USING PREFERENCE ESTIMATES TO CUSTOMIZE INCENTIVES: AN APPLICATION TO POLIO VACCINATION DRIVES IN PAKISTAN

James Andreoni
Michael Callen
Yasir Khan
Karrar Jaffar
Charles Sprenger

Working Paper 22019
http://www.nber.org/papers/w22019

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2016

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 from financial support from the International Growth Center (IGC) 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, Matthew Kahn, Asim Khwaja, Aprajit Mahajan, Dilip Mookherjee, Andrew Newman, Robert Powell, Matthew Rabin, Gautam Rao, Martin Rotemberg, Joshua Tasoff, 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. Andreoni acknowledges financial support from the National Science Foundation, grant SES-1427355. The study is registered in the AEA RCT Registry as AEARCTR-0000417. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Using Preference Estimates to Customize Incentives: An Application to Polio Vaccination Drives in Pakistan
James Andreoni, Michael Callen, Yasir Khan, Karrar Jaffar, and Charles Sprenger
NBER Working Paper No. 22019
February 2016
JEL No. D03,I1,O1

ABSTRACT

We use structural estimates of time preferences to customize incentives for a sample of polio vaccinators during a series of door-to-door immunization drives in Pakistan. Our investigation proceeds in three stages. First, we measure time preferences using intertemporal allocations of vaccinations. Second, we derive the mapping between these structural estimates and individually optimal incentives given a specific policy objective. Third, we experimentally evaluate the effect of matching contract terms to individual discounting patterns in a subsequent experiment with the same vaccinators. This exercise provides a test of the specific point predictions given by structural estimates of time preference. We document present bias among vaccinators and find that tailored contracts achieve the intended policy objective of smoothing intertemporal allocations of effort. The benefits of customized incentives in terms of achieving the policy objective are largest for vaccinators allocating when present bias is relevant to the decision.

James Andreoni
Department of Economics
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0508
and NBER
andreoni@ucsd.edu

Michael Callen
Harvard Kennedy School
79 JFK Street
Cambridge, MA 02138
and NBER
michael_callen@hks.harvard.edu

Yasir Khan
UC Berkeley Haas School of Business
yasir.khan@berkeley.edu

Karrar Jaffar
USC Department of Economics
karrar.jaffar@gmail.com

Charles Sprenger
UCSD Rady School of Business
c.sprenger@gmail.com

A randomized controlled trials registry entry is available at
https://www.socialscienceregistry.org/trials/417

1 Introduction

Many economic decisions are intertemporal. From consumption and savings, to task performance and human capital accumulation, choice often entails tradeoffs through time. Characterizing such choices with structural models of discounting has retained the interest of economists for much of the last century (leading contributions include Samuelson, 1937; Koopmans, 1960; Laibson, 1997; O'Donoghue and Rabin, 2001). The preference parameters governing such models are of unique value for understanding a broad range of behaviors, and so have received much empirical attention as well. In both field and laboratory settings, substantial efforts have been made to structurally estimate the level and shape of discounting (examples include Hausman, 1979; Lawrance, 1991; Warner and Pleeter, 2001; Cagetti, 2003; Laibson, Repetto and Tobacman, 2005; Harrison, Lau and Williams, 2002; Andersen, Harrison, Lau and Rutstrom, 2008; Shapiro, 2005; Kuhn, 2013; Andreoni and Sprenger, 2012a).

This paper investigates the possibility of structurally estimating discounting parameters and testing corresponding point predictions for behavior out of sample. Interestingly, very little research makes use of such potential value in the articulation and estimation of structural models of discounting. Structural estimates of discount factors or discount rates have been used for comparison to market interest rates (e.g., Hausman, 1979), for comparison across samples, time, or elicitation and estimation strategies (e.g., Coller and Williams, 1999; Frederick, Loewenstein and O'Donoghue, 2002; Meier and Sprenger, 2015; Andersen et al., 2008), to assess differences in patience across subpopulations (e.g., Kirby, Petry and Bickel, 1999; Tanaka, Camerer and Nguyen, 2010; Harrison et al., 2002; Dohmen, Falk, Huffman and Sunde, 2010; Lawrance, 1991; Warner and Pleeter, 2001), and to conduct welfare analyses (e.g., Laibson, 1997). When structural estimates are used in out-of-sample exercises, the analysis has been unstructured, linking differences in measured patience to differences in other behaviors without an articulated model for the relationship (e.g., Chabris, Laibson, Morris, Schuldt and Taubinsky, 2008b; Meier and Sprenger, 2008, 2012, 2010; Ashraf, Karlan and Yin, 2006; Dohmen, Falk, Huffman and Sunde, 2006). Though such correlational exercises yield valuable insights, they could be conducted without appeal to a structural model by linking a non-parametric measure of discounting to other behaviors. That is, such exercises do not rely on the precise point predictions that structural estimation allows.

In an important policy-relevant environment, the provision of door-to-door polio vaccinations in Pakistan, we elicit and structurally estimate intertemporal preference parameters for polio vaccinators. These preference parameters are subsequently used to construct tailored intertemporal incentive contracts for the same set of vaccinators. The contract maps each vaccinator's discount factor to an interest rate designed to achieve a specific policy objective. The structural estimates are critical as without an estimated discount factor, we would not know how to tailor each vaccinator's contract.

Our project engages government health workers, termed Lady Health Workers (LHWs), associated with polio eradication efforts for the Department of Health in Lahore, Pakistan. 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.[1] The disease largely affects children under five. The function of LHWs is to provide oral polio vaccine to children during government organized vaccination drives, which usually last two or more days and are conducted approximately every month. LHWs are given a supply of oral vaccine and a neighborhood map and are asked to travel door-to-door vaccinating children with a suggested target for vaccination. Prior to our project, LHWs self-reported achievement and no technology existed for monitoring vaccinations. Consistent with the large literature on public sector absence (Banerjee and Duflo, 2006; Banerjee, Duflo and Glennerster, 2008; Chaudhury, Hammer, Kremer, Muralidharan and Rogers, 2006; Callen, Gulzar, Hasanain and Khan, 2015), LHWs often fall short of their suggested targets, but rarely report doing so.

We begin by introducing a high-resolution monitoring technology. Each vaccinator in our sample is provided a smart-phone, equipped with a reporting application. This application, discussed in detail below, permits precise observation of when and where vaccinations are

[1]Between 95 percent and 99 percent of individuals carrying polio disease are asymptomatic. One infection is therefore enough to indicate a substantial degree of ambient wild polio virus.

conducted.

Monitoring technology in place, we introduce intertemporal bonus contracts designed to elicit the time preferences of the LHWs performing door-to-door vaccinations. We offer a fixed bonus of 1000 rupees (around \$10) for completing a total of $V = 300$ vaccinations over a two-day drive. Vaccinators 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, are not met, the bonus is not received. Vaccinators are randomly assigned an interest rate, R, such that a single vaccination that is allocated to day 2 reduces by R the number of vaccinations required on day 1. That is, v_1 and v_2 satisfy the intertemporal budget constraint

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

Intertemporal allocations, (v_1, v_2), can be used to structurally estimate discounting parameters for LHWs. An additional experimental variation permits identification of an important behavioral aspect of intertemporal choice: the existence of present-biased preferences (Laibson, 1997; O'Donoghue and Rabin, 1999). Subjects are randomly assigned to make their allocation decision either in advance of the first day of the drive, day 0, or on day 1 itself. Additionally, subjects are randomly assigned an interest rate, R. Under specific structural assumptions, the experiment identifies a set of aggregate discounting parameters (for similar estimation strategies see Andreoni and Sprenger, 2012a; Augenblick, Niederle and Sprenger, 2015; Augenblick and Rabin, 2015). Further, under additional assumptions, each LHW's allocation identifies her individual discount factor.

We use the individual discounting parameters from this first drive to tailor incentive contracts in a follow-up drive. Our exercise has a clear policy rationale. We envision a policymaker, endowed with information on workers' time preferences, who would like to achieve a specific policy objective. We formalize the problem as maximizing policy preferences subject to a worker's offer curve. The specific policy preferences we study are Leontief, with a policymaker wishing to ensure smooth provision of service through time.[2] The optimal policy is simple. To ensure smooth provision of service, $v_1 = v_2$, the policymaker must give each worker an interest rate equal to their discount factor. To provide proof-of-concept for such tailored contracting, we randomly provide half of LHWs with a tailored contract and half with a random contract in the follow-up drive.

In a sample of 349 LHWs we document three principal results. First, on aggregate, a present bias exists in vaccination behavior. Vaccinators allocating three days in advance of a drive allocate significantly fewer vaccinations to v_1 than those allocating on the morning that the drive actually commences. Corresponding estimates of present bias accord with those of prior laboratory exercises. Second, substantial heterogeneity in discounting is observed. This heterogeneity is important as it points to possible gains from individually-tailored contracts. Third, tailored contracts work. Relative to random contracts, LHWs with tailored contracts provide significantly smoother service.

Our results have several implications. First, and most importantly, our results show the value of structural estimates of time preferences. The corresponding point predictions for individual behavior have empirical content and are potentially actionable. Our tailored contracts provide a proof-of-concept for leveraging structural estimates for policy purposes. Such evidence may prove valuable when considering extensions to other contract forms, firm or policy objectives, and other domains. Second, our exercise uses field behavior to identify time preferences. Indeed, ours are the first results identifying present bias in field choices about effort.[3] This evidence is 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, 2012a; Augenblick et al., 2015; Carvalho, Meier

[2]Our experiment also allows us to explore the viability of alternative policy preferences and tailored contracts such as a policy maker who wishes to maximize the total number of vaccinations.

[3]Other examples of present bias or dynamic inconsistency in field choices include Read and van Leeuwen (1998); Sadoff, Samek and Sprenger (2015) for food choices, Read, Loewenstein and Kalyanaraman (1999) for highbrow and lowbrow movie choices, Sayman and Onculer (2009) for cafe reward choices, and Duflo, Kremer and Robinson (2011) for fertilizer purchase decisions. Hallmarks of dynamic inconsistency in the decisions of data entry workers are investigated by Kaur, Kremer and Mullainathan (2010). For discussion of this literature, see Sprenger (2015).

and Wang, 2014). As with Augenblick et al. (2015); Carvalho et al. (2014), and Augenblick and Rabin (2015) our results show that when investigating non-monetary choices, present bias may well have empirical support. Third, the results show a path by which government service may falter. Government workers are critical to service delivery and dynamic inconsistency of workers may be a mechanism for bad delivery. This is a largely undiscussed avenue by which governments fail to provide. In this vein, our results may show ways by which service could be brought into line with policymaker preferences. Beyond the standard policy levers such as detecting shirking and increasing pay, our data indicate that the intertemporal nature of incentives may be important for service delivery.

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

2 Experimental Design and Structural Estimation

Our experiment has three core objectives: (i) eliciting the time preferences of vaccinators; (ii) testing for present-biased behavior in the intertemporal allocation of work; and (iii) testing the point predictions of corresponding structural estimates using tailored contracts.

In order to achieve these objectives, the effort of vaccinators must be observable. Hence, the first component of our study is a smartphone monitoring application for tracking vaccinations. We describe this technology in a first sub-section, then we describe the design elements which allow for the identification of preferences and present bias, and then we proceed to tailored contracts. A fourth sub-section is dedicated to design details.

2.1 Vaccinations and Smartphone Monitoring

The Department of Health in Lahore, Pakistan employs Lady Health Workers (LHWs) throughout the city to conduct polio vaccination drives. Every month there is a vaccination drive that is at least two days long. LHWs are organized into teams of one senior worker and one junior assistant. These teams work together throughout the drive. Our experiment focuses on the incentives of the senior LHW.

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").[4] No explicit incentives for completing vaccinations were provided and LHWs received a fixed daily wage of 100 rupees (around $1). LHWs were asked to walk their map, knocking on each compound door, and vaccinating each child for whom parental permission was granted.[5] 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 in chalk on the compound wall information related to the attempt (e.g., number of children vaccinated, whether or not all children were available for vaccination).[6]

At the end of each day of the drive, LHWs in each neighborhood convened at a local clinic with their supervisor, and self-reported their vaccinations for the day.[7] Though, in principle, a monitor could investigate the neighborhood chalk markings, in practice, this system provided for virtually no monitoring. There are reasons to believe that many reported vaccinations never actually took place.[8]

In collaboration with the Department of Health, we designed a smartphone-based monitoring system allowing us to track the door-to-door activities of polio vaccinators. Each LHW in our study was given a smartphone equipped with a vaccination monitoring application. The

[4]Appendix Figure A.1 provides an example of such a map.

[5]Appendix Figure A.2 provides an example of such a vaccination.

[6]Appendix Figure A.3 provides a picture of such a chalk marking.

[7]Appendix Figure A.4 provides a picture of the form capturing the self-reports. The second column records the number of vaccinations for the day. The seventh column reports the number of vials of vaccine used in the process.

[8]We attempted to independently audit vaccinators by following the trail of chalk markings, but our enumerators found the process to difficult to produce a reliable of houses visited. We do, however, know the targets associated with each micro-plan prior to our monitoring intervention and that vaccinators almost never reported failing to meet a target. Even with a bonus incentive and smartphone monitoring in place, we find that vaccinators on average achieve only 62 percent (s.d. = 58 percent) of the target given by their micro-plans. There is reason to expect that vaccinators would achieve a smaller share of their target with no incentive.

Panel A: Splash Page　　　　Panel B: Slider Bar

Figure 1: Vaccination Monitoring Smarthphone App

LHW was asked to record a set of information related to each vaccination attempt identical to what she had written in the chalk mark. Then she was asked to take a picture of the chalk mark 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. This dashboard is depicted in Appendix Figure A.5.[9]

The smartphone system allows us to register vaccination attempts and provides a basis for creating intertemporal bonus contracts designed to elicit LHW time preferences. We next provide an outline of the bonus contracts.

2.2 Intertemporal Bonus Contracts

Vaccination drives in Lahore occur approximately once every month. Vaccination drives are at least two days long. Given the intertemporal nature of effort for these drives, we worked

[9]This dashboard system is based on the technology described in Callen et al. (2015).

with the Department of Health to implement intertemporal bonus contracts in two-day drives in September, November and December of 2014.

The intertemporal bonus contracts required workers to complete a present value total of $V = 300$ vaccinations in exchange for a fixed bonus of 1000 rupees. Vaccinators 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 an interest rate 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 intertemporal bonus contract is identical to an experimental device termed a Convex Time Budget used to investigate time preferences (Andreoni and Sprenger, 2012a,b; Augenblick et al., 2015; Carvalho et al., 2014; Gine, Goldberg, Silverman and Yang, 2010).[10]

The intertemporal allocation (v_1, v_2) potentially carries information on the time preferences of each LHW. A more patient worker may wish to allocate more vaccinations to day 1, while a less patient worker may wish to allocate more vaccinations to day 2. In Section 4, we discuss alternative rationales for specific patterns of allocation behavior. Under the assumption that choices are driven by discounting, we next describe the relevant experimental variation and structural assumptions that permit us to identify discounting parameters at the aggregate and individual level.

[10] 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 and Structural Identification

Our design generates two sources of experimental variation. First, each LHW is randomly assigned an interest rate, 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 an assigned interest rate, 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 beginning vaccinations.[11]

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 term the first of these 'Advance' decisions and the second 'Immediate' decisions. The assignment to either the Advance or Immediate group was orthogonal to the interest rate assignment. Section 2.4 describes the efforts taken to make everything else besides allocation timing equal between these conditions.

Random assignment to Advance or Immediate choice and random assignment of R are both critical design elements for identifying discounting parameters of interest. In particular, we assume that individuals minimize the discounted costs of effort subject to the intertemporal budget constraint provided by their bonus contract. We make two further structural assumptions. 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 the convex costs of effort. Second, we assume that individuals discount the future quasi-hyperbolically (Laibson, 1997; O'Donoghue and Rabin, 1999). Hence, the worker's preferences can be written as

$$v_1^\gamma + \beta^{\mathbf{1}_{d=1}} \delta \cdot v_2^\gamma.$$

The indicator $\mathbf{1}_{d=1}$ captures whether the decision is made in advance or immediately on day 1. The parameters β and δ summarize individual discounting with β capturing the degree of

[11] One advantage of using slider bars over conventional price lists is potentially assisting subject comprehension.

present bias, active for LHWs who make Immediate decisions, i.e. $\mathbf{1}_{d=1} = 1$. If $\beta = 1$, the model nests exponential discounting with discount factor δ, while if $\beta < 1$ the decisionmaker exhibits a present bias, being less patient in Immediate relative to Advance decisions.

Minimizing discounted costs subject to the intertemporal budget constraint of the experiment yields intertemporal Euler equation

$$\left(\frac{v_1}{v_2}\right)^{\gamma-1} \frac{1}{\beta^{\mathbf{1}_{d=1}}\delta} = \frac{1}{R}. \tag{1}$$

Taking logs and rearranging yields

$$log\left(\frac{v_1}{v_2}\right) = \frac{log\delta}{\gamma-1} + \frac{log\beta}{\gamma-1}\mathbf{1}_{d=1} - \frac{1}{\gamma-1}logR.$$

If we assume allocations satisfy the above equation subject to an additive error term, ϵ, we arrive at the linear regression equation

$$log\left(\frac{v_1}{v_2}\right) = \frac{log\delta}{\gamma-1} + \frac{log\beta}{\gamma-1}\mathbf{1}_{d=1} - \frac{1}{\gamma-1}logR + \epsilon, \tag{2}$$

which can be estimated with standard techniques. This formulation provides intuition for the identification of structural parameters from LHW allocations, and make clear the purpose of our experimental variation in R and $\mathbf{1}_{d=1}$. Variation in the interest rate, R, identifies the shape of the cost function, γ, while variation in $\mathbf{1}_{d=1}$ identifies β. Note that δ is identified from the average level of v_1 relative to v_2 when decisions are made in advance (i.e. identified from the constant). An identical strategy for structurally estimating time preferences was put to use in Andreoni and Sprenger (2012a) and Augenblick et al. (2015) and has precedents in a body of macroeconomic research identifying aggregate preferences from consumption data (see e.g., Shapiro, 1984; Zeldes, 1989; Lawrance, 1991)

The above development delivers aggregate estimates of discounting parameters with each LHW's allocation contributing a single observation to the aggregate. Exercises exploring heterogeneity in time preferences document substantial differences across people, even from relatively homogeneous populations (see e.g., Harrison et al., 2002; Ashraf et al., 2006; Meier and Sprenger, 2015). Given only a single observation per LHW, estimation of all parameters at the individual level is infeasible. However, we can calculate each LHWs discount factor, either δ for those who make Advance decisions or $\beta\delta$ for those who make Immediate decisions. To make such a calculation, two further structural assumptions are required. First, we assume every LHW shares a common cost function, $\gamma = 2$, corresponding to quadratic cost. Second, we assume the intertemporal Euler equation (1) is satisfied with equality. Under these assumptions,

$$\frac{R \cdot v_1}{v_2} = \beta^{\mathbf{1}_{d=1}}\delta, \tag{3}$$

such that the interest rate-adjusted ratio of allocated vaccinations identifies a discount factor for each LHW.

The structural assumptions required for identification of aggregate and individual discount factors are potentially quite restrictive. Our research design, which involves tailoring contracts to individual discount factors, requires an ex-ante commitment to the specific functional forms of equations (1) and (3). Hence, for the purposes of estimation and tailoring we analyze only these previously determined functional forms. In sub-section 4.1.2, we assess the validity of a set of required assumptions and present exploratory analysis related to alternative functional forms.

2.3 Tailored Contracts

Under the set of structural assumptions above, each LHW's allocation in an intertemporal bonus contract identifies her discount factor for vaccinations, either δ for those who make Advance decisions or $\beta\delta$ for those who make Immediate decisions. We consider a policymaker who knows such preferences and wishes to achieve a specific policy objective. The policymaker has only one policy lever: manipulation of the interest rate, R. We formalize the problem as maximizing policy preferences, $P(v_1(R), v_2(R))$, subject to the LHW's offer curve. The problem is stated as

$$\begin{aligned} max_R \ P(v_1^*(R), v_2^*(R)) \ s.t \\ (v_1^*(R), v_2^*(R)) = argmin \ (v_1)^\gamma + \beta^{\mathbf{1}_{d=1}} \delta \cdot (v_2)^\gamma \ s.t. \\ v_1 + R \cdot v_2 = V. \end{aligned}$$

The solution maps the policy preferences into an interest rate for each LHW. One can consider many potential forms of policy preference, with policymakers desiring a variety of intertemporal patterns of effort. As proof-of-concept, we consider first a policy maker with one extreme form of preference, $P(v_1(R), v_2(R)) = min[v_1(R), v_2(R)]$.[12] Such Leontief preferences correspond to a policymaker who desires smooth provision of service. This problem has an intuitive solution. The worker's intertemporal Euler equation (3) yields smooth provision, $v_1 = v_2$, when $R = \beta^{\mathbf{1}_{t=1}} \delta$. Hence, the tailored contracts give each vaccinator a value of R equal to their discount factor. Note that the structural discounting parameters are critical in this development. Without information on discount factors, the policymaker has no precise idea how to tailor contracts for each worker to achieve her specific objectives.

In a second two-day drive we investigate the promise of 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 interest rate in the second drive.

Half of LHWs were given a tailored intertemporal bonus contract,

$$v_1 + R^* \cdot v_2 = V,$$

[12]The ability of our data to speak to alternative policy preferences is discussed in section 4.2.3. Leontief preferences in this environment are extreme, but there is general interest in understanding mechanisms to drive smooth behavior, particularly for saving and for avoiding procrastination.

where $R^* = \beta^{\mathbf{1}_{d=1}}\delta$, either $\beta\delta$ or δ depending on whether they made Immediate or Advance decisions.[13] Some LHW's allocation behavior in the first drive implied extreme discount factors and hence extreme values of R^*. Our tailoring exercise focused only on LHWs with discount factors between 0.75 and 1.5. LHWs outside of these bounds were given either the upper or lower bound accordingly.

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

$$v_1 + \tilde{R} \cdot v_2 = V,$$

where $\tilde{R}$ was drawn from a random uniform distribution $U[0.75, 1.5]$. The bounds on the distribution of $\tilde{R}$ were determined to match the bounds on R^*, while the choice of a random control (rather than a single value of $\tilde{R}$) was chosen to assess whether tailoring is effective conditional on the interest rate assigned.

Random assignment to tailoring in Drive 2 is stratified on the measure of absolute distance to equal provision $|\frac{v_1}{v_2} - 1|$, based on allocations from Drive 1.[14] This measure of distance to equal provision also serves as our eventual measure of distance from equal provision when analyzing the effect of assignment to 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.

2.4 Design Details

Our experiment is divided into two drives. A first drive measures LHW time preferences and a second drive explores the potential of tailored contracts. Each drive is two days in length, beginning on Monday morning and ending on Tuesday evening. Our team trained LHWs on operating the smartphone application and submitting vaccinations on the Friday before the

[13]Note that this tailoring exercise requires that LHWs remain in either in the Immediate or Advance assignment across drives.

[14]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.

preceding drive. 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.

2.4.1 Training and Allocation Decisions

On November 7, all LHWs participating in the November 10-11 drive were required to attend training at one of three separate central locations in Lahore. In a two hour session, they were trained on the smartphone vaccination application and its correct use.

Independent of treatment status, LHWs were also trained on the intertemporal bonus contracts and the process by which allocations were made and submitted. Hence, both Advance and Immediate LHWs were given identical training and information on how to submit their allocations. Additionally, LHWs with different interest rates were given identical information.

At the end of the training, LHWs assigned to Advance decision were asked to select their allocations. The decision screen for this group was activated and their allocations were recorded. Assistance was available from training staff for those who required assistance. LHWs assigned to Immediate decision were told they would select their allocations on Monday morning before beginning work. The decision screen for this group was activated on Monday morning and their allocations were recorded. 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

Figure 2 summarizes our experimental timeline and the sample for each vaccination drive of our study.

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 intro treatments and trained. Advance allocation decisions were collected from half of subjects on Friday, September 26th. On Monday, September 29th, when we attempted to collect immediate allocation decisions, a problem with the application prevented 82 of 168 Immediate decision LHWs from submitting their allocations. This was partly due to a general problem affecting the wireless network in our study location during these two days. This substantial treatment-related attrition pushed 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. Figure 2 provides sample details. From all appearances the failure of the application for these 82 individuals was random and related to the interaction of the application and the mobile network.[15] For completeness, we present data from Drive 0, but do not use Drive 0 for the purposes of tailoring contracts. The smartphone application issue was fixed quickly and we targeted the next vaccination drive for study.

Drive 1, November 7-11, 2014: Of the original 336 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 interest rate and allocation timing conditions. 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. 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. Because 11 LHWs attrited from the Immediate Choice condition, we

[15] Appendix Table A.1 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.

Objectives: 1. Measure preferences
2. Test for dynamic inconsistency

Sample: 336

Notes: 82 vaccinators of the 336 could not select task allocations because of a problem with the app.

Sample Allocation:

	R=0.9	R=1	R=1.1	R=1.25
Advance Choice	42	42	42	42
Immediate Choice	42	42	42	42

Objectives: 1. Measure preferences
2. Test for dynamic inconsistency

Sample: 349

Notes: Preferences are estimated for 338 of the 349 vaccinators recruited. A panel for Drive 0 and Drive 1 is available for 232 vaccinators.

Sample Allocation:

	R=0.9	R=1	R=1.1	R=1.25
Advance Choice	43	46	40	45
Immediate Choice	41	46	38	39

Objectives: 1. Test tailored contracts
2. Test tailoring by decision timing

Sample: Tailored (169), Untailored (168)

Notes: 337 of the 338 vaccinators participating in Drive 1 also participated in Drive 2 and were assigned to either Tailored or Untailored.

Sample Allocation:

	Tailored	Untailored
Advance Choice	85	88
Immediate Choice	84	80

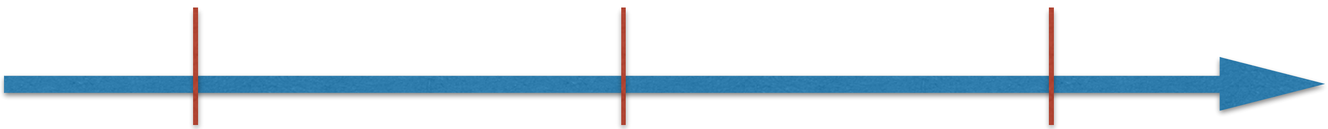

Failed Drive 0: September 26 - 30, 2014

Drive 1: November 7 - 11, 2014

Drive 2: December 5 - 9, 2014

Figure 2: Experiment Overview

Notes: Assignment to the advance choice and immediate choice condition in Drive 2 is inherited from vaccination Drive 1. Among vaccinators in the bonus+phone treatment: (i) 57 vaccinators participated only in Failed Drive 0; (ii) 6 vaccinators participated in Drive 1 only; (iii) 1 vaccinator participated in Failed Drive 0 and Drive 1, but not in Drive 2 (iii) 67 vaccinators participated in drives 2 and 3 only; (iv) 271 vaccinators participated in all three rounds.

also provide bounds on the estimated effect of decision timing using the method of Lee (2009). In addition, for 232 LHWs, we have allocation decisions in both the failed drive, Drive 0, and Drive 1, forming a potentially valuable panel of response. Figure 2 provides sample details.

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 tailored or untailored in their Drive 2 bonus contracts. Importantly, LHWs retained their Advance or Immediate assignment, such that Drive 2 delivers a 2x2 design for tailoring and allocation timing. This allows us not only to investigate the effect of tailoring, but whether tailoring is more or less effective depending on when allocations are made (and hence whether present bias is active). Figure 2 provides sample details.

2.4.3 Sample Details

Table 1 summarizes our sample of LHWs from Drive 1 and provides tests of experimental balance on observables. Column (1) presents the mean and standard deviation for each variable; columns (2) - (9) present the mean and standard error for each of our eight treatment arms, and column 10 presents a p-value corresponding to a joint test 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, only the one performed on a dummy variable equal to one for Punjabi subjects suggests baseline imbalance.

3 Results

Our project has two phases. The first phase focused on measuring intertemporal preferences. The second phase aimed specifically to test the effect of tailoring contracts.[16] We first report results related to the elicitation of intertemporal preference parameters, then evaluate the possibility of tailoring incentives based on individual preferences.

3.1 Elicitation of Time Preferences

3.1.1 Aggregate Behavior

Figure 3 presents mean behavior in the elicitation phase of our experiment, graphing the average allocation to the sooner work date, v_1, for each interest rate. Separate series are provided for Advance and Immediate choice. For completeness, we present the data from our failed drive,

[16]In addition, to test just the effect of providing the \$10 bonus, we randomly assigned 85 LHWs in Drive 0 to carry a phone but not receive an incentive. 73 of these LHWs also participated in Drive 1, retaining the same 'phone only' treatment status. In Drive 0, LHWs in the 'phone only' group attempted 143.12 vaccinations (s.e. = 12.85) while LHWs in the phone plus incentives group attempted 162.73 vaccinations (s.e. = 6.87), yielding an estimated increase of 19.61 attempts (s.e. = 15.11, p = .195). In Drive 1, LHWs in the 'phone only' group attempted 170.48 vaccinations (s.e. = 15.78) and LHWs in the phone plus incentives group attempted 201.32 vaccinations (se = 7.73) yielding an estimated increase of 30.84 attempts (s.e. = 18.37, p = .09).

Table 1: Summary Statistics and Covariates Balance

	Full Sample	Advance Decision				Immediate Decision				p-value
		R=0.9	R=1	R=1.1	R=1.25	R=0.9	R=1	R=1.1	R=1.25	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Demographics										
Gender (Female = 1)	0.985	1.000	1.000	1.000	0.978	0.975	0.978	0.947	1.000	0.284
	[0.121]	(0.000)	(0.000)	(0.000)	(0.022)	(0.025)	(0.022)	(0.037)	(0.000)	
Years of Education	10.415	10.767	10.652	10.650	10.279	9.850	10.565	10.184	10.282	0.500
	[2.291]	(0.416)	(0.273)	(0.462)	(0.330)	(0.298)	(0.368)	(0.238)	(0.395)	
Number of Children	3.424	3.419	3.422	3.538	3.286	3.605	3.391	3.421	3.333	0.997
	[1.826]	(0.279)	(0.301)	(0.309)	(0.296)	(0.286)	(0.274)	(0.243)	(0.294)	
Punjabi (=1)	0.952	0.930	0.932	1.000	0.955	0.950	0.978	0.917	0.947	0.022
	[0.215]	(0.039)	(0.038)	(0.000)	(0.032)	(0.035)	(0.022)	(0.047)	(0.037)	
Financial Background										
Has a Savings Account (=1)	0.269	0.310	0.250	0.275	0.302	0.350	0.283	0.189	0.179	0.630
	[0.444]	(0.072)	(0.066)	(0.071)	(0.071)	(0.076)	(0.067)	(0.065)	(0.062)	
Participated in a Rosca (=1)	0.389	0.349	0.378	0.425	0.350	0.500	0.289	0.351	0.487	0.482
	[0.488]	(0.074)	(0.073)	(0.079)	(0.076)	(0.080)	(0.068)	(0.079)	(0.081)	
Health Work Experience										
Years in Health Department	10.520	10.605	10.578	10.211	11.549	9.050	10.678	10.395	11.026	0.456
	[4.961]	(0.777)	(0.695)	(0.685)	(0.792)	(0.695)	(0.846)	(0.867)	(0.808)	
Years as Polio Vaccinator	10.428	10.209	10.728	11.050	11.143	9.238	9.935	10.447	10.692	0.581
	[4.727]	(0.758)	(0.689)	(0.668)	(0.743)	(0.689)	(0.713)	(0.858)	(0.751)	
# Vaccinators	338	43	46	40	45	41	46	38	39	

Notes: Column 1 presents the mean for each variable based on our sample of 338 vaccinators . These 338 vaccinators comprise the estimation sample in Table 2, which reports tests of dynamic inconsistency. Standard deviations are in brackets. Columns 2 to 9 report the mean level of each variable, with standard errors in parentheses, for each treatment cell. For each variable, Column 10 reports the p-value of a joint test that the mean levels are the same for all treatment cells (Columns 2–9). The last row presents the number of observations in each treatment condition. A ROSCA is an informal Rotating Savings and Credit Association. Some calculations used a smaller sample size due to missing information. The proportion of subjects with missing information for each variable is never greater than 3.5 percent (8 vaccinators did not report whether they had participated in a ROSCA).

Drive 0, alongside that of our primary elicitation drive, Drive 1, in Panel A, and present Drive 1 alone in Panel B. Two features of Figure 3 are notable. First, subjects appear to respond to the between-subject variation in interest rate. As the value R increases, vaccinations allocated to v_1 count relatively less towards reaching the two-day target of $V = 300$. Subjects respond to this changing incentive by reducing their allocation of v_1. Second, there is an average tendency of present bias. Subjects appear on average to allocate fewer vaccinations to v_1 when making immediate choice.

Table 2 presents corresponding regression analysis for aggregate behavior in the combined drives and Drive 1 alone. We regress v_1 on R and whether the allocation decision is immediate.

Column (1) echoes the findings from Figure 3, Panel A, demonstrating that LHWs allocated significantly fewer vaccinations to v_1 as R increases and when the allocation decision is immediate. Across Drives 0 and 1, LHWs assigned to Immediate choice allocate 5.62 (s.e. = 2.20) fewer vaccinations to v_1 than those assigned to Advance choice. In columns (2)-(4) we analyze the robustness of this result to changes in estimation method and sample restrictions. Column (2) conducts quantile regression identifying a median effect of 2.00 (s.e. = 0.95) fewer vaccinations in Immediate choice. Column (3) restricts the sample by trimming subjects with extreme allocation behavior that would imply individual discount factors from equation (3) outside of the range of [0.75, 1.5]. Column (4) restricts the sample by focusing only on the 232 LHWs who participated in both Drive 0 and Drive 1. For both of these restricted samples, we find median effects of 3 fewer vaccinations in immediate choice, indicating that the general tendency of present bias in the data is not driven by extreme allocation behavior or LHWs who attrit from the sample. In columns (5)-(8) we conduct the same analysis focusing only on Drive 1 behavior and reach quite similar conclusions. It should be noted that though the median effects remain significant when analyzing Drive 1 alone, the mean effect of column (5) falls outside of the range of significance. In addition, as discussed in Section 2.4.2 above, 11 LHWs attrited from the sample in the immediate choice condition in drive 1. Bounding the effect of being assigned to the immediate choice condition on v_1 allocations using the method of Lee (2009) provides a lower bound of -3.78 tasks (s.e. = 2.06) and an upper bound of 0.205 tasks (s.e. = 2.06). The less dramatic mean present bias in Drive 1 relative to the combined drives echoes Augenblick et al. (2015) who finds that dynamic inconsistency is reduced with experience.

3.1.2 Aggregate Preference Parameters

The raw data of Figure 3 and analysis of Table 2 indicate responsiveness of LHW behavior to our experimentally varied parameters, R and whether allocations are Immediate or Advance. Equation (2) links allocation behavior to these experimental parameters via a structural model

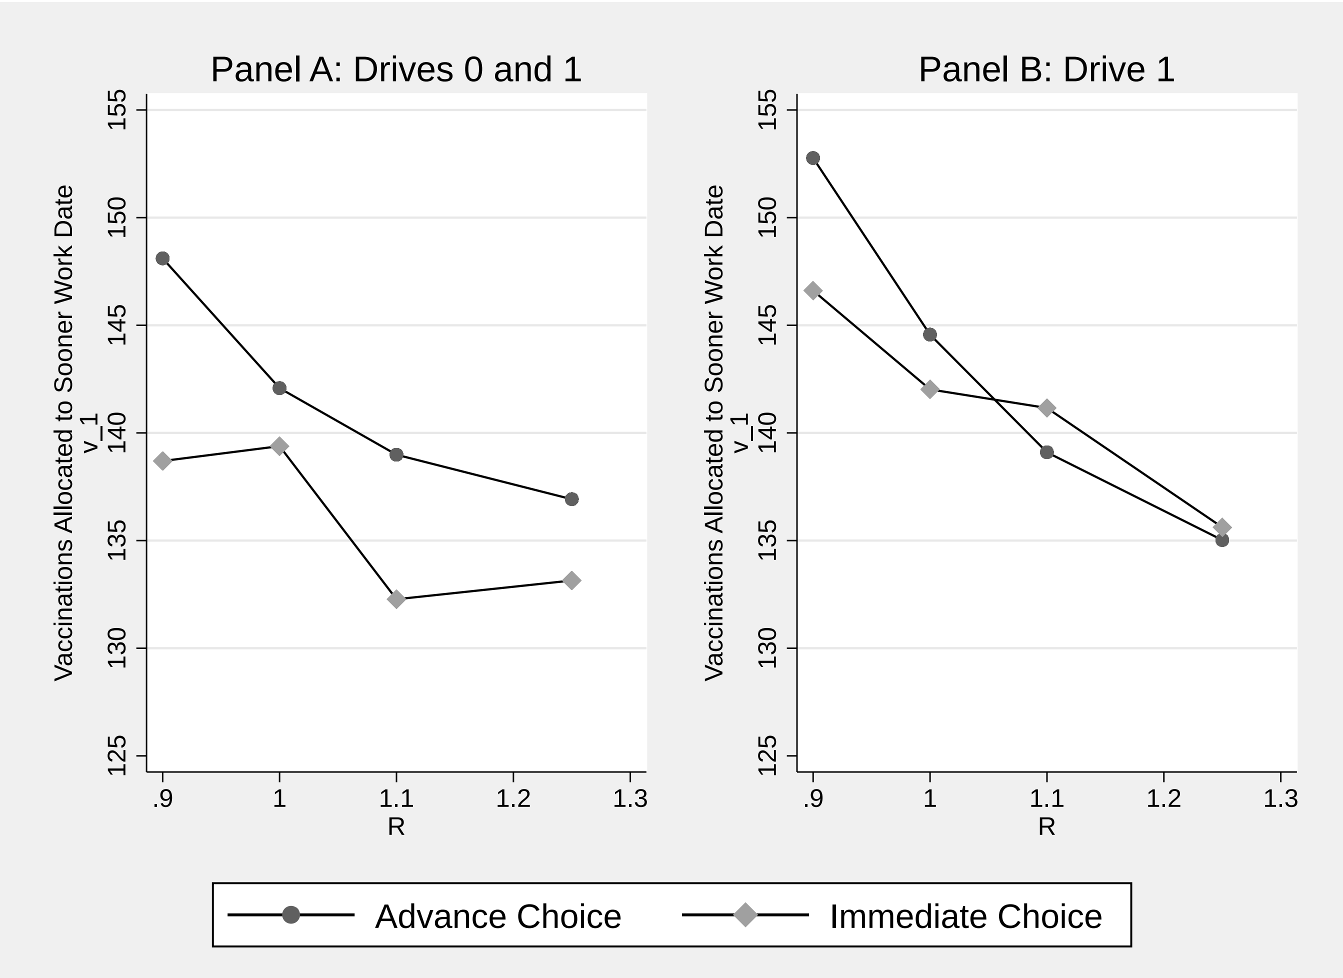

Figure 3: Discounting Behavior

Notes: Mean behavior in Drives 0 and 1 combined (Panel A), and Drive 1 alone (Panel B).

Table 2: Aggregate Behavior

Dependent variable:	Tasks Allocated to the First Day of the Drive (v_1)							
	Drives 0 & 1 Combined				Drive 1			
	(1) OLS	(2) Median	(3) Median	(4) Median	(5) OLS	(6) Median	(7) Median	(8) Median
Immediate Decision (=1)	-5.62**	-2.00**	-3.00***	-3.00**	-1.63	-2.00*	-3.00***	-2.14*
	(2.20)	(0.95)	(0.88)	(1.16)	(1.95)	(1.13)	(0.91)	(1.23)
Interest Rate	-25.91***	-40.00***	-60.00***	-37.14***	-39.92***	-54.29***	-66.67***	-54.29***
	(7.78)	(6.04)	(4.12)	(7.91)	(7.42)	(4.38)	(3.66)	(4.68)
Constant	169.07***	188.00***	210.00***	186.43***	185.30***	201.86***	216.33***	201.86***
	(8.41)	(6.06)	(4.34)	(7.93)	(7.91)	(4.72)	(3.93)	(5.03)
Include Boundary Sample	Yes	Yes	No	Yes	Yes	Yes	No	Yes
Only Panel Sample	No	No	No	Yes	No	No	No	Yes
R-Squared	0.027				0.081			
Mean/Median Advance Choice	141.54	150.00	150.00	150.00	142.87	146.50	148.00	146.50
# Observations	622	622	475	464	338	338	281	232

Notes: This table reports on the effect of making a decision on the morning of the drive on vaccinations allocated to the first day of the drive. Standard errors are reported in parentheses. Standard errors in column 1 - 4 are clustered at the vaccinator level. Clustered standard errors for quantile regressions are calculated using the approach in Parente and Santos Silva (2016). Immediate Decision is a dummy variable equal to one for vaccinators selecting their allocations on the morning of the vaccination drive. The interest rate R takes the values $R \in \{0.9, 1, 1.1, 1.25\}$. Estimates in columns 1 and 5 are estimated using OLS and estimates in the remaining columns are estimated using quantile regressions evaluated at the median. Table 3 provides corresponding between-subject structural preference parameter estimates of β, δ, and γ. *Levels of Significance:* $*p < 0.1$, $**p < 0.05$, $***p < 0.01$.

of choice. In Table 3 we present parameter estimates from maximum likelihood estimations of equation (2) with $\epsilon \sim N(0, \sigma^2)$ under a set of structural assumptions and sample restrictions.

In columns (1) and (5) of Table 3 we present estimates restricting $\gamma = 2$ (quadratic costs) for Drives 0 and 1 combined and Drive 1 alone. These estimates serve as an aggregate benchmark for our individual analysis which calculates individual discount factors under the assumption of quadratic costs. In column (1), we estimate a daily δ of 0.989 (0.019), β of 0.880 (0.041), and we reject the null hypothesis of zero present bias $\chi^2(1) = 8.57$, $(p < 0.01)$. These estimates for discounting parameters are close to prior estimates using effort tasks in the laboratory (Augenblick et al., 2015; Augenblick and Rabin, 2015). Following average behavior, present bias is found to be more limited when examining Drive 1 alone in column (5), though β is still estimated to be less than one in all specifications.

In columns (2) and (6) of Table 3, we relax our restriction of quadratic costs and attempt to estimate the shape of the cost function. In these estimates we restrict $\gamma \in [1, 3]$ by estimating the parameter of a box constraint, a, such that $\gamma = 1 + 2 \cdot \frac{1}{1+exp(a)}$. This restriction ensures allocations are indeed minima (i.e., $\gamma > 1$) and allows costs to be substantially convex (i.e.,

Table 3: Aggregate Parameter Estimates

	Drives 0 & 1 Combined				Drive 1			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
β	0.880	0.775	0.824	0.817	0.955	0.916	0.967	0.955
	(0.041)	(0.070)	(0.064)	(0.074)	(0.037)	(0.066)	(0.029)	(0.060)
δ	0.989	0.928	1.010	0.986	1.013	0.973	1.006	0.974
	(0.019)	(0.032)	(0.024)	(0.030)	(0.018)	(0.029)	(0.019)	(0.033)
a	-	-19.860	-15.415	-16.554	-	-19.827	-134.455	-20.045
	-	(0.492)	(0.408)	(0.445)	-	(1.189)	(8.736)	(0.190)
$\gamma = 1 + 2 \cdot \frac{1}{1+exp(a)}$	2	3	3	3	2	3	3	3
$ln(\sigma)$	-0.589	-0.623	-0.965	-0.767	-1.049	-1.122	-2.081	-1.453
	(0.089)	(0.094)	(0.146)	(0.124)	(0.150)	(0.176)	(0.039)	(0.089)
# Observations	622	622	475	464	338	338	281	232
Log-Likelihood	-516.137	-494.997	-215.567	-302.737	-125.092	-100.423	186.012	7.876
$H_0 : \beta = 1$	$\chi^2(1) = 8.57$	$\chi^2(1) = 10.26$	$\chi^2(1) = 7.67$	$\chi^2(1) = 6.08$	$\chi^2(1) = 1.50$	$\chi^2(1) = 1.62$	$\chi^2(1) = 1.26$	$\chi^2(1) = 0.56$
	$(p < 0.01)$	$(p < 0.01)$	$(p < 0.01)$	$(p < 0.05)$	$(p = 0.22)$	$(p = 0.20)$	$(p = 0.26)$	$(p = 0.56)$

Notes: This reports structural estimates of β, δ, and γ obtained using Maximum Likelihood Estimation based on Equation (2). Standard errors are reported in parentheses. Estimates in columns (1) and (5) are obtained imposing the restriction $\gamma = 2$. Estimates in columns (2)-(4) and in columns (6) - (8) remove this restriction. In columns (3), (4), (7) and (8) of Table 3 we restrict the sample by excluding extreme discount factors and by focusing only on our panel of LHWs who participated in both Drive 0 and Drive 1 as in Table 2.

up to cubic). Relaxing our restriction on γ alters somewhat the conclusions with respect to discounting. In particular, β and δ are both estimated to be further from 1 in columns (2) and (6) relative to columns (1) and (5). In columns (3), (4), (7) and (8) of Table 3, we restrict the sample by excluding extreme discount factors and by focusing only on our panel of LHWs who participated in both Drive 0 and Drive 1. As in Table 2, the conclusions are not much altered by these sample restrictions.

It is important to note that in both columns (2) and (6) the cost parameter a is estimated to be around -20, implying γ extremely close to 3. Estimating close to the edge of the box constraint may suggest some mis-specification in functional form or constraint choice. In section 4.1.2, we attempt to evaluate the appropriateness of our functional form assumptions and assess plausible alternative formulations. The conducted exercises identify an important issue with respect to the estimates of Table 3: the estimated parameters predict more sensitivity to R than truly exists in the data. This potential mis-specification presents a clear challenge for using individual preference parameters for tailored contracts. Having committed to a possibly mis-specified functional form ex-ante, any success in tailoring contracts should likely be viewed as a lower bound on the potential benefits of such initiatives.

3.1.3 Individual Preference Parameters

The aggregate estimates of Table 3 mask substantial heterogeneity across subjects. Following equation (3), we calculate individual discount factors for each LHW assuming quadratic costs. For those LHWs assigned to Advance choice this discount factor corresponds to δ, while for those assigned to Immediate choice it corresponds to $\beta\delta$. In Drive 1, the median [25th-75th %-ile] discount factor in Advance choice is 1.015 [0.88, 1.18], while the median discount factor in Immediate choice is 1 [0.84, 1.21].

An important minority of subjects have extreme discount factor calculations. Fifty-seven of 338 subjects in Drive 1 have implied discount factors either above 1.5 or below 0.75. Such extreme behavior is slightly more pronounced in Immediate choice (34 LHWs) relative to Advance choice (23 LHWs), ($t = 1.84,\ p = 0.07$). We term such LHWs the 'boundary sample.' As our tailoring exercise focuses on individuals with discount factors between 0.75 and 1.5, we restrict our individual analysis to the remaining 281 LHWs and discuss the boundary sample in robustness tests (see section 4.2). Figure 4 presents histograms of implied discount factors for these 281 LHWs in Advance and Immediate decisions. Two features are notable. First, in both contexts substantial heterogeneity in discount factors is observed. The 25th to 75th percentile ranges from 0.92 to 1.15 in Advance choice and from 0.88 to 1.15 in Immediate choice. Second, a present bias is observed in the shape of the distributions. The one period discount factors are skewed below 1 in Immediate relative to Advance choice. A Kolmogorov-Smirnov test marginally rejects equality of distributions ($D_{KS} = 0.15,\ p = 0.09$).

The observed heterogeneity in discount factors across LHWs resonates with prior exercises demonstrating heterogeneity of preferences even with relatively homogeneous samples (see e.g., Harrison et al., 2002; Ashraf et al., 2006; Meier and Sprenger, 2015). Further, this heterogeneity carries some promise for the possibility of individually-tailored contracts. That is, without

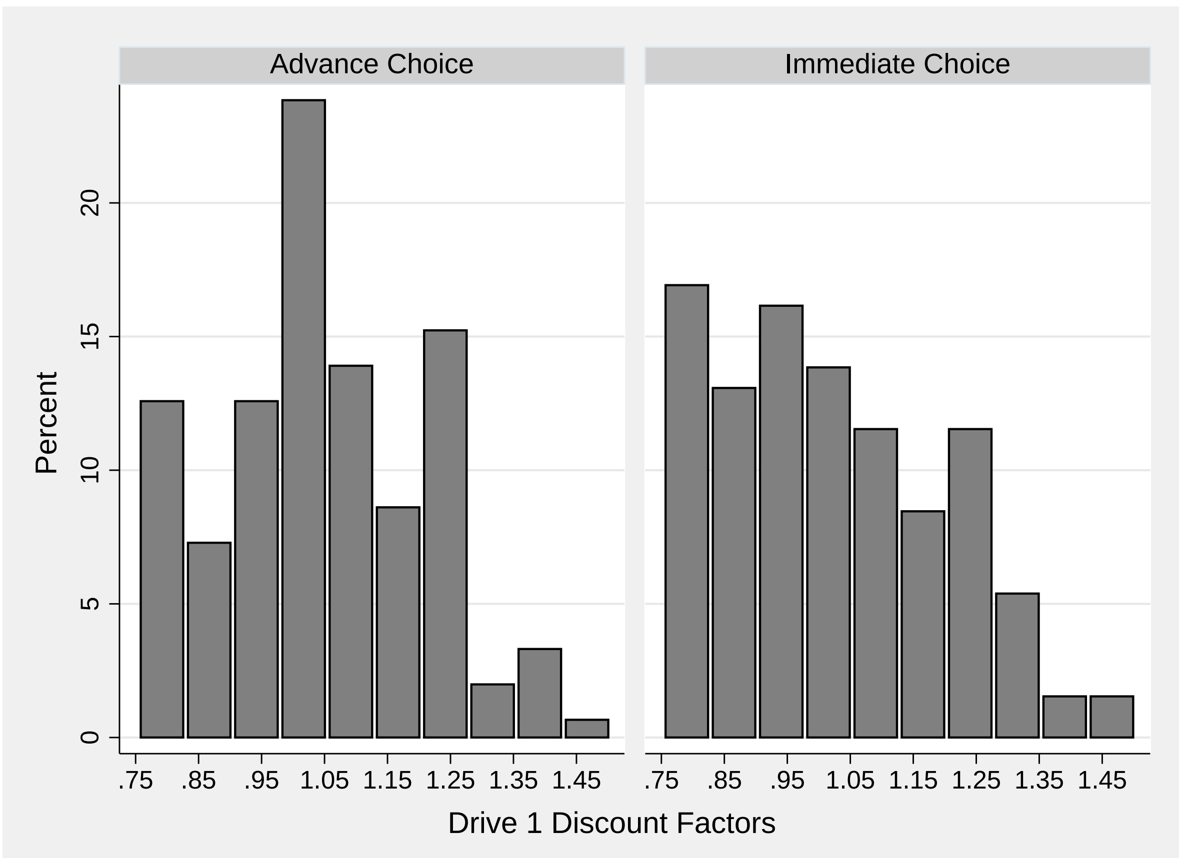

Figure 4: Individual Discount Factors

heterogeneity, there is no scope for customizing incentives.

3.2 Tailored Contracts

Individual discount factors from Drive 1 in hand, we turn to the possibility of tailoring intertemporal contracts to individual preferences. Of the 281 LHWs for whom we have individual discount factors in the range of 0.75 to 1.5 from Drive 1, 280 participated in Drive 2.[17] Of these 142 LHWs were assigned a value of R equal to their discount factor. That is, tailored LHWs were assigned $R^* = \beta^{\mathbf{1}_{d=1}}\delta$, which should induce equal provision of effort through time, $v_1 = v_2$. The remaining 138 LHWs serve as control and were assigned a uniform random interest rate $\tilde{R} \in U[0.75, 1.5]$.[18]

Panel A of Figure 5 plots the individual discount factor measured during Drive 1 against the assigned R in Drive 2 separately for the tailored and untailored groups. In the left panel, by design, there is no clear relationship between discount factors and the assigned interest rates. In the right panel, there is a strict one-to-one mapping along the 45 degree line between discount factors and the Drive 2 interest rates, consistent with the assignment $R^* = \beta^{\mathbf{1}_{t=1}}\delta$ in the tailored group.

Panel B of Figure 5 plots vaccinations allocated to the first day of the drive against vaccinations allocated to the second day of the drive separately for the tailored and the untailored group. Notable from Figure 5 is the relative dispersion of the untailored controls around the 45-degree line of equal provision relative to the tailored treatments.

We examine differences in the distance from the 45-degree line of equal provision using the metric $|\frac{v_1}{v_2} - 1|$. The mean distance for the untailored group is 0.61 (s.d. $= 3.64$) while the mean distance for the tailored group is 0.14 (s.d. $= 0.23$), $t_{278} = 1.53$, $(p = 0.13)$. The lack of statistical significance is due primarily to several substantial distance outliers. Trimming the

[17] LHWs from the boundary sample were allowed to participate in Drive 2 and were either assigned $R \in U[0.75, 1.5]$ if they were in the untailored control group (31 subjects) or assigned $R = 0.75$ or $R = 1.5$ if they were in the tailored group and had $R^* < 0.75$ (15 subjects) or $R^* > 1.5$ (11 subjects). See section 4.2 for analysis of the boundary sample.

[18] 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.

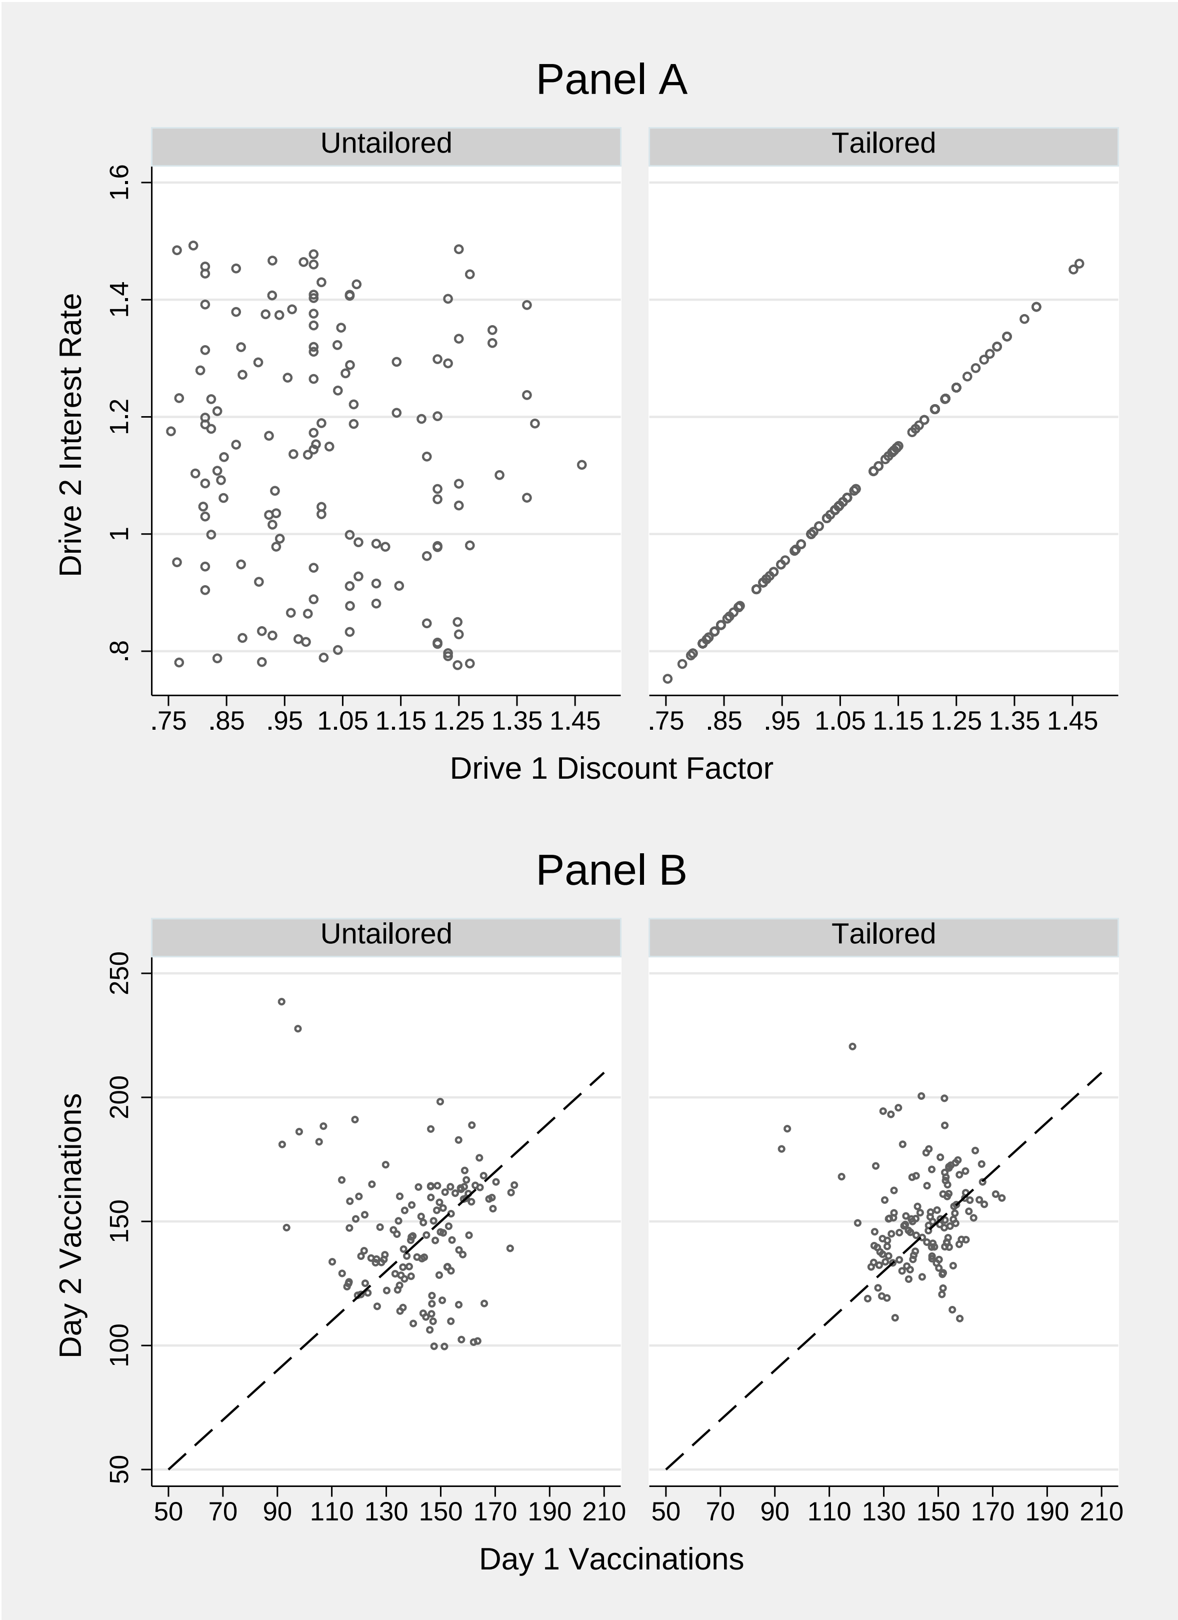

Figure 5: Discounting and Tailoring

top and bottom 1% of the sample in terms of Drive 2 allocations, the mean distance for the untailored group is 0.15 (s.d. $= 0.19$), while the mean distance for the tailored group is 0.10 (s.d. $= 0.11$), $t_{265} = 3.10$, $(p < 0.01)$.

In Table 4, we provide corresponding regression analysis. Tobit specifications are provided because 26 of 280 LHWs have allocations with v_1 exactly equal to v_2 and hence distance measures of zero. Following best practice for such analysis (Bruhn and McKenzie, 2009), we control for fixed effects for each stratum in the stratified randomization. In column (1), we analyze all 280 subjects and note a sizable reduction in distance under tailoring that falls just outside the range of significance. Echoing our raw results, when excluding outliers in column (2), we find that tailoring serves to reduce distance from equal provision significantly by around six percentage points. Relative to the untailored controls, tailoring reduces distance from equal provision by around one-third, indicating substantial benefits to our tailored policy initiative. In column (3), we additionally control for the value of R assigned in Drive 2. This regression identifies whether tailoring generates more equal provision for a given R, and hence controls for any differences in interest rates across tailored and untailored groups. Again, tailoring serves to reduce distance significantly.

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 columns (3)-(6) of Table 4, we examine differential effects of tailoring across these two groups. Given that the individual discount factors skew lower in Immediate choice, one might expect larger distance measures in Immediate controls (and hence greater benefits to tailoring). This is precisely what is observed. Untailored Immediate choice is associated with significantly larger distance measures and tailoring for Immediate choice significantly reduces these distances. In columns (5) and (6), excluding outliers, we find that tailoring in Immediate choice reduces distance from equal provision by around one-half. Note that this effect size (8.7 percentage points) is similar to the effect of moving an LHW from advance to immediate choice in the untailored group (8 percentage points). Tailoring in Advance choice appears to directionally reduce distance as well, but the effect is not significant, potentially due to the relatively small average distance measure identified in untailored Advance choice.

The findings of Table 4 indicate substantial benefits to tailored policy intitiatives. For a given interest rate, matching this interest rate to individual preferences generates significantly smoother service provision. The effects can be substantial, reducing distance measures by around one-third on average and around one-half in Immediate choice. Policy makers wishing to smooth service provision through time may well better achieve their objective by tailoring contracts to preferences.

Table 4: The Effect of Tailoring Intertemporal Incentives

Dependent variable:	$\lvert\frac{v_1}{v_2} - 1\rvert$					
	(1)	(2)	(3)	(4)	(5)	(6)
Tailored (=1)	-0.544	-0.063***	-0.052**	-0.086	-0.020	-0.015
	(0.352)	(0.021)	(0.020)	(0.142)	(0.021)	(0.021)
Immediate Choice (=1)				1.089*	0.130***	0.123***
				(0.616)	(0.037)	(0.036)
Tailored x Immediate				-0.953	-0.093**	-0.087**
				(0.643)	(0.042)	(0.042)
Constant	1.408	0.159***	0.025	0.809	0.090***	-0.003
	(0.859)	(0.023)	(0.061)	(0.534)	(0.028)	(0.060)
Stratum FEs	Yes	Yes	Yes	Yes	Yes	Yes
Exclude 99th and 1st Percentiles	No	Yes	Yes	No	Yes	Yes
Drive 2 R	No	No	Yes	No	No	Yes
Log-Likelihood	-625.629	69.504	72.117	-622.792	80.722	82.204
Mean in Untailored Contract	0.612	0.153	0.153	0.612	0.153	0.153
Mean in Untailored Advance				0.089	0.089	0.089
Mean in Untailored Immediate				0.701	0.169	0.169
# Vaccinators	280	267	267	280	267	267

Notes: This table reports the effects of tailoring on the equality of effort provision over time. The measure $\lvert\frac{v_t}{v_{t+1}} - 1\rvert$ (the percentage difference between tasks allocated to day 1 and day 2 of the drive) reflects the distance of the task allocation (v_1,v_2) from equality ($v_1 = v_2$). Column (1) reports a regression of this measure on a dummy equal to one for subjects in the tailored group. Column (2) reports estimates from the same specification excluding outliers. Column (3) controls for the interest rate assignment in round 2. Column (4) provides estimates on the same sample as column (1) interacting treatment with being in the immediate choice condition. Columns (4) and (6) apply the same restrictions to the sample as columns (2) and (3) respectively. Heteroskedasticity robust White standard errors reported in parentheses. *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$.

4 Robustness Tests and Additional Exercises

The analysis to here indicates three key findings. First, there appears to be a present bias in LHW allocation behavior. Those individuals making Immediate choice allocate fewer vaccinations to v_1 than those making Advance choice. Second, despite the general tendency towards less patience in Immediate choice, substantial heterogeneity in discounting is observed. Both of these effects resonate with prior experimental findings and highlight the potential for policy interventions tailored to individual preferences. Third, tailored contracts work. Those individuals given a tailored interest rate equal to their previously measured discount factor provide smoother service than untailored controls. In the following sub-sections, we explore robustness to a set of plausible alternative interpretations and provide a set of natural additional examinations.

4.1 Identification of Time Preferences and Present Bias

Our data show a general tendency towards greater impatience when choice is Immediate. Corresponding parameter estimates at the aggregate level and parameter calculations at the individual level reflect these patterns with present bias parameters that are less than one. We first examine the robustness of these results to estimation using only within-subject variation from our panel of subjects that participated in Drives 0, and 1, and changed from Advance to Immediate choice. We then examine the validity of a set of structural assumptions required to infer discounting parameters from LHW allocation behavior.

4.1.1 Within-Subject Variation

Using our failed Drive 0 data together with Drive 1, we identify a present bias both in allocation behavior and in corresponding parameter estimates. However, when we examine Drive 1 alone, though the median data are suggestive of present bias and the aggregate estimate of β is less than one, we cannot reject $\beta = 1$ in any specification. This suggests some potential sensitivity in our findings of present bias. Importantly, the variation identifying present bias from Drive

1 alone is exclusively between subjects. Given the wide heterogeneity in observed patience regardless of decision timing, one may fail to statistically identify present bias even if it exists on average. Indeed, most studies of present bias and dynamic inconsistency are conducted as within-subject exercises, 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 within-subject for LHWs who changed from Advance to Immediate choice across drives. In Table A.2, we re-conduct the analysis of Table 3, column (4), focusing on the panel of 232 LHWs who participated in both Drive 0 and Drive 1. Column (1) reproduces Table 3, column (4) indicating substantial present bias for our 232 panel LHWs. In column (2), we present results for 106 subjects who did not change from Advance to Immediate, identifying present bias only from between-subject variation. Though $\beta < 1$ is estimated, we cannot reject $\beta = 1$ using only between-subject variation. In column (3) we rely only on within-subject variation from the 126 LHWs who change from Advance to Immediate choice across Drives. Using such within-subject variation we reject the null hypothesis of no present bias at the 5% level.

In columns (4) and (5) of Table A.2, we break the sample into those who transition from Advance to Immediate and vice versa. Estimates for both sub-samples indicate $\beta < 1$. That estimated values of β lie below one regardless of whether a subject moves from Advance to Immediate or from Immediate to Advance through time is an important finding. One critique of our between-subject design is that all subjects in Immediate choice made their decisions on the same day. Some idiosyncratic shock on Day 1 of Drive 1 could lead to present bias. These findings indicate that if something specific is occurring on Day 1 of Drive 1, a similar shock must occur several weeks previously on Day 1 of Drive 0.

The results of Table A.2 demonstrate that when relying on within-subject variation, present bias is substantial at the aggregate level. The 126 LHWs that change from Advance to Immediate choice also provide an opportunity to investigate present bias at the individual level. Following equation (3), we calculate a discount factor for each condition the LHW faces. The parameter δ is identified as the discount factor from Advance choice while β is identified as the discount factor from Immediate choice divided by that of Advance choice. Figure A.6 presents histograms of δ and β for these 126 individuals.[19] Two features of Figure A.6 warrant attention. First, as in our analysis of discount factors in Advance choice, δ is centered around 1 with a median value of 1.04. Second, β is skewed below 1 with a median value of 0.95. Sixty-nine (54.8%) of 126 LHWs have $\beta < 1$, 6 (4.8%) have $\beta = 1$ and 51 (40.5%) have $\beta > 1$. A sign test for the null hypothesis that the median β is equal to 1 yields a p-value of 0.12 (two-sided test).[20] Together these results show that general patterns of present bias are observed at the aggregate and individual level when investigating only within-subject variation in a sub-sample of 126 LHWs who change from Advance to Immediate choice across conditions.

4.1.2 Structural Assumptions

As in any structural exercise, a set of assumptions are required to infer discounting parameters from LHW allocation behavior. Six assumptions are relevant for the present discussion. First, we assume a stationary cost function, constant across periods. Second, we assume that the only argument of the cost function is allocated vaccinations, ignoring idiosyncratic costs across time or individual. Third, we assume a homogeneous cost function, constant across individuals, governed by exponential parameter $\gamma = 2$. Fourth, we assume a deterministic environment with individuals being able to correctly and precisely forecast costs through time. Fifth, we assume that allocations provide insight into vaccination preferences without the possibility that subjects renege on their stated allocations. Sixth, we assume that the decision environment creates no biases in choice, such that preferences are directly revealed by allocation behavior. Violation of any of these assumptions would lead to misspecification and could lead to the misattribution of dynamically consistent preferences as present-biased. Each of these assumptions is investigated in turn, and, where appropriate, appeal is made to additional data sources.

[19] One LHW with a calculated β in excess of 19 is removed from the figure to ease reading

[20] For the one-sided test with an alternative of $\beta < 1$, the p-value is 0.06. Excluding the single subject with β in excess of 19 reduces the two-sided (one-sided) p-value to 0.10 (0.05).

Stationarity of the Cost Function: Our exercise assumes stationary costs. The extent to which LHWs deviate from stationarity may influence their measured preferences. 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 measured by β is conflated with non-stationarity.

Importantly, our monitoring technology provides time-stamps and geo-stamps for vaccination activity. Time stamps are recorded 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.[21] Of our 338 LHWs, measures for median time lapse between vaccination attempts are available for 277 on either Day 1 or Day 2 and for 228 LHWs on both days of Drive 1.[22] Of our 338 LHWs, measures for median distance traveled each 15 minutes are available for 274 on either Day 1 or Day 2 and for 224 LHWs on both days of Drive 1.[23]

[21]We focus only on the distance traveled and time taken for vaccinations between 8 am and 6pm each day. 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.

[22]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.

[23]260 LHWs have Day 1 distance data while 238 have Day 2 distance data. Of the 78 LHWs with missing Day 1 data, 72 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 100 LHWs with missing Day 2 data, 96 of them completed four or fewer vaccination attempts on Day 2 and the remaining 4 did not have phones that interacted with the server to report location or had faulty GPS.

Appendix Table A.3 demonstrates that 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 with Advance choice and uncorrelated with discount factors within condition. Time and distance are also uncorrelated with Advance choice and discount factors 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 discount factors within condition. These data indicate stability in required average effort per vaccination 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 patience or our finding of present bias.[24]

Appendix Table A.3 explores stationarity across Day 1 and Day 2 indirectly by examining time use. Another violation of stationarity would be if Advance and Immediate choices were governed by different, but constant, cost functions. This question can be investigated directly in the data. Because we have experimental variation in interest rates in both Advance and Immediate choice, the cost parameter, γ, can be estimated for both conditions. In Appendix Table A.4 we estimate a single discount factor and γ for both Advance and Immediate choice. While costs are estimated to be similarly convex in both Advance and Immediate choice, discount factors vary in an economically meaningful way. Furthermore, while we fail to reject that the discount factor is equal to one in Advance choice, we reject the null hypothesis of one in Immediate choice. Appendix Table A.4 indicates that non-stationarity in discount factors is more apparent than non-stationarity in costs.

Unobserved Idiosyncratic Costs: We assume that vaccinations are the only argument of costs when identifying time preferences. However, there may be idiosyncratic costs across time or

[24]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 vaccinators, and received a half day's training on the vaccination monitoring application.

individuals that could influence measured patience. For example, an LHW with an appointment lasting 2 hours on Day 1 and no appointment on Day 2 may find it extremely costly to allocate vaccinations to Day 1. This may appear to the researcher as impatience, but only reflects the LHW's idiosyncratic costs across days. Further, if such idiosyncratic events are easier to re-organize when making Advance choice, present bias may be conflated with ease of scheduling.

Here, again, the additional data on LHW time use available from the monitoring application is potentially valuable. We can investigate whether extended periods of non-vaccination exist and are correlated with measured preferences and allocation timing. As in the example above, an LHW with an extended period of non-vaccination may well be experiencing forecasted idiosyncratic costs unrelated to vaccinations. Appendix Table A.5 repeats the analysis from Appendix Table A.3, with dependent variables of the maximum daily time lapse between vaccination attempts and whether the longest daily break is in excess of two hours. Longest daily breaks are on average around 59 minutes on Day 1 with around 13% of LHWs taking longest breaks in excess of 2 hours. Focusing on individuals with measures on both days of the drive, we find that the length of longest breaks and the probability of 2 hour breaks are uncorrelated with Advance choice and uncorrelated with discount factors within condition. Almost identical patterns are observed on Day 2 of the drive. Differences in break behavior across days are statistically indistinguishable from zero, uncorrelated with allocation timing, and uncorrelated with discount factors within condition. These data suggest that idiosyncratic costs identified from taking extended breaks do not explain the extent of impatience in the sample, and that potential difficulties in rescheduling do not explain observed present bias.

Deterministic Environment: Our exercise assumes that individuals can perfectly forecast their own future costs of vaccination. A plausible alternative is that LHWs don't precisely know how costly vaccinations will be. Taken to a logical conclusion, one may expect LHWs in Advance choice to have less information than LHWs in Immediate choice. The natural evolution of uncertainty through time may lead to differences in measured preference parameters across groups. Though the resolution of uncertainty may lead to apparent dynamic inconsistency, the direction is not clear. Some LHWs may grow more patient as uncertainty is resolved, some less so.[25]

One potential, albeit imperfect, test for the effect of uncertainty in the environment is to examine differences in experience across individuals.[26] More experienced LHWs may have a more complete understanding of the difficulty of tasks and may face less uncertainty when making allocation decisions. As noted before, LHWs have an average of more than 10 years of experience in health work. In Table A.6, we separate our sample either at the median of experience (9 years) or at 15 years of experience.[27] Separating at the median, we find virtually no differences in estimated preferences. Separating at 15 years of experience, we find that, if anything, more experienced subjects are somewhat more present-biased than their less experienced counterparts. This indicates that experienced individuals who are likely to have the least surprises in costs through time neither yield dramatically different estimates of patience, nor are disproportionately likely to be dynamically consistent.

Homogeneity of Cost Function: Our aggregate exercise assumes homogeneous costs constant across subjects, and our individual elicitation assumes homogeneous quadratic costs. 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 allocation timing with differences in costs. One natural view would be to assume that individuals do not discount at all, $\delta = 1$ and $\beta = 1$, 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 $v_1 = v_2 = 150$ for all values of γ.[28] Examining the Drive 0 and Drive 1 data, we find that for 163 LHWs who were assigned $R = 1$, the mean allocation is $v_1 = 140.84$ ($s.d. = 24.76$).[29]

[25] Naturally, if shocks to costs do underly our observed differences in patience across individuals, one might not expect to be able to tailor contracts at the individual level over time with the success that we have.

[26] This test is imperfect as experience may correlate with other differences across LHWs.

[27] Experience measures are available for 329 of 338 LHWs in Drive 1.

[28] This is because the Euler equation reduces to $(\frac{v_1}{v_2})^\gamma = R = 1$, which implies $\frac{v_1}{v_2} = 1$.

[29] 42 of 163 LHWs allocated exactly $v_1 = v_2 = 150$.

Though the median allocation is indeed 150, responses range widely with 5th-95th percentiles of response being 103 to 160. If heterogeneity in costs were driving response and discounting was not a key feature 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 does not easily rationalize the observed present bias in the data.

One additional point that warrants attention is that whenever we estimate the homogeneous cost function parameter, γ, substantial convexity in costs is estimated. Indeed, even with the relatively extreme curvature estimates, our estimated models predict more sensitivity to changing interest rates than exist in the data. We wondered whether one could identify subsets of LHWs for whom behavior was more sensitive to price and hence curvature estimates were less convex. In Table A.8, we show that a subset of subjects, in particular those with more than 15 years of experience that are working in an urban environment, have substantially less convex cost functions relative to their less experienced or non-urban counterparts.[30] Albeit exploratory, such analysis is comforting as it suggests that LHWs for whom understanding of vaccinations and the ease with which they can move them through time are high, are indeed more sensitive to price changes and hence have less convex estimated costs. Though such a finding helps to understand the relative lack of price sensitivity in the broad sample[31], it also indicates a potential misspecification in costs. Not only are costs likely to be heterogeneous, they are also likely to be more convex than our quadratic assumption for large swathes of our sample. Viewed in this light our relative success at tailoring contracts assuming quadratic costs and ignoring heterogeneity should likely be viewed as a lower bound on the promise of such exercises.

Inability to Renege: The contracts we implement to elicit preferences feature a completion

[30] We identify a vaccinator as working in an urban environment if the average distance traveled per vaccination attempt over the two days is less than 0.1 miles. Average distance traveled is the ratio of total distance traveled over the two days of the drive divided by the total number of vaccination attempts during the drive. We determine a vaccinators position in space every 15 minutes while they are working as the average GPS coordinate recorded by their smartphone over that 15 minute interval. We calculate the total distance traveled as the sum of distances traveled between these average coordinates.

[31] That is, it may just be particularly difficult to move tasks through time for relatively inexperience and non-urban subjects.

bonus of 1000 rupees if both targets, v_1 and v_2, are met. The bonus is sizable relative to daily earnings of 100 rupees prior to our study. The design choice of such a large bonus was made purposefully to prevent LHWs from reneging on their contract choices. At any given point in time, not completing allocated vaccinations generates a 1000 rupee penalty. As such, LHWs should forecast that they will indeed complete required vaccinations and allocate according to their preferences. If LHWs forecast not being able to complete the required vaccinations, then allocation choices may not reflect their true preferences, and hence be effectively cheap talk.[32]

Our monitoring application allows us to assess, ex-post, the extent to which LHWs succeed or fail to reach their allocated targets. We have comprehensive completion data for 288 of 338 LHWs in Drive 1.[33] Of these 288, 142 (49.3%) met or exceeded their targets of v_1 and v_2. The data indicate that many who fall short come close to completion. Two-hundred seven LHWs (71.9%) met or exceeded their targets on one of the two days and an average of 75.5% of total target vaccinations were completed. Figure A.7 presents the histogram of average completion proportions across subjects.[34] Two hundred and two LHWs (70.1%) completed an average of 70% or more of their target vaccinations. These ex-post completion rates indicate that a majority of LHWs likely forecasted reaching their targets.

In Appendix Table A.7, we investigate the relationship between allocation timing, measured patience and completion rates. Two general patterns arise. First, LHWs making advance choice are somewhat more likely to complete their targets; and, second, higher measured discount factors are related to higher completion rates, particularly for the second day of the drive. The observed correlation is consistent with the view that impatient LHWs delay vaccinations believing they will complete them but subsequently fail to follow through. Interestingly, the correlation between completion on both days and measured patience is somewhat stronger in Immediate choice ($\rho = 0.24,\ p < 0.01$) relative to Advance choice ($\rho = 0.10,\ p = 0.21$), suggesting that some failure to complete is linked to delay associated

[32]Sizable bonuses to limit the possibility of reneging are also implemented in Augenblick et al. (2015).

[33]The remaining 50 LHWs did not activate did their application or submitted zero vaccinations on both days of the drive.

[34]Average completion proportions are calculated as $1/2(min(Completed_1/v_1, 1) + min(Completed_2/v_2, 1))$

with present bias. For such patterns to exist without choices being informative of preference, some aspect of the decision environment would have to lead uninformative responses to appear less patient and disproportionately so in Immediate choice.

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 $v_1 = v_2$, or choosing an allocation in the middle of their budget constraint, $v_1 = Rv_2$. Only 31 of 338 LHWs (9%) exhibit $v_1 = v_2$. Taking a less conservative measure of $v_2 - 2.5 \leq v_1 \leq v_2 + 2.5$, we find that still only 58 of 338 LHWs (17%) are within 5 vaccinations of $v_1 = v_2$.[35] Only 35 of 338 LHWs (10.3%) exhibit $v_1 = Rv_2$. Taking a less conservative measure of $Rv_2 - 2.5 \leq v_1 \leq Rv_2 + 2.5$, we find that 83 of 338 LHWs (25%) are within 5 vaccinations of $v_1 = Rv_2$.[36] Taken together, these 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 contracts tailored to individual preferences can yield policy benefits. LHWs who are given a value of R equal to their discount factor provide significantly smoother service. Here we examine robustness of this result to alternative measures for smoothness in service provision and potential issues related to elicitation of preferences. A set of additional exercises are also provided to assess the possibility for alternative interventions based on different policy preferences.

[35] As an even less conservative measure, 145 of 338 (43%) satisfy $v_2 - 10 \leq v_1 \leq v_2 + 10$.
[36] As an even less conservative measure, 137 of 338 (40.5%) satsify $Rv_2 - 10 \leq v_1 \leq Rv_2 + 10$.

4.2.1 Alternative Measures for Smooth Provision

Our analysis measures the distance to equal provision using the metric $|\frac{v_1}{v_2} - 1|$. In Table A.9, we reconduct the analysis of Table 4, using five alternate measures for smoothness. Panel A presents the Euclidean distance to the 45 degree line, $\frac{|v_1 - v_2|}{\sqrt{2}}$. Panel B presents the Euclidean distance normalized by the total number of vaccinations allocated, $\frac{|v_1 - v_2|}{\sqrt{2}(v_1+v_2)}$. Panel C presents the number of sooner vaccinations that would need to be reallocated to reach the 45 degree line, $|v_1 - \frac{300}{1+R}|$. Panel D presents probit regressions for needing to reallocate 10 or fewer vaccinations, $|v_1 - \frac{300}{1+R}| \leq 10$. And finally, Panel E presents probit regressions for needing to reallocate more than 25 vaccinations, $|v_1 - \frac{300}{1+R}| > 25$. Across all specifications, the main conclusions are reproduced. However, the results with respect to additional tailoring benefits in Immediate choice fall at times outside the range of statistical significance. These alternative measures of smooth provision indicate that our results on the potential benefits of tailoring are not an artifact of how one measures the outcome of interest.

4.2.2 Alternative Sample Restrictions and Treatment Measures

Our tailoring exercise focused on LHWs with discount factors between 0.75 and 1.5. Of 337 LHWs in Drive 2, 280 satisfied this requirement. Those LHWs whose discount factors fell outside of this range were given either $R = 0.75$ or 1.5 depending on which bound they were closest to. For such individuals tailoring isn't a binary treatment, but rather a continuous difference between their discount factor and the exogenously given one. Indeed, for all LHWs in the untailored group treatment is also a continuous measure. In Table A.10, Panel A we reconduct the analysis of Table 4 using as the measure of treatment the absolute difference between each LHW's discount factor and their assigned interest rate, *Tailor Intensity.* The main results are reproduced that the closer discount factors are to interest rates, the smoother is provision.[37] In Panel B we include those individuals in the boundary sample with discount factors that lie outside of the bounds of interest rate assignment. Including these individuals does not alter the

[37] Restricting attention only to the untailored group reveals directionally similar, though insignificant results across all specifications.

conclusions, however, it should be noted that treatment is no longer orthogonal to individual preferences as extremely patient and impatient LHWs will receive larger treatment intensity on average.[38]

4.2.3 Alternative Policy Preferences

While our results suggest tailored contracts can improve success in achieving a Leontief policy objective, a natural question is whether this approach could be put to use to achieve other policy objectives. Ultimately, any attempt to tailor contracts will rely on whether initially elicited preferences are stable. If LHW time preferences are stable, then changes in incentives will have predictable effects on behavior.

To provide a general assessment of the promise of alternative policy objectives, we examine the stability of identified discount factors across Drives 1 and 2 by calculating

$$\frac{R \cdot v_1}{v_2} = \beta^{\mathbf{1}_{d=1}}\delta,$$

for each LHW in each Drive. Figure A.8 presents the calculated discount factors for Drive 1 and Drive 2 along with 45 degree line for 317 of 337 LHWs.[39] The correlation in discount factors across rounds is $\rho = 0.41,\ (p < 0.01)$, indicating stability in preferences.[40] Our findings correspond with those of Meier and Sprenger (2015), who investigate subjects participating in an identical monetary discounting experiment approximately one year apart and identify a one-year correlation of around 0.5 for monetary choices. The level of correlation in discount factors across drives indicates stability in preferences such that alternative policy objectives

[38]Using the indicator for tailoring would not be an appropriate solution to this problem as tailored LHWs with extreme patience or impatience may actually receive interest rates that are further from their policy-optimal interest rates than those in the untailored condition.

[39]Eliminated from the figure and from our calculations of stability are twenty LHWs with discount factors in excess of 2 in one or both drives.

[40]Including the remaining twenty-one extreme LHWs, the correlation changes substantially to $\rho = 0.01,\ (p = 0.87)$. It should be noted that the correlation in identified discount factors is substantially higher in the tailored condition, $\rho = 0.67,\ (p < 0.01)$, relative to the untailored condition, $\rho = 0.17,\ (p < 0.05)$. We believe this is due to some sensitivity of behavior to extreme values of R in Drive 2. For untailored subjects who coincidentally receive a value of R within 0.25 of their value of R^*, the correlation in discount factors is $\rho = 0.53,\ (p < 0.01)$.

may also be achievable with tailored contracts.[41]

5 Conclusion

This paper seeks to understand whether the point predictions given by structural models of time preferences are empirically valid. In the context of polio vaccination in Pakistan, we examine whether preferences elicited from intertemporal bonus contract choice can be used subsequently to effectively tailor contracts. We document three primary results. First, a present bias in allocation behavior exists. Second, substantial heterogeneity in elicited preferences is observed. Third, and most importantly, corresponding point predictions are indeed empirically valid. Vaccinators who receive tailored contracts are one-third closer to the policy objective relative to their untailored counterparts.

To date, little research makes use of the predictive value of structural discounting estimates. Our results show not only that estimates are predictive, but also that valuable preference estimates are identifiable from a very limited number of experimental choices. This suggests that the substantial effort of articulating and estimating structural models in this domain has been well-invested.

There a number of clear limitations to our study which should be addressed by future research. First, and most importantly, 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 tailored contracts, it is important to know how much more can be achieved. Second, 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 contracts, multi-period settings, and alternative policy targets that are worthy of study. Third, our study sidesteps the critical issue

[41] One natural alternative is to maximize performance, regardless of timing. In such a case, we consider a policymaker with linear preferences, $P(v_1, v_2) = v_1 + v_2$, who wishes to maximize the total number of completed vaccinations regardless of timing. Maximizing this objective function subject to the vaccinator's offer curve, yields an optimal $R^*_{max} = \sqrt{\beta^{\mathbf{1}_{t=1}}\delta(1 + \beta^{\mathbf{1}_{t=1}}\delta)} - \beta^{\mathbf{1}_{t=1}}\delta$. Unfortunately, our assigned values of R are generally quite far from R^*_{max} making it difficult to test for the possibility of a maximizing contract.

of incentive compatibility by not informing subjects of Drive 2 when Drive 1 preferences are elicited. The mechanism design problem of eliciting preferences and tailoring on said preferences with complete information will be critical if one wishes to implement tailored contracts repeatedly in the field.

References

Andersen, Steffen, Glenn W. Harrison, Morten I. Lau, and Elisabet E. Rutstrom, "Eliciting Risk and Time Preferences," *Econometrica*, 2008, *76* (3), 583–618.

Andreoni, James and Charles Sprenger, "Estimating Time Preferences with Convex Budgets," *American Economic Review*, 2012, *102* (7), 3333–3356.

_ and _ , "Risk Preferences Are Not Time Preferences," *American Economic Review*, 2012, *102* (7), 3357–3376.

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.

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.

Banerjee, Abhijit and Esther Duflo, "Addressing Absence," *The Journal of Economic Perspectives*, 2006, *20* (1), 117–132.

Banerjee, Abhijit V., Esther Duflo, and Rachel Glennerster, "Putting a Band-Aid on a Corpse: Incentives for Nurses in the Indian Public Health Care System," *Journal of the European Economic Association*, 04-05 2008, *6* (2-3), 487–500.

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.

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, and Yasir Khan, "The Political Economy of Public Employee Absence: Experimental Evidence from Pakistan," 2015.

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.

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.

Chaudhury, Nazmul, Jeffrey Hammer, Michael Kremer, Karthik Muralidharan, and F. Halsey Rogers, "Missing in Action: Teacher and Health Worker Absence in Developing Countries," Winter 2006, *20* (1).

Coller, Maribeth and Melonie B. Williams, "Eliciting individual discount rates," *Experimental Economics*, 1999, *2*, 107–127.

Cubitt, Robin P. and Daniel Read, "Can Intertemporal Choice Experiments Elicit Preferences for Consumption," *Experimental Economics*, 2007, *10* (4), 369–389.

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde, "Dynamic inconsistency predicts self-control problems in humans," *Working Paper*, 2006.

_ , _ , _ , and _ , "Are Risk Aversion and Impatience Related to Cognitive Ability," *American Economic Review*, 2010, *120* (54), 256–271.

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

Frederick, Shane, George Loewenstein, and Ted O'Donoghue, "Time discounting and time preference: A critical review," *Journal of Economic Literature*, 2002, *40* (2), 351–401.

Gine, Xavier, Jessica Goldberg, Dan Silverman, and Dean Yang, "Revising Commitments: Time Preference and Time-Inconsistency in the Field," *Working Paper*, 2010.

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.

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.

Kirby, Kris N., Nancy M. Petry, and Warren K. Bickel, "Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls," *Journal of Experimental Psychology: General*, 1999, *128*, 78–87.

Koopmans, Tjalling C., "Stationary Ordinal Utility and Impatience," *Econometrica*, 1960, *28* (2), 287–309.

Kuhn, Michael A, "Curing the Calorie Crunch: The Effect of EBT on Household Present-Bias," 2013.

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.

Lee, David S., "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects."," *Review of Economic Studies*, 2009, *76* (3), 1071–1102.

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.

_ and _ , "Temporal Stability of Time Preferences," *Review of Economics and Statistics*, 2015, *97* (2), 273–286.

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.

Parente, Paulo and João Santos Silva, "Quantile Regression with Clustered Data," *Journal of Econometric Methods*, 2016, *5* (1), 1–15.

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 M, "Is there a daily discount rate? Evidence from the food stamp nutrition cycle," *Journal of Public Economics*, 2005, *89* (2), 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.

Tanaka, Tomomi, Colin Camerer, and Quang Nguyen, "Risk and time preferences: Experimental and household data from Vietnam," *American Economic Review*, 2010, *100* (1), 557–571.

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.

A Appendix

A.1 Appendix Figures

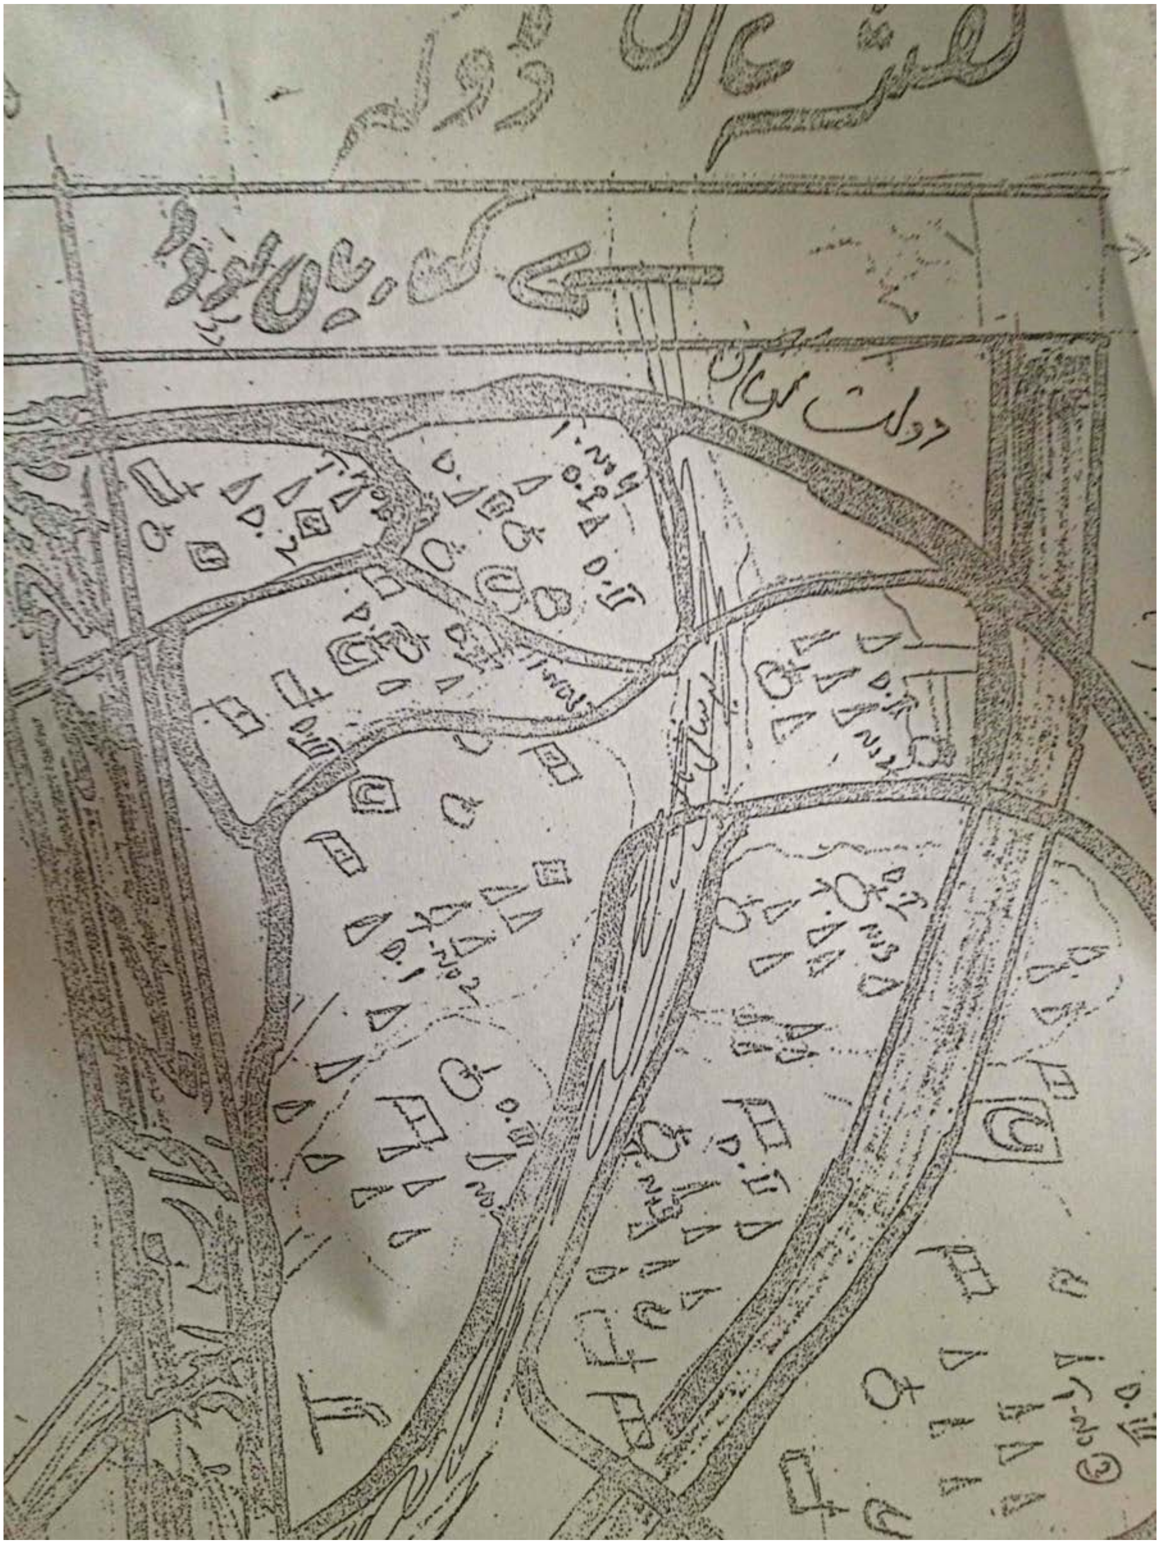

Figure A.1: Map Given to Vaccinators to Plan Route

Notes:

Figure A.2: Picture of a Door-to-Door Vaccination During a Drive

Notes:

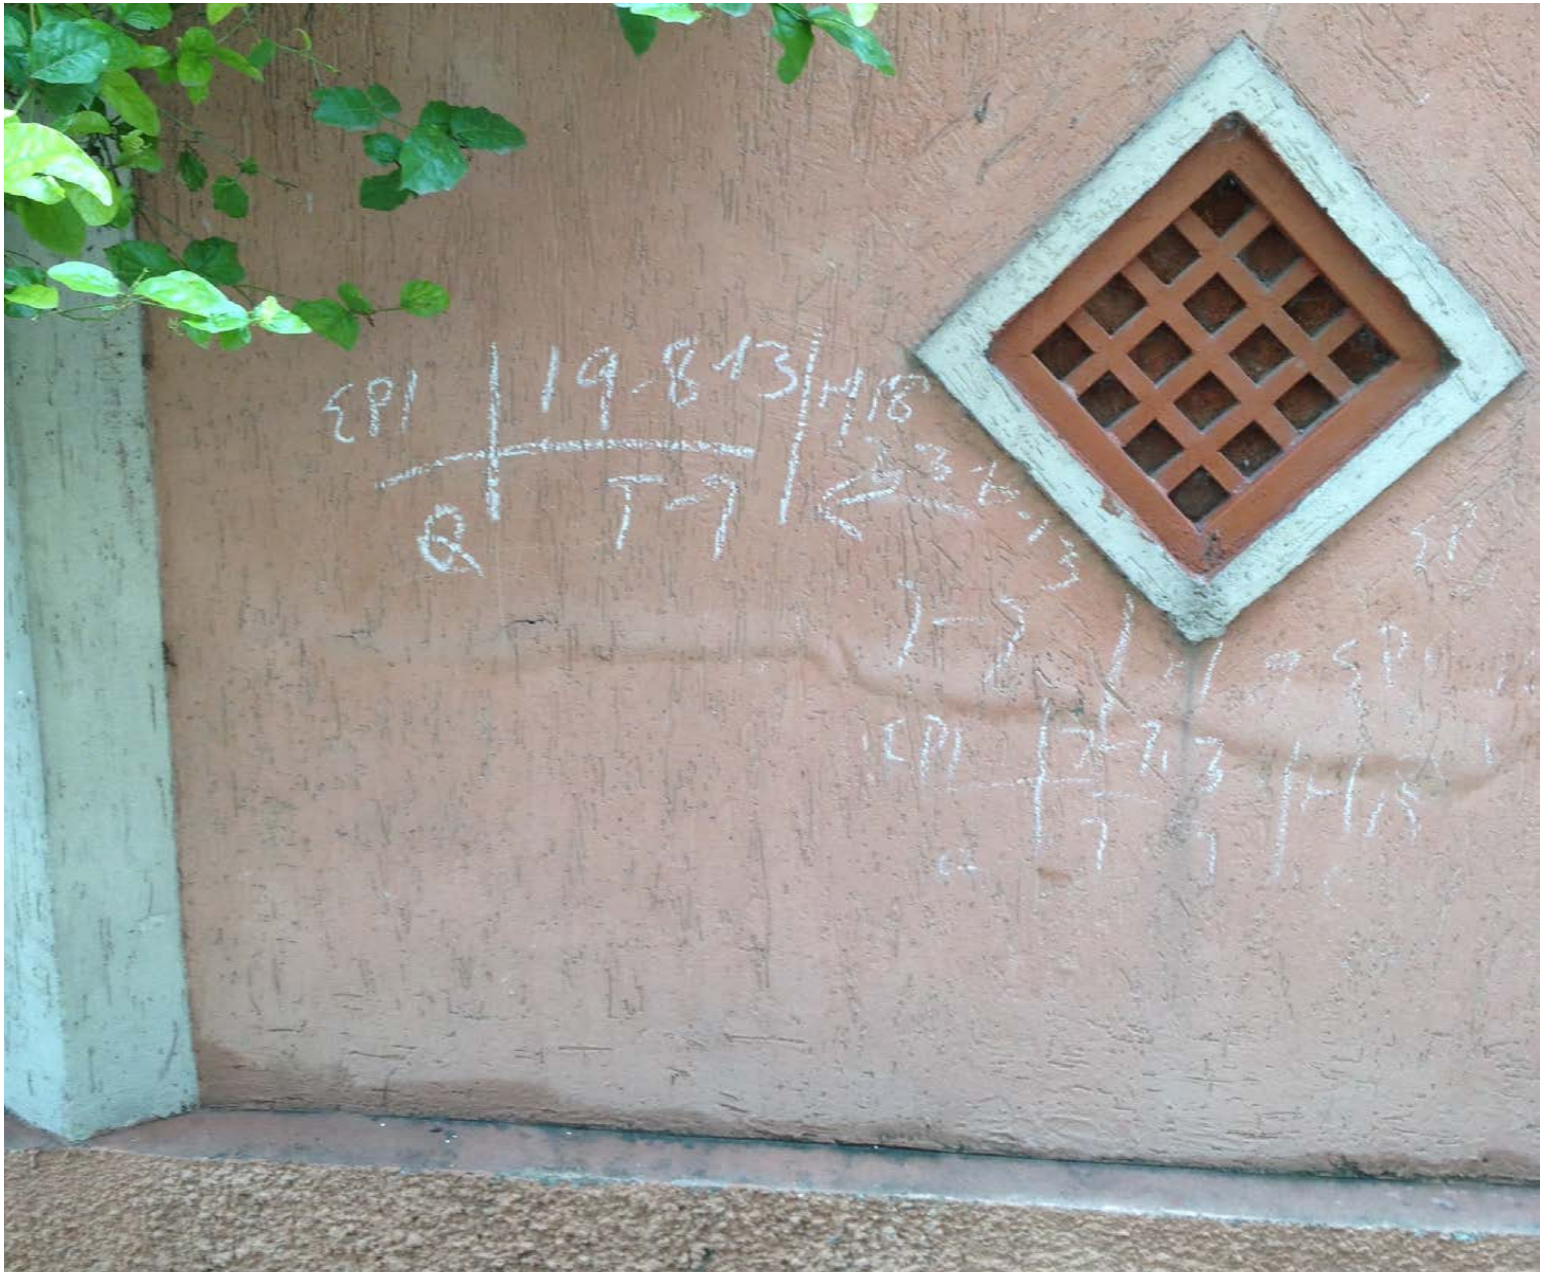

Figure A.3: Chalk Marking to Record Visit by Vaccination Team

Notes:

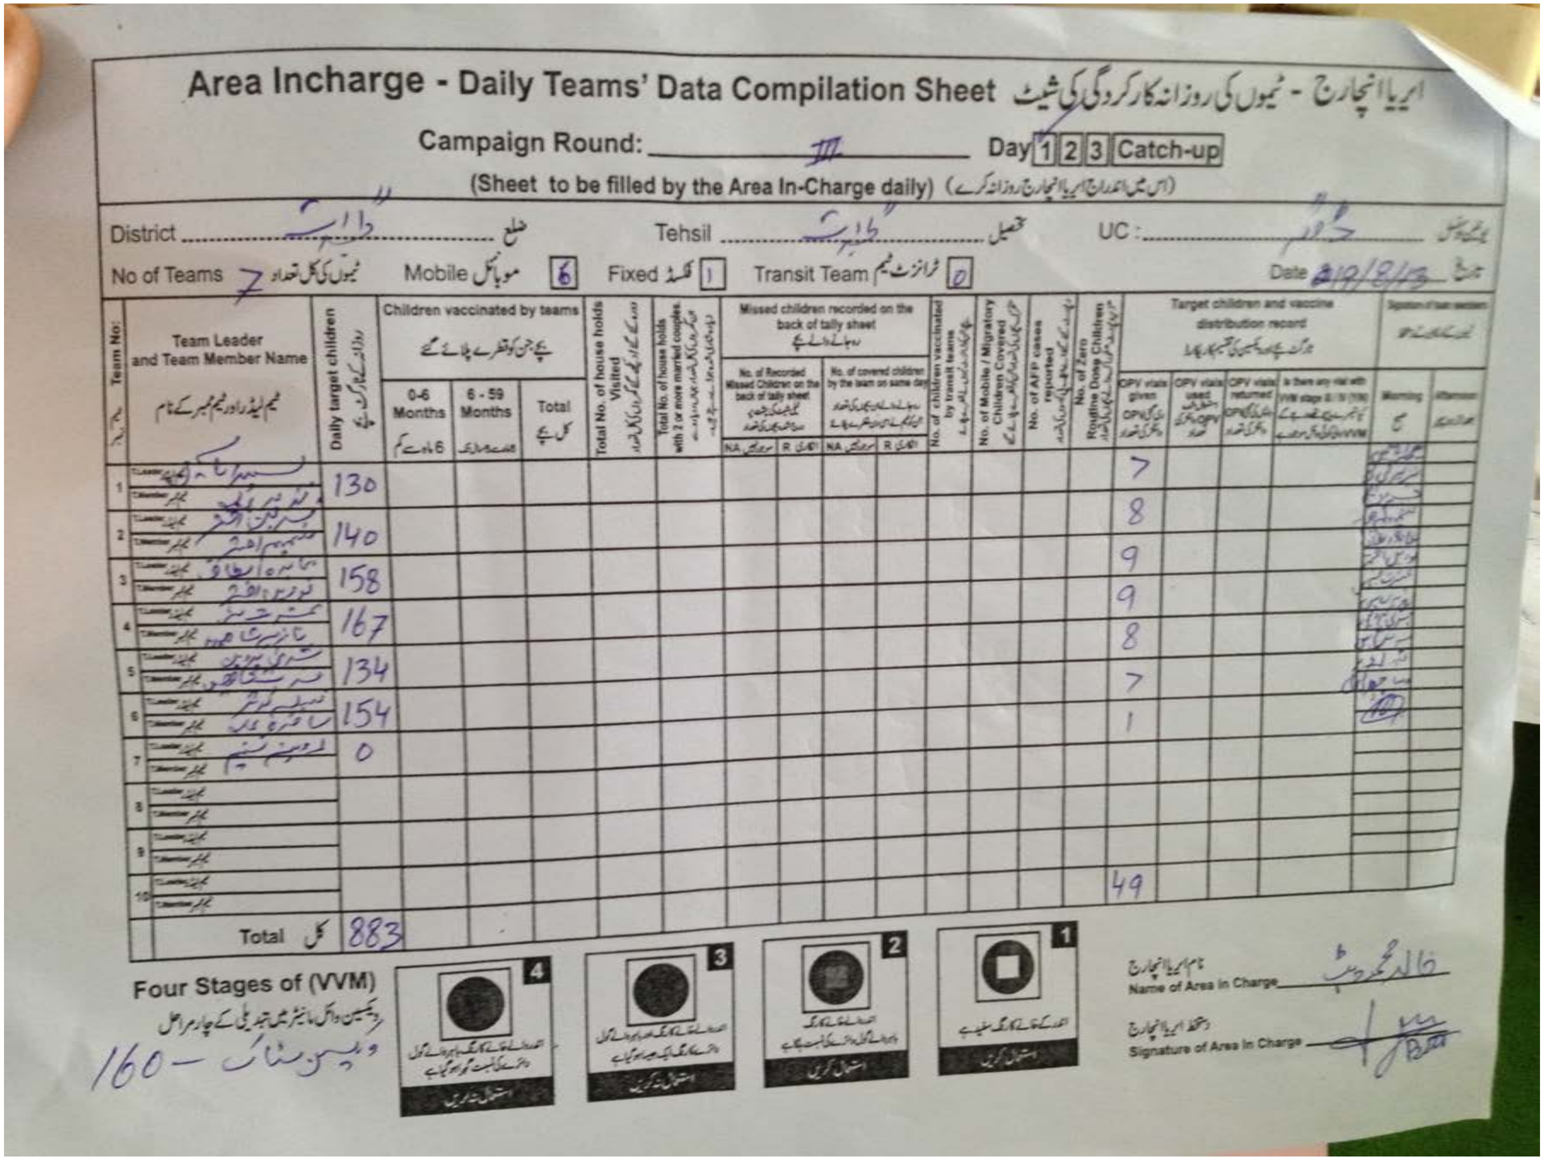

Figure A.4: End-of-Day Compilation of Self-Reports by Vaccination Teams

Notes:

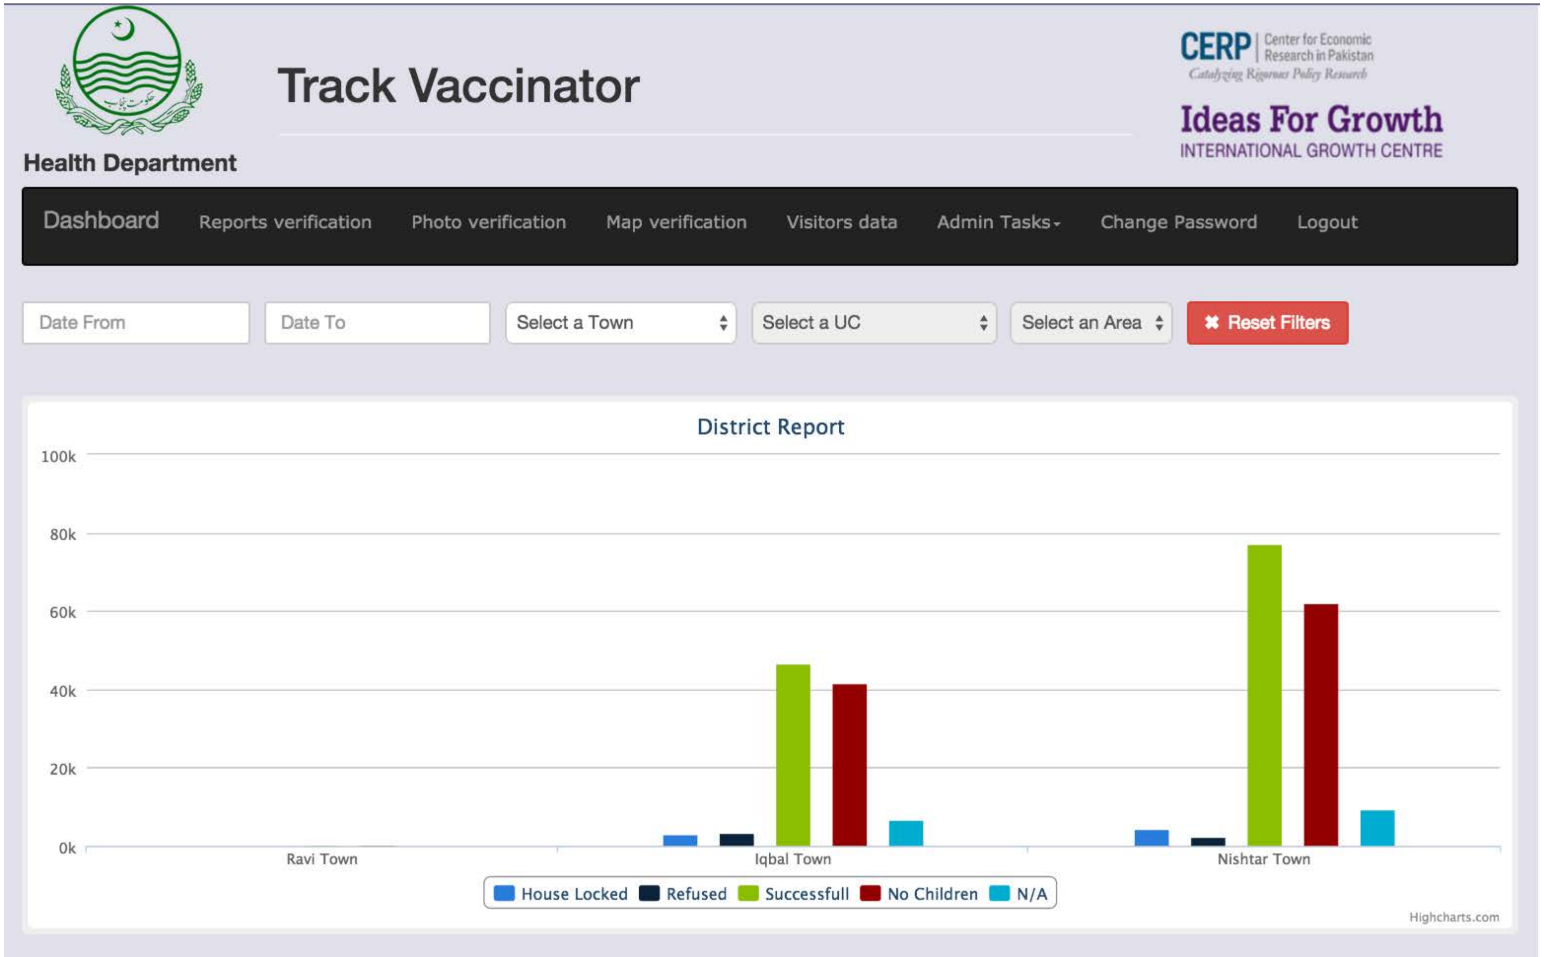

Town	House status				
	Successful	Refused	House Locked	No Children	N/A
Ravi Town	333	77	36	187	84
Iqbal Town	46483	3401	2932	41560	6689
Nishtar Town	77289	2451	4235	61962	9272

Figure A.5: Screenshot of the Track Vaccinator Dashboard

Notes:

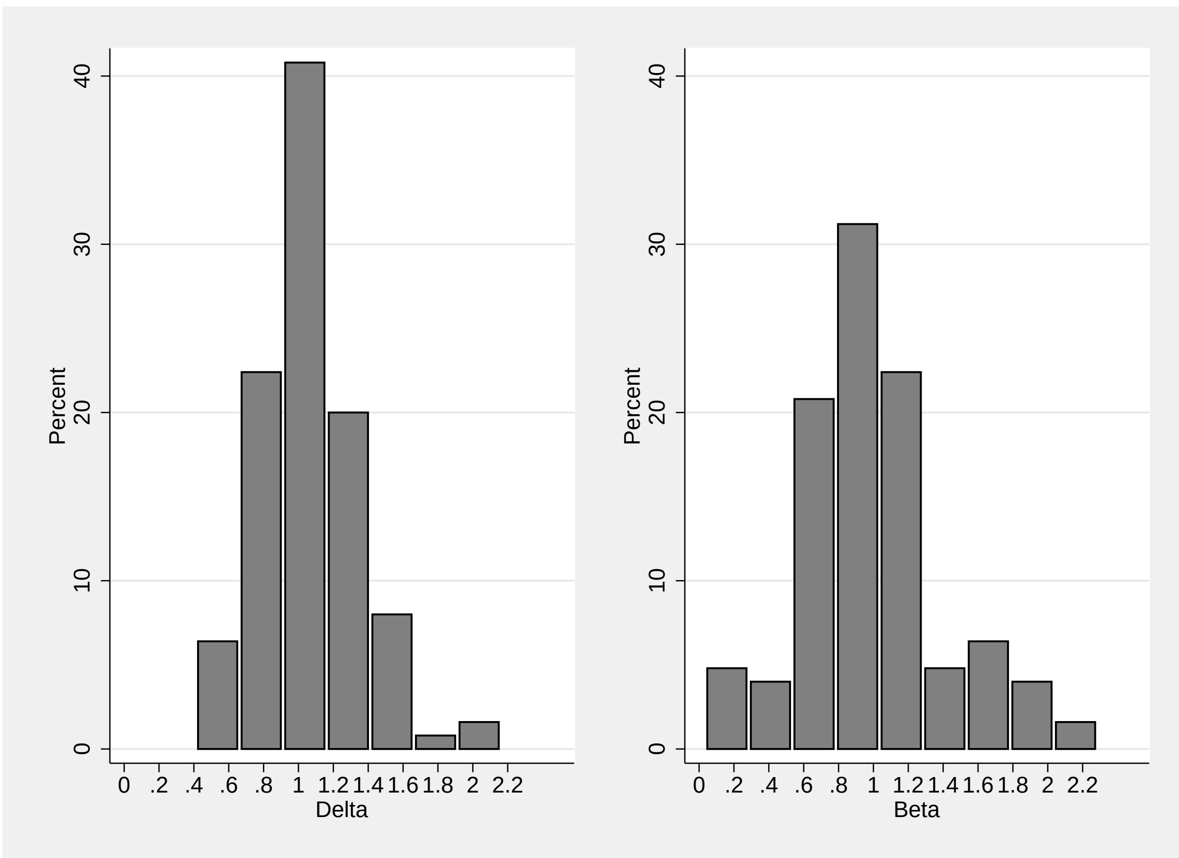

Figure A.6: Within-Subject Parameter Calculations

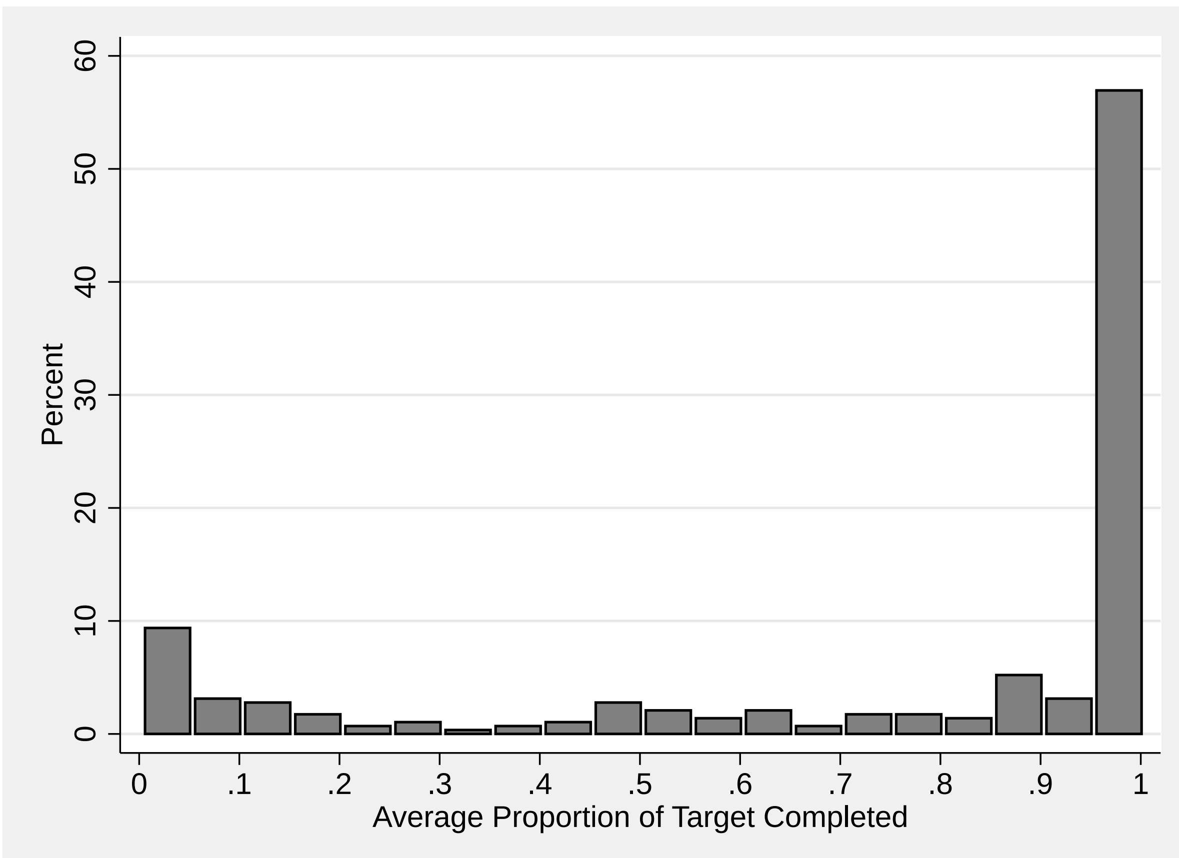

Figure A.7: Individual Completion Rates

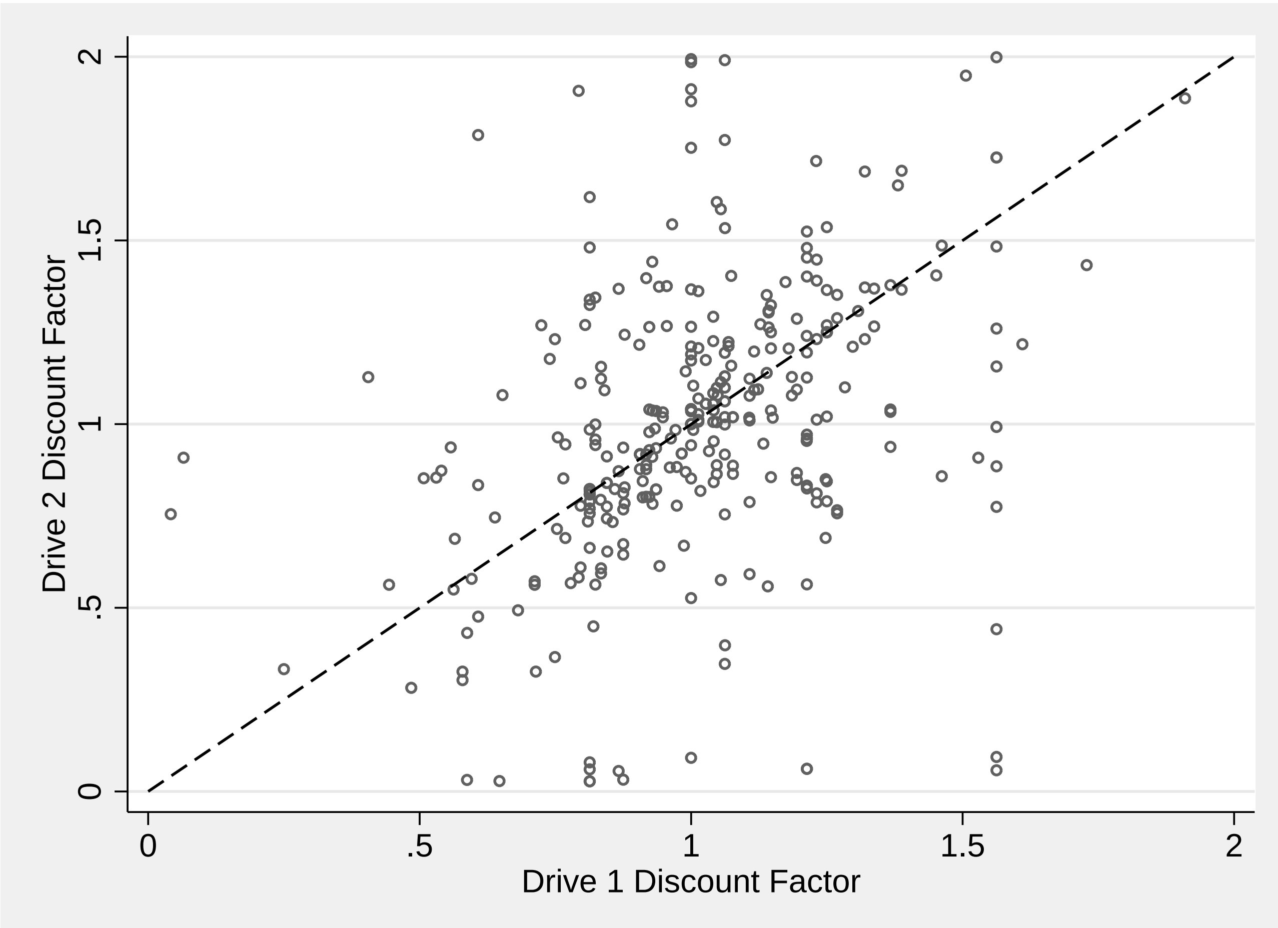

Figure A.8: Stability of Preferences

A.2 Appendix Tables

Table A.1: Testing Whether Failure to Set in Drive 0 Was Systematic

	Did Not Fail (1)	Failed (2)	p-value (3)
Gender (Female = 1)	0.965	1.000	0.082
	(0.020)	(0.000)	
Years of Education	10.294	10.146	0.608
	(0.220)	(0.185)	
Number of Children	3.268	3.388	0.695
	(0.239)	(0.188)	
Punjabi (=1)	0.952	0.975	0.440
	(0.023)	(0.018)	
Has a Savings Account (=1)	0.317	0.305	0.867
	(0.052)	(0.051)	
Participated in a Rosca (=1)	0.446	0.378	0.380
	(0.055)	(0.054)	
Years in Health Department	10.135	10.886	0.337
	(0.554)	(0.547)	
Years as Polio Vaccinator	9.994	10.531	0.467
	(0.538)	(0.502)	
# Vaccinators	254	82	

Notes: This table tests whether the failure of the smartphone app during Drive 0 was systematic. Standard errors reported in parentheses. Column 3 reports a p-value corresponding to the null that the mean in the Did Not Fail group is equal to the Failed group.

Table A.2: Within-Subject Parameter Estimates

	Complete Panel	No Change	Change	Immediate → Advance	Advance → Immediate
	(1)	(2)	(3)	(4)	(5)
β	0.817	0.836	0.798	0.700	0.886
	(0.074)	(0.113)	(0.097)	(0.178)	(0.073)
δ	0.986	0.965	1.007	1.038	0.983
	(0.030)	(0.043)	(0.040)	(0.047)	(0.061)
a	-16.554	-15.788	-16.609	-17.550	-21.683
	(0.445)	(0.814)	(0.549)	(0.970)	(0.835)
$\gamma = 1 + 2 \cdot \frac{1}{1+exp(a)}$	3	3	3	3	3
$ln(\sigma)$	-0.767	-0.862	-0.699	-0.389	-1.289
	(0.124)	(0.179)	(0.167)	(0.196)	(0.099)
# Observations	464	212	252	112	140
# Vaccinators	232	106	126	56	70
Log-Likelihood	-302.737	-118.101	-181.540	-115.316	-18.227
$H_0 : \beta = 1$	$\chi^2(1) = 6.076$	$\chi^2(1) = 2.081$	$\chi^2(1) = 4.348$	$\chi^2(1) = 2.858$	$\chi^2(1) = 2.444$
	$(p < 0.01)$	$(p = 0.14)$	$(p < 0.05)$	$(p < 0.10)$	$(p = 0.12)$

Notes:

Table A.3: Testing Stationarity of Costs Across Days

Panel A: Time Lapse Between Vaccinations								
Dependent variable:	Day 1 Med. Time Lapse			Day 2 Med. Time Lapse			Day 1 - Day 2 Med. Time Lapse	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Advance Choice (=1)	0.519	1.134	1.011	-0.910	-1.161	-0.829	2.295	1.840
	(2.492)	(1.163)	(1.045)	(3.164)	(3.324)	(3.182)	(3.527)	(3.343)
Discount Factor			-3.697			10.004		-13.701
			(3.504)			(8.247)		(9.000)
Constant	3.370*	1.422***	5.337	4.447*	4.540*	-6.053	-3.118	11.390
	(1.851)	(0.084)	(3.708)	(2.372)	(2.501)	(6.558)	(2.501)	(7.581)
R-Squared	0.000	0.004	0.016	0.000	0.001	0.013	0.002	0.022
# Observations	265	228	228	240	228	228	228	228
Panel B: Distance Walked Between Vaccinations								
Dependent variable:	Day 1 Med. Distance			Day 2 Med. Distance			Day 1 - Day 2 Med. Distance	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Advance Choice (=1)	0.039	0.065	0.058	-0.173	-0.168	-0.153	0.233	0.211
	(0.075)	(0.078)	(0.071)	(0.154)	(0.163)	(0.146)	(0.181)	(0.162)
Discount Factor			-0.210			0.515		-0.726
			(0.236)			(0.518)		(0.571)
Constant	0.058**	0.037***	0.260	0.215	0.198	-0.348	-0.161	0.607
	(0.026)	(0.010)	(0.249)	(0.153)	(0.163)	(0.390)	(0.163)	(0.465)
R-Squared	0.001	0.003	0.012	0.006	0.005	0.020	0.008	0.031
# Observations	260	224	224	238	224	224	224	224

Notes: $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.4: Testing Stationarity of Costs Across Conditions

	Drive 1		
	Advance and Immediate	Advance	Immediate
	(1)	(2)	(3)
Discount Factor	0.932	0.973	0.891
	(0.033)	(0.029)	(0.059)
a	-16.846	-18.613	-17.281
	(0.472)	(0.186)	(1.145)
$\gamma = 1 + 2 \cdot \frac{1}{1+exp(a)}$	3	3	3
$ln(\sigma)$	-1.120	-1.616	-0.867
	(0.177)	(0.077)	(0.216)
# Observations	338	174	164
Log-Likelihood	-101.193	34.226	-90.571
H_0 : Discount Factor = 1	$\chi^2(1) = 4.20$	$\chi^2(1) = 0.86$	$\chi^2(1) = 3.47$
	($p < 0.05$)	($p = 0.35$)	($p = 0.06$)

Notes: This reports structural estimates of β, δ, and γ obtained using Maximum Likelihood Estimation based on Equation (2). Standard errors are reported in parentheses.

Table A.5: Testing for Idiosyncratic Shocks

Panel A: Maximum Daily Time Lapse Between Vaccinations								
Dependent variable:	Max Day 1 Time Lapse			Max Day 2 Time Lapse			Day 1 - Day 2 Max Time Lapse	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Advance Choice (=1)	-0.702	0.578	0.739	8.067	5.274	5.574	-4.695	-4.836
	(9.783)	(9.104)	(9.026)	(8.885)	(9.114)	(9.021)	(12.319)	(12.211)
Discount Factor			4.831			9.054		-4.223
			(14.139)			(15.570)		(19.824)
Constant	59.258***	54.880***	49.764***	53.437***	54.362***	44.774***	0.518	4.990
	(7.920)	(7.254)	(15.266)	(5.178)	(5.404)	(15.351)	(8.724)	(20.662)
R-Squared	0.000	0.000	0.000	0.003	0.001	0.003	0.001	0.001
# Observations	265	228	228	240	228	228	228	228
Panel B: Maximum Time Lapse > 2 hours								
Dependent variable:	Max Day 1 Lapse > 2hr.			Max Day 2 Time Lapse > 2hr.			Day 1 > 2hr. - Day 2 > 2hr.	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Advance Choice (=1)	0.051	0.042	0.044	0.026	0.001	0.002	0.041	0.042
	(0.045)	(0.047)	(0.047)	(0.041)	(0.042)	(0.041)	(0.060)	(0.060)
Discount Factor			0.053			0.032		0.021
			(0.077)			(0.071)		(0.100)
Constant	0.133***	0.127***	0.071	0.103***	0.109***	0.075	0.018	-0.004
	(0.029)	(0.032)	(0.085)	(0.028)	(0.030)	(0.075)	(0.043)	(0.111)
R-Squared	0.005	0.004	0.005	0.002	0.000	0.001	0.002	0.002
# Observations	265	228	228	240	228	228	228	228

Notes: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.6: Testing for Deterministic Environment

	Drive 1			
	Experience < 9 Years	Experience ≥ 9 Years	Experience < 15 Years	Experience ≥ 15 Years
	(1)	(2)	(3)	(4)
β	0.925	0.924	0.953	0.855
	(0.087)	(0.104)	(0.071)	(0.166)
δ	0.985	0.943	0.964	0.946
	(0.042)	(0.040)	(0.033)	(0.065)
a	-17.181	-16.997	-19.654	-17.993
	(1.088)	(0.450)	(1.116)	(1.363)
$\gamma = 1 + 2 \cdot \frac{1}{1+exp(a)}$	3	3	3	3
$ln(\sigma)$	-1.189	-1.059	-1.238	-0.860
	(0.284)	(0.227)	(0.195)	(0.324)
# Observations	149	180	248	81
Log-Likelihood	-34.289	-64.742	-44.968	-45.288
$H_0 : \beta = 1$	$\chi^2(1) = 0.75$	$\chi^2(1) = 0.54$	$\chi^2(1) = 0.45$	$\chi^2(1) = 0.76$
	$(p = 0.39)$	$(p = 0.46)$	$(p = 0.50)$	$(p = 0.39)$

Notes: This reports structural estimates of β, δ, and γ obtained using Maximum Likelihood Estimation based on Equation (2). Standard errors are reported in parentheses.

Table A.7: Testing for Completion

Dependent variable:	$Completed_1 \geq v_1$ & $Completed_2 \geq v_2$		$Completed_1 \geq v_1$		$Completed_2 \geq v_2$		Average Proportion Completed	
Advance Choice (=1)	0.097	0.098*	0.038	0.039	0.085	0.087	0.098**	0.098**
	(0.059)	(0.058)	(0.057)	(0.057)	(0.058)	(0.057)	(0.042)	(0.042)
Discount Factor		0.323***		0.110		0.436***		0.173**
		(0.093)		(0.097)		(0.087)		(0.069)
Constant	0.445***	0.111	0.616***	0.503***	0.534***	0.084	0.706***	0.527***
	(0.041)	(0.104)	(0.040)	(0.111)	(0.041)	(0.100)	(0.031)	(0.080)
R-Squared	0.009	0.042	0.002	0.006	0.007	0.068	0.019	0.037
# Observations	288	288	288	288	288	288	288	288

Notes: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.8: Parameter Estimates by Urban and Experienced

	(1)	(2)	(3)
β	1.002	1.003	1.013
	(0.049)	(0.049)	(0.050)
Experienced Urban Group (=1)	0.024		
	(0.132)		
δ	0.981	0.993	0.978
	(0.031)	(0.032)	(0.031)
Experienced Urban Group (=1)		-0.060	
		(0.072)	
a	-182.100	-18.056	-15.733
	(25.909)	(0.335)	(0.276)
Experienced Urban Group (=1)			14.357
			(2.195)
$\gamma = 1 + 2 \cdot \frac{1}{1+exp(a)}$ (Rural Sample:)	3	3	3
$\gamma = 1 + 2 \cdot \frac{1}{1+exp(a)}$ (Urban Sample:)			2.597
$ln(\sigma)$	-2.170	-2.172	-2.170
	(0.052)	(0.052)	(0.052)
# Vaccinators	203	203	203
Log Likelihood	152.424	152.904	152.521
$H_0 : \beta = 1$	$\chi^2(1) = 0.00$	$\chi^2(1) = 0.000$	$\chi^2(1) = 0.060$
	$(p = 0.979)$	$(p = 0.944)$	$(p = 0.799)$

Notes:

Table A.9: Robustness Tests for Tailoring Intertemporal Incentives

Panel A: Dependent variable $\frac{\lvert v_1 - v_2 \rvert}{\sqrt{2}}$						
	(1)	(2)	(3)	(4)	(5)	(6)
Tailored (=1)	-2.338	-4.773**	-4.821**	-2.723	-1.813	-2.027
	(4.867)	(2.107)	(2.244)	(3.087)	(2.399)	(2.461)
Immediate Choice				22.566***	10.883***	11.158***
				(6.288)	(3.498)	(3.608)
Tailored x Immediate				1.759	-6.323	-6.551
				(10.221)	(4.335)	(4.377)
Constant	31.309***	16.390***	16.978**	18.441***	10.673***	14.464**
	(7.380)	(2.356)	(7.229)	(6.564)	(2.901)	(7.007)
Panel B: Dependent variable $\frac{\lvert v_1 - v_2 \rvert}{\sqrt{2}(v_1+v_2)}$						
	(1)	(2)	(3)	(4)	(5)	(6)
Tailored (=1)	-0.016	-0.019***	-0.017**	-0.010	-0.008	-0.007
	(0.016)	(0.007)	(0.007)	(0.010)	(0.008)	(0.008)
Immediate Choice				0.086***	0.039***	0.039***
				(0.023)	(0.012)	(0.012)
Tailored x Immediate				-0.010	-0.025*	-0.024
				(0.034)	(0.015)	(0.015)
Constant	0.103***	0.057***	0.035	0.055***	0.036***	0.026
	(0.023)	(0.008)	(0.023)	(0.021)	(0.010)	(0.023)
Panel C: Dependent variable $\lvert v_1 - \frac{300}{1+R} \rvert$						
	(1)	(2)	(3)	(4)	(5)	(6)
Tailored (=1)	-3.174	-4.016***	-3.663**	-3.078*	-1.918	-1.809
	(3.273)	(1.417)	(1.462)	(1.777)	(1.597)	(1.609)
Immediate Choice				16.619***	7.978***	7.824***
				(4.436)	(2.414)	(2.466)
Tailored x Immediate				0.257	-4.553	-4.424
				(6.861)	(2.913)	(2.943)
Constant	23.030***	12.131***	7.568	13.649***	7.970***	5.888
	(5.275)	(1.671)	(4.675)	(4.459)	(2.020)	(4.547)
Panel D: Dependent variable $\lvert v_1 - \frac{300}{1+R} \rvert \leq 10$						
	(1)	(2)	(3)	(4)	(5)	(6)
Tailored (=1)	0.267*	0.309*	0.206	0.216	0.215	0.146
	(0.158)	(0.165)	(0.168)	(0.234)	(0.234)	(0.234)
Immediate Choice				-0.737***	-0.625***	-0.573**
				(0.230)	(0.237)	(0.238)
Tailored x Immediate				0.104	0.189	0.144
				(0.325)	(0.337)	(0.340)
Constant	0.220	0.361*	1.530***	0.653***	0.704***	1.705***
	(0.184)	(0.195)	(0.513)	(0.239)	(0.248)	(0.524)
Panel E: Dependent variable $\lvert v_1 - \frac{300}{1+R} \rvert > 25$						
	(1)	(2)	(3)	(4)	(5)	(6)
Tailored (=1)	-0.454**	-0.654***	-0.578***	-0.282	-0.282	-0.235
	(0.191)	(0.221)	(0.216)	(0.317)	(0.314)	(0.309)
Immediate Choice				0.920***	0.720**	0.697**
				(0.272)	(0.283)	(0.286)
Tailored x Immediate				-0.275	-0.676	-0.653
				(0.408)	(0.462)	(0.460)
Constant	-0.639***	-0.887***	-1.568**	-1.231***	-1.312***	-1.812***
	(0.206)	(0.239)	(0.637)	(0.304)	(0.335)	(0.665)
Log-Likelihood	-110.271	-83.129	-82.513	-101.934	-79.678	-79.351
# Vaccinators	280	267	267	280	267	267
Stratum FEs	Yes	Yes	Yes	Yes	Yes	Yes
Exclude 99th and 1st Percentiles	No	Yes	Yes	No	Yes	Yes
Drive 2 R	No	No	Yes	No	No	Yes

Notes: This table reports the effects of tailoring on the equality of effort provision over time. The measure $\lvert \frac{v_t}{v_{t+1}} - 1 \rvert$ reflects the distance of the task allocation (v_1,v_2) from equality ($v_1 = v_2$). Column (1) reports a regression of this measure on a dummy equal to one for subjects in the tailored group. Column (2) reports the
report*$p < 0.1$, **$p < 0.05$, ***$p < 0.01$.

Table A.10: Alternate Treatment Measures and Sample Restrictions

Panel A: Tailoring Intensity						
	(1)	(2)	(3)	(4)	(5)	(6)
Tailor Intensity	2.745	0.175**	0.122*	0.198	0.133	0.098
	(2.141)	(0.068)	(0.067)	(0.373)	(0.085)	(0.081)
Immediate Choice				-0.017	0.074***	0.072***
				(0.321)	(0.024)	(0.024)
Tailor Intensity x Immediate				4.681	0.073	0.061
				(3.842)	(0.138)	(0.135)
Constant	0.720*	0.101***	-0.007	0.678	0.062***	-0.017
	(0.435)	(0.018)	(0.060)	(0.452)	(0.021)	(0.060)
# Vaccinators	280	267	267	280	267	267
Panel B: Tailoring Intensity and Boundary Sample						
	(1)	(2)	(3)	(4)	(5)	(6)
Tailor Intensity	1.318	0.159**	0.133**	-0.007	0.057	0.029
	(1.051)	(0.070)	(0.066)	(0.244)	(0.063)	(0.058)
Immediate Choice				0.142	0.070***	0.068**
				(0.215)	(0.027)	(0.026)
Tailor Intensity x Immediate				2.185	0.153	0.159
				(1.888)	(0.121)	(0.116)
Constant	0.651*	0.148***	0.045	0.565*	0.113***	0.015
	(0.348)	(0.026)	(0.067)	(0.340)	(0.026)	(0.065)
# Vaccinators	337	320	320	337	320	320
Stratum FEs	Yes	Yes	Yes	Yes	Yes	Yes
Exclude 99th and 1st Percentiles	No	Yes	Yes	No	Yes	Yes
Drive 2 R	No	No	Yes	No	No	Yes

Notes: This table reports the effects of tailoring on the equality of effort provision over time. The measure $|\frac{v_t}{v_{t+1}} - 1|$ reflects the distance of the task allocation (v_1,v_2) from equality ($v_1 = v_2$). Column (1) reports a regression of this measure on a dummy equal to one for subjects in the tailored group. Column (2) reports the report$^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.