CIRJE-F-1039

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

James Andreoni
UC San Diego, NBER
Michael Callen
UC San Diego, NBER
Muhammad Yasir Khan
UC Berkeley
Charles Sprenger
UC San Diego

February 2017

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

James Andreoni
UC San Diego, NBER

Michael Callen
UC San Diego, NBER

Karrar Hussain
USC

Muhammad Yasir Khan
UC Berkeley

Charles Sprenger
UC San Diego

February 6, 2017

Abstract

We use structural estimates of time preferences to customize incentives for a sample of polio vaccinators in Lahore, Pakistan. We measure time preferences using intertemporal allocations of effort, and derive the mapping between these structural estimates and individually optimized incentives. We evaluate the effect of matching contract terms to discounting parameters in a subsequent experiment with the same vaccinators. This exercise provides a test of the specific point predictions given by structural estimates of discounting parameters. We demonstrate that tailoring contract terms to individual discounting moves allocation behavior significantly towards the intended objective.

JEL classification: D1, D3, D90

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

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

1 Introduction

Nearly every economic decision people make entails some tradeoff through time, whether it is consumption versus savings, doing a task now or later, building human capital, or investing in one's career. Characterizing such choices with structural models of discounting has been a core challenge for economists for much of the last century, with important contributions by Samuelson (1937); Koopmans (1960); Laibson (1997) and 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 a great deal of attention in the empirical literature on intertemporal choice.[1]

This paper seeks to understand whether the out-of-sample predictions given by structural estimates of discounting are empirically valid. We conduct a field experiment on workers' allocation of effort through time, a decision where evidence suggests that present-biased models of decision-making may be particularly relevant.[2] We first estimate the individual discounting parameters of our workers, and then use these estimates to customize each worker's contract to their identified preferences with the intent of reaching a specific intertemporal pattern of work. That is, we tailor incentives within-subject with the objective of reaching a predicted out-of-sample target. Our core test of predictive validity compares tailored workers to a control group which receives untailored, random contract terms.

Interestingly, very little research makes use of the predictive value gained from the articulation and estimation of structural models of discounting.[3] When structural estimates or related

[1] Examples include Hausman (1979); Lawrance (1991); Warner and Pleeter (2001); Cagetti (2003); Laibson, Repetto and Tobacman (2005); Mahajan and Tarozzi (2011); Fang and Wang (2015); Harrison, Lau and Williams (2002); Andersen, Harrison, Lau and Rutstrom (2008); Andreoni and Sprenger (2012a).

[2] For recent experimental examples, see Kaur, Kremer and Mullainathan (2010, 2015); Augenblick, Niederle and Sprenger (2015); Carvalho, Meier and Wang (2014) and Augenblick and Rabin (2015).

[3] What structural models have been used for is for comparison to market interest rates (Hausman, 1979), for comparison across samples, time, or elicitation and estimation strategies (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 (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), to conduct welfare analyses (Laibson, 1997), and to conduct standard counterfactual exercises without out-of-sample testing (e.g., for how price changes should alter demand (Mahajan and Tarozzi, 2011)).

measures are used in out-of-sample prediction exercises, the analysis has often been indirect, linking differences in measured patience to differences in other behaviors without an articulated model for the precise relationship between the two (Chabris, Laibson, Morris, Schuldt and Taubinsky, 2008b; Meier and Sprenger, 2008, 2012, 2010; Ashraf, Karlan and Yin, 2006; Dohmen, Falk, Huffman and Sunde, 2006; Castillo, Ferraro, Jordan and Petrie, 2011).[4] Though such correlational exercises yield valuable insights, they could potentially be made more powerful by directly employing the theoretical parameters in developing the out-of-sample prediction.[5]

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.[6] The disease largely affects children under five. The function of LHW vaccinators 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. Vaccinators 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 vaccinations. Prior to our project there was no technology for monitoring vaccinators, and vaccinators self reported their achievements. As one might imagine, vaccinators often fell short of their suggested targets, but rarely reported doing so. This behavior is consistent with the large literature on public sector absenteeism (Banerjee and Duflo, 2006; Banerjee, Duflo and Glennerster, 2008; Chaudhury, Hammer, Kremer, Muralidharan and Rogers, 2006; Callen, Gulzar, Hasanain and Khan, 2015).

[4]One exception is Mahajan and Tarozzi (2011) who use monetary measures for time inconsistency and purchase and treatment decisions for insecticide treated bednets together to estimate the extent of present bias and 'sophistication' thereof. This exercise can be thought of as articulating the relationship between the experimental measures of time inconsistency and contract choice to deliver estimates of present bias. One point noted by Mahajan and Tarozzi (2011) is that the experimental measures wind up having limited predictive power for estimates of present bias that result from their structural exercise.

[5]Indeed, such exercises could by-and-large be conducted without appeal to structural estimation. Linking either non-parametric measures of discounting or structural parameters thereof to other behavior yields largely the same correlational insights if one does not articulate precisely how how the structural parameters should predict behavior.

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

Since our study requires implementing performance-based incentives, it hinges fundamentally on an accurate measure of productivity.[7] To this end, each vaccinator in our sample is provided a smartphone, equipped with a precise real-time reporting application developed expressly for this project.

Our tool for both measuring intertemporal preferences and tailoring intertemporal incentives is a special bonus contract. In this contract, vaccinators set daily work targets, and, conditional on reaching these targets, receive a sizable bonus. In particular, vaccinators set daily targets of v_1 and v_2 vaccination attempts on day 1 and day 2 of the drive, respectively. Vaccinators face 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 constraint

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

where $V = 300$. The bonus contract offers a fixed bonus of 1000 rupees (around \$10) for meeting *both* of their v_1 and v_2 vaccination target attempts over a two-day drive.[8] If either daily target, v_1 or v_2, is not met, the bonus is not received.

Chosen allocations, (v_1, v_2), can be used to structurally estimate discounting parameters for vaccinators. 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). Vaccinators are randomly assigned to make their allocation decision either in advance of the first day of the drive or immediately on day 1 itself. Additionally, vaccinators are randomly assigned an interest rate, R. Under specific structural assumptions, the experi-

[7]This links our work to a substantial body of recent research in development economics examining the role of incentives and monitoring in improving public sector performance (Bertrand, Burgess, Chawla and Xu, 2016; Basinga, Gertler, Binagwaho, Soucat, Sturdy and Vermeesch, 2011; Miller, Luo, Zhang, Sylvia, Shi, Foo, Zhao, Martorell, Medina and Rozelle, 2012; Olken, Onishi and Wong, 2014; Khan, Khwaja and Olken, 2015; Muralidharan and Sundararaman, 2011).

[8]Our bonus program paid vaccinators for attempted rather than successful vaccinations to avoid concerns that vaccinators, motivated by the incentives, would coerce individuals to receive vaccination. Details on the incentive program are provided in Section 2.2. Slightly more than half of vaccination attempts are successful, with slightly less than half of vaccination attempts reporting no child present. Appendix Figure A.6 reports vaccination behavior for each half-hour of the study, demonstrating limited variation in the proportion of successful and failed vaccination attempts throughout the work day.

ment identifies a set of aggregate discounting parameters (for similar estimation strategies see Andreoni and Sprenger, 2012a; Augenblick et al., 2015). And, under additional assumptions, each vaccinator's allocation identifies her individual discount factor.

We use the individual discounting parameters from an initial drive to tailor incentive contracts in a follow-up drive. The tailoring policy we adopt is one which uses the interest rate, R, to induce smooth provision of vaccinations through time, $v_1 = v_2$, for every vaccinator. The optimal choice of interest rate is simple: to ensure smooth provision of service, the policy must give each vaccinator an interest rate equal to their (appropriately defined) discount factor. The policy objective we adopt is admittedly arbitrary. As the importance of the exercise is specifically in choosing a target and measuring predictive success relative to it, many policy objectives could provide a reasonable test.[9] Hence, we view our exercise as a proof of concept for the potential of such tailored contracting.[10]

In a sample of 337 vaccinators, we document three principal findings. First, on aggregate, a present bias exists in vaccination behavior. Vaccinators allocating in advance of day 1 of the drive allocate significantly fewer vaccinations to v_1 than those allocating on the morning 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, vaccinators with tailored contracts provide significantly smoother service.

This paper makes three contributions. First, our exercise uses field behavior about effort

[9]However, there may be specific reason to be interested in smooth provision induced via tailoring of interest rates. First, there are likely many settings where smooth provision is desirable if, for example, work deteriorates if done all at once or planning of other tasks and work flow is facilitated by smooth provision. Second, giving the workers a choice from the menu of options induced by the tailored interest rate retains flexibility on the part of the worker, which may be desirable if workers are likely to renege from precise prescriptions on work at every point in time, or there is uncertainty in either the environment or the preference measurement technology. Third, given our interest in present-biased preferences, one may think that a person choosing in the moment may generate less smooth intertemporal allocations, delaying costly work, relative to this same person choosing in advance. Hence, a policy objective of smooth provision may effectively 'de-bias' the worker's choices, making their allocations look more like their advance choices even if these choices are taken immediately.

[10]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. See section 4.2.5 for detail.

to examine non-standard time preferences, joining a growing literature which identifies present bias from non-monetary choices in field settings (Read and van Leeuwen, 1998; Sadoff, Samek and Sprenger, 2015; Read, Loewenstein and Kalyanaraman, 1999; Sayman and Onculer, 2009; Kaur et al., 2010, 2015; Carvalho et al., 2014).[11] Most closely related is the work of Kaur et al. (2010) and Kaur et al. (2015), who investigate the decisions of Indian data entry workers, also finding evidence of dynamic inconsistency and self-control problems at work.[12]

Second, we find that the predictions given by structural preference estimates are, to a large degree, accurate. While a substantial literature provides evidence that preference measures correlate with economic behavior, this investigation is the first, to our knowledge, to test whether the specific point predictions for behavior given by an experimental preference measure are accurate. Further, we find evidence in Section 4.2.4 that the predictions given by our structural model may offer advantages beyond atheoretic prediction models optimized using machine learning techniques.

The third contribution of our paper is to show that preference measures can be used by principals to improve agents' incentives.[13] Much of the contract theory literature points to the central role of preferences in determining the optimal design of incentives.[14] The unobservability of preferences poses a key obstacle to testing the optimality of implemented contracts.[15]

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

[12] Documenting dynamic inconsistency outside of the laboratory and outside of the standard experimental domain of time dated monetary payments is particularly valuable given recent discussions on the elicitation of present-biased preferences using potentially fungible monetary payments (Cubitt and Read, 2007; Chabris, Laibson and Schuldt, 2008a; Andreoni and Sprenger, 2012a; Augenblick et al., 2015; Carvalho et al., 2014). As in prior research, our results show that when investigating non-monetary choices, dynamic inconsistency may well have empirical support (see, e.g., Augenblick et al., 2015; Carvalho et al., 2014; Augenblick and Rabin, 2015).

[13] A classic literature in personnel economics documents the potential benefits of implementing piece rates relative to lump sum payments (Lazear, 2000; Paarsch and Shearer, 1999; Shearer, 2004) and the elasticity of effort with respect to the piece rate (Paarsch and Shearer, 2009). Our focus is not on the incentive effects of piece rates, but rather the benefits of using preference estimates to customize incentives.

[14] The relevant empirical literature points to a role for risk preference in contractual settings (Jensen and Murphy, 1990; Haubrich, 1994; Ackerberg and Botticini, 2002; Dubois, 2002; Bellemare and Shearer, 2013) Chiappori and Salanié (2003) provide a review of empirical tests of predictions from contract theory about the design of incentives.

[15] In a review of the literature, Prendergast (1999) summarizes this point: "While the conclusions taken from

We measure a normally unobservable preference parameter and then use such measures in subsequent contract design. Rather than examining whether existing contracts are optimal, we derive the optimal compensation scheme given a policy instrument (the interest rate), a policy objective (smooth provision of service), and the measures of preferences, and then experimentally test whether the optimized scheme improves performance. Our tailored contracts demonstrate that structural preference estimates based on experimental procedures can assist the design of incentives.

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 components: eliciting the time preferences of vaccinators; identifying individual discounting parameters, and, after assigning individually tailored contracts to workers, testing whether these tailored contracts deliver on their specific objective.

In the first subsection below, we describe the smart-phone monitoring application we developed to track the productivity of our workers. We then describe how we identify discounting parameters, and how we use these to design our tailored contracts. A fourth subsection provides details of our experimental design.

2.1 Vaccinations and Smartphone Monitoring

The Department of Health in Lahore, Pakistan, employs Lady Health Worker vaccinators throughout the city to conduct polio vaccination drives. Every month there is a vaccination drive that is at least two days long. Vaccinators are organized into teams of one senior worker

this literature could be correct, this seems a poor method of testing agency theory...because many of the factors relevant for choosing the level of compensation are unobserved; the optimal piece rate depends on risk aversion and the returns to effort, both of which are unknown to the econometrician...it is a little like claiming that prices are too high without knowing costs." (p. 19)

and one junior assistant. These teams work together throughout the drive. Our experiment focuses on the incentives of the senior vaccinator.

Prior to our study, the standard protocol for vaccination drives was to provide each vaccinator a fixed target for total vaccinations over the drive and a map of potential households (called a "micro-plan"). No explicit incentives for completing vaccinations were provided and vaccinators received a fixed daily wage of 100 rupees (around $1). Vaccinators were asked to walk their map, knocking on each compound door, and vaccinating each child for whom parental permission was granted.[16] At the end of each day, vaccinators in each neighborhood convened with their supervisor and self-reported their vaccination activity for the day.[17] In principle, a monitor could verify the claims.[18] In practice, however, there was virtually no monitoring, and strong reasons to suspect over-reporting.[19]

In collaboration with the Department of Health, we designed a smartphone-based monitoring system. Each vaccinator in our study was given a smartphone equipped with a vaccination monitoring application. The vaccinator was asked to record information related to each vaccination. Then, she was asked to take a picture of the home visited and her current vial of vaccine. An image of the main page of the application is provided as Figure 1, Panel A. Data from the smartphone system were aggregated in real-time on a dashboard available to senior

[16]Vaccinating a child consists of administering a few drops of oral vaccine. As there is no medical risk of over-vaccination, vaccinators are encouraged to vaccinate every child for whom permission is granted. For each attempted vaccination, vaccinators were asked to mark information related to the attempt (number of children vaccinated, whether or not all children were available for vaccination, etc.) in chalk on the compound wall. Appendix Figure A.1 provides an example of neighborhood micro-plan, Appendix Figure A.2 provides an example of a vaccination attempt, and Appendix Figure A.3 provides a picture of a chalk marking on a compound wall.

[17]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.

[18]This could potentially be done by walking the micro-plan and examining the chalk markings on each compound wall.

[19]We attempted to independently audit vaccinators by following the trail of chalk markings, but our enumerators found the process too difficult to produce a reliable audit of houses visited. We do, however, know the targets associated with each micro-plan prior to our monitoring intervention and that 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 Smartphone App

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

health administrators.[20]

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

2.2 Intertemporal Bonus Contracts

We worked 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 $V = 300$ vaccination attempts in exchange for a fixed bonus of 1000 rupees. Vaccinators

[20]This dashboard system is based on the technology described in Callen et al. (2015) and is depicted in Appendix Figure A.5.

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 vaccinator would receive only her standard wage.

Each vaccinator 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).[21] The intertemporal allocation (v_1, v_2) potentially carries information on the time preferences of each vaccinator. We next describe the relevant experimental variation and structural assumptions that permit us to identify discounting parameters at the aggregate and individual level.

2.2.1 Experimental Variation and Structural Identification

Our design generates two sources of experimental variation. First, each vaccinator 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 vaccinator 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 vaccinator was asked to pull the slider bar to their desired allocation (v_1, v_2) and then submit. The allocation was required to be submitted before commencing vaccination.

Second, each vaccinator was randomly assigned to either submit their allocation in advance

[21] For applications to field studies and effort allocations, see Augenblick et al. (2015); Carvalho et al. (2014); Gine, Goldberg, Silverman and Yang (2010). We also borrow an additional design element from such studies—minimum allocation requirements—from such studies. In order to avoid vaccinators 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 vaccinators to work on both days of the drive, we avoid confounding extreme patience or extreme impatience with vaccinators simply not wishing to come to work on one of the two days.

of day 1 of the drive or on the morning of day 1. We refer to the first of these as 'Advance' decisions and the second as 'Immediate' decisions. The assignment to either the Advance or Immediate group was independent of 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 the discounting parameters of interest. 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 disutility of effort 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 present bias, active for vaccinators who make Immediate decisions, that is, $\mathbf{1}_{d=1} = 1$. 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 that 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 vaccinator 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 introduced in controlled experiments by Andreoni and Sprenger (2012a), and has precedents in a body of macroeconomic research identifying aggregate preferences from consumption data.[22]

The above development delivers aggregate estimates of discounting parameters with each vaccinator'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 vaccinator, estimation of all parameters at the individual level is infeasible. However, we can calculate each vaccinator's discount factor, which is either δ_i for those who make Advance decisions or $(\beta\delta)_i$ for those who make Immediate decisions. To make such a calculation, two further structural assumptions are required. First, we assume every vaccinator shares a common cost function, $\gamma = 2$, corresponding to quadratic cost. Second, we assume the intertemporal Euler equation (1) is satisfied with equality. Let R_i be the value of R assigned to individual i, let $\mathbf{1}_{d=1,i}$ be their assignment to Advance or Immediate choice, and let $(v_{1,i}, v_{2,i})$ be their allocation of vaccinations. Under the above assumptions, this equation holds:

$$\frac{R_i \cdot v_{1,i}}{v_{2,i}} = (\beta^{\mathbf{1}_{d=1,i}}\delta)_i. \tag{3}$$

[22]See, for example, Shapiro (1984); Zeldes (1989); Lawrance (1991).

The interest rate-adjusted ratio of allocated vaccinations identifies a discount factor for each individual, i.

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 vaccinator's allocation in an intertemporal bonus contract identifies her discount factor for vaccinations, either δ_i for those who make Advance decisions or $(\beta\delta)_i$ 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_i, at the individual level. We formalize the problem as maximizing policy preferences, $P(v_{1,i}(R_i), v_{2,i}(R_i))$, subject to the vaccinator's offer curve. The problem is stated as

$$max_{R,i}\ P(v^*_{1,i}(R_i), v^*_{2,i}(R_i)),$$

where $(v^*_{1,i}(R_i), v^*_{2,i}(R_i))$ are defined as the solution to the vaccinator's minimization problem,

$$min_{v_{1,i},v_{2,i}}\ (v_{1,i})^\gamma + (\beta^{\mathbf{1}_{d=1,i}}\delta)_i \cdot (v_{2,i})^\gamma\ s.t.$$
$$v_{1,i} + R_i \cdot v_{2,i} = V.$$

The solution maps the policy preferences into an interest rate for each vaccinator. 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, $P(v_{1,i}(R_i), v_{2,i}(R_i)) = min[v_{1,i}(R_i), v_{2,i}(R_i)]$.[23] Such Leontief preferences correspond to a policymaker who desires perfectly smooth provision of service. This problem has an intuitive solution. The worker's intertemporal Euler equation (3) yields smooth provision, $v_{1,i} = v_{2,i}$, when $R_i = (\beta^{\mathbf{1}_{d=1,i}}\delta)_i$. Hence, the tailored contracts give each vaccinator a value of R equal to their (appropriately defined) discount factor. Note that the structural discounting parameters are critical in this development. With information on discount factors, contracts can be tailored for each worker to achieve specific policy objectives.

In a second two-day drive, we investigate the promise of tailored contracts. All vaccinators from the first drive were invited to participate in a second intertemporal bonus contract. Vaccinators were unaware that their previously measured behavior would be used to potentially inform their subsequent contracts. This sidesteps an important possibility that vaccinators might alter their first drive behavior in order to receive a more desirable interest rate in the second drive.

Half of vaccinators were given an individually tailored intertemporal bonus contract,

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

where $R_i^* = (\beta^{\mathbf{1}_{d=1,i}}\delta)_i$, either $(\beta\delta)_i$ or δ_i depending on whether they made Immediate or Advance decisions.[24] Some vaccinators' allocation behavior in the first drive implied extreme discount factors and hence extreme values of R_i^*. Our tailoring exercise focused only on a Tailoring Sample of vaccinators with discount factors between 0.75 and 1.5.[25] Vaccinators outside of these bounds were given either the upper or lower bound accordingly.

[23]The ability of our data to speak to alternative policy preferences is discussed in section 4.2.5. 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.

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

[25]Of our sample of 338 vaccinators, 57 exhibit discount factors outside of this range. The Tailoring Sample consists of the remaining 281 vaccinators.

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

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

where $\tilde{R}_i$ was drawn from a random uniform distribution $U[0.75, 1.5]$. The bounds on the distribution of $\tilde{R}_i$ were determined to match the bounds on R_i^*, while the choice of a random uniform control—rather than a single value of $\tilde{R}_i$ or some alternative distribution—was chosen to provide flexible scope for constructing a range of comparison groups for tailored interest rates by drawing subsets of vaccinators assigned to the $\tilde{R}_i$ condition (see section 4.2.3 for details). We find our results are robust to a range of comparison groups.

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.[26] 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. 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 vaccinators participating in the November 10-11 drive received two hours of training at one of three locations in central Lahore on using the monitoring features of the smartphone application. Both Advance and Immediate vaccinators were given identical

[26] 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.

training on the intertemporal bonus contracts and the process by which allocations were made and submitted.

At the end of the training, vaccinators assigned to Advance decision were asked to select their allocations by using the page on their smartphone application. Assistance was available from training staff for those who required it. Vaccinators assigned to Immediate decision were told they would select their allocations using their smartphone application on Monday morning before beginning work. A hotline number was provided if assistance was required for those in the Immediate condition.

The training activities on December 5, for the December 8-9 drive were identical. However, because vaccinators 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 vaccinators had been recruited, were randomized into treatments, and trained. Advance allocation decisions were collected from half of the subjects on Friday, September 26th. On Monday, September 29th, when we attempted to collect immediate allocation decisions, there was apparently a disruption in the mobile network that prevented 82 of 168 Immediate decision vaccinators from submitting their allocations. This caused us to abandon this drive for the purposes of measuring preferences for subsequent tailoring of contracts. The drive, however, was completed and intertemporal bonuses were paid. For the 82 individuals who did not make their allocations, we contacted them, allowed them to continue working, and paid bonuses for all. Figure 2 provides sample details. For completeness, we present data from Drive 0 in Appendix Table A.3, but do not

use Drive 0 for the purposes of tailoring contracts.[27]

Drive 1, November 7-11, 2014: Of the original 336 vaccinators 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 vaccinators 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 vaccinators were assigned to the Advance Choice condition and 175 were assigned to the Immediate Choice condition. While all 174 vaccinators in the Advance Choice condition provided an allocation decision, only 164 of 175 in the Immediate Choice condition provided an allocation. Because 11 vaccinators attrited from the Immediate Choice condition, we also provide bounds on the estimated effect of decision timing using the method of Lee (2009). In addition, for 232 vaccinators, 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 vaccinators who participated in Drive 1 and provided an allocation, 337 again participated in Drive 2. These vaccinators were randomly assigned to be tailored or untailored in their Drive 2 bonus contracts. Importantly, vaccinators 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 also whether the effect of tailoring depends on whether present bias is active.

[27]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.

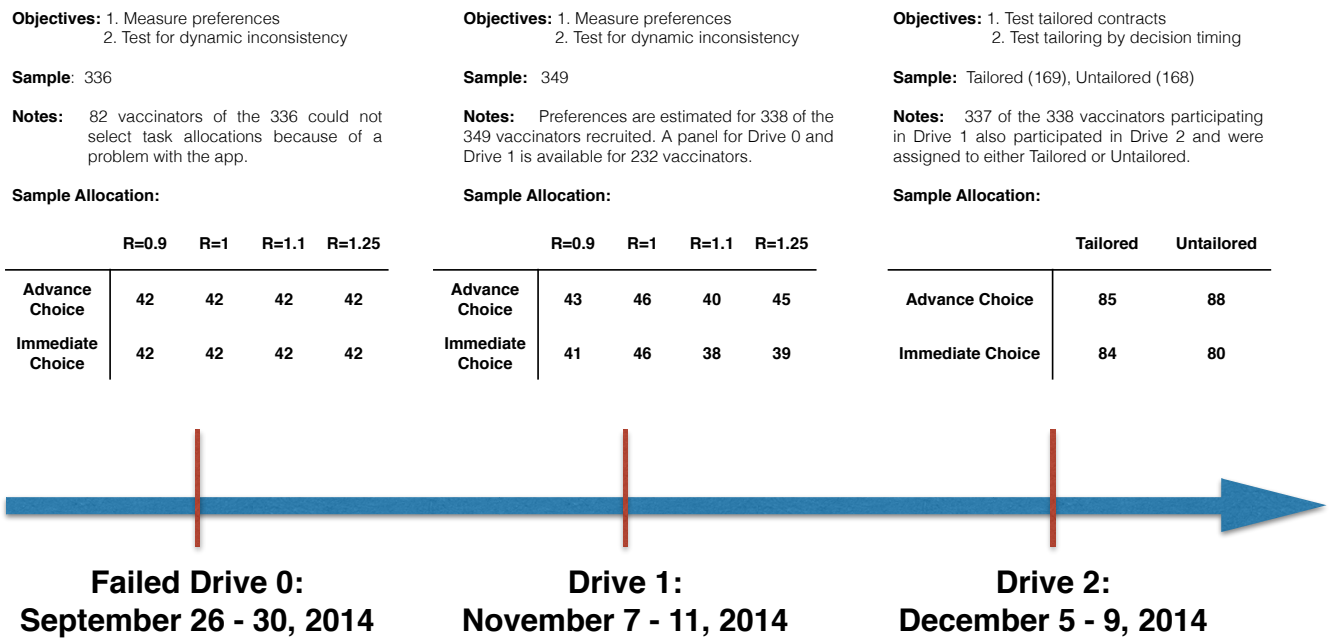

Figure 2: Experiment Overview

Notes: This figure provides an overview of the timing and sample breakdown of the experiment. Assignment to the advance choice and immediate choice condition in Drive 2 is inherited from vaccination Drive 1. Note that: (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.

2.4.3 Sample Details

Table 1 summarizes our sample of vaccinators from Drive 1 and provides tests of experimental balance on observables. Column (1) presents the mean and standard deviation for each variable; columns (2) to (9) present the mean and standard error for each of our eight treatment arms, and column 10 presents a p-value corresponding to joint tests of equality. Our sample is almost exclusively female, more than 90 percent Punjabi in all treatment arms, and broadly without access to formal savings accounts. Vaccinators 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 test performed on an indicator variable equal to one for Punjabi subjects suggests baseline imbalance.

3 Results

We first report results related to the elicitation of intertemporal preference parameters, then evaluate the possibility of tailoring incentives based on individual preferences.[28]

3.1 Elicitation of Time Preferences

3.1.1 Aggregate Behavior

Figure 3 presents median behavior in the elicitation phase of our experiment, graphing the allocation to the sooner work date, v_1, for each interest rate.[29] Separate series are provided for Advance and Immediate choice. In Panel A we provide data for our Full Sample of 338 vacci-

[28]In addition, to test just the effect of providing the $10 bonus, we randomly assigned 85 vaccinators in Drive 0 to carry a phone but not receive an incentive. 73 of these vaccinators also participated in Drive 1, retaining the same 'phone only' treatment status. In Drive 1, vaccinators in the 'phone only' group attempted 169.472 vaccinations (s.e. = 15.98) and vaccinators in the phone plus incentives group attempted 201.32 vaccinations (se = 7.73) yielding an estimated increase of 31.85 attempts (s.e. = 18.49, p = .09). 49.3% of vaccination attempts were successful for the 'phone only' group while 49.1% of vaccinations were successful for the 'phone plus incentives' group. The difference in success rates between the two groups is small (0.2 percentage points) and statistically insignificant (p=0.69). We discuss data on task completion in detail in Section 4.1.2 below.

[29]We opt to provide medians as the average data are influenced by several extreme outliers in allocation behavior. Qualitatively similar patterns are, however, observed.

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: This table checks balance across the eight treatment groups. 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).

nators who provided allocations in Drive 1. In Panel B we focus only on our Tailoring Sample of 281 vaccinators, trimming 57 vaccinators with extreme allocation behavior that would imply individual discount factors from equation (3) outside of the range of [0.75, 1.5]. 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 $V = 300$. Vaccinators respond to this changing incentive by reducing their allocation of v_1. Second, there is a tendency of present bias. Vaccinators appear to allocate fewer vaccinations to v_1 when making Immediate choice.

Table 2 presents corresponding median regression analysis for aggregate behavior in Drive 1.[30] We regress v_1 on R and whether the allocation decision is immediate. Column (1) echoes the findings from Figure 3, Panel A: in our Full Sample, vaccinators assigned to Immediate choice allocate a median of 2.00 (s.e. = 1.13) fewer vaccinations to v_1 than those assigned to Advance choice. Similar patterns are observed in column (2), focusing only on our Tailoring Sample. Vaccinators in the Tailoring Sample allocate a median of 3 fewer vaccinations to v_1 when making immediate choice. As discussed in Section 2.4.2 above, 11 vaccinators 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 (2.06) tasks.

Table 2: Aggregate Drive 1 Behavior

Dependent variable:	Tasks Allocated to the First Day of the Drive (v_1)	
	Full Sample	Tailoring Sample
	(1) Median	(2) Median
Immediate Decision (=1)	-2.00*	-3.00***
	(1.13)	(0.91)
Interest Rate (R)	-54.29***	-66.67***
	(4.38)	(3.66)
Constant	201.86***	216.33***
	(4.72)	(3.93)
Median Advance Choice	146.5	148
# Observations	338	281

Notes: This table reports on the effects of decision timing and interest rate variation on vaccinations allocated to the first day of the drive during Drive 1. Median regression coefficients with standard errors are reported in parentheses. Immediate Decision is an indicator equal to one for vaccinators selecting their allocations on the morning of the vaccination drive. The interest rate R takes the values $R \in \{0.9, 1, 1.1, 1.25\}$. Table 3 provides corresponding between-subject structural preference parameter estimates of β, δ, and γ. *Levels of Significance:* $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

[30] Appendix Table A.3 presents identical analysis incorporating data from failed Drive 0, and identifies qualitatively similar effects.

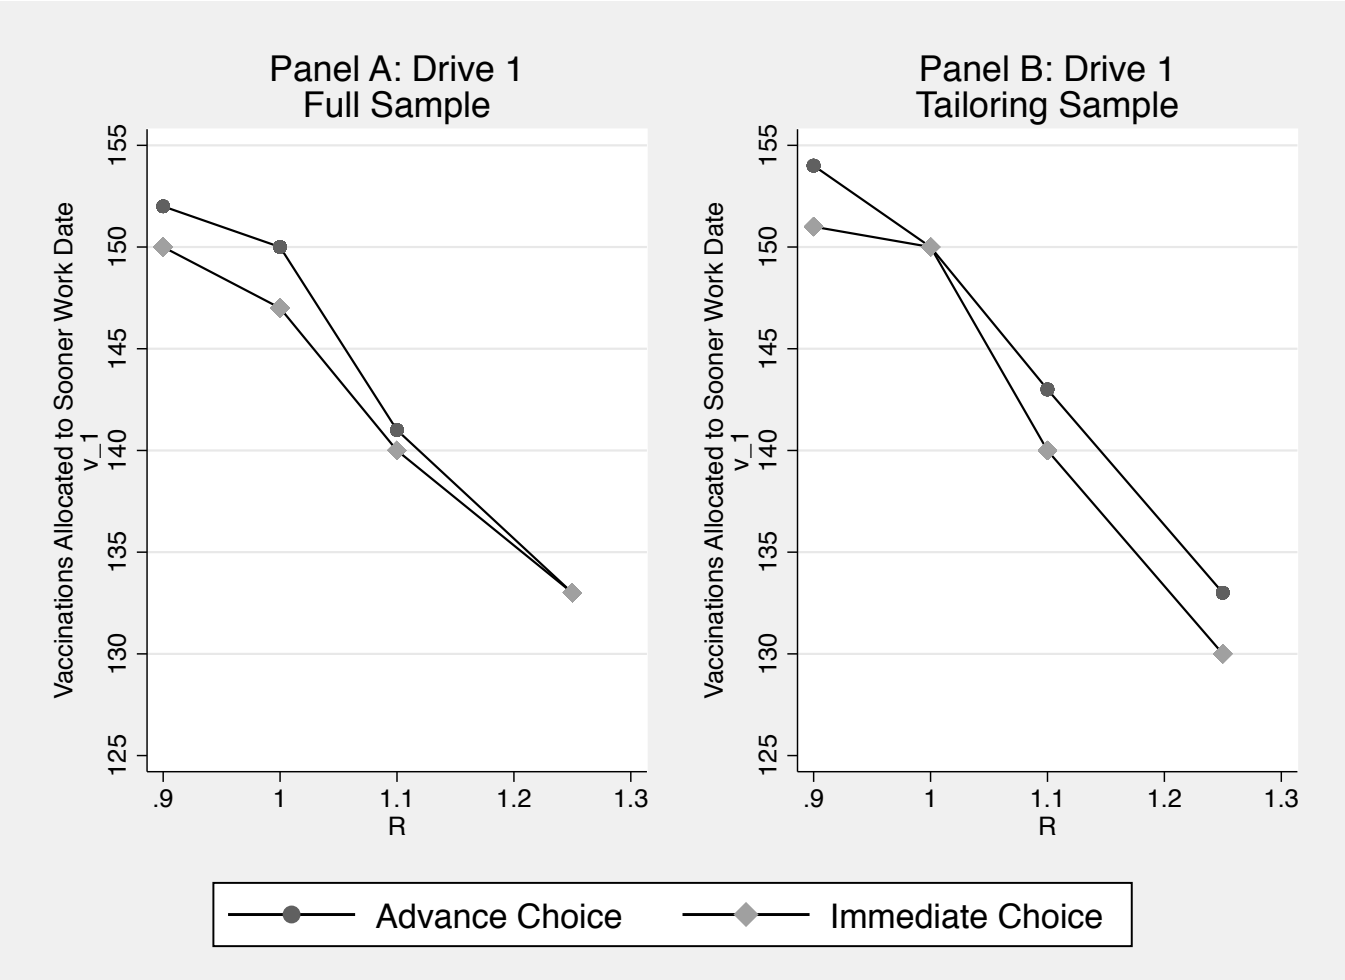

Figure 3: Aggregate Experimental Response

Notes: This figure examines whether tasks assigned to the sooner work date respond to the experimental variation in the interest rate R and in decision timing. Points in the plots are medians for each of the eight treatment groups respectively. Panel A depicts the Full Sample and Panel B depicts the tailoring sample (vaccinators with $R^* < 0.75$ or $R^* > 1.5$). Circles are advance choice groups and diamonds are immediate choice groups.

3.1.2 Aggregate Preference Parameters

The raw data of Figure 3 and analysis of Table 2 indicate responsiveness of vaccinators to our experimental parameters, R and whether allocations are Immediate or Advance. Equation (2) links allocation behavior to these experimental parameters via a structural model 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 (3) of Table 3, we present estimates restricting $\gamma = 2$ (quadratic costs) for the Full Sample and the Tailoring Sample. 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 1.013 (0.018), β of 0.955 (0.037). Though we find an aggregate present bias parameter less than one, echoing the raw results, our estimate of β cannot be statistically distinguished from 1, $\chi^2(1) = 1.50,\ (p = 0.22)$. A similar finding is obtained when focusing only on the Tailoring Sample in column (3).

In columns (2) and (4) 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 $\gamma \in [1, 4]$ by estimating the parameter of a box constraint, a, such that $\gamma = 1 + 3 \cdot \frac{1}{1+exp(a)}$. This restriction ensures allocations are indeed minima (i.e., $\gamma > 1$) and allows costs to be substantially convex (i.e., up to quartic). 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 (4) relative to columns (1) and (3). In column (2), we estimate an aggregate β of 0.878 (0.094), close to the estimates from controlled laboratory choices over effort (Augenblick et al., 2015; Augenblick and Rabin, 2015). Once again, however, our estimate of β cannot be statistically differentiated from 1, $\chi^2(1) = 1.69,\ (p = 0.19)$.

It is important to note that in both columns (2) and (4) the cost parameter a is estimated to be in excess of -15, implying γ extremely close to 4. 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

Table 3: Aggregate Parameter Estimates, Drive 1

	Full Sample		Tailoring Sample	
	(1)	(2)	(3)	(4)
β	0.955	0.878	0.977	0.957
	(0.037)	(0.094)	(0.019)	(0.041)
δ	1.013	0.934	1.029	0.983
	(0.018)	(0.040)	(0.013)	(0.027)
a	-	-16.423	-	-33.452
	-	(0.493)	-	(1.917)
$\gamma = 1 + 3 \cdot \frac{1}{1+exp(a)}$	2	4	2	4
$ln(\sigma)$	-1.049	-1.141	-1.838	-2.141
	(0.150)	(0.183)	(0.032)	(0.044)
# Observations	338	338	281	281
Log-Likelihood	-125.092	-94.034	117.861	202.774

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 (3) are obtained imposing the restriction $\gamma = 2$. Estimates in columns (2)-(4) remove this restriction.

and assess plausible alternative formulations. These 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.[31] 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 vaccinator assuming quadratic costs. For those vaccinators assigned to Advance choice, this discount factor corresponds to δ_i, while for those assigned to Immediate choice it corresponds to $(\beta\delta)_i$. 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].

As noted above, an important minority of vaccinators 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.[32] We term such vaccinators 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 281 vaccinators in the Tailoring Sample and discuss the Boundary Sample in robustness tests (see section 4.2). Figure 4 presents histograms of implied discount factors for the 281 Tailoring Sample vaccinators 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

[31] Appendix Figure A.7 reproduces Figure 3, with in-sample predictions from Table 3, columns (2) and (4). Though the estimates do match the responsiveness of behavior from $R = 1$ to $R = 1.1$, they do not generate the lack of sensitivity for other changes in R.

[32] Such extreme behavior is slightly more pronounced in Immediate choice (34 vaccinators) relative to Advance choice (23 vaccinators), ($t = 1.84,\ p = 0.07$).

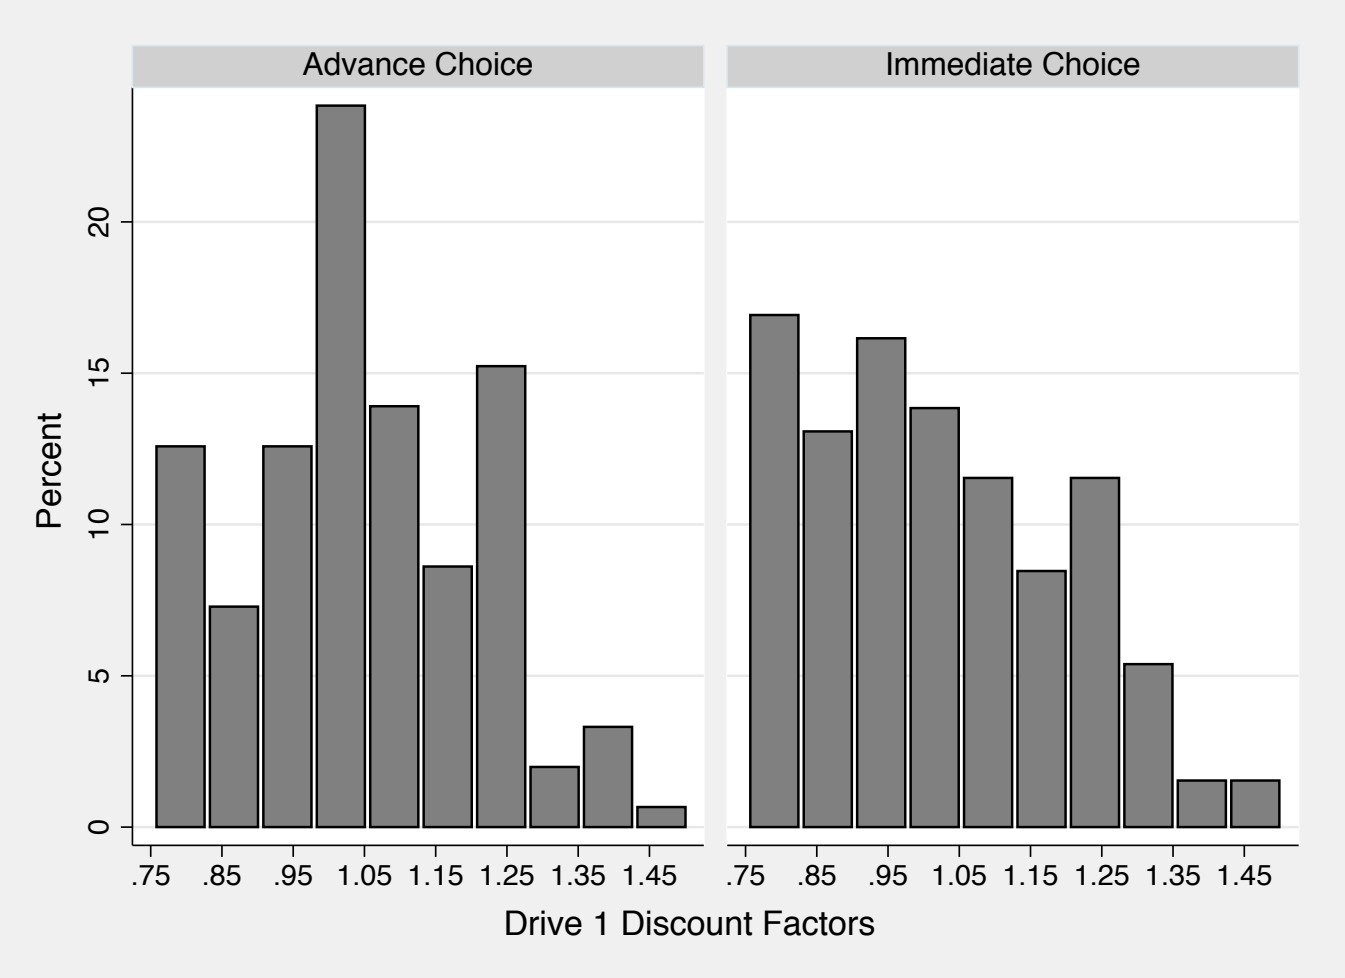

Figure 4: Individual Discount Factors in the Tailoring Sample

Notes: This figure provides histograms of one period discount factors separately for subjects in the Advance Choice condition (left panel) and the Immediate Choice condition (right panel). The sample is restricted to vaccinators in the Tailoring Sample (vaccinators with $R_i^* \geq 0.75$ or $R_i^* \leq 1.5$).

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 vaccinators 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.

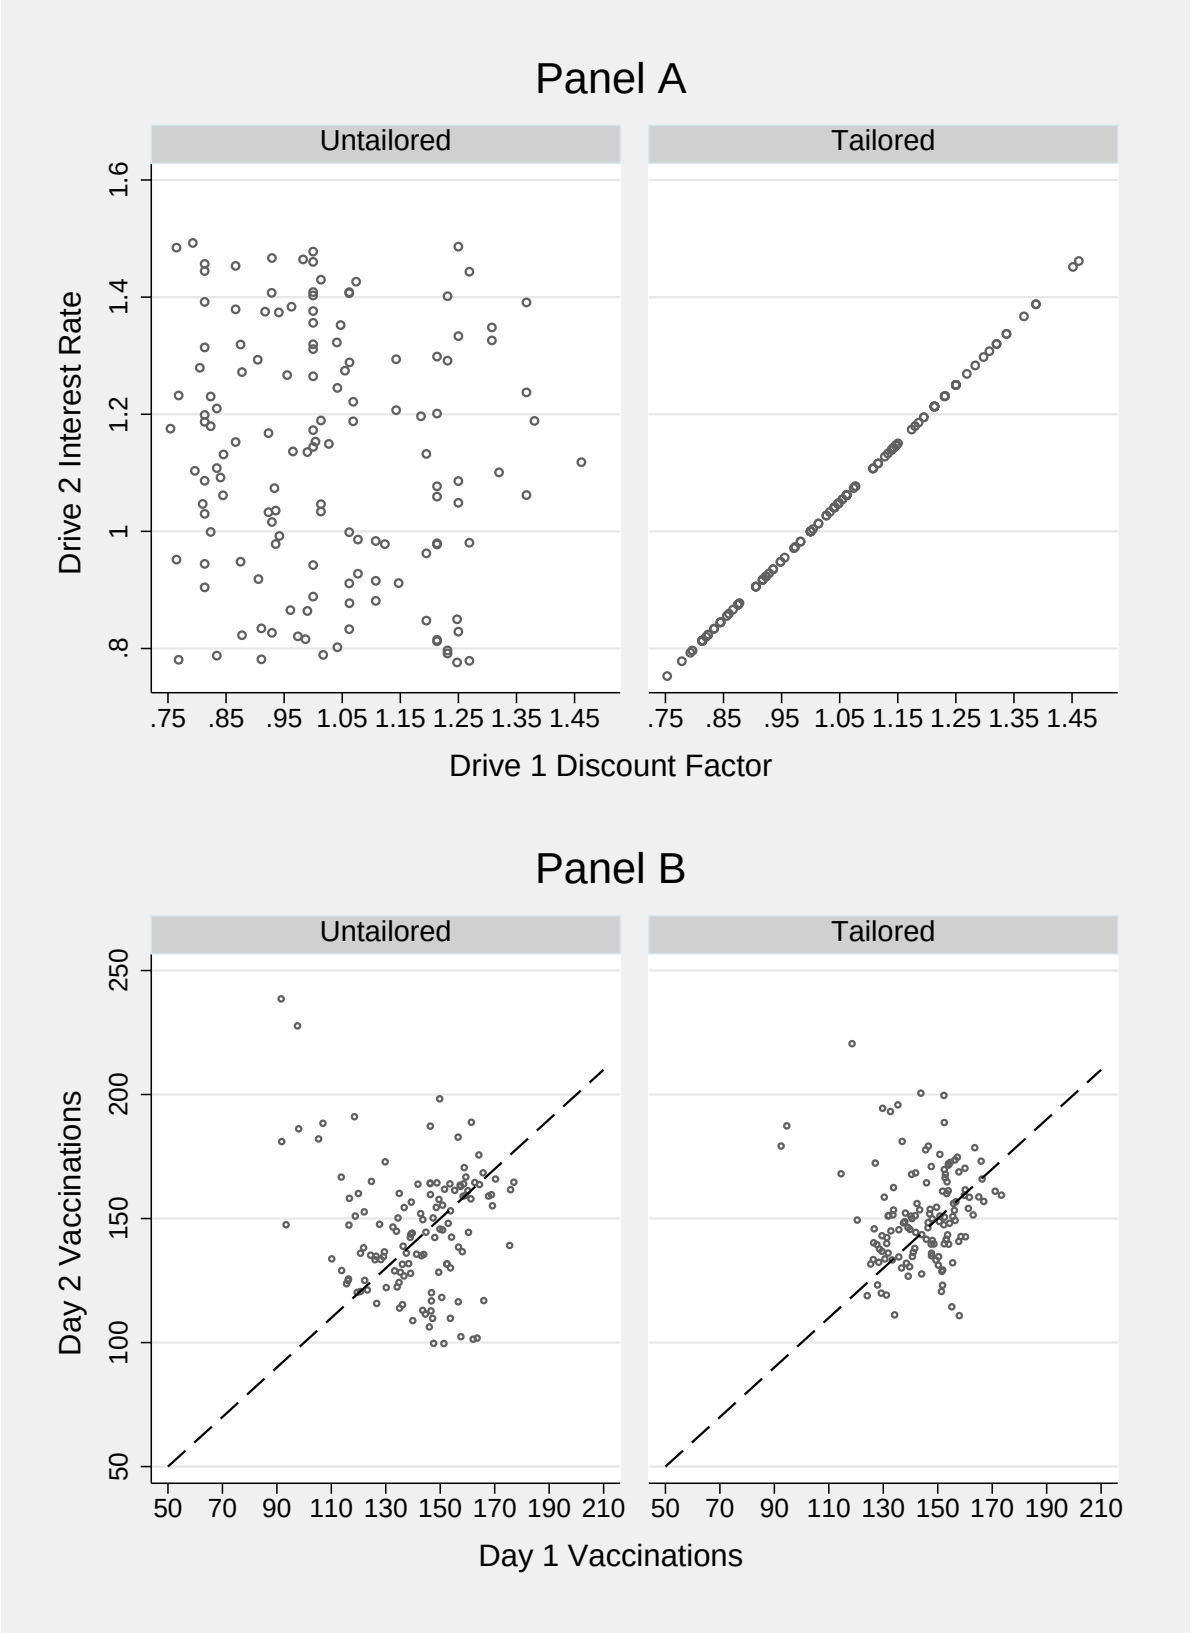

Figure 5: Discounting and Tailoring

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 vaccinators in the Tailoring Sample, 280 participated in Drive 2.[33] Of these, 142 vaccinators were assigned a value of R equal to their discount factor. That is, tailored vaccinators were assigned $R_i^* = (\beta^{\mathbf{1}_{d=1,i}}\delta)_i$, which should induce equal provision of effort through time, $v_{1,i} = v_{2,i}$. The remaining 138 vaccinators serve as control and were assigned a uniform random interest rate $\tilde{R}_i \in U[0.75, 1.5]$.[34]

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 $R_i^* = \beta^{\mathbf{1}_{t=1,i}}\delta_i$ 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 Drive 2 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 $|\frac{v_{1,i}}{v_{2,i}} - 1|$. 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 top and bottom 1% of the sample 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), $t_{265} = 3.10,\ (p < 0.01)$.

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

[34] 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.

In Table 4, we provide corresponding least squares regression analysis. 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_i^* or $\tilde{R}_i$ assigned in Drive 2. This regression identifies whether tailoring generates more equal provision for a given value of R, and hence controls for any differences in interest rates across tailored and untailored groups. Again, tailoring serves to reduce distance significantly.

Vaccinators 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 (8.4 percentage points) is similar to the effect of moving a vaccinator 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 potential benefits to tailored policy initiatives. For a given interest rate, matching this interest rate to individual preferences generates smoother

service provision. The effects are to reduce distance measures by around one-third on average and around one-half in Immediate choice. This suggests that policy makers wishing to alter intertemporal choices of workers may be able to achieve their goals through tailored contracts. Our findings show particular promise for tailoring contracts to change behavior in settings where present bias is relevant.

Table 4: The Effect of Tailoring Intertemporal Incentives

Dependent variable:	$\lvert\frac{v_{1,i}}{v_{2,i}} - 1\rvert$					
	(1)	(2)	(3)	(4)	(5)	(6)
Tailored (=1)	-0.489	-0.059***	-0.049***	-0.070	-0.019	-0.014
	(0.321)	(0.019)	(0.018)	(0.069)	(0.019)	(0.019)
Immediate Choice (=1)				0.982*	0.125***	0.117***
				(0.573)	(0.035)	(0.035)
Tailored x Immediate				-0.888	-0.090**	-0.084**
				(0.604)	(0.040)	(0.040)
Constant	1.407	0.159***	0.022	0.873	0.094***	-0.004
	(0.860)	(0.023)	(0.058)	(0.552)	(0.026)	(0.057)
Stratum FEs	Yes	Yes	Yes	Yes	Yes	Yes
Exclude 99th and 1st Percentiles	No	Yes	Yes	No	Yes	Yes
Drive 2 R_i^* or $\tilde{R}_i$	No	No	Yes	No	No	Yes
R-Squared	0.035	0.060	0.082	0.053	0.142	0.154
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 an indicator 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 (5) and (6) apply the same restrictions to the sample as columns (2) and (3) respectively. Ordinary least squares regressions. Heteroskedasticity robust White standard errors reported in parentheses. *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$.

4 Robustness Tests and Additional Exercises

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. We first examine the robustness of these results to estimation using only within-subject variation from our panel of subjects who participated in Drives 0 and 1, and who changed from Advance to Immediate choice. We then examine the validity of a set of structural assumptions required to infer discounting parameters from vaccinator allocation behavior.

4.1.1 Within-Subject Variation

As shown above, in Drive 1, relying on between subjects tests, we cannot reject $\beta = 1$ in any specification. 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 vaccinators who changed from Advance to Immediate choice (or vice versa) across drives. We re-conduct the analysis of Table 3, focusing on a panel of 126 vaccinators who changed between Advance and Immediate choice across Drives. Using such within-subject variation, we reject the null hypothesis of no present bias at the 5% level. Moreover, the effect is robust to just examining those who transitioned from Advance to Immediate, or just those who transitioned from Immediate to Advance. See Appendix Table A.2 for details. 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.[35]

These 126 vaccinators also provide an opportunity to investigate present bias at the individual level. Following equation (3), we calculate a discount factor for each condition the vaccinator 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. Interestingly, as in as in our analysis of discount factors, we find δ is centered around 1 with a median value of 1.04, and that β is skewed below 1 with a median value of 0.95. Sixty-nine (54.8%) of 126 vaccinators 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).[36] 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 vaccinators who change between Advance to Immediate choice across conditions. See Figure A.8 for graphical detail.

4.1.2 Structural Assumptions

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

Assumption 1: Stationarity of the Cost Function: We assume the cost function is the same for day 1 and day 2. If sooner costs are forecasted to be more severe than later costs, vaccinators 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

[35]More concretely, if vaccinator A is assigned to advanced choice in Drive 0 and immediate choice in Drive 1, and vaccinator B is assigned to immediate choice in Drive 0 and advanced choice in Drive 1, then an idiosyncratic shock (say it is local to immediate choice, such as a revelation that it is raining on the first day of the drive) needs to affect vaccinator A when the drive starts during Drive 0 (in September) and vaccinator B when the drive starts during Drive 1 (in November). Randomization at the level of the vaccinator, among vaccinators all working in two neighborhoods in Lahore, narrows the set of plausible shocks that would satisfy these conditions.

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

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 vaccinator, we identify the median time lapse between vaccination attempts and the median distance covered per 30 minute window each day.[37] Of our 338 vaccinators, measures for median time lapse between vaccination attempts are available for 277 on either Day 1 or Day 2 and for 228 vaccinators on both days of Drive 1.[38] Of our 338 vaccinators, measures for median distance traveled every 15 minutes are available for 274 on either Day 1 or Day 2 and for 224 vaccinators on both days of Drive 1.[39]

Vaccinators take around 3.4 minutes between vaccination attempts and walk around 0.06 miles per 15 minutes on Day 1. Focusing on individuals with measures on both days of the drive, we find that time taken and distance traveled are uncorrelated both with Advance choice and with discount factors within condition. Time and distance are also uncorrelated with Advance

[37] 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.

[38] 265 vaccinators have Day 1 lapse data while 240 have Day 2 lapse data. Of the 73 vaccinators 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 vaccinators 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 vaccinators who completed vaccinations but did not have interaction with the server had their vaccination records pulled manually from their phones after the drive.

[39] 260 vaccinators have Day 1 distance data while 238 have Day 2 distance data. Of the 78 vaccinators 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 vaccinators 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.

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 which is unrelated to assignment to Advance or Immediate choice, and that changes in efficacy are unrelated to measured preferences. This suggests that perceived changes in costs likely do not drive our measures of patience or our finding of present bias.[40] These results are all presented in Appendix Table A.4.

If Advance and Immediate choices were governed by different, but constant, cost functions, this could result in another violation of stationarity. This possibility 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.5, we report estimates of 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.5 indicates that non-stationarity in discount factors is more apparent than non-stationarity in costs.

Assumption 2: 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 individuals that could influence measured patience. For example, a vaccinator with an appointment lasting 2 hours on Day 1 and no appointments 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 vaccinator's idiosyncratic costs across days. Further, if such idiosyncratic events are easier to re-organize when making Advance choice, present bias may

[40] Ultimately, such stationarity is likely to be expected given that vaccinators 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.

be conflated with ease of scheduling.

Here, again, the additional data on vaccinator time use available from the monitoring application is potentially valuable. We can investigate whether extended periods of non-vaccination exist and if they are correlated with measured preferences and allocation timing. As in the example above, a vaccinator with an extended period of non-vaccination may well be experiencing forecasted idiosyncratic costs unrelated to vaccinations. Appendix Table A.6 repeats the analysis from Appendix Table A.4, 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 vaccinators 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.

Assumption 3: Deterministic Environment: Our exercise assumes that individuals can perfectly forecast their own future costs of vaccination. A plausible alternative is that costs are uncertain. 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 vaccinators may grow more patient as uncertainty is resolved, some less so.[41]

One potential, albeit imperfect, test for the effect of uncertainty in the environment is

[41]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.

to examine differences in experience across individuals.[42] More experienced vaccinators may have a more complete understanding of the difficulty of tasks and may face less uncertainty when making allocation decisions. As noted, vaccinators have an average of more than 10 years of experience in health work. In Table A.7, we separate our sample either at the median of experience (9 years) or at 15 years of experience.[43] 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 they disproportionately likely to be dynamically consistent.

Assumption 4: Identical Cost Functions: Our aggregate exercise assumes identical costs across subjects, and our individual elicitation assumes identical 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 vaccinators, regardless of allocation timing, should exhibit $v_1 = v_2 = 150$ for all values of γ.[44] Examining the Drive 0 and Drive 1 data, we find that for 163 vaccinators who were assigned $R = 1$, the mean allocation is $v_1 = 140.84$ ($s.d. = 24.76$).[45] 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.

[42] This test is imperfect as experience may correlate with other differences across vaccinators.
[43] Experience measures are available for 329 of 338 vaccinators in Drive 1.
[44] This is because the Euler equation reduces to $(\frac{v_1}{v_2})^\gamma = R = 1$, which implies $\frac{v_1}{v_2} = 1$.
[45] 42 of 163 vaccinators allocated exactly $v_1 = v_2 = 150$.

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 vaccinators for whom behavior was more sensitive to price and hence curvature estimates were less convex. In Table A.9, 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.[46] Albeit exploratory, such analysis suggests that vaccinators who are experienced vaccinators and who are in an environment where they can vaccinate rapidly, 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[47], 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.

Assumption 5: Inability to Renege: The contracts we implement to elicit preferences feature a completion 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 vaccinators from reneging on their contract choices. At any given point in time, not completing allocated vaccinations generates a 1000 rupee penalty. As such, vaccinators should forecast that they will indeed complete the required vaccinations and

[46] 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 vaccinator's 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.

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

allocate them according to their preferences. If vaccinators forecast not being able to complete the required vaccinations, then allocation choices may not reflect their true preferences, and hence be effectively cheap talk.[48]

Our monitoring application allows us to assess, ex-post, the extent to which vaccinators succeeded or failed to reach their allocated targets. We have comprehensive completion data for 288 of 338 vaccinators in Drive 1.[49] Of these 288, 142 (49.3%) met or exceeded their targets of v_1 and v_2. The data indicate that many who fell short came close to completion. 207 vaccinators (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.9 presents the histogram of average completion proportions across subjects.[50] Two hundred and two vaccinators (70.1%) completed an average of 70% or more of their target vaccinations. These ex-post completion rates indicate that a majority of vaccinators likely forecasted reaching their targets. Additionally, in Appendix Figure A.6 we plot for each half-hour of Drive 1 the total number of attempted vaccinations along with the probability of successful vaccination and the probability that no child was reported as present. Reporting that no child was present is likely to be less time consuming than a successful vaccination and easier to falsify. The vast majority of vaccination activity occurs before 3:00pm, and we find evidence that vaccinators' proportion of successful or failed vaccination attempts remains largely steady throughout the workday. This further suggests that allocated vaccination attempts are conducted with due diligence.

In Appendix Table A.8, we investigate the relationship between allocation timing, measured patience and completion rates. Two general patterns arise. First, vaccinators 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 vaccinators delay vaccinations believing they will complete them, but subsequently fail to follow through.

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

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

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

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

Assumption 6: No Biases in Choice: Our study assumes that the allocation environment itself induces no biases in choice such that vaccinator 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 vaccinators (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 vaccinators (17%) are within 5 vaccinations of $v_1 = v_2$.[51] Only 35 of 338 vaccinators (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 vaccinators (25%) are within 5 vaccinations of $v_1 = Rv_2$.[52] Taken together, this suggests that biases towards the middle of the budget constraint or towards equal allocation are unlikely to be driving substantial portions of allocation behavior.

4.2 Tailoring Robustness Tests

Our Drive 2 data show that vaccinators who are given bonus contracts with a value of R equal to their estimated discount factors provide significantly smoother service. Here we examine robustness of this result to alternative comparison groups, alternative measures for smoothness

[51] As an even less conservative measure, 145 of 338 (43%) satisfy $v_2 - 10 \leq v_1 \leq v_2 + 10$.

[52] As an even less conservative measure, 137 of 338 (40.5%) satsify $Rv_2 - 10 \leq v_1 \leq Rv_2 + 10$.

in service provision, and alternative measures for treatment. We conclude the section by providing results from a set of additional exercises assessing the value of atheoretic approaches to tailoring, and the possibility for alternative interventions based on different policy preferences.

4.2.1 Alternative Comparison Groups

Our results demonstrate that, relative to a comparison group with uniform random values of R, tailoring serves to reduce distance to the policy target by around one-third. A natural question is whether these tailoring benefits are observed relative to alternative controls. In Appendix Table A.10 and Appendix Figures A.10 and A.11, we present three additional analyses. A first natural control is the use of a single value of R applied to all individuals. Unfortunately, it is not possible to compare only a single value of R given our uniform random assignment protocol. However, we can examine a section of the untailored group around a given value. In column (1) of Table A.10, we repeat the analysis of Table 4, column (3), but use as the comparison group only those untailored vaccinators who received a value of R within one standard deviation of their group's mean R_i^* of 1.036. Tailoring continues to decrease the distance from smooth provision relative to this more limited control group. Appendix Figure A.10 provides corresponding graphical analysis. In column (2), we repeat this analysis excluding those individuals from the untailored group who randomly received a value of $\tilde{R}_i$ within 0.10 of their true value of R_i^*. The benefits of tailoring are observed with increased precision.[53] A second potential control group would be a subset of the untailored group who receive the same distribution of R as those in the tailored group. Matching on the 1st, 5th, 10th, 25th, 50th, 75th, 95th and 99th percentiles, column (3) and Appendix Figure A.11 demonstrate that relative to a control group receiving a matched distribution of R, tailoring continues to significantly reduce the distance from smooth provision.

[53]Comparing tailored individuals to those who received close to the untailored group's mean value of R_i^* is important because, in principle, the mean value of R^*i should yield smooth provision for the average subject. Not only is the average distance for these comparison groups substantial, 0.132 to 0.150, but the tailoring yields additional benefits at the individual level by leveraging the heterogeneity in discount factors across vaccinators.

4.2.2 Alternative Measures for Smooth Provision

Our analysis measures the distance to equal provision using the metric $|\frac{v_{1,i}}{v_{2,i}} - 1|$. In Table A.11, we reconduct the analysis of Table 4, using five alternate measures for smoothness. Panel A presents the Euclidean distance to the 45 degree line, $\frac{|v_{1,i}-v_{2,i}|}{\sqrt{2}}$. Panel B presents the Euclidean distance normalized by the total number of vaccinations allocated, $\frac{|v_{1,i}-v_{2,i}|}{\sqrt{2}(v_{1,i}+v_{2,i})}$. Panel C presents the number of sooner vaccinations that would need to be reallocated to reach the 45 degree line, $|v_{1,i} - \frac{300}{1+R}|$. Panel D presents probit regressions for needing to reallocate more than 25 vaccinations, $|v_{1,i} - \frac{300}{1+R}| > 25$. And finally, Panel E presents the value of the policymaker's objective function, $min[v_{1,i}, v_{2,i}]$. 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.

4.2.3 Alternative Sample Restrictions and Treatment Measures

Our tailoring exercise focused on vaccinators with discount factors between 0.75 and 1.5. Of 337 vaccinators in Drive 2, 280 satisfied this requirement. Those vaccinators 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 is not a binary treatment, but rather a continuous difference between their discount factor and the exogenously given one. Indeed, for all vaccinators in the untailored group, treatment is also a continuous measure. In Table A.12, Panel A, we reconduct the analysis of Table 4 using as the measure of treatment the absolute difference between each vaccinator's discount factor and their assigned interest rate, which we label *Tailor Intensity*. The main results are reproduced; the closer discount factors are to interest rates, the smoother is provision.[54] In Panel B, we include those individuals in the Boundary Sample with discount factors that lie outside of the bounds of interest rate

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

assignment. Including these observations does not alter the conclusions; however, it should be noted that treatment is no longer orthogonal to individual preferences as extremely patient and impatient vaccinators will receive larger treatment intensity on average.[55]

4.2.4 Atheoretic Approaches

Our exercise demonstrates that structural estimates of time preferences can be useful in predicting subsequent behavior. An alternative to developing a structural model of choice would rely on the researcher recovering the relationship between key parameters of interest, R and $\mathbf{1}_{d=1}$, and behavior, and developing a subsequent prediction without filtering the relationship through the structural model.

We examine one example of such an exercise. For Drive 1 behavior, we recover the relationship between v_1 and R and $\mathbf{1}_{d=1}$ by conducting a Least Absoluate Shrinkage and Selection Operator (LASSO) regression with penalty parameter chosen via 10-fold cross validation (Tibshirani, 1996) of v_1 on a cubic polynomial in R interacted with $\mathbf{1}_{d=1}$.[56] The corresponding selected lasso coefficients deliver an atheoretic representation of the most predictive (in terms of cross-validated mean squared error) relationship between v_1 and key parameters of interest in Drive 1. Given the values of R and $\mathbf{1}_{d=1}$ provided in Drive 2 and the Drive 1 lasso coefficients, we predict the Drive 2 value of v_1 for each individual.[57] Similarly, we predict Drive 2 value of v_1 from the assigned values of R and $\mathbf{1}_{d=1}$ and the individual discount factor identified in Drive 1. In Appendix Figure A.12 we examine the predictive validity of the structural versus the non-structural approaches by plotting predicted and actual values of v_1 in Drive 2 for the 337 subjects who participated in both Drive 1 and Drive 2. We also provide non-parametric

[55] Using the indicator for tailoring would not be an appropriate solution to this problem as tailored vaccinators 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.

[56] The provided regressors are a constant and normalized values of $R, R^2, R^3, \mathbf{1}_{d=1}, R \times \mathbf{1}_{d=1}, R^2 \times \mathbf{1}_{d=1}$, and $R^3 \times \mathbf{1}_{d=1}$. The lasso regression and cross-fold validation procedure were implemented using the glmnet package in R.

[57] Note that this prediction will be identical for all vaccinators given the same value of R and $\mathbf{1}_{d=1}$, and hence will mispredict any heterogeneity across vaccinators.

lowess curves for the relationship between predicted and actual values.[58] Notable from Appendix Figure A.12 is the generally close adherence between structurally predicted and actual values. The lowess line follows the 45 degree line of perfect prediction through the majority of the space. Less adherence is observed between non-structural predictions and actual behavior. The structural predictions deliver higher correlations with real behavior ($\rho = 0.18$) than do the non-structural predictions ($\rho = 0.16$) and lower bias in prediction (27.9 vs. 32). However, the root mean squared error is lower for the non-structural predictions (32.8 vs 39.1). Hence, on the basis of bias and correlation, our structural exercise outperforms the machine learning lasso algorithm trained on Drive 1.

There is an additional important difference between the LASSO and structural predictions presented here. The structural prediction for a given vaccinator's behavior in Drive 2 derives from their allocation in Drive 1 and the interest rate assigned in Drive 2. It therefore relies on a single data point. By contrast, the LASSO prediction for a given vaccinator is derived by training a model on the entire cross-section of data from Drive 1, and then obtaining a fitted value based on the interest rate assigned in Drive 2. This means the LASSO prediction will be identical for all subjects assigned the same interest rate while the structural prediction may vary depending on individually measured preferences. Hence, the two methods are not only different in the sense of being structural or atheoretic, they also differ both in the amount of data informing the prediction and in the extent of heterogeneity they can predict.

4.2.5 Alternative Policy Preferences

While our results suggest that 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 vaccinator time preferences are stable, then changes in incentives will have predictable effects on behavior.

[58] The lasso procedure on Drive 1 selects an intercept and R as delivering coefficients of sufficient size given the cross-validated constraint choice.

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 vaccinator in each Drive. Figure A.13 presents the calculated discount factors for Drive 1 and Drive 2 along with the 45 degree line for 317 of 337 vaccinators.[59] The correlation in discount factors across rounds is $\rho = 0.41,\ (p < 0.01)$, indicating stability in preferences.[60] 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 may also be achievable with tailored contracts.[61]

5 Conclusion

Structural parameters for intertemporal preferences have been at the center of theoretical and empirical research modeling intertemporal choice for much of the last century. This paper seeks to understand a heretofore unexplored question: are the out-of-sample predictions given by structural estimates of discounting empirically valid? We couch this question in an effort to customize contracts for 337 vaccination workers who spend two days each month attempting to deliver polio vaccines in the neighborhoods of Lahore, Pakistan.

[59] Eliminated from the figure and from our calculations of stability are 20 vaccinators with discount factors in excess of two in one or both drives.

[60] Including the remaining 20 extreme vaccinators, 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 $\tilde{R}_i$ within 0.25 of their value of R_i^*, the correlation in discount factors is $\rho = 0.53,\ (p < 0.01)$.

[61] 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.

We monitor our workers' efforts using a smartphone application developed especially for our project. Workers in the Advanced condition state their targets for both days of the vaccination drive on the Friday before the drive, which is conducted on the following Monday and Tuesday. Those in the Immediate condition wait until Monday morning to state their targets. Those who reach their targets get a 1000 rupee bonus (around $10 US). As anticipated, subjects stating targets on Monday morning are skewed to delaying vaccinations until day 2 of the drive. That is, vaccinators exhibit a present bias in effort allocations. With assumptions on costs of delivering vaccines we are able to identify (somewhat rough) estimates of discounting parameters for each of our workers. In the second stage of our study, conducted a month later, half of workers were offered a contract tailored to their own discounting parameters, designed to induce equal provision of vaccinations on both days of the drive. The initial preference measures are critical to the design of these contracts as, without a measure of preferences, there would be no prescription for contract terms. The policy objective of equal provision of vaccinations on both days is admittedly arbitrary. Hence, we view our exercise as a proof-of-concept for the possibility of tailored incentives.

Our finding is encouraging. Those workers who receive effort contracts that were tailored to their individual discounting parameters were significantly more likely to meet the policy objective relative to untailored workers. That is, using structurally identified estimates of discounting parameters to form a new incentive contract can indeed have a predictable effect on behavior. To date, little research makes use of such predictive value of structural discounting estimates. Our results show not only that estimates are predictive, but also that useful parameter estimates are identifiable from a very limited number of experimental choices. This suggests that the substantial effort of articulating and estimating structural models in this domain has been well-invested.

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

settings (Acemoglu, 2010; Banerjee, Chassang and Snowberg, 2016).[62] In our setting, translating from our reduced form experimental treatment effects to a structural model of choice requires a set of potentially strong (and implausible) assumptions.[63] Nonetheless, the finding of predictive validity in this case suggests there is indeed potential for using structure as a means of increasing the external validity of results obtained from a single sample.

Separately, our results link to the growing literature on the personnel economics of the state (Ashraf, Bandiera and Lee, 2015; Bertrand et al., 2016; Finan, Olken and Pande, Forthcoming; Dal Bó, Finan and Rossi, 2013; Deseranno, 2016; Callen, Gulzar, Hasanain, Khan and Rezaee, 2016). This literature emphasizes the idea that states play a vital role in delivering services and facilitating economic growth, and so their internal dynamics should be studied with the same degree of attention as has been applied to firms. Within this literature, there is interest in understanding whether heterogeneity in competencies and in motivation of state actors is linked to meaningful differences in state performance or service provsion (Ashraf et al., 2015; Dal Bó et al., 2013; Deseranno, 2016; Callen et al., 2016). We take the additional step of asking not only whether this heterogeneity matters for outcomes, but also whether it can be acknowledged and reflected in the design of individual incentives.

There are a number of clear limitations to our study which should be addressed by future research. First, our study sidesteps the critical issue of incentive compatibility by not informing subjects of 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. Second, future research should seek to gain more precise estimates of preferences. Our exercise requires restrictive assumptions that could be relaxed in the presence of more data. If our results point to a lower bound in the promise of tailored contracts, it is important to know how much more can be achieved. Third, alternative policy objectives and contract types should be investigated to

[62] Attanasio and Meghir (2012), Duflo, Hanna and Ryan (2012), and Duflo, Greenstone, Pande and Ryan (2016) provide examples in development of using experiments to estimate key policy parameters.

[63] Banerjee et al. (2016) discuss how the plausibility of such identifying assumptions might limit external validity.

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.

References

Acemoglu, Daron, "Theory, General Equilibrium, and Political Economy in Development Economics," *Journal of Economic Perspectives*, 2010, *24* (3), 17–32.

Ackerberg, Daniel A. and Maristella Botticini, "Endogenous Matching and the Empirical Determinants of Contract Form," *Journal of Political Economy*, 2002, *110* (3), 564 – 591.

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.

_ , Oriana Bandiera, and Scott S. Lee, "Do-Gooders and Go-Getters: Career Incentives, Selection, and Performance in Public Sector Service Delivery," 2015.

Attanasio, Orazio and Costas Meghir, "Education Choices in Mexico: Using a Structural Model and a Randomized Experiment to Evaluate PROGRESA," *Review of Economic Studies*, 2012, *79* (1), 37–66.

Augenblick, Ned and Matthew Rabin, "An Experiment on Time Preference and Misprediction in Unpleasant Tasks," 2015.

_ , Muriel Niederle, and Charles Sprenger, "Working Over Time: Dynamic Inconsistency in Real Effort Tasks," *Quarterly Journal of Economics*, 2015, *130* (3), 1067–1115.

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

_ , Sylvain Chassang, and Erik Snowberg, "Decision Theoretic Approaches to Experiment Design and External Validity," 2016.

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.

Basinga, Paulin, Paul Gertler, Agnes Binagwaho, Agnes Soucat, Jennifer Sturdy, and Vermeesch, "Effects of Maternal and Child Health Services in Rwanda of Payment to Primary Health-Care Providers for Performance: An Impact Evaluation," *The Lancet*, 2011, *377* (9775), 1421–1428.

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.

Bellemare, Charles and Bruce Shearer, "Multidimensional Heterogeneity and the Economic Importance of Risk and Matching: Evidence from Contractual Data and Field Experiments," *Rand Journal of Economics*, 2013, *44* (2), 361–389.

Bertrand, Marianne, Robin Burgess, Arunish Chawla, and Guo Xu, "Determinants and Consequences of Bureaucratic Effectiveness: Evidence from the Indian Administrative Service," 2016.

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.

_ , _ , _ , _ , **and Arman Rezaee**, "Personalities and Public Sector Performance: Evidence from a Health Experiment in Pakistan," 2016.

Carvalho, Leandro S., Stephan Meier, and Stephanie W. Wang, "Poverty and Economic Decision-Making: Evidence from Changes in Financial Resources at Payday," *Working Paper*, 2014.

Castillo, Marco, Paul Ferraro, Jeffrey Jordan, and Ragan Petrie, "The Today and Tomorrow of Kids: Time Preferences and Educations Outcomes of Children," *Journal of Public Economics*, 2011, *95* (11-12), 1377–1385.

Chabris, Christopher F., David Laibson, and Jonathon P. Schuldt, "Intertemporal Choice," in Steven N. Durlauf and Larry Blume, eds., *The New Palgrave Dictionary of Economics*, London: Palgrave Macmillan, 2008.

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

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).

Chiappori, Pierre-André and Bernard Salanié, "Testing Contract Theory: a Survey of Some Recent Work," in "Advances in Economics and Econometrics - 8th World Congress, Econometric Society Monographs" 2003.

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.

Dal Bó, Ernesto, Frederico Finan, and Martín A. Rossi, "Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service," *Quarterly Journal of Economics*, 2013, *forthcoming.*

Deseranno, Erika, "Financial Incentives as Signals: Experimental Evidence From the Recruitment of Health Promoters," 2016.

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde, "Dynamic Inconsistency Predicts Self-Control Problems in Humans," *Working Paper*, 2006.

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

Dubois, Pierre, "Moral Hazard, Land Fertility, and Sharecropping in a Rural Area of the Philippines," *Journal of Development Economics*, 2002, *68* (1), 35 – 64.

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

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

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

Fang, Hanming and Yang Wang, "Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions," *International Economic Review*, 2015, *56* (2), 565–596.

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

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.

Haubrich, Joseph, "Risk Aversion, Performance Pay, and the Principal-Agent Problem," *Journal of Political Economy*, 1994, *102* (2), 258 – 276.

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

Jensen, Michael and Kevin Murphy, "Performance Pay and Top-Management Incentives," *Journal of Political Economy*, 1990, *98* (2), 225 – 264.

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

_ , _ , and _ , "Self-Control at Work," *Journal of Political Economy*, 2015, *123* (6), 1227–1277.

Khan, Adnan, Asim Khwaja, and Benjamin Olken, "Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors," *Quarterly Journal of Economics*, 2015, *131* (1), 219–271.

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.

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.

Lazear, Edward, "Performance Pay and Productivity," *American Economic Review*, 2000, *90*, 1346 – 1361.

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

Mahajan, Aprajit and Alessandro Tarozzi, "Time Inconsistency, Ecpectations and Technology Adoption: The Case of Insecticide Treated Nets," *Working Paper*, 2011.

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.

Miller, Grant, Renfu Luo, Linxiu Zhang, Sean Sylvia, Yaojiang Shi, Patricia Foo, Qiran Zhao, Reynaldo Martorell, Alexis Medina, and Scott Rozelle, "Effectiveness of Provider Incentives for Anaemia Reduction in Rural China: a Cluster Randomized Trial," *British Medical Journal*, 2012, *345.*

Muralidharan, Karthik and Venkatesh Sundararaman, "Teacher Performance Pay: Experimental Evidence from India," *Journal of Political Economy*, 2011, *119* (1), 39–77.

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.

Olken, Benjamin, Junko Onishi, and Susan Wong, "Should Aid Reward Performance? Evidence from a File Experiment on Health and Education in Indonesia," *American Economic Journal: Applied Economics*, 2014, *6* (4), 1–34.

Paarsch, Harry and Bruce Shearer, "The Response of Worker Effort to Piece Rates: Evidence from the British Columbia Tree-Planting Industry," *Journal of Human Resources*, 1999, *34*, 643–667.

_ **and** _ , "The Response to Incentives and Contractual Efficiency: Evidence From a Field Experiment," *European Economic Review*, 2009, *53* (481 - 494).

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

Prendergast, Canice, "The Provision of Incentives in Firms," *Journal of Economic Literature*, 1999, *1*, 7 – 63.

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, Matthew D., "The Permanent Income Hypothesis and the Real Interest Rate: Some Evidence from Panel Data," *Economics Letters*, 1984, *14* (1), 93–100.

Shearer, Bruce, "Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment," *Review of Economic Studies*, 2004, *71*, 513–524.

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.

Tibshirani, Robert, "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, 1996, *58* (1), 267–288.

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

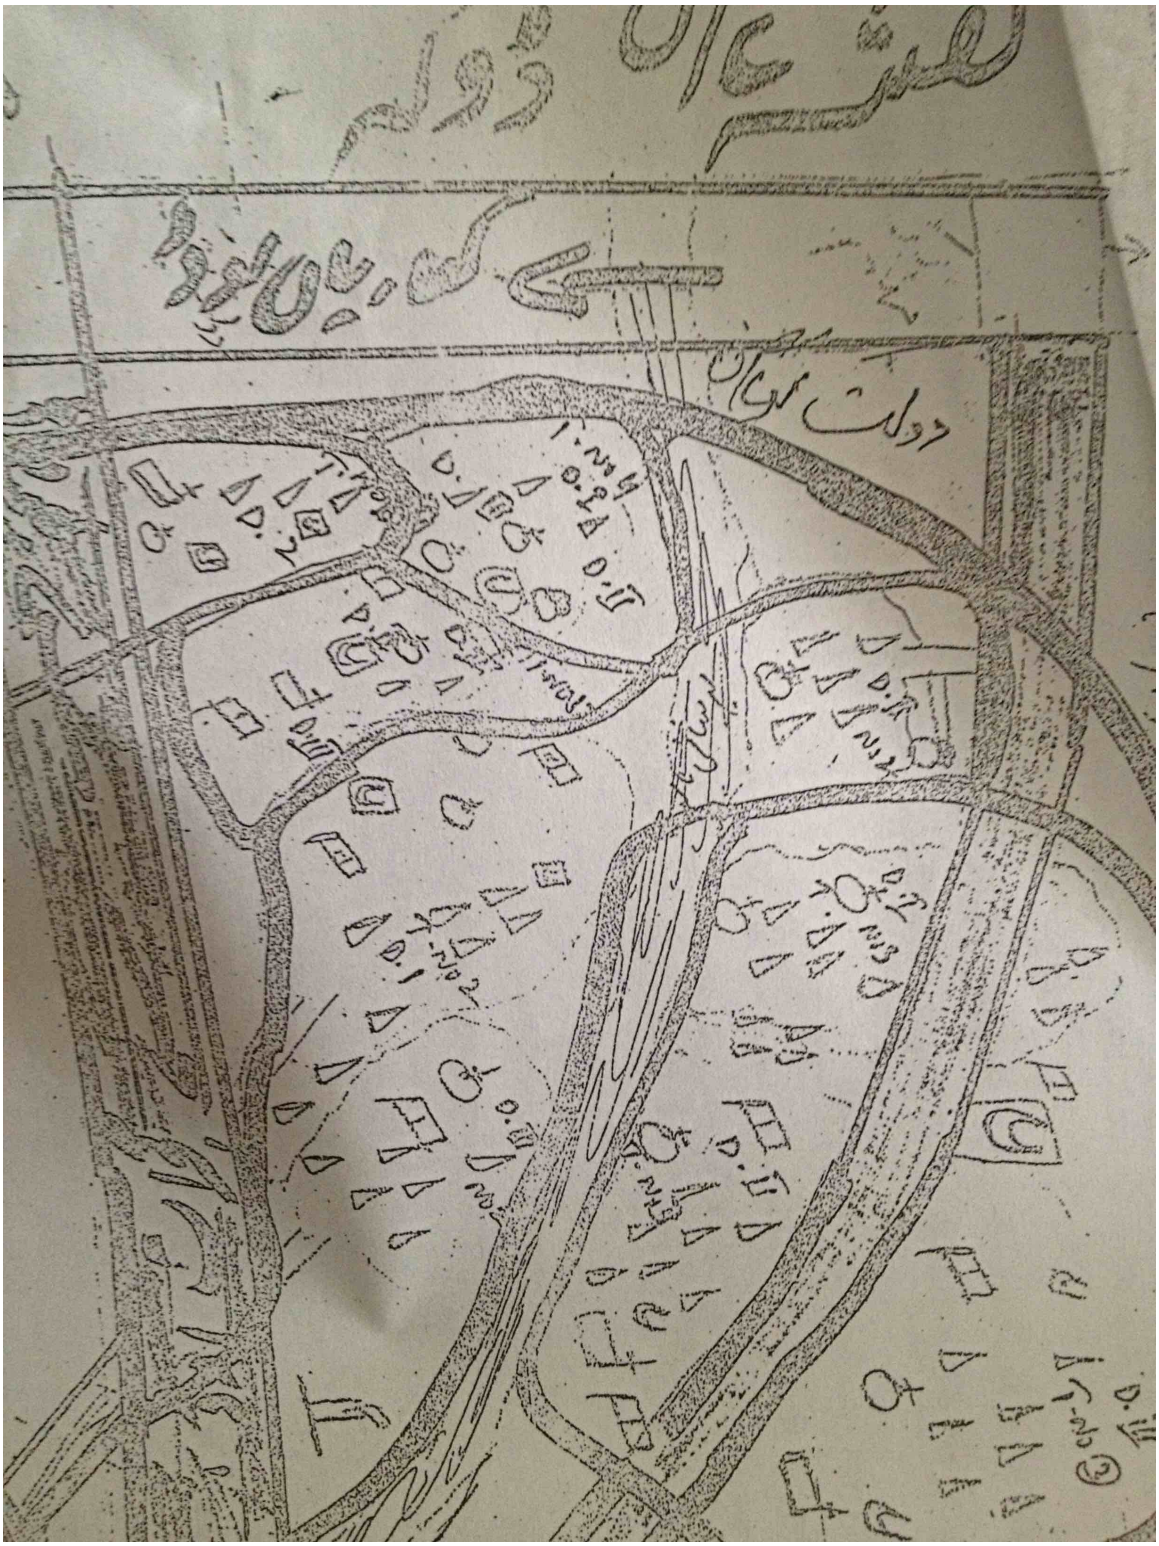

Figure A.1: Map Given to Vaccinators to Plan Route

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

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

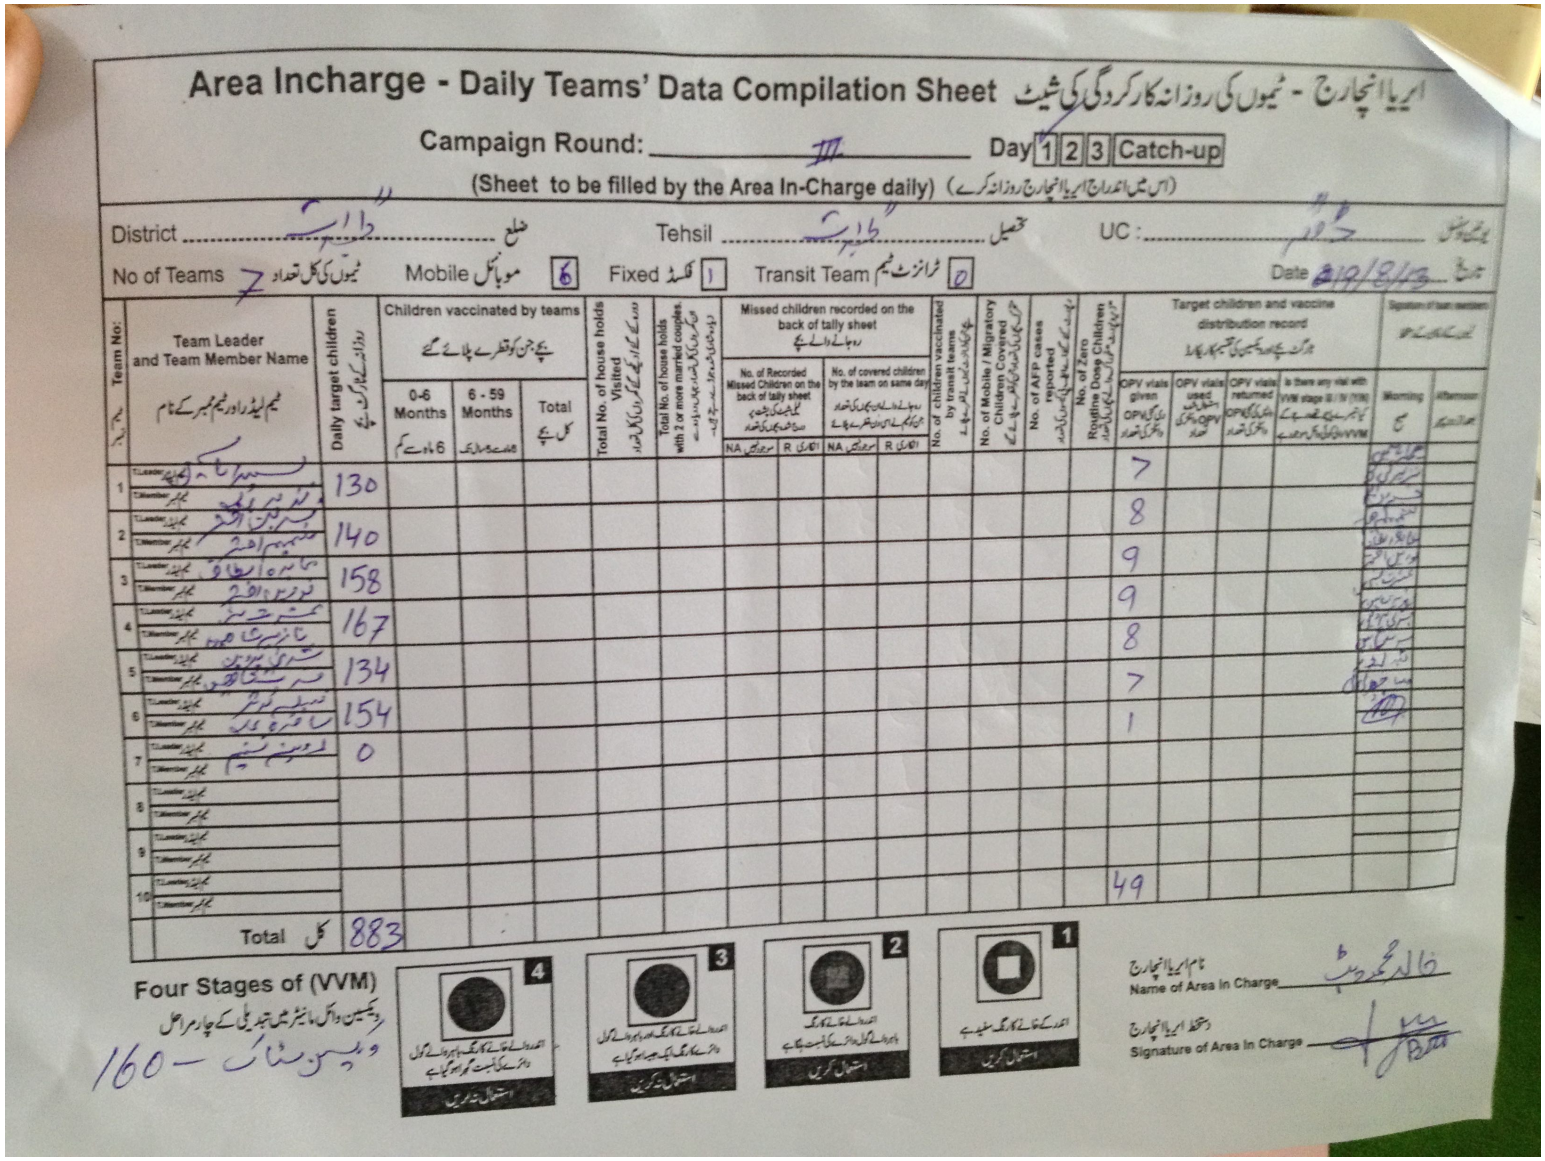

Area Incharge - Daily Teams' Data Compilation Sheet ایریا انچارج - ٹیموں کی روزانہ کارکردگی کی شیٹ

Campaign Round: III Day 1 2 3 Catch-up

(Sheet to be filled by the Area In-Charge daily) (اس میں اندراج ایریا انچارج روزانہ کرے)

District ڈاڈر ضلع Tehsil ڈاڈر تحصیل UC: ڈاڈر یونین کونسل

No of Teams 7 ٹیموں کی کل تعداد Mobile موبائل 6 Fixed فکسڈ 1 Transit Team ٹرانزٹ ٹیم 0 Date 19/8/13 تاریخ

Team No: ٹیم نمبر	Team Leader and Team Member Name ٹیم لیڈر اور ٹیم ممبر کے نام	Daily target children روزانہ بچوں کے ہدف	Children vaccinated by teams بچے جن کو قطرے پلائے گئے: 0-6 Months 6 ماہ سے کم	6 - 59 Months 6 ماہ سے 59 ماہ تک	Total کل بچے	Total No. of house holds Visited دورے کیے گئے گھروں کی کل تعداد	Total No. of house holds with 2 or more married couples. جن گھروں میں 2 یا زیادہ شادی شدہ جوڑے ہیں	Missed children recorded on the back of tally sheet: No. of Recorded Missed Children on the back of tally sheet NA	R	No. of covered children by the team on same day NA	R	No. of children vaccinated by transit teams	No. of Mobile / Migratory Children Covered	No. of AFP cases reported	No. of Zero Routine Dosg Children	Target children and vaccine distribution record: OPV vials given OPV کی دی گئی شیشیوں کی تعداد	OPV vials used OPV کی استعمال شدہ شیشیوں کی تعداد	OPV vials returned OPV کی واپس کی گئی شیشیوں کی تعداد	Is there any vial with VVM stage II / IV (Y/N) کیا ان میں سے کوئی شیشی VVM II یا IV کی ہے	Signature of team members: Morning صبح	Afternoon
1	Leader / Member	130														7					
2	Leader / Member	140														8					
3	Leader / Member	158														9					
4	Leader / Member	167														9					
5	Leader / Member	134														8					
6	Leader / Member	154														7					
7	Leader / Member	0														1					
8	Leader / Member																				
9	Leader / Member																				
10	Leader / Member																				
	Total کل	883														49					

Four Stages of (VVM) وکسین وائل مانیٹر میں تبدیلی کے چار مراحل

160 — ویکسین سٹاک

4 — استعمال نہ کریں

3 — استعمال نہ کریں

2 — استعمال کریں

1 — استعمال کریں

نام ایریا انچارج Name of Area In Charge

دستخط ایریا انچارج Signature of Area In Charge

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

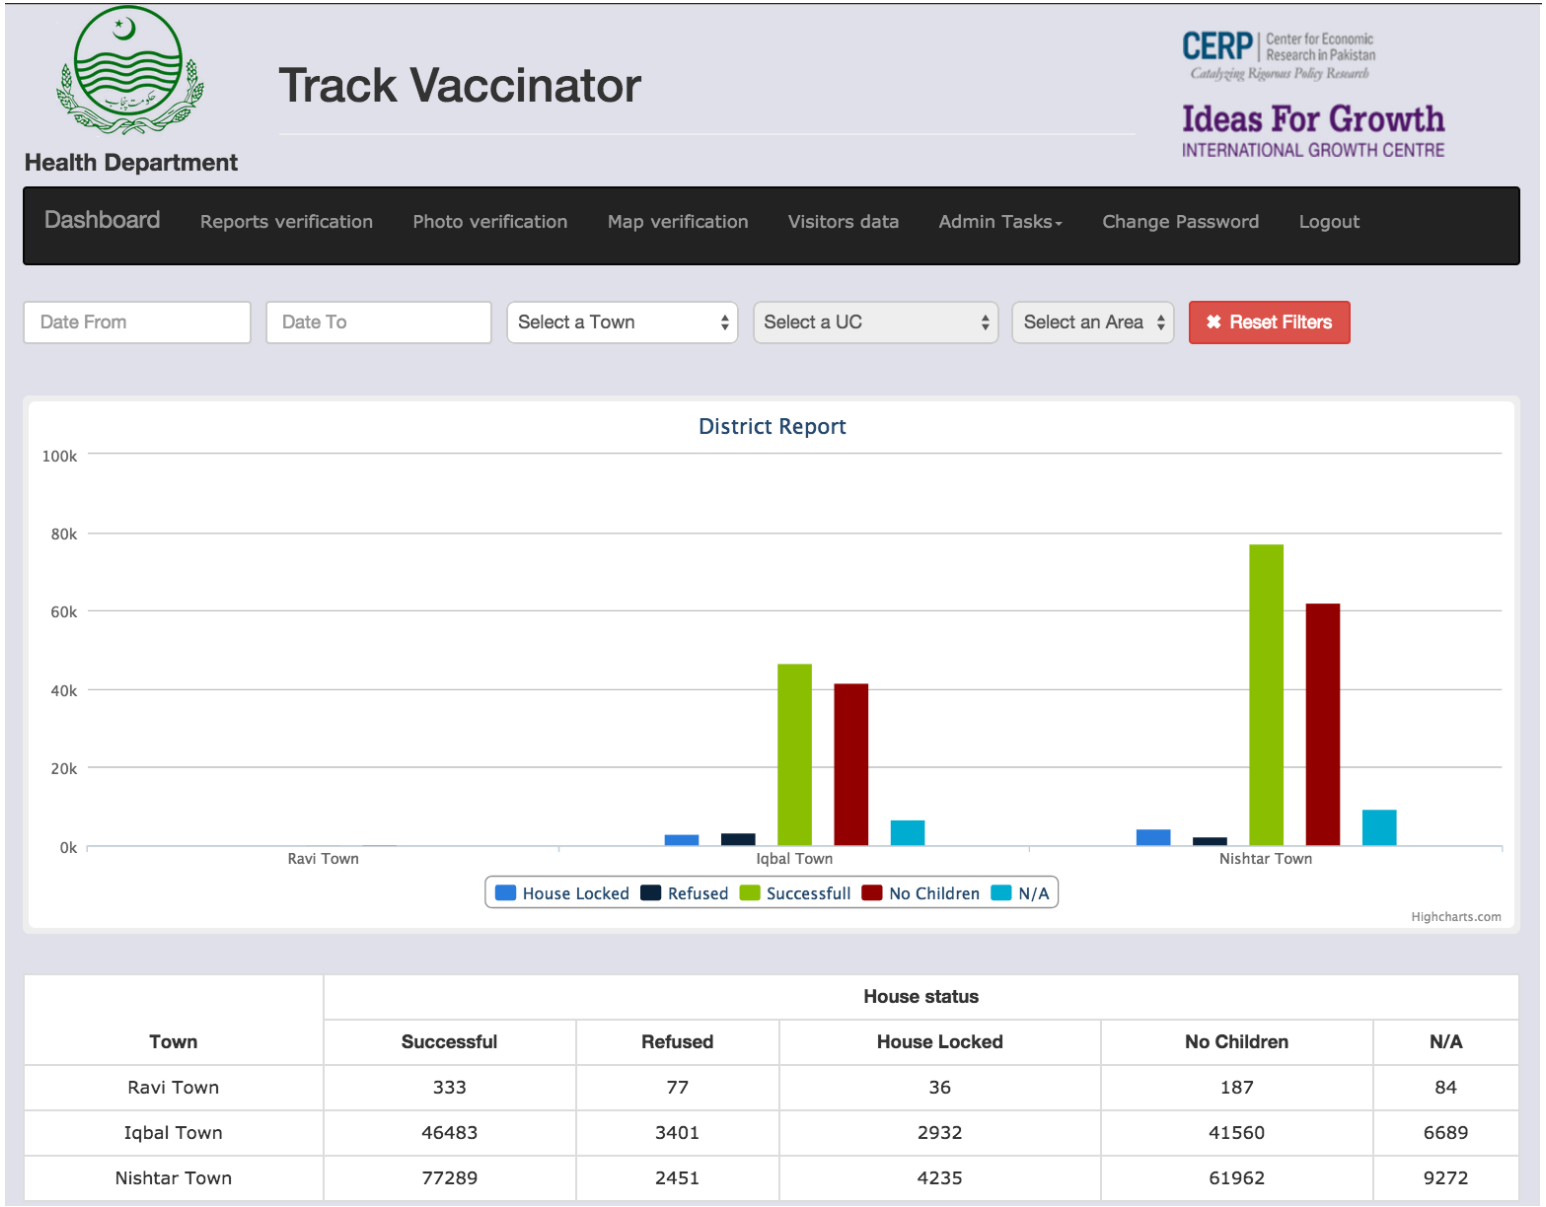

Town	House status				
	Successful	Refused	House Locked	No Children	N/A
Ravi Town	333	77	36	187	84
Iqbal Town	46483	3401	2932	41560	6689
Nishtar Town	77289	2451	4235	61962	9272

Figure A.5: Screenshot of the Track Vaccinator Dashboard

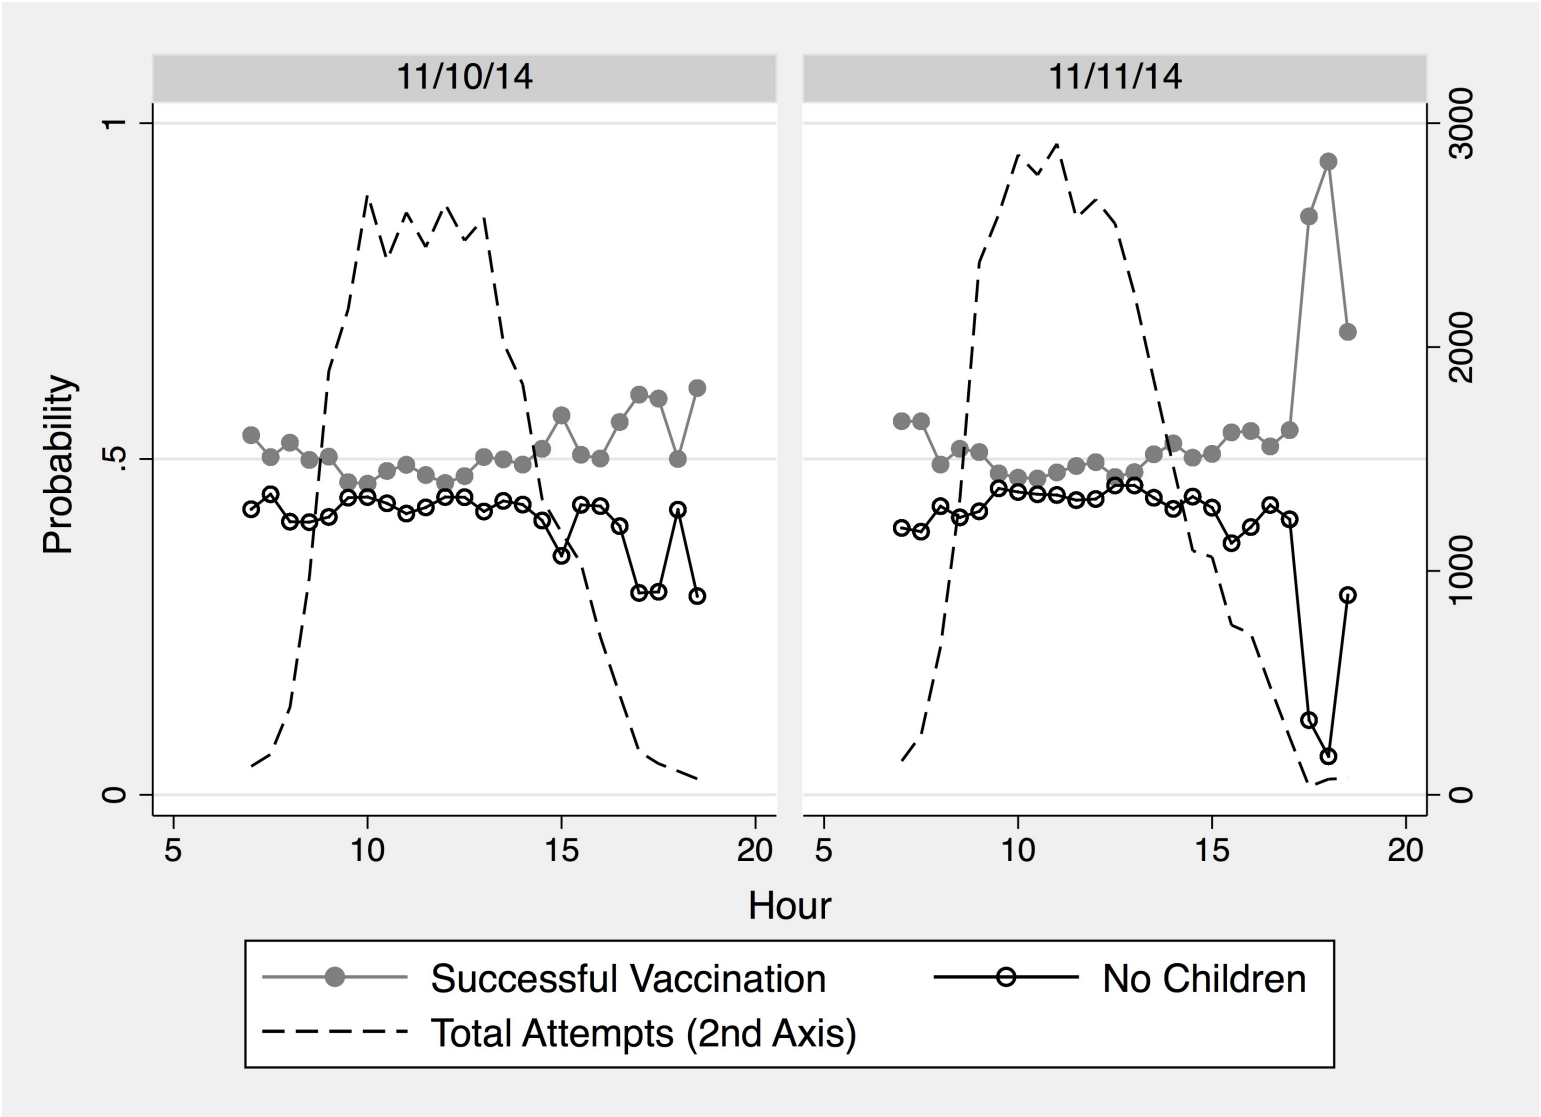

Figure A.6: Drive 1 Vaccination Activity

Notes: The solid light grey circles are the share of all vaccination attempts that reflect a successful vaccination during the indicated hour. The hollow dark black circles are the share of all vaccination attempts that report no children being available during the attempt. These quantities are compared against the left axis. The dotted line indicates the total number of vaccination attempts for all vaccinators in the sample. This quantity is compared against the right axis.

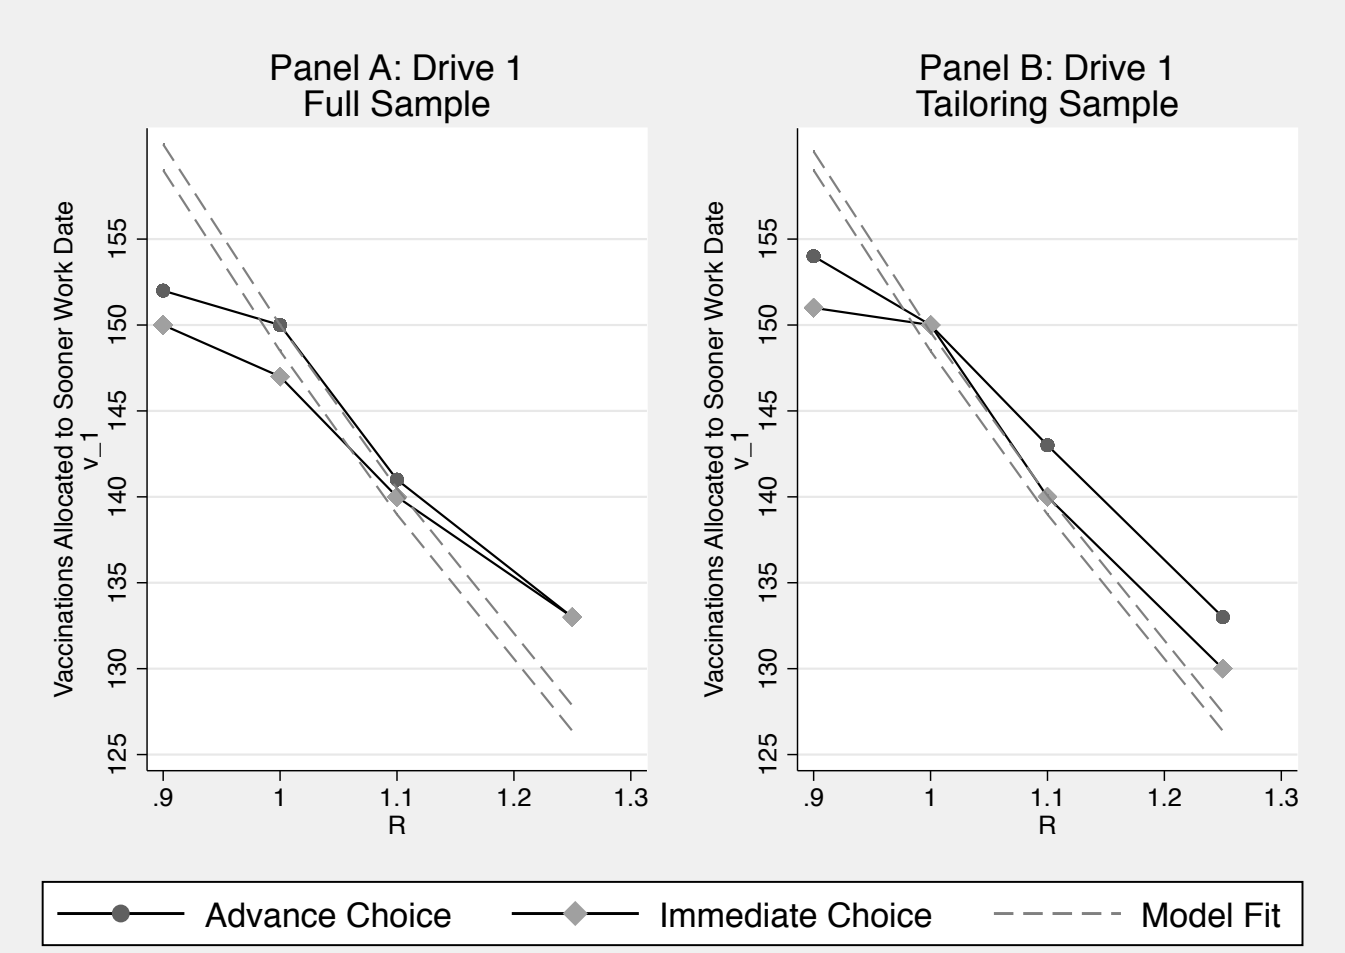

Figure A.7: Predicted and Actual Experimental Response

Notes: Points in the plots are medians for each of the eight treatment groups respectively. Panel A depicts the Full Sample and Panel B depicts the tailoring sample (vaccinators with $R^* < 0.75$ or $R^* > 1.5$). Circles are advance choice groups and diamonds are immediate choice groups. The series for model fit corresponds to predictions from Table 3, columns (2) and (4).

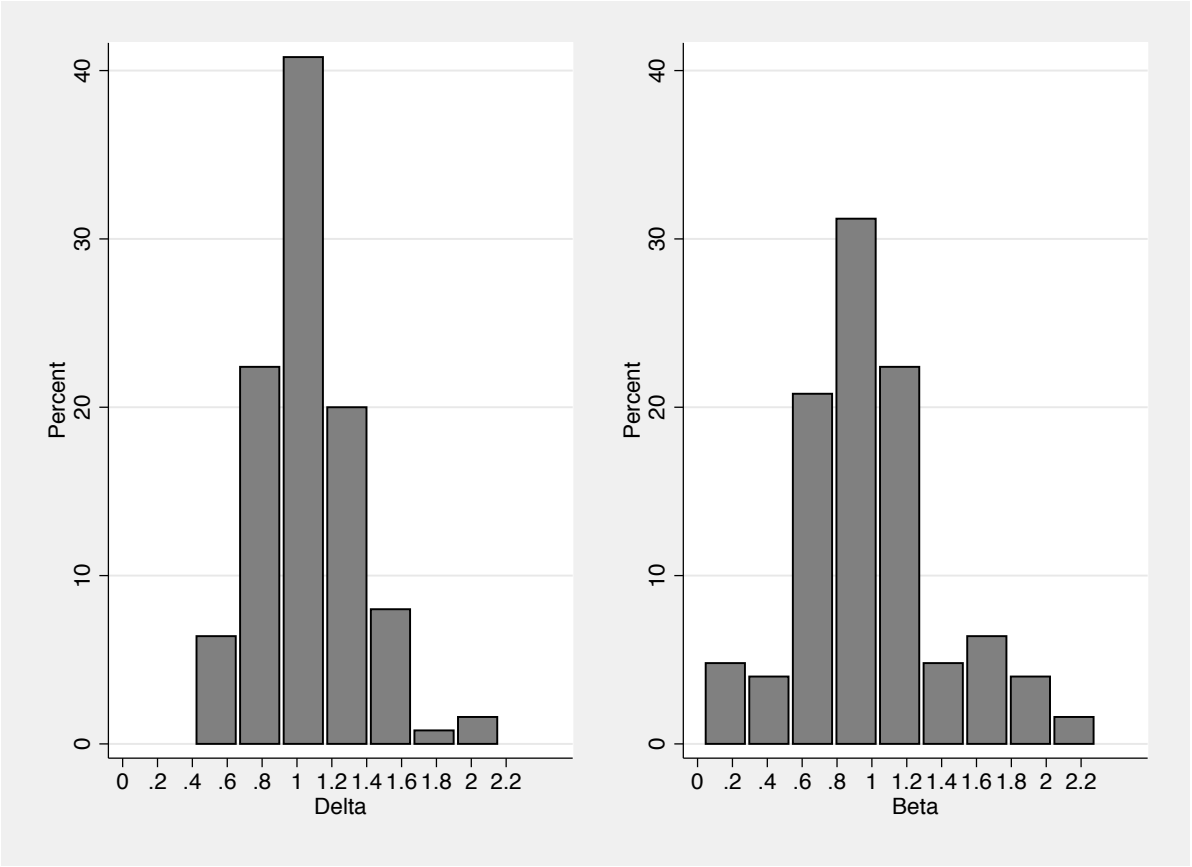

Figure A.8: Within-Subject Parameter Calculations

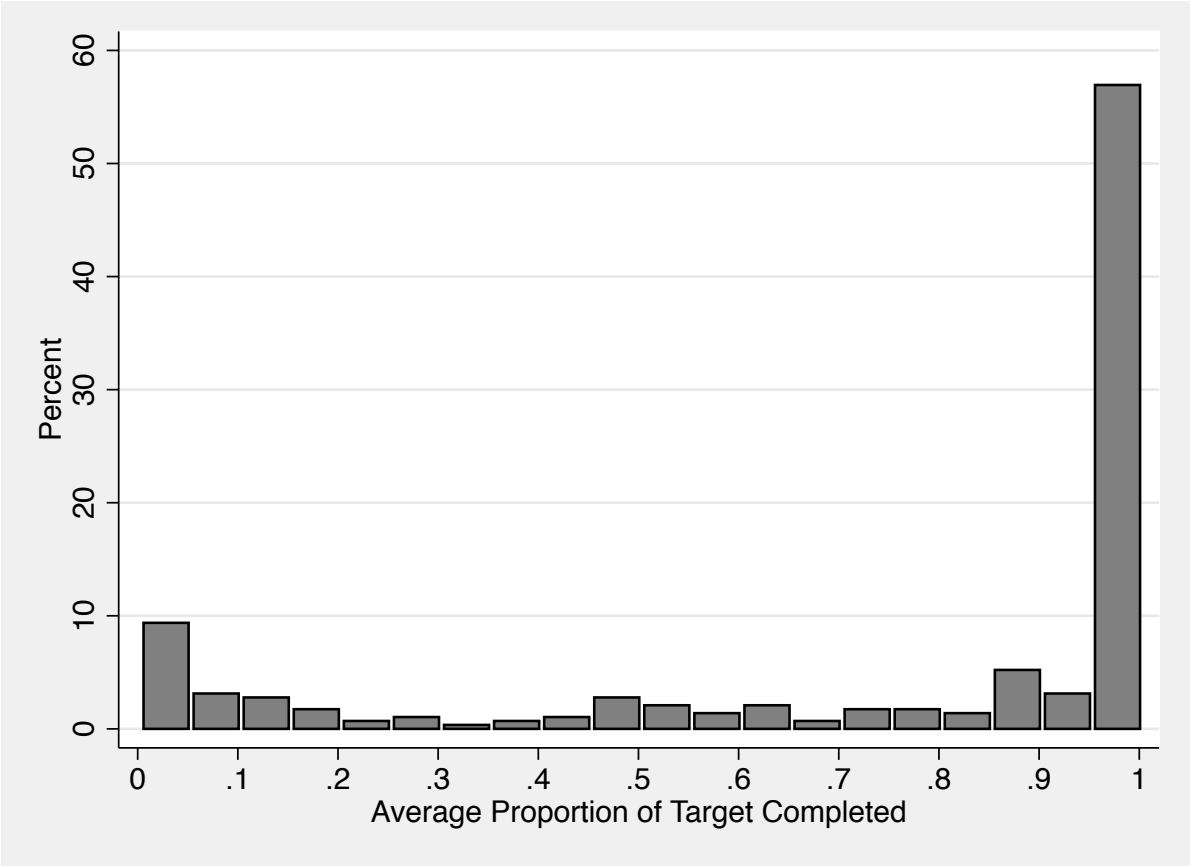

Figure A.9: Individual Completion Rates

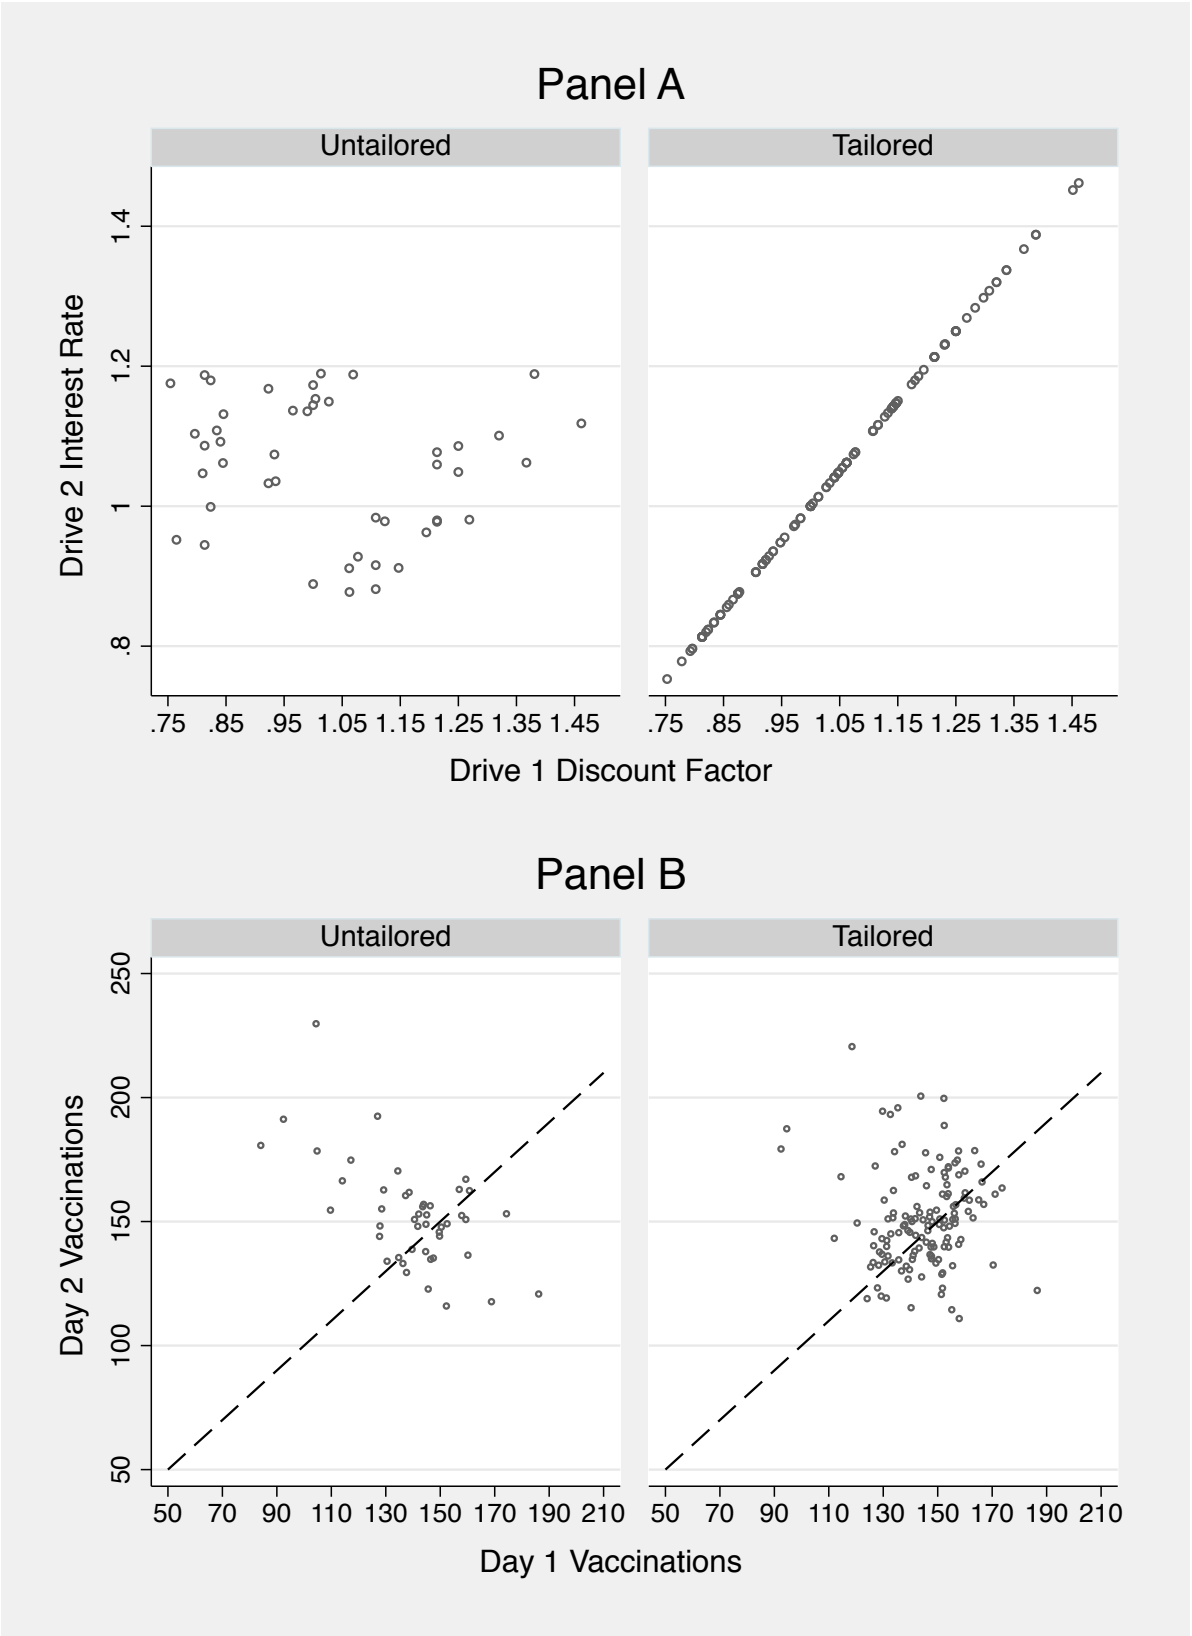

Figure A.10: Alternative Control Group: $\tilde{R}_i \in +/-$ 1 s.d of Mean R_i^*

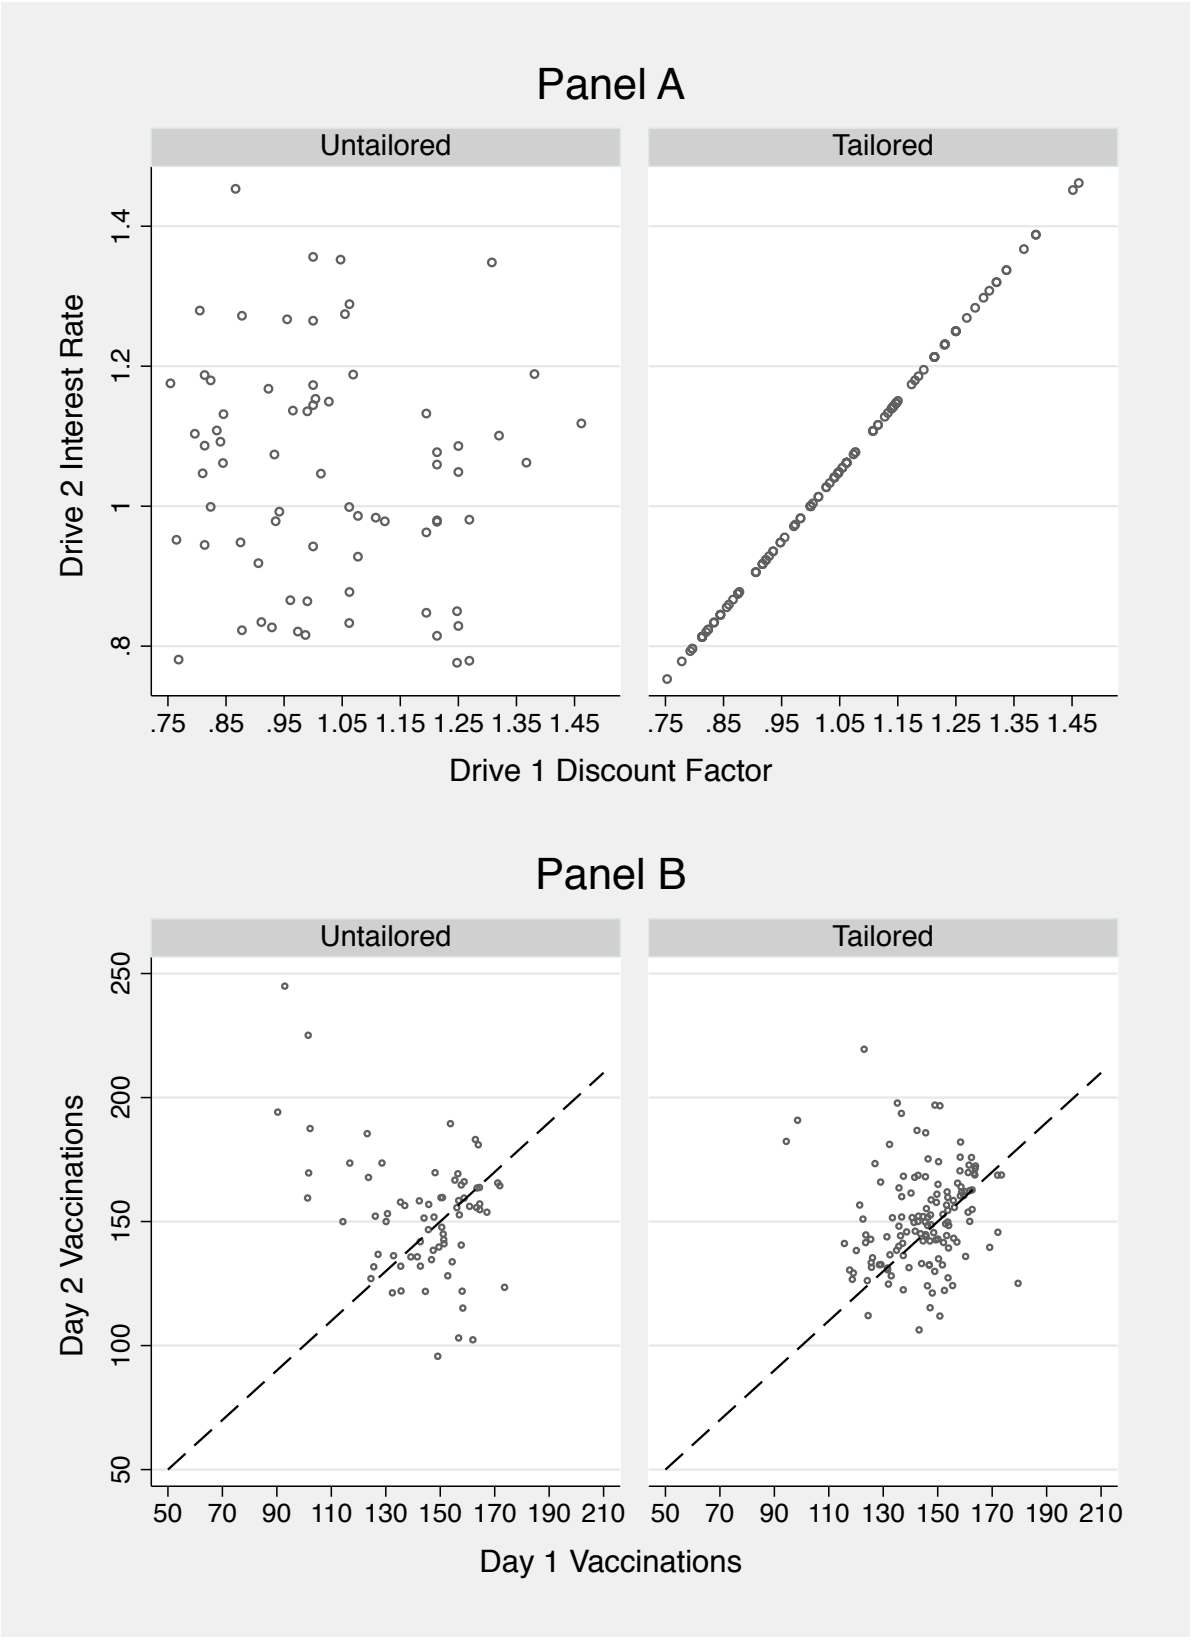

Figure A.11: Alternative Control Group: Matched Distribution

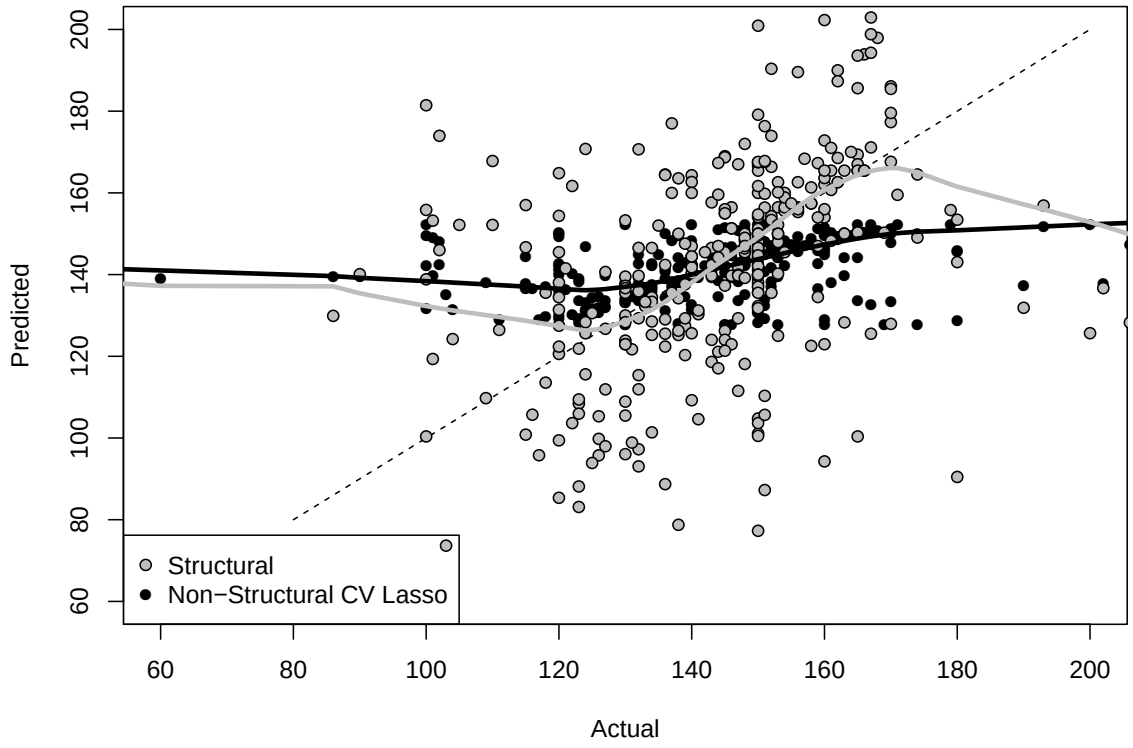

Figure A.12: Structural and Non-Structural Prediction

Notes: Predicted and actual Drive 2 value of v_1 for structural and non-structural models. Structural prediction based on individual discount factor calculated from Drive 1. Non-structural prediction based on lasso regression of v_1 on cubic polynomial in R interacted with $\mathbf{1}_{d=1}$ from Drive 1. Lowess curves for non-parametric adherence to 45-degree line of perfect prediction.

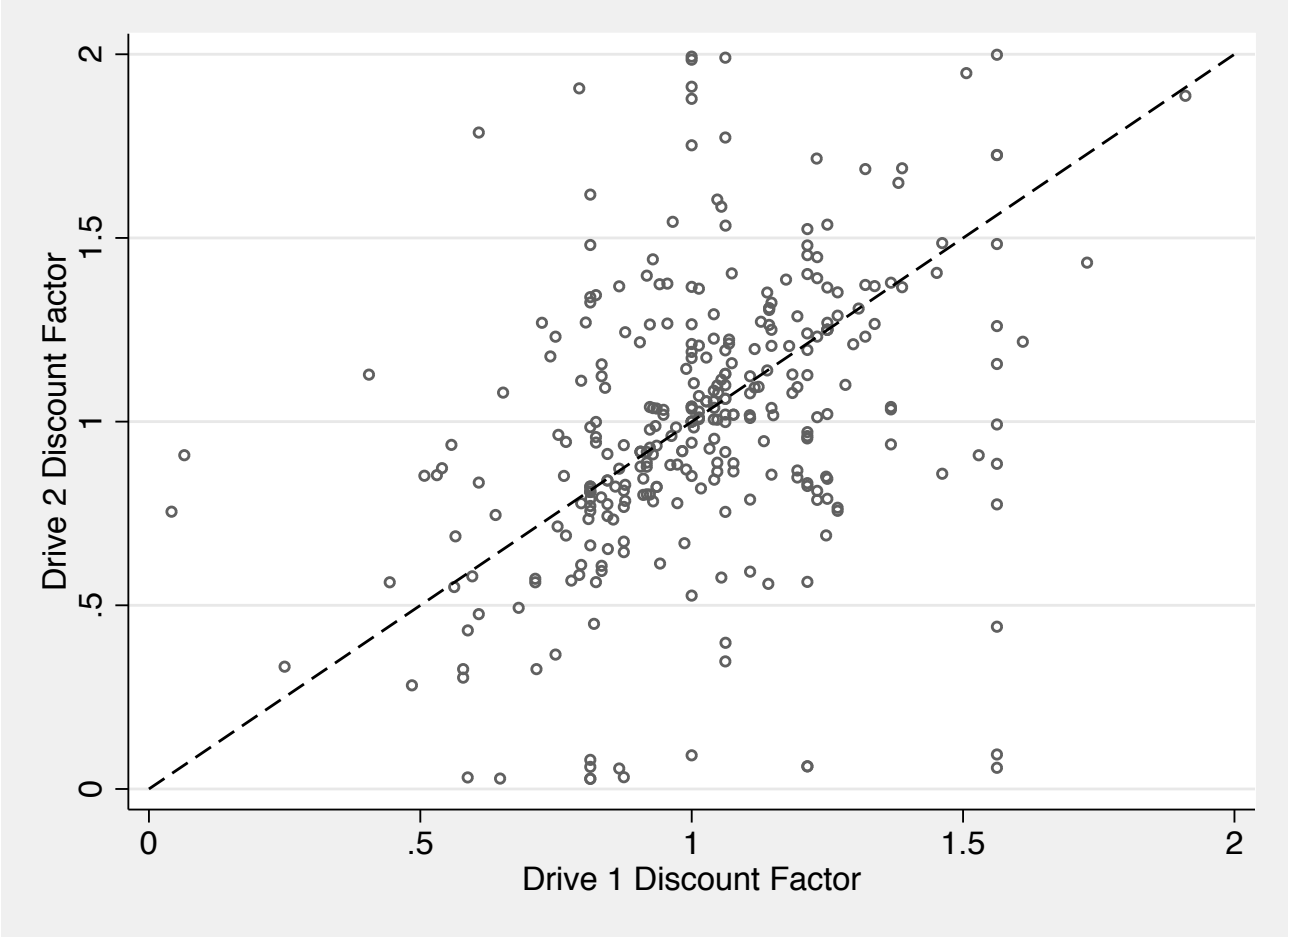

Figure A.13: Stability of Preferences

Notes: Drive 1 and Drive 2 discount factors calculated from equation (3) for each allocation. Figure includes 317 of 337 vaccinators present in both drives. Excluded are 20 vaccinators with calculated discount factors in excess of 2 in one or both drives. Correlation: $\rho = 0.41,\ (p < 0.01)$.

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	86	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	0.737	0.755	0.719	0.585	0.848
	(0.100)	(0.153)	(0.130)	(0.221)	(0.100)
δ	0.955	0.931	0.981	1.029	0.944
	(0.041)	(0.060)	(0.056)	(0.064)	(0.085)
a	-15.715	-14.049	-15.516	-30.333	-17.487
	(0.477)	(1.003)	(0.362)	(2.923)	(0.263)
$\gamma = 1 + 3 \cdot \frac{1}{1+exp(a)}$	4	4	4	4	4
$ln(\sigma)$	-0.778	-0.872	-0.711	-0.397	-1.325
	(0.127)	(0.184)	(0.170)	(0.197)	(0.105)
# Observations	464	212	252	112	140
# Vaccinators	232	106	126	56	70
Log-Likelihood	-297.467	-116.020	-178.460	-114.456	-13.099
$H_0 : \beta = 1$	$\chi^2(1) = 6.945$	$\chi^2(1) = 2.585$	$\chi^2(1) = 4.707$	$\chi^2(1) = 3.529$	$\chi^2(1) = 2.303$
	$(p < 0.01)$	$(p = 0.11)$	$(p < 0.05)$	$(p < 0.10)$	$(p = 0.13)$

Notes: This table reports structural estimates of β, δ, and γ obtained using Maximum Likelihood Estimation based on Equation (2) and using data both from Drive 0 and Drive 1. Standard errors are reported in parentheses. Column (1) reports estimates for all vaccinators who participated in both Drive 0 and Drive 1. Column (2) restricts the sample to vaccinators who participated in both rounds but whose decision timing assignment did not switch. Column (3) restricts the sample to vaccinators who participated in both rounds but whose decision timing switched. Column (4) uses the subset of vaccinators who switched from Immediate to Advance Choice. Column (5) uses the subset of vaccinators who swtiched from Advance to Immediate Choice.

Table A.3: Aggregate Drive 0 and 1 Behavior

Dependent variable:	Tasks Allocated to the First Day of the Drive (v_1)	
	Full Sample	Tailoring Sample
	(1) Median	(2) Median
Immediate Decision (=1)	-2.00**	-3.00***
	(0.95)	(0.88)
Interest Rate (R)	-40.00***	-60.00***
	(6.04)	(4.12)
Constant	188.00***	210.00***
	(6.06)	(4.34)
Median Advance Choice	150	150
# Observations	622	475

Notes: This table reports on the effects of making Immediate allocation decisions and interest rate on Drive 0 and Drive 1 vaccinations allocated to the first day of the drive. Median regression coefficients with standard errors reported in parentheses. Standard errors 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 an indicator equal to one for vaccinators selecting their allocations on the morning of the vaccination drive. The interest rate R takes the values $R \in \{0.9, 1, 1.1, 1.25\}$. *Levels of Significance:* $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.4: Testing Stationarity of Costs Across Days

Panel A: Time Lapse Between Vaccinations (in minutes)

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 (in Kilometers)

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: This table reports on the relationship between decision timing and the one period discount factor with two proxies of the cost of performing a vaccination (the amount of time that lapses between vaccinations and the distance traveled between vaccinations). Heteroskedasticity robust White standard errors reported in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.5: Testing Stationarity of Costs Across Conditions

	Drive 1		
	Advance and Immediate	Advance	Immediate
	(1)	(2)	(3)
Discount Factor	0.877	0.934	0.820
	(0.046)	(0.040)	(0.080)
a	-23.416	-18.828	-15.385
	(1.906)	(0.322)	(0.670)
$\gamma = 1 + 3 \cdot \frac{1}{1+exp(a)}$	4	4	4
$ln(\sigma)$	-1.138	-1.662	-0.879
	(0.185)	(0.080)	(0.222)
# Observations	338	174	164
Log-Likelihood	-94.806	42.226	-88.490
H_0 : Discount Factor = 1	$\chi^2(1) = 7.17$	$\chi^2(1) = 2.66$	$\chi^2(1) = 5.04$
	$(p < 0.01)$	$(p = 0.10)$	$(p = 0.02)$

Notes: This table reports structural estimates of an exponential discount factor, δ, and γ obtained using Maximum Likelihood Estimation based on Equation (1). Standard errors are reported in parentheses.

Table A.6: Testing for Idiosyncratic Shocks

Panel A: Maximum Daily Time Lapse Between Vaccinations (in Minutes)								
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: This table reports on the relationship between decision timing and the one period discount factor with two proxies for experiencing a shock during the drive (the maximum time lapse between vaccinations and whether a lapse of more than 2 hours occurred). Heteroskedasticity robust White standard errors reported in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.7: Testing for Deterministic Environment

	Drive 1			
	Experience < 9 Years	Experience ≥ 9 Years	Experience < 15 Years	Experience ≥ 15 Years
	(1)	(2)	(3)	(4)
β	0.883	0.890	0.932	0.784
	(0.123)	(0.147)	(0.102)	(0.226)
δ	0.965	0.886	0.924	0.899
	(0.058)	(0.054)	(0.045)	(0.089)
a	-19.367	-24.158	-16.058	-14.380
	(3.371)	(1.963)	(0.657)	(0.737)
$\gamma = 1 + 3 \cdot \frac{1}{1+exp(a)}$	4	4	4	4
$ln(\sigma)$	-1.203	-1.081	-1.259	-0.874
	(0.298)	(0.236)	(0.207)	(0.332)
# Observations	149	180	248	81
Log-Likelihood	-32.129	-60.907	-39.645	-44.123
$H_0 : \beta = 1$	$\chi^2(1) = 0.90$	$\chi^2(1) = 0.56$	$\chi^2(1) = 0.45$	$\chi^2(1) = 0.91$
	$(p = 0.34)$	$(p = 0.45)$	$(p = 0.50)$	$(p = 0.34)$

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.8: 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: This table reports on the relationship between decision timing and the one period discount factor with whether vaccinators completed their vaccination tasks. Heteroskedasticity robust White standard errors reported in parentheses. $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.9: 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)
$\gamma = 1 + 3 \cdot \frac{1}{1+exp(a)}$ (Rural Sample:)	4	4	4
$\gamma = 1 + 3 \cdot \frac{1}{1+exp(a)}$ (Urban Sample:)			3.395
$ln(\sigma)$	-2.170	-2.172	-2.170
	(0.052)	(0.052)	(0.052)
# Vaccinators	203	203	203
Log Likelihood	152.424	152.904	152.521

Notes: This reports structural estimates of β, δ, and γ obtained using Maximum Likelihood Estimation based on Equation (2). Experienced Urban Group is a dummy equal to one for vaccinators whose average travel per vaccination attempt is less than 0.1 kilometers and who have more than 15 years of experience in the health department. Standard errors are reported in parentheses.

Table A.10: Tailoring with Alternative Controls

Dependent variable:	$\lvert\frac{v_{1,i}}{v_{2,i}} - 1\rvert$		
Comparison group:	Untailored with $\tilde{R}_i \in +/-$ 1 s.d of Mean R_i^*	Untailored with $\tilde{R}_i \in +/-$ 1 s.d of Mean R_i^* $\lvert R_i^* - \tilde{R}_i\rvert > 0.10$	Untailored with Matched Distribution of $\tilde{R}_i$
	(1)	(2)	(3)
Tailored (=1)	-0.039*	-0.052**	-0.047**
	(0.021)	(0.025)	(0.022)
Constant	0.200***	0.231***	0.085
	(0.071)	(0.074)	(0.072)
Stratum FEs	Yes	Yes	Yes
Exclude 99th and 1st Percentiles	Yes	Yes	Yes
Drive 2 R_i^* or $\tilde{R}_i$	Yes	Yes	Yes
R-Squared	0.081	0.086	0.071
Mean in Untailored Contract	0.132	0.150	0.143
# Vaccinators	194	181	207
# Untailored Vaccinators	59	46	72

Notes: This table reports the effects of tailoring on the equality of effort provision over time. The measure $|\frac{v_t}{v_{t+1}} - 1|$ (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 ($v_1 = v_2$). Mean (standard deviation) R_i^* in untailored group = 1.036 (0.165). Heteroskedasticity robust White standard errors reported in parentheses. *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$.

Table A.11: Robustness Tests for Tailoring Intertemporal Incentives

Panel A: Dependent variable $\frac{\lvert v_{1,i}-v_{2,i}\rvert}{\sqrt{2}}$						
	(1)	(2)	(3)	(4)	(5)	(6)
Tailored (=1)	-1.758	-4.480**	-4.481**	-2.401	-1.703	-1.868
	(4.588)	(1.954)	(2.068)	(2.302)	(2.176)	(2.229)
Immediate Choice				20.994***	10.365***	10.597***
				(5.856)	(3.335)	(3.449)
Tailored x Immediate				2.127	-6.026	-6.220
				(9.793)	(4.084)	(4.136)
Constant	31.386***	16.399***	16.412**	19.464***	10.992***	14.128**
	(7.334)	(2.296)	(6.857)	(6.433)	(2.777)	(6.671)
Panel B: Dependent variable $\frac{\lvert v_{1,i}-v_{2,i}\rvert}{\sqrt{2}(v_{1,i}+v_{2,i})}$						
	(1)	(2)	(3)	(4)	(5)	(6)
Tailored (=1)	-0.014	-0.018***	-0.016**	-0.009	-0.007	-0.007
	(0.015)	(0.007)	(0.007)	(0.008)	(0.007)	(0.008)
Immediate Choice				0.080***	0.038***	0.037***
				(0.022)	(0.012)	(0.012)
Tailored x Immediate				-0.008	-0.024*	-0.023*
				(0.032)	(0.014)	(0.014)
Constant	0.103***	0.057***	0.033	0.058***	0.037***	0.025
	(0.023)	(0.008)	(0.022)	(0.020)	(0.010)	(0.022)
Panel C: Dependent variable $\lvert v_{1,i} - \frac{300}{1+R}\rvert$						
	(1)	(2)	(3)	(4)	(5)	(6)
Tailored (=1)	-2.612	-3.789***	-3.445**	-1.959	-1.509	-1.405
	(3.282)	(1.410)	(1.459)	(1.643)	(1.575)	(1.591)
Immediate Choice				16.650***	7.990***	7.844***
				(4.496)	(2.451)	(2.509)
Tailored x Immediate				-0.867	-4.972*	-4.850
				(6.891)	(2.938)	(2.974)
Constant	22.736***	12.014***	7.571	13.340***	7.849***	5.871
	(5.301)	(1.682)	(4.735)	(4.501)	(2.048)	(4.622)
Panel D: Dependent variable $\lvert v_{1,i} - \frac{300}{1+R}\rvert > 25$						
	(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)
Panel E: Dependent variable $min[v_{1,i}, v_{2,i}]$						
	(1)	(2)	(3)	(4)	(5)	(6)
Tailored (=1)	7.210**	8.112***	2.540*	4.764*	4.381*	0.843
	(3.512)	(2.051)	(1.416)	(2.650)	(2.643)	(1.567)
Immediate Choice				-20.562***	-11.804***	-6.815***
				(5.042)	(3.492)	(2.332)
Tailored x Immediate				4.672	8.194**	4.037
				(7.206)	(4.052)	(2.806)
Constant	126.824***	136.728***	208.758***	138.343***	142.870***	210.228***
	(5.181)	(2.645)	(4.541)	(4.905)	(3.359)	(4.433)
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
# Vaccinators	280	267	267	280	267	267

Notes: This table reports the effects of tailoring on the equality of effort provision over time using several different measures of 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 an indicator equal to one for subjects in the tailored group. Column (1) reports a regression of this measure on an indicator 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 (5) and (6) apply the same restrictions to the sample as columns (2) and (3) respectively. Heteroskedasticity robust White standard errors reported in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.12: Alternate Treatment Measures and Sample Restrictions

Panel A: Tailoring Intensity						
Dependent variable:	$\lvert\frac{v_{1,i}}{v_{2,i}} - 1\rvert$					
	(1)	(2)	(3)	(4)	(5)	(6)
Tailor Intensity	2.521	0.164**	0.110*	0.078	0.124	0.089
	(2.057)	(0.065)	(0.063)	(0.188)	(0.081)	(0.076)
Immediate Choice				-0.065	0.071***	0.068***
				(0.312)	(0.023)	(0.022)
Tailor Intensity x Immediate				4.537	0.071	0.057
				(3.809)	(0.133)	(0.131)
Constant	0.780*	0.104***	-0.009	0.765	0.067***	-0.018
	(0.457)	(0.018)	(0.058)	(0.482)	(0.020)	(0.058)
# Vaccinators	280	267	267	280	267	267
Panel B: Tailoring Intensity and Boundary Sample						
Dependent variable:	$\lvert\frac{v_{1,i}}{v_{2,i}} - 1\rvert$					
	(1)	(2)	(3)	(4)	(5)	(6)
Tailor Intensity	1.200	0.151**	0.124*	-0.038	0.054	0.025
	(1.006)	(0.069)	(0.065)	(0.127)	(0.060)	(0.054)
Immediate Choice				0.094	0.066***	0.064**
				(0.202)	(0.025)	(0.025)
Tailor Intensity x Immediate				2.075	0.148	0.154
				(1.848)	(0.119)	(0.114)
Constant	0.712*	0.152***	0.044	0.652*	0.119***	0.016
	(0.368)	(0.026)	(0.065)	(0.369)	(0.026)	(0.063)
# 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 an indicator 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 (5) and (6) apply the same restrictions to the sample as columns (2) and (3) respectively. Heteroskedasticity robust White standard errors reported in parentheses. $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

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