholds does not depend on time preferences, although the demand for static thresholds may (Kaur et al., 2015).

¹³As described in detail in Section 4, we measure both discount factors (over walking and over recharges) in our sample; the correlation between discount factors over recharges and consumption is low and not significant, and the sample-average discount rate over walking is much higher than over recharges, suggesting that the discount rate over recharges mainly represents the interest rate here.

¹⁴With negative costs, she also walks in period t if $e_t < 0$ regardless of whether the other conditions are satisfied.

compliance in the 2-Day dynamic threshold contract is thus:

$$\begin{aligned}
E \left[\sum_{t=1}^2 (w_t) | \text{Dynamic Threshold} \right] &= 2P(e_1 + d_c(1)e_2 < d_m(1)2m \text{ AND } e_2 < 2m) \\
&= 2 \int_{-\infty}^{2m} \int_{-\infty}^{d_m(1)2m - d_c(1)e_2} f(e_1)f(e_2)de_1de_2 \\
&= 2 \int_{-\infty}^{2m} F(d_m(1)2m - d_c(1)e)f(e)de \tag{4}
\end{aligned}$$

Whether total compliance with a 2-day dynamic threshold (Equation 4) is larger than the compliance in the linear contract (Equation 3) depends on the distribution of walking costs; thus, the effect of adding a threshold to a linear contract on overall compliance with the step target is theoretically ambiguous. However, we can show the following prediction:

Prediction 1. *Compliance in the dynamic threshold contract relative to in the linear contract is increasing in the discount rate over walking (i.e., decreasing in the discount factor over walking, $d_c(k)$).*

This follows directly from inspection of equations 3 and 4. As the discount factor decreases, the present discounted cost of walking on days 1 and 2 decreases, increasing the probability of walking. In other words, individuals who discount future walking heavily have a lower total discounted cost of reaching the dynamic threshold, and thus higher compliance.

In our experiment, we test prediction 1 by randomly varying whether the contract has a dynamic threshold, and testing for heterogeneity based on the discount rate over walking.

Note that prediction 1 holds for both time consistent and time inconsistent time preferences. Although this might seem like an artifact of our focus on a 2-period model, that statement also holds in longer models. Note also that equation 4 assumes that agents are “sophisticated” about the fact that the relative value of their future effort compared with their future incentive payment may be different from the point of view of their period 2 selves than it is for their period 1 selves. In particular, from period 1’s perspective, the agent would want her period 2 self to comply if $e_2 < \frac{d_m(1)}{d_c(1)}2m$ whereas in period 2, the agent will in fact comply if $e_2 < 2m$; these are only equivalent if $d_m(1) = d_c(1)$, for example if there is no borrowing. However, even if agents were “naive” and assumed that their period 2 relative tradeoff between effort and payoffs would be the same as in period 1, Prediction 1 would still hold: equation 4 would become $E \left[\sum_{t=1}^2 (w_t) | \text{Dynamic Threshold} \right] = P(e_1 + d_c(1)e_2 < d_m(1)2m \text{ AND } e_2 < 2m) + P(e_1 + d_c(1)e_2 < d_m(1)2m)$ which is also decreasing in $d_c(1)$. Interestingly, not only does Prediction 1 hold for naives, but it holds for them even more strongly than for sophisticates: with the dynamically non-separable contract, for a given set of discounting parameters d_m and d_c , the naif’s compliance is higher than the sophisticate’s

as her overoptimism about future compliance makes her more likely to comply today.¹⁵ The naive agent will thus also have a relatively higher gap between non-separable and separable than sophisticates, as sophistication and naivete do not affect compliance under separable contracts.

Note finally that, for many cost distributions, Prediction 1 also holds for dynamic threshold contracts exhibiting threshold levels less than 100% (e.g., if the agent has to comply at least 5 days out of a 7-day payment period to receive payment, as in our experiment).¹⁶ The intuition is the same as above: those who discount future walking cost still have a lower discounted total cost to achieve the threshold level. Again, the intuition holds for both time-consistents and time-inconsistent, and for both sophisticates and naives, although for thresholds less than 100% there is no longer a clean prediction that naives will do better than sophisticates with the dynamic non-separable (just that for both sophisticates and naives relative compliance in the non-separable contract decreases in $d_c(k)$).¹⁷

Predictions regarding payment frequency We now return to the linear contract setup to analyze the effects of changing the frequency of payment. Using equation 2, we can make three predictions, all quite intuitive.

Prediction 2. *If agents are “impatient” over the receipt of financial rewards (i.e., if $d_m(k) < 1$ and is decreasing in k), compliance is increasing in the payment frequency. If agents are patient ($d_m(k) \approx 1$), payment frequency does not affect compliance.*¹⁸

This follows from equation (2): the likelihood of walking is increasing in the discount factor over rewards $d_m(k)$, and increasing payment frequency weakly decreases the delay to payment k_t on each day t .

Prediction 3. *The quantitative effect of increasing the payment frequency depends not just on average discount rates but on the shape of the agent’s discount factor over time.*

¹⁵Note that we are following O’Donoghue and Rabin (1999b) in defining a sophisticated agent as one who “knows exactly what her future selves’ preferences will be” and a naive one as one who “believe(s) her future selves’ preferences will be identical to her current self’s.”

¹⁶We are in the process of proving under which specific cost functions it holds.

¹⁷Interestingly, one might expect that, when the threshold level decreases, the dynamic thresholds would stop working for naives because they would start to procrastinate. However, in simulations and analytical examples, this does not appear to be the case. The reason is as follows. The classic situation in which naives procrastinate is when agents have to complete just one task; in that case, current effort is always a substitute with future effort and so naivete (which increases perceived future effort) always decreases current effort. In contrast, with dynamic thresholds, in some periods current effort is a substitute with future effort, but in many periods it is a complement with future effort. When it is a complement, naives do better; when it is a substitute, sophisticates do better. In simulations, these two forces often cancel, leading dynamic thresholds to work similarly for naives and sophisticates.

¹⁸The prediction for patient agents relies on the linearity assumption: if utility were concave and there were no storage, then more frequent payments could still increase the likelihood of walking through a concavity channel, as higher frequency would mean the rewards were broken up into smaller tranches. However, linearity does not affect the important comparative static of payment frequency with respect to patience.

Figure 1 shows how discount factors might change with the lag length for four different potential shapes used in the literature: Quasi-hyperbolic (“beta-delta”), hyperbolic, exponential impatient, and exponential patient (with the former two time-inconsistent and the latter two time-consistent). One can see that under the models where discount factors decay gradually over time (hyperbolic or time-consistent impatient), there could be large gains to switching from low-frequency (e.g., monthly) to medium-frequency (e.g., weekly) payments. In contrast, in a quasi-hyperbolic model, where the biggest difference is between “the present” and “the future,” there would only be big gains to increasing frequency if payment could be made within the “beta window” (often modeled as 1 day, which would require daily payments.) Given that paying within the “beta window” could be costly or infeasible in some settings, it is important to distinguish between these scenarios.

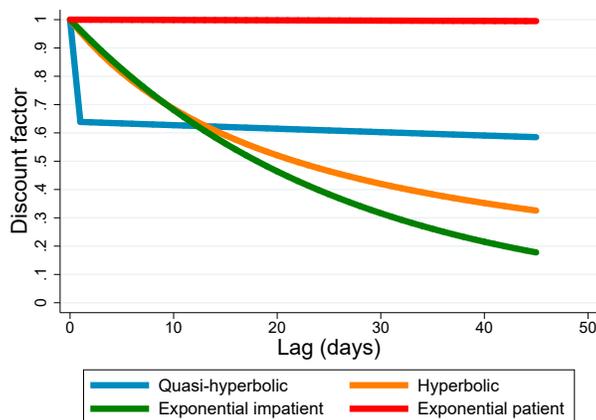


Figure 1: Hypothetical discount factors

Note: Figure displays hypothetical discount factors as a function of lag length under different models of discounting.

As a result, our experimental design will test the efficacy of three payment frequencies – monthly, weekly, and daily – in order to answer the question of whether and what type of increases in payment frequency improve compliance. Our three frequencies also allow us to explore which discount factor model for payments best fits the data, with the overall magnitude of frequency effects informing our understanding of the overall level of discounting, and the relative effects of moving from monthly to weekly frequency, and from weekly to daily frequency informing our understanding of the shape of time-preferences. A final prediction allows us to use our experiment to shed further light on the model of discounting.

Prediction 4. *If the discount factor over payments is decreasing in k and agents are paid every X days with $X > 1$, then compliance will increase as the “payday” (e.g., the end of the week if agents are paid weekly) approaches.*

This follows from equation (1) since the time to payment decreases as the payment date approaches.

3 Study Setting and Experimental Design

India is facing a diabetes epidemic. The presence of 60 million diabetics and 77 million pre-diabetics in the country has large economic and social implications. In 2010, diabetes imposed an estimated cost of \$38 billion – 2 percent of India’s GDP – on the healthcare system, and led to the death of approximately 1 million individuals (Tharkar et al., 2010).

There is widespread agreement that lifestyle changes are essential for managing the burden of diabetes, but existing strategies to promote change have had limited success. In particular, increased physical activity can prevent diabetes, and help the diagnosed avert serious (and expensive) long-term complications such as amputations, heart disease, kidney disease, and stroke. Despite the Indian government’s current efforts to address the epidemic, through its “National Programme on Prevention and Control of Diabetes, Cardiovascular diseases and Stroke (NPCDS),” adoption of the recommended lifestyle changes is low, according to physicians. Our study was conducted in partnership with and partially funded by the Government of Tamil Nadu, one of India’s southern states, who were interested in evaluating and scaling up effective strategies to promote disease management among diabetics.

3.1 Sample Selection and Pre-Intervention Period

Our project is a collaboration with the government of the state of Tamil Nadu to encourage lifestyle changes, specifically exercise, among those with or at risk of Type 2 diabetes. We selected our sample through a series of public screening camps in the city of Coimbatore, Tamil Nadu. In order to recruit diverse socioeconomic groups, the camps were held in locations ranging from the government hospital to markets, mosques, temples, and parks. During the camps, trained surveyors took health measurements; discussed each individual’s risk for diabetes, hypertension, and obesity; and conducted a brief eligibility survey. In order to be included in the study, individuals were required to have elevated blood sugar or have been diagnosed with diabetes, have low risk of injury or complications from regular walking, be capable with a mobile-phone, and be able to receive personal rewards in the form of mobile recharges.¹⁹ Within a week of attending a screening camp, eligible individuals were contacted by phone and invited to participate in a program to encourage walking.

¹⁹The full list of eligibility criteria was that the respondent must: either be diabetic or have elevated Random Blood Sugar, or RBS, (> 130 if haven’t eaten, > 150 if have eaten in previous 2 hours); be 30-65 years of age; have a prepaid mobile number which is used solely by them and without an unlimited calling pack; be literate in Tamil; be physically capable of walking half an hour; be currently living in Coimbatore city; not be pregnant; not be currently receiving insulin injections for diabetes; not be suffering from blindness, kidney disease or foot ulcers; not have had medical conditions such as stroke or heart attack; and not have been diagnosed with Type 1 diabetes.

Surveyors visited the potential participants at their homes or workplaces in order to conduct an initial baseline health survey and enroll participants in a one-week phase-in period. During the baseline health survey, surveyors collected detailed health, fitness, and lifestyle information. Surveyors then prepared respondents for the phase-in period, which was designed to collect baseline walking data and to familiarize participants with pedometer-wearing and step-reporting. Surveyors first demonstrated how to properly wear and read a pedometer. Next, they demonstrated how to report steps to our database by either responding to an automated call or directly calling into the system, and how to check text messages sent by the reporting system (as explained in Section 3.3, we created this automated calling system for respondents, who typically lack internet access, to self-report their daily steps). After the demonstration, respondents were asked to consistently wear a pedometer, and to report their steps each day through the automated call system for the weeklong phase-in period.²⁰

Following the phase-in period, surveyors again visited respondents to sync the data from the pedometers, and conducted a baseline time-preference survey.²¹ In the time-preference survey, surveyors elicited time preferences with a series of choices in the two domains relevant for our intervention: walking and mobile recharges (the financial reward we used to incentivize walking). The choices follow the Convex Time Budget (CTB) methodology pioneered by Andreoni and Sprenger (2012a) and Andreoni and Sprenger (2012b), and applied in many studies, including Andreoni et al. (2016); Augenblick et al. (2015); Augenblick and Rabin (2017); Carvalho et al. (2016); Giné et al. (2017).

Finally, the participants were randomly assigned to participate in one of two comparison groups (a monitoring group that received pedometers during the intervention period and a control group that did not), or to one of six incentive contracts for walking. All participants who withdrew or were found ineligible for the study prior to randomization were excluded from the sample, leaving a final experimental sample of 3192 individuals.

3.2 Experimental Design

3.2.1 The Daily Step Target

Our interventions center around encouraging participants to walk at least 10,000 steps a day. We chose this daily step target to match exercise recommendations for diabetics. The choice of a daily target reflects the fact that research organizations like the Center for Disease Control (CDC) and American Diabetes Association (ADA) recommend daily exercise sessions with no more than two consecutive days of rest. The target choice of

²⁰Respondents received a small cash reward of 50 INR at the end of the phase-in period for consistently wearing their pedometers and reporting their steps.

²¹Surveyors first used the Fitbit web application to automatically sync the actual walking data from the phase-in week to an online step database. They compared actual steps to reported steps, and reviewed the step-reporting processes as needed, before administering the time-preference survey.

10,000 steps approximates the number of steps that our average participant would take if he added the exercise routine recommended by the CDC and ADA to his existing behavior.²² In addition, 10,000 steps per day is a widely quoted target among health advocates and a common benchmark in health studies, making our choice consistent with existing literature and standard advice.

3.2.2 Treatment Groups

Participants were randomized into the incentives group or one of two comparison (non-incentive) groups:

1. Incentives: Receive a pedometer and incentives to reach a daily step target of 10,000 steps.
2. Monitoring: Receive a pedometer but no incentive contract.
3. Control: Receive neither a pedometer nor an incentive contract.

Within the incentives group, we randomized participants into one of six incentive contracts for walking. All treatments are summarized in Figure 2 and further elaborated below. The randomization was stratified by baseline Hba1c (a measure of blood sugar control) and a simple survey-based measure of impatience, using a randomization list generated in Stata.²³ Treatment groups were not of equal size: the size of each treatment group was chosen to ensure power to detect health impacts of the pooled incentives treatments relative to the comparison treatment, and the interactions between particular baseline characteristics and incentive contract features.

²²In particular, daily exercise recommendations for diabetics translate into approximately 3,000 steps of brisk walking per day (Marshall et al., 2009). In our sample, the average participant does not walk for exercise, but completes 7,000 steps per day. Our daily target is the sum of average daily pre-intervention steps plus the steps needed for daily recommended exercise.

²³Specifically, participants were stratified into four cells according to whether their baseline Hba1c was greater than 8 mmol/mol, and whether the average of their answer to the question “On a scale of 1 to 10, how patient are you?” at screening and baseline is greater than 6.5.

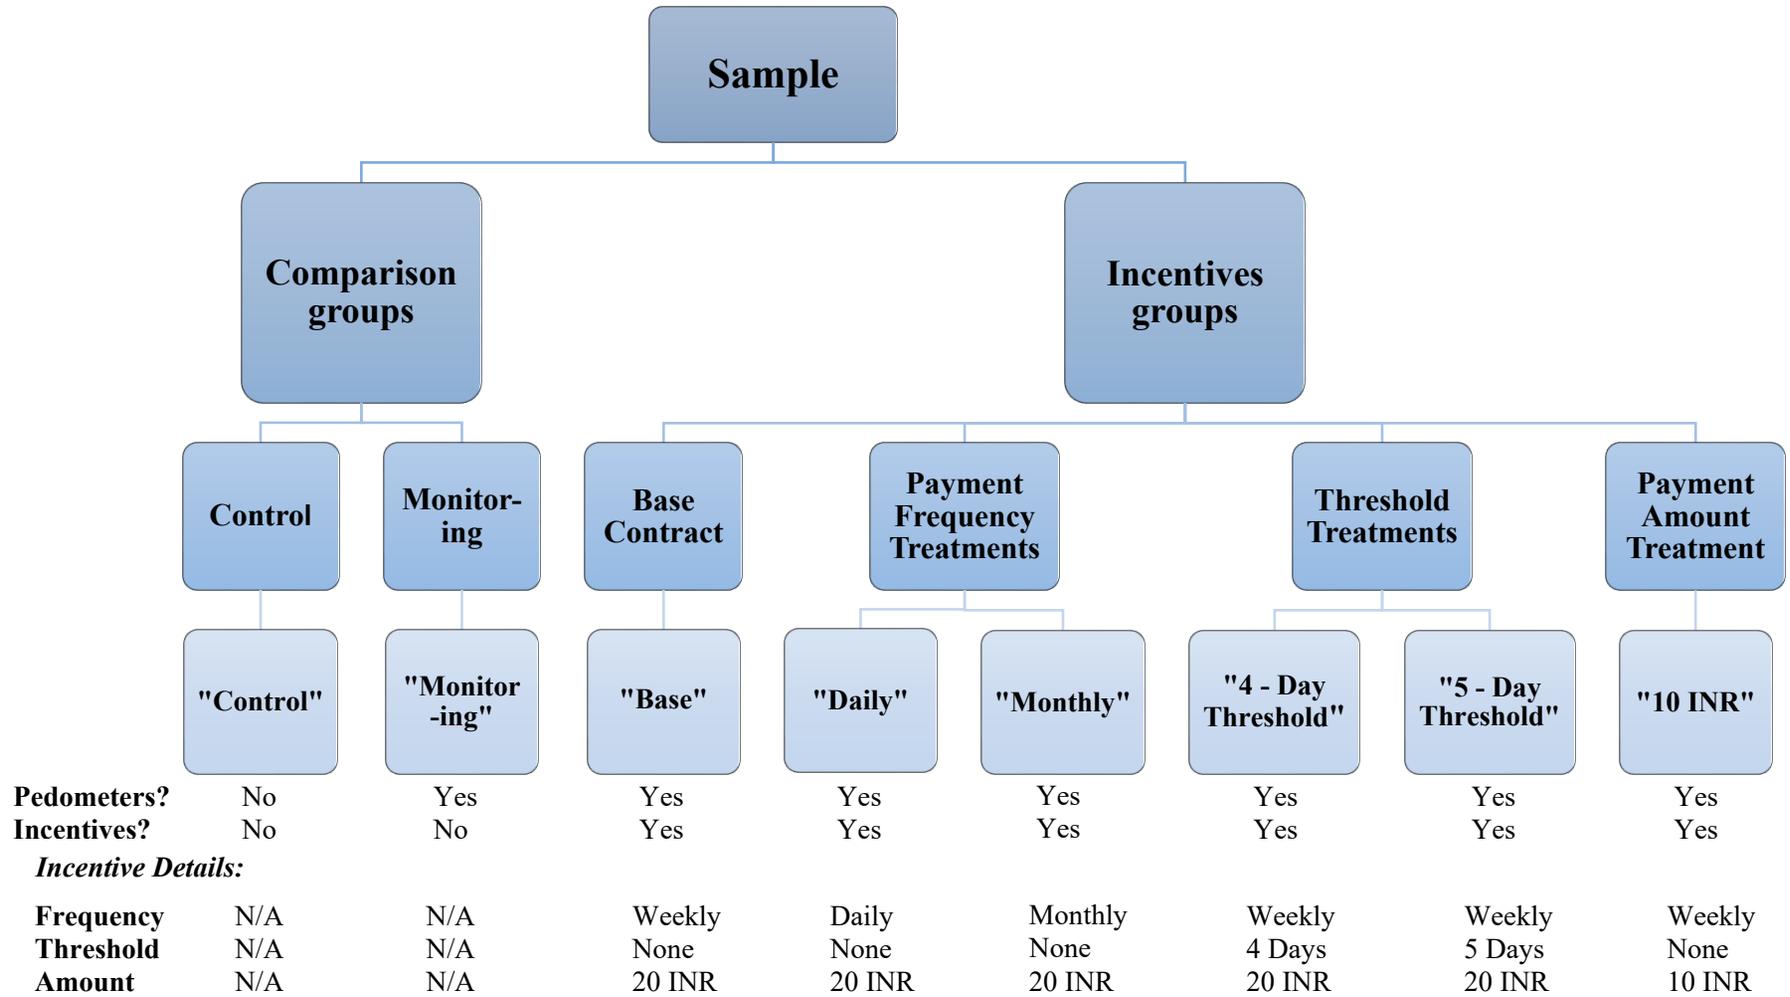


Figure 2: Experimental Design

Incentives Groups

All incentives groups were rewarded for accurately reporting steps above the daily 10,000 step target through the automated step-reporting system. As in the phase-in period, this step-reporting system called participants every evening (participants could choose a call time at the beginning of the intervention period), and prompted them to enter their daily steps as shown on the pedometer. Participants also had the option to call in their steps to a dedicated phone line at any time. The step-reporting system sent immediate text-message confirmations of each step report, and weekly text messages summarizing walking behavior.

During the explanation of the incentive contract, surveyors explained the step target to participants in the context of health recommendations, saying: “Remember that doctors recommend that you walk at least 10,000 steps a day, and more is always better! We recommend that you try to walk at least 10,000 steps a day and build up.”

The threshold treatments implicitly gave participants a goal of how many days to walk per week. To control for these goal effects, surveyors verbally encouraged participants in all treatment groups to walk at least 4 or 5 days per week at contract launch.

The Base Case Incentives Group We vary three dimensions of the payment: frequency, linearity/non-separability, and amount. The base case incentives group serves as our “base contract” or comparison group for all other incentives groups. To assess the responses to variation on each dimension, we compare the base case incentives group to a treatment group differing only along that dimension.

The base case incentives group was offered a separable, linear incentive contract awarding them mobile recharges worth 20 INR for each day they reported complying with the daily 10,000 step target. Recharges were delivered at a weekly frequency for each day the participant complied with the step target in the previous week.

Our next treatment groups differ from the base case incentive group in one of the following two dimensions that we predict will interact with time preference: payment frequency and whether the contract has a dynamic threshold.

Payment Frequency Two other treatment groups, the daily and monthly groups, differed from the base case incentives group only by the frequency of incentive delivery. In the daily group, recharges were delivered at 1am the same night participants reported their steps. In the monthly group, recharges were delivered every four weeks for all days of compliance in the previous four weeks. Comparing walking behavior in these two groups with the base case incentives group allows us to assess the role of payment frequency and time preferences over mobile recharges in incentive effectiveness. See Figure 2.

Dynamic Thresholds Two other treatment groups, the 4-day threshold and the 5-day threshold groups, differed from the base case incentives group only by the minimum threshold of weekly step-target compliance required before an incentive was paid. The base case incentives group’s contract was separable across days: participants received 20 INR for each day of compliance, while the threshold contracts were dynamically non-separable. The 4-day threshold group received mobile recharges worth 20 INR for each day of compliance if they exceeded the target at least 4 days in the weeklong payment period. So, a 4-day threshold participant who exceeded the step target on only three days in a payment period would receive no reward, but a participant who exceeded the step target on four days would receive mobile recharges worth 80 INR at the end of the week. Similarly, the 5-day threshold group received mobile recharges worth 20 INR for each day of compliance if they exceeded the target at least 5 days in the week.

Recall that, to control for goal effects, at the start of the intervention period, surveyors verbally encouraged participants in all treatment groups to walk at least 4 or 5 days per week. For those in the threshold groups, the target days-per-week was the same as their assigned threshold levels; for those in the other groups, the target days-per-week was randomly assigned in the same proportion as the threshold participants are divided between the 4- and 5-day threshold groups.

Payment Amount Finally, we included a small-payment treatment group that differed from the base case incentive group only by the amount of incentive paid. The 10-INR group was offered an incentive contract awarding them mobile recharges worth 10 INR, instead of the base-case 20INR, for each day they reported exceeding the daily step target. This treatment was included to help us learn about the distribution of walking costs, and to benchmark the magnitude of our other treatments effects (following for example Bertrand et al. (2005) and Kaur et al. (2015)).

Control Groups

We include two control groups in our experiment, a monitoring group and a pure control. In order to measure the overall health effects of the incentives program, we compare outcomes from endline surveys between the pooled incentives treatments to outcomes in the pure control treatments.²⁴

Monitoring The monitoring group allows us to isolate the effects of incentives alone. The monitoring group was treated identically to the incentives groups, but for the fact that monitoring participants did not receive incentives. In particular, they received pedometers and were encouraged to wear the pedometers and report their steps every day through the

²⁴Our experiment was not powered to detect differences in health outcomes between the control and monitoring groups, but we report these comparisons nonetheless.

step reporting system.²⁵ To control for the possibility that incentives may increase the salience of walking behavior, monitoring participants received daily confirmations of their step reports, and weekly text messages summarizing their walking behavior. In order to control for the effect of step goal-setting that an incentive for 10,000 daily steps may bring, monitoring and incentive treatment participants are given the same verbal step target of 10,000 daily steps at contract launch, and the same encouragement to walk at least 4 or 5 days per week.

Pure Control The pure control group allows us to measure the impact of all those aspects of the incentive treatments that were necessary for operating a walking incentives program in our setting, excluding the incentives themselves. Participants in the pure control group returned their pedometers at randomization (after the one-week phase-in period), but, because we wanted to net out any effects due to survey visits related only to research needs, still received regular visits from the survey team at the same frequency of the pedometer sync visits. Thus, the difference between the pure control and incentives groups includes the effect of incentives bundled with the effect of receiving a pedometer, but excludes the effect of the regular survey visits. Because any feasible incentive program would bundle the “monitoring” effect of a pedometer with the effect of incentives, the pure control group is a useful benchmark from a policy perspective: the difference between the pure control group and the incentives groups measures the total effect of a walking incentives program, including the effects that come simply from participants utilizing a step monitoring technology.²⁶

3.3 The Intervention Period and After

After randomization, all participants in the experiment were given a contract that detailed the specifics of the treatment group they had been assigned to, and also outlined the evaluation activities entailed for the rest of the study. A trained surveyor walked them through the contract and answered any of their questions to make sure it was clear.

In order to determine the number of steps taken, we gave those assigned to the monitoring and incentive groups Fitbit Zip pedometers for the duration of the intervention.²⁷ Although

²⁵Participants in all incentives groups and the monitoring group received a cash bonus of 200 INR for regularly wearing the pedometer and reporting their steps at the endline survey. In addition, if participants did not report steps for a number of days, the system would send them messages asking them to please report their steps regularly.

²⁶To accommodate a request from our government partners, we also cross-randomized one additional intervention a small sub-sample. In particular, 10% of the sample, cross-randomized across all other treatments, received the “SMS treatment,” which consisted of weekly text-message-based reminders to engage in healthy behaviors for diabetes such as eating right and exercising, adapted from another SMS program that had been shown to be successful in the Tamil Nadu region for diabetes prevention (Ramachandran, 2013). We control for the presence of the SMS in our main regressions and the results of this treatment are shown in the online appendix.

²⁷We chose Fitbit Zip pedometers due to their wearability, long memory, and relatively simple process for

these pedometers could be synced to a central database with an internet connection, most participants did not have regular internet access and so these data were not available in real time. Instead, we asked participants to report their daily step count to an automated calling system every evening. Incentive deliveries, i.e., mobile credits, were based on these reports. To verify the reports, we visited participants every two to three weeks to manually sync their pedometers and discuss any discrepancies with them. Anyone found to be chronically over-reporting was suspended from the program. All empirical analysis is based on the synced data from the Fitbits, not the reported data.

We visited all participants three times during the twelve-week intervention period. The primary purpose was to sync pedometers, but we also conducted short surveys to collect biometric and mobile phone usage data (we conducted these visits even with those participants who did not have a pedometer). We conducted a slightly longer midline survey at the second sync visit. Following the twelve-week intervention period, we conducted an endline survey. At endline, surveyors again collected detailed health, fitness, and lifestyle information. The timeline of the full intervention is outlined in Figure 3.

4 Data and Summary Statistics

4.1 Baseline Data: Health, Walking, and Time Preference

In the paper, we use three datasets of baseline characteristics: a baseline health survey, a week of baseline walking data, and a time-preference survey. The baseline health survey, conducted at the first household visit, contains information on respondent demographics, as well as health, fitness, and lifestyle information. Health measures include Hba1c, a measure of blood sugar control over the previous three months and the most commonly used measure of diabetes risk; random blood sugar, a measure of more immediate blood sugar control; BMI and waist circumference, two measures of obesity; blood pressure, a measure of hypertension; and a short mental health assessment. The baseline also includes two fitness measures (time to complete 5 stands from a seated position, and time to walk 4 meters), and lifestyle information including information on dietary, exercise, and substance use habits. Then, during the phase-in period between the baseline health survey and randomization, we collected one week of pedometer data, consisting of daily step counts.

Following the phase-in period, we conducted a baseline time-preference survey.²⁸ The survey adapts the convex time budget (CTB) methodology of Andreoni and Sprenger (2012a) to measure time preferences in two domains, walking and mobile recharges, which correspond

syncing data to a central database.

²⁸This survey was split temporally from the baseline survey both to avoid survey fatigue and because it was easier to measure time preferences over walking with participants who had used pedometers already so had a sense of what steps mean.

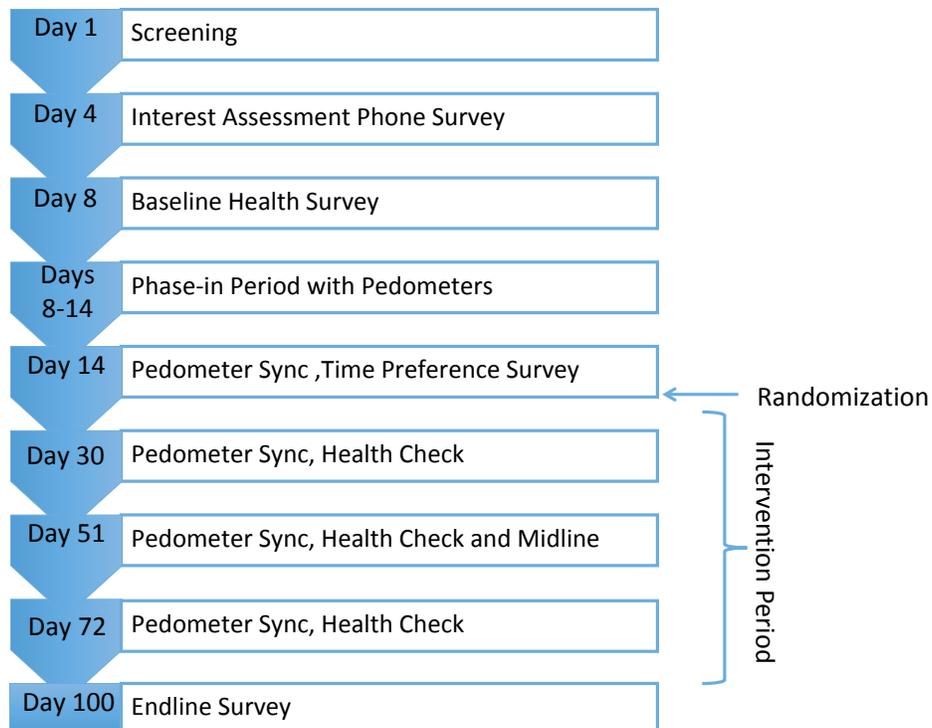


Figure 3: Experimental Timeline for a Sample Participant

Notes: This figure shows a representative experimental timeline for a participant in the experiment. Screening camps occurred throughout Coimbatore, Tamil Nadu from January 2016 to October 2017. In practice, visits were scheduled according to the availability of the respondent, leading to variation in the exact number of days between each visit. In addition, we intentionally introduced random variation into the timing of incentive delivery by randomly delaying the start of the intervention period by one day for selected participants. However, the intervention period was 12 weeks for all participants.

to $d_c(k)$ and $d_m(k)$ from Section 2. We asked participants to make a series of decisions allocating either recharges or steps on two dates: a “sooner” and “later” date. Each decision satisfies a budget set of the form

$$c_t + \frac{1}{r}c_{t+k} = m$$

where (c_t, c_{t+k}) are the chosen recharge amounts to be received or steps to be taken on the sooner and later dates, respectively. For example, when the interest rate is 1, participants are simply dividing a fixed budget of recharges or steps between two different dates. In general, those who are impatient over recharges will allocate more recharges towards the sooner date, and those who are impatient over walking costs will push more steps towards the later date. The sooner date t , the time lag between the sooner and later date k , and the interest rate r vary across decisions. Within a domain for a given respondent, the total budget m is fixed across allocations.

We used a participant’s decisions to construct two individual-specific parameter estimates of discount rates: one in the walking domain and one in the mobile recharge domain. In particular, in each domain we construct an estimate of the *Daily discount rate*, $\frac{1}{\delta} - 1$. This is a one-parameter estimate of the daily discount rate that is increasing both in time-consistent impatience and in present bias. Following Augenblick et al. (2015), our estimate is from a two-limit Tobit specification of the standard intertemporal Euler equation for an agent with an exponential daily discount factor δ , and concavity over recharges (or convexity over steps) α . Further details of the estimation methodology are described in Appendix A.^{29,30} However, because the discount rate can only be estimated for individuals making interior choices, we cannot estimate our structural parameter for all individuals in our sample. As a result, we supplement the estimation with several survey-based measures of impatience and time preference taken from the psychology literature in order to demonstrate whether the effects we see in the structural parameter sample extend beyond it; note that these are not specific to the domain of walking but are meant to proxy for discount rates over consumption.³¹

²⁹Our predictions are generally about overall impatience, not about whether an individual is time-consistent, and so we want one summary measure capturing impatience over the time horizon. Estimating just one parameter has the advantage of avoiding overfitting, which is relevant in our setting because we use fewer CTB allocations than some of the US-based methodological papers on CTB.

³⁰Other papers have also used the CTB data to estimate reduced-form measures (e.g., Giné et al. (2017) use the number of present-biased reversals). The structural measures have several advantages and so we choose to focus on them. First, the structural measure of impatience is increasing both in present-bias and in overall myopia, both of which are theoretically relevant to the performance of the contracts we offer. Second, whereas the standard reduced-form measures treat all preference reversals as equal regardless of magnitude, the structural measure takes into account the magnitude of impatience indicated by each decision. Third, by estimating the concavity (convexity) of preferences over recharges (steps), the structural measure avoids bias in the estimated discount rate from assuming linear utility (Andreoni and Sprenger, 2012a).

³¹Note that we only began measuring these latter measures partway through the data collection (at the point when we realized it was common for participants to choose non-interior solutions) and so the measures

Our CTB environment builds on a number of features from previous studies. First, the choices are made after the one-week phase-in period in which all participants have pedometers and report their daily steps, ensuring that participants are familiar with the costs of walking. This allows for meaningful allocations of steps between sooner and later dates. Second, the responses are designed to be incentive compatible; all respondents were informed that we would implement their choice from a randomly selected survey question. We set the probabilities such that for most respondents, the randomly selected survey question was a multiple price list of lotteries over money (which measures risk preferences), but for a few, a CTB allocation was selected. Because the allocations might have interfered with any walking program offered, we excluded these respondents from the experimental sample. To ensure that participants complete the allocated steps, we offer a large cash completion bonus of 500 INR in the step domain if the allocation is selected to be implemented, and the steps are completed as allocated, with the bonus to be delivered 15 days from the date of the survey (which is 1 day after the latest “later” day used).

We take a number of precautions to avoid various potential confounds, including confounds reflecting fixed costs or benefits of taking an action, or confounds due to the time of day of measurement, as described in more detail in Appendix A. However, we were not able to fully address one potential confound to our estimates of time-preferences across individuals: variation across people in the cost of walking over time, or in the benefit of receiving a recharge over time. For example, an individual with a particularly busy week after the time-preference survey, and therefore relatively high costs to steps in the near-term relative to the distant future, will appear to be particularly impatient over steps in our data (he will wish to put off walking). An individual with a relatively free week just after the time-preference survey will instead appear particularly forward-looking (he will not wish to put off walking). The same concerns can also arise with recharges, but because recharges are storable and thus may be consumed on a different day than when they are received, we expect less variation (and heterogeneity) in the utility of recharge receipt over time.

4.2 Summary Statistics

The baseline characteristics of the full experimental sample are reported in the first column of Table 1. Our sample is on average 49.42 years old, and has slightly more males than females. Their average monthly household income is approximately 16,000 INR (about 200 USD) per month; for comparison, in 2015 the median urban household in India earned between 10,000 and 20,000 INR per month (Labor Bureau of India). Panel B shows that our sample is at high risk for diabetes and its complications: 67 % of the sample has been diagnosed with diabetes by a doctor, and 82 % have Hba1c levels which are strongly indicative

are available for only part of the sample.

of diabetes. The random blood sugar concentrations are also indicative of high diabetes risk. Note that Hba1c above 6.5 is considered diabetic, and RBS above 180 (even just after eating) is unlikely except among diabetic individuals; average Hba1c and RBS in our sample surpass both of these cut-offs. The sample also has high rates of common diabetes comorbidities: 41 % have hypertension (defined as systolic blood pressure above 140 or diastolic blood pressure above 90), and 61 % are overweight (defined as BMI above 25) at baseline.

Panel C shows that although baseline walking levels are below international daily walking recommendations of 10,000 steps per day, they are comparable to the average steps taken in many developed countries. On average, participants walked just under 7000 steps per day in the phase-in period. For comparison, Japanese adults also take approximately 7,000 steps per day, whereas adults in the United States take approximately 5,000 steps per day, and adults in western Australia take about 9,000 steps per day (Bassett et al., 2010).

Panel D of Table 1 reports measures of impatience measured using the CTB survey questions. First, we do not see evidence of impatience over recharges on aggregate. In particular, the average estimated daily discount rate over recharges is only 0.01, which is similar to monetary discount rates estimated using the CTB methodology in other settings (e.g. Andreoni and Sprenger (2012a) and Augenblick et al. (2015)). Second, individuals are quite impatient over steps: present-biased preference reversals are more common than future-biased reversals, and the average estimated daily discount rate over steps is 0.37. Our estimate of the average discount rate over steps is somewhat larger than effort discount rates estimated by Augenblick et al. (2015) using a similar CTB methodology. However, because the present time period in our CTB questions was only two hours whereas our future time periods were entire days, inflating the costs of present relative to future steps, it is likely that our measures of daily discount rates over steps are biased away from zero. In addition, the discount rates over steps show wide variation across individuals compared to Augenblick et al. (2015), suggesting that these estimates may be confounding heterogeneous differences in walking costs in the CTB allocation dates. Such a confound would add noise to our discount rate estimates over walking and attenuate the observed relationship between time preferences and incentive contract design. Appendix Table A.2 has more detailed summary statistics on our discount rate measures.

Baseline health and time preferences are similar across treatment groups. Columns 1 and 2 of Table 1 show means for the pure control and monitoring groups, and Columns 5-10 show means separately for each incentive group, with standard deviations in parentheses. To explore whether randomization provided balance in these characteristics across the different groups, we test that all characteristics are jointly orthogonal to treatment assignment relative to the pure control group (Hansen and Bowers, 2008). We fail to reject that the coefficients on

all characteristics are 0 in regressions of treatment assignment on characteristics, suggesting that balance was achieved.

4.3 Outcomes

Our outcomes are gathered from two datasets. The first is a time-series dataset of daily steps walked for each participant with a pedometer during the twelve-week intervention period. Because surveyors collect pedometers back from pure control participants after the phase-in period, we do not have daily steps for this group. Surveyors collect pedometer data at three separate “pedometer sync” visits during the intervention period and at the endline survey.³²

A potential issue with the daily step data is that we only observe steps taken while participants wear the pedometer. Because participants in the incentives groups are rewarded for taking 10,000 steps in a day with the pedometer,³³ they have an additional incentive to wear the pedometer on days that they expect to walk more. This could lead to a potential selection issue: if the incentives group selectively makes an effort to wear the pedometer when they think they will walk more but the monitoring group does not, then we will see a spurious positive relationship between incentives and observed daily steps.

In order to minimize selective pedometer-wearing, we incentivize all monitoring and incentives participants to wear their pedometers even on days with few steps. We do this by offering a cash bonus of 200 INR (About 3 USD) if participants wear their pedometer (i.e., have non-zero recorded steps) on at least 70% of days in the intervention period. This approach largely addressed the issue: Figure 4 shows that the rates of pedometer-wearing are high and similar between treatment groups. Although the difference between them statistically insignificant, to address the potential for differential selection, we report Lee (2009) bounds when comparing pedometer data between the incentives and monitoring groups.³⁴

The second outcomes dataset – the endline survey – gathered health, fitness, and lifestyle information similar to the baseline health survey, as well as information about dietary and exercise behavior changes made during the intervention period. These data are available for participants in all treatment groups, including the pure control group. The primary purpose

³²In order to collect pedometer data, surveyors ask to see the pedometer, open the Fitbit web application on a wifi-enabled tablet computer, sign into a respondent-specific account, and upload the previous 30 days of daily pedometer step data to the Fitbit database. We later pull these data through the Fitbit application program interface (API) using a web application we designed for this study.

³³Although incentives are delivered for steps reported, we cross-check step reports with actual pedometer data after every pedometer sync visit. Anyone found to be over-reporting is initially warned, and is eventually suspended from the program if the behavior continues.

³⁴Lee bounds assume that selection is monotonic, i.e. that being in the incentives group makes every participant weakly more likely to wear their pedometer on any given day or weakly less likely to do so. This assumption is likely justified in our setting, as being in the incentives group weakly increases the payoff to pedometer wearing by increasing the probability of reward.

Table 1: Baseline summary statistics in full sample and by treatment group.

	Averages of Baseline Characteristics by Treatment Group									
	Full Sample	Control	Monitoring	Incentives Pooled	Daily	Base Case	Monthly	4-Day TH	5-Day TH	10 INR
A. Demographics										
Age (from BL)	49.54 (8.52)	49.78 (8.19)	50.28 (8.95)	49.44 (8.55)	49.57 (8.60)	49.60 (8.33)	48.80 (8.94)	49.31 (8.68)	49.67 (8.77)	49.11 (7.84)
Female (=1)	0.42 (0.49)	0.46 (0.50)	0.43 (0.50)	0.41 (0.49)	0.44 (0.50)	0.41 (0.49)	0.38 (0.49)	0.42 (0.49)	0.38 (0.49)	0.48 (0.50)
Labor force participation	0.75 (0.44)	0.73 (0.45)	0.72 (0.45)	0.75 (0.43)	0.75 (0.43)	0.74 (0.44)	0.81 (0.39)	0.74 (0.44)	0.77 (0.42)	0.70 (0.46)
Daily mobile usage (INR)	6.61 (8.79)	7.22 (10.14)	6.47 (8.95)	6.44 (8.36)	5.86 (6.25)	6.58 (8.77)	7.67 (9.19)	6.43 (8.05)	6.01 (8.87)	4.94 (5.77)
Mobile balance (INR)	29.26 (49.42)	30.80 (48.79)	29.48 (48.68)	28.98 (49.88)	28.61 (38.54)	29.69 (52.08)	28.55 (63.65)	28.57 (49.10)	28.14 (44.98)	30.05 (36.59)
Per capita income (INR/month)	4463 (3638)	4488 (4483)	4620 (3160)	4447 (3447)	4068 (2765)	4477 (3496)	4599 (3235)	4454 (3590)	4480 (3525)	4341 (2615)
Private water source	0.67 (0.47)	0.65 (0.48)	0.68 (0.47)	0.67 (0.47)	0.66 (0.48)	0.69 (0.46)	0.63 (0.48)	0.65 (0.48)	0.70 (0.46)	0.67 (0.48)
Household Size	3.91 (1.62)	3.94 (1.54)	3.82 (1.51)	3.91 (1.64)	3.92 (1.45)	3.89 (1.70)	3.74 (1.59)	3.96 (1.64)	3.96 (1.68)	3.58 (1.29)
B. Health										
Diagnosed diabetic	0.67 (0.47)	0.67 (0.47)	0.68 (0.47)	0.66 (0.47)	0.62 (0.49)	0.68 (0.47)	0.62 (0.49)	0.67 (0.47)	0.68 (0.47)	0.59 (0.50)
Hba1c (mmol/mol)	8.68 (2.33)	8.67 (2.36)	8.76 (2.40)	8.68 (2.32)	8.58 (2.36)	8.72 (2.29)	8.66 (2.44)	8.68 (2.32)	8.69 (2.38)	8.35 (2.14)
RBS (mmol/L)	192.42 (89.39)	191.32 (88.73)	196.07 (86.67)	192.51 (89.87)	195.58 (91.54)	193.26 (88.25)	193.30 (98.14)	192.12 (89.96)	192.50 (91.75)	177.38 (77.00)
Systolic BP (mmHg)	133.35 (19.15)	133.33 (20.34)	134.06 (17.68)	133.34 (18.99)	135.25 (21.55)	133.27 (19.07)	134.18 (19.13)	132.49 (18.00)	133.71 (19.20)	135.62 (21.42)
Diastolic BP (mmHg)	88.47 (11.11)	88.54 (11.50)	88.53 (10.10)	88.46 (11.09)	89.30 (12.79)	88.19 (10.75)	88.60 (10.10)	88.23 (10.73)	89.01 (11.96)	90.00 (13.19)
BL BMI	26.42 (4.35)	26.52 (4.34)	26.47 (3.67)	26.40 (4.39)	26.41 (5.35)	26.47 (4.53)	26.39 (4.81)	26.34 (4.21)	26.19 (3.70)	26.99 (4.10)
HbA1c: Diabetic (=1)	0.82 (0.38)	0.82 (0.38)	0.81 (0.39)	0.82 (0.38)	0.77 (0.42)	0.84 (0.36)	0.79 (0.41)	0.81 (0.39)	0.82 (0.38)	0.77 (0.42)
BP: Hypertensive (=1)	0.41 (0.49)	0.36 (0.48)	0.42 (0.50)	0.42 (0.49)	0.45 (0.50)	0.41 (0.49)	0.44 (0.50)	0.42 (0.49)	0.43 (0.50)	0.29 (0.46)
Overweight	0.61 (0.49)	0.62 (0.48)	0.66 (0.47)	0.60 (0.49)	0.57 (0.50)	0.60 (0.49)	0.58 (0.50)	0.61 (0.49)	0.59 (0.49)	0.67 (0.48)
C. Walking										
Pr(exceeded step) target (phase-in)	0.25 (0.32)	0.25 (0.31)	0.24 (0.32)	0.25 (0.32)	0.25 (0.32)	0.23 (0.31)	0.27 (0.33)	0.26 (0.32)	0.25 (0.34)	0.27 (0.34)
Average daily steps (phase-in)	6960 (4004)	7015 (3972)	6868 (3729)	6960 (4038)	6961 (4240)	6795 (3984)	7403 (3889)	7069 (4047)	6916 (4111)	7018 (4195)
D. Time Preferences										
<i>i. Mobile Recharges</i>										
Discount rate	0.01 (0.09)	0.01 (0.08)	0.01 (0.07)	0.01 (0.09)	-0.01 (0.09)	0.01 (0.09)	0.02 (0.09)	0.01 (0.08)	0.01 (0.10)	0.01 (0.06)
<i>ii. Steps</i>										
Discount rate	0.36 (3.76)	0.21 (0.71)	0.28 (1.03)	0.41 (4.34)	0.18 (0.56)	0.47 (4.96)	0.27 (1.02)	0.49 (5.29)	0.29 (1.66)	0.11 (0.47)
F-test for Joint Orthogonality										
F-statistic			0.87	1.17	1.19	1.05	1.16	1.13	1.16	0.88
Sample size										
Number of individuals	3,192	585	203	2404	166	902	164	794	312	66
Percent of sample	100.0	18.3	6.4	75.3	5.2	28.3	5.1	24.9	9.8	2.1

Notes: Hba1c (glycated haemoglobin) and RBS (random blood sugar) respectively measure medium- and short-term blood sugar concentrations. Systolic and Diastolic BP measure blood pressure. Hba1c, RBS, and weight are risk factors for developing diabetes and diabetes-related complications. High blood pressure and weight are cardiovascular risk factors. In each domain (mobile recharges and steps), the discount rate $\frac{1}{\delta_i} - 1$ is an individual-level measure of impatience estimated from a two-limit Tobit regression with the restriction that $\beta_i = 1$ and $\alpha_i = \alpha$. The F-statistic tests the joint orthogonality of all characteristics to treatment assignment relative to control.

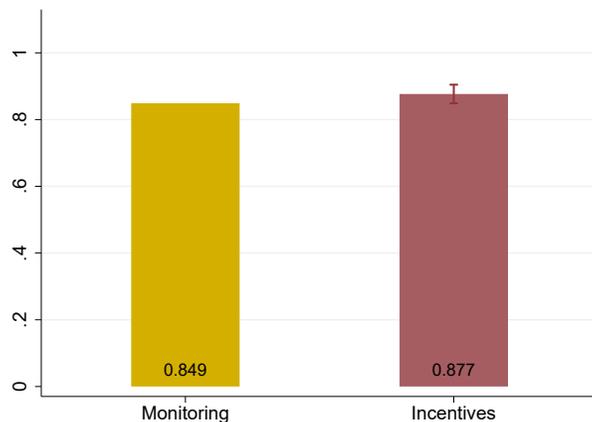


Figure 4: Fraction of days participants wore Fitbits

of these data is to allow us to study the impacts of the programs on health and healthy behaviors. Many of the measures (such as biometric health measures) have little margin for manipulation. There is no differential incentive for attending the endline survey across treatments, and although some respondents did withdraw from the study prior to the endline survey, Table D. 1 shows that endline attrition rates are not statistically distinguishable between the pure control, monitoring, and incentives groups.

5 Results: Incentive Design

This section examines the effects of our incentive contract variations on exercise, and explores the implications of our results for incentive design in the presence of impatience. We begin by establishing that providing incentives increases exercise in this context; if it did not, then this would not be a good laboratory to explore the effects of varying the contract. We then explore the implications of time preferences for incentive design, first exploring the role of frequency and then of dynamic thresholds.

5.1 Incentives and exercise

We begin by establishing that our incentives program impacts exercise. This finding not only suggests that we can use the experiment to explore the design of incentives, but is also of independent policy interest: exercise has been shown to benefit health for diabetics (Hill, 2005; Manders et al., 2010; Praet and van Loon, 2009; Qiu et al., 2014; Shenoy et al., 2010; Thomas et al., 2009; Zanuso et al., 2009) and ours is the first randomized trial of exercise incentives in a developing country.

We use intent-to-treat (ITT) estimates to assess the impact of incentives for daily walking on exercise. In particular, we compare average exercise outcomes in the pooled incentives treatment groups to those in the monitoring treatment group. We focus here on the Fitbit

exercise data as it is both less noisy than self-reported exercise and less prone to bias; as a result, we only evaluate the effect of incentives relative to monitoring, not control. Because monitoring may have an independent impact, this likely understates the policy impact of incentives overall, and these estimates should be interpreted as lower bounds on the effects of incentives relative to control. We return to exploring the effects of the monitoring group relative to the control group in Section 6.2.

For pedometer outcomes, which are measured at a daily frequency during the intervention period, we compare averages on the person-day level across treatment groups using regressions of the following form:

$$y_{it} = \alpha + \beta \times incentives_i + \mathbf{X}'_i \gamma + \mathbf{X}'_{it} \lambda + \varepsilon_{it}, \quad (5)$$

where y_{it} is either daily steps or an indicator for whether the individual surpassed the 10,000-step target, for individual i on day t during the intervention period; $incentives_i$ is an indicator for being in the incentives group; \mathbf{X}_i and \mathbf{X}_{it} are vectors of individual-level and day-level controls, respectively; and the standard errors ε_{it} are clustered at the individual level. Control vectors are described in the notes to Table 2. The coefficient of interest, β , is the ITT effect of incentives relative to the monitoring group. The results are shown in Panel A of Table 2, and without controls in Table C.2. The results are also shown graphically in Figure 5, where the confidence interval shown on the incentives bar is the 95% confidence interval for the gap between the incentives and monitoring groups (as is the case for all other graphs in this section).

We find that incentives have large and sustained impacts on walking during the intervention period. Incentives increase the number of days that participants reach their 10,000 step target, and the size of the effect is large: Column 1 of Table 2 shows that incentivized participants exceed their step target on 20 % more days than those in the monitoring group (the Lee bounds are 19.1% to 23.8%). These effects are not simply a result of participants shifting steps from one day to another: Column 2 shows that incentives increase walking by 1291 steps per day (the Lee bounds are 934 to 1509 steps). This effect is equivalent to approximately 13 minutes of extra brisk walking daily on average.

To examine the impacts of incentives on walking routines, Figure 6a shows histograms of the number of days the step target was met per week (i.e., each data point is a respondent \times week) in the monitoring and incentives groups. Relative to the monitoring group, the incentives group has a striking reduction in the number of weeks where the step target is never met and an equally striking increase in the number of weeks where the target is met on every day.

Figure 6b shows the impact of incentives on the distribution of daily steps. There is

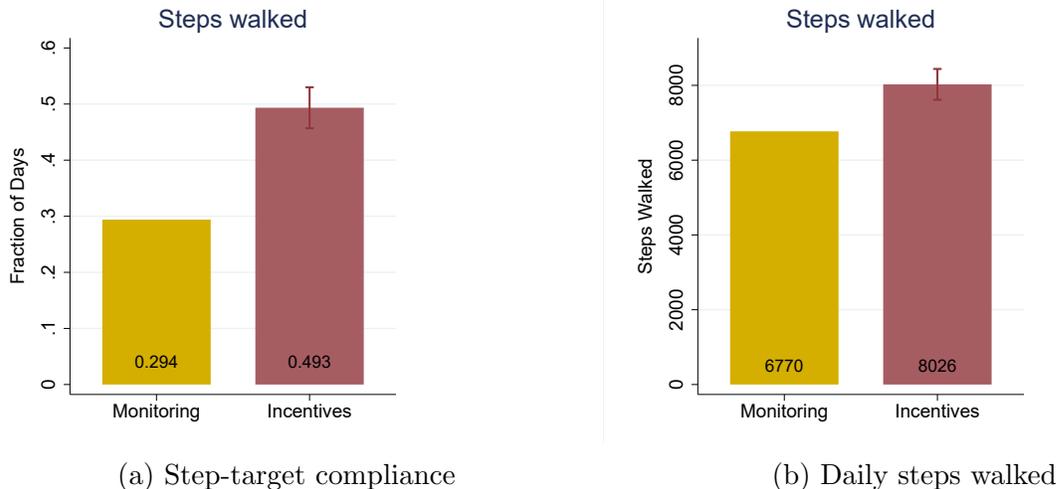


Figure 5: Incentives increase average walking

Note: Figure displays the impact of the pooled incentives treatments on steps during the intervention period, and the confidence interval for the test of equality between the incentives and monitoring groups. The dependent variable for the first panel is whether the participant met the daily step target of 10,000 steps on that day. The dependent variable in the second panel is the average steps walked.

bunching at 10,000 steps in both groups, but the bunching in the incentive group is more severe. Encouragingly, providing incentives also appear to shift the entire distribution of daily steps, rather than simply pushing marginal participants who would otherwise walk nearly 10,000 steps in a day over the 10,000-step target. There is less mass everywhere below the 10,000 step target, and more mass everywhere above.

Figure 7 shows that, after an initial spike at week 1, the effect of incentives on walking also does not attenuate over time, a rare finding in the literature with, for example, Patel et al. (2016) finding that physical activity drops steeply 5-7 weeks into a 12-week walking-incentive program.

Having established that our incentives affects behavior, we next use the experiment to explore the effectiveness of incentive contract variations designed to improve performance in the face of impatience over consumption and over rewards, respectively.

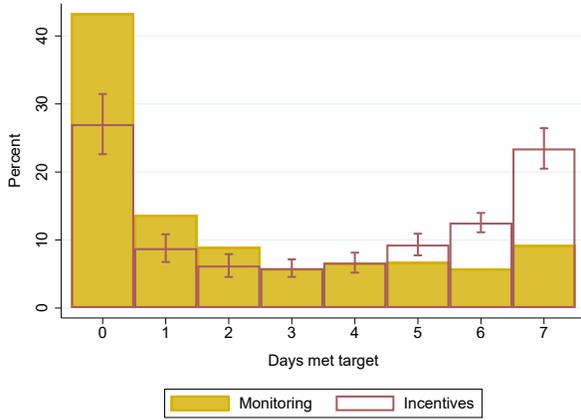
5.2 Dynamic Thresholds

Our primary prediction for the effect of the dynamic threshold contracts is that relative to linear contracts, they should improve performance for those who are impatient over consumption; we do not have any predictions for the average effect. However, it is still useful to analyze the average effect of thresholds, most notably because our experiment represents (to our knowledge) the first randomized evidence comparing a dynamically non-separable contract to one that is not. We thus begin by analyzing the average effect of dynamic threshold contracts relative to linear, before turning to test for heterogeneity in their impacts by time

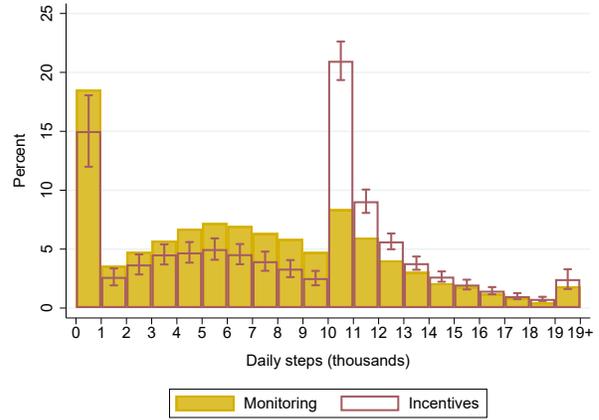
Table 2: Impacts of incentives on exercise

	Pedometer Data (Intervention Period)		
	Proportion of Days Achieved 10K Steps	Daily Steps	Daily Steps (conditional on being > 0)
<i>A. Pooled Incentives</i>			
Incentives	0.198*** [0.0181]	1291.1*** [208.2]	1160.9*** [186.5]
<i>B. Unpooled Incentives</i>			
Base Case	0.207*** [0.0198]	1405.6*** [221.0]	1200.6*** [197.5]
Daily	0.204*** [0.0306]	1165.9*** [326.2]	1211.8*** [273.4]
Monthly	0.182*** [0.0284]	1297.1*** [304.5]	1222.3*** [264.4]
4-Day Threshold	0.190*** [0.0204]	1222.5*** [228.8]	1118.4*** [203.1]
5-Day Threshold	0.210*** [0.0250]	1309.9*** [261.9]	1228.2*** [230.1]
10 INR	0.128*** [0.0382]	737.1* [382.7]	515.5 [327.8]
Monitoring mean	0.302	6822.658	7986.126
Controls	Yes	Yes	Yes
<i>P-value for Base Case vs</i>			
Daily	.93	.39	.96
Monthly	.31	.66	.92
4-Day Threshold	.27	.22	.52
5-Day Threshold	.87	.63	.87
10 INR	.03	.05	.02
# Individuals	2,563	2,563	2,559
Observations	203,245	204,614	180,418

Notes: We report pooled incentive effects in Panel A, and separately by incentive treatment group in Panel B. The columns show coefficient estimates from regressions based on Equations 5 (Panel A) and 6 (Panel B), using daily panel data from pedometers during the intervention period. The sample includes the incentives and monitoring groups. Individual-level controls include a second order polynomial of age, gender, a second order polynomial of weight, and the average of the dependent variable during the phase-in period (before randomization). Day-level controls include fixed effects for the month-year and day-of-week. The omitted category in all columns is the monitoring group. Standard errors, in brackets, are clustered at the individual level.



(a) Days per week exceeded step target



(b) Daily steps walked

Figure 6: Incentives shift the distributions of days walked per week and steps walked per day

Note: Figure displays the impact of the pooled incentives treatments relative to the monitoring group during the intervention period, with the confidence intervals representing the confidence interval for the test of equality between the incentives and monitoring groups.

preferences.

5.2.1 Average Effectiveness

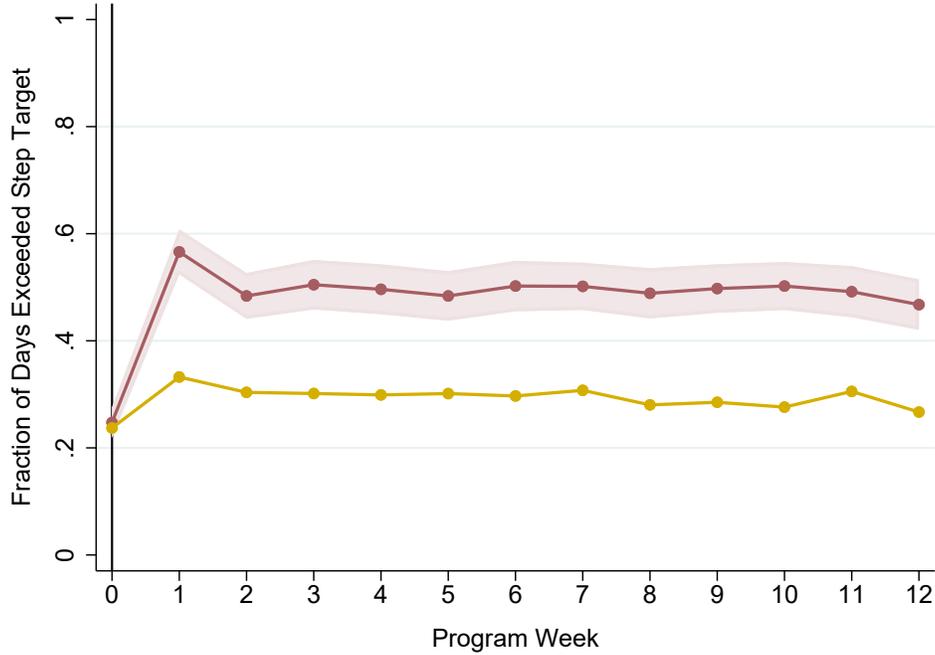
Panel B of Table 2 evaluates the ITT effects of all of our incentive contract variations, estimating regressions of the following form:

$$y_{it} = \alpha + \beta_j \times (\text{incentives}^j)_i + \mathbf{X}'_i \gamma + \varepsilon_{it}, \quad (6)$$

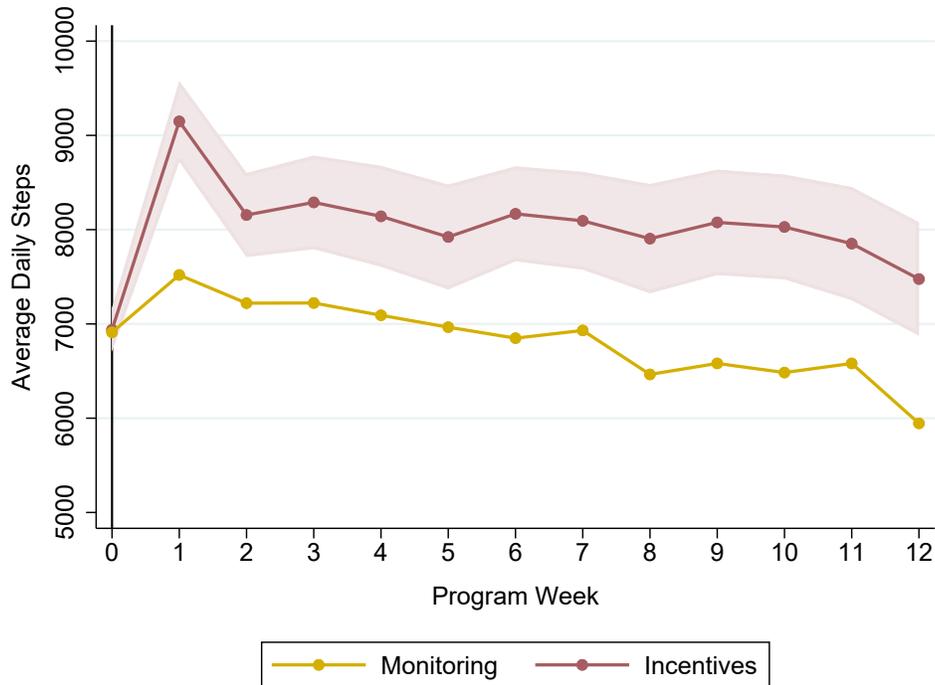
where y_{it} are daily walking outcomes and $(\text{incentives}^j)_i$ is an indicator for whether individual i is enrolled in incentive treatment group $j \in (\text{daily, base case, monthly, 4-day threshold, 5-day threshold, 10 INR})$. Recall that all treatments besides the base case incentive vary from the base case contract on exactly 1 dimension (dynamic separability, payment frequency, or payment amount); the bottom rows of the table thus show the p-values for the significance of the difference between each treatment group and the base case group.

We find that adding a dynamic threshold does not affect the level of exercise. Figure 8 and the fourth and fifth rows of Table 2 show that individuals in the 4-day threshold and 5-day threshold treatment groups exceed the 10,000 daily step target roughly as frequently as individuals in the base case incentive group (which has no threshold). For both threshold treatments, compliance with the step target is within 2 percentage points of compliance in the base case incentive (linear) treatment, with the difference not statistically significant. The number of steps are similar across treatments as well.

Figure 7: The effects of incentives persist throughout the 12-week program



(a) Step-target compliance



(b) Daily steps walked

Notes: Panel A shows the average weekly probability of exceeding the step target over time for the monitoring and pooled incentives groups, and Panel B shows the steps walked per day averaged over each weekly period. Week 0 is the phase-in period, before randomization. The intervention period runs from Week 1 - Week 12. The confidence intervals represent the confidence interval for the test of equality between the incentives and monitoring groups from a regression controlling for the same control variables as in Table 2.

However, the dynamic threshold contracts have an advantage. Individuals in the two threshold groups only receive a reward for exceeding the step target if they do so on at least 4 or 5 days in a weeklong payment cycle; when they walk on fewer than the threshold days, they are not rewarded. Because individuals with threshold contracts do not reduce overall walking, but are paid for a lower fraction of days walked, the threshold contracts we offer are more cost-effective than base case incentive contracts without a threshold.

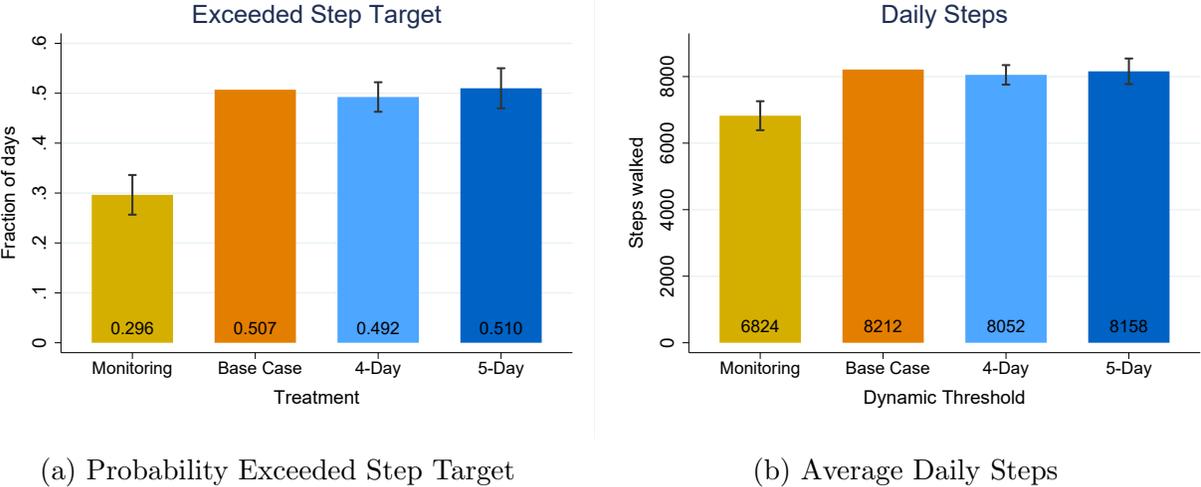


Figure 8: Adding a dynamic threshold does not significantly affect average walking

Notes: Figures compare the effects of the dynamic threshold treatments with the “base case” (linear) incentive treatment. Panel A shows the average probability of exceeding the daily 10,000-step target during the intervention period; Panel B shows average daily steps walked during the intervention period. The confidence intervals show the 95% confidence interval for a test of equality between the base case incentive group and each other treatment group, controlling for the same things as Table 2.

Table 3 quantifies the cost-effectiveness of all contracts in two ways. Column 5 shows the average incentive delivered on a day the participant exceeded the daily 10,000 step target. For all contracts other than the 4- and 5-day thresholds, which all pay out linearly, this is just the incentive amount. However, as Column 4 shows, in the 4- and 5- day threshold groups, participants are paid 91% and 86% of the days the achieve the step target, respectively. Thus, the incentive paid per day the target is reached is lower than in the base case (linear) group: 18 INR and 17 INR per day of compliance as compared to 20 INR. These cost savings of 9% and 14% are made while participants achieve nearly the same amount of walking. Column 6 shows the average incentive cost per additional day the step target was reached above the monitoring group. According to this metric, the 4-day threshold and 5-day threshold achieve cost savings of 6.1% and 11.0%, respectively. For comparison, the incentive amount per day walked is mechanically lower in the 10-INR treatment group, but this comes at the cost of reduced walking overall.

Table 3: Cost Effectiveness of Monitoring and Incentive Treatments

Cost-effectiveness of Incentive Contracts						
	Walking		Rewards		Cost-effectiveness	
	Compliance Treatment - Proportion of days met step target	effect on compliance relative to monitoring	Incentive Amount (INR)	Proportion Compliance Incentivized	INR per Day Complied	INR per Day Complied above Monitoring
	(1)	(2)	(3)	(4)	(5)	(6)
Monitoring	0.3	N/A	0	0	0	N/A
Daily	0.51	0.21	20	1	20	49.04
Base Case	0.51	0.21	20	1	20	49.44
Monthly	0.5	0.2	20	1	20	50.34
4-Day Threshold	0.51	0.2	20	0.9	18.08	44.92
5-Day Threshold	0.52	0.21	20	0.85	17.07	41.04
10 INR	0.45	0.15	10	1	10	30.71

Notes: INR per day complied represents Incentive Amount \times Proportion compliance incentivized (i.e., Column (5) = Column(3) \times Column(4)). INR per day complied above monitoring represents the total INR paid to a person (INR per day complied \times Compliance) divided by the non-inframarginal compliance (Compliance - Monitoring Avg.) (i.e., Column(6) = Column(5) \times Column(1)/(Column(1) - 0.3)).

One potential explanation for similar average walking in threshold and base case groups is that individuals do not notice the thresholds. However, the threshold contracts lead to markedly different walking patterns than the base case non-threshold group, showing that individuals clearly understand and respond to the thresholds. Figure 9 shows that the threshold contracts have a large bimodal effect on walking: more individuals in the threshold contracts achieve their step target 7 days in a week or 0 days in a week. The bimodal treatment effect from thresholds is not simply a feature of behavior across weeks, but also appears across individuals. Figure 10 plots the density of each individual’s probability of exceeding her step target, and mean daily steps, over the entire intervention period. The results across individuals mirror the results across weeks: the distribution of individual walking habits has thicker tails under the threshold treatments, with more people walking at the high and low ends. Appendix Table C.3 substantiates these conclusions using quantile regressions. In sum, although thresholds do not work well for everyone, they work very well for some people, inducing them to walk with more consistency across days than non-threshold contracts.

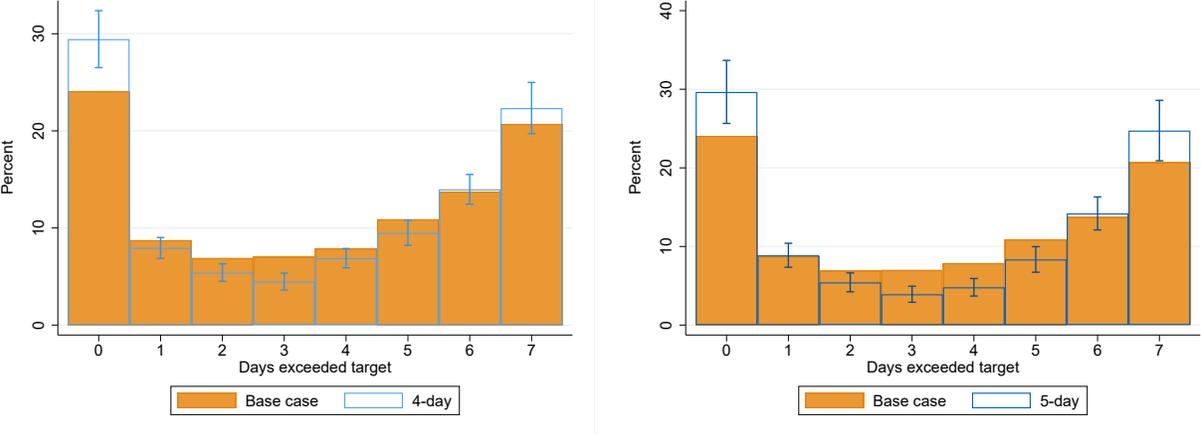


Figure 9: Days walked per week in threshold and base case (linear) contracts.

Notes: These figures show the distribution of the number of days walked each week in the 4-day and 5-day threshold contracts and the base case (linear) contract during the intervention period. All data is at the respondent \times week level. Orange bars show raw base case means; blue bars add on the coefficient on a dummy for being in the 4-day or 5-day threshold treatment from a regression where the dependent variable is a dummy for whether the days on which the participant exceeded the step target is equal to the value on the x-axis. Regressions control for the same things as Table 2. Confidence intervals are for the test of equality between the base case and 4-day or 5-day treatment from the same regression.

From a policy perspective, since threshold contracts create more extreme outcomes, we might be concerned if there are diminishing returns to behavior. In this setting, diminishing returns to exercise seem plausible, although the medical evidence is not definitive. If so, instituting a dynamic threshold creates a tradeoff, decreasing the cost per day of exercise

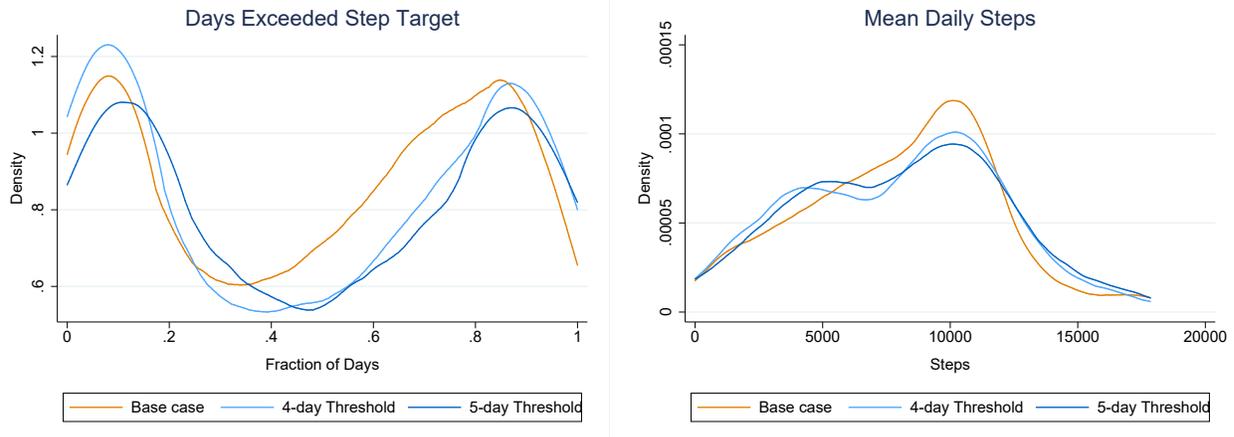


Figure 10: Fraction of days walked and average steps at the participant level, for base case vs. threshold

Notes: This figure shows the distribution of the fraction of days walked and average steps for each participant over the entire intervention period in each of the threshold contracts compared with the base case (linear) contract.

induced, but perhaps also diminishing the health benefit per day of exercise induced.³⁵ The bimodal effects of thresholds also highlight the importance of understanding for whom they work best. We next proceed to test our theoretical prediction about one type of individual for whom the threshold contracts will work better: those who are impatient over walking (Prediction 1 from Section 2).

5.2.2 Heterogeneity in Dynamic Threshold Effects by Time Preferences

To test for heterogeneity in the effects of dynamic thresholds by impatience, we regress an indicator for walking 10,000 steps (i.e., compliance with the incentivized activity) on our structural measure of impatience in the step domain, an indicator for being in a threshold treatment, and their interaction. We restrict the sample to only the base case incentive group and the 4- and 5-day threshold groups so that the only dimension that varies between groups is whether their contract has a dynamic threshold. Following our ex ante analysis plan, we pool the threshold treatments for power purposes (so *Threshold* is an indicator for being in either threshold group). The key coefficient of interest is on the interaction between the threshold treatment and the measure of impatience over steps.

Consistent with our theoretical predictions, impatience over steps appears to play a role in the success of thresholds. Column 1 of Table 4 shows that those with higher impatience over

³⁵On a similar note, one might think there would be greater value in inducing exercise among those who have low levels of baseline exercise rather than those who already exercised a lot at baseline. Appendix Table C.5 and Figure D.1 explore heterogeneity by baseline walking, showing suggestive but weak evidence that thresholds increase exercise more for those who walk more at baseline, providing one other potential downside of threshold contracts.

Table 4: Dynamic thresholds increase walking more for those who are more impatient.

Dependent variable:	Met step target ($\times 100$)				
	Discount Rate	Above-median disc. rate	Above-75th-perc. disc. rate	Standardized self-control index	Self-control index correlates
Impatience measure:	Has disc. rate	Has disc. rate	Has disc. rate	Late sample	Full
Impatience \times Threshold	0.382*** [0.138]	3.817 [3.387]	8.810** [3.830]	5.316* [2.994]	5.133** [2.003]
Threshold	-1.218 [1.698]	-2.852 [2.447]	-3.288* [1.972]	-1.205 [1.904]	-0.987 [1.399]
Impatience \times Incentives	-0.0138 [1.008]	-1.805 [4.454]	-4.541 [5.046]	-10.82** [4.344]	-2.903 [2.749]
Incentives	21.56*** [2.314]	22.27*** [3.310]	22.63*** [2.612]	21.71*** [2.653]	20.74*** [1.990]
Impatience	-0.226 [1.002]	2.910 [3.711]	2.883 [4.216]	5.549 [3.858]	-0.853 [2.329]
# Individuals	1,807	1,807	1,807	1,377	2,524
Base Case mean	48.5	48.5	48.5	50.8	50.7

Notes: This table shows heterogeneity by time preferences in the effect of threshold contracts relative to linear contracts. All discount rates are over walking (not payments). The sample includes all groups who received a fitbit. Control variables include all incentive subtreatment dummies (main effects and interaction with impatience), gender, age, weight, time fixed effects, and the baseline value of the dependent variable. The base case group is the omitted category. The first 3 columns all are based on the structural estimate of impatience in the step domain. Standard errors in brackets clustered at the respondent level. Significance levels: * 10%, ** 5%, *** 1%.

exercise have higher compliance under the threshold treatment relative to the linear. Because the structural discount rate measure often falls outside of the 0/1 range – as often happens with discount rate measurement, which is noisy at the individual level (Augenblick et al., 2015) – to aid in interpretation of the magnitude, columns 2 and 3 use indicators for whether impatience is above the 50th percentile and the 75th percentile, respectively. We lose some precision, but the columns show that the effect is large in magnitude: having above-median (above 75th percentile) increases compliance with the threshold contract relative to the base case contract by 4 percentage points (9 percentage points), which are large impacts relative to the average effect of incentives (20 percentage points). Note that only the second coefficient is significant. Figure 11 shows the threshold effects visually by quartile of impatience: the effectiveness of the threshold seems to be relatively flat until the final quartile.

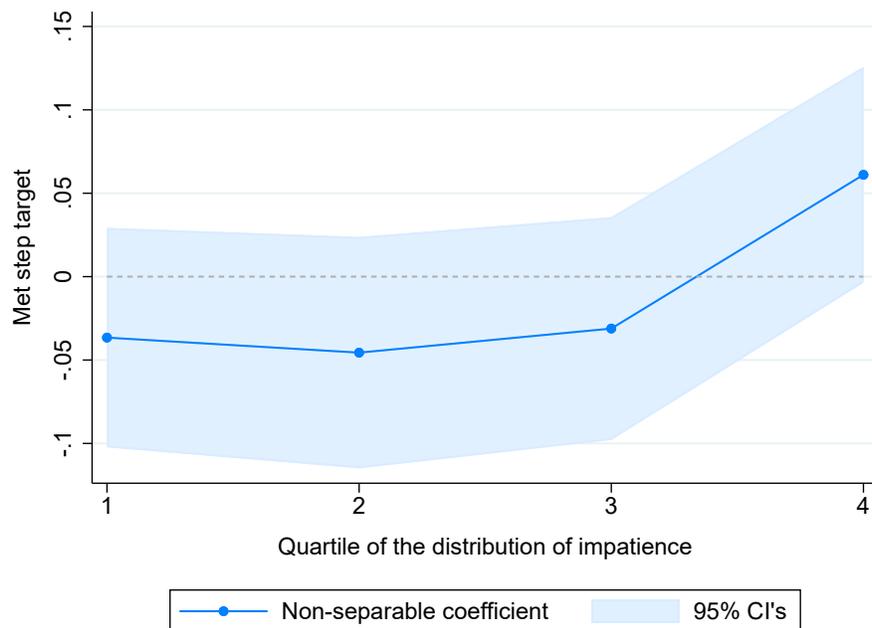


Figure 11: Threshold effect, by quartile of time preference

Notes: The sample is restricted to the weekly groups – i.e., the base case incentive group, which has a linear contract, and the 2 threshold groups, 4-day threshold and 5-day threshold, pooled here together as “Threshold.” The chart shows the coefficient on a dummy for being in a threshold treatment from a regression of compliance on a dummy for being in a threshold treatment group, estimated separately by quartiles of the distribution of impatience. The base case group is the omitted category. All regressions are estimated using the same controls as Table 2.

One concern with the structural impatience measure is that it can only be estimated for individuals making interior choices in the CTB allocation, and thus is missing for about 30% of the sample. We thus might be concerned that those who make interior choices are different than those who do not. Several weeks into the data collection for the project, we realized

the prevalence of non-interior choices, and added to our baseline survey a set of questions on impatience and self-control from the psychology literature that proxy for impatience and which we use to create a standardized index of impatience/self-control issues. Although these data are only available for participants who were enrolled after that point in the data collection (and thus precision is low), unlike for our structural measure, the sample with data available should be representative of the full sample. Reassuringly, we obtain a consistent finding using this standardized index, with the estimated coefficient in column 4 suggesting that the threshold works 5 pp better for those who are one standard deviation higher in the index. Finally, to get data from the full sample, we also create an index of procrastination-style questions that were included in the baseline survey from the beginning of data collection and that correlate with the index from column 4, and again find consistent results (see column 5).

Of course, impatience is not randomly assigned and could correlate with other variables that influence the effectiveness of the threshold treatment. Luckily, Appendix Table C.6 shows that the impatience measure is not correlated with most other variables (e.g., risk aversion), somewhat assuaging this concern. We have two other approaches to address this concern. First, Appendix Table C.7 controls for other baseline covariates and their interactions with the threshold treatments. Reassuringly, the tables shows that the threshold interactions are relatively robust. Note that this test is in many senses too stringent (“over-controlling”), as some of these other control variables could also be downstream outcomes affected by impatience.

Second, we calibrate a model using the empirical distribution of walking costs to show that, in this setting, the performance of the threshold treatments should indeed increase meaningfully with impatience over exercise. We first extend the simple framework from Section 2 to cover a 7-day model with 4-day and 5-day thresholds. To calibrate the average compliance in the threshold and base case (linear) contracts, we need to estimate the distribution of walking costs, $F(\cdot)$. We do this by fitting a normal distribution to several moments from the data. Average walking in the monitoring treatment uncovers $F(0)$; average payday walking in the 10 INR treatment uncovers $F(10)$; and average walking in the base case group on payday and in the daily group uncover $F(20)$. We also use two additional moments: the probability of walking for the 4-day (5-day) threshold group when one had already walked 3 days (4 days) and it is the last day of the contract period uncovers $F(80)$ ($F(100)$); these final moments improve fit, but do involve an assumption that costs are IID across people.

We then use this normal distribution to estimate how relative compliance in the base case and threshold contracts would vary with the discount rate over walking, d_c . The results are displayed visually in Figure 12, with the discount factor over walking on the x-axis, the gap

between performance in the threshold and base case on the y-axis (shown separately for the 4-day and 5-day thresholds), and the figure shown separately for different scenarios of the discount rate over payments d_m . The figure confirms that, given the sample’s distribution of walking costs, increase in performance of the threshold contract as impatience increases should be quantitatively important. Note that the calibration overestimates the average effect of the threshold, likely at least in part because our simple model does not incorporate uncertainty over future walking costs and risk aversion, which would decrease the average performance of the dynamic threshold. However, these other factors should primarily affect the average effect of the dynamic threshold relative to the base case, and should not we believe meaningfully affect the heterogeneity by impatience, which is our main interest here.

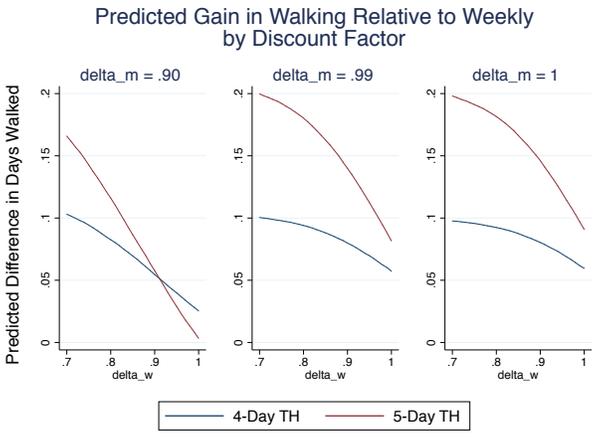


Figure 12: Calibration: Threshold compliance relative to the base case linear compliance, by the discount factor over walking

5.3 Payment Frequency

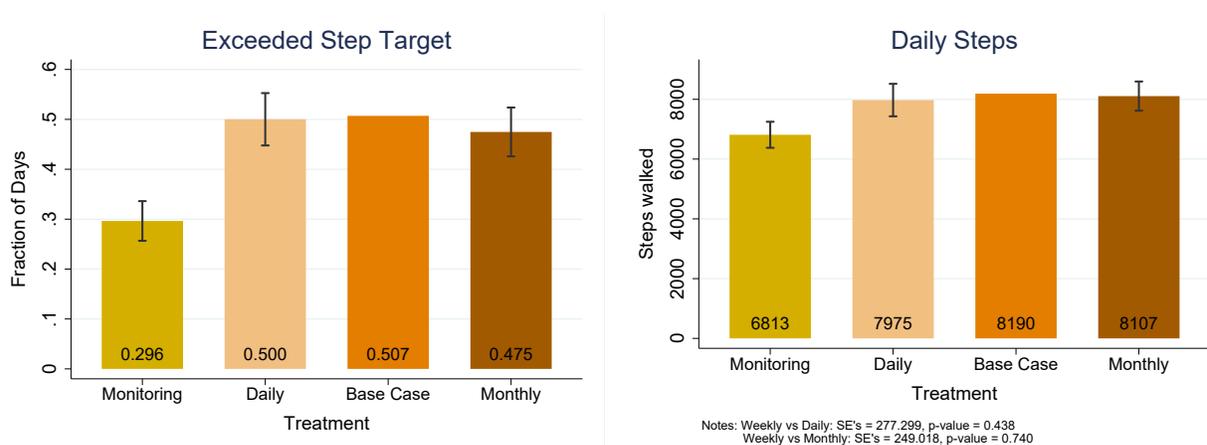
Motivated by Section 2, we conduct two primary tests:

1. Between-treatment: We compare average compliance between the daily, weekly (base case), and monthly payment groups in order to assess how payment frequency affects compliance and shed light on the level and shape of discount rates via Predictions 2 and 3.
2. Within-treatment: Within the base case and monthly groups, we examine whether compliance increases as the payday approaches to shed light on discount rates via Prediction 4. Similar variation has been used in previous studies to shed light on discount rates (Kaur et al., 2015; Oyer, 1998).

The approaches are complementary. The between-treatment approach directly answers

the policy question of whether payment frequency matters, while the within-treatment approach exhibits higher statistical power and can shed additional light on discount rates.

We begin with the between-treatment comparisons. In addition to Panel B of Table 2, Figure 13 shows the compliance in the frequency treatments visually. The impacts of the three frequency treatments on both the likelihood of exceeding the step target and on average steps walked are statistically indistinguishable. In addition, the differences between the point estimates are relatively small and there is no clear pattern by frequency. The weekly (base case) group has slightly higher steps and compliance with the step target than the daily group, while the monthly group has 2% lower compliance but slightly higher average steps than the daily group.



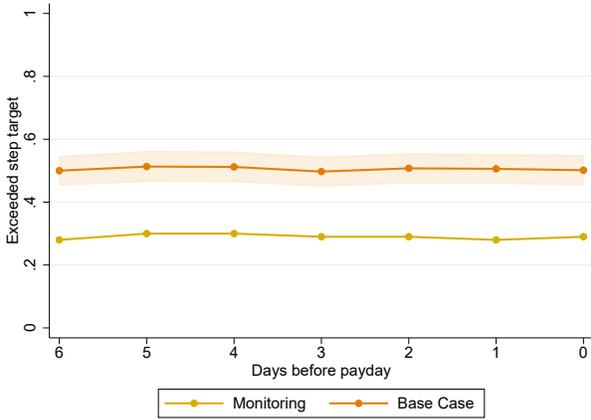
(a) Probability Exceeded Step Target

(b) Average Daily Steps

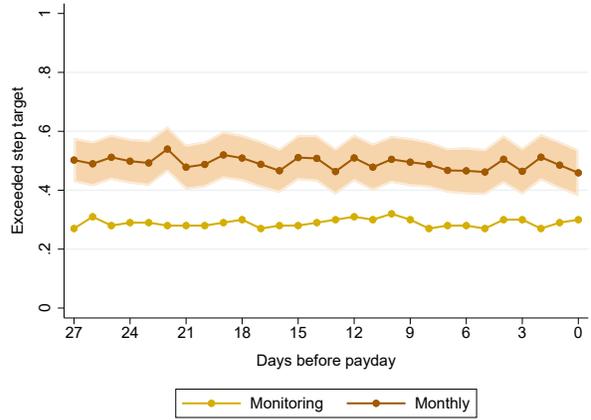
Figure 13: Payment frequency does not significantly impact walking.

Notes: Panel A shows the average probability of exceeding the daily 10,000-step target during the intervention period for the 3 different frequency treatments (note that the “base case” treatment pays with weekly frequency); Panel B shows average daily steps walked during the intervention period. Confidence interval bars show tests for equality between each group and the base case incentive group, and come from regressions that control for the same things as Table 2. We also control for different contract period start dates completed surveys.

We thus do not find any meaningful evidence that increasing payment frequency in the range from daily to monthly affects compliance, implying that the discount rate over financial payments may be relatively small over this range. Note that the implication for discount rates depends on whether there are other confounds besides discount rates that also affect individual responsiveness to payment frequency; however, most of the potential confounds one can think of (e.g., concavity of utility in payment, higher salience of higher frequency payments) would also bias higher frequency payments towards working better, thus making it unlikely that confounds (as opposed to lack of discounting of payments) are driving the null result. Note also that precision is a caveat to these conclusions: we cannot rule out



(a) Weekly Payment Cycle



(b) Monthly Payment Cycle

Figure 14: The probability of exceeding the step target is stable over the payment cycle

Notes: Figures show the probability of exceeding the daily 10,000-step target among individuals receiving the base case, i.e., weekly, incentive (Panel A) and a monthly incentive (Panel B) relative to the monitoring group, according to days remaining until payday. Effects control for payday day-of-week fixed effects, as well as the same things in Table 2. There is no evidence of a spike in walking on the day of incentive delivery for incentivized participants; in contrast, the slope of walking as the payday approaches is negative for participants in both the base case and monthly treatment groups.

that daily has an effect 5 percentage points higher than the base case, or that monthly has one 8 percentage points lower. We thus turn next to the within-treatment test, which has somewhat higher statistical power, to confirm the evidence from this analysis that discount rates over payment are low and that payment frequency does not meaningfully improve compliance.

Figure 14 shows how compliance within the base case weekly (Panel A) and monthly (Panel B) treatments changes as the payment day approaches. The prediction of impatience over payments would be that compliance increases as the payday approaches. Instead, walking behavior is remarkably steady across the payment cycle. Table 5 estimates the slopes of each line as the payment day approaches, conditional on day of week fixed effects.³⁶ The estimates are not significant and suggest that, if anything, the slopes decrease as the payment day approaches: for each day closer to the payday, base case (weekly) participants are 0.1 percentage points less likely to comply, and monthly participants are 0.06 percentage points less likely. Our confidence intervals are also tighter here: if we assume linearity of compliance in lag to payment, the bottom end of the confidence interval for the slope of the base case weekly treatment can rule out that weekly payment efficacy falls by 0.05

³⁶There is variation in day of week relative to payday since the intervention launch visits were done on all days of the week. Since the launch visit dates were endogenous to participants' schedules, survey day of week may also be endogenous; as a result, we also randomized the delay between the survey date and the intervention start date. We thus also control for fixed effects for the day-of-week relative to survey date.

pp with each day further from payment, which would correspond to ruling out that daily payments are on average more than 0.2pp more effective than weekly. Although these results are consistent with recent work by Augenblick et al. (2015), the absence of payday spikes conflicts with Kaur et al. (2015); an open question for further work is to further explore the differences (e.g., whether they reflect different countries, different payment amounts, or different settings in health vs. workplace).

Table 5: Walking does not vary significantly across the paycycle

Dependent variable:	Met step target ($\times 100$)				
Payment Frequency:	Weekly		Monthly		
	(1)	(2)	(3)	(4)	(5)
Days before payday	0.145 [0.101]		0.0578 [0.183]		
Payday		-0.868 [0.604]		-0.347 [1.100]	
Payweek					0.347 [1.100]
# Observations	68,334	68,334	12,908	12,908	12,908
Sample mean	50.71	50.71	50.10	50.10	50.10

Notes: Columns show the effect of paydays on the probability of meeting the step target, in the weekly and monthly frequency groups. The sample in columns (1) and (2) is restricted to the weekly treatment group and the sample in columns (3) and (4) is restricted to the monthly treatment group. All regressions control for the same things as Table 2. Standard errors in brackets clustered at the individual level. Significance levels: * 10%, ** 5%, *** 1%.

Thus, the evidence suggests that, on aggregate, the discounting model that best describes our participants is one of patience over mobile recharges within a 1-month time horizon, and that increasing frequency between daily and monthly does not have meaningful effects on average compliance. These are all average effects. Appendix B tests for heterogeneity by discount rates over recharges. We do not find significant heterogeneity in either within-treatment or between-treatment effects. Even for those with high measured impatience over recharges, increasing payment frequency does not appear to be effective in this context.

6 Results: Program Evaluation

In this section, we evaluate the effects of the incentives and monitoring treatments on health and behavioral outcomes, and investigate whether monitoring alone affects exercise.

6.1 Health and Lifestyle Effects

The impacts of an incentives program on health and healthy behaviors are of independent policy interest, especially among a population at high risk for complications from non-

communicable disease such as ours. Regular exercise such as walking can help prevent complications from diabetes, as well as hypertension. In addition, exercise may have coincident benefits for physical fitness and mental health. Finally, walking incentives programs may also impact other behaviors, either encouraging them (e.g., by increasing the salience of good health), or discouraging (e.g. if people substitute between healthy behaviors). We now assess the impacts of our programs on health and healthy behavior.

For looking at health outcomes – our primary outcomes of interest – our experiment was primarily powered to detect the difference between incentives groups (pooled) and the pure control group. Table 6 reports results from regressions of the following form:

$$y_i = \alpha + \beta_1 \times incentives_i + \beta_2 \times monitoring_i + \mathbf{X}'_i \gamma + \varepsilon_i \quad (7)$$

where y_i is a health or lifestyle outcome at endline for individual i ; $incentives_i$ is an indicator for being in the incentives group; $monitoring_i$ is an indicator for being in the monitoring group; and \mathbf{X}_i is a vector of controls, shown in the table notes. β_1 is the ITT effect of incentives relative to the pure control group, β_2 is the ITT effect of the monitoring relative to the pure control group, and, for benchmarking, β_3 is the ITT effect of being in the SMS reminders group.

This section reports ITT effects on outcomes in five categories: physical health (our primary outcomes), anaerobic fitness, mental health, diet, and addictive substance use. In order to maximize our power to detect overall effects on each category, we create a single index of all variables in each category by taking the simple average of each variable, standardized by the mean and standard deviation in the pure control group.³⁷ While we report regression estimates for each outcome individually, we focus on the category indices for inferring effectiveness.

Panel A of Table 6 shows that the incentives program improves health indicators in this population ³⁸ Column 1 presents the treatment effect on the “Health Risk Index”, which averages the five health risk factors displayed in the table. We find a moderately sized reduction in health risk of 0.04 SD’s, significant at the 10% level. Turning to the components of the index, we see average reductions in two measures of blood sugar at endline: Hba1c, a three-month weighted average of blood sugar levels; and random blood sugar (RBS), a measure of instantaneous blood sugar levels, although only the latter is significant, and

³⁷We follow Kling et al. (2007), by imputing missings (for individuals who have non-missing responses to at least one component of the index) for each component using the sample mean.

³⁸All physical health outcome variables are trimmed according to the WHO flexible exclusion method for dealing with biologically implausible values for health outcomes; specifically, for each measure, observations for which z-scores were greater than 4 units from the mean z-score are trimmed. Note that the WHO guidelines are designed for infants and children; the WHO does not provide guidelines for adults. Results are robust to other methods of trimming outliers (e.g., winsorizing at the 1st percentile).

only at the 10% level. Of the other 3 components, only waist circumference is individually significant (at the 10% level). The table also shows that monitoring alone did not seem to impact health.

To assess the size of our incentives treatment on health, we can also compare our effect sizes to the effects of other interventions in the literature. Although the Incentive treatment effects appear small, they are, in fact, relatively reasonable when compared to other interventions in terms of scalability, intensity, and cost. Online Appendix Table D. 2 shows the effects sizes and intervention details of other SMS and exercise interventions. Note that the majority of studies that find larger effects on HbA1c utilize more intensive interventions that are both costly and are unlikely to be scalable.

We next examine the ITT impacts of incentives on physical fitness. Our survey collected two measures of anaerobic fitness at baseline and endline: 4-meter timed walk, and standing five times from a sitting position (e.g., from a chair). Thus, a smaller value of either of these measures indicates greater anaerobic fitness. Panel B of Table 6 shows that participants in the incentives groups are not meaningfully faster or slower at either the times walk or sit-stands, nor on our index of the two measures. Although it is surprising that walking does not have any detectable impacts on our measure of fitness, while having large impacts on exercise, this may partly be explained by the fact that our intervention motivated a low-intensity form of exercise, while we were only able to implement time trials of high-intensity, short-duration exercise activities in our surveys.

We next turn to the ITT impacts of incentives on mental health. We measure mental health using seven questions adapted from the Rand 36-Item Short Form Survey (SF-36), a standard quality-of-life survey which has been validated for measuring emotional wellbeing in India (Rajeswari et al., 2005; Sinha et al., 2013). We selected questions related to emotional health. Each question asks for the frequency of a feeling or event in the previous four weeks.³⁹ Answers are then recoded so that larger values indicate better mental health.

Panel C of Table 6 shows that the incentives program significantly improve our index of mental health. In addition, the monitoring program has an even larger positive effect on mental health, although we cannot reject equality. Although many studies have found a positive association between exercise and mental health (Biddle, 2016), experimental evidence that exercise causes improvements in mental health is scarce and mixed. Our result is novel experimental evidence that exercise can improve mental health.

Finally, Table C.1 examines effects two dimensions of a healthy lifestyle: diet, and addictive good consumption. We do not find significant incentive effects on either.

³⁹For example, the “Felt happy” question asks: “In the previous four weeks, how often have you felt happy? All of the time, most of the time, a good bit of the time, some of the time, a little of the time, or none of the time?”

Table 6: Impacts of incentives and monitoring on health.

A. Health risk factors						
	Health Risk Index	HbA1c	RBS	Mean Arterial BP	BMI	Waist Circumference
Incentives	-0.046* [0.025]	-0.075 [0.071]	-5.94* [3.42]	0.075 [0.42]	-0.049 [0.042]	-0.19 [0.27]
Monitoring	0.014 [0.044]	-0.14 [0.13]	2.02 [6.06]	1.11 [0.75]	0.068 [0.074]	0.0079 [0.48]
Control mean	0.00	8.44	193.83	103.02	26.45	94.44
P-value: M = I	0.13	0.55	0.14	0.12	0.08	0.65
# Individuals	3,068	3,066	3,067	3,056	3,058	3,059

B. Fitness			
	Fitness Time Trial Index	Seconds to Walk 4m	Seconds for 5 Sit-Stands
Incentives	0.031 [0.029]	0.042 [0.043]	-0.11 [0.12]
Monitoring	0.069 [0.051]	0.082 [0.076]	-0.089 [0.20]
Control mean	0.00	3.88	13.18
P-value: M = I	0.41	0.56	0.92
# Individuals	2,890	2,825	2,793

C. Mental Health								
	Mental Health Index	Felt Happy	Less Nervous	Peaceful	Energy	Less Blue	Less Worn	Less harm to Social Life
Incentives	0.072** [0.033]	0.088** [0.045]	0.028 [0.044]	0.056 [0.047]	0.061 [0.048]	0.016 [0.043]	0.089** [0.039]	0.050* [0.030]
Monitoring	0.11* [0.058]	0.074 [0.079]	0.13 [0.077]	0.096 [0.083]	0.033 [0.084]	0.13* [0.077]	0.17** [0.069]	0.048 [0.053]
Control mean	0.00	3.06	3.48	3.35	3.30	3.86	4.40	4.71
P-value: M = I	0.51	0.84	0.15	0.60	0.70	0.11	0.17	0.95
# Individuals	3,068	3,068	3,068	3,068	3,068	3,068	3,068	3,068

Notes: Standard errors in brackets. Controls for all outcomes are the same as Table 2, along with second order polynomials of all relevant health, fitness and mental health variables underlying the index at baseline. We follow Kling et al. (2007) in constructing the index by imputing missings for each component using the sample mean. The Health Risk Index is an index created by the average of endline HbA1c, RBS, MAP, BMI, and waist circumference standardized by their average and standard deviation in the control group. HbA1c is the average plasma glucose concentration (%), RBS is the blood glucose level (mg/dL), MAP is the mean arterial blood pressure (mm Hg), and BMI is the body mass index. A large value of Fitness Time Trial Index indicates low fitness. The Mental Health Index is the average values of seven questions adapted from the Rand 36-Item Short Form Survey (SF-36). The omitted category in all columns is the pure control group.

6.2 Monitoring and Exercise

The previous results suggest that the monitoring group had limited impact, although the results are somewhat imprecise. Did the monitoring treatment not affect exercise, or were the exercise impacts too small to translate into measurable health impacts? In this section, we evaluate the impact of monitoring on walking.

Because we do not have pedometer walking data from the control group and the self-reported exercise data appears to be biased (see Online Appendix C for more detail), we evaluate the effect of monitoring using a before-after design, comparing pedometer-measured walking in the monitoring group during the phase-in period (during which we had not given participants a walking goal and just told them to walk the same as they normally do) to their behavior during the intervention period. This strategy will be biased either in the presence of within-person time trends in walking, or if the phase-in period directly effects walking behavior. We control for year-month fixed effects to help address time trends, but the latter concern is more difficult, as the phase-in period likely did increase walking above normal, either because of Hawthorne effects or because participants received a pedometer and a step-reporting system, which are two of the elements of the monitoring treatment itself (the other three remaining that we can still evaluate are (a) a daily 10,000 step goal, (b) positive feedback for meeting the step goal through SMS messages and the step-reporting system, and (c) periodic walking summaries). Thus, we consider a pre-post comparison of walking in the monitoring group to be a lower bound of the monitoring program treatment effect.

One can visualize the variation used for our pre-post estimate in Figure 7. Walking increases immediately during the intervention period for the monitoring group, although the effects decay over time. We estimate the effect in Online Appendix Section E, controlling for date effects. The monitoring group achieves the 10,000 step target on approximately 7% more days in the intervention period than in the phase-in period, an effect significant at the 1% level and equal to roughly 35% of the estimated impact of incentives. However, the estimated effect on steps is only 62 steps (statistically indistinguishable from zero), which is only 10% of the additional impact of incentives. The monitoring treatment thus appears to do more to make walking consistent across days than it does to increase total steps.

7 Conclusion

This paper investigates incentive design for impatient agents. From a model of agents with separate time preferences over consumption and financial rewards, we identify a feature of incentive contracts that will interact with impatience in each domain. In particular, our model predicts that the frequency of incentive delivery will interact with impatience over

financial rewards, and whether the payment function is dynamically separable over time will interact with impatience over consumption.

In order to test our predictions, we implement an RCT to incentivize walking among approximately 3000 individuals with diabetes and prediabetes in India. Overall, the incentives program leads to a large increase in walking among the study population and leads to improvements in diabetes- and mental- health risk factors. This is encouraging evidence that exercise-incentives programs can be successful in a developing country setting.

We find evidence that individuals are impatient over consumption - that is, they prefer to put off the effort of walking. Consistent with our model, the dynamically non-separable contract works better for those who are more impatient over consumption. However, we find limited evidence of impatience over payments in our sample. Neither more frequent nor more immediate payment leads to increased walking behavior in our sample, and we cannot reject a null relationship between a survey-based measure of impatience over rewards and the effectiveness of more frequent payment. The finding that impatience is more prevalent in the consumption domain than the financial domain is consistent with previous experimental work (Augenblick et al., 2015). However, our finding that dynamically non-separable contracts can be used to motivate time-inconsistent individuals is a new and policy-relevant insight. The fact that these contracts have heterogeneous treatment effects opens up a key question, which we hope to address in future work: can we tailor incentive contracts to more cost-effectively encourage exercise?

References

- Andreoni, J., Callen, M., Khan, Y., Jaffar, K., and Sprenger, C. (2016). Using Preference Estimates to Customize Incentives: An Application to Polio Vaccination Drives in Pakistan. *NBER Working Paper*, 22019:1–66.
- Andreoni, J. and Sprenger, C. (2012a). Estimating Time Preferences from Convex Budgets. *American Economic Review*, 102(7):3333–3356.
- Andreoni, J. and Sprenger, C. (2012b). Risk Preferences Are Not Time Preferences. *American Economic Review*, 102(7):3357–3376.
- Ariely, D. and Wertenbroch, K. (2002). Procrastination, Deadlines, and Performance: Self-Control by Precommitment. *Psychological Science*, 13(3):219–224.
- Ashraf, N., Karlan, D. S., and Yin, W. (2006). Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines. *The Quarterly Journal of Economics*, 121(2):635–672.
- Augenblick, N., Niederle, M., and Sprenger, C. (2015). Working Over Time: Dynamic Inconsistency in Real Effort Tasks. *The Quarterly Journal of Economics*, 130(3):1067–1115.
- Augenblick, N. and Rabin, M. (2017). An Experiment on Time Preference and Misprediction in Unpleasant Tasks. *Working Paper*, pages 1–81.
- Bachireddy, C., Joung, A., John, L. K., Gino, F., Tuckfield, B., Foschini, L., and Milkman, K. L. (2018). Comparing Financial Incentive Structures for Promoting Physical Activity: A Randomized Controlled Trial. *Working Paper*.
- Bai, L., Handel, B. R., Miguel, E., and Rao, G. (2017). Self-Control and Demand for Preventive Health: Evidence from Hypertension in India. *Working Paper*.
- Bassett, D. R., Wyatt, H. R., Thompson, H., Peters, J. C., and Hill, J. O. (2010). Pedometer-measured physical activity and health behaviors in U.S. adults. *Medicine and Science in Sports and Exercise*, 42(10):1819–1825.
- Bertrand, M., Karlan, D. S., Mullainathan, S., Shafir, E., and Zinman, J. (2005). What’s Psychology Worth? A Field Experiment in the Consumer Credit Market. *NBER Working Paper*, 11892.
- Biddle, S. (2016). Physical activity and mental health: evidence is growing. *World Psychiatry*, 15(2):176–177.

- Bommer, C., Heesemann, E., Sagalova, V., Manne-Goehler, J., Atun, R., Bärnighausen, T., and Vollmer, S. (2017). The global economic burden of diabetes in adults aged 20–79 years: a cost-of-illness study. *The Lancet Diabetes and Endocrinology*, 5(6):423–430.
- Carrera, M., Royer, H., Stehr, M. F., and Sydnor, J. R. (2017). The Structure of Health Incentives: Evidence from a Field Experiment. *NBER Working Paper Series*, page 49.
- Carvalho, L. S., Meier, S., and Wang, S. W. (2016). Poverty and Economic Decision - Making: Evidence from Changes in Financial Resources at Payday. *The American Economic Review*, 106(2):260–284.
- Charness, G. and Gneezy, U. (2009). Incentives to Exercise. *Econometrica*, 77(3):909–931.
- Clark, G. (1994). Factory Discipline. *The Journal of Economic History*, 54(1):128–163.
- Cutler, D. M. and Everett, W. (2010). Thinking Outside the Pillbox - Medication Adherence and a Priority for Health Care Reform. *The New England Journal of Medicine*, 362(17):1553–1555.
- Davis, L. W., Fuchs, A., and Gertler, P. J. (2014). Cash for coolers: Evaluating a large-scale appliance replacement program in Mexico. *American Economic Journal: Economic Policy*, 6(4):207–238.
- Dellavigna, S. and Malmendier, U. (2004). Contract design and self-control: Theory and evidence. *The Quarterly Journal of Economics*, 119(2):135–146.
- Duflo, E., Greenstone, M., Pande, R., and Ryan, N. (2010). Towards an Emissions Trading Scheme for Air Pollutants in India.
- Duflo, E., Kremer, M., and Robinson, J. (2011). Nudging Farmers to Use Fertilizer : Theory and Experimental Evidence from Kenya †. *American Economic Review*, 101(October):2350–2390.
- Finkelstein, E. A., Brown, D. S., Brown, D. R., and Buchner, D. M. (2008). A randomized study of financial incentives to increase physical activity among sedentary older adults. *Preventive Medicine*, 47:182–187.
- Fryer, R. G. (2011). Financial incentives and student achievement: Evidence from randomized trials. *Quarterly Journal of Economics*, 126(4):1755–1798.
- Giné, X., Goldberg, J., Silverman, D., and Yang, D. (2017). Revising Commitments: Field Evidence on the Adjustment of Prior Choices. *Economic Journal*, Forthcomin.

- Giné, X., Karlan, D., and Zinman, J. (2010). Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation. *American Economic Journal: Applied Economics*, 2(4):213–235.
- Hansen, B. B. and Bowers, J. (2008). Covariate Balance in Simple, Stratified and Clustered Comparative Studies. *Statistical Science*, 23(2):219–236.
- Hill, J. O. (2005). Walking and Type 2 Diabetes. *Diabetes Care*, 28(6):1524–1525.
- Jayachandran, S., De Laat, J., Lambin, E. F., Stanton, C. Y., Audy, R., and Thomas, N. E. (2017). Cash for carbon: A randomized trial of payments for ecosystem services to reduce deforestation. *Science*, 357(6348):267–273.
- Kaur, S., Kremer, M., and Mullainathan, S. (2015). Self-Control at Work. *Journal of Political Economy*, 123(6):1227–1277.
- Kling, J. R., Liebman, J. B., and Katz, L. F. (2007). Experimental Analysis of Neighborhood Effects. *Econometrica*, 75(1):83–119.
- Kullgren, J. T., Troxel, A. B., Loewenstein, G., Asch, D. A., Norton, L. A., Wesby, L., Tao, Y., Zhu, J., and Volpp, K. G. (2013). Individual- Versus Group-Based Financial Incentives for Weight Loss. *Annals of Internal Medicine*, 158:505–514.
- Labhardt, N. D., Balo, J. R., Ndam, M., Manga, E., and Stoll, B. (2011). Improved retention rates with low-cost interventions in hypertension and diabetes management in a rural African environment of nurse-led care: A cluster-randomised trial. *Tropical Medicine and International Health*, 16(10):1276–1284.
- Laibson, D. (2015). Why Don't Present-Biased Agents Make Commitments? *American Economic Review*, 105(5):267–272.
- Lazear, E. P. (1979). Why Is There Mandatory Retirement? *Journal of Political Economy*, 87(6):1261–1284.
- Lee, D. S. (2009). Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *Review of Economic Studies*, 76(3):1071–1102.
- Manders, R. J. F., Van Dijk, J. W. M., and Van Loon, L. J. C. (2010). Low-intensity exercise reduces the prevalence of hyperglycemia in type 2 diabetes. *Medicine and Science in Sports and Exercise*, 42(2):219–225.

- Marshall, S. J., Levy, S. S., Tudor-Locke, C. E., Kolkhorst, F. W., Wooten, K. M., Ji, M., Macera, C. A., Ainsworth, B. E., Ming, J., Macera, C. A., and Ainsworth, B. E. (2009). Translating Physical Activity Recommendations into a Pedometer-Based Step Goal: 3000 Steps in 30 Minutes. *American Journal of Preventive Medicine*, 36(5):410–415.
- Martins, N., Morris, P., and Kelly, P. M. (2009). Food incentives to improve completion of tuberculosis treatment: randomised controlled trial in Dili, Timor-Leste. *BMJ (Clinical research ed.)*, 339:b4248.
- Morris, S. S., Flores, R., Olinto, P., and Medina, J. M. (2004). Monetary incentives in primary health care and effects on use and coverage of preventive health care interventions in rural Honduras: cluster randomised trial. *Lancet (London, England)*, 364(9450):2030–7.
- O’Donoghue, T. and Rabin, M. (1999a). Incentives for Procrastinators. *The Quarterly Journal of Economics*, 114(3):769–816.
- O’Donoghue, T. and Rabin, M. (1999b). Incentives for Procrastinators. *The Quarterly Journal of Economics*, 10(2):130–153.
- Opp, M. M. and Zhu, J. Y. (2015). Impatience versus Incentives. *Econometrica*, 83(4):1601–1617.
- Oyer, P. (1998). Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality. *The Quarterly Journal of Economics*, 113(1):149–185.
- Patel, M. S., Asch, D. A., Rosin, R., Small, D. S., Bellamy, S. L., Eberbach, K., Walters, K. J., Haff, N., Lee, S. M., Wesby, L., Hoffer, K., Shuttleworth, D., Taylor, D. H., Hilbert, V., Zhu, J., Yang, L., Wang, X., and Volpp, K. G. (2016). Individual Versus Team-Based Financial Incentives to Increase Physical Activity: A Randomized, Controlled Trial. *Journal of General Internal Medicine*, 31(7):746–754.
- Praet, S. F. E. and van Loon, L. J. C. (2009). Exercise therapy in Type 2 diabetes. *Acta Diabetol*, 46:263–278.
- Prendergast, C. (1999). The Provision of Incentives in Firms. *Journal of Economic Literature*, 37(1):7–63.
- Qiu, S., Cai, X., Schumann, U., Velders, M., Sun, Z., and Steinacker, J. M. (2014). Impact of walking on glycemic control and other cardiovascular risk factors in type 2 diabetes: A meta-analysis. *PLoS ONE*, 9(10).

- Rajeswari, R., Muniyandi, M., Balasubramanian, R., and Narayanan, P. (2005). Perceptions of tuberculosis patients about their physical, mental and social well-being: a field report from south India. *Social Science & Medicine*, 60(8):1845–1853.
- Ramachandran, A. (2013). Effectiveness of mobile phone messaging in prevention of type 2 diabetes by lifestyle modification in men in India. *THE LANCET Diabetes & Endocrinology*, 1(3):191–198.
- Royer, H., Stehr, M., and Sydnor, J. (2015). Incentives, Commitments, and Habit Formation in Exercise: Evidence from a Field Experiment with Workers at a Fortune-500 Company. *American Economic Journal: Applied Economics*, 7(3):51–84.
- Schilbach, F. (2017). Alcohol and Self-Control: A Field Experiment in India. *Mimeo*.
- Shenoy, S., Guglani, R., and Sandhu, J. S. (2010). Effectiveness of an aerobic walking program using heart rate monitor and pedometer on the parameters of diabetes control in Asian Indians with type 2 diabetes. *Primary Care Diabetes*, 4(1):41–45.
- Sinha, R., van den Heuvel, W. J. A., and Arokiasamy, P. (2013). Validity and Reliability of MOS Short Form Health Survey (SF-36) for Use in India. *Indian journal of community medicine : official publication of Indian Association of Preventive & Social Medicine*, 38(1):22–6.
- Tanaka, T., Camerer, C. F., and Nguyen, Q. (2010). Risk and time preferences: Experimental and household survey data from Vietnam. *American Economic Review*, 100(1):557–571.
- Tharkar, S., Devarajan, A., Kumpatla, S., and Viswanathan, V. (2010). The socioeconomics of diabetes from a developing country: A population based cost of illness study. *Diabetes Research and Clinical Practice*, 89(3):334–340.
- Thomas, D., Elliott, E. J., and Naughton, G. A. (2009). Exercise for type 2 diabetes mellitus (Review). *Cochrane Database of Systematic Reviews*, 2006(3).
- Thornton, R. L. (2008). The Demand for, and Impact of, Learning HIV Status. *The American economic review*, 98(5):1829–1863.
- Volpp, K. G., John, L. K., Troxel, A. B., Norton, L. A., Fassbender, J., and Loewenstein, G. (2008). Financial Incentive Based Approaches for Weight Loss: A Randomized Trial. *JAMA : the journal of the American Medical Association*, 300(22):2631–2637.
- Zanuso, S., Jimenez, A., Pugliese, G., Corigliano, G., and Balducci, S. (2009). Exercise for the management of type 2 diabetes: a review of the evidence. *Acta Diabetol*, 47:15–22.

Appendix for:

Incentivizing Behavioral Change: The Role of Time Preferences

Table of Contents

Appendix

A. Baseline Time Preference Measurement

A.1 This table summarizes the parameters of the six CTB allocations made over recharges, and the eight CTB allocations made over steps.	57
A.2 Time preference summary statistics.	58

B. Heterogeneity in Frequency Effects by Impatience over Recharges

B.1 High-frequency treatments are not more effective for those who are impatient over recharges	60
B.2 Payday effects are not bigger for those with higher measured impatience over recharges	61

C. Supplementary Tables and Figures

C.1 Impacts of incentives and monitoring on diet and addictive consumption. . .	62
C.2 Impacts of incentives and monitoring on exercise outcomes, without baseline controls.	63
C.3 Impacts of incentives contracts, compared to the base case non-threshold contract, on the probability of being in 6 quantiles of average exercise outcomes.	64
C.4 Differential Incentive Effects according to Impatience over Recharges and Steps	65
C.5 Heterogeneity in threshold impacts by baseline walking	66
C.6 The structural measure of impatience over exercise is not significantly correlated with other variables	67
C.7 Time preference heterogeneity robust to including other controls	68

Supplemental Online Appendix

Note: The Online Appendix is a separate document that can be found at:
<http://faculty.chicagobooth.edu/rebecca.dizon-ross/research/papers/IncentiveDesignOnlineApp.pdf>

D. Additional Tables and Figures

D. 1 Endline attendance in full sample and by treatment group.	1
D. 2 Other studies of interventions to improve health outcomes among diabetics	3
D. 3 Frequency and Payday Effects, Benchmarked against Monitoring	4

E. SMS Treatment Impacts

E.1 Impacts of incentives, monitoring, and SMS treatments on health risk factors. 6
E.2 Impacts of incentives, monitoring, SMS treatments, and their interactions on health risk factors 7
E.3 Impacts of incentives, monitoring, and SMS treatments on fitness. 8
E.4 Impacts of incentives, monitoring, and SMS treatments on mental health. 9
E.5 Impacts of incentives, monitoring, and SMS treatments on Healthy Diet, relative to pure control. 10
E.6 Impacts of incentives, monitoring, and SMS treatments on Addictive Substance Consumption, relative to pure control. 11

F. Monitoring Treatment Impacts

F.1 Impacts of monitoring (pre-post) and incentives (difference-in-differences) on exercise outcomes. 13
F.2 Impacts of incentives and monitoring on self-reported exercise. 14

A Baseline Time Preference Measurement

We follow earlier work and use a Convex Time Budget (CTB) methodology to estimate time preferences. In each CTB choice of the time-preference survey, the participant is asked to allocate a fixed budget of either steps or mobile recharges between a “sooner” and a “later” date using a slider bar. In particular, each choice allows the respondent to choose an allocation of consumption on the sooner and later dates, c_t, c_{t+k} that satisfies the budget constraint

$$c_t + \frac{1}{r}c_{t+k} = m \quad (8)$$

where the sooner date t , the later date $t + k$, the interest rate r , and the budget m change between each choice. A sample slider screen allowing for such choices is shown in Figure A.1.

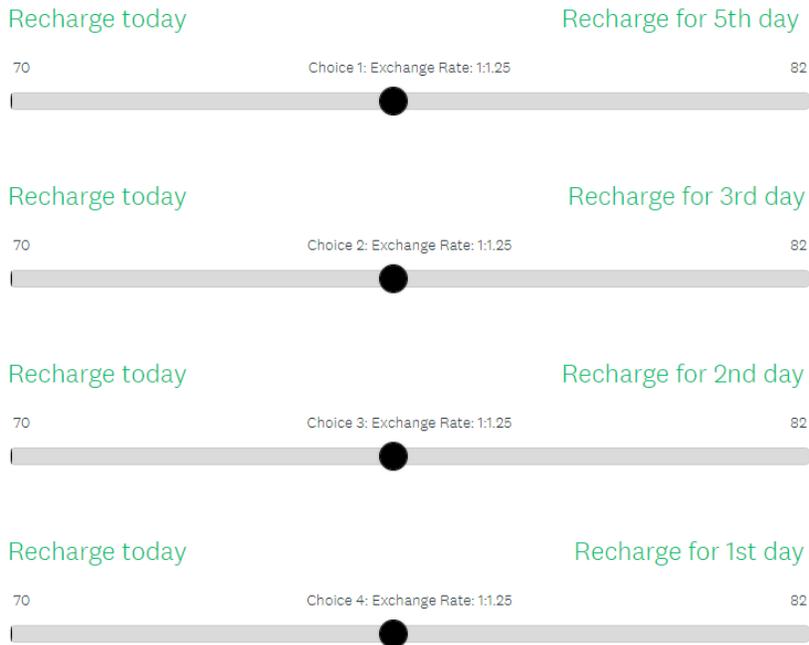


Figure A.1: Sample decision screen for mobile recharges. In this example, the interest rate, r , is 1.25; the total budget, m , is 140; the “sooner” date is Today; and the “later” date decreases from 5 days from today in the first choice to 1 day from today in the final choice. The sliders are shown positioned at the choice ($c_t = 70, c_{t+k} = 82$).

We asked participants to make six allocations in the recharge domain, and eight allocations in the step domain, as summarized in Table A.1. We assume a time-separable and good-separable CRRA utility function with quasi-hyperbolic discounting. In the domain of recharges, individuals will then seek to maximize utility,

$$U(c_t, c_{t+k}) = \frac{1}{\alpha}(c_t - \omega)^\alpha + \beta\delta^k \frac{1}{\alpha}(c_{t+k} - \omega)^\alpha \quad (9)$$

and in the step domain, individuals will seek to minimize costs of effort

$$C(c_t, c_{t+k}) = \frac{1}{\alpha} (c_t + \omega)^\alpha + \beta \delta^k \frac{1}{\alpha} (c_{t+k} + \omega)^\alpha \quad (10)$$

The variation between consumption choices given different parameters of the budget constraint identify the daily discount factor δ , the present-bias parameter β , and the concavity or convexity of preferences α in each domain. We recover structural estimates of time preference and concavity parameters from the allocations, (c_t, c_{t+k}) , using a two-limit Tobit specification of the intertemporal Euler condition following Augenblick et al. (2015).

Appendix Table A.1: This table summarizes the parameters of the six CTB allocations made over recharges, and the eight CTB allocations made over steps.

Summary of Convex Time Budget allocations					
Question no.	t	k	r	Recharge Domain	Step Domain
1	7	7	1	X	X
2	0	7	1	X	X
3	0	5	1.25	X	X
4	0	3	1.25	X	X
5	0	2	1.25	X	X
6	0	1	1.25	X	X
7	7	7	1.25		X
8	0	7	1.25		X

We follow Augenblick et al. (2015) and use our CTB allocations in each domain to estimate a structural time-preference measure of the daily discount rate. We use two-limit Tobit specifications of the standard intertemporal Euler equation for an agent with an exponential daily discount factor δ , and concavity over recharges, or convexity over steps, α .⁴⁰ Summary statistics can be found in Appendix A.2. About 40% of our observations have a negative discount rate over recharges while 37% have a negative discount rate over steps; it is common to estimate negative discount rates (i.e., deltas over 1) during experiments measuring monetary discount rates.

$$\log\left(\frac{c_t}{c_{t+1}}\right) = \frac{\log(\delta)}{\alpha - 1} k + \frac{1}{\alpha - 1} \log(r) \quad (11)$$

We take the following precautions to try to avoid confounds. To avoid confounds related to fixed costs or benefits, such as the effort of wearing a pedometer or the psychological benefit of receiving a free recharge, we include minimum allocations on both sooner and later days in each domain. The minimum allocations were chosen to be high enough that any fixed costs would be included (e.g. one could not easily achieve the minimums by simply shaking the pedometer), but low enough to avoid corner solutions. In the step domain, this required a novel modification of the CTB methodology: individual-specific minimum allocations. Our step allocations also featured individual-specific total step budgets m , which were chosen to be large enough that achieving them would require some effort beyond simply wearing the

⁴⁰Details on the estimation strategy can be found in the Online Appendix of Augenblick et al. (2015).

Appendix Table A.2: Time preference summary statistics.

	Discount rate over recharges	Discount rate over steps
Mean	0.011	0.363
Median	0.013	0.021
Standard Deviation	0.086	3.762
10 th Percentile	-0.103	-0.173
90 th Percentile	0.102	0.719
Observations	1723	2296

pedometer, but small enough that participants would certainly achieve them in exchange for the completion bonus. Specifically, minimum steps on each day are calculated as $\frac{X}{10}$, and the total step budget m is $X + 2\frac{X}{10}$, respectively, where $X \in \{3000, 4000, 5000\}$ is the element closest to the participant's average daily walking during the phase-in period. That is, minimum steps are one of 300, 400, or 500 on each day, and the total step budget is one of 3,600, 4,800, or 6,000. To avoid confounding impatience with the time of day that the baseline time-preference survey was administered (which could influence the desirability of walking and/or recharges delivered in the next 24 hours), as well as to capture heterogeneity in time preferences including any present-bias for very short beta-windows, we required that all walking on any date be conducted within a 2 hour period, which was chosen to start at the time immediately after the time-preference survey would end (e.g., if the survey ended at 4pm, the time period for any day's walking would be 5-7pm). The short window could potentially bias our overall measures of impatience downwards, as uncertainty about future schedules in a short time window could lead participants to want to get their walking done early when they had more certainty over their schedule. However, our primary purpose was to capture heterogeneity in time-preferences, and we considered the potential loss in validity of aggregate time preference estimates to be worth the ability to capture heterogeneity in time preferences in the time frames near to the present.

B Heterogeneity in frequency effects by impatience over recharges

This appendix explores the possibility that immediate incentive delivery is a driver of incentive effectiveness among the subset of more impatient participants. If so, we expect a positive interaction between more immediate incentive delivery and our measure of baseline impatience over mobile recharges. We test this interaction using both between-treatment and within-treatment variation in immediacy of payment.

Our first test is whether daily incentives are relatively more effective, and monthly relatively less effective than the base case of weekly payments, for those who display more impatience for recharges in their baseline CTB allocations. For simplicity, we restrict the sample to those who were in the daily, weekly, and monthly groups, and run the following

regression:

$$y_{it} = \alpha + \beta_0 \text{Impatience}_i + \beta_1 \text{daily}_i + \beta_2 \text{monthly}_i + \beta_3 \text{Impatience}_i \times \text{daily}_i + \beta_4 \text{Impatience}_i \times \text{monthly}_i + \mathbf{X}'_i \gamma + \varepsilon_{it}, \quad (12)$$

where y_{it} is a daily walking outcome; Impatience_i is either the daily discount rate estimated using CTB allocations over recharges at baseline or an indicator for having above-median daily discount rate; and daily_i and monthly_i are indicators for being assigned to the daily and monthly treatments, respectively. β_1 and β_2 represent the effects of daily and monthly relative to the base case weekly payment (respectively). The coefficients of interest are β_3 and β_4 , showing whether the effects of daily or monthly relative to Weekly are differentially large for those who are more impatient. If impatience over recharges is a mechanism through which more immediate incentive delivery increases effectiveness, then we expect the daily treatment to be more effective ($\beta_3 > 0$) and the monthly treatment to be less effective ($\beta_4 < 0$) for more impatient individuals.⁴¹ Our results are reported in Table B.1. We see no evidence that suggests that sooner payments work better for those with higher measured impatience, with the one marginally significant effect going the wrong direction.

Our second test is whether individuals who display more impatience for recharges in their baseline CTB allocations are more likely to increase step-target compliance on their payday. We perform this test among individuals in the base case incentive and monthly incentives groups. Following Kaur et al. (2015), we define individual-specific walking “payday effects” as the difference in the probability of exceeding 10,000 steps on paydays compared to all other days. The walking payday effect is a revealed-preference measure of impatience over rewards. We estimate the interaction between individual payday effects and our structural measure of baseline impatience over recharges using regressions of the following form:

$$y_{it} = \alpha + \beta_0 (\text{Impatience Measure})_i + \beta_1 (\text{Payday})_{it} + \beta_2 (\text{Payday})_{it} \times (\text{Impatience Measure})_i + \mathbf{X}'_i \gamma + \varepsilon_{it}, \quad (13)$$

where y_i , $(\text{Impatience Measure})_i$, and \mathbf{X}_i are defined as in equation 12; and $(\text{Payday})_{it}$ is an indicator for whether day t is a payday for individual i . To test whether more impatient individuals respond more to more immediate payment, we test whether $\beta_2 > 0$.

Our results are shown in Table B.2. Both measures of impatience predict larger payday effects, but the coefficients are not statistically different from zero. Thus, we find no strong evidence that even those individuals who are most impatient over rewards react to more immediate reward delivery over the payment cycle.

⁴¹Note that we do not have predictions for the interactions of the other incentive contracts with impatience over recharges; nonetheless, for completeness, Appendix Table C.4 shows regressions where the *Impatience* variable is interacted with all separate incentive treatments.

Appendix Table B.1: High-frequency treatments are not more effective for those who are impatient over recharges

Dependent variable:	Met step target	
	Discount Rate	Above-median disc. rate
Impatience measure:	(1)	(2)
Daily \times Impatience	-0.332 [0.40]	-0.125* [0.07]
Monthly \times Impatience	-0.230 [0.37]	0.00284 [0.07]
Daily	0.00300 [0.04]	0.0590 [0.05]
Monthly	-0.0197 [0.03]	-0.0247 [0.05]
Impatience	-0.0632 [0.14]	0.00185 [0.03]
Domain of Impatience	Recharges	Recharges
Base Case: mean	0.51	0.51
# Individuals	675	675
Observations	53,587	53,587

Notes: This table shows heterogeneity in the effect of the frequency subtreatments by treatment effects of each incentive non-threshold treatment, interacted with measures of impatience over steps; the base case incentive group is omitted. Standard errors clustered at the individual level in brackets. Controls are the same as Table 2. “Discount Rate” indicates a structural measure of the daily discount rate $\frac{1}{\delta_i} - 1$ estimated from a two-limit Tobit model of CTB allocations with individual discount-rate fixed effects, restricting the present-bias parameter β to be one. Larger values of each impatience measure indicates more impatience. The unit of observation is a respondent \times day. Standard errors in brackets clustered at the respondent level. Significance levels: * 10%, ** 5%, *** 1%.

Appendix Table B.2: Payday effects are not bigger for those with higher measured impatience over recharges

Dependent variable:	Met step target	
	Discount Rate	Above- median disc. rate
Impatience measure:	(1)	(2)
Impatience	-0.148 [0.16]	-0.00592 [0.03]
Payday	-0.00594 [0.01]	-0.0117 [0.01]
Impatience \times Payday	0.0906 [0.10]	0.0137 [0.02]
Controls	X	X
Domain of Impatience	Recharges	Recharges
Methodology	Tobit FE	Tobit FE
Dep. var. mean	0.51	0.51
# Base Case	462	462
# Monthly	87	87
# Individuals	549	549
Observations	44,519	44,519

Notes: This table shows heterogeneity in the “payday” effects for those in the base case incentive and the monthly incentive groups, by impatience in the recharge domain. Payday effects are defined as the difference in a daily exercise behavior on paydays compared to all other days. Standard errors clustered by individual are in brackets.

C Supplementary Tables and Figures

Appendix Table C.1: Impacts of incentives and monitoring on diet and addictive consumption.

A. Healthy diet									
	Healthy Diet Index	Wheat meals	Meals with vegetables	Servings fruit	Negative of rice meals	Negative of junk-food pieces	Negative of spoons sugar in coffee	Negative of sweets yesterday)	Avoid unhealthy food
Incentives	0.032 [0.028]	0.028 [0.030]	0.060* [0.031]	0.040 [0.038]	0.029 [0.033]	-0.024 [0.064]	-0.021 [0.047]	-0.027 [0.038]	0.0045 [0.018]
Monitoring	0.012 [0.050]	0.017 [0.053]	0.078 [0.054]	0.061 [0.066]	-0.0086 [0.059]	0.12 [0.11]	-0.028 [0.083]	-0.047 [0.067]	-0.040 [0.031]
Control mean	0.00	0.49	0.58	0.53	-2.34	-0.91	-1.12	-0.35	0.83
P-value: M = I	0.65	0.82	0.71	0.73	0.47	0.14	0.93	0.75	0.11
# Individuals	3,068	3,068	3,068	3,068	3,068	3,068	3,068	3,068	3,068
B. Addictive consumption									
	Addictive Good Consumption Index	Average Daily Areca	Average Daily Alcohol	Average Daily Cigarettes					
Incentives	-0.0076 [0.022]	0.035 [0.042]	-0.035* [0.020]	-0.058 [0.11]					
Monitoring	-0.0017 [0.038]	0.016 [0.074]	-0.016 [0.036]	-0.022 [0.19]					
Control mean	0.00	0.13	0.11	1.02					
P-value: M = I	0.86	0.78	0.53	0.83					
# Individuals	3,068	3,068	3,068	3,068					

Notes: Standard errors in brackets. For the two indices, controls are the same as Table 2, along with second order polynomials of all questions underlying the indices at baseline. We follow Kling et al. (2007) in constructing the indices by imputing missings (for individuals who have non-missing responses to at least one component of the index) for each component using the sample mean. Controls for all other outcomes are the same as Table 2. The Healthy Diet Index is an index created by the average values of eight diet questions, standardized by their average and standard deviation in the control group. The Addictive Good Consumption Index is an index created by the average self-reported average daily consumption of areca, alcoholic drinks, and cigarettes, standardized by their average and standard deviation in the control group. A larger value indicates more consumption. A larger value indicates a healthier diet. The omitted category in all columns is the pure control group.

Appendix Table C.2: Impacts of incentives and monitoring on exercise outcomes, without baseline controls.

	Pedometer Data (Intervention Period)		Self-Reported Data (at Endline)	
	Fraction Days Achieved 10K Steps	Daily Steps	Fraction Days Exercised in Previous Week	Minutes Walked for Exercise Yesterday
A. Pooled Incentives				
Incentives	0.20*** [0.023]	1349.7*** [260.7]	0.057* [0.034]	4.72** [2.34]
Pure Control			-0.13*** [0.037]	-7.94*** [2.45]
B. Unpooled Incentives				
10 INR	0.15*** [0.049]	819.5 [519.0]	-0.0043 [0.065]	2.91 [4.89]
Daily	0.21*** [0.035]	1207.9*** [386.5]	0.027 [0.047]	6.24 [3.89]
Weekly	0.21*** [0.024]	1367.4*** [276.1]	0.054 [0.036]	4.46* [2.52]
Monthly	0.20*** [0.035]	1571.8*** [392.7]	-0.029 [0.048]	4.48 [5.09]
4-Day Threshold	0.20*** [0.025]	1342.1*** [286.5]	0.090** [0.036]	6.19** [2.57]
5-Day Threshold	0.21*** [0.030]	1385.2*** [334.9]	0.053 [0.041]	1.37 [2.89]
Pure Control			-0.13*** [0.037]	-7.94*** [2.45]
Monitoring mean	0.30	6822.66	0.50	22.33
Controls	No	No	No	No
# Monitoring	199	200	195	195
# 10 INR	64	64	62	62
# Daily NTH	163	163	161	161
# Base Case	891	890	867	867
# Monthly NTH	163	163	160	160
# 4-Day TH	778	778	759	759
# 5-Day TH	305	305	293	293
# Control	0	0	568	568
# Individuals	2,563	2,563	3,065	3,065
Observations	203,245	204,614	3,065	3,065

Notes: Standard errors in brackets. The first two columns use daily panel data from pedometers, and standard errors are clustered at the individual level. The second two columns use a cross-section of self-reported data at endline. The omitted category in all columns is the monitoring group.

Appendix Table C.3: Impacts of incentives contracts, compared to the base case non-threshold contract, on the probability of being in 6 quantiles of average exercise outcomes.

Differential Effects of Incentive Contracts on the Distribution of Exercise						
Outcome Quantile	1	2	3	4	5	6
A. Average Step-Target Compliance						
Incentives	-0.080*** [0.022]	-0.029 [0.023]	-0.11*** [0.021]	-0.026 [0.029]	0.017 [0.029]	0.24*** [0.036]
Incentives X (10 INR)	-0.0057 [0.037]	-0.012 [0.039]	0.040 [0.036]	0.029 [0.049]	-0.043 [0.050]	-0.025 [0.061]
Incentives X (Daily NTH)	0.032 [0.024]	-0.017 [0.025]	0.0064 [0.023]	0.036 [0.031]	-0.016 [0.032]	-0.037 [0.039]
Incentives X (Monthly NTH)	-0.018 [0.024]	0.013 [0.025]	0.036 [0.023]	0.017 [0.031]	-0.076** [0.032]	-0.0040 [0.039]
Incentives X (4- or 5-Day TH)	0.023* [0.013]	-0.0031 [0.013]	0.011 [0.012]	-0.0069 [0.017]	-0.038** [0.017]	0.014 [0.021]
Monitoring mean	0.17	0.13	0.16	0.18	0.16	0.16
Controls	Yes	Yes	Yes	Yes	Yes	Yes
B. Average Daily Steps						
Incentives	-0.043* [0.025]	-0.062** [0.025]	-0.095*** [0.023]	-0.019 [0.027]	0.13*** [0.035]	0.071** [0.031]
Incentives X (10 INR)	0.046 [0.042]	-0.00021 [0.043]	0.039 [0.038]	-0.065 [0.047]	-0.044 [0.060]	-0.011 [0.052]
Incentives X (Daily NTH)	0.051* [0.027]	-0.021 [0.027]	0.038 [0.025]	-0.011 [0.030]	-0.070* [0.038]	0.0014 [0.033]
Incentives X (Monthly NTH)	-0.015 [0.027]	0.029 [0.027]	0.034 [0.024]	-0.017 [0.030]	-0.089** [0.038]	0.037 [0.033]
Incentives X (4- or 5-Day TH)	0.027* [0.014]	0.018 [0.015]	0.0027 [0.013]	-0.034** [0.016]	-0.040** [0.020]	0.026 [0.018]
Monitoring mean	0.17	0.17	0.17	0.17	0.17	0.17
Controls	Yes	Yes	Yes	Yes	Yes	Yes
# Monitoring	191	191	191	191	191	191
# Daily NTH	157	157	157	157	157	157
# Base Case	844	844	844	844	844	844
# Monthly NTH	158	158	158	158	158	158
# 4-Day TH	746	746	746	746	746	746
# 5-Day TH	286	286	286	286	286	286
# 10 INR	58	58	58	58	58	58
Observations	2,440	2,440	2,440	2,440	2,440	2,440

Notes: This table presents a series of regressions of dummies for six average daily step quantiles, and six fraction-of-days-step-target-exceeded quantiles. The fifth row of each panel shows the additional effect of the threshold incentives contracts compared to the base case contracts. The coefficients are positive for the highest and lowest quintiles, but negative for intermediate quintiles, showing that the threshold treatments push more people to the extremes of walking behavior. Standard errors are in brackets. Controls are the same as Table 2. The omitted interaction with incentives is the base case incentive treatment group.

Appendix Table C.4: Differential Incentive Effects according to Impatience over Recharges and Steps

Dependent variable:	Met step target	Average daily steps	Met step target	Average daily steps
	(1)	(2)	(3)	(4)
Incentives	0.243*** [0.03]	1822.4*** [276.53]	0.215*** [0.02]	1611.1*** [262.57]
10 INR	-0.0673 [0.05]	-900.7* [466.21]	-0.0863* [0.05]	-702.9* [424.77]
Daily	-0.000998 [0.04]	-358.1 [379.87]	-0.0186 [0.03]	-401.3 [318.42]
Monthly	-0.0207 [0.03]	-54.48 [361.36]	-0.0479 [0.03]	-365.3 [293.62]
4-Day TH	-0.0187 [0.02]	-329.4 [208.85]	-0.0157 [0.02]	-164.0 [190.28]
5-Day TH	0.0168 [0.03]	-187.1 [269.46]	-0.0248 [0.03]	-370.1 [235.58]
Impatience	0.155 [0.31]	1605.3 [4183.04]	-0.0000873 [0.02]	123.4 [214.94]
Incentives × Impatience	-0.233 [0.35]	-1997.1 [4484.77]	0.00246 [0.03]	-323.0 [290.21]
10 INR × Impatience	0.391 [0.63]	-239.8 [5605.97]	-0.0785 [0.05]	-728.9 [700.87]
Daily × Impatience	-0.338 [0.42]	-8458.3** [3730.69]	0.0346 [0.06]	354.9 [534.94]
Monthly × Impatience	-0.192 [0.40]	-709.1 [4603.34]	-0.0257 [0.04]	430.0 [609.24]
4-Day TH × Impatience	-0.179 [0.26]	-2008.7 [2437.41]	0.0168 [0.03]	235.6 [256.20]
5-Day TH × Impatience	0.0718 [0.36]	1300.4 [3516.21]	0.0988** [0.04]	990.4** [436.83]
Controls	X	X	X	X
Domain of Impatience	Recharges	Recharges	Steps	Steps
Base Case: mean	.50712114	8189.998	.50712114	8189.998
# Monitoring	97.00	98.00	144.00	145.00
# 10-INR	36	36	47	47
# Daily NTH	92	92	119	119
# Base Case	483	483	623	624
# Monthly NTH	90	90	116	116
# 4-day TH	428	428	558	558
# 5-day TH	159	159	218	218
# Individuals	1,385	1,386	1,825	1,827
Observations	109,903	110,753	145,103	146,163

Notes: This table shows the treatment effects of each Incentive treatment, interacted with measures of impatience over recharges. The dummy for the base case incentive sub treatment is omitted; the "Incentives" coefficient along with other incentive sub treatment dummies are interpreted relative to the base case contract. Standard errors clustered at the individual level in brackets. Controls are the same as Table 2. "Discount Rate" indicates a structural measure of the daily discount rate $\frac{1}{\delta_i} - 1$ estimated from a two-limit Tobit model of CTB allocations with individual discount-rate fixed effects, restricting the present-bias parameter β to be one. Larger values of each impatience measure indicates more impatience.

Appendix Table C.5: Heterogeneity in threshold impacts by baseline walking

	Exceeded Daily Step Target		Average Daily Steps	
	(1)	(2)	(3)	(4)
4- or 5-day TH \times Walking Measure	0.057 [0.04]	0.0000036 [0.00]	629.2 [448.73]	0.033 [0.04]
4- or 5-day TH	-0.026 [0.02]	-0.036 [0.03]	-336.2* [190.02]	-395.1 [318.47]
Walking Measure	0.45*** [0.03]	0.000028*** [0.00]	5563.6*** [339.92]	0.47*** [0.04]
Walking Measure	Step Target Compliance	Average Steps	Step Target Compliance	Average Steps
# Individuals	2,438	2,438	2,438	2,438

Notes: This table shows the treatment effects of the 4- and 5-day threshold treatments, interacted with measures of walking at baseline; the base case incentive group is the omitted group. The sample is limited to the base case, 4-day threshold, and 5-day threshold treatment groups. Standard errors clustered at the individual level in brackets. Controls are the same as Table 2, along with the average phase-in period value of the dependent variable and its square.

Appendix Table C.6: The structural measure of impatience over exercise is not significantly correlated with other variables

Dependent variable: Specification:	Impatience measure: Delta	
	Separate regressions	Pooled regression
Age	0.009 [0.006]	0.004 [0.006]
Female (=1)	0.066 [0.161]	0.078 [0.168]
Previously diagnosed diabetic (=1)	0.158 [0.123]	0.196 [0.175]
HbA1c	-0.020 [0.028]	-0.036 [0.040]
Mean blood pressure	0.001 [0.006]	0.001 [0.006]
Risk aversion (higher means more risk loving)	0.014 [0.063]	0.018 [0.066]
Discount rate (recharges)	0.285 [0.331]	0.364 [0.390]
Above-median baseline steps (=1)	0.057 [0.158]	0.076 [0.152]

Notes: The unit of observation is a respondent. Robust standard errors in brackets. Significance levels: * 10%, ** 5%, *** 1%.

Appendix Table C.7: Time preference heterogeneity robust to including other controls

Dependent variable:	Met step target ($\times 1000$)										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
A. Above 75th Percentile Discount rate											
Above-75th-perc. disc. rate \times Threshold	86.3** [38.5]	87.0** [38.2]	86.4** [38.6]	86.7** [38.5]	87.2** [38.5]	88.9** [38.4]	84.5** [38.6]	86.3** [38.7]	72.6* [37.6]	76.8** [37.7]	81.7** [37.0]
Above-75th-perc. disc. rate	-19.1 [28.0]	-19.4 [28.0]	-19.0 [28.0]	-19.3 [28.0]	-19.5 [28.1]	-21.0 [27.9]	-17.9 [28.2]	-14.9 [28.5]	-12.7 [27.3]	-18.6 [27.3]	-24.6 [26.2]
Threshold	-32.1 [19.8]	-170.6* [96.6]	-28.9 [23.9]	-40.8 [32.5]	34.4 [66.7]	-69.4*** [22.4]	20.6 [39.2]	-23.2 [23.2]	-54.7** [27.0]	-66.4 [46.0]	-22.7 [48.9]
Threshold \times Covariate		2.80 [1.94]	-7.58 [34.7]	12.6 [36.8]	-7.62 [7.25]	0.36*** [0.12]	-16.1 [10.00]	0.023 [0.035]	57.7* [33.1]	0.0065 [0.0064]	- [0.011]
Covariate		1.77 [1.40]	4.71 [25.2]	-4.46 [26.2]	5.60 [5.12]	0.60 [0.68]	7.73 [7.26]	-244.3* [129.9]	152.7*** [27.8]	0.039*** [0.0057]	0.065*** [0.0048]
Threshold \times Covariate ²											0.000000039 [0.00000055]
Covariate ²											- 0.0000012*** [0.0000012]
B. Discount Rate											
Disc. rate \times Threshold	3.65*** [1.36]	3.23** [1.37]	3.69*** [1.36]	3.63*** [1.36]	3.62*** [1.37]	3.79*** [1.39]	3.70*** [1.35]	3.81*** [1.38]	2.09 [1.37]	1.79 [1.23]	1.39 [1.18]
Discount rate	-2.26** [1.13]	-2.29** [1.12]	-2.30** [1.13]	-2.25** [1.14]	-2.26** [1.13]	-2.32** [1.16]	-2.46** [1.12]	-2.28** [1.14]	-1.99* [1.18]	-2.09** [0.99]	-1.88** [0.94]
Threshold	-11.8 [17.1]	-148.5 [96.7]	-10.0 [21.9]	-17.9 [30.7]	50.1 [66.3]	-48.3** [20.0]	43.0 [37.3]	-3.43 [21.5]	-38.7 [25.5]	-51.0 [46.1]	-8.21 [48.9]
Threshold \times Covariate		2.77 [1.94]	-4.07 [34.7]	8.83 [36.8]	-7.07 [7.24]	0.36*** [0.13]	-16.9* [10.0]	0.022 [0.035]	60.9* [33.1]	0.0070 [0.0064]	0.00057 [0.011]
Covariate		1.76 [1.40]	6.32 [25.2]	-2.91 [26.3]	5.56 [5.12]	0.61 [0.68]	8.41 [7.24]	-223.6* [128.3]	152.3*** [27.9]	0.039*** [0.0056]	0.065*** [0.0049]
Threshold \times Covariate ²											5.1e-09 [0.00000056]
Covariate ²											- 0.0000012*** [0.0000013]
Covariate used	-	Age	Female	Prev. diag- nosed dia- betic	HbA1c	Mean arterial blood pres- sure	Risk aver- sion	Discount rate (recharges)	Above- median base- line steps	Baseline steps	Baseline steps
# Observations	110,942	110,942	110,942	110,942	110,942	110,942	110,942	110,942	110,942	110,942	110,942
Base case mean	507.12	507.12	507.12	507.12	507.12	507.12	507.12	507.12	507.12	507.12	507.12

Notes: The sample is restricted to the weekly groups – i.e., the base case (linear) group, and the 2 threshold groups, 4-day threshold and 5-day threshold, pooled here together as “Threshold.” All columns control for the baseline value of the dependent variable. The unit of observation is a respondent \times day. Standard errors in brackets clustered at the respondent level. Significance levels: * 10%, ** 5%, *** 1%