

# Using Preference Estimates to Customize Incentives: An Application to Polio Vaccination Drives in Pakistan\*

James Andreoni                      Michael Callen                      Karrar Hussain  
UC San Diego, NBER              London School of Economics              FAU Erlangen-Nuremberg

Muhammad Yasir Khan              Charles Sprenger  
University of Pittsburgh              Caltech

April 27, 2022

## Abstract

We use estimates of time preferences to customize incentives for polio vaccinators in Lahore, Pakistan. We measure time preferences using intertemporal allocations of effort, and use these estimates to construct individually-tailored incentives. We evaluate the effect of matching contract terms to discounting parameters in a subsequent experiment with the same vaccinators. Our tailored policy is compared to alternatives that either rely on atheoretic reduced-form relationships for policy guidance or apply the same policy to all individuals. We find that contracts tailored to individual discounting outperform this range of policy alternatives.

*JEL classification:* D1, D3, D90

*Keywords:* Structural estimation, Out-of-sample prediction, Discounting, Present Bias

---

\*We are grateful to Rashid Langrial (former Commissioner of Lahore), Muhammad Usman (District Coordinating Officer Lahore), and Zulfiqar Ali (Executive District Officer Health Lahore) for championing the project. We are grateful for financial support from the International Growth Center (IGC), the Pershing Square Fund for Research on the Foundations of Human Behavior, and from the Department for International Development (DFID) Building Capacity for Research Evidence (BCURE) pilot program initiative. We are grateful to Evidence for Policy Design (EPoD), the Center for Economic Research Pakistan (CERP), Katy Doyle, and Sarah Oberst, for excellent logistical support. We are indebted to Michael Best, Eli Berman, Leonardo Bursztyn, Ernesto Dal Bó, Torun Dewan, Asim Fayaz, Frederico Finan, Ed Glaeser, Matthew Kahn, Supreet Kaur, Asim Khwaja, Aprajit Mahajan, Dilip Mookherjee, Andrew Newman, Gerard Padró-i-Miquel, Robert Powell, Matthew Rabin, Gautam Rao, Martin Rotemberg, Frank Schilbach, Erik Snowberg, Joshua Tasoff, Juan Vargas, Pedro Vicente, Nico Voigtlander, Romain Wacziarg, Noam Yuchtman and seminar participants at UC San Diego, UC Berkeley, Claremont Graduate University, Harvard, the London School of Economics, Nova de Lisboa, Stanford, the Pacific Development Conference, the New England Universities Development Consortium, UCLA Anderson, Wharton, Princeton, and Boston University for extremely helpful comments. We are especially grateful to Rohini Pande for many conversations about this project. Hashim Rao and Danish Raza provided excellent research assistance. Research is approved by the Institutional Review Board at Harvard and at UC San Diego. The study is registered in the AEA RCT Registry as AEARCTR-0000417.

# 1 Introduction

The preference parameters governing intertemporal decisions affect a broad range of outcomes, justifying the considerable theoretical and empirical investments made to describe the level and shape of discounting.<sup>1</sup> An understanding of intertemporal preference parameters also provides valuable policy guidance. Indeed, a number of recent policies are motivated by empirical research on time preferences: commitment savings products, default retirement allocations, and the Save More Tomorrow retirement savings program are all partly motivated by the insight that time preferences may be ‘present-biased’ (for discussion and examples see, e.g., Laibson, 1997; Benartzi and Thaler, 2004; Beshears, Choi, Laibson and Madrian, 2009; Ashraf, Karlan and Yin, 2006; Blumenstock, Callen and Ghani, 2018).

Policy interventions may be further enhanced by using individualized, rather than broad, information on time preferences. Differences in experimental measures of time preferences correlate with differences in a number of policy-relevant behaviors such as take-up of commitment devices and credit card borrowing (examples include Chabris, Laibson, Morris, Schuldt and Taubinsky, 2008b; Meier and Sprenger, 2008, 2012, 2010; Ashraf et al., 2006; Dohmen, Falk, Huffman and Sunde, 2006; Castillo, Ferraro, Jordan and Petrie, 2011).<sup>2</sup> These correlations suggest that interventions could leverage individual information on time preference to tailor unique policies for each person. The purpose of tailoring policies is to shape real-world behavior. Therefore, we study the potential to deploy lab protocols for measuring preferences and structurally tailoring contracts in the field.

This paper studies the promise of theoretically-informed, individually-tailored policy inter-

---

<sup>1</sup> Central examples of theoretical work include Samuelson (1937); Koopmans (1960); Laibson (1997) and O’Donoghue and Rabin (2001). Empirical exercises in field and laboratory settings focusing on parameter estimation include Hausman (1979); Lawrance (1991); Warner and Pleeter (2001); Cagetti (2003); Laibson, Repetto and Tobacman (2005); Mahajan, Michel and Tarozzi (2020); Harrison, Lau and Williams (2002); Andreoni and Sprenger (2012).

<sup>2</sup>It should be noted that none of these examples linking structural estimates of time preference to other behaviors provide an articulated model for what the precise correlation between the two values should be. Unlike our own efforts, such exercises could be conducted without appeal to structural estimation.

ventions in a real-world policy setting. Our project engages government health workers—termed Lady Health Workers (LHWs)—associated with polio eradication efforts for the Department of Health in Lahore, Pakistan.<sup>3</sup> The function of LHWs is to provide oral polio vaccine to children during monthly vaccination drives, which usually last two days. We introduce a monitoring and incentive system to measure intertemporal preferences via effort choices at work. Closely following the Convex Time Budget (CTB) design of Andreoni and Sprenger (2012); Augenblick, Niederle and Sprenger (2015), LHWs are asked to trade off work between the two days of a vaccination drive. Unlike standard laboratory measurements of time preferences, for empirical realism, each LHW only makes a single CTB choice. Completion of the allocated work is tied to a bonus of 10 times the standard LHW daily wage. Under a set of structural assumptions, the distribution of work allocations identifies the distribution of time preferences accounting for shocks to the marginal costs of effort. Each LHW’s choice can, in turn, be linked to an expected level of discounting at the individual level under the estimated distributions of time preferences and shocks.

We then tailor policy based upon these measured preferences in a subsequent work decision. The structural, tailored policy we examine attempts to equalize vaccinations over time by changing relative prices for each LHW.<sup>4</sup> We compare this tailored policy to alternatives that span the policy space on two dimensions: broad vs. tailored and structural vs. atheoretic. A broad policy is one that is applied uniformly to all individuals, while a tailored policy is one individualized to each based on some characteristic. In our setting, a broad policy is one where all vaccinators face the same price for trading vaccinations off across the two days, while an individualized policy is one where each vaccinator faces their own price. A structural policy

---

<sup>3</sup>Polio is endemic in Pakistan. Of 350 new worldwide cases in 2014, 297 occurred in Pakistan, constituting a ‘global public health emergency’ according to the World Health Organization. Between 95 percent and 99 percent of individuals carrying polio are asymptomatic. One infection is therefore enough to indicate a substantial degree of ambient wild polio virus. The disease largely affects children under five.

<sup>4</sup>Our tailored policy was pre-specified prior to conducting the study, and was chosen to precisely equalize vaccinations in the absence of shocks. Accounting for shocks to marginal costs in our distributional estimates alters the policy target slightly from equal vaccinations over time.

is one that draws policy guidance from a theoretical model of preferences, while an atheoretic policy draws policy guidance from some reduced-form relationship.<sup>5</sup> In our setting, we use a model to translate LHWs' allocations of vaccinations to an estimate of their preferences. Our structural policy uses information on measured preferences, while atheoretic policies do not. Our structural, tailored policy is compared to broad policies which set relative prices to achieve the same objective based on aggregate values, or the reduced-form price-sensitivity of effort; and a tailored, but atheoretic policy, which sets relative prices based on a simple rule of giving higher relative prices to plausibly more patient individuals. These comparisons are facilitated by an overarching control group that receives a uniform random price, from which we draw subsets for comparison purposes.

In a sample of 338 LHWs, we document three principal findings. First, on aggregate, we estimate effectively no present bias in between-subjects comparisons. LHWs choosing their work in advance of the first day of the vaccination drive allocate slightly more of their vaccinations to the first day of the drive than those allocating on the morning the drive actually commences, but corresponding estimates of the distribution of time preferences indicate only a small and insignificant degree of present bias. Interestingly, when examining within-subjects data, comparing individual LHWs across drives, our estimates of present bias grow more substantial, and accord with estimates from prior laboratory studies of present bias in effort. Second, there is a large degree of heterogeneity in time preferences across subjects and this heterogeneity is markedly more pronounced when LHWs make immediate choices. This sizable cross-sectional variation also resonates with prior experimental exercises. Third, and most importantly, our structural, tailored policy works. Relative to a range of policy alternatives, our intervention generates behavior around 30% closer to the policy target of equal allocation. Interestingly, when focusing on conditions where subjects are asked to make allocations which take effect immediately – that is, when present bias may be relevant – the tailored policy generates a

---

<sup>5</sup>Though far from an exhaustive labeling of potential policies, this 2-by-2 labeling helps to organize the comparisons we investigate.

roughly 50% improvement.

This paper makes two contributions. First, our exercise uses field behavior about effort to examine time preferences, reflecting, to our knowledge, the first attempt to use the Augenblick et al. (2015) methods for measurement based on an actual real-world work (vaccinations).<sup>6</sup> This effort yields some valuable lessons for future field implementations. While lab experiments have elicited preferences based on a few hours of effort, participants in our study make decisions over days' worth of work. Shocks to the marginal cost of effort have plausibly greater consequences when allocating several hundred vaccinations attempts in the field rather than fifty greek transcriptions online as in Augenblick et al. (2015). We, correspondingly, find an important role for shocks in our estimations. The volume of work required in our field setting generates a further difference from lab implementations. While roughly 90% of Augenblick et al.'s subjects completed all allocated tasks and requirements of the study and received a \$100 completion bonus, fewer of our subjects completed all their allocated vaccinations. While most LHWs appear to have made substantial efforts to meet their targets, completing about 75% of the required vaccinations, only about 50% of our sample successfully completed all vaccinations, even with a bonus of 10 times the daily LHW wage.<sup>7</sup> This difference between the lab and the field requires us to account for probabilistic completion and the potential distortionary effects of noncompletion in our analysis. We develop empirical tools for estimating the distribution of time preferences from a limited number of CTB choices, accounting for shocks and probabilistic completion, which may be valuable for researchers conducting future field efforts in this vein.

---

<sup>6</sup> Documenting dynamic inconsistency outside of the laboratory and outside of the standard experimental domain of time-dated monetary payments is particularly valuable given recent discussions on the elicitation of present-biased preferences using potentially fungible monetary payments (Cubitt and Read, 2007; Chabris, Laibson and Schuldt, 2008a; Andreoni and Sprenger, 2012; Augenblick et al., 2015; Carvalho, Meier and Wang, 2014).

<sup>7</sup>The lack of uniform completion was not a feature of the data we initially expected, but, in retrospect, is something we could have anticipated. Data from drives prior to our intervention, which were subject to almost no monitoring or scrutiny, showed that LHWs almost without exception hit their prescribed targets exactly. We believe these reports are at least partially driven by the fact that polio is a politicized issue in Pakistan, with a number of stakeholders and international donors being eager to demonstrate high numbers of vaccinations. Given our lack of foresight, the functional forms accounting for probabilistic non-completion were in our study registration. As such, they should be viewed with appropriate caveats.

The consistencies between our findings and prior lab experimental work helps to support the growing literature which identifies present bias from CTBs and non-monetary choices in the field (Read and van Leeuwen, 1998; Sadoff, Samek and Sprenger, 2015; Read, Loewenstein and Kalyanaraman, 1999; Sayman and Onculer, 2009; Kaur, Kremer and Mullainathan, 2010, 2015; Carvalho et al., 2014).<sup>8</sup>

Second, we provide the first empirical evidence to date of the value in structural, tailored policies for intertemporal choice in a field setting (or otherwise). Given these results, there is clear opportunity to expand the scope of interventions beyond uniform policy strategies. The policy objective we consider attempts to implement a smooth allocation over time. In our setting equalizing vaccine provision is important for the logistics of the vaccine supply chain. Polio vaccine requires cold storage. Managers therefore prefer that the amount of vaccine delivered every day meets a consistent and manageable target. Smoothing is also a natural objective to consider for other intertemporal decisions but may, of course, not be appropriate in all settings.<sup>9</sup> There are many other policy alternatives that could be considered, including the alternative of maximizing total vaccine provision<sup>10</sup> Hence, our results provide a simple

---

<sup>8</sup>These studies include examination of present bias or dynamic inconsistency for food choices (Read and van Leeuwen, 1998; Sadoff et al., 2015); for highbrow and lowbrow movie choices (Read et al., 1999); for cafe reward choices (Sayman and Onculer, 2009); for completing survey items (Carvalho et al., 2014); and for fertilizer purchase decisions (Duflo, Kremer and Robinson, 2011). For discussion of this literature, see Sprenger (2015).

<sup>9</sup>Evidence suggests that consumption of Supplemental Nutritional Assistance Program (SNAP) benefits may be subject to present bias, leading to declining consumption during the benefit period (see, e.g., Shapiro, 2005; Hastings and Washington, 2010). Our results indicate that a tailored policy along the lines implemented here could help smooth benefit consumption. Such a policy could complement alternatives that have been discussed, including increasing the frequency of benefit payments (Shapiro, 2005).

<sup>10</sup> One can consider a range of tailoring exercises using interest rate,  $R$ , to maximize a policymaker's objective function subject to each LHW's offer curve. If  $P(w_{1,i}(R), w_{2,i}(R))$  is the objective function at LHW allocation  $(w_{1,i}, w_{2,i})$ , then under the structural assumptions and experimental design of this paper the general problem is

$$\begin{aligned} \max_R P(w_{1,i}^*(R), w_{2,i}^*(R)) \text{ s.t.} \\ (w_{1,i}^*(R), w_{2,i}^*(R)) = \operatorname{argmin} (w_{1,i})^\gamma + \beta_i^{1-d} \delta_i \cdot (w_{2,i})^\gamma \text{ s.t.} \\ w_{1,i} + R \cdot w_{2,i} = V. \end{aligned}$$

Our exercise assumes a Leontief policymaker,  $P(w_{1,i}, w_{2,i}) = \min[w_{1,i}, w_{2,i}]$ . If, instead, one assumed a maximizing policymaker,  $P(w_{1,i}, w_{2,i}) = w_{1,i} + w_{2,i}$ , then when  $\gamma = 2$ , the optimal  $R^* = \sqrt{\beta_i^{1-t=1} \delta_i (1 + \beta_i^{1-d=1} \delta_i)} - \beta_i^{1-d=1} \delta_i$ . We estimate average values of  $\beta_i \delta_i$  around 0.95, and so  $R^*$  for maximizing vaccinations is roughly 0.4. This was an unrealistic value for implementation in our setting, but such a policy should be investigated in future work.

proof-of-concept for tailoring to a range of targets.

The paper proceeds as follows: Section 2 presents our experimental design and corresponding theoretical considerations for estimating time preferences and tailoring contracts, Section 3 presents results, Section 4 provides robustness tests, and Section 5 concludes.

## 2 Experimental Design

Our experiment has three components: implementing a high resolution smartphone monitoring system similar to that described in Callen, Gulzar, Hasanain, Khan and Rezaee (2019), eliciting individual discounting parameters using the Convex Time Budget (CTB) technique (Andreoni and Sprenger, 2012; Augenblick et al., 2015), and, after assigning tailored contracts to LHWs, testing whether these tailored contracts outperform comparison policies.

### 2.1 Vaccinations, Smartphone Monitoring, and Evaluation of Monitoring Technology

The Department of Health in Lahore, Pakistan, employs LHWs throughout the city to conduct polio vaccination drives. Every month there is a vaccination drive that is at least two days long. Prior to our study, the standard protocol for vaccination drives was to provide each LHW a fixed target for total vaccinations over the drive and a map of potential households (called a “micro-plan”). LHWs received no explicit benefits for reaching targets; they received a fixed daily wage of 100 rupees (around \$1 at contemporary exchange rates). LHWs mapped their walk with pen and ink, knocking on each compound door, and vaccinating each child if their parents granted permission. Vaccinating a child consists of administering a few drops of oral vaccine. As there is no medical risk of over-vaccination, LHWs are encouraged to vaccinate every child for whom permission is granted. For each attempted vaccination, LHWs were asked to mark information related to the attempt (number of children vaccinated, whether or not

all children were available for vaccination, etc.) in chalk on the compound wall. Appendix Figure A.1 provides an example of neighborhood micro-plan, Appendix Figure A.2 provides an example of a vaccination attempt, and Appendix Figure A.3 provides a picture of a chalk marking on a compound wall. At the end of each day, LHWs in each neighborhood convened with their supervisor and self-reported their vaccination activity for the day (see Appendix Figure A.4 for an example of the form). In principle, a monitor could verify the claims. In practice, however, there was virtually no monitoring, and reasons to suspect over-reporting.<sup>11</sup>

In collaboration with the Department of Health, we designed a smartphone-based monitoring system. The Department of Health provided a smartphone equipped with a vaccination monitoring application to all LHWs in our sample in order to record information related to each vaccination. For each vaccination attempt, the LHW was asked to take a picture of the home visited and her current vial of vaccine. An image of the main page of the application is provided as Figure 1, Panel A. Data from the smartphone system were aggregated in real-time on a dashboard available to senior health administrators (see Appendix Figure A.5 for an example of the dashboard). In order to separate the effects of this smartphone monitoring system from those of our incentive program discussed next, a sample of vaccinators (157 total between two vaccination drives) were given a smartphone equipped with the application and instructed on its use, but were not given any additional incentives beyond their daily wage.

---

<sup>11</sup>We attempted to independently audit LHWs by following the trail of chalk markings, but our enumerators found the process too difficult to produce a reliable audit of houses visited. We do, however, know the targets associated with each micro-plan prior to our monitoring intervention and that LHWs almost always reported meeting their targets exactly. Even with a bonus incentive and smartphone monitoring in place, we find that LHWs on average achieve only 62 percent (s.d. = 58 percent) of the target given by their micro-plans. LHWs likely would achieve a smaller share of their target in the absence of both monitoring and financial incentives.



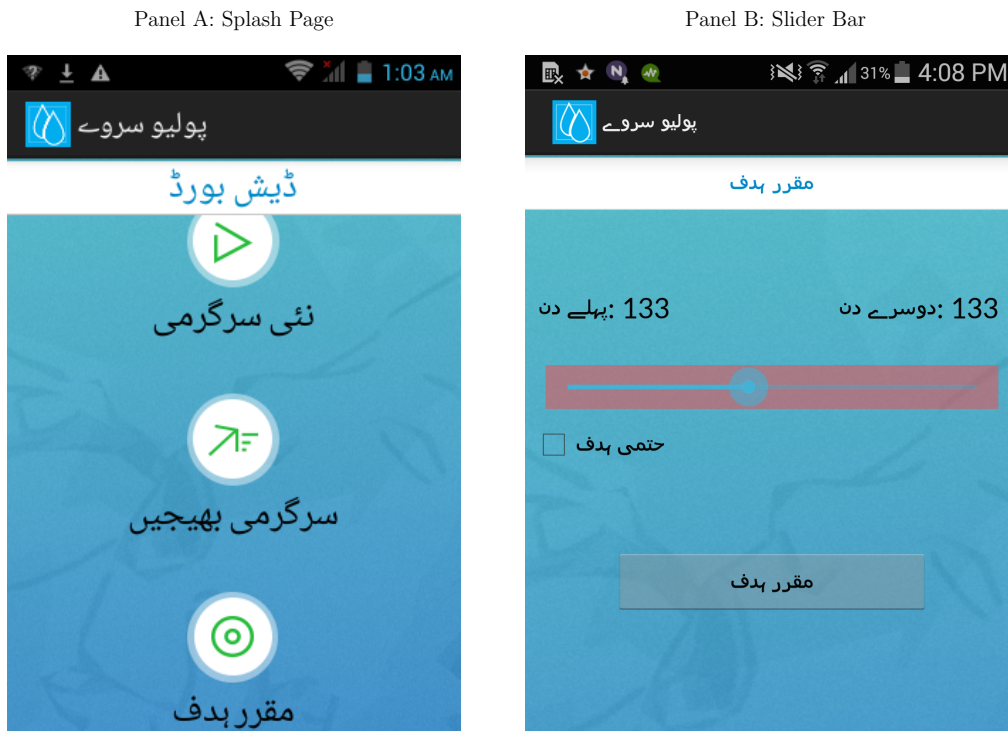


Figure 1: Vaccination Monitoring Smartphone App

*Notes:* The picture is of two screenshots from the smartphone app used by Vaccinators. Panel A is depicted after partially scrolling down. The top bar in Panel A (white letters) translates to “polio survey.” The next panel down (blue letters) translates to “Dashboard” (literally transliterated). The black letters under the top button translate to “new activity”, the letters under the second button translate to “send activity” and the letters under the lowest button translate to “set target”. The blue letters in panel B translate to “set target”. The next line translates to “First day: 133; Second day: 133”. The text next to the box translates to “finalize target” and the black letters on the bar translate to “set target.”

## 2.2 Drive 1: Intertemporal Bonus Contracts and the Measurement of Preferences

The smartphone monitoring application was equipped with additional functionality for implementing intertemporal bonus contracts. These contracts had LHWs set daily work targets in order to receive a substantial bonus. The chosen targets provide critical intertemporal preference information that we use to estimate the distribution of preferences, and, subsequently, to tailor future contracts.

We worked with the Department of Health to implement these contracts in two-day drives in September, November and December of 2014. The contracts required workers to complete a present value total of  $V = 300$  vaccination attempts in exchange for a fixed bonus of 1000 rupees. LHWs set daily targets,  $v_1$  and  $v_2$ , corresponding to vaccinations on day 1 and day 2 of the drive, respectively. If either of the vaccination targets,  $v_1$  or  $v_2$ , were not met, the 1000 rupees would not be received, and the LHW would receive only her standard wage.

Each LHW was randomly assigned a relative price,  $R$ , translating vaccinations on day 1 to vaccinations on day 2. For each vaccination allocated to day 2, the number of vaccinations allocated to day 1 would be reduced by  $R$ . Hence, the targets  $v_1$  and  $v_2$  satisfy the intertemporal budget constraint

$$v_1 + R \cdot v_2 = V.$$

This bonus contract is identical to an experimental device termed a Convex Time Budget used to investigate time preferences (Andreoni and Sprenger, 2012; Augenblick et al., 2015).<sup>12</sup> Intertemporal allocations,  $(v_1, v_2)$ , carry information on the time preferences of each LHW that can be used to estimate the distribution of time preferences.

---

<sup>12</sup>We also borrow an additional design element from such studies: minimum allocation requirements. In order to avoid LHWs allocating all their vaccinations to a single day of the drive, we placed minimum work requirements of  $v_1 \geq 12$  and  $v_2 \geq 12$ . The objective of minimum allocation requirements is to avoid confounds related to fixed costs. That is, by requiring LHWs to work on both days of the drive, we avoid confounding extreme patience or extreme impatience with LHWs simply not wishing to come to work on one of the two days.

### 2.2.1 Experimental Variation

Our design leverages two sources of experimental variation. First, each LHW is randomly assigned a single relative price,  $R$ , from the set  $R \in \{0.9, 1, 1.1, 1.25\}$ . These values were chosen following Augenblick et al. (2015). Operationally, experimental variation in  $R$  was implemented by providing each LHW with a slider bar on the introduction screen of the smartphone application. Figure 1, Panel B depicts the slider bar with  $R$  equal to 1.25. The LHW was asked to pull the slider bar to their desired allocation  $(v_1, v_2)$  and then submit. The allocation was required to be submitted before commencing vaccination.

Second, each LHW was randomly assigned to either submit their allocation in advance of day 1 of the drive or on the morning of day 1. We refer to the first of these as the ‘Advance’ treatment arm and the second as ‘Immediate’ treatment arm. The assignment to either the Advance or the Immediate group was cross-randomized with the assignment of  $R$ , creating a 2 x 4 design within our incentive treatments. Section 2.4 describes the efforts taken to make everything else besides allocation timing equal between the Advance and Immediate conditions.

Variation in the timing and interest rate of allocations provides relevant experimental variation for estimating the distribution of time preferences and evaluating whether the distribution is dynamically consistent. The distribution of LHW allocations in these conditions is mapped to a distribution of discounting parameters accounting for the influence of shocks. Distributional estimates in hand, we calculate an expected discount factor for each LHW based on their allocation under the estimated distributions of time preferences and shocks.

### 2.2.2 Structural Assumptions for Estimating Time Preferences

We make a number of structural assumptions in order to map from LHW allocations to discounting parameter distributions. First, we assume a stationary power cost of effort function  $c(v) = v^\gamma$ , where  $v$  represents vaccinations performed on a given day, and  $\gamma > 1$  captures convex costs of effort. In our pre-specified analysis plan, we posited  $\gamma$  to be constant across individuals

and our tailoring exercise was conducted under the assumption of  $\gamma = 2$ , quadratic costs. While the field implementation of intertemporal discounting experiments is likely to differ substantially from laboratory examples, quadratic costs are close to the prior laboratory findings for effort (see, e.g., Augenblick et al., 2015). While the cost function is constant across individuals and over time, we assume below that individual allocations are subject to random shocks in relative marginal costs across period, and we estimate the distribution of those shocks. Hence, in calculating an individual's expected discount factor we take into account this potential source of variation in costs.

Second, we assume that individuals discount the future quasi-hyperbolically (Laibson, 1997; O'Donoghue and Rabin, 1999). For a given LHW,  $i$ , making allocation  $(v_{1,i}, v_{2,i})$  the discounted disutility of effort can be written as

$$v_{1,i}^\gamma + \beta_i^{\mathbf{1}_{d=1}} \delta_i \cdot v_{2,i}^\gamma.$$

The indicator  $\mathbf{1}_{d=1}$  captures whether the decision is made in advance or immediately on day 1. The parameters  $\beta_i$  and  $\delta_i$  summarize individual discounting. The parameter  $\beta_i$  captures the degree of present bias active for LHWs who make Immediate decisions when  $\mathbf{1}_{d=1} = 1$ . If  $\beta_i = 1$ , the vaccinator adheres to exponential discounting with discount factor  $\delta_i$ , while if  $\beta_i < 1$  the vaccinator exhibits a present bias, being less patient in Immediate relative to Advance decisions.

Third, we assume that the distribution of one period discount factors is normal in each condition

$$\begin{aligned} \delta_i &\sim N(\mu_\delta, \sigma_\delta^2) \quad \text{if } d = 0 \\ \beta_i \delta_i &\sim N(\mu_{\beta\delta}, \sigma_{\beta\delta}^2) \quad \text{if } d = 1. \end{aligned}$$

Our objective is to estimate the relevant distributional parameters of the population:  $\mu_\delta, \sigma_\delta^2$  in the Advance condition, and  $\mu_{\beta\delta}, \sigma_{\beta\delta}^2$  in the Immediate condition.

Fourth, we assume that LHWs minimize the discounted costs of effort subject to the intertemporal budget constraint provided by their bonus contract. This yields a marginal condition

$$R\gamma v_{1,i}^{\gamma-1} - \beta_i^{1_{d=1}} \delta_i \gamma v_{2,i}^{\gamma-1} = 0. \quad (1)$$

Note that in the absence of shocks, equation 1 provides a deterministic mapping from allocation behavior to individual discount factors,

$$R * \left( \frac{v_{1,i}}{v_{2,i}} \right)^{\gamma-1} = \beta_i^{1_{d=1}} \delta_i, \quad (2)$$

which can be used to provide initial guidance on the extent of heterogeneity in choice.

Fifth, we assume vaccinators' behavior is not the deterministic function of their preferences provided in equations 1 and 2. Rather, their optimization is subject to an additive shock to relative marginal costs in the two periods,  $\epsilon_i$ . We assume these relative marginal cost shocks to be normal and mean zero,

$$\epsilon_i \sim N(0, \sigma_\epsilon^2).$$

Minimizing discounted costs subject to the intertemporal budget constraint of the experiment and adding the relative marginal cost shock yields an adjusted marginal condition:

$$\gamma v_{1,i}^{\gamma-1} - \frac{\beta_i^{1_{d=1}} \delta_i}{R} \gamma v_{2,i}^{\gamma-1} = \epsilon_i. \quad (3)$$

---

<sup>13</sup>This approach for structurally estimating time preferences by assuming a marginal condition is satisfied up to a random shock was introduced in controlled experiments by Andreoni and Sprenger (2012), and has precedents in a body of macroeconomic research identifying aggregate preferences from consumption data. See, for example, Shapiro (1984); Zeldes (1989); Lawrance (1991).

Under the above assumptions, the conditional likelihood of an allocation  $(v_{1,i}, v_{2,i})$  given  $\beta_i^{1_{d=1}} \delta_i$  is

$$L(v_{1,i}, v_{2,i} | \delta_i) = \phi \left( \frac{\gamma v_{1,i}^{\gamma-1} - \frac{\delta_i}{R} \gamma v_{2,i}^{\gamma-1}}{\sigma_\epsilon} \right) \quad \text{if } d = 0$$

$$L(v_{1,i}, v_{2,i} | \beta_i \delta_i) = \phi \left( \frac{\gamma v_{1,i}^{\gamma-1} - \frac{\beta_i \delta_i}{R} \gamma v_{2,i}^{\gamma-1}}{\sigma_\epsilon} \right) \quad \text{if } d = 1$$

where  $\phi(\cdot)$  is the density of the standard normal distribution. Integrating over the relevant distribution of preference parameters in Advance and Immediate conditions gives the likelihoods

$$L(v_{1,i}, v_{2,i}) = \int \phi \left( \frac{\gamma v_{1,i}^{\gamma-1} - \frac{\delta_i}{R} \gamma v_{2,i}^{\gamma-1}}{\sigma_\epsilon} \right) f(\delta_i) d\delta_i \quad \text{if } d = 0 \quad (4)$$

$$L(v_{1,i}, v_{2,i}) = \int \phi \left( \frac{\gamma v_{1,i}^{\gamma-1} - \frac{\beta_i \delta_i}{R} \gamma v_{2,i}^{\gamma-1}}{\sigma_\epsilon} \right) g(\beta_i \delta_i) d\beta_i \delta_i \quad \text{if } d = 1, \quad (5)$$

where  $f(\cdot)$  and  $g(\cdot)$  are the normal densities from which  $\delta_i$  and  $\beta_i \delta_i$  are drawn. We deploy the Method of Simulated Likelihood (MSL) with 1000 random Halton draws from the simulated distributions,  $f(\cdot)$  and  $g(\cdot)$ , at each observation to estimate simulation analogs of equations (4) and (5). Appendix A.1 provides additional detail on the estimation methodology and robustness of corresponding results.

In our setting, vaccinators receive a bonus of 1000 rupees paid the day after the drive if they meet both targets  $v_1$  and  $v_2$ .<sup>14</sup> Not completing allocated vaccinations creates a sizable penalty at any given point in time. This design choice was made to encourage LHWs to forecast that they will complete the required vaccinations and so allocate them according to their true preferences. Nonetheless, the total present value of vaccinations,  $V = 300$ , is ambitious. Even with substantial effort, LHWs may not meet their targets, and so may need to account for

<sup>14</sup>The choice of large bonuses (around 10 times daily wages) followed the design logic discussed in Augenblick et al. (2015), who implemented a \$100 completion bonus.

the potential for failure when making their allocations. To incorporate potential failure, we introduce a sixth assumption: LHWs forecast completing  $v$  allocations with probability

$$p(v) = \frac{1}{1 + \alpha v}.$$

Forecasted non-completion, introduces a wedge in the LHW's marginal condition, as they must recognize the marginal impact of an allocation on failure and loss of the completion bonus. This, in turn, changes the estimation procedure in two ways: the completion probabilities,  $p(v_{1,i})$  and  $p(v_{2,i})$ , are a component of the total likelihood of an observation, and the component of the likelihood associated with the marginal condition is adjusted for the non-completion loss of the bonus. In Appendix A.1 we present the extended methodology that incorporates probabilistic non-completion in estimation and provide additional empirical results. For completeness we present our distributional results with and without the incorporation of the completion parameter,  $\alpha$ , and find little difference in estimated time preferences.

Our structural exercise links the distributions of behavior in each of our 8 experimental conditions of  $R$  and Advance vs Immediate Choice in Drive 1 to a heterogeneous distribution of discount factors and a distribution shocks. Absent heterogeneity in discount factors or shocks, behavior would be deterministic, and all LHWs in a given condition should provide identical allocations. Permitting shocks alone, the distributions of behavior would reflect a constant shock variance, but be displaced in a proportional way across conditions (e.g., from  $\phi\left(\frac{\gamma v_{1,i}^{\gamma-1} - \frac{\delta}{R} \gamma v_{2,i}^{\gamma-1}}{\sigma_\epsilon}\right)$  to  $\phi\left(\frac{\gamma v_{1,i}^{\gamma-1} - \frac{\delta}{R'} \gamma v_{2,i}^{\gamma-1}}{\sigma_\epsilon}\right)$  between  $R$  and  $R'$ ). A substantial role for heterogeneous preferences would be recovered if this restriction provided a relatively poor fit to the data. Identifying heterogeneity in preferences in this way may be challenging if assumptions like a constant distribution for  $\epsilon$  are inaccurate. Importantly, our project is predicated on using the resulting measures to tailor contracts and predict subsequent Drive 2 allocations. As such, if we mischaracterize behavior that should be attributable to shocks or other forces as evidence

of heterogeneous discounting, our exercise should do a notably poor job of tailoring incentives.

### 2.2.3 Deriving Individual Expected Discount Factors From Allocation Behavior

The methodology discussed to here yields estimated distributions of preference parameters and additive marginal cost shocks  $N(\hat{\mu}_\delta, \hat{\sigma}_\delta^2)$ ,  $N(\hat{\mu}_{\beta\delta}, \hat{\sigma}_{\beta\delta}^2)$ , and  $N(0, \hat{\sigma}_\epsilon^2)$ . These distributions in hand, we map back to expected values of discounting for each individual. We do this by simulating behavior (allocations of  $v_{1,i}$  and  $v_{2,i}$ ) using the estimated distributions –  $N(\hat{\mu}_\delta, \hat{\sigma}_\delta^2)$ ,  $N(\hat{\mu}_{\beta\delta}, \hat{\sigma}_{\beta\delta}^2)$ , and  $N(0, \hat{\sigma}_\epsilon^2)$  – and then calculating the expected value of  $\delta_i$  or  $\beta_i\delta_i$  associated with a given allocation,  $E[\delta_i|v_{1,i}, v_{2,i}, R]$  if  $d = 0$  and  $E[\beta_i\delta_i|v_{1,i}, v_{2,i}, R]$  if  $d = 1$ . Specifically, we take 1 million draws from the distribution of preference parameters and shocks separately for the Advance and Immediate Conditions. We then simulate the distribution of allocations at each value of  $R \in \{0.9, 1, 1.1, 1.25\}$ . We construct  $E[\delta_i|v_{1,i}, v_{2,i}, R]$  if  $d = 0$  and  $E[\beta_i\delta_i|v_{1,i}, v_{2,i}, R]$  if  $d = 1$  as the expected discount factor associated with simulated allocations within 2.5 tasks of the actual allocation’s value of  $v_{1,i}$ . Such simulation-based methodology for calculating individual level parameters is analogous to the approach taken for discrete choice models after MSL estimation (See Train (2009) for examples).

The assumptions required for the calculation of expected discount factors are potentially quite restrictive. Our research design, which involves tailoring contracts to discounting, required commitment to specific functional forms for costs and discounting. Our empirical exercise estimates the distribution of discounting parameters accounting for shocks to relative marginal costs and the probability of non-completion. In sub-section 4.1, we assess the validity of our assumptions and present further exploratory analysis related to alternative functional forms.

## 2.3 Drive 2: Test of Structural, Tailored Contracts

In a second two-day drive, we investigate tailored contracts. All LHWs from the first drive were invited to participate in a second intertemporal bonus contract. LHWs were unaware that their



previously measured behavior would be used to potentially inform their subsequent contracts. This sidesteps an important possibility that LHWs might alter their first drive behavior in order to receive a more desirable contract in the second drive.

Half of LHWs received an individually-tailored intertemporal bonus contract,

$$w_{1,i} + R_i^* \cdot w_{2,i} = V,$$

where  $w_{1,i}$  and  $w_{2,i}$  are Drive 2 allocations. The value

$$R_i^* = \frac{R \cdot v_{1,i}}{v_{2,i}} \tag{6}$$

is their individually-tailored interest rate based upon their Drive 1 allocations. Absent shocks and with perfect completion, when  $\gamma = 2$  in (1),  $\frac{R \cdot v_{1,i}}{v_{2,i}} = \beta_i \delta_i$  or  $\delta_i$ , depending on whether the LHW made an Immediate or Advance decision in Drive 1.<sup>15</sup> Hence, the tailored contract was pre-specified to be equal to the LHW's discount factor under deterministic choice and perfect completion. Absent shocks and probabilistic completion, setting the relative price,  $R_i^*$ , in this way should lead LHWs to allocate an equal number of vaccinations to each day of the drive,  $w_{1,i} = w_{2,i}$ . Accounting for shocks and probabilistic completion alters this prediction slightly, but the prediction of roughly equal allocations to each period is maintained. Though LHWs in this group receive different relative prices, the contract is designed for each of them to achieve the same objective of smoothing vaccinations through time. Some LHWs' allocation behavior in the first drive implied extreme values of  $R_i^*$ . Our tailoring exercise focused only on a Tailoring Sample of LHWs with values of  $R_i^*$  between 0.75 and 1.5.<sup>16</sup> LHWs outside of these bounds were given either the upper or lower bound accordingly.

---

<sup>15</sup>Note that this tailoring exercise requires that LHWs remain in either the Immediate or Advance assignment across drives.

<sup>16</sup>Of our sample of 338 LHWs, 57 exhibit  $\frac{R \cdot v_{1,i}}{v_{2,i}}$  outside of this range. The Tailoring Sample consists of the remaining 281 LHWs. One LHW, out of 281 does, not report vaccination activity via mobile phone in Drive 2.

The other half of LHWs were given a random intertemporal bonus contract,

$$w_{1,i} + \tilde{R}_i \cdot w_{2,i} = V,$$

where  $\tilde{R}_i$  was drawn from a random uniform distribution  $U[0.75, 1.5]$ . The bounds on the distribution of  $\tilde{R}_i$  were determined to match the bounds on  $R_i^*$ , while the choice of a random uniform control—rather than a single value of  $\tilde{R}_i$  or some alternative distribution—was chosen to provide flexible scope for constructing a range of comparison policies by drawing subsets of LHWs assigned to the  $\tilde{R}_i$  condition. Relevant subgroups that we draw from this group of LHWs are: 1) structural, broad: those with values of  $\tilde{R}_i$  close to the average value of  $R_i^*$ ; 2) atheoretic, broad: those with values of  $\tilde{R}_i$  that are close to the optimal value for achieving  $w_{1,i} = w_{2,i}$  implied by a reduced form exercise; 3) atheoretic, tailored: those with values of  $\tilde{R}_i$  that are generally increasing in patience but not required to be linear as in the structurally tailored policy,  $R_i^* = \frac{R \cdot v_{1,i}}{v_{2,i}}$ . These comparisons span the policy space of being either atheoretic vs. structural and tailored vs. broad. Comparison is also provided for the full group of LHWs who received random bonus contracts.

Random assignment to structural tailoring in Drive 2 is stratified on the measure of absolute distance to equal allocations  $|\frac{v_{1,i}}{v_{2,i}} - 1|$ , based on allocations from Drive 1.<sup>17</sup> This measure of distance to equal provision also serves as our eventual outcome measure when analyzing the effect of assignment to structural tailoring in Drive 2. Stratifying assignment on key outcomes of interest is standard practice in the field experimental literature (Bruhn and McKenzie, 2009), as it generally increases precision in estimating treatment effects.

---

<sup>17</sup>Specifically, subjects are divided into terciles by this measure, with a roughly even number in each bin being assigned to the tailoring and to the control condition.

## 2.4 Design Details

Our experiment is divided into two drives. The first drive took place November 10-11, 2014 with training on November 7. The second drive took place December 8-9, 2014 with training on December 5. These drives are denoted Drive 1 and Drive 2 and are used for measuring preference, and tailoring contracts, respectively. We had actually anticipated beginning our field study in September, 2014. However, as noted below, a disruption to the mobile network prevented us from measuring the preferences for all LHWs in this drive. We call this failed drive Drive 0, and leverage the data that was collected for within-subjects panel analysis.

### 2.4.1 Training and Allocation Decisions

On November 7, all LHWs participating in Drive 1 received two hours of training at one of three locations in central Lahore on using the monitoring features of the smartphone application and the process by which allocations were made. Both Advance and Immediate LHWs were given identical training. At the end of the training, LHWs assigned to Advance decision were asked to select their allocations by using the page on their smartphone application. Assistance was available from training staff for those who required it. LHWs assigned to the Immediate condition were told they would select their allocations using their smartphone application on Monday morning before beginning work. A hotline number was provided if assistance was required for those in the Immediate condition. The training activities on December 5, for the December 8-9 drive were identical. However, because LHWs had previously been trained on the smartphone application, this portion of the training was conducted as a refresher.

### 2.4.2 Experimental Timeline

*Drive 0, Failed Drive, September 26-30, 2014:* We had hoped to begin our study on Friday, September 26th, 2014 with a training session. 336 LHWs had been recruited, were randomized into treatments, and trained. Advance allocation decisions were collected from half of the

subjects on Friday, September 26th. On Monday, September 29th, when we attempted to collect immediate allocation decisions, there was a disruption in the mobile network that prevented 82 of 168 Immediate decision LHWs from submitting their allocations. This caused us to abandon this drive for the purposes of measuring preferences for subsequent tailoring of contracts. The drive, however, was completed and intertemporal bonuses were paid. For the 82 individuals who did not make their allocations, we contacted them, allowed them to continue working, and paid bonuses for all.<sup>18</sup> We present data from Drive 0 in our robustness exercises, but do not use Drive 0 for the purposes of tailoring contracts. In addition to the 336 LHWs who experienced intertemporal bonus contracts, Drive 0 also had a separate a sample of 85 LHWs who received a phone but no bonus contract to study the differential effect of our incentives.

*Drive 1, November 7-11, 2014:* Of the original 336 bonus contract LHWs in our failed drive, 57 did not participate in the next drive organized for November 7 - 11. We recruited replacements with the help of the Department of Health, identifying a total of 349 LHWs to participate in the intertemporal bonus program. The entire sample was re-randomized into  $R$  and allocation timing conditions. We again included in our study a sample of 72 vaccinators in Drive 1 who received a phone but no financial bonus. Training was conducted on November 7, and Advance allocation decisions were collected. The drive began on November 10, and Immediate allocation decisions were collected. 174 LHWs were assigned to the Advance Choice condition and 175 were assigned to the Immediate Choice condition. Bonuses were paid on November 12. While all 174 LHWs in the Advance Choice condition provided an allocation decision, only 164 of 175 in the Immediate Choice condition provided an allocation, generating a usable sample of 338 allocations in Drive 1. For 232 LHWs, we have allocation decisions in both the failed drive, Drive 0, and Drive 1, forming a potentially valuable panel of response. In addition

---

<sup>18</sup>Appendix Table A.5 checks for balance by failure of the smartphone application in Drive 0. Only one of the eight comparison of means hypothesis tests reject equality at the 10 percent level.

to the 338 LHWs who experienced intertemporal bonus contracts and provided allocations, Drive 1 also had a separate a sample of 73 LHWs who received a phone but no bonus contract to study the differential effect of our incentives.

*Drive 2, December 5-9, 2014:* Of the 338 LHWs who participated in Drive 1 and provided an allocation, 337 again participated in Drive 2. These LHWs were randomly assigned to be structurally tailored or receive a random price in their Drive 2 bonus contracts. Importantly, LHWs retained their Advance or Immediate assignment, such that Drive 2 delivers a 2x2 design for structural tailoring and allocation timing. This allows us to investigate the effect of structural tailoring in general, and if the effects depend on whether present bias may be active.

### 2.4.3 Sample Details

Table 1 summarizes our sample of LHWs from Drive 1 and provides tests of experimental balance on observables. Columns (1) through (4) present the mean and standard deviation of a number of demographic characteristics for LHWs assigned to intertemporal bonus contracts in Drive 1. These values are separated by Advance and Immediate Choice and whether the LHW was subsequently assigned to be Tailored or Untailored in Drive 2. Column (5) presents the means and standard deviations for LHWs that participated in our study (including Drive 0), but did not receive incentives.<sup>19</sup> Column (6) presents a  $p$ -value corresponding to joint tests of equality. Our sample is almost exclusively female, more than 90 percent Punjabi in all treatment arms, and broadly without access to formal savings accounts. LHWs are generally highly experienced with an average of 10.5 years of health work experience and 10.4 years of polio work experience. Consistent with randomization, of the 8 tests performed on demographic balance, only the test performed on an indicator variable equal to one for female subjects suggests baseline imbalance.

---

<sup>19</sup>13 LHWs were in the Phone Only group in Drive 0 and re-randomized into the incentive group in Drive 1. They are included in the Table 1, column (5).

Table 1: Drive 1 Summary Statistics and Covariate Balance

	Intertemporal Bonus Contracts				Phone Only (5)	p-value (6)
	Advance Choice		Immediate Choice			
	Tailored (1)	Untailored (2)	Tailored (3)	Untailored (4)		
<i>Demographics</i>						
Gender (Female = 1)	0.989 (0.011)	1.000 (0.000)	0.976 (0.017)	0.978 (0.015)	0.958 (0.024)	0.089
Years of Education	10.713 (0.300)	10.459 (0.214)	10.232 (0.200)	10.207 (0.245)	10.528 (0.290)	0.613
Number of Children	3.593 (0.212)	3.229 (0.202)	3.438 (0.191)	3.451 (0.195)	3.470 (0.198)	0.803
Punjabi (=1)	0.954 (0.023)	0.952 (0.023)	0.938 (0.027)	0.944 (0.025)	0.957 (0.024)	0.984
<i>Financial Background</i>						
Has a Savings Account (=1)	0.247 (0.047)	0.321 (0.051)	0.268 (0.049)	0.242 (0.045)	0.243 (0.052)	0.773
Participated in a ROSCA (=1)	0.388 (0.053)	0.361 (0.053)	0.400 (0.055)	0.391 (0.051)	0.429 (0.060)	0.947
<i>Health Work Experience</i>						
Years in Health Department	10.826 (0.535)	10.648 (0.509)	10.256 (0.556)	10.319 (0.555)	11.197 (0.546)	0.742
Years as Polio Vaccinator	11.041 (0.504)	10.506 (0.504)	10.323 (0.529)	9.647 (0.500)	10.831 (0.530)	0.340
<i>Vaccination Data</i>						
Drive 1 Behavior: $\frac{R \cdot v_{1,i}}{v_{2,i}}$	1.039 (0.024)	1.043 (0.029)	1.045 (0.036)	1.024 (0.036)		0.975
Proportion: $0.75 \leq \frac{R \cdot v_{1,i}}{v_{2,i}} \leq 1.5$	0.876 (0.035)	0.859 (0.038)	0.783 (0.045)	0.706 (0.047)		0.020
Registered Vaccination Activity	0.809 (0.042)	0.824 (0.042)	0.916 (0.031)	0.815 (0.041)	0.833 (0.044)	0.148
Proportion of Targets Completed	0.797 (0.040)	0.811 (0.038)	0.682 (0.047)	0.731 (0.042)		0.119
Total Vaccinations Attempted	254.333 (13.506)	261.143 (14.032)	212.513 (15.337)	233.587 (14.810)	203.367 (15.847)	0.021
# Vaccinators	89	85	83	92	72	

*Notes:* Mean and standard error in parentheses from Drive 1 observations in columns (1) through (5). Statistical significance level of  $F$ -test for constant means in column (6). Some calculations use smaller sample than that reported in final row due to missing information. The proportion of subjects with missing demographic information for is never greater than 3.5 percent (8 Vaccinators did not report whether they had participated in a ROSCA). A ROSCA is an informal Rotating Savings and Credit Association. Some calculations used a smaller sample size due to missing information.

Table 1 also provides some additional information on vaccination activity in Drive 1 for each group. Average allocations are roughly equal in Advance Choice and Immediate Choice, but the Immediate condition exhibits greater variation in behavior. Correspondingly slightly fewer LHWs in the Immediate condition exhibit behavior within the bounds  $0.75 \leq \frac{R \cdot v_{1,i}}{v_{2,i}} \leq 1.5$ . Of 338 LHWs in our Drive 1 incentive conditions, 288 registered activity in their cell phone application during the drive, while 50 generated no data.<sup>20</sup> Table 1 demonstrates slightly more

<sup>20</sup>Appendix Table A.9 suggests that this data is unrelated to experimental variation in  $R$ , but that LHWs in

registered activity in the Immediate Choice condition. This may be due to the fact that LHWs in Immediate Choice made their allocations in the application on the morning the drive began and so could be more likely to remember to register their activity. Conditional on registering vaccination activity, LHWs in our incentive conditions completed around 75% of their target vaccination attempts, with somewhat greater completion rates in the Advance condition.<sup>21</sup> As noted in section 2.2.2, we account for forecasted non-completion in our empirical approach to estimating the distributions of discount factors and in subsection 4.1 we explore the determinants of target completion. Conditional on registering vaccination activity, LHWs in our incentive conditions attempted an average of around 235 vaccinations in Drive 1. LHWs assigned to the Phone Only condition without incentives completed only around 203 vaccinations. In subsection 4.4 we further explore the separate question of differential effects of incentives relative to monitoring alone for increasing vaccinations.

### 3 Results

Our project has two phases. The first phase measures intertemporal preferences. The second phase evaluates the effects of structural, tailored contracts. These results of these two phases are presented in the next two subsections. A third subsection provides additional analyses and robustness tests.

---

the Advance condition are slightly more likely to have not registered activity.

<sup>21</sup>Completion rates are calculated as  $1/2(\min(\text{Completed}_{1,i}/v_{1,i}, 1) + \min(\text{Completed}_{2,i}/v_{2,i}, 1))$ . Appendix Figure A.6 presents the histogram of average completion percentages across subjects, showing a bimodal distribution of success and failure. In Appendix Table A.9, we examine the determinants of completion with linear probability models and an indicator for completion,  $\text{Complete}(= 1)$ , as dependent variable. Within the Advance and Immediate groups we find no discernible relationship between  $R$  and completion. However, individuals assigned to Immediate choice are significantly less likely to satisfactorily complete their allocated vaccinations. This evidence may have a present-biased interpretation. If subjects in the Immediate condition postpone more work due to present bias, they may be subsequently unable to satisfactorily complete said work.

### 3.1 Drive 1: Individual Behavior and Distributions of Intertemporal Preferences

Each LHW in Drive 1 provides a single allocation,  $(v_{1,i}, v_{2,i})$ , at their randomly assigned value of  $R$  and Advance versus Immediate Choice assignment. To provide a comparable value of these allocations across conditions, Figure 2, Panel A plots the distribution of  $\frac{Rv_{1,i}}{v_{2,i}}$  separately for the 174 LHWs in the Advance condition and 164 LHWs in the Immediate condition. Figure 2, Panel A also provides indicators for how our Drive 2 Tailoring Sample is constructed. In total 57 LHWs with values of  $\frac{Rv_{1,i}}{v_{2,i}}$  outside of the region  $[0.75, 1.5]$  are excluded from the Tailoring Sample, while the remaining 281 are the focus of our Drive 2 analysis.

Absent shocks to marginal costs and probabilistic non-completion, the quantity  $\frac{Rv_{1,i}}{v_{2,i}}$  would be equivalent to each LHW's one period discount factor. Two features of Figure 2, Panel A warrant attention. First, the median value of  $\frac{Rv_{1,i}}{v_{2,i}}$  in the Advance condition is 1.02, while the median in the immediate condition is 1. LHWs allocate slightly fewer vaccination attempts to day 1 of the drive when making immediate allocations, but this difference in medians is not statistically significant, Pearson's  $\chi^2 = 1.18$ , ( $p = 0.28$ ). Second, LHWs in both conditions exhibit substantial variation in allocation behavior. The 25th-75th percentiles of  $\frac{Rv_{1,i}}{v_{2,i}}$  range from  $[0.88 - 1.18]$  in the Advance condition, and from  $[0.84 - 1.21]$  in the Immediate condition. While this wide variation is suggestive of heterogeneous time preferences, preferences are unlikely to be the sole source of heterogeneity. Given that in the absence of shocks  $\frac{Rv_{1,i}}{v_{2,i}}$  identifies an individual discount factor, this would imply daily discount factors well beyond empirical rates of interest. A plausible additional driver of these individual differences is shocks to the marginal costs of vaccination.

The structural exercise developed in section 2.2.2 links the distributions of behavior presented in Figure 2, Panel A to distributions of time preferences accounting for both shocks and probabilistic completion. Table 2 provides the implementations of this methodology for subjects in Drive 1. Columns (1) through (3) of Table 2 presents the distributional discount-



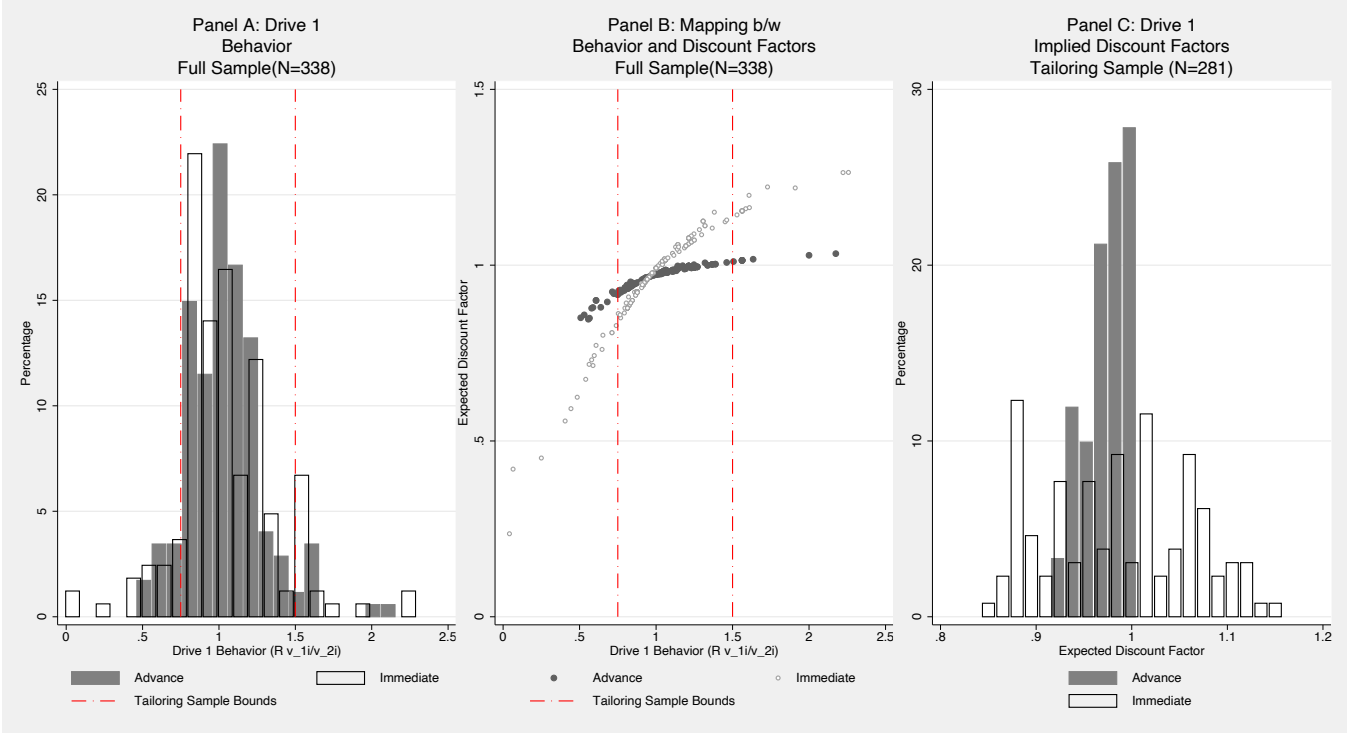


Figure 2: Individual Behavior and Expected Discount Factors

Notes: Panel A reports individual allocation behavior,  $\frac{Rv_{1,i}}{v_{2,i}}$ , vaccinators in Advance and Immediate conditions for Full Sample of 338 vaccinators. Dashed red lines correspond to the boundaries  $\frac{Rv_{1,i}}{v_{2,i}} = 0.75$  and  $\frac{Rv_{1,i}}{v_{2,i}} = 1.5$  which define the Tailoring Sample. Panel B presents the mapping between individual behavior and expected discount factors in Advance and Immediate conditions based on the structural models estimated in columns (3) and (6) of Table 2. Dashed red lines correspond to the boundaries  $\frac{Rv_{1,i}}{v_{2,i}} = 0.75$  and  $\frac{Rv_{1,i}}{v_{2,i}} = 1.5$  which define the Tailoring Sample. Panel C presents the expected discount factors in Advance and Immediate condition specifically for the Tailoring Sample of 281 vaccinators in Drive 1.

ing estimates for LHWs in the Advance Choice condition,  $N(\mu_\delta, \sigma_\delta^2)$ . In column (1), we set  $\ln(\sigma_\delta) = -10$  such that there is effectively zero heterogeneity in preferences, and estimate a homogeneous value  $\mu_\delta = 0.954$  (0.016) and a substantial role for shocks to relative marginal costs,  $\ln(\sigma_\epsilon) = 4.155$  (0.054). In column (2), we permit heterogeneity in discount factors and estimate  $\mu_\delta = 0.965$  (0.022) and  $\ln(\sigma_\delta) = -2.449$  (0.587), implying a standard deviation of  $\delta$  of 0.09. This specification identifies a slightly smaller role for shocks to relative marginal costs. However, heterogeneous preferences provides only a minor improvement in the estimated likelihood relative to the homogeneous model, suggesting that many LHWs have values of  $\delta_i$  close to the aggregate value. Columns (4) and (5) of Table 2 repeat this analysis for all subjects in the Immediate Choice condition. Absent heterogeneity we obtain  $\mu_{\beta\delta} = 0.895$  (0.024) and  $\sigma_\epsilon = 4.529$  (0.055). Incorporating heterogeneity, we estimate  $\mu_{\beta\delta} = 0.959$  (0.029) and  $\sigma_{\beta\delta} = -1.526$  (0.164), implying a standard deviation of  $\beta_i\delta_i$  of 0.22, more than twice as large as that of  $\delta_i$  alone. In the Immediate Choice condition, heterogeneous preferences provides a substantial improvement in the estimated likelihood relative to the homogeneous model.

The differential discounting estimates between the Advance and Immediate conditions are suggestive of limited present bias in discounting:  $\mu_\delta/\mu_{\beta\delta} = 0.938$  and 0.992 without and with accounting for heterogeneous preferences, respectively. Moreover, the estimated distribution of preferences in the Immediate condition suggests a greater degree of heterogeneity, perhaps reflecting that it incorporates two dimensions of individual difference, both  $\delta$  and  $\beta$ . Table 2 identifies the distributions of discount factors using only the between-subject variation in  $R$  and Advance vs. Immediate choice of Drive 1. In subsection 4.3, we reconduct the analysis making using of our failed Drive 0 data, and the corresponding within-subject variation of LHWs transitioning between conditions across drives. These estimates show a quantitatively larger degree of present bias,  $\mu_\delta/\mu_{\beta\delta} \approx 0.90$  and continue to demonstrate substantial heterogeneity in discounting.

Of the 338 LHWs in Drive 1, 288 registered activity in their cell phone application during the

Table 2: Distributional Estimates

	Advance Choice			Immediate Choice		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Discounting Parameters:</i>						
$\mu_\delta$	0.954 (0.016)	0.966 (0.022)	0.966 (0.022)			
$\ln(\sigma_\delta)$	-10 (-)	-2.449 (0.587)	-2.450 (0.582)			
$\mu_{\beta\delta}$				0.895 (0.024)	0.959 (0.029)	0.959 (0.029)
$\ln(\sigma_{\beta\delta})$				-10 (-)	-1.526 (0.164)	-1.527 (0.163)
<i>Completion Parameter:</i>						
$\alpha$			0.001 (0.000)			0.003 (0.000)
<i>Shock Parameter:</i>						
$\ln(\sigma_\epsilon)$	4.155 (0.054)	4.067 (0.120)	4.067 (0.118)	4.529 (0.055)	4.146 (0.148)	4.146 (0.147)
# Vaccinators Allocation Obs.	174	174	174	164	164	164
# Cell Phone Completion Obs.	-	-	142	-	-	146
Log-Likelihood	-969.78	-969.40	-1057.98	-975.41	-967.05	-1068.03

*Notes:* Parameters from maximum simulated likelihood estimation for Drive 1 allocations and completion data where noted. Allocation data provided by Full Sample of 338 vaccinators. Estimates in column (1) assume  $\ln(\sigma_\delta) = -10$  such that there is zero heterogeneity in preferences. Completion data provided by 288 vaccinators. Standard errors in parentheses.

drive. The cellular network in Lahore is known to have some coverage gaps, and so we consider a subject to have successfully completed their work if they registered activity and completed an average of 90% or more of their required tasks on each day of the drive. By this measure, one-hundred seventy-four LHWs successfully completed their Drive 1 allocations. Columns (3) and (6) of Table 2 conduct estimates accounting for probabilistic completion using the completion data from the 288 subjects with registered cell-phone activity. Missing completion data is given a likelihood contribution of zero. Accounting for probabilistic completion has little influence on the estimated discounting distributions in both the Advanced and Immediate conditions, and we estimate a completion parameter of  $\alpha = 0.001 - 0.003$ . In Appendix A.1, we provide additional analyses related to probabilistic completion, demonstrating the robustness of these results to altered assumptions on the utility impact of forecasted non-completion and different sample pools.

The distributional estimates of Table 2, columns (3) and (6) facilitate the simulation exercise described in section 2.2.3 for calculating each LHW's expected discount factor:  $\delta_i$  for LHWs in the Advance condition, and  $\beta_i \delta_i$  for LHWs in the Immediate condition. For each value of  $R$  we simulate 1 million allocations under the assumed distributions for preferences and shocks accounting for probabilistic completion estimated in Table 2, columns (3) and (6). For each LHW, we assign a value of  $\delta_i$  or  $\beta_i \delta_i$  as the expected value of corresponding simulated discount factors that yield allocations within 2.5 vaccinations of the LHWs allocation,  $(v_{1,i}, v_{2,i})$ . That is, we calculate the posteriors  $E[\delta_i | v_{1,i}, v_{2,i}, R]$  and  $E[\beta_i \delta_i | v_{1,i}, v_{2,i}, R]$  using the simulated distributions of behavior. Figure 2, Panel B provides the mapping in Advance and Immediate conditions from behavior,  $\frac{Rv_{1,i}}{v_{2,i}}$ , to individual expected discount factor calculations. In both conditions, the approach deployed substantially shrinks the degree of individual differences. Extreme allocations are estimated to be due to shocks, not preferences, and so are mapped back to more reasonable expected discount factors. Given the greater heterogeneity in estimated preferences in Immediate Choice, the mapping for  $E[\beta_i \delta_i | v_{1,i}, v_{2,i}, R]$  permits wider variation in

calculated expected discount factors. Panel B also provides the indicators for how our Drive 2 Tailoring Sample is constructed, eliminating those 57 LHWs with values of  $\frac{Rv_{1,i}}{v_{2,i}}$  outside of the region  $[0.75, 1.5]$  and discounting estimates outside of the corresponding bounds for  $E[\delta_i|v_{1,i}, v_{2,i}, R]$  and  $E[\beta_i\delta_i|v_{1,i}, v_{2,i}, R]$ .

Figure 2, Panel C focuses on the 281 LHWs in the Tailoring Sample and provides the the obtained distributions of  $E[\delta_i|v_{1,i}, v_{2,i}, R]$  and  $E[\beta_i\delta_i|v_{1,i}, v_{2,i}, R]$ . The median value of  $E[\delta_i|v_{1,i}, v_{2,i}, R]$  in the Advance condition is 0.976, while the median value of  $E[\beta_i\delta_i|v_{1,i}, v_{2,i}, R]$  in the Immediate condition is 0.991. The 25th-75th percentiles of expected discount factors range from  $[0.960 - 0.991]$  in the Advance condition, and from  $[0.925 - 1.055]$  in the Immediate condition, echoing the substantial heterogeneity in reduced form behavior. The LHWs in the Tailoring Sample are focus of our evaluation of structural tailored contracts in the next section.

### 3.2 Drive 2: Evaluating Structural, Tailored Contracts

Individual expected discount factors from Drive 1 in hand, we evaluate contracts tailored to individual discounting parameters. Of the 338 LHWs who participated in Drive 1, 281 provided allocations that were within our Tailoring Sample,  $0.75 \leq \frac{Rv_{1,i}}{v_{2,i}} \leq 1.5$ , and 280 participated in Drive 2.<sup>22</sup> Within this Tailoring Sample, 142 LHWs were assigned a value of  $R_i^* = \frac{Rv_{1,i}}{v_{2,i}}$ , which should induce approximately equal allocation of vaccinations through time,  $w_{1,i} = w_{2,i}$ . The remaining 138 LHWs provide the basis for the different comparison policies that we consider, and were assigned a uniform random price  $\tilde{R}_i \in U[0.75, 1.5]$ .<sup>23</sup>

Structurally tailored LHWs who received  $R_i^* = \frac{Rv_{1,i}}{v_{2,i}}$  have an average values of  $|\frac{w_{1,i}}{w_{2,i}} - 1| = 0.14$  (0.02), while those who received  $\tilde{R}_i \in U[0.75, 1.5]$  have an average value of 0.61 (0.31),

<sup>22</sup>LHWs with allocations outside of these bounds were allowed to participate in Drive 2 and were either assigned  $\tilde{R}_i \in U[0.75, 1.5]$  if they were in the untailored control group (31 subjects) or assigned  $R_i = 0.75$  or  $R_i = 1.5$  if they were in the tailored group and had  $R_i^* < 0.75$  (15 subjects) or  $R_i^* > 1.5$  (11 subjects). See subsection 4.2.2 for analysis of these observations.

<sup>23</sup>As noted in section 2.3, assignment to the tailored or the untailored group was conducted via stratified randomization with strata based upon the tercile of differences from equal provision of effort in Drive 1.

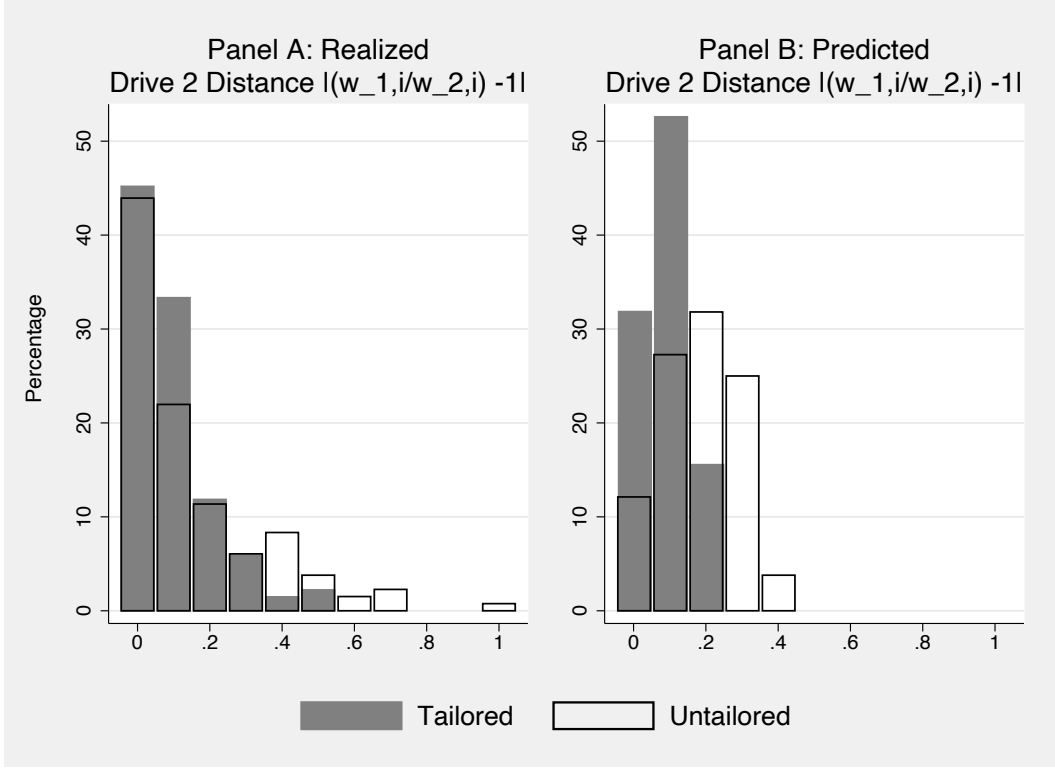


Figure 3: Realized and Predicted Distances

*Notes:* Panel A presents the realized distance to the policy target,  $|\frac{w_{1,i}}{w_{2,i}} - 1|$  in Drive 2 for 267 Tailored and Untailored vaccinators in the Tailoring Sample. Thirteen of 280 Tailoring Sample vaccinators with extreme Drive 2 allocations are excluded. Panel B presents the predicted distances for these same vaccinators in the Tailored and Untailored conditions. Predictions generated by simulating allocations for each LHW associated with their Drive 2 assigned interest rate and new shock realizations.

$t_{278} = 1.53$ , ( $p = 0.13$ ). The presence of a few extreme outliers skews the central measures of the data somewhat. Trimming the top and bottom 1% of the sample of Drive 2 allocations, mean distance for the tailored group is 0.10 (0.01), while the mean distance for the untailored control group is 0.15 (0.02),  $t_{265} = 3.07$ , ( $p = 0.002$ ).<sup>24</sup> We focus our analysis on this trimmed sample, but provide results corresponding to the full sample in Appendix Table A.6 and also report results using the same specifications but winsorizing the sample at the top and bottom percentile in Appendix Table A.7. Figure 3, Panel A presents the distributions of distances from equal allocation,  $|\frac{w_{1,i}}{w_{2,i}} - 1|$  for the trimmed sample. Tailored LHWs provide systematically lower distance measures than untailored LHWs.

<sup>24</sup>This trimming eliminates 13 total observations, 7 from tailored group, and 6 from the untailored group.

Table 3 provides corresponding regression. We find that tailoring reduces distance to equal provision by around 5%-age points. Recall that LHWs assigned to Advance Choice in Drive 1 remain in Advance Choice in Drive 2, while those assigned to Immediate Choice remain in Immediate Choice. In the even columns of Table 3, we examine differential effects across these two groups. One might expect larger distance measures in Immediate conditions (and hence greater benefits to structural tailoring) if LHWs were present-biased. This is what is observed. Immediate Choice is associated with significantly larger distance measures and structural tailoring in Immediate Choice significantly reduces these distances. Structural tailoring in Immediate Choice reduces distance from equal provision by around one-half. Note that the effect size is similar to the effect of moving a comparison LHW from Advance to Immediate choice. That is, structural tailoring effectively eliminates present bias in allocations.

Figure 3, Panel B presents the corresponding distributions of expected allocations in the tailored treatment group and in the untailored controls. To calculate these distributions we simulate new allocations of  $w_{1,i}$  and  $w_{2,i}$  for each LHW at their Drive 2 assigned value of  $\tilde{R}_i$  or  $R_i^*$  and new potential shock realizations. Specifically, we take the Drive 1 window of simulated allocations used to calculate each LHW's expected discount factor, and for every simulant in the window we re-simulate behavior with the new interest rate and a new shock realization. We then construct the average simulated allocation and the corresponding simulated average distance to equal allocations,  $|\frac{\widehat{w_{1,i}}}{w_{2,i}} - 1|$ . The distance to equal allocations should, indeed, be substantially higher for those individuals receiving a random value of  $\tilde{R}_i \in U[0.75, 1.5]$  compared to those with structural tailored values,  $R_i^* = \frac{Rv_{1,i}}{v_{2,i}}$

### 3.2.1 Alternative Counterfactuals for Structural Tailored Contracts

In Drive 2 of our study, the group receiving a uniform random interest rate,  $\tilde{R}_i$ , provides the counterfactual for evaluating the effect of our intervention. While this is one reasonable counterfactual to consider, a stricter evaluation of structural tailored interventions is possible.

We can draw subsets from our random control group to approximate counterfactuals that are atheoretic, broad, or both. Comparison of our treatment group to these various controls allows us to evaluate the source of any potential benefits of the structural, tailored policy relative to alternative policy approaches. Table 3 provides additional analysis considering the following three policy alternatives:

1. *Structural, Broad Policy:* This corresponds to a case where the policymaker knows workers' discount factors, but can offer only one interest rate to all workers. Absent shocks and probabilistic noncompletion, the quantity  $\frac{Rv_{1,i}}{v_{2,i}}$  identifies individual discount factors. In the Tailoring Sample, the average value of  $\frac{Rv_{1,i}}{v_{2,i}}$  for Drive 1 is 1.041 for the Advanced condition and 1.019 for the Immediate condition. These correspond to the broad values of  $\tilde{R}$  required to approximately equalize allocations in a similar manner as our structural tailored policy, but without each individual getting a unique interest rate. To approximate this broad, structural policy we select the 54 individuals from the condition who received  $\tilde{R}_i \in U[0.75, 1.5]$  within one standard deviation of these prices. Appendix Figure A.8, Panel A indicates the exact assignments for this subgroup. The relative prices implied by the aggregate model are structural, informed by preferences in Drive 1 in the same manner as our treatment group, but not tailored to each individual. Table 3, Columns (3) and (4) repeat the analysis of columns (1) and (2), but using this alternative control group. We find that our structural, tailored policy induces allocations that are directionally closer to equality than this structural, broad alternative, but the effect is not statistically significant.

2. *Atheoretic, Broad Policy:*

To develop an atheoretic broad policy, we first estimate the reduced-form relationship between Drive 1 allocations and the experimentally varied parameters  $R$  and  $d = 0$  or  $d = 1$ . Appendix Table A.8 provides regression coefficients indicating the sensitivity of the



allocation  $v_1$  to  $R$  and whether the decision is immediate. The reduced-form relationship estimated for the Tailoring Sample in Drive 1 is  $v_1 = 216.33 - 3.00 \times \mathbf{1}_{d=1} - 66.67 \times R$ . In order to equate  $w_1 = w_2$  under the constraint  $w_1 + R w_2 = 300$ , one requires  $(1 + R)w_1 - 300 = 0$ . Substituting in for the reduced-form relationship, one obtains  $f(R) = (1 + R)(216.33 - 3.00 \times \mathbf{1}_{d=1} - 66.67 \times R) - 300 = 0$ . Note that  $f(R)$  is quadratic in  $R$ . In Advance Choice, it obtains the value of zero at  $R = 1.05$  and  $R = 1.19$ . In Immediate Choice,  $f(R)$  does not achieve the value zero, but has a maximum value of  $f(R) = -6.01$  at  $R = 1.10$ . To approximate this broad atheoretic policy, we select the 54 individuals from the subsample who received  $\tilde{R}_i \in U[0.75, 1.5]$  within one standard deviation of  $R = 1.1$ . Appendix Figure A.8, Panel B indicates the exact assignments for this subgroup. The relative prices implied in this case are atheoretic, informed only by the estimated reduced-form sensitivities, and not tailored to each individual. Table 3, Columns (5) and (6) repeat the analysis of columns (1) and (2), but using this alternative control group. We find that our structural, tailored policy induces allocations closer to equality than this atheoretic broad alternative.

3. *Atheoretic, Tailored Policy:* From the subsample who received  $\tilde{R}_i \in U[0.75, 1.5]$ , random assignment generates a match between the random price received,  $\tilde{R}_i$ , and Drive 1 allocation behavior. Even without structural guidance on the correct value of  $R$  to achieve equal allocations, random assignment will at times assign higher values of  $\tilde{R}_i$  to individuals with higher values of  $\frac{Rv_{1,i}}{v_{2,i}}$ . These assignments give higher prices to more patient LHWs, but do not require that the relationship between prices and patience be linear as in the structural, tailored policy, which gives  $R_i^* = \frac{Rv_{1,i}}{v_{2,i}}$ . For each LHW who received a random price in Drive 2, we count the percent of LHWs who were more patient in Drive 1 but received a lower value of  $\tilde{R}_i$ . From the subsample who received  $\tilde{R}_i \in U[0.75, 1.5]$ , we select the 49 LHWs for whom this number is less than or equal to 10%, being effectively in order with at least 90% of the sample. Appendix Figure A.8, Panel C indicates the exact

assignments for this subgroup. The relative prices implied in this case are atheoretic — loosely related to patience, but not designed to achieve a specific objective beyond giving more patient LHWs higher prices — and tailored to each individual. Table 3, Columns (7) and (8) repeat the analysis of columns (1) and (2), but using this alternative control group. We find that our structural, tailored policy induces allocations closer to equality than this atheoretic, tailored alternative.

Table 3: The Effect of Tailoring Intertemporal Incentives

Dependent variable: Policy Comparison Group	$\left  \frac{w_{1,i}}{w_{2,i}} - 1 \right $							
	Random Price		Structural, Broad		Atheoretic, Broad		Atheoretic, Tailored	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Structural Tailored (=1)	-0.049*** (0.018)	-0.014 (0.019)	-0.034 (0.021)	-0.021 (0.024)	-0.051** (0.025)	-0.010 (0.022)	-0.054* (0.029)	-0.014 (0.023)
Immediate Choice		0.117*** (0.035)		0.082* (0.044)		0.137*** (0.051)		0.169*** (0.062)
Structural Tailored x Immediate		-0.084** (0.040)		-0.045 (0.047)		-0.102* (0.055)		-0.132** (0.064)
Constant	0.022 (0.058)	-0.004 (0.057)	0.202*** (0.071)	0.185** (0.073)	0.168** (0.065)	0.143** (0.063)	0.083 (0.066)	0.032 (0.066)
Stratum FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Exclude 99th and 1st Percentiles	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Drive 2 $R_i^*$ or $\tilde{R}_i$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-Squared	0.082	0.154	0.076	0.122	0.114	0.194	0.085	0.193
Mean in Untailored Contract	0.153	0.153	0.129	0.129	0.145	0.145	0.148	0.148
Mean in Untailored Advance		0.098		0.103		0.088		0.091
Mean in Untailored Immediated		0.222		0.177		0.223		0.265
# Vaccinators	267	267	191	191	194	194	184	184
# Comparison Vaccinators	132	132	56	56	59	59	49	49

*Notes:* Table reports the effect of structural, tailored policy relative to alternatives on realized distance to the policy target,  $\left| \frac{w_{1,i}}{w_{2,i}} - 1 \right|$  in Drive 2 for 267 Tailored and Untailored vaccinators in the Tailoring Sample. Thirteen of 280 Tailoring Sample vaccinators with extreme Drive 2 allocations are excluded. Ordinary least squares regressions. Heteroskedasticity robust White standard errors reported in parentheses. \* $p < 0.1$ , \*\* $p < 0.05$ , \*\*\* $p < 0.01$ .

The analysis to this point indicates two key findings. First, heterogeneity in discounting is observed in both Advance and Immediate Choice, with substantially greater heterogeneity when present bias is implicated. This highlights the potential for policy interventions tailored

to individual preferences. Second, structural, tailored contracts work. Those LHWs given a tailored price equal to their previously measured expected discount factor provide smoother service than a set of alternatives that span the policy space of structural vs. atheoretic and broad vs. tailored. In the following section, we explore robustness and provide a set of additional examinations.

## 4 Robustness Tests and Additional Exercises

### 4.1 Evaluating Model Assumptions

As in any structural exercise, a number of assumptions are required to infer discounting parameters from LHW allocation behavior. Four assumptions are relevant for the present discussion, which we discuss below.

*Stationarity of Deterministic Costs:* We assume the deterministic portion of marginal costs is the same for day 1 and day 2 of each drive. If sooner costs are forecasted to be more severe than later costs, LHWs may appear disproportionately impatient, while if later costs are forecasted to be more severe, they may appear disproportionately patient. Further, if perceived costliness of vaccinations changes from Advance to Immediate choice, present bias is conflated with non-stationarity.

Importantly, our monitoring technology provides time-stamps and geo-stamps for vaccination activity. Time stamps are recorded for every vaccination attempt, while geo-stamps are collected approximately every 10 vaccination attempts. This may provide independent means for assessing the costliness of tasks from time use. For each LHW, we identify the median time lapse between vaccination attempts and the median distance covered per 30 minute window each day.<sup>25</sup> Of our 338 LHWs, measures for median time lapse between vaccination attempts

---

<sup>25</sup>We focus only on the distance traveled and time taken for vaccinations between 8 am and 6pm each day.

are available for 277 on either Day 1 or Day 2 and for 228 LHWs on both days of Drive 1.<sup>26</sup> Of our 338 LHWs, measures for median distance traveled every 15 minutes are available for 274 on either Day 1 or Day 2 and for 226 LHWs on both days of Drive 1.<sup>27</sup>

LHWs take around 3.4 minutes between vaccination attempts and walk around 0.06 miles per 15 minutes on Day 1. Focusing on individuals with measures on both days of the drive, we find that time taken and distance traveled are uncorrelated both with Advance choice and with allocation behavior within condition. Time and distance are also uncorrelated with Advance choice and allocation behavior on Day 2 of the drive. Further, differences in time taken or distance walked are statistically indistinguishable from zero, uncorrelated with allocation timing, and uncorrelated with allocation behavior within condition. These data indicate stability in required average effort per vaccination which is unrelated to assignment to Advance or Immediate choice, and that changes in efficacy are unrelated to measured preferences. This suggests that perceived changes in costs likely do not drive our measures of discounting.<sup>28</sup> These results are all presented in Appendix Table A.10.

---

The distribution of time taken and distance traveled carried some extreme outliers for some subjects. As such, we felt the median was an appropriate summary statistic. Though we had expected to receive geo-stamp data approximately every 10 vaccination attempts, when the monitoring data arrived we noted substantial variance in the number of vaccinations with common geo-stamps and sequences of geo-stamps which ‘bounced’ back and forth between geographic coordinates. In order to not overstate subject movements, we opted to take average coordinates within a 15 minute window and calculate direct-line distance between window-average coordinates as our measures of distance.

<sup>26</sup>265 LHWs have Day 1 lapse data while 240 have Day 2 lapse data. Of the 73 LHWs with missing Day 1 data, 68 completed either zero or one vaccination on Day 1 such that time lapse between vaccination attempts is not calculable. The remaining 5 conducted vaccinations but did not have phones that interacted with the server to report time use. Of the 98 LHWs with missing Day 2 data, 92 of them completed either zero or one vaccination on Day 2 and the remaining 6 did not have phones that interacted with the server to report time use. Those LHWs who completed vaccinations but did not have interaction with the server had their vaccination records pulled manually from their phones after the drive.

<sup>27</sup>257 LHWs have Day 1 distance data while 240 have Day 2 distance data. Of the 81 LHWs with missing Day 1 data, 75 completed four or fewer vaccination attempts on Day 1 such that distance traveled between 15 minute windows is not calculable. The remaining 6 conducted vaccinations but either did not have phones that interacted with the server to report location or had faulty Global Position Systems (GPS) in their phones. Of the 98 LHWs with missing Day 2 data, 96 of them completed four or fewer vaccination attempts on Day 2 and the remaining 2 did not have phones that interacted with the server to report location or had faulty GPS.

<sup>28</sup>Ultimately, such stationarity is likely to be expected given that LHWs are already well-versed in vaccination procedures, have an average of 10.5 years of experience as LHWs, and received a half day’s training on the vaccination monitoring application.

*Identical Cost Functions:* Our calculation of expected discount factors assumes identical quadratic cost functions with additive cost shocks. Though these assumptions allow for straightforward estimation and calculation of time preferences, any violation would lead us to confound differences in patience across individuals or across Advance/Immediate timing with differences in costs. One natural view would be to assume that individuals do not discount at all,  $\delta = 1$  and  $\beta = 1$  and allocations are deterministic, such that allocations identify only the shape of the cost function. In this case, when  $R = 1$ , all LHWs, regardless of allocation timing, should exhibit  $v_{1,i} = v_{2,i} = 150$  for all values of  $\gamma$ .<sup>29</sup> Examining the Drive 0 and Drive 1 data, we find that for 163 LHWs who were assigned  $R = 1$ , the mean allocation is  $v_1 = 140.84$  (*s.d.* = 24.76).<sup>30</sup> Though the median allocation is indeed 150, responses range widely with 5th-95th percentiles of response being 103 to 160. If heterogeneity in costs alone were driving response, and discounting and shocks were not key features of the data, one would not expect to see this extent of variation in response when  $R = 1$ . Further, given random assignment to allocation timing, heterogeneity in costs alone does not easily rationalize the difference in discounting distributions between Advance and Immediate conditions.

*Only Failure, No Shirking:* Our structural exercise assumes individuals know their likelihood to succeed and should work only some minimal amount if their target is not attainable. Appendix Figure A.6 demonstrates the plausibility of this assumption with a bimodal pattern of almost complete success and almost complete failure. Another possibility is that subjects find an alternate way to renege on their contracts by shirking and still receiving pay. Not all vaccination attempts are equally challenging. In Appendix Figure A.7 we plot for each half-hour of Drive 1 the total number of attempted vaccinations along with the probability of successful vaccination and the probability that no child was reported as present. Reporting

---

<sup>29</sup>This is because the Euler equation reduces to  $(\frac{v_{1,i}}{v_{2,i}})^\gamma = R = 1$ , which implies  $\frac{v_{1,i}}{v_{2,i}} = 1$ .

<sup>30</sup>42 of 163 LHWs allocated exactly  $v_{1,i} = v_{2,i} = 150$ .

that no child was present is likely to be less time consuming than a successful vaccination and easier to falsify. The vast majority of vaccination activity occurs before 3:00pm, there exists no sharp uptick in activity as days end, and we find evidence that LHWs' proportion of successful or failed vaccination attempts remains largely steady throughout the workday. This suggests that allocated vaccination attempts are conducted with due diligence.

*No Biases in Choice:* Our study assumes that the allocation environment itself induces no biases in choice such that LHW allocations are directly informative of preferences. A substantial literature in experimental economics suggests that aspects of the decision environment may deeply influence measures of preferences (for recent examples, see Harrison, Lau, Rutstrom and Sullivan, 2005; Beauchamp, Benjamin, Chabris and Laibson, 2015). One common view is that subjects are biased towards the middle of a choice set. In our environment, this could involve subjects opting for either equal allocations of  $v_{1,i} = v_{2,i}$ , or choosing an allocation in the middle of their budget constraint,  $v_{1,i} = Rv_{2,i}$ . Only 31 of 338 LHWs (9%) exhibit  $v_{1,i} = v_{2,i}$  in Drive 1. Taking a less conservative measure of  $v_{2,i} - 2.5 \leq v_{1,i} \leq v_{2,i} + 2.5$ , we find that still only 58 of 338 LHWs (17%) are within 5 vaccinations of  $v_{1,i} = v_{2,i}$ .<sup>31</sup> Only 35 of 338 LHWs (10.3%) exhibit  $v_{1,i} = Rv_{2,i}$ . Taking a less conservative measure of  $Rv_{2,i} - 2.5 \leq v_{1,i} \leq Rv_{2,i} + 2.5$ , we find that 83 of 338 LHWs (25%) are within 5 vaccinations of  $v_{1,i} = Rv_{2,i}$ .<sup>32</sup> Taken together, this suggests that biases towards the middle of the budget constraint or towards equal allocation are unlikely to be driving substantial portions of allocation behavior.

## 4.2 Tailoring Robustness Tests

Our Drive 2 data show that LHWs who are given bonus contracts with an interest rate linked to their expected discount factors provide significantly smoother service than a number of

---

<sup>31</sup>As an even less conservative measure, 145 of 338 (43%) satisfy  $v_{2,i} - 10 \leq v_{1,i} \leq v_{2,i} + 10$ .

<sup>32</sup>As an even less conservative measure, 137 of 338 (40.5%) satisfy  $Rv_{2,i} - 10 \leq v_{1,i} \leq Rv_{2,i} + 10$ .

policy comparison groups. Here we examine robustness of this result to alternative measures for smoothness in service provision and alternative measures for tailoring. We also provide an analysis of tailoring by completion.

#### 4.2.1 Alternative Measures for Smooth Provision

Our analysis measures the distance to equal provision using the metric  $|\frac{w_{1,i}}{w_{2,i}} - 1|$ . In Table A.11, we reconduct the analysis of Table 3, using five alternate measures for smoothness. Panel A presents the Euclidean distance to the 45 degree line,  $\frac{|w_{1,i} - w_{2,i}|}{\sqrt{2}}$ . Panel B presents the Euclidean distance normalized by the total number of vaccinations allocated,  $\frac{|w_{1,i} - w_{2,i}|}{\sqrt{2}(w_{1,i} + w_{2,i})}$ . Panel C presents the number of sooner vaccinations that would need to be reallocated to reach the 45 degree line,  $|w_{1,i} - \frac{300}{1+R}|$ . Panel D presents probit regressions for needing to reallocate more than 10 vaccinations,  $|w_{1,i} - \frac{300}{1+R}| > 10$ . And finally, Panel E presents the value,  $\min[w_{1,i}, w_{2,i}]$ . Across all specifications, the main conclusions are reproduced. However, the results with respect to additional structural tailoring benefits in Immediate Choice fall, at times, outside the range of statistical significance. These alternative measures of smooth provision indicate that our results are not an artifact of how one measures the outcome of interest.

#### 4.2.2 Alternative Sample Restrictions and Treatment Measures

Our exercise focuses on LHWs with  $\frac{Rv_{1,i}}{v_{2,i}}$  between 0.75 and 1.5. Of 338 LHWs in Drive 1, 280 satisfied this requirement and participated in Drive 2. Those LHWs who were assigned to be tailored, but exhibited behavior that fell outside of these bounds behavior were given either  $R_i^* = 0.75$  or 1.5 depending on which boundary they crossed. For such individuals, structural tailoring is not a binary treatment, but rather a continuous difference between their expected discount factor and the exogenously given value of  $R_i^* = 0.75$  or 1.5. Indeed, for all LHWs in the untailored group, treatment is also a continuous measure. In columns (1) and (2) of Table A.12, we reconduct the analysis of Table 3, columns (1) and (2) using as the measure of treatment

the absolute difference between each LHW’s value of  $\frac{Rv_{1,i}}{v_{2,i}}$  in Drive 1 and their assigned value of  $R$  in Drive 2, which we label *Structural Tailoring Intensity*. Structural Tailoring Intensity takes value zero for all tailored LHWs in the Tailoring Sample and is a positive number for all others. Hence, larger values of intensity should be associated with greater distances to equal provision. The main results of Table 3 are reproduced with in the Tailoring Sample; the closer the connection between patience and the assigned value of  $R$ , the smoother is provision. In columns (3) and (4) of Table A.12, we include those LHWs in the Boundary Sample. Including these observations does not alter the conclusions; however, it should be noted that treatment is no longer orthogonal to individual preferences as extremely patient and impatient LHWs will receive larger intensity measures on average. Using the indicator for ‘Structurally Tailored’ would not be an appropriate solution to this problem as tailored LHWs with extreme patience or impatience may actually receive relative prices that are further from their policy-optimal values than those in the untailored condition. In columns (5) and (6) of Table A.12, we include both the Boundary Sample and those LHWs with extreme Drive 2 allocations. As in our other analyses including outliers, we find similar directional effects but very wide confidence intervals.

### 4.2.3 Tailoring and Completion

Our analysis of probabilistic completion evaluates completion through the lens of a model and attempts to assess the trade-off between marginal completion probabilities and discounted marginal costs. Though this analysis seems both tractable and yields valuable predictive insights, an alternative interpretation for non-completion exists. If the *outcome* of failure rather than its probability is perfectly forecasted by the LHW, there is no incentive to respond truthfully. As such, the targets set in Drive 1, and our corresponding inference on time preferences, would be systematically inaccurate for individuals expecting to fail. In effect, successful LHWs are allocating according to equation (2), while unsuccessful LHWs are providing only noisy response. Under this assumption, we should be dramatically less able to predict allocation



behavior for LHWs who fail in Drive 1.

Table A.13 repeats the analysis of Table 3, columns (1) and (2) separately for LHWs who completed and failed to complete their Drive 1 targets. Similar magnitude effects are observed for both sets of LHWs, with structural tailoring serving to reduce distance from the equal provision by around one third. Focusing only on the completing subjects, we would reach effectively the same conclusion as our initial analysis.

### 4.3 Repeated Measurement and Estimation of Present Bias

In Drive 1, when relying on between-subject differences in behavior, aggregate presents bias appears limited. Given the wide heterogeneity in behavior in both Advance and Immediate Choice, one may fail to identify present bias due to sampling variation even if it exists. Indeed, most studies of present bias and dynamic inconsistency are conducted as within-subject exercises with more choices, potentially because of such wide heterogeneity.

Fortunately, our failed Drive 0 and the corresponding re-randomization in Drive 1 allows us to identify present bias using both more data and within-subject variation for LHWs who changed from Advance to Immediate choice (or vice versa) across drives. Appendix Table A.3, re-estimates the specifications of Table 2, columns (2) and (5) using both Drive 0 and Drive 1 behavior, drawing from 622 choices made by 390 LHWs in either Drive 0 or Drive 1. We find mean estimates  $\mu_\delta = 0.970$  (0.019) and  $\mu_{\beta\delta} = 0.889$  (0.027), indicating a present bias of  $\beta = 0.916$ . When focusing only within subject variation by estimating on the 126 LHWs who transition from Advance to Immediate choice across drives, we estimate  $\mu_\delta = 0.982$  (0.029) and  $\mu_{\beta\delta} = 0.886$  (0.039), indicating an average present bias of  $\beta = 0.902$ . These values are closely consistent with recent laboratory work eliciting time preferences over effort (Augenblick et al., 2015; Augenblick and Rabin, 2015). Figure 4, plots expected present bias for the panel of 125 of 126 LHWs who transition from Advance to Immediate choice across




Figure 4: Within-Vaccinator Present Bias

*Notes:* Figure presents within-vaccinator measures of present bias for 125 vaccinators who transitioned from Advance to Immediate conditions across Drive 0 and 1. Measure of present bias calculated as the ratio of conditional expected discount factors given behavior in the two drives,  $E[\beta_i \delta_i | \cdot] / E[\delta_i | \cdot]$ . Median value presented as dashed red line.

drives vice versa.<sup>33</sup> We find a median estimate of present bias of  $E[\beta_i \delta_i | \cdot] / E[\delta_i | \cdot] = 0.943$ , and a mean estimate of 0.919 ( $s.d. = 0.143$ ). As in laboratory work identifying present bias within-subject from longitudinal designs (Augenblick et al., 2015; Augenblick and Rabin, 2015), these individual estimates of present bias can be provided only for the selection of LHWs who provided allocations in both drives. Of the 254 LHWs who successfully provided an allocation in Drive 0, 232 provide an allocation in Drive 1. This level of attrition over time compares favorably to that of (Augenblick et al., 2015; Augenblick and Rabin, 2015).

<sup>33</sup>One LHW provided an extreme observation in one of the drives for which our simulation exercise yielded zero simulants in the relevant window of allocations.

## 4.4 Monitoring Without Incentives

We include in our study a sample of 85 LHWs in Drive 0 and 72 LHWs in Drive 1 who received a phone but no financial bonus. Importantly, we can measure the number of vaccination attempts for LHWs who were and were not assigned to received an additional financial incentive. Of the 157 LHWs who received a smartphone monitoring device, but no financial bonus in Drives 0 or 1, 134 register any vaccination information. We combine these 134 observations with 573 observations from Drives 0 and 1 from our incentive treatments who register vaccination information. Table A.4 reports the impact of providing financial incentives and monitoring relative to monitoring alone. Without incentives, LHWs register 182 vaccination attempts over each two day drive. With incentives, this quantity grows to 218 in Drives 0 and 1. This large increase in total vaccinations reflects the impact of the substantial completion bonus provided in our incentive treatments. Additionally Table A.4 examines whether within Drive 2, being assigned to the tailored or untailored group influences total recorded vaccinations. Of the 267 LHWs noted in Table 3, 225 registered vaccination information in Drive 2. For these LHWs tailoring is associated with an insignificant reduction of about 5% in the total number of vaccinations recorded. As our tailoring policy was not designed to increase total vaccinations, but rather to equalize vaccinations over time, such limited effects should be expected. Hence, incentives increase vaccination attempts above monitoring alone, while tailoring for equalization of vaccinations provides no further increase.

## 5 Conclusion

This paper examines the potential for policy interventions to be tailored to individual time preferences. We couch this question in an effort to customize contracts for 337 vaccination workers who spend two days each month attempting to deliver polio vaccines in the neighborhoods of Lahore, Pakistan.

We monitor workers' efforts using a smartphone application developed especially for our project, and elicit preferences using a Convex Time Budget design (Andreoni and Sprenger, 2012; Augenblick et al., 2015). Workers in an Advance condition allocate vaccinations over a two day drive prior to the beginning of the drive, while workers in an Immediate condition state their allocations at the beginning of the first day. Each worker also faces a randomized relative price for converting vaccinations across days. Worker behavior in this drive is used to identify individual time preferences. In a subsequent drive, we tailor contract terms to individual time preferences for half of the workers. This is done by choosing a relative price designed to encourage equal provision of effort over the two days of the drive. The other half of workers is given a random uniform price. We contrast our structural, tailored policy with three alternatives drawn from the random uniform condition that span the policy space in two dimensions: atheoretic vs. structural and broad vs. tailored.

Our findings are encouraging. Those workers who receive structural, tailored contract terms are substantially closer to the policy objective than the alternate policies considered. Using individual discounting parameter estimates to form a new incentive contract does indeed have the predicted effect on allocation behavior. To date, little research makes use of such predictive value of discounting estimates. Our results show not only that estimates are predictive, but also that useful parameter estimates are identifiable from a very limited number of experimental choices. This suggests that the substantial effort of articulating and estimating models in this domain has been well-invested. Policymakers should be encouraged by these findings to consider such tailored interventions. In the domain of intertemporal choice, the specific intervention we consider may be of interest for policymakers wishing to achieve smoothness in allocation behavior or consumption over time.

This paper also speaks to a recent discussion on the external validity of randomized control trials. Developing structural models through which to interpret experimental treatment effects potentially provides a means for generalizing results to other settings (Acemoglu, 2010;

Banerjee, Chassang and Snowberg, 2016).<sup>34</sup> In our setting, translating from our reduced form experimental treatment effects to a structural model of choice requires a set of potentially strong (and implausible) assumptions.<sup>35</sup> Nonetheless, the findings of predictive validity in this case suggests there is indeed potential for using structure as a means of increasing the external validity of results obtained from a single sample.

Separately, our results link to the growing literature on the personnel economics of the state (Ashraf, Bandiera and Lee, 2015; Bertrand, Burgess, Chawla and Xu, 2016; Finan, Olken and Pande, Forthcoming; Dal Bó, Finan and Rossi, 2013; Deseranno, 2016; Callen et al., 2019). Within this literature, there is interest in understanding whether heterogeneity in competencies and in motivation of state actors is linked to meaningful differences in state performance or service provision (Ashraf et al., 2015; Dal Bó et al., 2013; Deseranno, 2016; Callen et al., 2019). We take the additional step of asking not only whether this heterogeneity matters for outcomes, but also whether it can be acknowledged and reflected in the design of individual incentives.

There are a number of clear limitations to our study which should be addressed by future research. First, our study sidesteps the critical issue of incentive compatibility by not informing subjects of the possibility that their initial behavior would potentially be subsequently used to inform their own contract terms. The mechanism design problem of eliciting preferences and tailoring on said preferences with complete information will be critical if one wishes to implement such contracts repeatedly in the field. Second, future research should seek to gain more precise estimates of preferences. Our exercise requires restrictive assumptions that could be relaxed in the presence of more data. If our results point to a lower bound in the promise of structural, tailored contracts, it is important to know how much more can be achieved. Third, alternative policy objectives and contract types should be investigated to ensure robustness of the identified predictive validity. Our findings have natural extensions to piece rate con-

---

<sup>34</sup>Attanasio and Meghir (2012), Duflo, Hanna and Ryan (2012), and Duflo, Greenstone, Pande and Ryan (2016) provide examples in development of using experiments to estimate key policy parameters.

<sup>35</sup>Banerjee et al. (2016) discuss how the plausibility of such identifying assumptions might limit external validity.

tracts, multi-period settings, and alternative policy targets that are worthy of study. Notable contributions in this vein include the recent work of Bai, Handel, Miguel and Rao (2019) and Aggarwal, Dizon-Ross and Zucker (2019).

## References

- Acemoglu, Daron**, “Theory, General Equilibrium, and Political Economy in Development Economics,” *Journal of Economic Perspectives*, 2010, *24* (3), 17–32.
- Aggarwal, Shilpa, Rebecca Dizon-Ross, and Ariel Zucker**, “Incentivizing Behavioral Change: The Role of Time Preferences,” 2019.
- Andreoni, James and Charles Sprenger**, “Estimating Time Preferences with Convex Budgets,” *American Economic Review*, 2012, *102* (7), 3333–3356.
- Ashraf, Nava, Dean Karlan, and Wesley Yin**, “Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines,” *Quarterly Journal of Economics*, 2006, *121* (1), 635–672.
- , **Oriana Bandiera, and Scott S. Lee**, “Do-Gooders and Go-Getters: Career Incentives, Selection, and Performance in Public Sector Service Delivery,” 2015.
- Attanasio, Orazio and Costas Meghir**, “Education Choices in Mexico: Using a Structural Model and a Randomized Experiment to Evaluate PROGRESA,” *Review of Economic Studies*, 2012, *79* (1), 37–66.
- Augenblick, Ned and Matthew Rabin**, “An Experiment on Time Preference and Misprediction in Unpleasant Tasks,” 2015.
- , **Muriel Niederle, and Charles Sprenger**, “Working Over Time: Dynamic Inconsistency in Real Effort Tasks,” *Quarterly Journal of Economics*, 2015, *130* (3), 1067–1115.
- Bai, Liang, Benjamin Handel, Edward Miguel, and Gautam Rao**, “Self-Control and Demand for Preventiv Health: Evidence from Hypertension in India,” 2019.
- Banerjee, Abhijit, Sylvain Chassang, and Erik Snowberg**, “Decision Theoretic Approaches to Experiment Design and External Validity,” 2016.

- Beauchamp, Jonathan P., Daniel J. Benjamin, Christopher F. Chabris, and David I. Laibson**, “Controlling for Compromise Effects Debiases Estimates of Preference Parameters,” *Working Paper*, 2015.
- Benartzi, Shlomo and Richard H. Thaler**, “Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving,” *Journal of Political Economy*, 2004, *112* (1), 164–87.
- Bertrand, Marianne, Robin Burgess, Arunish Chawla, and Guo Xu**, “Determinants and Consequences of Bureaucratic Effectiveness: Evidence from the Indian Administrative Service,” 2016.
- Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian**, “The Importance of Default Options for Retirement Savings Outcomes: Evidence from the United States,” *National Bureau of Economic Research Working Paper 12009*, 2009.
- Blumenstock, Joshua, Michael Callen, and Tarek Ghani**, “Why Do Defaults Affect Behavior? Experimental Evidence from Afghanistan,” *American Economic Review*, 2018, *108* (10), 2868–2901.
- Bruhn, Miriam and David McKenzie**, “In Pursuit of Balance: Randomization in Practice in Development Field Experiments,” *American Economic Journal: Applied Economics*, 2009, *1* (4), 200–232.
- Cagetti, Marco**, “Wealth Accumulation Over the Life Cycle and Precautionary Savings,” *Journal of Business and Economic Statistics*, 2003, *21* (3), 339–353.
- Callen, Michael, Saad Gulzar, Ali Hasanain, Yasir Khan, and Arman Rezaee**, “Data and Policy Decisions: Experimental Evidence from Pakistan,” 2019.



- Carvalho, Leandro S., Stephan Meier, and Stephanie W. Wang**, “Poverty and Economic Decision-Making: Evidence from Changes in Financial Resources at Payday,” *Working Paper*, 2014.
- Castillo, Marco, Paul Ferraro, Jeffrey Jordan, and Ragan Petrie**, “The Today and Tomorrow of Kids: Time Preferences and Educational Outcomes of Children,” *Journal of Public Economics*, 2011, *95* (11-12), 1377–1385.
- Chabris, Christopher F., David Laibson, and Jonathon P. Schuldt**, “Intertemporal Choice,” in Steven N. Durlauf and Larry Blume, eds., *The New Palgrave Dictionary of Economics*, London: Palgrave Macmillan, 2008.
- , – , **Carrie Morris, Jonathon Schuldt, and Dmitry Taubinsky**, “Individual Laboratory-Measured Discount Rates Predict Field Behavior,” *Journal of Risk and Uncertainty*, 2008, *37* (2-3), 237–269.
- Cubitt, Robin P. and Daniel Read**, “Can Intertemporal Choice Experiments Elicit Preferences for Consumption?,” *Experimental Economics*, 2007, *10* (4), 369–389.
- Dal Bó, Ernesto, Frederico Finan, and Martín A. Rossi**, “Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service,” *Quarterly Journal of Economics*, 2013, *forthcoming*.
- DellaVigna, Stefano and Ulrike Malmendier**, “Paying Not to Go to the Gym,” *American Economic Review*, 2006, *96* (3), 694–719.
- Deseranno, Erika**, “Financial Incentives as Signals: Experimental Evidence From the Recruitment of Health Promoters,” 2016.
- Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde**, “Dynamic Inconsistency Predicts Self-Control Problems in Humans,” *Working Paper*, 2006.

**Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan**, “The Value of Regulatory Discretion: Estimates from Environmental Inspections in India,” 2016.

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

– , **Rema Hanna, and Stephen P. Ryan**, “Incentives Work: Getting Teachers to Come to School,” *The American Economic Review*, 2012, *102* (4), 1241–1278.

**Finan, Frederico, Benjamin A. Olken, and Rohini Pande**, “The Personnel Economics of the State,” *Annual Review of Economics*, Forthcoming.

**Harrison, Glenn W., Morten I. Lau, and Melonie B. Williams**, “Estimating Individual Discount Rates in Denmark: A Field Experiment,” *American Economic Review*, 2002, *92* (5), 1606–1617.

– , – , **Elisabet E. Rutstrom, and Melonie B. Sullivan**, “Eliciting Risk and Time Preferences Using Field Experiments: Some Methodological Issues,” in Jeffrey Carpenter, Glenn W. Harrison, and John A. List, eds., *Field experiments in economics*, Vol. Vol. 10 (Research in Experimental Economics), Greenwich and London: JAI Press, 2005.

**Hastings, Justine and Ebonya Washington**, “The First of the Month Effect: Consumer Behavior and Store Responses,” *American Economic Journal - Economic Policy*, 2010, *2* (2), 142–162.

**Hausman, Jerry A.**, “Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables,” *The Bell Journal of Economics*, 1979, *10* (1).

- Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan**, “Self-Control and the Development of Work Arrangements,” *American Economic Review, Papers and Proceedings*, 2010, *100* (2), 624–628.
- , —, and —, “Self-Control at Work,” *Journal of Political Economy*, 2015, *123* (6), 1227–1277.
- Koopmans, Tjalling C.**, “Stationary Ordinal Utility and Impatience,” *Econometrica*, 1960, *28* (2), 287–309.
- Laibson, David**, “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, 1997, *112* (2), 443–477.
- , **Andrea Repetto, and Jeremy Tobacman**, “Estimating Discount Functions with Consumption Choices Over the Lifecycle,” *Working Paper*, 2005.
- Lawrance, Emily C.**, “Poverty and the Rate of Time Preference: Evidence from Panel Data,” *Journal of Political Economy*, 1991, *99* (1), 54–77.
- Mahajan, Aprajit, Christian Michel, and Alessandro Tarozzi**, “Identification of Time-Inconsistent Models: The Case of Insecticide Treated Nets,” *Working Paper*, 2020.
- Meier, Stephan and Charles Sprenger**, “Discounting Financial Literacy: Time Preferences and Participation in Financial Education Programs,” *IZA Discussion Paper Number 3507*, 2008.
- and —, “Present-Biased Preferences and Credit Card Borrowing,” *American Economic Journal - Applied Economics*, 2010, *2* (1), 193–210.
- and —, “Time Discounting Predicts Creditworthiness,” *Psychological Science*, 2012, *23*, 56–58.

- O'Donoghue, Ted and Matthew Rabin**, "Doing It Now or Later," *American Economic Review*, 1999, 89 (1), 103–124.
- and – , "Choice and Procrastination," *The Quarterly Journal of Economics*, 2001, 116 (1), 121–160.
- Read, Daniel and Barbara van Leeuwen**, "Predicting Hunger: The Effects of Appetite and Delay on Choice," *Organizational Behavior and Human Decision Processes*, 1998, 76 (2), 189–205.
- , **George Loewenstein, and Shobana Kalyanaraman**, "Combining the Immediacy Effect and the Diversification Heuristic," *Journal of Behavioral Decision Making*, 1999, 12, 257–273.
- Sadoff, Sally, Anya Samek, and Charles Sprenger**, "Dynamic Inconsistency in Food Choice: Experimental Evidence from a Food Desert," *Working Paper*, 2015.
- Samuelson, Paul A.**, "A Note on Measurement of Utility," *The Review of Economic Studies*, 1937, 4 (2), 155–161.
- Sayman, Serdar and Ayse Onculer**, "An Investigation of Time Inconsistency," *Management Science*, 2009, 55 (3), 470–482.
- Shapiro, Jesse**, "Is There a Daily Discount Rate? Evidence From the Food Stamp Nutrition Cycle," *Journal of Public Economics*, 2005, 89 (2-3), 303–325.
- Shapiro, Matthew D.**, "The Permanent Income Hypothesis and the Real Interest Rate: Some Evidence from Panel Data," *Economics Letters*, 1984, 14 (1), 93–100.
- Sprenger, Charles**, "Judging Experimental Evidence on Dynamic Inconsistency," *American Economic Review, Papers and Proceedings*, 2015.

**Train, Kenneth**, *Discrete Choice Methods with Simulation*, Cambridge University Press, 2009.

**Warner, John and Saul Pleeter**, “The Personal Discount Rate: Evidence from Military Downsizing Programs,” *American Economic Review*, 2001, *91* (1), 33–53.

**Zeldes, Stephen B.**, “Consumption and Liquidity Constraints: An Empirical Investigation,” *The Journal of Political Economy*, 1989, *97* (2), 305–346.