Time Inconsistency, Expectations and Technology Adoption: The case of Insecticide Treated Nets

Aprajit Mahajan
Alessandro Tarozzi*

March 11, 2011

Abstract

Economists have recently argued that time inconsistency may play a central role in explaining intertemporal behavior, particularly among poor households. However, time-preference parameters are typically not identified in standard dynamic choice models and little is known about the fraction of inconsistent agents in the population. We formulate a dynamic discrete choice model in an unobservedly heterogeneous population of possibly time-inconsistent agents motivated by specifically collected information combined with a field intervention in rural India. We identify and estimate all time-preference parameters as well as the population fractions of time-consistent and "naïve" and "sophisticated" time-inconsistent agents. We estimate that time-inconsistent agents account for more than half of the population and that "sophisticated" inconsistent agents are considerably more present-biased than their "naïve" counterparts. We also examine whether there are other differences across types (e.g. in risk and cost preferences) and find that these differences are small relative to the differences in time preferences.

JEL: I1,I3

Key words: Malaria, Expectations, Bednets, Identification, Dynamic Programming, Discrete Choice, Time Inconsistency

*We thank seminar participants at SITE, UC Davis, UBC, Simon Fraser, University of Victoria, Oxford University, PACDEV, Duke University, University of Pennsylvania, Northwestern University, the ESWC (Shanghai) meetings, Columbia University, RAND, Harvard-MIT, UC Berkeley and the AEA Winter 2011 meetings for useful comments. We are also grateful to Nathan Hendren, Hanming Fang and Han Hong for valuable comments and suggestions. The usual disclaimer applies.
Aprajit Mahajan, Dept. of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305, amahajan@stanford.edu. Alessandro Tarozzi, Dept of Economics, Duke University, Social Sciences Building, PO Box 90097, Durham, NC 27708, taroz@econ.duke.edu.

1 Introduction

One of the constitutive tenets of standard neoclassical economics is that individuals pursue constrained utility maximization. In models where agents take decisions over time, it is usually assumed that individuals maximize expected future utility flows under an intertemporal budget constraint. Such models have provided invaluable insights in understanding economic decisions such as savings, asset allocation or investment in health and education. On the other hand, a number of studies have proposed alternative models to explain behavior that is hard to reconcile with standard models of individual optimization. Examples of such behavior are addiction, preference reversals in intertemporal choices and under-investment in activities with apparent low cost and high expected returns.[1] In several cases, insights from psychology and behavioral economics have suggested that such behavior may be better explained by models where individuals exhibit self-control or time inconsistency problems.

These theories have played an increasing role in explaining "inefficient" choices among poor individuals in developing countries, a context where such choices may have particularly dire consequences (Mullainathan, 2004). In recent work, non-standard preferences displaying bias towards the present have been proposed to explain poverty traps (Banerjee and Mullainathan, 2010), the existence of demand for commitment devices in savings or health-protecting technologies (Ashraf et al., 2006; Tarozzi et al., 2009, 2011) and low demand for immunization and fertilizer (Banerjee et al., 2010; Duflo et al., 2010).

Present bias is typically modeled assuming that preferences are characterized by "hyperbolic discounting" (Laibson, 1997). In such models, utility is modified so that, at each time t, future utility at time s ($> t$) is discounted not by the usual geometric factor δ^{s-t} but by a factor $\beta\delta^{s-t}$. As a consequence, while δ is the only discount factor entering the intertemporal rate of substitution between any two future periods, the rate of substitution between current time t and any future period also depends on β. This model generates a declining rate of time preference and has been used to explain the "preference reversal" that is commonly observed in laboratory experiments: individuals choose a reward at date t over a larger one at date $t + s$, but instead choose the later reward if the two dates are shifted forward by an equal time period.[2] Such choices are not consistent with standard expected utility models.

A consequence of hyperbolic preferences is that an individual who maximizes intertemporal utility at time t will have an incentive to deviate from this solution at time $t + 1$, when present-bias will induce an increase in consumption relative to what was previously decided. While such models promise to help in explaining the often observed inability of the poor to save or invest even when the budget constraint would allow it, structural estimation of the discount factors that characterize hyperbolic preferences is non-trivial.

In fact, time preference parameters are generically not identified in standard dynamic choice models (Rust, 1994; Magnac and Thesmar, 2002). In this paper, we propose a dynamic discrete choice model with unobserved types and time varying utilities, and provide identification results for all time preference parameters. We overcome the previous non-identification results in the literature by (i) adding more information in the form of elicited beliefs about state occurrences and elicited responses to time preference questions and (ii) designing a product appealing to particular types of agent and offering it for sale in a field intervention.

[1] See Frederick et al. (2002) and DellaVigna (2009) for reviews.

[2] See Andreoni and Sprenger (2010) for an alternative explanation for these findings.

We first show that the model is identified and then estimate it to test several hypotheses of interest. First, we ask whether time-inconsistent preferences provide a better fit for the data than alternative proposed explanations. In particular, we examine alternative explanations that stress differences in per-period utilities[3] as well as information-based explanations. We find that while per-period utilities do vary across agent types,[4] they are not substantively important in explaining outcomes in our sample. Second, we also identify and estimate the distribution of types in the population which provides a quantitative measure of the importance of time-inconsistent agents in the economy. We estimate that approximately 40% of the population from which our sample is drawn is time consistent, while 50% are "naïve" inconsistent and the remaining 10% are "sophisticated" inconsistent. Further, we find that "sophisticated" agents are considerably more present-biased than "naïve" agents. This finding is possible because we show identification for separate hyperbolic parameters for each type. In particular, we find that "naïve" agents have a hyperbolic parameter close to 1 and that in a set of counterfactual simulations, "naïve" agent choices are similar to those made by consistent agents.

Finally, we evaluate to what extent "sophisticated" agents are more likely to choose specially designed (commitment) products by comparing the type distribution among product purchasers to that in the general population. We find that commitment products are not particularly appealing to "sophisticated" agents and that the purchase of these products is in fact higher among wealthier (and even "naive") households. Note that this finding contradicts a deterministic mapping from the take-up of commitment products into agent type. In this latter framework, the choice of a commitment product reveals an agent to be "sophisticated." One of our key identification results allows product choice to only imperfectly predict type which in turn allows for a much richer analysis of preferences.[5]

In drawing links to the extensive literature on time-inconsistency and on structural estimation with unknown types we restrict ourselves here to work that is closest in spirit to our approach.[6] Our work is closely related to Fang and Wang (2010) who outline methods for estimating time preference parameters by imposing exclusion restrictions on the standard model (in our context, elicited beliefs are a natural candidate for such restrictions). Our approach differs from theirs in that we allow for unobserved agent types who differ in both time-preference as well as per-period utility parameters (though we do not allow for partially "naïve" agents). Van der Klaauw (2000) and van der Klaauw and Wolpin (2008) use information about expected future choices to improve precision in the context of a structural dynamic model. Our work instead uses expectations about state transitions and focuses on using this information to achieve identification. Our identification results are also closely related to Kasahara and Shimotsu (2009), who consider an environment with unknown agent types. However, our approach differs in that we achieve identification by imposing an exclusion restriction (by requiring a variable that affects type probabilities

[3] e.g. Banerjee and Mullainathan (2010), propose a model without hyperbolic time preferences as an alternative to standard models of time-inconsistency

[4] See p.10 for a description of the types of agent considered. Briefly, "naïve" inconsistent agents do not take their future present-bias into consideration while formulating their dynamic plans. In contrast, "sophisticated" inconsistent agents incorporate their future present-bias in their planning.

[5] Previous work (e.g Fang and Silverman, 2009; Paserman, 2008) does not address these questions directly since agent type heterogeneity is typically ruled out by assumption. In these models agents usually have identical ("sophisticated" inconsistent) preferences. Fang and Wang (2010) also deal with identical preferences across agents but allow for partially "naïve" agents.

[6] See Aguirregabiria and Mira (2010) or Arcidiacono and Ellickson (2011) for a recent survey of work on dynamic discrete choice structural models.

but not the choice probabilities). In contrast, Kasahara and Shimotsu (2009) place assumptions on the length of the panel available to the researcher. Our approach is also similar to Ashraf et al. (2006) as we use elicited time preferences to predict behavior and we design a product that should appeal to "sophisticated" inconsistent agents. In addition, our estimation can be viewed as a structurally based field version of the laboratory approach of Andersen et al. (2008) since we jointly estimate discount and utility curvature parameters based on multiple price lists as well as elicited risk preferences.[7]

The paper is organized as follows. Section 2 provides an overview of the project design and data collected. Section 3 outlines the basic elements of the dynamic discrete choice model with different types and describes the model primitives in some detail. Section 4 provides the identification results for the model, first for the simple case where observables reveal type completely and then for the more realistic case where type is only imperfectly observed. Section 5 reports the results from a series of Monte-Carlo simulations detailing the finite sample properties of the estimation methods that follow from the identification strategy. Section 6 outlines the estimation on the intervention data and reports the results. Section 7 reports the results of a set of counterfactual exercises and Section 8 concludes.

2 Data

The data used in this paper have been collected as part of the evaluation of a large-scale randomized controlled trial (RCT) carried out in the Indian state of Orissa between spring 2007 and winter 2009. The main purpose of the study was the evaluation of alternative mechanisms of providing insecticide treated nets (ITNs) on health and socio-economic outcomes of potential users (see Tarozzi et al., 2011). The broader project has been conducted in 166 villages in five districts in Orissa: Bargarh, Balangir, Keonjhar, Kandhamal (Phulbani), and Sambalpur. Figure 1 shows the location of the districts within India. Orissa is the most malaria endemic state in the country: despite containing less than 4% of India's population, Orissa accounts for 25% of reported malaria cases, 40% of *P. falciparum malaria* cases, and 30% of malaria-related deaths (Kumar et al., 2007, citing figures from the Indian National Vector Borne Disease Control Programme). Study locations were selected randomly from a list of 878 villages provided by our local partner, Bharat Integrated Social Welfare Agency (BISWA), a micro-lender with a large presence in Orissa and elsewhere in India.[8] This paper uses data collected in a randomly determined subset of 47 villages where, in the fall of 2007, BISWA offered all its clients the opportunity to purchase high quality ITNs on credit, with repayment over one year.

A baseline, pre-intervention survey of 627 randomly sampled households was carried out in March-April 2007. Households were selected from client rosters provided by BISWA. In each village, up to 15 households were selected at random from these lists, while interviews were attempted with all households in villages with fewer than 15 listed clients. Between September and November, 2007, all villages were exposed to a short but intensive community-based information campaign (IC) about the importance and rationale for ITN use, advice on proper use and retreatment, and a clear explanation of the micro-loans offered for ITN purchase. BISWA clients were offered the opportunity to purchase as many ITNs as they wished, choosing

[7]The risk preferences measures are based on Holt and Laury (2002). This "dual elicitation" is referred by Andreoni and Sprenger (2010) as the Double Multiple Price List approach.

[8]For more information about BISWA see `www.biswa.org`.

one of two alternative loan contracts (described in detail below). Net distribution and recording of loan contracts for BISWA members who decided to purchase were completed 2–3 days after the IC. The time interval between the IC and the purchase decision was introduced to ensure that the households had an opportunity to consider the offer carefully. A second visit was scheduled approximately one month later, and nets were offered again with the same contracts. No further ITNs were offered after this second visit was completed. The first net re-treatment was completed approximately six months after the ITN sale, in March–April 2008, while the second and final retreatment took place approximately another six months later, in September–November 2008.

Two alternative contracts were offered to BISWA clients. With the first offer (referred to as C1 henceforth), single (double) nets were sold on credit for Rs. 173 (223), and repaid with twelve monthly installments of Rs. 16 (21).[9] Nets were immediately treated with insecticide, with a chemical concentration that makes re-treatment optimal after approximately six months. Survey personnel would re-visit the villages after six and twelve months and offer retreatment for Rs. 15 (single) or Rs. 18 (double). With the second offer (referred to as C2 henceforth), the household purchased not only the treated net but also a sequence of two re-treatments. The price in this case was Rs. 203 (259), to be paid as twelve monthly installments of Rs. 19 (23). With this second option, no cash payment would be required for re-treatment as the price of the chemicals was already included in the loan amount. In all cases, the interest rate was the standard annual 20 percent charged by BISWA. BISWA microcredit operations are based on group lending: loans are offered to borrowers organized in small self-help group averaging 15 members. Each member is responsible for the repayment of all loans granted to the group, which diffuses responsibility to all group members. There is no collateral for the loans, but (as is standard in micro-finance) defaulting borrowers were informed that they would be denied further loans from BISWA. Default is only determined at the end of the loan period, so clients are allowed some flexibility in the repayment schedule. For instance, a borrower may miss a few monthly repayments during the "lean" agricultural season while paying current and past dues after the harvest, and early repayments are allowed.

Table 1 show summary statistics for the households sampled from the 47 villages where ITNs were offered on credit. Mean monthly total expenditure per head is approximately twice as large as the official poverty line for rural Orissa in 2004–5.[10] Existing net ownership was not uncommon, with a mean of one bednet for every three persons, although one third of households did not own any nets. Treated net ownership was minimal, with only 0.06 ITNs per head on average. 16% of individuals slept under a net the night before the survey, and the corresponding figure was 3% for ITNs. Reports about bednet use the night before the interview are unlikely to suffer from significant recall bias. On the other hand, the baseline survey was completed during the hot and dry season, when mosquitoes are less of a nuisance and malaria rates are relatively low. For this reason, we also asked about bednets use in periods of high mosquito activity. During such periods, more than half of the members were reported as sleeping "regularly" under a bednet. Note, however, that the vast majority of nets in the area were not treated with insecticide, so that even during the mosquito season the protective power of the available nets remained suboptimal.

[9] For perspective, daily wages for agricultural labor in the area were around Rs. 50 and the price of one kilogram of rice was approximately Rs. 10. The official rural poverty line for Orissa in 2004–5 was Rs. 326 per person, per month (Government of India, 2007).

[10] Household expenditure was measured asking about usual consumption of eighteen broad item categories, including self-production.

Fingerprick blood testing was also completed for a sample of household members and the results show high prevalence of malaria (11 %) and anemia (46 %), with the latter defined here as occurring when hemoglobin levels were below $11g/dl$.[11] Given that the baseline survey was completed during the dry season, malaria prevalence was expected to be even higher during the rainy and post-monsoon season, a result confirmed in the winter 2008–9 follow-up survey. The next rows of Table 1 show that respondents were aware of the role of mosquitoes in transmitting malaria as well as of the high economic cost of malaria episodes.

Awareness about the protective power of bednets is also reflected in the beliefs elicited directly from respondents. This was done by asking respondents to hold up a number of fingers increasing in the perceived likelihood that an event will happen, with no fingers representing "no chance" and ten fingers indicating certainty. We then estimated subjective probabilities by dividing the number of fingers held by ten.[12] Given that most respondents are illiterate and unfamiliar with the concept of probability, the interviewer discussed first hypothetical examples of certain and uncertain events to explain the rationale. At baseline, the survey instrument included questions about the probability for an adult, a child under the age of six (U6) or a pregnant woman (PW) of contracting malaria in the next year depending on whether the individual slept regularly under an untreated net, an ITN or no net.[13]

The graphs in Figure 2 show the distribution of elicited beliefs of falling sick with malaria within one year for individuals who sleep regularly under a treated bednet (top row), an untreated net (middle) or who do not use a net (bottom). Beliefs are very similar for individuals who belong to different demographic groups. Importantly, both bednets and re-treatment with insecticide appear to be widely recognized as very effective at reducing malaria risk. The histograms also show that a majority of elicited beliefs are concentrated over the focal figures 0, 5 and 10. About three quarters of respondents believe that if nets are not used one will certainly get malaria, and approximately the same fraction believes that regular use of treated nets will virtually wipe out all risk. According to about half of respondents, there is instead a 50% chance of developing malaria if an untreated net is used. On the other hand, there remains a degree of variation in the beliefs which can be exploited in the structural model outlined in Section 3.

The baseline survey instrument also included twelve questions intended to gauge respondents' intertemporal preferences and the extent of time inconsistency in these preferences. In a first group of four questions, the respondent choose between an *actual* Rs. 10 sum to be paid one month later and an equal or larger sum (Rs. 10, 12, 14 or 15) to be paid four months later. In a second group of questions the choice was between Rs. 10 one month later and Rs. 10, 15, 20 or 25 seven months later. Finally, in a third set of

[11] Malaria infection and hemoglobin levels were measured via fingerprick blood specimens requiring less than $0.5ml$ of blood. Consent was requested to measure malaria infection and hemoglobin levels for all pregnant women, all children aged 5 and under (U5) and their mothers and one randomly selected adult (age 15-60). Malaria infection was determined using the Binax Now malaria rapid diagnostic test (RDT). This RDT is well validated internationally in comparison to blood smears (Moody, 2002) and provides accurate diagnosis for current or very recent (2 to 4-week) malaria infection. The test is particularly accurate for the most severe form of malaria, caused by *P. falciparum*, and it can distinguish the infecting *Plasmodium* species to some extent. Hemoglobin levels were tested with the HemoCue 201 hemoglobin analyzer (a portable, accurate system for measuring hemoglobin). Results are available within about 15 minutes for these tests and communicated directly to the participants.

[12] Note that we do not attempt to measure ranges of probability, so that our data do not allow to identify the degree of uncertainty around the reports.

[13] For instance, one question asked: "imagine first that your household [or a household like yours] does not own or use a bed net. In your opinion, and on a scale of 0-10, how likely do you think it is that a child under 6 that does not sleep under a bed net will contract malaria in the next 1 year?" Questions for different demographic groups and bednet use were asked using analogous wording.

questions the same rewards described for the first group were offered, but with time horizons shifted by three months.[14] Standard expected utility models imply that if a respondent prefers, say, Rs. 15 four months later to Rs. 10 paid a month from today, s/he should also prefer Rs. 15 paid seven months in the future to Rs. 10 paid three months in the future. We interpret preference "reversals", whereby the former is true but the choice is reverted for the later rewards, to signal a form of inconsistency in time preference consistent with hyperbolic discounting.[15] In Table 2 we summarize the findings. As expected, in each group of four questions, the fraction of individuals who prefers the earlier and lower reward decreases when the time horizon of the later reward remains the same but the reward increases. Approximately one fourth of respondents exhibit at least one "hyperbolic preference reversal."

Table 3 includes a summary of the results of the ITN sale, completed in September–November, 2007. Slightly more than 50 percent of sample households purchased at least one net on credit (330 of 621). Of these, 153 chose to purchase only ITNs, while 165 opted for the "commitment" product whose price also included the cost of two re-treatments. Only twelve buyers purchased contracts of both types and we omit them from the analysis. Among the buyers, the mean number of ITNs purchases was close to two, regardless of the contract type chosen. Panel (B) of Table 3 shows that fewer than 3% chose to purchase ITNs for cash rather than on loan and we will not focus on these purchases.

3 Model

The agent chooses three actions over the time period of the project. In the first period, the agent chooses whether to purchase a net, and if so the type of contract. In the second period, the malaria status of the household is realized following which the agent chooses whether to re-treat the net. Then, in period three malaria status is realized and then the agent chooses whether to re-treat again. Period 4 is the terminal period and the agent takes no action in that period. We begin by defining and placing standard assumptions on the state space, the action space, the transition probabilities, the class of acceptable decision rules and finally the preferences and objective function the agent maximizes.

3.1 Primitives

State Space: $\mathcal{S}_t$

The state space $\mathcal{S}_t$ can be partitioned as $\mathcal{S}_t \equiv (\mathcal{X}_t, \mathcal{E}_t)$ where the first element denotes the domain of the observed (to the researcher) state variables and the second element denotes the domain of the unobserved state variables. In the empirical work we allow for a rich observable state space (including income, prices and other characteristics), but to focus attention on the key parts of the identification argument we simplify

[14] Interviewers told respondents that one of the twelve chosen rewards, selected at random, would be paid by our micro-lender partner BISWA at the chosen time horizon. In practice, to avoid logistical difficulties, we decided to pay immediately the selected reward at the end of the interview (we find no evidence that the responses varied for households interviewed later during the day). Note also that all options entailed rewards to be paid at least one month later. This was done so that choices would not depend on issues of trust, issues which were also likely to be made less relevant by the fact that all respondents belonged to households with at least one BISWA client.

[15] However, see Rubinstein (2003) or Andreoni and Sprenger (2010) for an alternative view.

the state space in the exposition to the bare minimum.[16]

In period 1, $x_1 \in \mathcal{X}_1$ is a binary variable equal to one if the respondent reported at least one case of malaria in the household in the past six months. In periods 2 and 3, $x_t \in \mathcal{X}_t$ takes on six possible values. We denote the possible values by (nm, nh, bm, bh, cm, ch) where the first lower case letter in each state value records whether the agent did not purchase a net (n), purchased a net using the first type of contract C1 (b) or purchased a net using the commitment contract C2 (c). The second letter captures whether anyone in the agent's household suffered from malaria in the last six months, with m denoting someone had malaria and h ("healthy") denoting that no-one contracted malaria.[17] In the sequel, it will be useful to define the sets $\mathcal{X}_C \equiv \{ch, cm\}$ and $\mathcal{X}_B \equiv \{bh, bm\}$ that partition $\mathcal{X}_t$ depending upon the first period purchase decision.

As commonly assumed in dynamic discrete choice models, the vector of unobservables $\epsilon_t \in \mathcal{E}_t$ has dimension equal to the number of actions available to the agent in period t. Such variables will represent a stochastic, choice-specific component in the utility function that is known to the agent but unobserved by the econometrician. In addition, we assume that the support of each element of ϵ_t is the entire real line although in the notation below we are somewhat informal and imply positive point probabilities for ϵ_t.

We can allow for a much richer discrete state space where, for instance, we keep track of the fraction of household members that are healthy in each period. The identification arguments remain the same and are presented in detail in Appendix B but for ease of exposition we present a simplified version of the model in the text.

Action Space: $\mathcal{A}_t$

The action space in period one ($\mathcal{A}_1$) has three elements denoted by (n, b, c), which are defined as above. In periods 2 and 3 ($t \in \{2, 3\}$), the action space is $\mathcal{A}_t = \{0, 1\}$, where 0 denotes that the agent did not re-treat a net and 1 denotes that an agent did re-treat the net. Note that if an agent did not purchase a net in period 1, she cannot take any more actions. Finally, we do not observe the state of the world in the terminal period and the agent takes no action in this period.

In Appendix B, we generalize the model to allow for more than two or three actions in each period. This would be useful for instance if we viewed the decision as being about the number (or fraction) of household members that can be covered with an ITN. The general arguments in this case are very similar and so we discuss the simpler action space in the main text.

Transition Probabilities: $\mathbb{P}(s_t|s_{t-1})$

Let $\mathbb{P}(s_t|s_{t-1})$ denote the distribution function of the random vector s_t conditional on s_{t-1} and refer to it as the transition probability distribution. We make the standard assumption that the transition probabilities are Markov (see e.g. Aguirregabiria and Mira, 2010) in the sense that

[16] Appendix B contains details on the identification argument with a general discrete state space and Appendix D contains details on the construction of the state space for the empirical model.

[17] The state space can be easily extended so that agents keep track of their entire history of malaria. The current specification is a convenient short-cut and is also undertaken for tractability since in the sequel we will consider a first-step non-parametric estimator at each point of the state space.

ASSUMPTION 1.

$$\mathbb{P}(s_t|s_{t-1},...,s_1,a_{t-1},...,a_1) = \mathbb{P}(s_t|s_{t-1},a_{t-1}) \tag{1}$$

where $\mathbb{P}(s_t|s_{t-1},a_{t-1})$ is the conditional distribution of the random variable s_t given s_{t-1} and a_{t-1}. This assumption, along with the simplest definition of the state space described above, rules out for instance the possibility that the probability of malaria infection in period 3 depends on malaria status in the first period given malaria status in the second period and the retreatment decision in period 2. Incorporating such dependencies is straightforward by suitably redefining the state variable at t to contain the complete malaria history up to t.

In addition, we assume (as is standard) that the vector of unobservables ϵ_t is independently distributed across time. This rules out serially correlated unobserved heterogeneity, such as if agents' decisions were driven by shocks (to income for instance) whose effects last for multiple periods. We mitigate this limitation in two ways. First, we allow for considerable heterogeneity across time and across agents by permitting time- and type-varying preferences (see below for details). Second, we include a number of observed variables, such as income, in the state space. The hope is that these two approaches (expanding the state space and allowing for unobserved types with flexible time-varying utility) will minimize the extent of serially correlated unobserved heterogeneity.

We also assume that ϵ_t has a known distribution and is independent of the whole path of observable state variables $\{x_t\}_{t=1}^{T}$ and also independent of $\{a_s\}_{s=1}^{t-1}$. This rules out for instance direct feedback from current shocks to future state variables: for instance, if a positive shock today leads to the household not only purchasing nets but also investing in other health improving technologies that reduce the likelihood of malaria incidence in the future. We deal with this limitation by including a comprehensive set of state variables directly modelling their evolution over time (see Appendix D for more details).

ASSUMPTION 2.

$$\mathbb{P}(x_t,\epsilon_t|x_{t-1},\epsilon_{t-1},a_{t-1}) = \mathbb{P}(x_t|x_{t-1},a_{t-1})\mathbb{P}(\epsilon_t), \tag{2}$$

where the distribution of the vector ϵ_t is absolutely continuous on the real line (w.r.t. Lebesgue measure), known and has mean zero. The dimension of ϵ_t is equal to the number of elements in $\mathcal{A}_t$.

Beliefs

Our data contains detailed information on subjective beliefs about the risk of contracting malaria elicited from respondents. Specifically, for each household we elicit the perceived probability of malaria infection within a year for a person never sleeping under a net, sleeping regularly under an untreated net, and sleeping regularly a treated net. Let these perceived probabilities be denoted by π, $\pi - \phi$ and $\pi - \phi - \gamma$ respectively so that ϕ and $\phi + \gamma$ measure the perceived reduction in malaria risk offered by nets and ITNs respectively. Define the vector of beliefs $z \equiv (\pi, \phi, \gamma)$.

In addition, we also collected post-intervention beliefs in the follow up survey. These allow us to directly evaluate the extent of belief evolution and to incorporate it into the model. This is potentially important since it allows us to explicitly incorporate issues about learning into the dynamic choice process. The empirical section describes our approach to modelling the evolution of beliefs in greater detail. However, to simplify exposition and focus attention on the key elements for identification, we assume here that agents' beliefs stay constant over time. This is only a convenient simplification and the identification results remain unaffected if we abandon it in favor of time-varying observed beliefs. In this special case, we can construct individual level transition probabilities $\mathbb{P}(x_t|x_{t-1}, a_{t-1})$ for each period using just baseline beliefs. In the sequel, we denote the transition probabilities as $\mathbb{P}(x_t|x_{t-1}, a_{t-1}, z)$. For instance, the probability of malaria infection in $t = 3$ for an individual who had malaria in the previous period and who is using regularly a treated net purchased with contract b, is $\mathbb{P}(x_3 = bm|x_2 = bm, a_2 = 1; z) = \pi - \phi - \gamma$. However, the transition probabilities could be represented by more complex functions of the beliefs z, depending on the richness of the state space. For instance this would be the case if x_3 denoted the fraction of members with malaria in a household.

Decision Rules: d_t

The decision rule in period t, d_t, is a mapping from $\mathcal{S}_t$ to $\mathcal{A}_t$. Note that we do not allow history dependent decision rules in the sense that we do not allow decision rules to be mappings from $\prod_{s=1}^{t-1}(\mathcal{S}_s, \mathcal{A}_s) \times \mathcal{S}_t$ to the action space. However, given the Markov property for the transition probabilities and the assumptions on preferences below, the optimal decision rule will indeed be a deterministic function only of the current state.

Types and Preferences

As is standard, we assume that preferences are additively time-separable, and parameterize time inconsistency using the tractable (β, δ) formulation described in Strotz (1955).[18] Then, for a given sequence of actions $\{a_t\}_{t=1}^{3}$, the utility of an agent of type τ is[19] :

$$\tilde{u}_t(s_t, a_t, \tau) + \beta_\tau \sum_{j=t+1}^{T} \delta^{j-t} \mathbb{E}_t(\tilde{u}_j(s_j, a_j, \tau)) \tag{3}$$

We allow for 3 different types of agents. Following O'Donoghue and Rabin (1999), time consistent agents $(\tau = \tau_C)$ have $\beta_{\tau_C} = 1$, which corresponds to the standard case of exponential discounting. Such agents will maximize (3) using standard dynamic programming methods (backward induction in this finite horizon case). We also allow for the existence of two types of time-inconsistent agents, for whom the hyperbolic parameter β is below one. Both types of time-inconsistent agents are aware of their current present-bias, but they differ in whether they recognize the present-bias their future selves will be subject to. Time-inconsistent "sophisticated" agents $(\tau = \tau_S)$ will use backward induction to solve the maximization problem, but do so while recognizing their future present-bias. On the other hand, time-inconsistent "naïve" agents $(\tau = \tau_N)$ ignore the present-bias of their future selves in their backward induction. For the econometrician, these differences will generate key identification issues that we address below.

The formulation in (3) allows for time-varying type-specific per-period utilities $\tilde{u}_t(\cdot, \tau)$. This flexibility is important since it allows us to examine how much of the difference in behavior across types can be attributed to different preferences over states and how much to the differing extent of present-bias (and its recognition). In addition, the time- and type-varying formulation provides a mechanism for dealing with serially correlated unobserved heterogeneity.

[18] This is certainly not the only possible formulation: see for instance Gul and Pesendorfer (2001, 2004). However, empirical implementations of time-inconsistency have worked almost exclusively using the (β, δ) formulation.

[19] With the understanding that in the terminal period $\mathcal{A}_T$ is the empty set.

Note that this formulation nests the typical model where types only differ in the degree of present-bias so that we can provide a quantitative judgement (using a standard goodness of fit test) of the relative strength of the present-bias explanation for observed behavior. In particular, we can compare the hyperbolic explanations for observed behavior against alternative explanations in a coherent fashion. In sum, types differ in their degree of present-bias, preference structure as well as in their solution to the utility maximization problem and we can test to what extent these differences help explain observed behavior.

Finally, we assume that within each period, utility is additively separable in the unobserved state variables. Formally, and recalling that no action is taken in period 4:

ASSUMPTION 3. *For each type $\tau \in \mathcal{T}$ where $\mathcal{T} = \{\tau_C, \tau_S, \tau_N\}$ the utility function in period $t \in \{1, 2, 3\}$ is given by*

$$\tilde{u}_t(s_t, a_t, \tau) = u_t(x_t, a_t, \tau) + \epsilon_t(a_t)$$

and utility in the terminal period is given by $u_4(x_4, \tau)$.

In general, time-varying observed characteristics will be included as state variables. The transition probabilities for such variables will be generated using elicited beliefs for variables for which we collected such information (malaria incidence, income) and by invoking a rational expectation assumption for variables for which we have no belief information. In addition to the state variables, per-period utility can also be a function of time-invariant characteristics (e.g. education level of the household head) that we collectively refer to as v. Since these play no role in the identification, we omit them as arguments in preferences. The empirical section contains a discussion of which variables are included in v for the estimation.

3.2 Observed versus Unobserved Types

We consider both the case where types are directly observed as well as the case where types are not observed. While the second model is more general, the identification arguments for it require showing identification for the directly identified types case, so it is useful to discuss both cases.

In both cases, we use two important pieces of information from the intervention: first, we collect information about whether individuals exhibit preference reversals in a series of questions designed to gauge the extent of consistency in time preferences. In previous work we have shown that these reversals are important predictors (in a reduced form sense) of subsequent decisions about bednet retreatment (Tarozzi et al., 2009). Agents who exhibited at least one preference reversal are referenced by the binary variable $r = 1$ and agents who exhibit no preference reversals have $r = 0$. Second, we designed a contract that should appeal to "sophisticated" time-inconsistent agents, and agents who purchased these contracts provide us some additional information on their type.

DIRECTLY IDENTIFIED TYPES

In the directly observed types case, we use both these pieces of information to directly identify three types of agent. Agents with $r = 0$ are classified as time consistent and agents with $r = 1$ are classified as time inconsistent. Further, agents with $r = 1$ and who purchase the "commitment" product ($a_1 = c$) are classified as time-inconsistent "sophisticated" types and agents with $r = 1$ and $a_1 = b$ are classified as time-inconsistent "naïve" agents. Time-inconsistent agents who do not purchase a net ($a_1 = n$) can be either "naïve" or "sophisticated", but we cannot directly assign these labels to them. We discuss identification of their type in greater detail in Section 4.3 below. Note that since the state variable x_t for $t > 1$ includes the choice of product in period 1, "sophisticated" and "naïve" types will have mutually exclusive state spaces which we denote by $\mathcal{X}_C \equiv \{ch, cm\}$ and $\mathcal{X}_B = \{bh, bm\}$ respectively.

UNOBSERVED TYPES

In this case, the researcher does not directly observe the type of any individual. We assume instead that the variables (r, a_1) are only imperfect proxies, as is likely the case. For instance, an agent may choose $r = 1$ due to an imperfect understanding of the choices offered rather than genuine time inconsistency. Alternatively, an agent who expects sufficiently high income at the time of re-treatment may not choose the commitment product regardless of time-inconsistency. In principle, the same decision not to commit could also depend on low perceived benefits of re-treatment. However, we will show that this is not a concern for identification to the extent that such perceptions are reflected in agents' elicited beliefs.

We next discuss identification of the model in case where types are directly observed, and then for the general case where types are unobservable.

4 Identification

4.1 Identification: Directly Observed Types

We observe an i.i.d. sample on $(\{a_t^*, x_t\}_{t=1}^{T-1}, w)$ where $\{a_t^*, x_t\}_{t=1}^{T-1}$[20] are the optimal actions and realized states of the world and $w = (z, r, v)$ denotes a vector of agent specific variables measured at baseline. These comprise the agent beliefs $z \equiv (\pi, \delta, \gamma)$ described in Section 3, the binary variable r equal to one in the presence of preference reversals, and other household characteristics v.

The key starting point for identification are the type-specific choice probabilities $\mathbb{P}_\tau(a_t^* = a|x_t, z)$. For $t > 1$, for agents who purchase a product, these type-specific probabilities are directly identified from the observed choice probabilities $\mathbb{P}(a_t^* = a|x_t, z, r)$ since the type is a deterministic function of the observed variables r and the contract choice is an element of x_t. Formally,

ASSUMPTION 4. *Choice probabilities for types that purchase a product are directly observed. In particular, for a time consistent agent*

$$\mathbb{P}_{\tau_C}(a_t^* = a|x_t, z) = \mathbb{P}(a_t^* = a|x_t, z, r = 0).$$

For a "naïve" time-inconsistent agent

$$\mathbb{P}_{\tau_N}(a_t^* = a|x_t, z) = \mathbb{P}(a_t^* = a|x_t, z, r = 1) \text{ for } t > 1 \text{ and } x_t \in \mathcal{X}_B$$

[20] for the exposition here we set $T = 4$

Finally for a "sophisticated" time-inconsistent agent (for $t > 1$)

$$\mathbb{P}_{\tau_S}(a_t^* = a|x_t, z) = \mathbb{P}(a_t^* = a|x_t, z, r = 1) \text{ for } t > 1 \text{ and } x_t \in \mathcal{X}_C$$

4.2 Backward Induction

We discuss identification using a backward induction argument. Since no action is taken in period 4, we begin by examining choice in period 3. In period 3, if the agent has previously purchased a net, s/he will choose to re-treat it if the expected gains from doing so outweigh the costs, that is, if

$$\tilde{u}_3(s_3, 1, \tau) + \beta_\tau \delta \int u(s_4, \tau)\, \mathrm{dF}(s_4|s_3, 1, z) > \tilde{u}(s_3, 0, \tau) + \beta_\tau \delta \int u(s_4, \tau)\, \mathrm{dF}(s_4|s_3, 0, z)$$

where $\mathrm{dF}(s_4|s_3, a_3, z)$ denotes the distribution function for the state variable in period 4 (s_4) conditional on the state of the world s_3, the action a_3 and beliefs $z \equiv (\pi, \delta, \gamma)$. Recall that at time $t = 3$ the action a_3 is binary and equal to one if the agent decides to re-treat the net. Under Assumptions 2 (independence), 3 (additive separability) and 4 (directly observed types) we can write

$$\mathbb{P}_\tau(a_3^* = 1|x_3, z) = G_\Delta\Big(u_3(x_3, 1, \tau) - u_3(x_3, 0, \tau) + \beta_\tau \delta \int u(x_4, \tau)(\mathrm{dF}(x_4|x_3, 1, z) - \mathrm{dF}(x_4|x_3, 0, z))\Big)$$

G_Δ is the distribution of $\epsilon_3(0) - \epsilon_3(1)$ which by assumption is known and has support over the real line. The expression $dF(x_4|x_3, 1, z) - \mathrm{dF}(x_4|x_3, 0, z)$ denotes the difference in the state probability functions as a result of retreating the net (formally, it is a finite signed measure obtained by differencing the two action specific measures). Note that we have assumed that the data is generated in period 3 by agents exhibiting present-bias in the sense that the discount rate applied to the next period is $\beta_\tau \delta$. However, we allow for time-consistent preferences as well, because in principle β_τ could be equal to one for some sub-population of agents. Next, we can invert this relation above to obtain the identified function (see Appendix C)

$$g_{\tau,3}(x_3, z) = u_3(x_3, 1, \tau) - u_3(x_3, 0, \tau) + \beta_\tau \delta \int u(x_4, \tau)(\mathrm{dF}(x_4|x_3, 1, z) - \mathrm{dF}(x_4|x_3, 0, z)). \quad (4)$$

We then explore which of the unknown elements on the right hand side – the utility functions and the discount rates – can be identified using knowledge of the function $g_{\tau,3}(\cdot)$. We use the specifics of the study design to place more structure on the transition probabilities. This has the advantage that for the application at hand we can obtain identification by making fewer (and more reasonable) normalizations. However, all the results stated below can be derived without the additional structure on the transition probabilities, and indeed all that is required is sufficient variability in the beliefs. In Appendix B we derive the identification results for the discrete state space, finite horizon case with a general structure on the transition probabilities. We omit that discussion here to focus on the empirical issue at hand.

ASSUMPTION 5. *Let $\mathcal{X}_B \equiv \{bm, bh\}$ and $\mathcal{X}_C \equiv \{cm, ch\}$ be subsets of the state space $\mathcal{X}_t$ for $t > 1$. Let $a_1 \in \{n, b, c\}$ denote the purchase decision in period 1. Then the transition probabilities from states t*

to $t+1$ can be written as:

$$\mathbb{P}(x_{t+1} = nm | x_t, a_1 = n, w_t) = \pi \qquad \textit{for } x_t \in \{nm, nh\}$$
$$\mathbb{P}(x_{t+1} = x | x_t, a_t, a_1 \in \{b, c\}, w_t) = \pi - \delta - \gamma a_t \qquad \textit{for } (x = bm, x_t \in \mathcal{X}_B), (x = cm, x_t \in \mathcal{X}_C).$$

This assumption rules out changes in beliefs over time but, as we pointed out earlier, this is not important for identification. We assume time-invariant beliefs as a simple starting point because first, we do not directly observe beliefs at each point in time (only at two points) and second, observed time-varying beliefs would in general strengthen identification since the assumptions on variation in beliefs (explicated below) would probably be more credible.

However, we do observe beliefs at the end of the project (i.e. in period 3) and we use these two sets of beliefs to estimate a (parametric) model for their evolution, which is therefore taken into account in the empirical work. In particular, for our empirical example, changes in beliefs are important to incorporate into the analysis as agents may be learning about the efficacy of ITNs once they acquire them and correspondingly updating their beliefs. Finally, note that no re-treatment is possible if $a_1 = n$ (i.e. the no-purchase decision leads to an absorbing state) so that the model assumes that net ownership cannot change after $t = 1$ (which is consistent with the data).

Using Assumption 5, the forward looking component of equation (4) (the integral) is a linear function of the belief variable γ. This allows us to use particularly simple variation in beliefs to identify some of the key unknown objects in (4). To ease notation, in what follows, given two sets A, B, $B \subset A$, let $A \backslash B$ denote the elements of A that do not belong to B. Then:

ASSUMPTION 6. *The distribution of γ conditional on $(x_3, w \backslash \gamma)$ has at least two points of support.*

This assumption requires that beliefs have sufficient residual variation even after conditioning on all state variables. This assumption would fail if, for instance, beliefs were perfectly predictable using state variables. In general, it seems most plausible when beliefs contain private information that affects (in this case) agents' susceptibility to illness. In the context of the ITN study, elicited beliefs showed considerable variation even after controlling for observables.
We can then state the first identification result directly.

LEMMA 1. *Consider an agent of type τ solving at $t = 1$ the problem (3), and that Assumptions 1–6 hold. The researcher observes an i.i.d. sample on $(\{a_t^*, x_t\}_{t=1}^{T-1}, w)$. Then,*

1. *The utility differentials $u_3(x_3, 1, \tau) - u_3(x_3, 0, \tau)$ are identified for $x_t \in \mathcal{X}_B \cup \mathcal{X}_C$*

2. *The object $\beta_\tau \delta \int u_4(x_4, \tau)(\mathrm{dF}(x_4 | x_3, a_3 = 1, z) - \mathrm{dF}(x_4 | x_3, a_3 = 0, z))$ is identified.*

This result describes sufficient conditions for the identification of the change in utility at time 3 associated with retreatment. The proof is in Appendix A and the key insight is equation (4). The underlying differences in latent utilities in this equation has a forward looking component that depends upon beliefs (γ). This component is additively separable from the penultimate period utility differential which can then be identified by using the variation in beliefs.[21] In addition, we assume (as is the case in our data) that at least some of the time consistent agents purchase product b and further that the utility differential from avoiding malaria in period 4 (conditional on exogenous household characteristics) is the same for both "naïve" and time consistent agents. The last assumption is restrictive but we do allow for considerably more variation in utilities by type than considered in the previous literature which usually assumes that agents have the same per-period preferences for all time periods.

ASSUMPTION 7. *Some time consistent agents choose to purchase product b and, in addition, $u_4(bh, \tau_C) - u_4(bm, \tau_C) = u_4(bh, \tau_N) - u_4(bm, \tau_N)$.*

This assumption, along with the previous results, allows us to identify the hyperbolic parameter for "naïve" inconsistent agents. Note that one could carry out a similar exercise if some time consistent agents also purchased the commitment product. However, we do not follow this approach since in subsequent sections we will separately examine these agents as potentially distinct types. Although the equality of utility differentials between the two types at time $t = 4$ cannot be tested, we show that separate differentials can be estimated for $t = 1, 2, 3$, so that the restrictiveness of the assumption can be informally gauged looking at these other identified differences.

LEMMA 2. *Consider an agent of type τ_N solving at $t = 1$ the problem (3) and that Assumptions 1—7 hold. Then, the parameter β_{τ_N} is identified.*

We next consider identification of utility differentials in $t = 2$. The arguments here are very similar to those in the previous step. In addition to the utility differentials, we also identify the exponential discount rate (δ) as well as the remaining hyperbolic parameters ($\beta_{\tau_C}, \beta_{\tau_S}$). First, using the same inversion argument as before we can identify the type-specific function $g_{\tau,2}(x_2, z)$ where

$$g_{\tau,2}(x_2, z) = u_2(x_2, 1, \tau) - u_2(x_2, 0, \tau) + \beta_\tau \delta \int v^*_{\tau,3}(s_3, z)(\mathrm{dF}(s_3|s_2, 1, z) - \mathrm{dF}(s_3|s_2, 0, z)) \tag{5}$$

and $v^*_\tau(s_3, z)$ is the type-specific value function (defined in Appendix A) incorporating the forward looking aspect of the problem. The key to the identification argument, as before, is noting that beliefs (γ) enter only the last part of the expression above (which in turn can be simplified considerably using the results from Lemma 1 and Lemma 2). As usual, we will need to have sufficient variation in beliefs conditional on other state variables and other agent characteristics. These assumptions can be verified directly so that their appropriateness is easily judged. Formally,

ASSUMPTION 8. *The distribution of γ conditional on $(x_2, w \backslash \gamma)$ has at least two points of support (denoted γ_1 and γ_2). Also,*

1. *For the identified quantity $\int v^*_{\tau_C,3}(s_3, z)\, \mathrm{dF}_\Delta(s_3|x_2, z)$.*

$$\int v^*_{\tau_C,3}(s_3, z)\, \mathrm{dF}_\Delta(s_3|x_2, z\backslash\gamma, \gamma_2) \neq \int v^*_{\tau_C,3}(s_3, z)\, \mathrm{dF}_\Delta(s_3|x_2, z\backslash\gamma, \gamma_1)$$

[21] See Fang and Wang (2010) for identification in dynamic models with exclusion restrictions in the transition probabilities. The explicit elicitation of beliefs here provides a natural candidate for exclusion restrictions, because beliefs do not enter the per-period utility function.

2. For the identified function $\mathbf{H}(x_2, z)$[22]

$$\mathbf{H}(x_2, z\backslash\gamma, \gamma_1) - \mathbf{H}(x_2, z\backslash\gamma, \gamma_2) \neq 0$$

Next, we need to make a standard normalization and assume that utility in all states of the world is known for a reference action (see e.g. Magnac and Thesmar, 2002). In particular, we assume that utility levels when the net is not retreated are known for all types.

ASSUMPTION 9. *We normalize utility levels by assuming that* $u_3(x_3, 0, \tau)$ *is known* $\forall x_3 \in \mathcal{X}_3, \tau \in \mathcal{T}$.

Alternative normalizations are possible and discussed in subsequent sections. We can then state the identification result for utility differentials in period 2.

LEMMA 3. *Consider an agent of type* τ *solving at* $t = 2$ *the problem (3) and that Assumptions 1–9 hold. Then, the utility differentials* $u_2(x_2, 1, \tau) - u_2(x_2, 0, \tau)$ *are identified for all* $x_2 \in \mathcal{X}_2$. *In addition* δ *and the hyperbolic parameter* β_{τ_S} *are identified.*

Therefore, under the assumptions stated above, we have now identified all the time-preference parameters. To recapitulate, there are three key features in the study design that enable us to identify the exponential parameter δ and the type-specific hyperbolic parameters $\{\beta_\tau : \tau \in \{\tau_S, \tau_N\}\}$ separately. First, we use baseline responses to a time discounting question to classify agents as time-inconsistent. Second, the experimental intervention offered a product designed to be attractive to "sophisticated" time-inconsistent agents. Third, we elicit beliefs from agents about the probability of malaria incidence conditional on action choice. These three key features allow us to separately identify type-specific hyperbolic parameters β_τ as well as the standard exponential discounting parameter δ.

4.3 Identification of Period 1 Utilities

There is a sharp distinction in period 1 relative to the later periods regarding direct type identification for individual agents. In particular, we cannot directly sub-classify time-inconsistent agents who do not purchase a product (i.e. agents with $r = 1$ and $a_1 = n$) into "naïve" or "sophisticated" types. These agents could either be "naïve" or "sophisticated" but their decision to not purchase a product is not informative of their type. To compound the problem, these agents make no further decisions.

We approach this problem by first noting that the key object required for the inversion argument is the type-specific choice probability $P_\tau(a_1|x_1, z)$. For $t > 1$ we identified $P_\tau(a_t|x_t, z)$ for agents who purchased a product (i.e for $x_t \in \mathcal{X}_B \cup \mathcal{X}_C$) since the agent's choice of product revealed his type perfectly (as ensured by Assumption 4). However, at $t = 1$ we cannot proceed as earlier since the non-purchase decision ($a_1 = n$) does not reveal an agent's type. Instead, we pursue an approach due to Kasahara and Shimotsu (2009) that allows us to identify these type-specific choice probabilities even though we do not observe each agent's type. This approach requires us to impose a set of exclusion restrictions and the argument previews the arguments in the next section on unobserved types.

As a starting point, consider the relationship between joint distribution of actions and states in the first two periods and the type-specific choice probabilities. We use the distribution for the initial two periods

[22] The value functions $v^*_{\tau,3}(\cdot)$ is defined in equation (36), and $\mathbf{H}(\cdot)$ in equation (31) respectively (all in Appendix A).

(rather than just the first) because given the Markov assumptions, it provides more information (and hence restrictions) on the unknown period 1 type-specific probabilities. For concreteness, consider the observed joint probability that an agent purchases contract C1 and re-treats the net in the second period. This probability is obtained by integrating the corresponding unobserved type-specific joint probabilities of the same event:

$$\mathbb{P}(a_1 = b, a_2 = 1, x_1, x_2|r, z) = \sum_{\tau \in \{\tau_C, \tau_N, \tau_S\}} \pi_\tau(r, z)\mathbb{P}_\tau(a_2 = 1, a_1 = b, x_2, x_1|r, z) =$$

$$\sum_{\tau \in \mathcal{T}} \pi_\tau(r, z)\mathbb{P}_\tau(a_2 = 1|a_1 = b, x_2, x_1, r, z)\mathbb{P}_\tau(x_2|a_1 = b, x_1, r, z)\mathbb{P}_\tau(a_1 = b, x_1|r, z) \tag{6}$$

where $\pi_\tau(r, z)$ denotes the probability that an agent with values (r, z) is of type τ and we refer to this as the type probabilities. We next simplify the right hand side of (6) . First, given the Markov nature of the state transition probabilities and the resultant optimal actions, we note that the conditioning set for the second period choice should be (x_2, r, z).

Next, we assume that conditional upon an agent's type, the preference reversal information (r) is not informative about choice. This is reasonable to the extent that one believes that r is only relevant for actions because it proxies for type. This would be implausible in situations where r provides information about other aspects of the decision process. For instance, if $r = 1$ indicates not just time inconsistency but also reflects a lack of numeracy or other flaws in an agent's cognitive processes, one might believe that it has an independent effect on choice, even after conditioning on type.

Further, we assume that the period 2 state transition probability is independent of type (τ) and the preference reversal variable (r). While this assumption is not directly testable, we can test for whether subsequent transition probabilities vary by (τ, r) since in future periods, types are observed for the sub-population that purchases a product. To the extent that future transitions are type invariant, there is perhaps some reason to believe that this is credible for the second period transition as well. Finally, the identification results only hold for states x_2 that can be reached from a given state x_1. In the assumptions below, we assume that there exist a set of (a_1, x_2, x_1, z, r) with positive probability such that the following statements are true.

ASSUMPTION 10. *(i)* $\mathbb{P}_\tau(a_2|x_2, z, r) = \mathbb{P}_\tau(a_2|x_2, z)$ *(ii)* $\mathbb{P}_\tau(x_2|a_1, x_1, r, z) = \mathbb{P}(x_2|a_1, x_1, z) \neq 0$. [23]

In what follows we suppress the dependence on beliefs z since variation in it plays no role in identifying the type-specific choice probabilities. Under the assumptions above, we can use (6) to define the directly identified quantities $\mathbf{F}$ as

$$\mathbf{F}^{1,2}_{a,x_1,x_2,r} \equiv \frac{\mathbb{P}(a_1 = a, a_2 = 1, x_1, x_2|r, z)}{\mathbb{P}(x_2|a_1 = b, x_1)} = \sum_\tau \pi_\tau(r)\mathbb{P}_\tau(a_2 = 1|x_2)\mathbb{P}_\tau(a_1 = a, x_1)$$

[23] Note that if z is continuously distributed, the statement should be modified to hold a.s..

and correspondingly define

$$\mathbf{F}^1_{a,x_1,r} \equiv \mathbb{P}(a_1 = a, x_1 | r) = \sum_\tau \pi_\tau(r)\mathbb{P}_\tau(a_1 = a, x_1)$$
$$\mathbf{F}^2_{x_2,r} \equiv \mathbb{P}(a_2 = 1 | x_2, r) = \sum_\tau \pi_\tau(r)\mathbb{P}_\tau(a_2 = 1 | x_2)$$

The objects of interest that we need to identify are the type-specific probabilities in period 1 $\{\mathbb{P}_\tau(a_1 = a, x_1)\}_{a\in\{n,b,c\},x_1\in\mathcal{X}_1}$ and the type probabilities $\{\pi_\tau(r)\}_{\tau\in\mathcal{T}}$. We define the period 1 object of interest as the type-specific joint distribution of (a_1, x_1) since that allows the initial state to be type-specific. This is important if, for instance, the past actions of "naïve" time-inconsistent agents makes them more likely to have malaria at the start of the project.

We now outline how the type-specific choice probabilities in the first period are identified. Note that some of these objects are identified by Assumption 4. Specifically the assumption implies that "sophisticated" agents only choose between not purchasing a product and purchasing a commitment product, and "naïve" agents only choose between no purchase and purchasing the standard contract. This was critical for the identification arguments in the previous periods and we can impose them in the following argument to ensure consistency. However, as an examination of the argument below reveals, these direct restrictions are not used for identification of first-period type-specific choice probabilities.[24] The inherent unsatisfactoriness of these assumptions partly motivates the unobserved types model in the next section.

Using the notation above, consider the 3×3 directly identified matrix $\mathsf{P}_{1,r}(a, x_2, x_2', x_1, x_1', r)$ (which we abbreviate to $\mathsf{P}^a_{1,r}$

$$\mathsf{P}^a_{1,r} = \begin{pmatrix} 1 & \mathbf{F}^2_{x_2,r} & \mathbf{F}^2_{x_2',r,f} \\ \mathbf{F}^1_{a,x_1,\tau} & \mathbf{F}^{1,2}_{a,x_1,x_2,r,} & \mathbf{F}^{1,2}_{a,x_1,x_2',r} \\ \mathbf{F}^1_{a,x_1',r} & \mathbf{F}^{1,2}_{a,x_1',x_2,r} & \mathbf{F}^{1,2}_{a,x_1',x_2',r} \end{pmatrix} \tag{7}$$

Then, we can write

$$\mathsf{P}^a_{1,r} = (\mathsf{L}^a_1)'\mathsf{V}_r\mathsf{L}_2$$

where $\mathsf{V}_r = \text{diag}(\pi_{\tau_C}(r), \pi_{\tau_N}(r), \pi_{\tau_S}(r))$ and

$$\mathsf{L}^a_1 \equiv \begin{pmatrix} 1 & \mathbb{P}_{\tau_C}(a_1 = a, x_1) & \mathbb{P}_{\tau_C}(a_1 = a, x_1') \\ 1 & \mathbb{P}_{\tau_N}(a_1 = a, x_1) & \mathbb{P}_{\tau_N}(a_1 = a, x_1') \\ 1 & \mathbb{P}_{\tau_S}(a_1 = a, x_1) & \mathbb{P}_{\tau_S}(a_1 = a, x_1') \end{pmatrix} \quad \mathsf{L}_2 \equiv \begin{pmatrix} 1 & \mathbb{P}_{\tau_C}(a_2 = 1|x_2) & \mathbb{P}_{\tau_C}(a_2 = 1|x_2') \\ 1 & \mathbb{P}_{\tau_N}(a_2 = 1|x_2) & \mathbb{P}_{\tau_N}(a_2 = 1|x_2') \\ 1 & \mathbb{P}_{\tau_S}(a_2 = 1|x_2) & \mathbb{P}_{\tau_S}(a_2 = 1|x_2') \end{pmatrix}. \tag{8}$$

Next, note that L_2 is directly identified by Assumption 4 for $x_2 \in \mathcal{X}_B \cup \mathcal{X}_C$. If it is invertible (which we will assume to be the case below), we can postmultiply both sides by its inverse. Next, the first row of $\mathsf{P}^a_{1,r}\mathsf{L}_2^{-1}$ identifies $(\pi_{\tau_C}(r), \pi_{\tau_N}(r), \pi_{\tau_S}(r))$. We obtain further simplification since under the directly observed type assumptions, $\pi_{\tau_C}(1) = 0$.[25] Once the type probabilities are identified, the type-specific choice probabilities

[24] Specifically, Assumption 4 implies that $\mathbb{P}_{\tau_S}(a_1 = b|x_1) = 0$ and $\mathbb{P}_{\tau_N}(a_1 = c|x_1) = 0$. Note that this assumption cannot be directly tested here since we use it to identify the type-specific choice probabilities for period 2.

[25] The case for $r = 0$ is not interesting since the directly identified types assumption implies that $(\pi_{\tau_C}(0), \pi_{\tau_N}(0), \pi_{\tau_S}(0)) = (1, 0, 0)$. We keep the more general motivation since it anticipates the notation in the next section.

are also identified and the proof is relegated to the appendix. We record the additional assumption required and the subsequent lemma below.

ASSUMPTION 11. *The matrix* L_2 *defined above is invertible for* $x_2, x_2' \in \mathcal{X}_2$.

We can then state the result for period 1:

LEMMA 4. *Consider an agent of type* τ *solving at* $t = 1$ *the problem (3) and that Assumptions 1–11 hold. Then, the first period utility differentials* $u(x_1, b, \tau) - u(x_1, n, \tau)$ *and* $u(x_1, c, \tau) - u(x_1, n, \tau)$ *are identified for all* $x_1 \in \mathcal{X}_1$ *and for all types* τ. *In addition the type probabilities* $\{\pi_\tau(\cdot)\}_{\tau\in\mathcal{T}}$ *are also identified.*

The lemma is useful for at least two reasons: First, we have now identified type-specific utilities for each time period, which along with the identified time parameters, can form the basis for standard model specification tests as well as conducting complete counterfactual analysis. Second, we also identify the relative size of all three different types of agent in the population. This is important because it provides us with the *unconditional* distribution of types whereas previous work (as well as the type classification by observed product choice) provides at best only the distribution of types conditional upon take-up of the offered product. To the extent that the purchase decision is affected by type (e.g. "naïve" agents may be more likely to purchase nets than "sophisticated" agents because they down-weight the future costs of retreatment in the present) the two distributions will be different. Further, this difference, *ceterus paribus*, provides us with a measure of how attractive the commitment contract is for the different types of agents. We explore each of these in the estimation section below.

4.4 Identification: Unobserved Types

We now turn to the case where agent type is not observed. This is reasonable to the extent that we believe that agent observables do not completely reveal type. In particular, in the likely case that the survey responses to the discounting questions and the choice of commitment product are not perfect predictors for type it would be useful to have another approach to identify type-specific preferences.

The primary advantage of this approach is that it is relatively agnostic about the extent to which observed time inconsistency (i.e. the survey responses to the discounting questions as well as the choice of commitment product) map into agent types. Since it nests the perfectly observed types model, one can also test whether the mapping of types in Assumption 4 is appropriate.

The main complication is that we first need to identify type-specific choice probabilities — the key to identification in the previous section — from the observed choice probabilities. Each observed choice probability is a mixture of all the type-specific choice probabilities.

We proceed by making the same kinds of arguments as we made for period one in Section 4.3, and show that both the type probabilities and the type-specific choice probabilities in each period are identified. Once these probabilities are identified we can recover preferences using the same arguments as in Section 4.2.

A key additional set of parameters is now the type probabilities. We define $\pi_\tau(r_u) \equiv \pi_\tau(r, a_1)$ as the probability that an agent is of type τ conditional on their response to the reversal question and their period

1 choice.[26] How information from (r, a_1) is mapped to types will be obtained by imposing assumptions on these probabilities which require imposing a sort of monotonicity (or monotone likelihood ratio) assumption on $\frac{\pi_\tau(r_u)}{\pi_\tau(r'_u)}$.

One set of sufficient conditions for identification is that first, the set of agents with responses $(r, a_1) = (1, c)$ are most likely to be "sophisticated" inconsistent agents and least likely to be time consistent agents. Second, the set of agents with $(r, a_1) = (0, b)$ are most likely to be time consistent agents and least likely to be "sophisticated" inconsistent agents. This implies an ordering on the ratios: $\left\{\frac{\pi_{\tau_C}(r_u)}{\pi_{\tau_C}(r'_u)} \geq \frac{\pi_{\tau_N}(r_u)}{\pi_{\tau_N}(r'_u)} \geq \frac{\pi_{\tau_S}(r_u)}{\pi_{\tau_S}(r'_u)}\right\}$ for $r_u = (0, b)$ and $r'_u = (1, c)$. This assumption appears reasonable in our empirical framework but more generally we only need the following weaker condition to hold (which in fact allows us to test the previous set of conditions):

ASSUMPTION 12. *For some $r_u \neq r'_u$, the three ratios $\left\{\frac{\pi_{\tau_C}(r_u)}{\pi_{\tau_C}(r'_u)}, \frac{\pi_{\tau_N}(r_u)}{\pi_{\tau_N}(r'_u)}, \frac{\pi_{\tau_S}(r_u)}{\pi_{\tau_S}(r'_u)}\right\}$ can be ordered ex-ante.*

Next, we examine the link between the observed choice probabilities and the type-specific choice probabilities. Consider (for $t > 1$ and suppressing dependence on z)

$$\mathbb{P}(a_t = 1, a_{t+1} = 1, x_t, x_{t+1}|r, a_1) = \sum_{\tau \in \mathcal{T}} \pi_\tau(r, a_1)\mathbb{P}_\tau(a_{t+1} = 1, a_t = 1, x_{t+1}, x_t|r, a_1)$$

We simplify the right hand side by assuming

ASSUMPTION 13. *(i) $\mathbb{P}_\tau(a_t = 1|x_t, r_u) = \mathbb{P}_\tau(a_t = 1|x_t)$ $\forall t > 1$ (ii) The transition probabilities do not vary by type and are independent of r_u: $\mathbb{P}_\tau(x_{t+1}|x_t, a_t, r_u) = \mathbb{P}(x_{t+1}|x_t, a_t)$.*

The first part of the assumption asserts that conditional on the state and type, information about an agent's responses to r_u are uninformative about actions. This seems reasonable to the extent that r_u is only informative about choices through its predictive power for agent type. Part of the assumption indeed is trivial, since x_t contains a_1. If, however, r_u provides information about other aspects of the decision process this assumption would be implausible. For instance, if $r = 1$ indicates not just time inconsistency but also reflects a lack of numeracy or other flaws in an agent's cognitive processes, one might believe that it has an independent effect on choice, even after conditioning on type. The second part of the assumption states that the observed transition probabilities do not vary by agent type. Although this assumption is not directly testable, we can examine the type-specific transition probabilities in the directly observed types model to assess its plausibility.

We show identification in two steps. In the first step, we use the observed joint distribution $\mathbb{P}(a_t = 1, a_{t+1} = 1, x_t, x_{t+1}|r_u)$ and the restrictions in Assumptions 12 and 13 to identify the type-specific choice probabilities $\mathbb{P}_\tau(a_t|x_t)$ and the type probabilities $\pi_\tau(r_u)$. The argument is similar to Kasahara and Shimotsu (2009) with the primary difference that we exploit the exclusion restrictions (Assumption 13) that permit identification with just two periods. In the second step, we recover the preference parameters from the type-specific choice probabilities using results from the previous section.

[26] In terms of this probability, the previous section assumed that $\pi_{\tau_S}(1, c) = 1$, $\pi_{\tau_C}(0, \cdot) = 1$, and $\pi_{\tau_N}(1, b) = 1$. The type-probabilities can also depend upon (z, v) but we ignore that dependence here since it is not used for identification.

We first introduce the notation required to state two central assumptions for the first step in the argument. Define the 3×3 matrices containing the objects of interest as[27]

$$\mathsf{L}_2 \equiv \begin{pmatrix} 1 & \mathbb{P}_{\tau_C}(a_2 = 1, x_2) & \mathbb{P}_{\tau_C}(a_2 = 1, x_2') \\ 1 & \mathbb{P}_{\tau_N}(a_2 = 1, x_2) & \mathbb{P}_{\tau_N}(a_2 = 1, x_2') \\ 1 & \mathbb{P}_{\tau_S}(a_2 = 1, x_2) & \mathbb{P}_{\tau_S}(a_2 = 1, x_2') \end{pmatrix} \mathsf{L}_3 \equiv \begin{pmatrix} 1 & \mathbb{P}_{\tau_C}(a_3 = 1|x_3) & \mathbb{P}_{\tau_C}(a_3 = 1|x_3') \\ 1 & \mathbb{P}_{\tau_N}(a_3 = 1|x_3) & \mathbb{P}_{\tau_N}(a_3 = 1|x_3') \\ 1 & \mathbb{P}_{\tau_S}(a_3 = 1|x_3) & \mathbb{P}_{\tau_S}(a_3 = 1|x_3') \end{pmatrix} \quad (9)$$

ASSUMPTION 14. *There exist (x_2, x_2', x_3, x_3') such that the matrices L_2 and L_3 defined above are invertible.*

This assumption requires that there is sufficient variation in the type-specific choice probabilities. It formalizes the intuition that if optimal choice probabilities do not differ very much across types then it will be difficult to identify them separately. In the context of the empirical section, it requires that the consistent, "naïve" and "sophisticated" agents indeed have different choice probabilities. This assumption is reasonable here to the extent that the model is only interesting if it is true.

ASSUMPTION 15. *All types exist with positive probability for at least two values of r_u: $\pi_\tau(r_u) > 0 \; \tau \in \mathcal{T}$*

This assumption requires that all types exist with positive probability conditional on r_u. This is reasonable since the motivation for the unobserved type model in the first place was that the proxies were imperfect predictors of type. Moreover, this is a convenient assumption and specific departures from it can be accommodated. For instance, the directly observed types model in the previous section does not satisfy this assumption but the model is nevertheless identified (in fact, the identification arguments are much easier).[28]

LEMMA 5. *Under Assumptions 13 and 14, the choice specific probabilities $\mathbb{P}_\tau(a_t = 1|x_t)$ are identified for all $x_t \in \mathcal{X}_B \cup \mathcal{X}_C$ and $t > 1$. In addition, the type probabilities conditional on r and first-period choice, $\{\pi_\tau(r_u)\}_{\tau\in\mathcal{T}}$, are also identified.*

LEMMA 6. *Consider an agent of type τ solving the problem (3) and suppose that Assumptions 1–3 and 5–15 hold. We observe an i.i.d. sample on $(\{a_t^*, x_t\}_{t=1}^{T-1}, w)$. Then, we can identify*

1. *The type-specific utility differentials $u_t(x_t, 1, \tau) - u_t(x_t, 0, \tau) \quad \forall\, \tau \in \mathcal{T},\; x_t \in \mathcal{X}_B \cup \mathcal{X}_C$*

2. *The exponential discount parameter δ and the hyperbolic parameters $\beta_\tau \quad \forall \tau \in \mathcal{T}$.*

The proof is a combination of applying Lemma 5 and the results from Section 4.1. These results have several interesting sets of implications. First, since we identify the unconditional (on first period choice) type probabilities we can make statements about the sizes of the different types in the population. This is important since it provides us with a quantitative measure of how large a fraction of the population is time inconsistent. Second, we also identify the conditional (on first period choice) type probabilities which allow us to make statements about what the relative weights of the different types are when we stratify

[27] We identify the joint distribution of (a_2, x_2). An alternative specification would be to write this joint distribution as $\mathbb{P}_\tau(a_2|x_2)\mathbb{P}_\tau(x_2)$ and then impose that $\mathbb{P}_\tau(x_2) = \mathbb{P}(x_2)$. We do not follow this approach here since it is somewhat less general.

[28] Importantly, we have assumed that the total number of types is known. In principle one could test this assumption using the arguments outlined in Kasahara and Shimotsu (2009).

by first period choice. This allows us to make statements about the relative appeal of the commitment contract versus the regular contract for different types.

Second, since we allowed significant flexibility in type preferences we can evaluate the strength of alternative explanations in explaining observed behavior. Concretely, we can compare a model in which different types have identical per-period preferences (but differ in their time preference parameters) to a model where agents are time consistent but differ along other dimensions of their preferences.

5 Monte Carlo Simulations

We illustrate the properties of these models with a set of Monte Carlo simulations, beginning with the case where types are directly observed and then moving on to the unobserved types case. In order to focus attention on the estimation of the time preference parameters, we provide a simple parametrization for per-period utilities, imposing that they are common across types and are linear in the unknown parameters.

We begin by specifying utility in each period as a function of the state variables and actions taken.

- **Period 4**: $x_4 \in \{0, 1\}$.

$$u(x_4) = -\theta_4 x_4$$

- **Period 3**: $x_3 \in \{bm, bh, cm, ch, nh, nm\} \equiv \{b, c, n\} \times \{h, m\} \equiv \{0, 1, 2, 3, 4, 5\}$ and $a \in \{0, 1\}$

$$u(x_3, a) = -3\mathbb{I}\{x_3 \in \{1, 3, 5\}\} - \theta_5 p_r \mathbb{I}\{x_3 \in \{0, 1\}, a = 1\}$$

  where p_r is the price of retreatment.

- **Period 2**: $x_2 \in \{bm, bh, cm, ch, nh, nm\} \equiv \{b, c, n\} \times \{h, m\} \equiv \{0, 1, 2, 3, 4, 5\}$ and $a \in \{0, 1\}$

$$u(x_2, a) = -3\mathbb{I}\{x_2 \in \{1, 3, 5\}\} - \theta_5 p_r \mathbb{I}\{x_2 \in \{0, 1\}, a = 1\}.$$

- **Period 1**: $x_1 \in \{h, m\} \equiv \{0, 1\}$ and $a \in \{b, c, n\}$

$$u(x_1, a) = -3\mathbb{I}\{x_1 = 1\} - \theta_5 p_b \mathbb{I}\{a_1 = b\} - \theta_5 p_c \mathbb{I}\{a_1 = c\},$$

  where p_b is the price of the no-commitment contract and p_c is the price of the commitment contract.

We assume that the unobserved state variables ϵ_t are independent Type I extreme-valued so that we obtain a simple characterization of the choice probabilities

$$\mathbb{P}_\tau(a_t = a | x_t, z) = \frac{\exp(v_\tau(x_t, a, z))}{\sum_{s=0}^{S} \exp(v_\tau(x_t, s, z))}$$

where the $v_\tau(\cdot)$ functions are constructed using backward induction. The results for the case where we assume that types are directly observed are presented in Table 4 and those for the unobserved types case are presented in Table 5. We present results for a range of sample sizes. The tables provide evidence that the model is identified and that at least for moderate sample sizes (300 and above) the estimators are reasonably close to the true values. (we present both the median and mean for the point estimates). In addition, examining the standard errors across sample sizes provides corroborating evidence that the estimators converge at the parametric rate as hypothesized.

6 Estimation

We turn next to solving and estimating the model using data from the intervention. We begin by specifying preferences and then discuss the transition probabilities and other key ingredients of the dynamic programming problem. As is standard in the literature, we assume that the unobserved components of preferences have the Type I generalized extreme value (GEV henceforth) distribution.

The central difference from standard analyses of dynamic models in what follows is the presence of the different types of time-inconsistent agents and the fact that these types are unobserved. These two deviations alter the standard results and we highlight these differences below.

6.1 Preferences

6.1.1 Period 4

In period 4, the state variables are income and health ($x_4 = (y_4, h_4)$) while v are household characteristics that enter preferences. In the base specification v includes household size (at baseline), a measure of households assets, and an indicator of risk aversion ($v = (v_{\texttt{hhs}}, v_{\texttt{assets}}, v_{\texttt{risk}})$).[29] Preferences in period 4 are given by

$$u_\tau(x_4; v) = \mathcal{C}(x_4)^{\alpha_\tau(v)} - c_\tau(x_4, v)$$

where $\mathcal{C}(x_4)$ is consumption in state x_4 and the exponent $\alpha_\tau(\cdot)$ captures risk aversion with respect to consumption. The function $c_\tau(\cdot)$ is the direct disutility from malaria (i.e. not including the loss in utility that arises from lower consumption due to potentially lower income when sick).

We allow for heterogeneity in the risk-aversion parameter along both observable (v) and unobservable dimensions (τ):[30]

$$\alpha_\tau(v) = \texttt{Logit}\left(\alpha_{0,\tau} + \alpha_1 v_{\texttt{hhs}} + \alpha_2 v_{\texttt{assets}} + \alpha_3 v_{\texttt{risk}}\right) \tag{10}$$

Note that the intercept differs by type; in principle, one could allow the entire function to vary by type, but the sample size requirements for such flexibility are onerous. We have also restricted the risk aversion parameters to be between 0 and 1. The latter restrictions can easily be relaxed but we keep them because they simplify estimation and are reasonable given the various empirical estimates in the literature.

[29] The survey based measure of attitudes towards risk is obtained by using an abbreviated version of the procedure proposed in Holt and Laury (2002). See Appendix D for details of the variables used.

[30] where Logit(x)= $\frac{\exp(x)}{1+\exp(x)}$

The malaria disutility function $c_\tau(\cdot)$ is also allowed to vary by agent type as well as observables v. Specifically, we model the disutility as

$$c_\tau(x_4, v) \equiv h_4 c_\tau(v) = \mathbb{I}\{h_4 = m\} \exp(\kappa_{0,\tau} + \kappa_1 v_{\texttt{hhs}} + \kappa_2 v_{\texttt{assets}}), \tag{11}$$

where h_4 is an indicator equal to 1 if someone in the household contracts malaria in period 4. Note that the survey based measure of risk does not enter the cost of malaria function and this exclusion restriction is useful in identifying the two functions.

6.1.2 Periods 2 and 3

The state variables in period $t \in \{2, 3\}$ are income (y_t), health status (h_t) and the choice of product in period 1 (a_1). We define utility in period t as

$$u_\tau(x_t, a_t; v) = (\mathcal{C}(x_t) - p_r a_t \mathbb{I}\{a_1 = b\})^{\alpha_\tau(v)} - c_\tau(v) h_t$$

where p_r is the price of retreatment. Note that agents who have purchased the commitment contract do not incur any additional cost in retreating the net.

6.1.3 Period 1

In period 1, preferences are given by

$$u_\tau(x_1, a_1; v) = (\mathcal{C}(x_1) - p_b \mathbb{I}\{a_1 = b\} - p_c \mathbb{I}\{a_1 = c\})^{\alpha_\tau(v)} - c_\tau(v) h_1$$

where p_b is the price of the ITN alone, and p_c is the price of the commitment product.

6.2 Solving the Model: Period 3

Given a finite horizon model, we can solve for the optimal decision rule using backward induction. We solve and estimate the model using the mapping between type-specific choice probabilities and type-specific value functions (defined below). This is justified because even though we do not observe types, the type-specific choice probabilities are identified using Lemma 5. Next, recall that in period 3 an agent who has purchased an ITN will choose to re-treat his net if

$$u_\tau(x_3, 1; v) + \epsilon_3(1) + \beta_\tau \delta \int u(x_4; v)\, \mathrm{dF}(x_4 | x_3, 1; z) > u_\tau(x_3, 0; v) + \epsilon_3(0) + \beta_\tau \delta \int u(x_4; v)\, \mathrm{dF}(x_4 | x_3, 0; z)$$

which we rewrite as

$$\epsilon_3(0) - \epsilon_3(1) < v_\tau(x_3, 1, z, v, \beta_\tau) - v_\tau(x_3, 0, z, v, \beta_\tau).$$

Recall that $v_\tau(\cdot)$ is the choice specific value function for type τ:

$$v_\tau(x_3, a, z, v, \beta_\tau) \equiv u_\tau(x_3, a; v) + \beta_\tau \delta \int u_\tau(x_4, v)\, \mathrm{dF}(x_4 | x_3, a, z) \tag{12}$$

We emphasize the dependence of these functions on the hyperbolic parameter β since it will be useful for future manipulations. We assume that $(\epsilon_3(0), \epsilon_3(1))/\sigma$ are i.i.d. standard GEV (to ease notation we set $\sigma = 1$ in what follows). Under the GEV assumption, the type-specific choice probability is given by

$$\mathbb{P}_\tau(a_3 = 1|x_3; z, v) = \frac{\exp(v_\tau(x_3, 1, z, v, \beta_\tau))}{\sum_{j=0}^{1} \exp(v_\tau(x_2, j, z, v, \beta_\tau))}. \tag{13}$$

6.2.1 Period 2 Choice

Under the GEV assumption on the errors, the choice probabilities are given as usual by

$$\mathbb{P}_\tau(a_2 = 1|x_2; z, v) = \frac{\exp(v_\tau(x_2, 1, z, v, \beta_\tau))}{\sum_{j=0}^{1} \exp(v_\tau(x_2, j, z, v, \beta_\tau))} \tag{14}$$

For period 2, the $v_\tau(\cdot)$ functions are defined similarly as in period 3 except that the calculation of the integral is more involved. As we show below, the form of these functions provides some insight into the time inconsistency problem

$$v_\tau(x_2, a, z, v, \beta_\tau) \equiv u_\tau(x_2, a; v) + \beta_\tau \delta \int v_\tau^*(s_3, z, v)\, \mathrm{dF}(s_3|x_2, a, z) \tag{15}$$

where $v_\tau^*(\cdot)$ represents the utility (from the point of view of period 2) of behaving optimally in period 3, given action a_2 and state x_2 and $s_3 = (x_3, \epsilon_3)$. To make the notation for this function easier, define the event

$$A_k^\tau \equiv \{v_\tau(x_3, k, z, v, \tilde{\beta}_\tau) + \epsilon_3(k) > v_\tau(x_3, s, z, v, \tilde{\beta}_\tau) + \epsilon_3(s) \quad \forall s \neq k\} \tag{16}$$

which is the event that action k is optimal in period 3 given a τ agent's future expected present-bias of $\tilde{\beta}_\tau$. Next, define[31]

$$v_\tau^*(s_3, z, v) \equiv \sum_{s=1}^{J} \left(v(x_3, s, z, v, 1) + \epsilon_3(s)\right) \mathbb{I}_{A_s^\tau}$$

[31] where $\mathbb{I}_{A_s^\tau}$ is an indicator for the event A_s^τ

Integrating with respect to the GEV error distribution yields

$$\begin{aligned} v_\tau^*(x_3, z, v) &\equiv \int v_\tau^*(x_3, \epsilon_3, z, v)\, \mathrm{dF}(\epsilon_3) \\ &= \sum_{s=1}^{J} \int (v_\tau(x_3, s, 1) + \epsilon_3(s)) \mathbb{I}_{A_s^\tau}\, \mathrm{dF}(\epsilon_3) \\ &= \sum_{s=1}^{J} \mathbb{P}(A_s^\tau) \Big(v_\tau(x_3, s, z, v, 1) + \mathbb{E}(\epsilon_3(s) | A_s^\tau) \Big) \\ &= \sum_{s=1}^{J} \mathbb{P}(A_s^\tau) \Big(v_\tau(x_3, s, z, v, 1) + \mathbb{E}(\epsilon_3(s) | A_s^\tau) \Big) \\ &= \sum_{s=1}^{J} \mathbb{P}(A_s^\tau) \Big(v_\tau(x_3, s, z, v, 1) + \gamma_{\texttt{euler}} - \log \mathbb{P}(A_s^\tau) \Big) \end{aligned}$$

where $\gamma_{\texttt{euler}}$ is Euler's constant, and

$$\begin{aligned} \mathbb{P}(A_s^\tau) &\equiv \int \mathbb{I}_{A_s^\tau(x_3, \epsilon_3, z, v)}\, \mathrm{dF}(\epsilon_3) \\ &= \frac{\exp(v_\tau(x_3, s, z, v, \tilde{\beta}_\tau))}{\sum_{j=1}^{J} \exp(v_\tau(x_3, j, z, v, \tilde{\beta}_\tau))}. \end{aligned}$$

After simplification we obtain

$$v_\tau^*(x_3, z, v) = \sum_{s=1}^{J} \mathbb{P}(A_s^\tau) \left[v_\tau(x_3, s, z, v, 1) - v_\tau(x_3, s, z, v, \tilde{\beta}_\tau) + \gamma_{\texttt{euler}} + \log \left(\sum_{j=1}^{J} \exp(v_\tau(x_3, j, z, v, \tilde{\beta}_\tau)) \right) \right]$$

The first two terms in the brackets give the difference between the net value of taking action s assuming no present-bias and the net value of the same action assuming a present-bias of $\tilde{\beta}_\tau$. For time consistent agents who are not present-biased and know this, this term disappears and the expression reduces to the standard expression for discrete choice models (see e.g. equation (12) of Aguirregabiria and Mira, 2010). For "naïve" time-inconsistent agents this expression is also zero and reflects the fact that such agents do not take their future present-bias into account while making choices. In this case such agents are ignoring their future present bias in period 3 while making choices in period 2. In contrast, "sophisticated" inconsistent agents recognize their future inconsistency and this expression is non-zero for them.

6.2.2 Period 1 Choice

The discussion is very similar to that for the previous period with the only substantive difference that there are three possible actions in period 1 and the choice probabilities for an agent of type τ are given by

$$\mathbb{P}_\tau(a_1 = a | x_1; z, v) = \frac{\exp(v_\tau(x_1, a, z, v, \beta_\tau))}{\sum_{j=0}^{2} \exp(v_\tau(x_1, j, z, v, \beta_\tau)}. \tag{17}$$

For period 1, the $v_\tau(\cdot)$ functions are

$$v_\tau(x_1, a, z, v, \beta_\tau) \equiv u(x_1, a; v) + \beta_\tau \delta \int v^*_\tau(s_2, z, v)\,\mathrm{dF}(s_2|x_1, a, z), \tag{18}$$

where $v^*_\tau(\cdot)$ is the value function for period 2 (for type τ).
As before, define

$$A^\tau_k \equiv \{v_\tau(x_2, k, z, v, \tilde{\beta}_\tau) + \epsilon_2(k) > v_\tau(x_2, s, z, v, \tilde{\beta}_\tau) + \epsilon_2(s) \quad \forall s \neq k\}, \tag{19}$$

which is the event that action k is optimal in period 2 given an agent's future expected present-bias of $\tilde{\beta}_\tau$ and where

$$v_\tau(x_2, a, z, v, c) = u_\tau(x_2, a; v) + c\delta \int v^*_\tau(s_3, z, v)\,\mathrm{dF}(s_3|x_2, a, z), \tag{20}$$

which is the choice specific value function in period 2 given an agent's hyperbolic parameter c. Next, as earlier

$$v^*_\tau(s_2, z, v) = \sum_{s=0}^{2} (v_\tau(x_2, s, z, v, 1) + \epsilon_2(s))\mathbb{I}_{A^\tau_s},$$

and as before, integrating with respect to the error distribution we obtain

$$\begin{aligned} v^*_\tau(x_2, z, v) &\equiv \int v^*_\tau(x_2, \epsilon_2, z, v)\,\mathrm{dF}(\epsilon_2) \\ &= \sum_{s=0}^{2} \mathbb{P}(A^\tau_s)(v_\tau(x_2, s, z, v, 1) + \gamma_{\mathtt{euler}} - \log \mathbb{P}(A^\tau_s)), \end{aligned}$$

where

$$\mathbb{P}(A^\tau_s) = \frac{\exp(v_\tau(x_2, s, z, v, \tilde{\beta}_\tau))}{\sum_{j=0}^{2} \exp(v_\tau(x_2, j, z, v, \tilde{\beta}_\tau))},$$

so that simplifying, we obtain

$$v^*_\tau(x_2, z, v) = \sum_{s=0}^{2} \mathbb{P}(A^\tau_s) \left[v_\tau(x_2, s, z, v, 1) - v_\tau(x_2, s, z, v, \tilde{\beta}_\tau) + \gamma_{\mathtt{euler}} + \log\left(\sum_{j=0}^{2} \exp(v_\tau(x_2, j, z, v, \tilde{\beta}_\tau)) \right) \right].$$

6.3 Estimation Method

Estimation is carried out in two steps. In the first step, we identify the type-specific choice probabilities $\mathbb{P}_\tau(a_t|x_t, z, v)$ using the methods outlined in Lemma 5. In the second step, we use these probabilities to conduct a minimum distance estimation exercises using equations (13), (14) and (17) as the mapping between the identified type-specific choice probabilities and the type-specific probabilities predicted by the model.

In the first step, we recover the type-specific choice probabilities for each state variable x_t at each value of the time-invariant exogenous variables (z, v). Note that the identification results require x_t to be discrete and so the state support comprises health status (binary valued), income levels (binary valued)[32] and for $t \in \{2,3\}$ also includes period 1 choice. The time-invariant variables comprise $(z, v) = (v_{\texttt{hhs}}, v_{\texttt{assets}}, v_{\texttt{risk}}, \pi, \delta, \gamma)$. Each such calculation follows the steps outlined in equations (40)–(43). Note that though the state space support is required to be discrete, this is not the case for (z, v). However, given the steps involved in the Kasahara and Shimotsu inversions and the sample size, we discretized (z, v) into 108 unique cells and implemented Lemma 5 in each cell to obtain three type-specific probabilities per cell. In addition, we estimated the quantities defined in (37) and used in Lemma 5 using flexible logit specifications. This procedure yields type probabilities (the V matrix) that are functions of (z, v). However, these were very imprecisely estimated within each cell, and so we averaged these weights over the distribution of (z, v) to arrive at the type probabilities reported in the paper.

In the second step, we begin by calculating the model choice probabilities — the right hands side of equations (13), (14) and (17) — for a given set of parameter values $\theta \equiv (\delta, \beta_{\tau_N}, \beta_{\tau_S} \boldsymbol{\alpha}, \boldsymbol{\kappa})$. δ is the usual exponential discounting parameter, $(\beta_{\tau_N}, \beta_{\tau_S})$ are the hyperbolic parameters for the "naïve" and "sophisticated" agents respectively, and $(\boldsymbol{\alpha}, \boldsymbol{\kappa})$ are the parameter vectors characterizing risk aversion and malaria disutility respectively, (see equations 10 and 11). These model probabilities are calculated for each candidate value θ starting with the value functions for the last period and then working backwards using equations (12), (15) and (18). In order to compute these value functions we also need estimates of the transition probabilities $\mathrm{dF}(x_{t+1}|x_t, a_t)$ used by households in solving the problem. We obtain these using the households' elicited beliefs about the two stochastic components of this distribution (income and health) along with information on the monetary costs of illness. Beliefs were elicited during the baseline as well as in the follow-up survey. We estimated the model using both sets of beliefs as well as only using the baseline beliefs, and the results were insensitive to the choice of beliefs.[33] In what follows we present the results using only baseline beliefs. Finally, we choose our estimate of θ_0 to be the value of θ that minimizes the distance between these model probabilities and the type-specific choice probabilities recovered in the first step using the Kasahara-Shimotsu inversion.

6.4 Reduced Form Regressions

Before discussing the estimates of the structural model, we first present some descriptive statistics and basic reduced form regressions intended to highlight the basic correlations in the data. In Table 6 we predict product take up as a function of the period 1 state variables as well as the exogenous variables listed above. In column 1 the dependent variable is equal to 1 if the household purchased at least one net using either loan product (C1 or C2) while in column 2 we report the probability of choosing contract C1 restricting attention to the sub-sample that purchased at least one net. Finally, in column 3 we estimate a multinomial logit with values of 0,1, and 2 for no purchase, C1 purchase and C2 purchase respectively. The results suggest that being sick with malaria in the baseline makes households substantially more likely (by seven percentage points or 16% of the take-up rate) to purchase one of the ITN products ($p = .04$). Conversely, higher period income in period 1 ($p = .09$) and greater asset ownership at baseline ($p = .07$) are both associated with substantial declines (8–10 percentage points) in the likelihood of purchasing a product

[32] see Appendix D for more details on variable construction.

[33] Households' stated beliefs about the efficacy of ITNs decreased by about 10 percentage points over the course of the intervention. The most significant change in beliefs was a 30% decrease in the probability of contracting malaria while sleeping unprotected by a net.

suggesting that poorer household were somewhat more likely to purchase a product. Greater beliefs in the efficacy of ITNs (captured by the γ variable) were associated with higher propensity to purchase a product, although in a non-linear fashion. Note that one implication of the forward-looking model is that the effect of the beliefs in the efficacy of nets γ on take up in the first period will be non-linear (since these effects will multiply each other due to the recursive nature of the value functions) and we use the quadratic form as a convenient approximation.

In order to see whether these variables have greater predictive power for any one of the contract choices, we next break down the dependent variable into three values and implement a multinomial logit. The results suggest that beliefs are stronger predictors, both substantively and statistically, for contract C2 which is consistent with the idea that households with stronger beliefs about ITN efficacy are more likely to purchase contract C2. In addition, there is some evidence that poor households are more likely to choose contract C1 (column 3a). Interestingly, net ownership (which was about 2/3 at baseline) is not a strong predictor for take-up once we condition on beliefs. Finally, it is also interesting to note that the survey measure of time-inconsistency is not strongly correlated with the choice of contract (although it does have predictive power for retreatments as we discuss below).

In Table 7 we predict (using the logit) retreatment rates as a function of the state variables (including contract choice) and time-invariant exogenous variables, restricting attention to the sub-sample that purchased one of the two products. The most striking result (column 1) is that retreatment rates for households that purchased contract the commitment product are much higher (in fact almost double) than for households that purchased C1 and this remains true for the second retreatment as well.

Turning next to the determinants of retreatment among households who choose contract C1, we see that the survey measure of time-inconsistency is highly significant: households that have at least one preference reversal are about 22 percentage points less likely to re-treat any ITNs ($p = .02$). These results are consistent with a situation where some buyers are naïvely time-inconsistent. Turning next to determinants of retreatment among purchasers of the commitment contract, we first note that the unconditional retreatment rates are very high (over 90%) so there is relatively little variation in the dependent variable. As was the case in the take-up regressions an increase in the beliefs about efficacy of ITNs is associated with substantially higher retreatment rates for this sub-sample of households as well.

In sharp contrast to the results from the first period, malaria incidence in the past six months is *negatively* associated with the decision to re-treat for both types of contract. In particular, the likelihood of the second retreatment decreases by 13 percentage points if someone in the household contracted malaria after the ITN purchase but before the first retreatment (p=.04). The sign remains the same for the second retreatment (columns 4 and 5) although it is no longer statistically significant. Interestingly, for the second retreatment, we find that the relationship between baseline beliefs and retreatment reverses for the two kinds of contract purchasers. For C1 holders, the relationship between beliefs and retreatment is increasing at high levels of belief while the relationship reverses for C2 holders. These relationships point to the need to impose more structure on the data to understand the relationship between purchase, retreatment and household characteristics.

The regressions in this section provide useful information about the importance of beliefs, past health status and income in driving health decisions as well as some suggestive evidence about the potential importance of time-inconsistency in retreatment decisions. In the following section we attempt to quantify these links in the context of the model described in section above.

6.5 Estimation Results

6.6 The Population Distribution of Time-Inconsistent Agents

We next report the results from the estimation exercise outlined in Section 6.3. We begin first with the estimates of the type probabilities. We first discuss the type probabilities conditional only on the survey responses r, $\{\mathbb{P}(\tau|r) : \tau \in \{\tau_C, \tau_N, \tau_S\}\}$, before further conditioning upon first period choice, $\{\mathbb{P}(\tau|r, a_1) : \tau \in \{\tau_C, \tau_N, \tau_S\}\}$. The former probabilities provide an estimate of the unconditional distribution of types in the population, something that has not previously been estimated in the literature. The latter probabilities estimate the distribution of types conditional on purchase and a comparison of these distributions provides us with an estimate of how likely "sophisticated" agents are to choose commitment products. Recall that we cannot observe agent types, so that direct choice probabilities are not informative about the selection problem here.

Note that our identification assumptions only imposed a monotonicity condition (Assumption 12) on the relationship between (r, a_1) and types. Therefore, we can actually test whether the more conventional, stronger, mappings hold. In particular, we can test whether (i) the sub-population with $r = 1$ has a greater fraction of time-inconsistent agents and (ii) whether the sub-population that purchase the commitment product (and had $r = 1$) is more likely to contain "sophisticated" time-inconsistent agents.

The results from Table 8 show that our estimate of the fraction of time consistent agents in the population is about 40% and the bulk of the time-inconsistent agents (about 80%) are "naïve". These figures are approximately the same regardless of whether the sub-population exhibited preference reversals (i.e $r = 1$). Assumption 12 is weak, in the sense that it does not require that the fraction of inconsistent agents be larger in the sub-population with $r = 1$. In fact, we find that the fraction of time-inconsistent agents is actually (statistically insignificantly) smaller in the sub-population with $r = 1$. These results suggest that, at least in our sample, the conventional deterministic mapping of time-consistency to standard time preference survey questions does not hold.

We next discuss the type probabilities conditional upon first period choice in Table 9 for each of the four combinations of the conditioning variables (r, a_1). First, as with the previous table, the questionnaire based measures of time-inconsistency and the choice of commitment product do not perfectly predict agent type. In fact, all three types of agent exist for *every* value of these variables. Recall that the directly observed types model assumes that $\pi_C(0, \cdot) = 1$, $\pi_N(1, b) = 1$, and $\pi_S(1, c) = 1$. Table 9 suggests that this is not the case with our data. For instance, among the sub-population of agents who exhibit preference reversals on discounting questions (and purchase either C1 or C2) around a third are estimated to be time-consistent agents. Even more starkly, this fraction is about the same (in fact somewhat lower) in the sub-population that did not exhibit any such preference reversals. In addition, across all combinations of of (r, a_1), time-consistent agents are a minority ranging roughly between $30 - 40\%$. "Naïve" inconsistent agents form the bulk — approximately $70 - 80\%$ — of the time-inconsistent agents. Note that in the context of the model, preference reversals are an imperfect proxy for time-inconsistency and are affected by measurement (misclassification) and cognitive issues as well as factors such as seasonality and other constraints.

Finally, comparing results from Tables 8 and 9, we conclude that among households who purchased the commitment product: (i) the proportion of time consistent agents is almost 30% lower in the sub-population with $r = 1$ relative to the overall mean; (ii) the proportion of "naive" agents is 19% higher than the overall fraction, while that of "sophisticated" agents is 22% higher.

6.7 Estimates of Time, Cost and Risk Preferences

Next, we describe the estimates of the preference parameters for each of the three types of agents, either assuming equal risk preferences (Table 10) or allowing for heterogeneity in risk-preferences across types (Table 11). It is comforting to note that the estimates for time-preference parameters remain stable across the different utility specifications.

First, note that the two types of time-inconsistent agent have different rates of time preference. In particular, in both specifications the "hyperbolic" parameters for the "sophisticated" agents are significantly smaller — about 70% lower — than those for the "naïve" agent and this difference is statistically significant at conventional levels. This result suggests that there is substantial heterogeneity across inconsistent agents beyond that usually considered in empirical work on the issue. Second, the "hyperbolic" parameter for the "naïve"-inconsistent type is quite close to 1, so that present-bias may be a relatively small problem. These results suggest the hypothesis that agents with strong present-bias problems are likely to become more aware of their problem and take actions to mitigate the problem (i.e. by becoming "sophisticated"). Conversely, "naïve" agents are less likely to take actions to mitigate their present-bias problem because it is relatively small.[34]

Finally, we deal with the hypothesis that types differ by more than just their time preferences. Specifically, we allow different types to have different rates of risk aversion[35] as well as different cost parameters associated with illness. Table 11 shows that there are differences in risk attitudes for different types but these differences are quantitatively small[36] and, as we shall see subsequently, also substantively unimportant. The point estimates for the cost parameters follow this pattern as well. These results together suggest that while there is some evidence that types differ in their attitudes towards risk or the disutility from illness, these differences are small relative to the differences in time preferences documented above.

7 Counterfactuals

In this section, we carry out a series of counterfactual exercises using the estimated model to (i) assess the effect of changes in the model's exogenous parameters and (ii) evaluate the relative importance of risk, cost and time preference parameters in explaining outcomes. We consider several exogenous changes in the price of retreatment and report their effects on ITN take-up and retreatment rates. In addition, we examine outcomes for the different types of agent and assess the relative importance of the four ways in which agents differ — risk, cost, degree of present-bias and awareness of future present-bias — in generating the variations in outcomes.

[34] There is relatively little work, theoretical or empirical, on self-control and learning. See Ali (2010) for a theoretical model of a decision-maker who is only partially aware about his self-control problems and learns about them over time.

[35] In addition risk aversion varies by observables

[36] Recall the risk aversion is parameterized by equation (10).

We first discuss the consequences of a known doubling in the price of retreatments (balanced by a corresponding increase in the price of the commitment contract).[37] Intuitively, the price change has several effects. First, the increases will reduce contemporaneous demand for retreatment through a substitution and income effect. Second, the price increase may reduce overall ITN adoption in the initial period, because the dynamic nature of the problem implies that agents predict that the cost of enhancing the protective power of the net with the treatment has increased. Third, a sophisticated hyperbolic agent who cares about the retreatment may switch from the standard to the commitment contract, anticipating that the higher cost of re-treatment will cause more temptation in the future. This latter effect is, however, moderated by the effect of the corresponding increase in the price of the commitment product. In practice, which effect dominates in the first period is an empirical question that the counterfactuals can answer.

7.1 Adoption and Retreatment: Averaging Across Types

Consider first a doubling of the retreatment price from Rs. 15 to Rs. 30. First, averaging across types, demographics and states, we find that retreatment rates under contract C1 fall by about 74% (see Table 12). These figures are large and incorporate the effects of prices via preferences (including time-inconsistency) and beliefs on retreatment. They suggest a price elasticity of demand of -0.72. Second, not surprisingly, we find no effect on retreatment decisions under contract C2, since retreatment price increases have no effect on those who have committed to retreatment. However, it could be the case that the increased price of retreatment increases demand for purchasing contract C2 in the first period. Weighing against this effect is the corresponding increase (by Rs. 30) in the price of the commitment product so that the final change in the demand of C2 is an empirical question. According to the model, demand for the commitment product C1 declines by 5% while demand for C2 increases by a little under 2%. Given that slightly more agents purchase C2 and retreatment rates for C2 remain unchanged, overall retreatment rates (averaging across all contracts) decline by approximately 12%.

We next examine (also in Table 12) changes in take-up and retreatment when the price of retreatment is halved (from Rs. 15 to Rs. 7.5). First, we find that retreatment rates for purchasers of contract C1 are predicted to increase by 82% which is similar to the proportional change in the exercise when prices were doubled. Further, we find that overall purchase of the standard contract increases by 17% while demand for the commitment contract declines by 8%. Incorporating the differences in take-up of contracts we find that overall retreatment rates averaged across the population increase by approximately 12% which is again similar in absolute terms to the corresponding figure in the previous counterfactual.

We carried out similar analyses for other price changes and found similar results. Broadly speaking, increasing retreatment prices generally leads to a significant decline in retreatment rates for households that purchased contract C1 with an elasticity of approximately -0.7 to -0.8. Further, these price increases also lead to some wealthier households switching towards the commitment contract, though this effect is not very large.

[37] Counterfactuals without the corresponding increase in the price of C2 imply that demand for C2 increases unambiguously. We omit these results here.

7.2 Adoption and Retreatment: Type Differences

Households differ in their time preferences, risk and cost preferences as well as awareness about their future present-bias. The structural model combined with the estimates from the previous section allows us to understand the relative contributions of each of these differences in explaining take-up rates. Concretely, in this sub-section we seek to use the structure to understand why even wealthy "sophisticated" households do not find the commitment contract appealing.

As might be anticipated from the empirical results, the commitment contract does not seem to appeal exclusively to "sophisticated" inconsistent agents. In particular, wealthier households are much more likely to purchase the commitment product (regardless of time preferences). Second, even among wealthy households, "sophisticated" households are actually less likely to purchase the commitment contract. We explore the reasons for this using the structural estimates, varying each of the model elements (time and risk preferences, costs, and awareness of future present-bias) systematically.

The results suggest that the primary reason for the lower adoption of contract C2 by "sophisticated" households is their degree of present-bias β_S. We find evidence for this in two ways. First, we carry out counterfactuals for a model where all types have the same risk and cost preferences but different hyperbolic parameters (Model 2 in Table 13). In this model, types differ only in their hyperbolic parameters and their awareness of their future present-bias. We find that the take-up results for contract C2 are identical to those obtained from the general model where risk and cost preferences as well as time preferences are allowed to vary by type (Model 4 in Table 13). This suggests that variations in risk and cost preferences across types are not important in explaining variation across types in take-up of contract C2. Note, however, that risk and cost preferences do seem to play a role in the take-up of Contract C1 by "naïve" agents (take-up rates for "naïve" agents rise about 20% when we set their risk and cost parameters equal to the estimated parameters).

We next evaluate the separate contribution of the hyperbolic parameters from that of the awareness of future present-bias by carrying out another set of counterfactual exercises. These were conducted using a model where all types have the same risk and cost preferences but the hyperbolic parameter for the "sophisticated" type is set equal to that estimated for the "naïve" type (Model 1 in Table 13). In effect, the only difference across inconsistent types is whether they recognize their future present-bias. In this case, the take-up and retreatment rates for the "naïve" agents are, as expected, the same as in Model 4. However, we find that the take-up rates for the "sophisticated" and "naïve" types are very close. This again suggests that it is the hyperbolic parameter for the "sophisticated" types (and its significant difference from those of the "naïve" and consistent types) that matters for adoption decisions. This is further confirmed by comparing the results for Models 1 and 2.

Further evidence of the role of the hyperbolic parameters can be found in the similarity of retreatment rates for "naïve" and time-consistent agents. In particular, it appears that since the hyperbolic parameter for the "naïve" agents is very close to one, their take-up decisions in the model are very similar to those of the time consistent agents. Finally, under the range of policy considerations considered here, the retreatment rates across types are extremely similar. This suggests that the main effect of time-inconsistent preferences in the model is working through the differential take-up rates across types rather than through differential retreatment rates.

To conclude, the results from the counterfactual exercise support the view that time consistent behavior is perhaps a reasonable approximation for the bulk of the population, with outcomes for only about 10% of the population (the "sophisticated" types) being affected by time-inconsistency issues.

8 Conclusions

Time inconsistency is often proposed as an explanation for observed intertemporal choice behavior, particularly in poor households. However, the relative occurrence of this phenomena (i.e the population prevalence of time-inconsistent agents) as well as its quantitative importance (i.e the extent of present-bias) are difficult to identify and estimate using standard methods and data. We develop a dynamic discrete choice model for time-inconsistent agents with unobserved types. We show identification for all key parameters — including separate hyperbolic parameters for different types and time- and type-varying per-period utilities — for a model based on a specifically collected dataset containing detailed information on beliefs combined with a field intervention. We also show how the identification results can be extended to a general state and action space.

We estimate that time-inconsistent agents account for more than half of the population and that "sophisticated" inconsistent agents are considerably more present-biased than their "naïve" counterparts. However, the hyperbolic parameter for "naïve" inconsistent agents — who we estimate comprise 40% of the population — is very close to one and, at least for the range of policy interventions considered, their behavior is very similar to that of the time-consistent agents. We also examined whether there are other differences across types (e.g. in risk and cost parameters) and found that these differences are small relative to the differences in time preferences. Finally, we can evaluate to what extent "sophisticated" agents are more likely to choose commitment products by comparing the type distribution among product purchasers to that of the general population. We find that commitment products are not particularly appealing to "sophisticated" agents and that the purchase of these products is in fact higher among wealthier (and even "naive") households.

References

Victor Aguirregabiria and Pedro Mira. Dynamic discrete choice structural models: A survey. *Journal of Econometrics*, 156(1):38–67, 2010. ISSN 0304-4076. URL http://linkinghub.elsevier.com/retrieve/pii/S0304407609001985. [3, 8, 26]

Syed N. Ali. Learning Self-Control. *Quarterly Journal of Economics*, Forthcoming, 2010. Available at http://www.dklevine.com/archive/refs4814577000000000384.pdf. [31]

S. Andersen, G.W. Harrison, M.I. Lau, and E.E. Rutström. Eliciting Risk and Time Preferences. *Econometrica*, 76 (3):583–618, 2008. ISSN 1468-0262. [4]

J. Andreoni and C. Sprenger. Estimating Time Preferences from Convex Budgets. Available at http://www.nber.org/papers/w16347.pdf, 2010. NBER Working Paper w16437. [2, 4, 7]

P. Arcidiacono and P.B. Ellickson. Practical Methods for Estimation of Dynamic Discrete Choice Models. *Annual Review of Economics*, 2011. Forthcoming. [3]

Nava Ashraf, Dean Karlan, and Wesley Yin. Tying Odysseus to the mast: evidence from a commitment savings product in the Philippines. *Quarterly Journal of Economics*, 121(2):635–672, 2006. [2, 4]

A.V. Banerjee and Sendhil Mullainathan. The Shape of Temptation: Implications for the Economic lives of the Poor. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1601724, 2010. NBER Working Paper w15973. [2, 3]

A.V. Banerjee, E. Duflo, R. Glennerster, and D. Kothari. Improving immunisation coverage in rural India: clustered randomised controlled evaluation of immunisation campaigns with and without incentives. *British Medical Journal*, 340(may17 1):c2220, 2010. [2]

S DellaVigna. Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2):315–372, 2009. ISSN 0022-0515. URL http://www.ingentaconnect.com/content/aea/jel/2009/00000047/00000002/art00001. [2]

Esther Duflo, Michael Kremer, and Jonathan Robinson. Nudging farmers to use fertilizer: Evidence from Kenya. *American Economic Review*, 2010. forthcoming. [2]

H. Fang and D. Silverman. Time-Inconsistency And Welfare Program Participation: Evidence From The NLSY. *International Economic Review*, 50(4):1043–1077, 2009. ISSN 1468-2354. [3]

Hanming Fang and Y. Wang. Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions. Available at http://www.econ.upenn.edu/~hfang/WorkingPaper/wang/hyperbolic6.pdf, 2010. NBER Working Paper 16438. [3, 15]

Shane Frederick, George Loewenstein, and Edward D. O'Donoghue. Time discounting: A critical review. *Journal of Economic Literature*, 40(2):351–401, 2002. [2]

Government of India. Poverty estimates for 2004-05, 2007. Available at http://planningcommission.nic.in/news/prmar07.pdf (accessed May 2008). [5]

Faruk Gul and Wolfgang Pesendorfer. Temptation and self-control. *Econometrica*, 69(6):1403–1435, 2001. [10]

Faruk Gul and Wolfgang Pesendorfer. Self-control and the theory of consumption. *Econometrica*, 72(1):119–158, 2004. [10]

Charles A Holt and Susan K Laury. Risk Aversion and Incentive Effects. *American Economic Review*, 92(5):1644–1655, December 2002. ISSN 0002-8282. doi: 10.1257/000282802762024700. URL http://www.atypon-link.com/AEAP/doi/abs/10.1257/000282802762024700. [4, 23, 49]

V.J. Hotz and R.A. Miller. Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3):497–529, 1993. URL http://www.jstor.org/stable/2298122. [43, 47]

Hiroyuki Kasahara and Katsumi Shimotsu. Nonparametric identification of finite mixture models of dynamic discrete choices. *Econometrica*, 77(1):135–175, 01 2009. URL http://ideas.repec.org/a/ecm/emetrp/v77y2009i1p135-175.html. [3, 4, 16, 20, 21, 28, 41]

Ashwani Kumar, Neena Valecha, Tanu Jain, and Aditya P. Dash. Burden of malaria in India: Retrospective and prospective view. *AJTMH*, 77(Suppl 6):69–78, 2007. [4]

David I. Laibson. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112(2):443–477, 1997. [2]

T. Magnac and D. Thesmar. Identifying Dynamic Discrete Decision Processes. *Econometrica*, 70(2):801–816, 2002. [2, 16]

A Moody. Rapid diagnostic tests for malaria parasites. *Clinical Microbiology Reviews*, 15(1):66–78, 2002. [6]

Sendhil Mullainathan. Development economics through the lens of psychology. Available at http://www.economics.harvard.edu/admin/person_pages/edit/html/files/LENS%20OF%20PSYCHOLOGY.pdf, 2004. [2]

Ted O'Donoghue and Matthew Rabin. Doing it now or later. *American Economic Review*, 89(1):103–124, 1999. [10]

MD Paserman. Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation. *The Economic Journal*, 118(531):1418–1452, August 2008. ISSN 00130133. doi: 10.1111/j.1468-0297.2008.02175.x. URL http://onlinelibrary.wiley.com/doi/10.1111/j.1468-0297.2008.02175.x/full. [3]

Ariel Rubinstein. "Economics and psychology"? The case of hyperbolic discounting. *International Economic Review*, 44(4):1207–1216, 2003. [7]

John Rust. Structural estimation of markov decision processes. In R. Engle and D. McFadden, editors, *Handbook of Econometrics*, pages 3081–3143. Elsevier Science, 1994. [2]

R. H. Strotz. Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23(3): 165–180, 1955. URL http://www.jstor.org/stable/2295722. [10]

Alessandro Tarozzi, Aprajit Mahajan, Joanne Yoong, and Brian Blackburn. Commitment mechanisms and compliance with health-protecting behavior: Preliminary evidence from Orissa, India. *American Economic Review (Papers and Proceedings)*, 99(2):231–235, 2009. [2, 11]

Alessandro Tarozzi, Aprajit Mahajan, Brian Blackburn, Daniel Kopf, Lakshmi Krishnan, and Joanne Yoong. Microloans, insecticide-treated nets and malaria: Evidence from a randomized controlled trial in orissa (india). Available at http://www.stanford.edu/people/amahajan/RCT.pdf, 2011. [2, 4]

W. van der Klaauw. On the Use of Expectations Data in Estimating Structural Dynamic Models: An Analysis of Career Choices. Working Paper, University of North Carolina at Chapel Hill, 2000. [3]

W. van der Klaauw and K.I. Wolpin. Social Security and the Retirement and Savings Behavior of Low-Income Households. *Journal of Econometrics*, 145(1-2):21–42, 2008. [3]

World Bank. Global purchasing power parities and real expenditures, 2005. Technical report, International Comparison Program, Washington DC: World Bank, 2008. Available at www.worldbank.org/data/icp. [51]

Appendix A

Proof of Lemma 1

Proof. Under the independence assumption (Assumption 2) and the additive separability assumption (Assumption 3) and the direct observation of types (Assumption 4), the probability that an agent of type τ who has purchased a net will re-treat the net is given by

$$\mathbb{P}_\tau(a_3^* = 1|x_3, z) = G_\Delta(u_3(x_3, 1, \tau) - u_3(x_3, 0, \tau) + \beta_\tau \delta \int u(x_4, \tau)(\mathrm{dF}(x_4|x_3, 1, z) - \mathrm{dF}(x_4|x_3, 0, z)) \tag{21}$$

G_Δ is the distribution of $\epsilon_0 - \epsilon_1$ which is known and has support over the real line by assumption. We can invert this function to obtain

$$g_{\tau,3}(x_3, z) = u_3(x_3, 1, \tau) - u_3(x_3, 0, \tau) + \beta_\tau \delta \int u(x_4, \tau)(\mathrm{dF}(x_4|x_3, 1, z) - \mathrm{dF}(x_4|x_3, 0, z)) \tag{22}$$

The left hand side is directly identified, and by Assumption 5 the integral on the right hand side simplifies to

$$\gamma((u_4(bh, \tau) - u_4(bm, \tau))\mathbb{I}_{(x_3 \in \mathcal{X}_B)} + (u_4(ch, \tau) - u_4(cm, \tau))\mathbb{I}_{(x_3 \in \mathcal{X}_C)})$$

Under Assumption 6 we can evaluate (22) at two different values of γ and take differences to obtain

$$\begin{aligned}&g_{\tau,3}(x_3, z \backslash \gamma_1, \gamma_1) - g_{\tau,3}(x_3, z \backslash \gamma_2, \gamma_2) = \\ &(\gamma_1 - \gamma_2)\beta_\tau \delta((u_4(bh, \tau) - u_4(bm, \tau))\mathbb{I}_{(x_3 \in \mathcal{X}_B)} + (u_4(ch, \tau) - u_4(cm, \tau))\mathbb{I}_{(x_3 \in \mathcal{X}_c)})\end{aligned}$$

so that the object on the right hand side below is identified.

$$\frac{(g_{\tau,3}(x_3, z \backslash \gamma_1, \gamma_1) - g(x_3, z \backslash \gamma_2, \gamma_2))}{\gamma_1 - \gamma_2} = \beta_\tau \delta \Big((u_4(bh, \tau) - u_4(bm, \tau))\mathbb{I}_{(x_3 \in \mathcal{X}_B)} + (u_4(ch, \tau) - u_4(cm, \tau))\mathbb{I}_{(x_3 \in \mathcal{X}_C)} \Big)$$

Therefore, we can identify the utility differential as

$$u_3(x_3, 1, \tau) - u_3(x_3, 0, \tau) = g_{\tau,3}(x_3, z, \gamma) - \frac{(g_{\tau,3}(x_3, z \backslash \gamma_1, \gamma_1) - g_{\tau,3}(x_3, z \backslash \gamma_2, \gamma_2))}{\gamma_1 - \gamma_2}\gamma. \tag{23}$$

□

Proof of Lemma 2

Proof. The proof is straightforward by comparing the identified quantities $\beta_{\tau_N} \delta \int u_4(x_4, \tau_N)\, \mathrm{dF}_\Delta(x_4|x_3, z)$ and $\delta \int u_4(x_4, \tau_C)\, \mathrm{dF}_\Delta(x_4|x_3, z)$ (where we have defined the finite signed measure $\mathrm{dF}_\Delta(x_4|x_3, z) \equiv \mathrm{dF}(x_4|x_3, 1, z) - \mathrm{dF}(x_4|x_3, 0, z)$), and noting that $\int u_4(x_4, \tau_N)\, \mathrm{dF}_\Delta(x_4|x_3, z) = \int u_4(x_4, \tau_C)\, \mathrm{dF}_\Delta(x_4|x_3, z)$ by Assumption 7. □

Proof of Lemma 3

Proof. Under the assumptions in the lemma, the probability that an agent of type τ who has purchased a net will choose to re-treat it is

$$\mathbb{P}_\tau(a_2^* = 1|x_2, z) = G_\Delta\Big(u_2(x_2, 1, \tau) - u_2(x_2, 0, \tau) + \beta_\tau \delta \textstyle\int v_\tau^*(s_3, z)\, \mathrm{dF}_\Delta(s_3|x_2, z) \Big), \tag{24}$$

where we have defined the finite signed measure $\mathrm{dF}_\Delta(s_3|x_2, z) \equiv \mathrm{dF}(s_3|x_2, 1, z) - \mathrm{dF}(s_3|s_2, 0, z)$. The value function $v_\tau^*(\cdot)$ is defined as

$$v_\tau^*(s_3, z) \equiv v_\tau(s_3, d^\tau(s_3, z), z), \tag{25}$$

where

$$v_\tau(s_3, a, z) \equiv u_3(x_3, a, \tau) + \epsilon_3(a) + \delta \int u(x_4, \tau)\, \mathrm{dF}(x_4|x_3, a, z) \tag{26}$$

and

$$d^{\tau}(s_3,z)=\mathbb{I}\Big\{u_3(x_3,1,\tau)+\epsilon_3(1)+\tilde{\beta}_{\tau}\delta\int u(x_4,\tau)\,\mathrm{dF}(x_4|x_3,1,z)\geq$$
$$u_3(x_3,0,\tau)+\epsilon_3(0)+\tilde{\beta}_{\tau}\delta\int u(x_4,\tau)\,\mathrm{dF}(x_4|x_3,0,z)\Big\}$$

so that an agent of type τ will evaluate his continuation utility from period 3 onwards using the decision rule $d^{\tau}(s_3,z)$. $\tilde{\beta}_{\tau}$ is the present bias that an agent (of type τ) in period 2 thinks she will have in period 3. A time consistent agent knows that he is not subject to present-bias and so $\tilde{\beta}_{\tau_C}=\beta_{\tau_C}=1$. "Naïve" inconsistent agents assume that their future selves are not subject to present-bias and will therefore set $\tilde{\beta}_{\tau_N}=1$. Conversely, "sophisticated" agents recognize that their future selves will be present biased and we assume that such agents will set $\tilde{\beta}_{\tau_S}=\beta_{\tau_S}$.[38]

To summarize, "naïve" agents will assume (in period 2) that their period 3 decision rule is unaffected by present-bias while "sophisticated" agents will recognize that it is and this will be reflected in the decision rule they think will be used in period 3.

Next, we examine $\int v^*_{\tau}(s_3,z)\,\mathrm{dF}_{\Delta}(s_3|x_2,z)$. Define the function

$$\begin{aligned}\tilde{g}_{\tau,3}(x_3,z,c)&\equiv u_3(x_3,1,\tau)-u_3(x_3,0,\tau)+c\delta\int u_4(x_4,\tau)\,\mathrm{dF}_{\Delta}(x_4|x_3,z)\\&=u_3(x_3,1,\tau)-u_3(x_3,0,\tau)+c\delta\gamma\Big((u_4(bh,\tau)-u_4(bm,\tau))\mathbb{I}_{(x_3\in B)}+(u_4(ch,\tau)-u_4(cm,\tau))\mathbb{I}_{(x_3\in C)}\Big)\end{aligned}$$

where the second equality follows from Assumption 5. We can rewrite

$$d^{\tau}(s_3,z)=\mathbb{I}\Big(\tilde{g}_{\tau,3}(x_3,z,\tilde{\beta}_{\tau})\geq\epsilon_3(0)-\epsilon_3(1)\Big)$$

and

$$\begin{aligned}v^*_{\tau}(s_3,z)=\;&\mathbb{I}\Big(\tilde{g}_{\tau,3}(x_3,z,\tilde{\beta}_{\tau})\geq\epsilon_3(0)-\epsilon_3(1)\Big)\Big(\tilde{g}_{\tau,3}(x_3,z,1)+\epsilon_3(1)-\epsilon_3(0)\Big)\\&+u_3(x_3,0,\tau)+\epsilon_3(0)+\delta\int u(x_4,\tau)\,\mathrm{dF}(x_4|x_3,0,z)\end{aligned}\tag{27}$$

so that we can write

$$\begin{aligned}\int v^*_{\tau}(s_3,z)\,\mathrm{dF}_{\Delta}(s_3|x_2,z)=&\iint^{\tilde{g}_{\tau,3}(x_3,z,\tilde{\beta}_{\tau})}\Big(\tilde{g}_{\tau,3}(x_3,z,1)-\Delta\epsilon\Big)\,\mathrm{dG}(\Delta\epsilon)\,\mathrm{dF}_{\Delta}(x_3|x_2,z)+\\&\int u_3(x_3,0,\tau)\,\mathrm{dF}_{\Delta}(x_3|x_2,z)+\delta\iint u(x_4,\tau)\,\mathrm{dF}(x_4|x_3,0,z)\,\mathrm{dF}_{\Delta}(x_3|x_2,z)\\=&\int G(\tilde{g}_{\tau,3}(x_3,z,\tilde{\beta}_{\tau}))\tilde{g}_{\tau,3}(x_3,z,1)\,\mathrm{dF}_{\Delta}(x_3|x_2,z)-\int^{\tilde{g}_{\tau,3}(x_3,z,\tilde{\beta}_{\tau})}\Delta\epsilon\quad\mathrm{dG}(\Delta\epsilon)\,\mathrm{dF}_{\Delta}(x_3|x_2,z)\\&+\int u_3(x_3,0,\tau)\,\mathrm{dF}_{\Delta}(x_3|x_2,z)\end{aligned}\tag{28}$$

where we ignore the last two terms in (27) since they integrate to zero. Next, note that the last term in the preceding is identified by the normalization (Assumption 9).

TIME-CONSISTENT AGENTS

We first consider time consistent agents where $\beta_{\tau_C}=\tilde{\beta}_{\tau_C}=1$ and the function $\tilde{g}_{\tau,3}(x_3,z,1)$ is identified since it is equal to the identified function $g_{\tau_C,3}(x_3,z)$. Therefore, for this type the function $\int v^*_{\tau}(s_3,z)\,\mathrm{dF}_{\Delta}(s_3|x_2,z)$ is

[38] We do not consider "partially sophisticated" agents that is to say agents for whom $\tilde{\beta}_{\tau_S}\neq\beta_{\tau_S}$ but also not equal to 1. Identification in this case would be considerably more difficult without further information.

identified. Returning now to the identified function for period 2 for this type,

$$g_{\tau_C,2}(x_2,z) = u_2(x_2,1,\tau_C) - u_2(x_2,0,\tau_C) + \delta \int v^*_{\tau_C}(s_3,z)\,\mathrm{dF}_\Delta(s_3|x_2,z). \tag{29}$$

Then, Assumption 8 guarantees that

$$\int v^*_{\tau_C,3}(s_3,z)\,\mathrm{dF}_\Delta(s_3|x_2,z\backslash\gamma,\gamma_2) \neq \int v^*_{\tau_C,3}(s_3,z)\,\mathrm{dF}_\Delta(s_3|x_2,z\backslash\gamma,\gamma_1) \tag{30}$$

Therefore, evaluating (29) at two different points of γ and differencing will identify the exponential discount parameter δ. Next, using δ we can identify the second period utility differentials $u_2(x_2,1,\tau_C) - u_2(x_2,0,\tau_C)$ for all $x_2 \in \mathcal{X}_B$.

"SOPHISTICATED" TIME-INCONSISTENT AGENTS

Next, consider the time-inconsistent "sophisticated" agents, i.e. those for whom $\tilde{\beta}_{\tau_S} = \beta_{\tau_S}$. For this type the function $\tilde{g}_{\tau,3}(x_3,z,\tilde{\beta}_{\tau_S})$ is identified as $g_{\tau_S,3}(x_3,z)$. This implies that the second term on the right-hand side in the expression below is identified (note that the last term is identified by the normalization).

$$\begin{aligned}\int v^*_\tau(s_3,z)\,\mathrm{dF}_\Delta(s_3|x_2,z) &= \int G(g_{\tau_S,3}(x_3,z))\tilde{g}_{\tau_S,3}(x_3,z,1)\,\mathrm{dF}_\Delta(x_3|x_2,z) - \int^{\tilde{g}_{\tau_S,3}(x_3,z)} \Delta\epsilon \quad \mathrm{dG}(\Delta\epsilon)\,\mathrm{dF}_\Delta(x_3|x_2,z)\\ &+ \int u_3(x_3,0,\tau)\,\mathrm{dF}_\Delta(x_3|x_2,z).\end{aligned}$$

Focusing therefore on the first term, we can rewrite it as

$$\begin{aligned}\int G(g_{\tau_S,3}(x_3,z)\tilde{g}_{\tau_S,3}(x_3,z,1)\,\mathrm{dF}_\Delta(x_3|x_2,z) &= \gamma\Big(G(g_{\tau_S,3}(ch,z))\tilde{g}_{\tau_S,3}(ch,z,1) - G(g_{\tau_S,3}(cm,z))\tilde{g}_{\tau_S,3}(cm,z,1)\Big)\\ &= \gamma\Big(G(g_{\tau_S,3}(ch,z))\Big(u_3(ch,1,\tau_S) - u_3(ch,0,\tau_S) + \delta\int u(x_4,\tau_S)\,\mathrm{dF}_\Delta(x_4|ch)\Big)\Big)\\ &+ \gamma\Big(G(g_{\tau_S,3}(cm,z))\Big(u_3(cm,1,\tau_S) - u_3(cm,0,\tau_S) + \delta\int u(x_4,\tau_S)\,\mathrm{dF}_\Delta(x_4|cm)\Big)\Big)\end{aligned}$$

Define the identified functions $\mathbf{H}(\cdot)$ and $\mathbf{Q}(\cdot)$ as

$$\begin{aligned}\mathbf{H}(x_2,z) = &\int u_3(x_3,0,\tau_S)\,\mathrm{dF}_\Delta(x_3|x_2,z)\\ &+ \int^{g(x_3,z,\tau_S)} \Delta\epsilon\,\mathrm{dG}(\Delta\epsilon)\,\mathrm{dF}_\Delta(x_3|x_2,z)\\ &+ \gamma G(g_{\tau_S,3}(ch,z))(u_3(ch,1,\tau_S) - u_3(ch,0,\tau_S))\\ &+ \gamma G(g_{\tau_S,3}(cm,z))(u_3(cm,1,\tau_S) - u_3(cm,0,\tau_S)).\end{aligned} \tag{31}$$

$$\mathbf{Q}(z) = \delta G(g_{\tau_S,3}(ch,z))\beta_{\tau_S}\delta\int u(x_4,\tau_S)\,\mathrm{dF}_\Delta(x_4|ch,z) - \delta G(g_{\tau_S,3}(cm,z))\beta_{\tau_S}\delta\int u(x_4,\tau_S)\,\mathrm{dF}_\Delta(x_4|cm,z)$$

In particular, $\mathbf{Q}(z)$ is identified since $\beta_{\tau_S}\delta\int u_4(x_4,\tau_S)\,\mathrm{dF}_\Delta(x_4|x_3,z)$ is identified from Lemma 1, δ is identified from the results above and $G(g_{\tau_S,3}(ch,z))$ is identified since $G(\cdot)$ is known and $g_{\tau_S,3}(ch,z)$ is identified. We can then write

$$\beta_{\tau_S}\delta\int v^*_{\tau_S,3}(s_3,z)\,\mathrm{dF}_\Delta(s_3|x_2,z) = \beta_{\tau_S}\delta\mathbf{H}(x_2,z) + \mathbf{Q}(z).$$

The identified function for the "sophisticated" types is then

$$g_{\tau_S,2}(x_2,z) = u_2(x_2,1,\tau_S) - u_2(x_2,0,\tau_S) + \beta_{\tau_S}\delta\mathbf{H}(x_2,z) + \mathbf{Q}(z). \tag{32}$$

Therefore, under Assumption 8, the hyperbolic parameter β_{τ_S} is identified as are the utility differentials for period 2: $u_2(x_2,1,\tau_S) - u_2(x_2,0,\tau_S)$ for $x_2 \in \mathcal{X}_C$.

INCONSISTENT "NAÏVE" AGENTS

Finally, consider agents for whom $\tilde{\beta}_{\tau_N} = 1$. For these agents, we identify the function

$$g_{\tau_N,2}(x_2, z) = u_2(x_2, 1, \tau_N) - u_2(x_2, 0, \tau_N) + \beta_{\tau_N}\delta \int v^*_{\tau_N}(s_3, z)\,\mathrm{dF}_\Delta(s_3|x_2, z). \tag{33}$$

Next, the function $\tilde{g}_{\tau_N,s}(x_3, z, 1)$ is identified since (i) the utility differentials $u_3(x_3, 1, \tau_N) - u_3(x_3, 0, \tau_N)$ are identified by Lemma 1, (ii) δ is identified by the first part of the proof above and (iii) $\int u_4(x_4, \tau_N)\,\mathrm{dF}_\Delta(x_4|x_3, z)$ is identified by Assumption 7. Looking at equation (8) we see that $\int v^*_{\tau_N}(s_3, z)\,\mathrm{dF}_\Delta(s_3|x_2, z)$ is therefore identified (given the normalization assumption). Since β_{τ_N} is identified by Lemma 2 and δ is identified above we conclude that the last term in the display (33) is also identified. Therefore, the utility differential $u_2(x_2, 1, \tau_N) - u_2(x_2, 0, \tau_N)$ is identified for all $x_2 \in \mathcal{X}_B$. □

Proof of Lemma 4

Proof. We first prove that the type-specific choice probabilities are identified. We then use the inversion argument as before (see Appendix C for the direct argument) to recover preference parameters. First, note that by the assumption of directly observed types, $\mathbb{P}_{\tau_C}(a_1|x_1) = \mathbb{P}(a_1|x_1, r = 0)$ so that the type-specific choice probabilities for time consistent agents are identified trivially. Similarly, since we assume that "sophisticated" types only choose between $\{n, c\}$ and "naïve" types only choose between $\{n, b\}$, $\mathbb{P}_{\tau_S}(a_1 = b|x_1) = \mathbb{P}_{\tau_N}(a_1 = c|x_1) = 0$.

In order to identify the type-specific choice probabilities for "naïve" and "sophisticated" agents separately we begin by considering the directly identified matrix $\mathsf{E} \equiv \mathsf{P}^b_{1,1}\mathsf{L}_2^{-1}$ for given values of (x_1, x_1', x_2, x_2') for $x_t \neq x_t'$. Note that under the assumptions stated, $\mathsf{E} = (\mathsf{L}^a_1)'\mathsf{V}_1$, which is given by

$$(\mathsf{L}^b_1)'\mathsf{V}_1 = \begin{pmatrix} 0 & \pi_{\tau_N}(1) & \pi_{\tau_S}(1) \\ 0 & \pi_{\tau_N}(1)\mathbb{P}_{\tau_N}(a_1 = b, x_1) & \pi_{\tau_S}(1)\mathbb{P}_{\tau_N}(a_1 = b, x_1') \\ 0 & \pi_{\tau_N}(1)\mathbb{P}_{\tau_S}(a_1 = b, x_1) & \pi_{\tau_S}(1)\mathbb{P}_{\tau_S}(a_1 = b, x_1') \end{pmatrix}$$

so that the type probabilities (π_τ) and $\mathbb{P}_{\tau_N}(a_1 = b|x_1)$ are identified. In addition, the type-specific distribution for the initial state $\mathbb{P}_{\tau_N}(x_1)$ is identified (note that by assumption $\mathbb{P}_{\tau_S}(a_1 = b|x_1) = 0$) . Following the same argument but for the choice of the commitment contract yields identification of $\mathbb{P}_{\tau_S}(a_1 = c|x_1)$ and $\mathbb{P}_{\tau_S}(x_1)$. Therefore, all the type-specific choice probabilities are identified. Next, using the same inversion argument as before we can invert the choice probability for type τ to identify the functions

$$g_{\tau,1,b}(x_1, z) = u(x_1, b, \tau) - u(x_1, n, \tau) + \beta_\tau\delta \int v^*_\tau(s_2, z)\,\mathrm{dF}_{\Delta,b}(s_2|x_1, z) \tag{34}$$

$$g_{\tau,1,c}(x_1, z) = u(x_1, c, \tau) - u(x_1, n, \tau) + \beta_\tau\delta \int v^*_\tau(s_2, z)\,\mathrm{dF}_{\Delta,c}(s_2|x_1, z) \tag{35}$$

where the signed measure is $\mathrm{dF}_{\Delta,a}(s_2|s_1, z) = \mathrm{dF}(s_2|x_1, a_1 = a, z) - \mathrm{dF}(s_2|x_1, a_1 = n, z)$. Define the period 2 type-specific value function as

$$v^*_\tau(s_2, z) \equiv v_\tau(s_2, d^\tau(s_2, z), z) \tag{36}$$

where

$$v_\tau(s_2, a, z) \equiv u_2(x_2, a, \tau) + \epsilon_2(a) + \delta \int v^*_\tau(s_3, z)\,\mathrm{dF}(s_3|x_2, a, z).$$

The decision rule for an agent of type τ in period 2 is

$$\begin{aligned} d^\tau(s_2, z) = \mathbb{I}\Big\{ & u_2(x_2, 1, \tau) + \epsilon_2(1) + \tilde{\beta}_\tau\delta \int v^*_\tau(s_3, z)\,\mathrm{dF}(s_3|x_2, 1, z) \geq \\ & u_2(x_2, 0, \tau) + \epsilon_2(0) + \tilde{\beta}_\tau\delta \int v^*_\tau(s_3, z)\,\mathrm{dF}(s_3|x_2, 0, z)\Big\}. \end{aligned}$$

As before, $\tilde{\beta}_\tau$ is the present-bias that agent in period 1 thinks she will have in period 2. A time consistent agent knows that he is not subject to present-bias and so $\tilde{\beta}_{\tau_C} = \beta_{\tau_C} = 1$. "Naïve" inconsistent agents assume that their future selves are not subject to present-bias and will therefore set $\tilde{\beta}_{\tau_N} = 1$. Conversely, "sophisticated" agents recognize that their future selves will be present biased and we assume that such agents will set $\tilde{\beta}_{\tau_S} = \beta_{\tau_S}$.

From the previous results the last term in each of the equations (34 and 35) is identified so that using the variation in beliefs assured by Assumption 8 we conclude that the utility differentials $u(x_1, b, \tau) - u(x_1, n, \tau)$ and $u(x_1, c, \tau) - u(x_1, n, \tau)$ are identified for all initial states x_1 and for all types τ. □

Proof of Lemma 5

Proof. The idea of the proof is based on Lemma 4 of Kasahara and Shimotsu (2009). The difference is that we use an exclusion restriction — that conditional on the state,[39] $r_u \equiv (r, a_1)$ only affect type probabilities — to generate identification instead of placing restrictions on the number of periods required.

We first simplify the joint probability of $(a_t, a_{t+1}, x_t, x_{t+1})$ (for $t > 1$) and suppress the conditioning on z throughout for convenience.[40]

$$
\begin{aligned}
&\mathbb{P}(a_t = 1, a_{t+1} = 1, x_t, x_{t+1}|r_u) \\
&= \sum_{\tau \in \mathcal{T}} \pi_\tau(r_u)\mathbb{P}_\tau(a_{t+1} = 1, a_t = 1, x_{t+1}, x_t|r_u) \\
&= \sum_{\tau} \pi_\tau(r_u)\mathbb{P}_\tau(a_{t+1} = 1|a_t = 1, x_{t+1}, x_t, r_u)\mathbb{P}_\tau(x_{t+1}|a_t = 1, x_t, r_u)\mathbb{P}_\tau(a_t = 1, x_t)
\end{aligned}
$$

We use the fact that the optimal action in any period only depends upon previous actions and states through the current state (see e.g. equation (21)), Assumption 13 and that $\mathbb{P}_\tau(a_t = 1, x_t|r_u) = \mathbb{P}_\tau(a_t = 1, x_t)$. Next, since by assumption $\mathbb{P}(x_{t+1}|x_t, a_t = 1, z)$ does not vary by type, we can take them on the left hand side (under the assumption that these are always strictly positive) and define

$$\frac{\mathbb{P}(a_t = 1, a_{t+1} = 1, x_t = x_a, x_{t+1} = x_b|r_u)}{\mathbb{P}(x_{t+1} = x_b|x_t = x_a, a_t = 1)} \equiv \mathbf{F}^{t,t+1}_{x_a,x_b,r_u}. \tag{37}$$

In what follows, we assume that there are M total types to highlight the generality of the argument.

$$
\begin{aligned}
\mathbf{F}^{t,t+1}_{x_a,x_b,r_u} &= \sum_{\tau=1}^{M} \pi_\tau(r_u)\mathbb{P}_\tau(a_{t+1} = 1|x_{t+1} = x_b, z)\mathbb{P}_\tau(a_t = 1, x_t = x_a) \\
\mathbf{F}^{t}_{x_a,r_u} &= \sum_{\tau=1}^{M} \pi_\tau(r_u)\mathbb{P}_\tau(a_t = 1, x_t = x_a) \\
\mathbf{F}^{t+1}_{x_b,r_u} &= \sum_{\tau=1}^{M} \pi_\tau(r_u)\mathbb{P}_\tau(a_{t+1} = 1|x_{t+1} = x_b)
\end{aligned}
$$

Define the $M \times M$ matrix

$$\mathsf{P}_{t,r_u} = \begin{pmatrix} 1 & \mathbf{F}^{t+1}_{x_a,\tau} & \cdots & \mathbf{F}^{t+1}_{x_{M-1},r_u} \\ \vdots & \mathbf{F}^{t,t+1}_{x_a,x_a,r_u} & \cdots & \mathbf{F}^{t,t+1}_{x_a,x_{M-1},r_u} \\ \vdots & \vdots & \cdots & \\ \mathbf{F}^{t}_{x_{M-1},r_u} & \mathbf{F}^{t,t+1}_{x_{M-1},x_a,r_u} & \cdots & \mathbf{F}^{t,t+1}_{x_{M-1},x_{M-1},r_u} \end{pmatrix} \tag{38}$$

Next, define $\mathbb{P}_\tau(a_{t+1} = 1|x_{t+1} = x_b) \equiv \lambda^\tau_{t+1,x_b}$ and $\mathbb{P}_\tau(a_{t+1} = 1, x_t = x_a) \equiv \lambda^\tau_{t,x_a}$. Define the $M \times M$ matrix for t and $t+1$

$$\mathsf{L}_t = \begin{pmatrix} 1 & \lambda^1_{t,x_a} & \lambda^1_{t,x_{M-1}} \\ \vdots & \vdots & \vdots \\ 1 & \lambda^M_{t,x_a} & \lambda^M_{t,x_{M-1}} \end{pmatrix} \tag{39}$$

Finally, define the $M \times M$ matrix $\mathsf{V}_{r_u} = \text{diag}(\pi_{\tau_1}(r_u)...\pi_{\tau_M}(r_u))$. It is easy to show the following factorization holds.

$$\mathsf{P}_{t,r_u} = \mathsf{L}'_t\mathsf{V}_{r_u}\mathsf{L}_{t+1} \tag{40}$$

[39] Defined to include first period choice.

[40] Therefore, all identified quantities, including the type probabilities can vary by observables z which is potentially important.

and by assumption each term on the right hand side is invertible. Next, consider the directly identified object A defined by

$$\mathsf{A} \equiv \mathsf{P}_{t,r_u}^{-1}\mathsf{P}_{t,r_u'} = \mathsf{L}_{t+1}^{-1}\mathsf{V}_{r_u}^{-1}\mathsf{V}_{r_u'}\mathsf{L}_{t+1} \tag{41}$$

so that

$$\mathsf{L}_{t+1}\mathsf{A} = \hat{V}_{r_u,r_u'}\mathsf{L}_{t+1}$$

where $\hat{\mathsf{V}}_{r_u,r_u'} \equiv \mathsf{V}_{r_u}^{-1}\mathsf{V}_{r_u'}$ is a diagonal matrix. The expression above asserts that the diagonal matrix $\hat{\mathsf{V}}_{r_u,r_u'}$ contains the eigenvalues of A and that the rows of L_{t+1} comprise its left eigenvectors. Therefore, these objects are identified by carrying out an eigenvalue decomposition of the identified matrix A. Note that the eigenvectors are only identified up to scale, so that we can identify the matrix $\mathsf{E} \equiv \mathsf{D}\mathsf{L}_{t+1}$ where D is a diagonal matrix (and we have $\mathsf{L}_{t+1} = \mathsf{D}^{-1}\mathsf{E}$) Next,

$$\mathsf{P}_{t,r_u}\mathsf{E}^{-1} = \mathsf{L}_t'\mathsf{V}_{r_u}\mathsf{D}^{-1}.$$

Since the first row of L_t' consists of ones, the first row of the identified matrix $\mathsf{P}_{t,r_u}\mathsf{E}^{-1}$ identifies the elements of the diagonal matrix $\mathsf{V}_{r_u}\mathsf{D}^{-1}$. Define $\mathsf{F} \equiv \mathsf{V}_{r_u}\mathsf{D}^{-1}$ to be the identified matrix from this analysis. Next, note that

$$\mathsf{L}_t' = \mathsf{P}_{t,r_u}\mathsf{L}_{t+1}^{-1}\mathsf{V}_{r_u}^{-1} = \mathsf{P}_{t,r_u}\mathsf{E}^{-1}\mathsf{D}\mathsf{V}_{r_u}^{-1} = \mathsf{P}_{t,r_u}\mathsf{E}^{-1}\mathsf{F}^{-1}$$

where all the terms on the right hand side are identified, so that L_t is identified.
Next,

$$\mathsf{V}_{r_u}\mathsf{L}_{t+1} = (\mathsf{L}_t')^{-1}\mathsf{P}_{t,r_u} = \mathsf{P}_{t,r_u}\mathsf{E}^{-1}\mathsf{F}^{-1}\mathsf{P}_{t,r_u}. \tag{42}$$

The first column on the left hand side consists of the diagonal elements of the matrix V_{r_u}. Therefore V_{r_u} is identified since all the matrices on the right hand side in (42) are identified. Denote by $\mathsf{G} \equiv \mathsf{V}_{r_u}$ the diagonal matrix obtained by this argument. Then,

$$\mathsf{L}_{t+1} = \mathsf{G}^{-1}\mathsf{P}_{t,r_u}\mathsf{E}^{-1}\mathsf{F}^{-1}\mathsf{P}_{t,r_u}, \tag{43}$$

where the matrix G is invertible since by assumption all its diagonal entries are non-zero. Finally, note that since V_{r_u} is identified, then $\mathsf{V}_{r_u'} = \mathsf{G}\hat{\mathsf{V}}_{r_u,r_u'}$ and so $\mathsf{V}_{r_u'}$ is also identified since both G and $\hat{\mathsf{V}}$ are identified. We can apply this result for $(t, t+1) = (2, 3)$ to identify the type-specific choice probabilities for these periods. We can then use Lemma 4 to recover the type-specific choice probabilities for period 1 as well as the type probabilities $\pi_\tau(r)$ (i.e. the type probabilities only conditioning on r). To see this, consider the matrices $\mathsf{P}_{1,r}^a$ (defined in (7)), $\mathsf{V}_r = \text{diag}(\pi_{\tau_C}(r), \pi_{\tau_N}(r), \pi_{\tau_S}(r))$ and L_1^a (defined in (8)) and the relationship $\mathsf{P}_{1,r}^a = (\mathsf{L}_1^a)'\mathsf{V}_r\mathsf{L}_2$. Since L_2 is identified from the previous arguments, we can identify $\mathsf{L}_2^{-1}\mathsf{P}_{1,r} \equiv \mathsf{E}$. Next, note that the first row of E contains the elements of the matrix V_r. Therefore, the matrix L_1^a is identified as $\mathsf{V}_r^{-1}\mathsf{E}$ for $a \in \{b, c\}$ and therefore L_1^n is also identified. Finally, since $\{\pi_\tau(r, a)\}_{\tau\in\mathcal{T},a\in\{b,c\}}$ and $\{\pi_\tau(r)\}_{\tau\in\mathcal{T}}$ are both identified, and since

$$\pi_\tau(r) = \sum_{a_1\in\mathcal{A}_1} \pi_\tau(r, a_1)\mathbb{P}(a_1|r),$$

we can recover $\pi_\tau(r, n)$ as well for $\tau \in \mathcal{T}$. □

Appendix B: Results for General Discrete State and Action Space

In this section, we consider the case where the action space $\mathcal{A}_t$ has cardinality K (the results extend easily to the case where the cardinality varies by time) so the agent can take one of K actions in each period, $a_t \in \{0, ...K-1\}$. In addition, we assume that the state space has general finite discrete support. The identification results will require a precise relationship between the support of the belief vector z and the number of points in the state space support and these are explicated below. In the interest of exposition, we consider identifying variation in the belief vector z directly (rather than the γ component of it as earlier) and assume z has finite support. This assumption is not necessary and can be modified straightforwardly if this is not the case. Finally, we suppress dependence upon v which is not used for identification.

We show identification for the case where the type-specific choice probabilities are identified. In the case of unobserved types, an analogue of Lemma 5 can be applied to yield the type-specific choice probabilities in the first step and the results below applied subsequently.

Final Period and Penultimate Period Identification

First, as before we use standards results (see Hotz and Miller, 1993 or the direct argument in Appendix C) to justify the inversion argument from the choice probabilities to utility differentials. From this argument we conclude that there exists a $K-1$ vector of directly identified functions $\mathbf{g}_{\tau,\mathbf{3}}(x_3, z)$ such that

$$\mathbf{g}_{\tau,\mathbf{3}}(x_3, z) = \begin{bmatrix} \tilde{u}_{\tau,1} \\ \vdots \\ \tilde{u}_{\tau,K-1} \end{bmatrix} \tag{44}$$

where

$$\tilde{u}_{\tau,k} \equiv u(x_3, k, \tau) - u_3(x_3, 0, \tau) + \beta_\tau \delta \int u(x_4, \tau)\, \mathrm{dF}_{\Delta,k}(x_4|x_3, z)$$

and we define the signed measure

$$\mathrm{dF}_{\Delta,k}(x_4|x_3, z) \equiv \mathrm{dF}(x_4|x_3, k, z) - \mathrm{dF}(x_4|x_3, 0, z)$$

The identification argument exploits the restrictions implied by equation (44). It will be useful to take differences of $g(\cdot)$ evaluated at two distinct points in the support of z. This eliminates the third period utility differentials and allows us to first focus on identifying the utility function in the final period.

$$\begin{aligned} &\mathbf{g}_{\tau,\mathbf{3}}(x_3, z_s) - \mathbf{g}_{\tau,\mathbf{3}}(x_3, z_1) \\ &= \beta_\tau \delta \begin{bmatrix} \int u(x_4, \tau)\, \mathrm{dF}^{\Delta,s}_{\Delta,1}(x_4) \\ \vdots \\ \int u(x_4, \tau)\, \mathrm{dF}^{\Delta,s}_{\Delta,K-1}(x_4) \end{bmatrix} \end{aligned} \tag{45}$$

where the signed measure $\mathrm{dF}^{\Delta,s}_{\Delta,k}(x_4)$ is defined as

$$\mathrm{dF}^{\Delta,s}_{\Delta,k}(x_4) \equiv \mathrm{dF}_{\Delta,k}(x_4|x_3, z_s) - \mathrm{dF}_{\Delta,k}(x_4|x_3, z_1)$$

Define $\Delta\mathbf{g}_{\tau,\mathbf{3}}(z_s) \equiv \mathbf{g}_{\tau,\mathbf{3}}(x_3, z_s) - \mathbf{g}_{\tau,\mathbf{3}}(x_3, z_1)$.

One can pursue at least two distinct identification strategies from this point on. We discuss each in turn. First, one direct route (that does not depend upon the cardinality of the action space) places strong assumptions on the variation in beliefs to identify final period utility. In the context of our application, this could either be variation in beliefs about future income or about malarial incidence. In this first approach, we require that there be as many points of support in the belief distribution as there are observable states in the final period (which we denote by $\#\mathcal{X}_4$).

ASSUMPTION 16. *The distribution of the belief vector z conditional on x_3 has at least $\#\mathcal{X}_4$ points of support.*

Consider, for a given z_s, the k^{th} element $\Delta g_{\tau,3,k}(z_s)$ of $\Delta\mathbf{g}_{\tau,\mathbf{3}}(z_s)$ and form the vector (of dimension $(\#\mathcal{X}_4 - 1) \times 1$) $\bar{\mathbf{\Delta}}\mathbf{g}_{\tau,\mathbf{3},\mathbf{k}} = \Big(\Delta g_{\tau,3,k}(z_1), \ldots, \Delta g_{\tau,3,k}(z_{\#\mathcal{X}_4-1})\Big)'$. Since by assumption the state space is discrete and since $\sum_{r=1}^{\#\mathcal{X}_4} \mathrm{dF}^{\Delta,s}_{\Delta,k}(x_{4,r}) = 0$, we can rewrite this vector as

$$\bar{\mathbf{\Delta}}\mathbf{g}_{\tau,\mathbf{3},\mathbf{k}} = \begin{bmatrix} \sum_{r=1}^{\#\mathcal{X}_4-1} \beta_\tau \delta (u(x_{4,r}, \tau) - u(x_{4,0}, \tau)) F^{\Delta,1}_{\Delta,k}(x_{4r}) \\ \vdots \\ \sum_{r=1}^{\#\mathcal{X}_4-1} \beta_\tau \delta (u(x_{4,r}, \tau) - u(x_{4,0}, \tau)) F^{\Delta,\#\mathcal{X}_4-1}_{\Delta,k}(x_{4r}) \end{bmatrix} \tag{46}$$

which is a system of $\#\mathcal{X}_4 - 1$ linear equations in the $\#\mathcal{X}_4 - 1$ unknowns $\{u(x_{4,r}, \tau) - u(x_{4,0}, \tau)\}_{r=1}^{\#\mathcal{X}_4-1}$. This system will always have a solution if

ASSUMPTION 17. *The matrix* $\mathrm{dF}^{\boldsymbol{\Delta},\cdot}_{\boldsymbol{\Delta},\mathbf{k}}$ *defined below is invertible a.e.* x_3

$$\mathrm{dF}^{\boldsymbol{\Delta},\cdot}_{\boldsymbol{\Delta},\mathbf{k}} \equiv \begin{bmatrix} \mathrm{dF}^{\Delta,1}_{\Delta,k}(x_{4,1}) & \dots & \mathrm{dF}^{\Delta,1}_{\Delta,k}(x_{4,\#\mathcal{X}_4-1}) \\ \vdots & \vdots & \vdots \\ \mathrm{dF}^{\Delta,\#\mathcal{X}_4-1}_{\Delta,k}(x_{4,1}) & \dots & \mathrm{dF}^{\Delta,\#\mathcal{X}_4-1}_{\Delta,k}(x_{4,\#\mathcal{X}_4-1}) \end{bmatrix}$$

The invertibility assumption implies that not only is there sufficient variation in beliefs, but that this variation in beliefs translates into sufficient independent variation in state probabilities. In the simple case with only 2 states in the final period and where $\mathrm{dF}(x_4)$ depends linearly on z, this reduces to the condition that z has at least two points of support. Note that the above assumption assumes that from every state x_3 every possible point in the support of x_4 is reached with positive probability. If this is not true, the above argument (and assumption) needs to be interpreted as referring to only those points in $\mathcal{X}_4$ that can be reached with positive probability from a given x_3

LEMMA 7. *Consider an agent of type* τ *solving at* $t = 1$ *the problem (3) and that Assumptions 1–4,16 and 17 hold. We observe an i.i.d. sample on* $(\{a_t^*, x_t\}_{t=1}^{T-1}, w)$*. Then,*

1. *Discounted normalized utility in the final period* $\{\beta_\tau\delta\,(u(x_{4,r},\tau) - u(x_{4,0},\tau))\}_{r=1}^{\#\mathcal{X}_4-1}$ *are identified.*
2. *The utility differentials in the penultimate period* $\{(u_3(x_3,k,\tau) - u_3(x_3,0,\tau))\}_{k=1}^{K-1}$ *are identified for* $x_3 \in \mathcal{X}_3$

The proof is straightforward and follows directly from the invertibility assumption and equation (46) and is omitted. An alternative approach that requires fewer points of support in the belief distribution would use the unexploited restrictions on $u(x_4)$ coming from all $K-1$ equations in (45).
Next, as before we can identify β_{τ_N} under the addition of Assumption 7. We state the result but omit the proof.

LEMMA 8. *Consider an agent of type* τ_N *solving at* $t = 1$ *the problem (3) and that Assumptions 1–4, 7,16 and 17 hold. Then,* β_{τ_N} *is identified.*

Period 2 Identification

We next show identification for normalized utility in period 2. We first state the additional assumptions (analogues of Assumptions 8 and 9 respectively).

ASSUMPTION 18. *The distribution of* z *conditional on* x_2 *has at least two points of support. Also,*

1. *For the identified function* $v^*_{\tau_C,3}(s_3,z)$ *defined in (48) below,* $\int v^*_{\tau_C,3}(s_3,z_2)\,\mathrm{dF}_{\Delta,k}(s_3|x_2,z_2) \neq \int v^*_{\tau_C,3}(s_3,z_1)\,\mathrm{dF}_{\Delta,k}(s_3|x_2,z_1)$
2. *For the identified function* $\mathbf{J}(x_2,z)$ $\mathbf{J}(x_2,z_1) \neq \mathbf{J}(x_2,z_2)$ *where* $\mathbf{J}(x_2,z)$ *is defined in Equation (52).*

ASSUMPTION 19. *We normalize utility levels by assuming that* $u_3(x_3,0,\tau)$ *is known for all* $x_3 \in \mathcal{X}_3, \tau \in \mathcal{T}$ *and that final period utility in a base state* $u(x_{4,0},\tau)$ $, \tau \in \mathcal{T}$ *is known.*

We can next state the result

LEMMA 9. *Consider an agent of type* τ *solving at* $t = 1$ *the problem (3) and that Assumptions 1–4, 7,16–19 hold. Then*

1. *Period 2 utility differentials* $\{u_2(x_2,k,\tau) - u_2(x_2,0,\tau)\}_{k=1}^{K-1}$ *are identified for all* $x_2 \in \mathcal{X}_2$
2. δ *and* β_{τ_S} *are identified.*

Proof. The probability that an agent of type τ will take action k in period 2 (corresponding, say to the number of household members he can potentially cover under a treated net) is

$$\mathbb{P}_\tau(a_2^* = j|x_2,z) = \mathbb{P}\left(j = \mathrm{argmax}_k \left\{u_2(x_2,k,\tau) + \epsilon_2(k) + \beta_\tau\delta\int v^*_\tau(s_3,z)\,\mathrm{dF}(s_3|x_2,k,z)\right\}\right) \quad (47)$$

where we define

$$
\begin{aligned}
&v^*_\tau(s_3,z) \equiv v_\tau(s_3, d^\tau(s_3,z), z, 1)\\
&v_\tau(s_3,a,z,c) \equiv u_3(x_3,a,\tau) + \epsilon_3(a) + c\delta \int u(x_4,\tau)\, \mathrm{dF}(x_4|x_3,a,z)\\
&d^\tau(s_3,z) \equiv \mathrm{argmax}_k \left\{ v_\tau(s_3,k,z,\tilde{\beta}_\tau) \right\}\\
&v_{\tau,\Delta}(s_3,a,z,c) \equiv v_\tau(s_3,a,z,c) - v_\tau(s_3,0,z,c).\\
&A^\tau_a(s_3,z) \equiv \mathbb{I}\left\{ a = \mathrm{argmax}_j \{ v_\tau(s_3,j,z,\tilde{\beta}_\tau)\} \right\}
\end{aligned}
$$

We can then rewrite

$$
v^*_\tau(s_3,z) = \sum_{a=1}^{K-1} v_{\tau,\Delta}(s_3,a,z,1) A^\tau_a(s_3,z) + v_\tau(s_3,0,z,1) \tag{48}
$$

We consider identification for the three different types of agent. We first focus on the identification of

$$
\beta_\tau \delta \int v^*_\tau(s_3,z) dF_{\Delta,k}(s_3|x_2,z) \tag{49}
$$

TIME-CONSISTENT AGENTS

For time consistent agents,

$$
v_{\tau_C,\Delta}(s_3,a,z,1) = v_{\tau_C,\Delta}(s_3,a,z,\tilde{\beta}) = \tilde{u}_{\tau_C,a}(x_3,z) + \epsilon_3(a) - \epsilon_3(0)
$$

where $\tilde{u}_{\tau_C,a}(x_3,z)$ is identified from equation (44). Since the distribution of ϵ_3 is known by assumption we conclude that $\int v^*_{\tau_C}(s_3,z) dF_{\Delta,k}(s_3|x_2,z)$ is identified.

Next, as before, using standard inversion arguments (see Appendix C for a direct argument) we can identify

$$
\mathbf{g}_{\tau,2}(x_2,z) = \begin{bmatrix} \tilde{u}_{\tau,2,1} \\ \vdots \\ \tilde{u}_{\tau,2,K-1} \end{bmatrix} \tag{50}
$$

where

$$
\tilde{u}_{\tau,2,k}(x_2,z) \equiv u_2(x_2,k,\tau) - u_2(x_2,0,\tau) + \beta_\tau \delta \int v^*_\tau(s_3,z)\, \mathrm{dF}_{\Delta,k}(s_3|x_2,z). \tag{51}
$$

For time-consistent agents, the last term in the expression above is identified from the previous argument upto δ. Under Assumption 18, evaluating $\tilde{u}_{\tau,k,z}$ at two different values of z allows us to identify δ. Substituting for δ into equation (51) yields identification of $u_2(x_2,k,\tau_C) - u_2(x_2,0,\tau_C)$.

"SOPHISTICATED" TIME-INCONSISTENT AGENTS

For "sophisticated" agents,

$$
\begin{aligned}
v_{\tau_S,\Delta}(s_3,a,z,\tilde{\beta}_{\tau_S}) &= v_{\tau_S,\Delta}(s_3,a,z,\beta_{\tau_S})\\
&= \tilde{u}_{3,\tau_S,a}(x_3,z) + \epsilon_3(a) - \epsilon_3(0)
\end{aligned}
$$

where the first term on the right-hand side is identified. However, $v_{\tau_S,\Delta}(s_3,a,z,1)$ is not identified so that showing that (49) is identified requires more work. First, we write

$$
\begin{aligned}
\beta_{\tau_S} v^*_{\tau_S}(s_3,z) = {} & \underbrace{\beta_{\tau_S} \sum_{a=1}^{K-1} \left(u(x_3,a,\tau_S) - u(x_3,0,\tau_S)\right) A^\tau_a(s_3,z)}_{\hat{\mathbf{J}}} + \underbrace{\sum_{a=1}^{K-1} \beta_{\tau_S}\delta \int \left(u(x_4,\tau_S) - u(x_{4,0},\tau_S)\right) \mathrm{dF}_{\Delta,a}(x_4|x_3,z) A^\tau_a(s_3,z)}_{\hat{\mathbf{K}}}\\
& + \beta_{\tau_S} \underbrace{u_3(x_3,0,\tau_S)}_{\hat{\mathbf{M}}} + \underbrace{\beta_{\tau_S}\delta \int \left(u(x_4,\tau_S) - u(x_{4,0},\tau_S)\right) \mathrm{dF}(x_4|x_3,0,z)}_{\hat{\mathbf{N}}}
\end{aligned} \tag{52}
$$

and examine each expression in turn. First, the term inside the summation does not pose a problem since the period 3 utility functions are identified and the distribution of ϵ is known so that $\mathbf{J} \equiv \int \hat{\mathbf{J}}\, \mathrm{dF}_{\Delta,k}(s_3|x_2, z)$ will be identified. Similarly, $\mathbf{K} \equiv \int \hat{\mathbf{K}}\, \mathrm{dF}_{\Delta,k}(s_3|x_2, z)$ and $\mathbf{N} \equiv \int \hat{\mathbf{N}}\, \mathrm{dF}_{\Delta,k}(s_3|x_2, z)$ are also identified since final period utilities (multiplied by $\beta_\tau\delta$) are identified by Lemma 7. Next, $u_3(x_3, 0, \tau_S)$ is assumed known (Assumption 19) so that $\mathbf{M} \equiv \int \hat{\mathbf{M}}\, \mathrm{dF}_{\Delta,k}(s_3|x_2, z)$ is also identified. Therefore we can write

$$\int \beta_{\tau_S} v^*_{\tau_S}(s_3, z)\, \mathrm{dF}_{\Delta,k}(s_3|x_2, z) = \beta_{\tau_S}(\mathbf{J}(x_2, z) + \mathbf{M}(x_2)) + \mathbf{K}(x_2, z) + \mathbf{N}(x_2, z)$$

so that we can rewrite (51) (for sophisticated types) as

$$\tilde{u}_{\tau_S,2,k}(x_2, z) \equiv u_2(x_2, k, \tau_S) - u_2(x_2, 0, \tau_S) + \delta\Big(\beta_{\tau_S}(\mathbf{J}(x_2, z) + \mathbf{M}(x_2)) + \mathbf{K}(x_2, z) + \mathbf{N}(x_2, z)\Big) \tag{53}$$

and under Assumption 18 we can conclude that β_{τ_S} is identified and consequently that $u_2(x_2, k, \tau_S) - u_2(x_2, 0, \tau_S)$ is identified for $x_2 \in \mathcal{X}_2$.

<u>"NAÏVE" TIME-INCONSISTENT AGENTS</u>

For "naïve" agents, the standard inversion argument identifies

$$\tilde{u}_{\tau_N,2,k}(x_2, z) \equiv u_2(x_2, k, \tau_N) - u_2(x_2, 0, \tau_N) + \beta_{\tau_N}\delta \int v^*_{\tau_N}(s_3, z)\, \mathrm{dF}_{\Delta,k}(s_3|x_2, z). \tag{54}$$

Recall that (β_{τ_N}, δ) are already identified. Further as we will show below, $\int v^*_\tau(s_3, z)\, \mathrm{dF}_{\Delta,k}(s_3|x_2, z)$ is also identified. Therefore, $u_2(x_2, k, \tau_N) - u_2(x_2, 0, \tau_N)$ is identified.

To see that $v^*_\tau(s_3, z)\, \mathrm{dF}_{\Delta,k}(s_3|x_2, z)$ is identified, first recall that

$$\begin{aligned}
v^*_{\tau_N}(s_3, z) &= \sum_{a=1}^{K-1} v_{\tau_N,\Delta}(s_3, a, z, 1) A^\tau_a(s_3, z) + v_{\tau_N}(s_3, 0, z, 1) \\
&= \sum_{a=1}^{K-1} \left(u(x_3, a, \tau_N) - u(x_3, 0, \tau_N)\right) A^\tau_a(s_3, z) + \sum_{a=1}^{K-1} \delta \int \left(u(x_4, \tau_N) - u(x_{4,0}, \tau_N)\right) \mathrm{dF}_{\Delta,a}(x_4|x_3, z) A^\tau_a(s_3, z) \\
&\quad + u(x_3, 0, \tau_N) + \delta \int \left(u(x_4, \tau_N) - u(x_{4,0}, \tau_N)\right) \mathrm{dF}(x_4|x_3, 0, z)
\end{aligned}$$

and each expression on the right hand side above is identified.

□

Period 1 Identification

Conditional on the identification of the type-specific choice probabilities, identification of utility in the first period follows the same arguments as previously. Identifying the type-specific choice probabilities for the first period uses the same arguments as in Lemma 4 (under Assumptions 10 and 11) and we omit the details here.

Once the type specific choice probabilities are identified, using standard inversion arguments (see Appendix C for a direct argument) we can show that the function

$$\mathbf{g}_{\mathbf{1},\tau}(x_1, z) = \begin{bmatrix} \tilde{u}_{\tau,1,1} \\ \vdots \\ \tilde{u}_{\tau,1,K-1} \end{bmatrix} \tag{55}$$

is identified and where

$$\tilde{u}_{\tau,1,k}(x_1, z) \equiv u(x_1, k, \tau) - u(x_1, 0, \tau) + \beta_\tau\delta \int v^*_\tau(s_2, z)\, \mathrm{dF}_{\Delta,k}(s_2|x_1, z) \tag{56}$$

and as before we define the signed measure

$$\mathrm{dF}_{\Delta,k}(s_2|s_1,z) \equiv \mathrm{dF}(s_2|s_1,k,z) - \mathrm{dF}(s_2|s_1,0,z)$$

$$\begin{aligned}
&v^*_\tau(s_2,z) \equiv v_\tau(s_2,d^\tau(s_2,z),z,1)\\
&v_\tau(s_2,a,z,c) \equiv u_2(x_2,a,\tau) + \epsilon_2(a) + c\delta \int v^*_\tau(s_3,\tau)\,\mathrm{dF}(s_3|x_2,a,z)\\
&d^\tau(s_2,z) \equiv \mathrm{argmax}_k \left\{ v_\tau(s_2,k,z,\tilde{\beta}_\tau) \right\}\\
&v_{\tau,\Delta}(s_2,a,z,c) \equiv v_\tau(s_2,a,z,c) - v_\tau(s_2,0,z,c).\\
&A^\tau_a(s_2,z) \equiv \mathbb{I}\left\{ a = \mathrm{argmax}_j \{ v_\tau(s_2,j,z,\tilde{\beta}_\tau) \} \right\}
\end{aligned}$$

We next make an explicit normalization on second period utility

ASSUMPTION 20. *$u_2(x_2,0,\tau)$ is known for all $x_2 \in \mathcal{X}_2$*

Then, $\beta_\tau\delta \int v^*_2(s_2,z)\,\mathrm{dF}_{\Delta,k}(s_3|x_{2,z})$ is identified. We can then directly identify the first period utility (upto a normalization)

$$\tilde{u}_{\tau,1,k}(x_1,z) - \beta\delta \int v^*_\tau(s_2,z)\,\mathrm{dF}_{\Delta,1}(s_2|s_1,z) = u(x_1,k,\tau) - u(x_1,0,\tau)$$

and we record the result

LEMMA 10. *Consider an agent of type τ solving at $t=1$ the problem (3) and that Assumptions 1–4, 7,10,11,16–20 hold. Then, the first period utility differentials $u(x_1,b,\tau) - u(x_1,n,\tau)$ and $u(x_1,c,\tau) - u(x_1,n,\tau)$ are identified for all $x_1 \in \mathcal{X}_1$ and for all types τ. In addition the type probabilities $\{\pi_\tau(\cdot)\}_{\tau\in\mathcal{T}}$ are also identified.*

Appendix C: Inversion Argument

In in the interest of keeping proofs self-contained we provide a simple direct argument for the inversion of choice probabilities that is used repeatedly in the previous proofs.[41] To simplify the exposition, we consider the case where the action space has 3 elements so that $a \in \{0,1,2\}$ although the general case follows exactly analogously. We maintain assumptions Assumptions 1-3 for the argument. The probability that an agent chooses action 0 will be given by

$$\mathbb{P}(a_2=0|x_2) = \mathbb{P}_{x_2}\begin{pmatrix} u(x_2,0)+\epsilon(0)+\beta\delta\int v^*(s_3)\,\mathrm{dF}(s_3|s_2,0) \geq \\ u(x_2,1)+\epsilon(1)+\beta\delta\int v^*(s_3)\,\mathrm{dF}(s_3|s_2,1), \\ u(x_2,0)+\epsilon(0)+\beta\delta\int v^*(s_3)\,\mathrm{dF}(s_3|s_2,0) \geq \\ u(x_2,2)+\epsilon(2)+\beta\delta\int v^*(s_3)\,\mathrm{dF}(s_3|s_2,2) \end{pmatrix}$$

Correspondingly, the probability that an agent will choose action 1 will be given by

$$\mathbb{P}(a_2=1|x_2) = \mathbb{P}_{x_2}\begin{pmatrix} u(x_2,0)+\epsilon(0)+\beta\delta\int v^*(s_3)\,\mathrm{dF}(s_3|s_2,0) \leq \\ u(x_2,1)+\epsilon(1)+\beta\delta\int v^*(s_3)\,\mathrm{dF}(s_3|s_2,1), \\ u(x_2,1)+\epsilon(1)+\beta\delta\int v^*(s_3)\,\mathrm{dF}(s_3|s_2,1) \geq \\ u(x_2,2)+\epsilon(2)+\beta\delta\int v^*(s_3)\,\mathrm{dF}(s_3|s_2,2) \end{pmatrix}$$

[41] see Hotz and Miller (1993) for the original (different) argument. Note that for our argument, we require that the distribution of the unobservable state variables conditional on the observed state variables has support over all of $\mathbb{R}^K$ where K is the number of possible actions

Next, define

$$\hat{u}_1 \equiv u(x_2, 1) - u(x_2, 0) + \beta\delta \int v^*(s_3)\, \mathrm{dF}_{\Delta,1}(s_3|s_2)$$
$$\hat{u}_2 \equiv u(x_2, 2) - u(x_2, 0) + \beta\delta \int v^*(s_3)\, \mathrm{dF}_{\Delta,2}(s_3|s_2)$$

and as usual, the signed measure is defined as

$$\mathrm{dF}_{\Delta,k}(s_3|s_2) \equiv \mathrm{dF}(s_3|s_2, k) - \mathrm{dF}(s_3|s_2, 0)$$

Using this notation, we can write the inequalities more compactly as

$$\mathbb{P}(a_2 = 0|x_2) = \mathbb{P}\left(\epsilon(0) - \hat{u}_1 \geq \epsilon(1), \epsilon(0) - \hat{u}_2 \geq \epsilon(2)|\, x_2\right)$$
$$\mathbb{P}(a_2 = 1|x_2) = \mathbb{P}\left(\epsilon(0) - \hat{u}_1 \leq \epsilon(1), \epsilon(1) + (\hat{u}_1 - \hat{u}_2) \geq \epsilon(2)|\, x_2\right)$$

Suppose that $(\hat{u}_1, \hat{u}_2)$ are not identified from these equations. Then, there exist (u_1^*, u_2^*) such that

$$\mathbb{P}\left(\epsilon(0) - \hat{u}_1 \geq \epsilon(1), \epsilon(0) - \hat{u}_2 \geq \epsilon(2)|\, x_2\right) - \mathbb{P}\left(\epsilon(0) - u_1^* \geq \epsilon(1), \epsilon(0) - u_2^* \geq \epsilon(2)|x_2\right) = 0 \tag{57}$$

$$\mathbb{P}\left(\epsilon(0) - \hat{u}_1 \leq \epsilon(1), \epsilon(1) + (\hat{u}_1 - \hat{u}_2) \geq \epsilon(2)|\, x_2\right) - \mathbb{P}\left(\epsilon(0) - u_1^* \leq \epsilon(1), \epsilon(1) + (u_1^* - u_2^*) \geq \epsilon(2)|\, x_2\right) = 0 \tag{58}$$

We will show that these inequalities are mutually contradictory. We will throughout assume that we are conditioning on x_2. First, assume first that $\hat{u}_1 > u_1^*$. Then, in order for the first equality to hold, we must have $\hat{u}_2 < u_2^*$. To see this, note that if instead $\hat{u}_2 \geq u_2^*$ then the set

$$\{\epsilon(0) - \hat{u}_1 \geq \epsilon(1), \epsilon(0) - \hat{u}_2 \geq \epsilon(2)\} \subset \{\epsilon(0) - u_1^* \geq \epsilon(1), \epsilon(0) - u_2^* \geq \epsilon(2)\} = 0$$

and as long as $\mathrm{dF}(\epsilon|x_2)$ had strictly positive measure on all of $\mathbb{R}^3$, the equality in (57) cannot hold.[42] Therefore, if $\hat{u}_1 > u_1^*$ we must have $\hat{u}_2 < u_2^*$. But, in turn, if this is true, then the equality (58) cannot hold because

$$\{\epsilon(0) - \hat{u}_1 \leq \epsilon(1), \epsilon(1) + (\hat{u}_1 - \hat{u}_2) \geq \epsilon(2)\} \subset \{\epsilon(0) - u_1^* \leq \epsilon(1), \epsilon(1) + (u_1^* - u_2^*) \geq \epsilon(2)\}$$

We can carry out similar arguments using the opposite inequalities to conclude that the $(\hat{u}_1, \hat{u}_2)$ are identified.

Appendix D: Data and Estimation Details

State Space

Malaria: In period one, the malaria indicator is equal to 1 if any one in the household tested positive for malaria using the rapid diagnostic test during the baseline. In period two, the malaria indicator is equal to one if the household reported someone contracting malaria in the period between the purchase of the nets and the first retreatment (this information was collected at the time of the first retreatment). Finally, malaria in period 3 is a binary indicator for malaria incidence in the household that is measured during the follow up survey.

Income: The income indicator is equal to one if a household's income level was high in that period and zero other wise. This variable was derived by first generating an income process and then choosing a cut-off value appropriately. The income process is generated as follows: First, we use household reports about their expectations of future income to construct a household specific income distribution using a triangular distribution. In particular, households at baseline report an upper and lower bound for expected future annual income as well as the probability that the realized income will be greater than the average of the lower and upper bounds. These reports (denoted by $[l, u, q]$)

[42]Note that in this argument all that is required is that the distribution of the ϵ vector conditional on x_2 has support over all of $\mathbb{R}^3$

and the parametric distribution assumption imply that the C.D.F. of income is

$$
\begin{aligned}
F(y) =& \mathbb{I}\{y \leq \frac{l+u}{2}\} \left(\frac{4q}{(u-l)^2}(y-l)^2 \right) + \\
& \mathbb{I}\{y \geq \frac{l+u}{2}\} \left(\frac{4(1-q)}{(u-l)^2}(u-y)^2 \right)
\end{aligned}
$$

Next, we draw from this distribution by inverting the CDF (for a uniformly distributed random variable u) to generate y as

$$
\begin{aligned}
u \leq q \Rightarrow y = l + \left(\frac{(u-l)}{2} \right) \sqrt{\frac{u}{q}} \\
u \geq q \Rightarrow y = u - \left(\frac{(u-l)}{2} \right) \sqrt{\frac{1-u}{1-q}}
\end{aligned}
$$

and we set y equal to the lower or upper bound if the above algorithm yields draws that violate the support condition. Denote the three draws from this distribution as $\{\epsilon_s\}_{s=1}^3$. We then generate income $\{y_t\}_{t=1}^3$ as $y_t = \alpha y_{t-1} + (1-\alpha)\epsilon_t$ where y_0 is baseline income and α is the autoregressive coefficient in the regression of follow-up income on baseline income. We then experimented with various discretizations of the income variable and given the sparseness of the data, settled on a two point distribution depending upon whether household income was above or below the median income for that period (using alternative cut-offs such as the mean did not alter the results). Finally, for periods 2 and 3 the state variables also include the kind of contract purchased in period 1.

Other Variables

Attitudes Towards Risk: We also measured household's attitudes towards risk using a version of the procedure proposed by Holt and Laury (2002). Each respondent was presented with a set of five choice problems. In each problem, the respondent was asked to choose between two lotteries (denoted A and B respectively). The lotteries were designed so that a risk-neutral agent would choose lottery A for the first two problems and switch to lottery B for the remaining 3 problems. We use as our measure of a household's attitude towards risk the number of times the household chose option A in response to the choice problems.

Household Assets: A baseline measure of household assets is used as a conditioning variable in the analysis. The measure is (a function of) the first principal component of the following baseline binary asset indicators: house ownership, motorbike ownership, bicycle ownership, radio ownership, clock ownership, car ownership television ownership, fan ownership, poultry ownership, livestock ownership (small and large), land ownership. In order to ease the first step inversion (which needs to be carried out at each value of the conditioning variables, we classify households into either a low or a high asset category if they were respectively below or above the median of the first principal component.

Beliefs The beliefs data is discussed in greater detail on page 6. In order to ease the first step inversion (which needs to be carried out at each value of the conditioning variables) we use a three point support for the belief data.

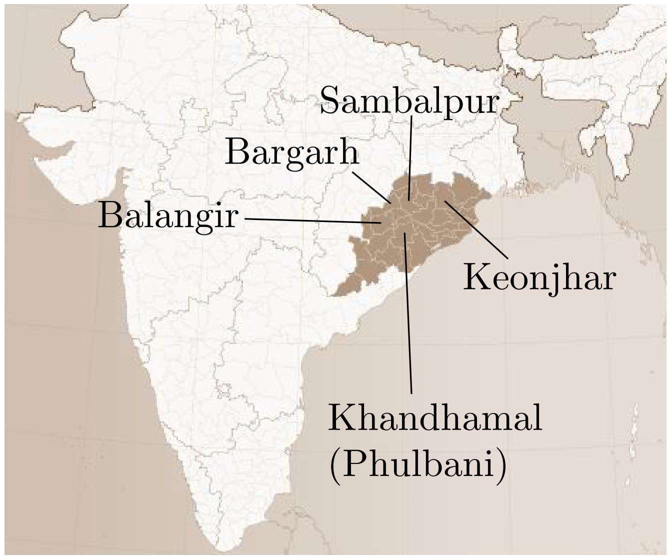

Figure 1: Study Areas

Notes: A total of 166 villages have been included in the study. The 47 communities studied within this paper include 8 villages in Sambalpur, 3 in Kandhamal (Phulbani), 10 in Keonjhar, 10 in Balangir and 16 in Bargarh.

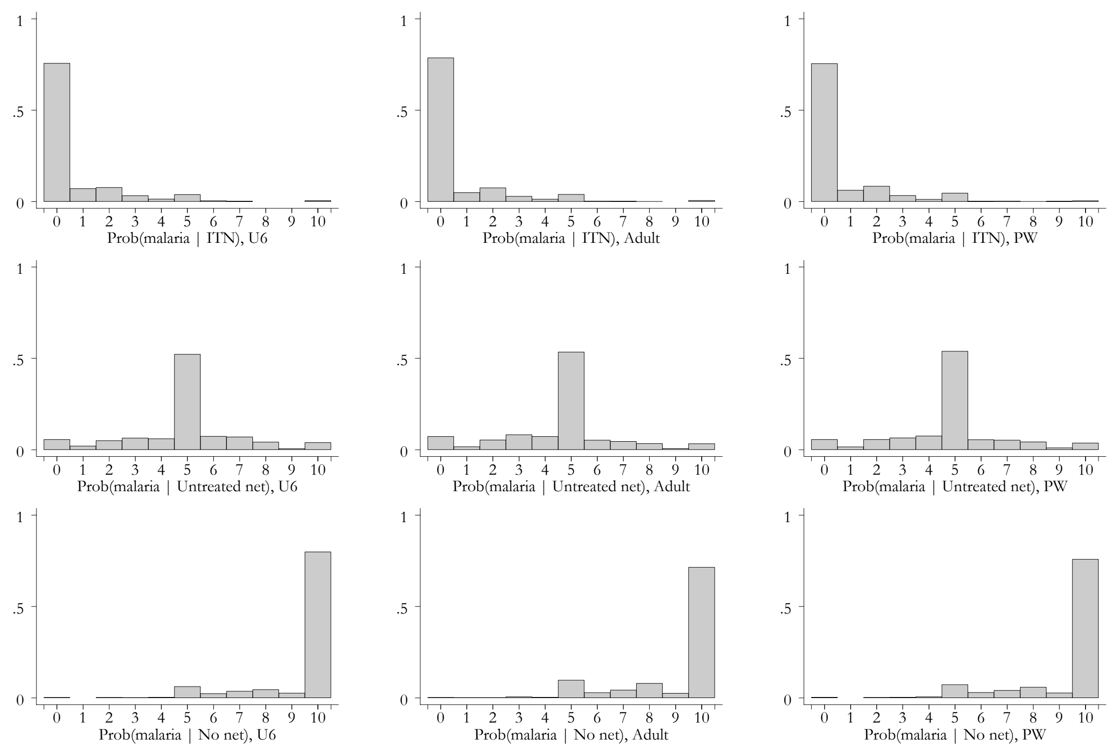

Figure 2: Perceived Protective Power of Bednets

Notes: Histograms of subjective beliefs about the protective power of bednets. Data from March-April 2007 baseline survey.

Table 1: Baseline Summary Statistics

	Mean	Median	S.d.
Household size	5.3	5	2.1
no. children under 5	.49	0	.7
Head is male	.93	1	.25
Age of head	45	45	12
Per capita monthly total expenditure (Rs.)	705	602	426
Highest no. of years of schooling in household	8.6	9	3.6
Nets per head (pre-intervention)	.31	.25	.31
ITNs per head (pre-intervention)	.055	0	.18
Owns at least one net	.67	1	.47
% protected by a net last night	.16	0	.32
% protected by an ITN last night	.032	0	.15
% usually sleeping under net in malaria season	.57	.8	.46
Malaria prevalence (RDT)	.11	0	.29
Anemia prevalence (RDT)	.46	.5	.46
Thinks mosquitoes are malaria carriers	.96	1	.2
Thinks bednets can protect against malaria	.95	1	.22
Expected cost of a malaria episode (working man) (Rs.)	2865	2300	2318
Expected cost of a malaria episode (working woman) (Rs.)	1874	1550	2167
Expected cost of a malaria episode (non-working) (Rs.)	1772	1400	1524
Mean cost of recent (actual) malaria episodes (Rs.)	1326	748	1637

Notes: Data from March-April 2007 survey. $n = 621$. All monetary values are in Rs. (PPP exchange rate $\approx 16 Rs/USD$, World Bank, 2008).

Table 2: Baseline Time Preferences

Prefers Rs. 10 to Rs. 10 in 4 months	0.84
Prefers Rs. 10 to Rs. 12 in 4 months	0.71
Prefers Rs. 10 to Rs. 14 in 4 months	0.65
Prefers Rs. 10 to Rs. 16 in 4 months	0.60
Prefers Rs. 10 to Rs. 10 in 7 months	0.82
Prefers Rs. 10 to Rs. 15 in 7 months	0.63
Prefers Rs. 10 to Rs. 20 in 7 months	0.53
Prefers Rs. 10 to Rs. 25 in 7 months	0.49
Prefers Rs. 10 in 4 months to Rs. 10 in 7 months	0.83
Prefers Rs. 10 in 4 months to Rs. 12 in 7 months	0.73
Prefers Rs. 10 in 4 months to Rs. 14 in 7 months	0.66
Prefers Rs. 10 in 4 months to Rs. 16 in 7 months	0.58
Always prefers earlier reward	0.27
At least one "hyperbolic" preference reversal	0.26
Mean no. of "hyperbolic" preference reversals (> 0)	1.3

Notes: Data from March-April 2007 survey. $n = 621$. "Hyperbolic" preference reversals are defined as cases when the respondent prefers the earlier reward at a short time horizon but switches to the later reward when both time horizons are shifted away from the present by a same time period. The means in the last two rows are calculated including only respondents who displayed at least one hyperbolic preference reversal.

Table 3: Summary of purchases

(A)	Mean	S.d.
Fraction who purchased at least one loan contract	.53	.5
no. ITNs purchased on credit	.51	1.2
no. ITNs + 2 re-treatments purchased on credit	.65	2.1
no. ITNs purchased on credit (conditional on purchase)	1.9	1.6
no. ITNs + 2 re-treatments purchased on credit (conditional on purchase)	2.3	3.5
no. of households who purchased ITNs only	153	
no. of households who purchased ITNs + retreatment contracts	165	
no. of households who purchased both types of contracts	12	
(B)	Mean	S.d.
Fraction who purchased at least one cash contract	.023	.15
no. ITNs purchased for cash	.024	.33
no. ITNs + 2 re-treatments purchased for cash	.0081	.09
no. ITNs purchased on loan (conditional on purchase)	2.5	2.5
no. ITNs + 2 re-treatments purchased on loan (conditional on purchase)	1	0

Notes: Data from September-November 2007.

Table 4: Monte Carlo Results: Directly Observed Types

	Mean	Median	Std.Dev	IQR
N=150				
δ	0.81	0.76	0.75	0.93
β_N	0.76	0.68	0.45	0.57
β_S	0.91	0.80	0.64	0.65
θ_4	3.82	2.86	8.27	3.49
θ_5	1.10	1.11	0.84	1.15
N=300				
δ	0.88	0.86	0.52	0.65
β_N	0.74	0.71	0.30	0.40
β_S	0.83	0.79	0.32	0.41
θ_4	4.37	3.09	4.92	3.11
θ_5	1.03	1.03	0.56	0.74
N=600				
δ	0.90	0.86	0.37	0.52
β_N	0.71	0.70	0.18	0.25
β_S	0.81	0.78	0.23	0.28
θ_4	3.71	3.09	2.10	1.93
θ_5	1.04	1.04	0.38	0.51
N=1200				
δ	0.88	0.88	0.25	0.35
β_N	0.71	0.70	0.13	0.20
β_S	0.81	0.80	0.16	0.21
θ_4	3.37	3.09	1.12	1.29
θ_5	1.02	1.02	0.27	0.39
N=2400				
δ	0.89	0.89	0.18	0.26
β_N	0.69	0.69	0.09	0.13
β_S	0.80	0.79	0.11	0.14
θ_4	3.17	3.03	0.73	0.94
θ_5	1.00	0.99	0.18	0.24
N=4800				
δ	0.91	0.92	0.13	0.18
β_N	0.70	0.70	0.07	0.10
β_S	0.80	0.79	0.08	0.10
θ_4	3.05	2.97	0.47	0.59
θ_5	1.01	1.01	0.13	0.18

Notes: Each model was simulated 500 times. The true values for the parameter vector are $(\delta, \beta_N, \beta_S, \theta_4, \theta_5) = (.9, .7, .8, 3, 1)$

Table 5: Monte Carlo Results: Unobserved Types

	Mean	Median	Std.Dev	IQR
N=150				
δ	0.4992	0.3229	0.4955	0.6216
β_N	1.6499	0.5385	2.9145	1.9593
β_S	0.7057	0.5981	0.6116	0.4903
θ_4	7.6032	6.2491	4.7227	4.1101
θ_5	1.0383	1.0037	0.6848	0.9602
N=300				
δ	0.6669	0.6309	0.3303	0.4147
β_N	0.4034	0.2795	0.4306	0.6809
β_S	0.9608	0.9315	0.4766	0.6875
θ_4	5.0283	4.0879	3.3146	3.0744
θ_5	1.0576	1.0462	0.5426	0.6934
N=600				
δ	0.7377	0.7051	0.3016	0.4182
β_N	0.4330	0.4020	0.4000	0.4674
β_S	0.9475	0.9263	0.3027	0.4387
θ_4	4.0817	3.6559	1.7953	2.2880
θ_5	1.0742	1.0695	0.3836	0.5152
N=1200				
δ	0.7200	0.7132	0.2873	0.4098
β_N	0.4374	0.4146	0.2944	0.3445
β_S	1.0216	0.9626	0.3825	0.4188
θ_4	4.0004	3.3144	2.1385	2.0337
θ_5	1.0135	1.0107	0.2779	0.3777
N=2400				
δ	0.7865	0.7751	0.2083	0.2920
β_N	0.4137	0.4096	0.1782	0.2229
β_S	0.9701	0.9552	0.2143	0.2611
θ_4	3.2838	3.0896	0.9665	1.2084
θ_5	1.0159	1.0165	0.2054	0.2545
N=4800				
δ	0.7921	0.7882	0.1464	0.2053
β_N	0.4071	0.4040	0.1217	0.1595
β_S	0.9623	0.9503	0.1479	0.1894
θ_4	3.1284	3.0296	0.5999	0.8060
θ_5	1.0093	1.0079	0.1401	0.1745

Notes: Each model was simulated 500 times. The true values for the parameter vector are $(\delta, \beta_N, \beta_S, \theta_4, \theta_5) = (.8, .4, .95, 3, 1)$

Table 6: Predicting Contract Choice (Period 1)

	(1)	(2)	(3)	
	Logit	Logit	Multinomial Logit	
	Any Contract	C1 vs C2	C1	C2
Malaria (t=1)	0.34**	0.20	0.42*	0.26
	(0.17)	(0.31)	(0.24)	(0.21)
Income (t=1)	-0.37*	0.055	-0.31	-0.43
	(0.22)	(0.33)	(0.28)	(0.27)
ITN Efficacy Beliefs (γ)	-1.62*	1.94	-0.70	-2.39**
	(0.98)	(1.20)	(1.29)	(1.05)
ITN Efficacy^2 (γ^2)	2.26*	-2.31*	1.16	3.18**
	(1.20)	(1.32)	(1.52)	(1.24)
Time Inconsistency	0.10	-0.071	0.046	0.15
	(0.19)	(0.27)	(0.21)	(0.26)
Household Size	0.028	-0.11	-0.017	0.065
	(0.047)	(0.078)	(0.061)	(0.055)
Baseline Assets	0.024	-0.35	-0.14	0.19
(Tercile 2)	(0.19)	(0.27)	(0.25)	(0.22)
Baseline Assets	-0.45*	-0.47	-0.69**	-0.21
(Tercile 3)	(0.26)	(0.36)	(0.35)	(0.29)
2 #A Choices	-0.24	0.55	0.012	-0.49
	(0.29)	(0.34)	(0.33)	(0.33)
3 #A Choices	-0.33	0.072	-0.30	-0.36
	(0.26)	(0.36)	(0.34)	(0.30)
4 #A Choices	-0.43	-0.19	-0.52*	-0.35
	(0.31)	(0.34)	(0.32)	(0.37)
Baseline Net	0.18	0.34	0.34	0.025
Ownership	(0.24)	(0.30)	(0.30)	(0.27)
Constant	0.45	0.097	-0.25	-0.25
	(0.46)	(0.73)	(0.67)	(0.49)
Observations	572	290	572	

Notes: Malaria is equal to 1 is anyone in the household suffered from Malaria in the six months before the offer of ITNs. ITN Efficacy beliefs are household reported probabilities about the efficacy of ITNs. Time Inconsistency is equal to 1 if the household exhibited at least one preference reversal to a set questions designed to elicit time preferences. Baseline Assets is a ternary valued variable (3 is the highest tercile) based on a principal component decomposition of asset ownership at baseline. The # A choices variable takes on 4 values equal to the number of times the respondent picked option A in the questions on risk aversion (see Section 2 and Appendix 8 for more details on the variables used). Standard errors are clustered at the village level. T-test significant at ***1%, **5%,*10%

Table 7: Predicting Retreatment Rates

	(1) Retreatment	(2) C1 Retreatment	(3) C2 Retreatment	(4) C1 -Period 3	(5) C2 - Period 3
Standard (C1)	-0.49***				
	(0.075)				
Malaria		-0.56	-2.40**	-0.11	-0.16
		(0.44)	(1.10)	(0.35)	(0.62)
Income		0.49	-0.60	-0.31	-0.34
		(0.41)	(0.66)	(0.61)(	(0.63)
ITN Efficacy Beliefs (γ)		-0.60	7.40***	-3.30*	4.86***
		(2.11)	(1.63)	(1.84)	(1.65)
ITN Efficacy^2 (γ^2)		1.44	-4.89*	6.05***	-3.20
		(2.01)	(2.54)	(2.21)	(2.03)
Time Inconsistency		-1.04**	0.45	-0.78	0.74
		(0.47)	(0.56)	(0.62)	(0.56)
Household Size		-0.064	-0.074	-0.20	0.22
		(0.11)	(0.20)	(0.18)	(0.22)
Baseline Assets		0.49	-1.14*	0.84	-0.14
(Tercile 2)		(0.47)	(0.60)	(0.61)	(0.47)
Baseline Assets		0.62	-1.12	0.49	0.64
(Tercile 3)		(0.60)	(0.84)	(0.62)	(0.57)
2 #A Choices		-0.090	-1.31	-0.38	0.30
		(0.40)	(1.21)	(0.50)	(0.98)
3 #A Choices		-0.073	-1.43	-1.79**	-0.29
		(0.54)	(1.16)	(0.82)	(0.77)
4 #A Choices		0.17	-0.78	-1.02*	-0.43
		(0.59)	(1.15)	(0.53)	(0.69)
Baseline Net		0.11	0.15	0.94*	-0.62
Ownership		(0.48)	(0.78)	(0.54)	(0.66)
Constant	0.93***	-0.23	4.06***	-0.26	0.36
	(0.041)	(1.04)	(1.57)	(1.13)	(1.81)
Observations	290	141	149	141	149

Notes: Malaria is equal to 1 is anyone in the household suffered from Malaria in the six months before the offer of ITNs. ITN Efficacy beliefs are household reported probabilities about the efficacy of ITNs. Time Inconsistency is equal to 1 if the household exhibited at least one preference reversal to a set questions designed to elicit time preferences. Baseline Assets is a ternary valued variable (3 is the highest tercile) based on a principal component decomposition of asset ownership at baseline. The # A choices variable takes on 4 values equal to the number of times the respondent picked option A in the questions on risk aversion (see Section 2 and Appendix 8 for more details on the variable used). Models are estimated using the Logit and standard errors are clustered at the village level. T-test significant at ***1%, **5%,*10%

Table 8: Type Probabilities

$\pi_\tau(r)$	Estimate	2.5	97.5
$\pi_C(0)$	0.3870	0.2894	0.4837
$\pi_N(0)$	0.5019	0.4172	0.6059
$\pi_S(0)$	0.1111	0.0593	0.1691
$\pi_C(1)$	0.4143	0.3092	0.5126
$\pi_N(1)$	0.4699	0.3851	0.5790
$\pi_S(1)$	0.1158	0.0639	0.1756

Notes: $\pi_\tau(r)$ is the probability that an agent is of type τ given their response to the time-inconsistency question r. The second and third columns are the 2.5^{th} and 97.5^{th} percentiles of the bootstrap distribution of the type-probabilities computed using a clustered bootstrap (at the village level) with 200 replications.

Table 9: Type Probabilities (Conditional on Contract Choices C1 or C2)

$\pi_\tau(r, a_1)$	Estimate	2.5	97.5
$\pi_C(0,1)$	0.3565	0.2776	0.4484
$\pi_N(0,1)$	0.5153	0.4409	0.5876
$\pi_S(0,1)$	0.1282	0.0729	0.2034
$\pi_C(0,2)$	0.2780	0.1394	0.4186
$\pi_N(0,2)$	0.5697	0.4644	0.6856
$\pi_S(0,2)$	0.1523	0.0486	0.2978
$\pi_C(1,1)$	0.3970	0.3196	0.4792
$\pi_N(1,1)$	0.4738	0.4052	0.5505
$\pi_S(1,1)$	0.1292	0.0801	0.1938
$\pi_C(1,2)$	0.2978	0.1530	0.4372
$\pi_N(1,2)$	0.5607	0.4512	0.6804
$\pi_S(1,2)$	0.1415	0.0438	0.2743

Notes: $\pi_\tau(r, a)$ is the probability that an agent is of type τ given their response to the time-inconsistency question r and their choice of contract a_1 (1 represents the standard contract C1 and 2 represents the commitment contract C2). The second and third columns are the 2.5^{th} and 97.5^{th} percentiles of the bootstrap distribution of the type-probabilities computed using a clustered bootstrap (at the village level) with 200 replications.

Table 10: Unobserved Types (Type Invariant Risk Aversion)

	Estimate	2.5	97.5
δ	0.7894	0.0010	0.9351
β_N	0.9727	0.9161	0.9798
β_S	0.5533	0.0060	0.7311
α_1	0.7219	0.6047	1.8436
α_4	0.5755	0.3725	1.6225
α_5	0.8424	0.6911	1.5438
α_6	0.9401	0.7935	2.0000
κ_C	-0.6870	-1.7168	1.0740
κ_N	-0.5277	-1.6417	0.8373
κ_S	-1.2090	-2.0000	1.3694
κ_4	-1.5486	-1.9933	0.3874
κ_5	-1.2644	-1.7997	1.7145

Notes: δ is the exponential discount parameter. β_N is the hyperbolic parameter for "naïve" time-inconsitent agents, β_S is the corresponding parameter for "sophisticated" time-inconsistent agents. the α vector parameterizes the risk-aversion parameter and the κ vector parameterizes the malaria cost function. The second and third columns are the 2.5^{th} and 97.5^{th} percentiles respectively of the clustered bootstrap distribution of the two-step estimation procedure outlined in section 6.3 with 200 replications.

Table 11: Unobserved Types (Types vary by risk aversion)

	Estimate	2.5	97.5
δ	0.7880	0.0000	0.9351
β_N	0.9757	0.9313	0.9798
β_S	0.5727	0.0007	0.7311
α_C	0.7230	0.6047	1.7890
α_N	0.4348	0.2935	1.9277
α_4	0.5513	0.3725	1.9736
α_5	0.8389	0.6911	2.0000
α_6	0.9205	0.7935	1.9445
κ_C	0.0070	-1.9950	1.0754
κ_N	-0.1998	-0.6869	0.8373
κ_S	-0.5314	-1.9951	1.3667
κ_S	-0.9613	-1.2298	0.3725
κ_5	-0.3721	-2.0000	1.6852

Notes: δ is the exponential discount parameter. β_N is the hyperbolic parameter for "naïve" time-inconsitent agents, β_S is the corresponding parameter for "sophisticated" time-inconsistent agents. the α vector parameterizes the risk-aversion parameter and the κ vector parameterizes the malaria cost function. The second and third columns are the 2.5^{th} and 97.5^{th} percentiles respectively of the clustered bootstrap distribution of the two-step estimation procedure outlined in section 6.3 with 200 replications.

Table 12: Counterfactuals: Change in Take up and Retreatment Rates

Outome	Retreatment Price Rs. 30	Retreatment Price Rs. 7
% Change C1 Take Up	-5.39	17.12
	[-7.19,2.67]	[0.51, 20.21]
% Change C2 Take Up	1.68	-7.98
	[-12.54,2.69]	[-9.02, 7.66]
% Change Retreatment C1	-74.15	81.91
	[-74.61,-74.04]	[81.10, 84.96]

Notes: All changes are relative to the retreatment price of Rs.15 offered during the intervention. All figures are arrived at by averaging over the estimated distribution of demographics, beliefs and types. Figures for retreatment are provided separately for each contract type. The figures in parentheses are the 2.5^{th} and 97.5^{th} percentiles respectively of the clustered bootstrap distribution of the estimates

Table 13: The Relative Importance of Risk, Cost and Time Preferences

Outcome	Model 1	Model 2	Model 3	Model 4
C1 Take Up (C)	.14	.14	.14	.14
	[.138,.144]	[.138,.144]	[.135,.143]	[.136,.143]
C1 Take Up (N)	.14	.14	.16	.17
	[.139,.145	[.139,.145]	[.056,.272]	[.057,.275]
C1 Take Up (S)	.14	.16	.15	.17
	[.139,.145]	[.145,.177]	[.015,.290]	[.015,.325]
C2 Take Up (C)	.31	.31	.31	.31
	[.129,.491]	[.129,.491]	[.128,.483]	[.128,.483]
C2 take Up (N)	31	.31	.31	.31
	[.129,.486]	[.129,.486]	[.041,.583]	[.041,.578]
C2 take Up (S)	.31	.21	.31	.21
	[.129,.486]	[.082,.342]	[.002,623]	[.002,.466]
Retreatment C1 (C)	. 19	.19	.19	.19
	[.181,.191]	[.181,.191]	[.181,.191]	[.181,.191]
Retreatment C1 (N)	. 19	.19	.19	.19
	[.181,.191]	[.181,.191]	[.174,.198]	[.174,.197]
Retreatment C1 (S)	.19	.18	.18	.18
	[.181,.191]	[.180,.186]	[.154,.208]	[.158,.203]
Retreatment C2 (C)	.51	.51	.51	.51
	[.500,.515]	[.500,.515]	[.501,.515]	[.500,.515]
Retreatment C2 (N)	.51	.51	.51	.51
	[.500,.515]	[.500,.515]	[.492,.526]	[.493,.526]
Retreatment C2 (S)	.51	.50	.50	.50
	[.500,.515]	[.498,.508]	[.471,.529]	[.478,.521]

Notes: In Model 1 all types have the same risk and cost preferences but β_S is set equal to β_N. Model 2 is the same as Model 1 but with $\beta_N \neq \beta_S$. In model 3, risk and cost parameters are as in Table 11 but $\beta_N = \beta_S = 1$. In Model 4, parameter values are as in Table 11 . Key: C = Time Consistent, N = "Naïve" Inconsistent, S= "Sophisticated" Inconsistent. The figures in parentheses are confidence intervals computed using the clustered bootstrap