NBER WORKING PAPER SERIES

IDENTIFICATION OF TIME-INCONSISTENT MODELS: THE CASE OF INSECTICIDE TREATED NETS

Aprajit Mahajan Christian Michel Alessandro Tarozzi

Working Paper 27198 http://www.nber.org/papers/w27198

NATIONAL BUREAU OF ECONOMIC RESEARCH 1050 Massachusetts Avenue Cambridge, MA 02138 May 2020

We are grateful to Øystein Daljord, Hanming Fang, Paul Heidhues, Nathan Hendren, Han Hong, Charlie Sprenger, Petra Todd, two exceptionally constructive anonymous referees and seminar participants at several conferences and institutions for valuable comments and suggestions. Florens Odendahl provided excellent research assistance. We also gratefully acknowledge financial support from the following sources for the completion of the field work that originated the data used in the paper: The Centre for Micro Finance (Chennai, India), the Stanford Presidential Fund for Innovation in International Studies, the Stanford Center for International Development, the Stanford OTL Research Incentive Fund, RAND Corporation, the Duke Arts and Sciences Committee on Faculty Research and the National Institute of Allergy and Infectious Diseases (Grant R03AI078119). A previous draft of the paper was titled "Time Inconsistency, Expectations and Technology Adoption: The case of Insecticide Treated Nets". Michel gratefully acknowledges support from the Spanish Ministerio de Economía y Competitividad (Grant ECO2016-76998) and the Jose Castillejo Mobility Grant (CAS19/00246). Tarozzi gratefully acknowledges support from the Marie Curie International Incoming Fellowship FP7-PEOPLE-2011-IXF, Proposal ID 298904 and the Spanish Ministerio de Economía y Competitividad (Grants ECO2012-33299 and ECO2015-71536-ERC). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Aprajit Mahajan, Christian Michel, and Alessandro Tarozzi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Identification of Time-Inconsistent Models: The Case of Insecticide Treated Nets Aprajit Mahajan, Christian Michel, and Alessandro Tarozzi NBER Working Paper No. 27198 May 2020 JEL No. D9,I1,I3

ABSTRACT

Time-inconsistency may play a central role in explaining inter-temporal behavior, particularly among poor households. However, little is known about the distribution of time-inconsistent agents, and time-preference parameters are typically not identified in standard dynamic choice models. We formulate a dynamic discrete choice model in an unobservedly heterogeneous population of possibly time-inconsistent agents. We provide conditions under which all population type probabilities and preferences for both time-consistent and sophisticated agents are point-identified and sharp set-identification results for naïve and partially sophisticated agents. Estimating the model using data from a health intervention providing insecticide treated nets (ITNs) in rural Orissa, India, we find that time-inconsistent agents account for almost 80 percent of our sample and that sophisticated and naïve agents are considerably present-biased. Counterfactuals show that the under-investment in ITNs attributable to present-bias leads to substantial costs that are about six times the price of an ITN.

Aprajit Mahajan Dept. of Agricultural & Resource Economics University of California, Berkeley 219 Giannini Hall Berkeley, CA 94720-3310 and NBER aprajit@gmail.com

Christian Michel Department of Economics and Business Universitat Pompeu Fabra and Barcelona GSE Jaume I building, 20-1E76 Ramon Trias Fargas, 25-27 08005 Barcelona Spain christian.michel@upf.edu Alessandro Tarozzi Department of Economics and Business Universitat Pompeu Fabra and Barcelona GSE Jaume I building, 20-2E12 Ramon Trias Fargas, 25-27 08005 Barcelona Spain alessandro.tarozzi@upf.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. However, 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 and under-investment in activities with apparent low costs and high expected returns.¹ 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 "sub-optimal" choices among poor individuals in developing countries, a context where such choices may have particularly dire consequences (Bernheim et al., 2015; Mullainathan, 2004; Carvalho et al., 2016). Non-standard preferences displaying bias towards the present have been proposed to explain poverty traps (Banerjee and Mullainathan, 2010; Ubfal, 2016), the existence of demand for commitment devices in savings or health-protecting technologies (Ashraf et al., 2006; Tarozzi et al., 2009, 2014; Schilbach, 2019), productivity (Kaur et al., 2014) and low demand for immunization and fertilizer (Banerjee et al., 2010; Duflo et al., 2011).

Present bias is typically modeled assuming that preferences are characterized by "hyperbolic discounting" (Laibson, 1997). In such models, at each time t, future utility at any time s (> t) is discounted not by the usual exponential discount 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 to receive a reward at current date t over a larger one to be received at date t + k, but instead choose the larger reward if the two reward dates are each shifted forward by some length of time s (i.e. to t + s and t + s + k).² Such choices are not consistent with standard models of inter-temporal preferences.

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. In addition, behavior typically differs between 'sophisticated' individuals who are aware of having such time-inconsistent preferences and 'naïve' individuals who are not. 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, the formal identification and credible estimation of the discount factors that characterize hyperbolic preferences is non-trivial.

In fact, the time preference parameter δ is generically not identified even in standard dynamic choice models (Rust, 1994; Magnac and Thesmar, 2002). This non-identification result applies *afortiori* to both β and δ in the hyperbolic " β - δ " formulations of time-inconsistent preferences which dominate empirical work on time-inconsistency. Equally importantly, and particularly for empirical work, it seems

¹See Frederick et al. (2002), DellaVigna (2009), Sprenger (2015), and Ericson and Laibson (2018) for reviews.

²See Andreoni and Sprenger (2012) for an alternative explanation for these findings and Augenblick et al. (2015) for a similar finding when agents' choices are over effort (rather than money).

unsatisfactory to assume that all agents in a population have the same time preferences. It is thus important to account for heterogeneity in time preferences—for instance allowing for both time-consistent and inconsistent agents. However, agent type is usually not directly observed by the researcher so that a model with unobserved types seems necessary.

This paper makes two contributions to the literature. First, we provide identification results for dynamic discrete choice models with time-inconsistent agents and unobserved types allowing for rich heterogeneity in per-period utility as well as time preferences. We also identify the population distribution of types, an object of direct policy interest. Second, we estimate a parametric version of the model using data from a field intervention to study the importance of present bias in explaining investment in a health preventive technology in a developing country. Specifically, we study demand for insecticide-treated nets (henceforth ITNs, a key product for the reduction of malaria risk), as well as for their recommended periodic re-treatment using specially collected data from malarious areas of rural Orissa, India.

In the general model we overcome the previous non-identification results by adding information in the form of two key exclusion restrictions. The first is the existence of a variable z that only affects utility via the perceived value of future states. The second is the presence of a variable r that acts as a (imperfect) signal of agent type but which, conditional on agent type and observables, provides no additional information about agent choices. In the empirical application, the role of z is played by elicited beliefs about the evolution of state variables, while r comprises elicited indicators of time preferences and agents' decisions to purchase products designed to appeal *apriori* to particular types.

In the general version of the model we assume that there exists an unknown (but finite) number of types with possibly time-inconsistent preferences. We begin by first identifying the total number of types in the population. Following this, we then identify the nature of the time-preferences of each type, classifying each type as either time-consistent, time-inconsistent "sophisticated" (if the agent is aware of the time-inconsistency implied by the preference structure) or time-inconsistent "naïve" (if the agent lacks such awareness). We allow for the possibility that there are multiple sub-types within each broad class of type of agent (i.e. that there are multiple types of time-consistent or sophisticated or naïve agents). Finally, for each type, we provide identification results for the preference parameters.

We show that in the most general version of the model where all types can have distinct time-preference parameters, all parameters are point-identified except for the time preferences of the naïve types. In this latter case we provide sharp bounds for the parameters of interest, and we show point identification under a further set of additional (but commonly assumed) conditions.

Next, we introduce our empirical application and carry out a set of Monte Carlo simulations that suggest that the time-preference parameters are well estimated with sample sizes similar to that in the application. We then estimate our model, and find that approximately 21% of the population is time-consistent, while 49% are naïve inconsistent and the remaining 30% are sophisticated inconsistent. Further, we find that naïve and sophisticated agents are considerably present-biased with our preferred estimates of β being 0.16 (for naïve agents) and 0.08 (for sophisticated agents). Both these sets of estimates (of the population distribution of types as well as of the separate β parameters), to our knowledge, are new to the literature. Second, we ask whether time-inconsistent preferences provide a better explanation for observed choices than alternative explanations that stress differences in per-period utilities. We find that while per-period utilities do vary across agent types, they are not substantively important in explaining outcomes in our sample.

Third, we evaluate the extent to which present-biased but sophisticated agents are more likely to choose specially designed "commitment" products (Bryan et al. 2010). These ITNs in our context require regular retreatment with insecticide in order to remain effective against mosquitoes. Households were offered the choice to buy ITNs either with a standard contract (with the option to purchase retreatment at a later time) or with a "commitment contract" which also included a bundle of two consecutive retreatments. The commitment contract was designed to mitigate the time-inconsistency problem associated with retreatment. We find that commitment products are not particularly appealing to sophisticated agents and that the purchase of these products is in fact higher among naïve households. Note that this contradicts a deterministic mapping (that is commonly assumed) from the take-up of commitment products into agent type whereby the choice of a commitment product reveals an agent to be sophisticated. This is possible within our model since we allow the choice to commit to only imperfectly predict type which in turn allows for a much richer analysis of preferences. 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 and agents have identical preferences.

Fourth, we quantify the relationship between the extent of present-bias and the expected cost of malaria. *Ceteris paribus*, a higher present-bias leads to lower ITN purchases and fewer retreatments. Since ITNs reduce the risk of malaria, fewer ITN purchases and retreatments increase the likelihood of contracting malaria. We find that the median (undiscounted) additional expected total cost of malaria during our study period exceeds the price of a treated net by a factor of around six. However, given the high fraction of time-inconsistent households and the high levels of present-bias, the discounted total costs of malaria are low for many inconsistent agents compared to the price of an ITN. This explains low demand, which is problematic from the perspective of a social planner given the strong evidence of positive externalities of ITNs (Lengeler, 2009).

In drawing links to the extensive literature on time-inconsistency and on structural estimation with unknown types we focus on work that is closest to our approach.³ Our identification results rely on the conditional choice probabilities approach pioneered by Hotz and Miller (1993). Our work is most closely related to Abbring and Daljord (2018), Abbring et al. (2019) and Fang and Wang (2015) but there are important differences. First, we consider a setting with multiple unobserved types while these papers consider the case of a single observed type. Second, as a consequence the distribution of these multiple types is a key parameter, and this allows us to assess the prevalence of the time-inconsistency problem in terms of proportion of each type in the population combined with the type-specific magnitude of time-inconsistency. Third, our model is motivated by our specific setting in which purposively collected data (beliefs about future state evolution) provide a natural candidate for the exclusion restrictions. This appears to be an important source of identifying variation and perhaps contributes to our Monte Carlo simulations being quite encouraging relative to the literature.

We provide identification results for cases that have not – to our knowledge – previously been covered in the literature. In the overlapping case of the single known type with known error distribution, our arguments and assumptions were inspired by those in Proposition 4 of Magnac and Thesmar (2002) –

³See, e.g., Aguirregabiria and Mira (2010) or Arcidiacono and Ellickson (2011) for a survey of work on dynamic discrete choice structural models, and DellaVigna (2017) for a survey of work on structural models in behavioral economics.

an exclusion restriction and a rank condition similar to the assumptions in Abbring et al. (2019). Our exclusion restrictions arise naturally as restrictions on elicited beliefs about future states and how they enter the choice problem. The special case of our results for a single known sophisticated type are closest to the model in Abbring and Daljord (2018) and Abbring et al. (2019). We combine the fact that we can identify final period utilities with an exclusion restriction and a rank condition (conditional on state variables that enter the per-period utility function) to identify earlier period utilities as well as certain combinations of time-preference parameters. Relative to Fang and Wang (2015) we use additional information and our identification argument is constructive (see also Abbring and Daljord (2019) who critique their identification results). Finally, we have a substantive empirical application to which we apply our identification results.

Our identification arguments for unknown types are closely related to those in Kasahara and Shimotsu (2009). We differ in that we achieve identification by imposing an exclusion restriction by requiring a variable that affects type probabilities but not the choice probabilities, while they place assumptions on the length of the panel available to the researcher. In addition, they do not consider identification and estimation of time preferences or time-inconsistency. Our work is also related to that of van der Klaauw (2012) and van der Klaauw and Wolpin (2008) though they 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.

Like Ashraf et al. (2006), we use elicited time preferences to predict behavior and we design a product that should appeal to sophisticated inconsistent agents, although they focus on reduced-form correlations between preference reversals and demand for commitment devices in savings markets and do not estimate time preference parameters. Augenblick et al. (2015) conduct a laboratory experiment where choices identify potentially heterogeneous time-preference parameters for agents who may be partially sophisticated. Bai et al. (2017) use a field experiment to estimate a structural model where per-period utility is parametric and time-inconsistency parameters are drawn from a parametric distribution. Unlike our study, they find low compliance rates among agents who chose commitment contracts, and they attribute this to partial naïveté. Heidhues and Strack (2019) provide identification results for partially naïveté in a stopping problem when data on both the stopping probabilities and on the continuation value are available. Martinez et al. (2019) adapt their model in the context of filing tax returns and find a non-negligible present-bias (assuming a per-period discount factor $\delta = 1$). Our paper is also related to Andreoni et al. (2016) who focus on estimating individual level time-preference parameters and using them to design incentive schemes for health workers.

The paper is organized as follows. Section 2 outlines the basic elements of the dynamic discrete choice model with different types and describes the model primitives in some detail. Section 3 provides the identification, first for the simpler case where observables reveal types completely, and then for the more realistic case where type is only imperfectly observed. Section 4 describes the data, the estimation methodology and the empirical results, followed by a set of counterfactual exercises. Section 5 concludes. Additional proofs related to the empirical application, alternative modeling assumptions, Monte Carlo simulations, and estimation details are relegated to the online appendix.

2 Model

We consider a three-period dynamic discrete choice model with a finite action and state space. Three periods are the minimum required to identify the time-preference parameters and the finite state and action space is required in order to use the deconvolution arguments required for the most general identification results. We begin by defining and placing assumptions on the state space, the action space, the transition probabilities, the class of acceptable decision rules and finally the preferences and objective function maximized by the agent.

STATE SPACE: S_t .

The state space S_t can be partitioned as $S_t \equiv (\mathcal{X}_t, \mathcal{Z}_t, \mathcal{E}_t)$ where $(\mathcal{X}_t, \mathcal{Z}_t)$ denote the domain of the state variables that are observed by both the researcher and the agent and \mathcal{E}_t denotes the domain of the state variables that are only observed by the agent. We distinguish between two kinds of observed state variables: $x_t \in \mathcal{X}_t$ enter the static payoff functions (or per-period utilities, defined below) while $z_t \in \mathcal{Z}_t$ do not – see Assumption 3 for a formal definition of this exclusion restriction. In the empirical application z_t comprises subjective beliefs elicited from the agent about the distribution of x_{t+1} and these are plausibly excludable from the period t payoff function (conditional on the observed state). More generally, we can allow for a rich observable state space with the substantive restriction that it is finite.

The vector of unobserved state variables $\epsilon_t \in \mathcal{E}_t$ is absolutely continuous (w.r.t. the Lebesgue measure) and has dimension equal to the number of actions available to the agent in period t. ACTION SPACE: \mathcal{A}_t .

In each period t, the agent takes one of a finite number of actions $a_t \in A_t$ and the number of possible actions in period t is denoted by K_t .

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 $s_t \in S_t$ 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

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 s_t given s_{t-1} and a_{t-1} . Incorporating dependencies across longer horizons requires redefining the state variable to include sufficient lags.

In addition, we assume (as is standard) that the vector ϵ_t is independently distributed across time. This rules out serially correlated unobserved heterogeneity, such as if agents' decisions were driven by shocks, unobserved to the econometrician, whose effects last for multiple periods. This limitation can be mitigated in two ways. First, one can allow for considerable heterogeneity across time and across agents by permitting time- and type-varying preferences (see below for details). Second, one can include a large number of observed time-varying variables in the state space, thereby reducing the serial correlation of the unobserved residual. We also assume that the preference shock ϵ_t has a known distribution and is independent of the whole path of observable state variables $\{x_t, z_t\}_{t=1}^3$ as well as past actions $\{a_s\}_{s=1}^{t-1}$. This rules out for instance direct feedback from current shocks to future state variables. We deal with this limitation by including a set of state variables and directly modeling their evolution over time. Formally, **ASSUMPTION 2.** (Independence and Known Error Distribution) $\mathbb{P}(x_t, z_t, \epsilon_t | x_{t-1}, z_{t-1}, \epsilon_{t-1}, a_{t-1}) = \mathbb{P}(x_t, z_t | x_{t-1}, z_{t-1}, a_{t-1})\mathbb{P}(\epsilon_t)$, where the distribution of the vector ϵ_t is known and is absolutely continuous on \mathbb{R}^{K_t} w.r.t. Lebesgue measure and independently distributed across t.

DECISION RULES: d_t .

The decision rule in period t, d_t , is a mapping from S_t to A_t . We do not allow history dependent decision rules so decision rules cannot be mappings from $\prod_{s=1}^{t-1} (S_s, A_s) \times S_t$ 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 (see e.g. Rust 1994). 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).⁴ Then, for a given sequence of actions $\{a_t\}_{t=1}^3$, the utility of an agent of type τ is:

$$\tilde{u}_t(s_t, a_t; \tau) + \beta_\tau \sum_{j=t+1}^3 \delta_\tau^{j-t} \mathbb{E}_t(\tilde{u}_j(s_j, a_j; \tau)).$$
(2)

Broadly, we deal with three types of agents: time-consistent agents (denoted by τ_C), time-inconsistent naïve agents (τ_N) and time-inconsistent sophisticated agents (τ_S) with two important qualifications: (a) within each type, we can allow for further heterogeneity in per-period and time preferences so that there could be multiple though finite consistent, sophisticated, and naïve types; (b) we can allow for partially sophisticated agents and provide set identification results for this case.

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 (2) using standard dynamic programming methods (backward induction in this finite horizon case). The other two types of agent are both time-inconsistent, with hyperbolic parameter $\beta_{\tau} < 1$. Both types of time-inconsistent agents are aware of their *current* present-bias and solve the maximization problem using backward induction. However, while sophisticated agents ($\tau = \tau_S$, $\beta_{\tau_S} < 1$) also recognize their *future* present-bias, naïve agents ($\tau = \tau_N$, $\beta_{\tau_N} < 1$) do not. For the econometrician, these differences will generate key identification issues that we address below.

The formulation in (2) allows for type-varying exponential and hyperbolic parameters. Previous empirical work assumes that $\beta_{\tau_N} = \beta_{\tau_S}$ and that $\delta_{\tau_N} = \delta_{\tau_S} = \delta_{\tau_C}$. We substantially relax these restrictions while still retaining point-identification for all parameters except the time-preference parameters for naïve agents. Finally, even within a type, our results allow for heterogeneity in the time-preference parameters so that for instance there could exist multiple sophisticated types each with a different value for their time-preferences β_{τ} and δ_{τ} or their per period utility.

The formulation in (2) also allows for time-varying type-specific per-period utilities $\tilde{u}_t(\cdot; \tau)$. This flexibility is important since we can then examine heterogeneity across three dimensions. First, within a given type one can assess how much of the difference in behavior across time can be attributed to evolving preferences over states and how much to time-preference issues. This matters because imposing

⁴This is not the only possible formulation: see for instance Gul and Pesendorfer (2001, 2004). See Toussaert (2018) for an experimental test of the Gul-Pesendorfer model and Giné et al. (2018) for a field-experimental test of commitment revisions.

time-invariant utilities may be too restrictive for some applications and this misspecification may overattribute changes in behavior over time to time-preference parameters. Second, holding time fixed, one can examine how much of the difference in behavior across types is driven by differing preferences over states versus different rates of time-preference. Third, the time- and type-varying formulation provides a mechanism for flexibly accounting for serially correlated unobserved heterogeneity.

This formulation nests the typical model where types only differ in the degree of present-bias so that we can gauge what the role of present bias is relative to alternative factors in explaining behavior. Next, we impose the exclusion restriction that the state variable z_t does not enter the per-period utility function.

ASSUMPTION 3. (Exclusion Restriction) $\tilde{u}_t(x_t, z_t, \epsilon_t, a_t; \tau) = \tilde{u}_t(x_t, \epsilon_t, a_t; \tau)$.

Intuitively, this restriction provides the basis for variation in the forward-looking component of the value function while keeping current period utility constant. It builds on the ideas (though not the precise assumption) in Magnac and Thesmar (2002) and is also used by Abbring and Daljord (2018) and Abbring et al. (2019). Abbring and Daljord (2018) point out that the Magnac and Thesmar (2002, Section 4.2) assumption imposes conditions on the value function (rather than the per-period utility) and is therefore not straightforward to interpret. In our context, elicited beliefs about the future evolution of state variables are a natural candidate for the exclusion restriction. The elicitation and use of expectational and belief data, as proposed most forcefully by Manski (2004), is becoming increasingly common, including in Development (Delavande et al., 2010; Delavande, 2014), Finance (Shleifer, 2019) and Macroeconomics Roth and Wohlfart (2018). Finally, we assume that within each period, utility is additively separable in the unobserved state variables.

ASSUMPTION 4. (Additive Separability) For each type $\tau \in \mathcal{T}$ the utility function in period $t \in \{1, 2, 3\}$ is given by $\tilde{u}_t(x_t, \epsilon_t, a_t; \tau) = u_t(x_t, a_t; \tau) + \epsilon_t(a_t)$.

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 will omit them as arguments in preferences for the most part. The empirical section contains a discussion of which variables are included in v for the estimation.

3 Identification

We consider both the case where types are directly observed as well as the case where they are not. 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 the first case we require that the researcher directly identifies the type for each individual by observing variables referred to as a type indicator or type proxy (collectively denoted by $r \in \mathcal{R}$). In the second case, we assume that r only imperfectly reveals the agent's type, for instance due to the agents' imperfect understanding of the choice problem, or to imperfectly chosen survey instruments or other differing circumstances of the agents.

3.1 Directly Observed Types

We observe an *i.i.d.* sample on $(\{a_t^*, x_t, z_t\}_{t=1}^T, w)$ where a_t^* is the optimal action chosen by the agent, (x_t, z_t) are observed state variables and w = (r, v) includes both the type proxy $r \in \mathcal{R}$ and other time-

invariant characteristics v. We set T = 3 because at least three periods are necessary to capture the notions of time-inconsistency popular in the literature (with only two periods, no time-inconsistency problem would arise), and extensions to a general T are straightforward. We allow for different specifications of \mathcal{R} : in the simplest case $\mathcal{R} = \{r_{\tau_C}, r_{\tau_S}, r_{\tau_N}\}$ so that each element corresponds to a unique time-preference type and there is no heterogeneity within each time-preference type. Alternatively, we also allow for sub-types within a particular class of time-inconsistent preferences, in which case $\mathcal{R} = \{r_{\tau_{C_1}}, \ldots, r_{\tau_{C_J}}, r_{\tau_{S_1}}, \ldots, r_{\tau_{S_K}}, r_{\tau_{N_1}}, \ldots, r_{\tau_{N_L}}\}$. This, for instance, allows for multiple time-consistent agents who may differ in their rates of time preferences or their per-period utility functions. In this section types are directly observed so we could equivalently have stated this as a condition on the set of possible types, but we prefer this formulation because it provides a natural generalization to the case with unobserved types.

The key starting point for identification are the type-specific choice probabilities $\mathbb{P}_{\tau}(a_t^* = a | x_t, z_t)$. These are directly observed since type is a known deterministic function of the observed type proxy r.

ASSUMPTION 5. (Directly Observed Types) Choice probabilities are directly observed for each type. For an agent of type τ , $\mathbb{P}_{\tau}(a_t^* = a | x_t, z_t) = \mathbb{P}(a_t^* = a | x_t, z_t, r = r_{\tau}).$

We now turn to identification of the preference parameters using backward induction, and we thus start from the terminal period. Note that for the initial set of results we do not need to know whether the identified type is consistent, naïve, or sophisticated; and this generality will prove useful in the sequel.

In the terminal period the agent will choose action k if and only if $\tilde{u}_3(s_3, k; \tau) > \tilde{u}_3(s_3, a; \tau) \quad \forall a \neq k$. Under Assumptions 2 (independence), 4 (additive separability) and 5 (observed types) we can write

$$\mathbb{P}_{\tau}(a_3^* = k | x_3, z_3) = \mathbb{P}\left(k = \operatorname*{argmax}_{a \in \mathcal{A}_3} \left\{ u_3(x_3, a; \tau) + \epsilon_3(a) \right\} | x_3, z_3 \right).$$

The decision in the terminal period is a standard static discrete choice model and with a known error distribution we can invert the relationship to directly identify the functions $u_3(x_3, k; \tau) - u_3(x_3, 0; \tau)$, where $a_3 = 0$ is the base action (see online Appendix C for an alternative and self-contained argument or Hotz and Miller, 1993). We can thus identify period 3 preferences subject to the following normalization:

ASSUMPTION 6. (Normalization) Utility in period 3 for a base action $a_3 = 0$ is known for all types and for all states, i.e. $u_3(x_3, 0; \tau)$ is known for all $x_3 \in \mathcal{X}_3$ for all types.

Recent work has emphasized the importance of such assumptions for counterfactual analysis. The main focus of our paper is on distinguishing types and time-preferences. One example of how to relax this assumption is given in Appendix A of Abbring and Daljord (2018) who consider a time-consistent model where the utility normalization is known only up to a constant shift.

We next consider the decision problem in period 2. In this period, the probability that an agent chooses action k is given by

$$\mathbb{P}_{\tau}(a_{2}^{*}=k|x_{2},z_{2}) = \mathbb{P}\left(k = \operatorname*{argmax}_{a \in \mathcal{A}_{2}} \left\{u_{2}(x_{2},a;\tau) + \epsilon_{2}(a) + \beta_{\tau}\delta_{\tau}\int v_{\tau,3}^{*}(s_{3})\,\mathrm{dF}(s_{3}|x_{2},z_{2},a)\right\} \middle| x_{2},z_{2}\right),\$$

where $dF(s_{t+1}|x_t, z_t, a)$ is a transition probability for s_{t+1} when action a is chosen in state (x_t, z_t) and

$$v_{\tau,3}^*(s_3) = \max_{a \in \mathcal{A}_3} \left\{ u_3(x_3, a; \tau) + \epsilon_3(a) \right\}.$$
 (3)

We can then use the standard Hotz-Miller inversion of the type-specific conditional choice probabilities to directly identify

$$g_{\tau,2,k}(x_2, z_2) \equiv u_2(x_2, k; \tau) - u_2(x_2, 0; \tau) + \beta_\tau \delta_\tau \int v_{\tau,3}^*(s_3) \,\mathrm{dF}_{\Delta,k}(s_3 | x_2, z_2) \tag{4}$$

where we define the signed measure $dF_{\Delta,k}(s_3|x_2, z_2) \equiv dF(\epsilon_3) (dF(x_3, z_3|x_2, z_2, k) - dF(x_3, z_3|x_2, z_2, 0))$. We then explore which of the unknown elements on the right hand side of (4)—the utility functions and the discount rates—can be identified using knowledge of the function $g_{\tau,2,k}(\cdot)$. First, the integral in (4) is directly identified since (a) $u_3(\cdot)$ is identified and (b) the distribution of ϵ_3 is known and $dF(x_3, z_3|x_2, z_2, k)$ is directly identified from the data so that the signed measure $dF_{\Delta,k}(s_3|x_2, z_2)$ is identified. Next, z_2 only enters the last term in (4) so that we can use variation in z_2 (conditional on x_2) to separately identify each term. This requires sufficient variation in z_2 and that this variation translates into variation in the expectation of the period 3 value function $v_{\tau,3}^*(\cdot)$. This is the rank condition in Proposition 4 of Magnac and Thesmar (2002) adapted to the context of our model.

ASSUMPTION 7. For all types, the distribution of z_2 conditional on x_2 has at least two points of support (z'_2, z''_2) such that the directly identified object $\int v^*_{\tau}(s_3) \left(dF_{\Delta,k}(s_3|x_2, z'_2) - dF_{\Delta,k}(s_3|x_2, z''_2) \right)$ is not equal to zero for at least one action k and one point x_2 .

With this condition, however, we can separately identify the per-period preferences and a function of the individual time-preference parameters. Finally, we need the usual normalization for t = 2 preferences:

ASSUMPTION 8. Utility for a base action $a_2 = 0$, $u_2(x_2, 0; \tau)$, is known for all types and states.

We can now state the first identification result.

LEMMA 1. Consider an agent of type τ maximizing (2) and suppose that the model satisfies Assumptions 1–8. Then

- 1. Period 3 utility $u_3(x_3, a_3; \tau) \quad \forall (a_3 \in \mathcal{A}_3, x_3 \in \mathcal{X}_3)$ is identified.
- 2. Period 2 utility $u_2(x_2, a_2; \tau) \quad \forall (a_2 \in \mathcal{A}_2, x_2 \in \mathcal{X}_2)$ is identified.
- 3. The product of the exponential parameter and the hyperbolic parameter $\{\beta_{\tau}\delta_{\tau} : \tau \in \mathcal{T}\}$ is identified and if $\tau = \tau_C$, $\beta_{\tau_C}\delta_{\tau_C} = \delta_{\tau_C}$.

All proofs are relegated to the appendix. We next provide results for identifying the time-preference parameters separately as well as the period 1 utility functions. As before, the (conditional) probability that an agent chooses action k in period 1 is given by

$$\mathbb{P}_{\tau}(a_{1}^{*}=k|x_{1},z_{1}) = \mathbb{P}\left(k = \operatorname*{argmax}_{a \in \mathcal{A}_{1}} \left\{u_{1}(x_{1},a;\tau) + \epsilon_{1}(a) + \beta_{\tau}\delta_{\tau}\int v_{\tau,2}^{*}(s_{2})\,\mathrm{dF}(s_{2}|x_{1},z_{1},a)\right\} \middle| x_{1},z_{1}\right)$$

The key difference between standard and hyperbolic dynamic programming problems is captured in the definition of the value function $v_{\tau,2}^*(s_2)$ which is defined as

$$v_{\tau,2}^{*}(s_{2}) = \sum_{a \in \mathcal{A}_{2}} v_{\tau,2}(s_{2}, a, 1, \delta_{\tau}) A_{\tau}(s_{2}, a, \tilde{\beta}_{\tau}, \delta_{\tau}), \text{ where}$$

$$v_{\tau,2}(s_{2}, a, \mathsf{b}, \mathsf{d}) = u_{2}(x_{2}, a; \tau) + \epsilon_{2}(a) + \mathsf{bd} \int v_{\tau,3}^{*}(s_{3}) \, \mathrm{dF}(s_{3}|x_{2}, z_{2}, a), \text{ and}$$

$$A_{\tau}(s_{2}, a, \tilde{\beta}_{\tau}, \delta_{\tau}) = \mathbb{I} \left\{ a = \operatorname*{argmax}_{j \in \mathcal{A}_{2}} v_{\tau,2}(s_{2}, j, \tilde{\beta}_{\tau}, \delta_{\tau}) \right\}.$$
(5)

The quantity $v_{\tau,2}^*(s_2)$ is the continuation value from period 2 onwards from the standpoint of period 1, assuming that the choice of action a in period 2 is governed by the indicator function $A_{\tau}(s_2, a, \tilde{\beta}_{\tau}, \delta_{\tau})$. In standard dynamic programming problems—i.e for time-consistent agents—the corresponding event would be given by $A_{\tau}(s_2, a, 1, \delta_{\tau})$. The function $v_{\tau,2}(s_2, a, \mathbf{b}, \mathbf{d})$ is the value of taking action a in state s_2 given a value of \mathbf{b} for the hyperbolic parameter and a value of \mathbf{d} for the exponential parameter. We introduce \mathbf{b}, \mathbf{d} explicitly as arguments since the identification results will involve evaluating this value function at different candidate values for the hyperbolic and exponential parameters.

The parameter $\tilde{\beta}_{\tau}$ is interpreted as the extent of present-bias that the agent in period 1 thinks his period 2 self will be subject to. For time-consistent agents $\tilde{\beta}_{\tau} = 1$ but for time-inconsistent agents in general $\tilde{\beta}_{\tau} \leq 1$. The value of this parameter is often mapped into notions of "sophistication" in the time-discounting literature. Time-inconsistent agents with values of $\tilde{\beta}_{\tau}$ that are close to β_{τ} are said to exhibit greater "sophistication" since they recognize more clearly the extent of the present-bias in their future behaviour while values of $\tilde{\beta}_{\tau}$ further from β_{τ} and closer to 1 reflect more "naivete" (since agents are failing to recognize the true extent of present-bias in their future behavior). For the main results in this paper we make the assumption that agents are either completely sophisticated or completely naive.

ASSUMPTION 9. The parameter $\hat{\beta}_{\tau}$ is equal to 1 for consistent and naïve agents and is equal to β_{τ_S} for sophisticated agents.

A weaker assumption, that of partial sophistication, would be that $\tilde{\beta}_{\tau} \in [\beta_{\tau}, 1]$ and then $\tilde{\beta}_{\tau}$ would be an additional parameter for type τ in the model. This is explored in Section 3.3 where we show that the partial sophistication parameter is typically only set-identified. Intuitively, with partial sophistication observed choices in later periods are effectively uninformative about $\tilde{\beta}_{\tau}$, which drives choice in period 1. This is because choice in period 2 involves β_{τ} not $\tilde{\beta}_{\tau}$. However, observed choice in period 2 is informative about $\beta_{\tau} = \beta_{\tau}$ for sophisticated agents.

We can now invert as usual the type-specific conditional choice probabilities to directly identify

$$g_{\tau,1,k}(x_1, z_1) = u_1(x_1, k; \tau) - u_1(x_1, 0; \tau) + \beta_\tau \delta_\tau \int h_\tau(x_2, z_2) \,\mathrm{dF}_{\Delta,k}(x_2, z_2 | x_1, z_1), \tag{6}$$

where

$$h_{\tau}(x_{2}, z_{2}) \equiv \int v_{\tau,2}^{*}(s_{2}) \,\mathrm{dF}(\epsilon_{2}) \\ = \sum_{a \in \mathcal{A}_{2}} \int \left(u_{2}(x_{2}, a; \tau) + \epsilon_{2}(a) + \delta_{\tau} \int v_{\tau,3}^{*}(s_{3}) \,\mathrm{dF}(s_{3}|x_{2}, z_{2}, a) \right) A_{\tau}(s_{2}, a, \tilde{\beta}_{\tau}\delta_{\tau}, 1) \,\mathrm{dF}(\epsilon_{2}).$$
(7)

We write $A_{\tau}(s_2, a, \beta_{\tau} \delta_{\tau}, 1)$ instead of the equivalent formulation $A_{\tau}(s_2, a, \beta_{\tau}, \delta_{\tau})$ to emphasize the insight that period 3 choices are only informative about the product $\tilde{\beta}_{\tau} \delta_{\tau}$ (for some types) and this information is used to identify $A_{\tau}(\cdot)$. Note also that in the expression above the expected value of $v_{\tau,3}^*(s_3)$ multiplies the discount factor δ_{τ} and not $\beta_{\tau} \delta_{\tau}$, because the hyperbolic parameter β_{τ} does not directly enter into the intertemporal rate of substitution between any two *future* periods (in this case, t = 2, 3) when seen from the point of view of the present (t = 1).

The function $h_{\tau}(x_2, z_2)$ represents how an agent in period 1 views the future after incorporating predictions about his own future behavior. Below, we show that h_{τ} is identified (up to a normalization) using only period 1 data. However, we will show that an analogous function \tilde{h} can also be constructed using data from periods 2 and 3, and a comparison of these two functions will provide a measure of the difference between an agent's beliefs at t = 1 about his future behavior and his actual choices. This difference is useful in allowing us to separately identify the exponential and the hyperbolic parameter in a number of important cases. We make three distinct arguments. First, we note that for consistent agents the last term in (6) is identified and hence period 1 preferences are identified (up to a normalization) without any further assumptions. Second, we provide conditions under which all the time-preference parameters for inconsistent sophisticated types are identified. Finally, we discuss partial and point identification results for naïve types under a set of stronger conditions. As in the previous period, we will require sufficient exogenous variation in z_1 to isolate the last term in (6).

ASSUMPTION 10. The distribution of z_1 conditional on x_1 has at least two points of support (z'_1, z''_1) and there exists at least one action k and one point in the support of \mathcal{X}_1 such that the identified object $\int h_{\tau}(x_2, z_2) \left(dF_{\Delta,k}(x_2, z_2 | x_1, z'_1) - dF_{\Delta,k}(x_2, z_2 | x_1, z''_1) \right) \neq 0.$

Next, we define the following function, where b denotes a given value of the identified product $\beta_{\tau}\delta_{\tau}$:

$$\begin{split} \tilde{h}_{\tau}(x_{2}, z_{2}, \mathsf{b}, \mathsf{d}) &\equiv \int \sum_{a \in \mathcal{A}_{2}} v_{\tau, 2}(s_{2}, a, 1, \mathsf{d}) A_{\tau}(s_{2}, a, \mathsf{b}, 1) \, \mathrm{dF}(\epsilon_{2}) \\ &= \sum_{a \in \mathcal{A}_{2}} \int \left(u_{2}(x_{2}, a; \tau) + \epsilon_{2}(a) + \mathsf{d} \cdot \int v_{\tau, 3}^{*}(s_{3}) \, \mathrm{dF}(s_{3} | x_{2}, z_{2}, a) \right) A_{\tau}(s_{2}, a, \mathsf{b}, 1) \, \mathrm{dF}(\epsilon_{2}) \\ &= \sum_{a \in \mathcal{A}_{2}} \int \left(u_{2}(x_{2}, a; \tau) + \epsilon_{2}(a) \right) A_{\tau}(s_{2}, a, \mathsf{b}, 1) \, \mathrm{dF}(\epsilon_{2}) + \mathsf{d} \cdot \sum_{a \in \mathcal{A}_{2}} q_{\tau}(x_{2}, z_{2}, a) \int A_{\tau}(s_{2}, a, \mathsf{b}, 1) \, \mathrm{dF}(\epsilon_{2}) \\ &\equiv \tilde{h}_{\tau, 1}(x_{2}, z_{2}, \mathsf{b}) + \mathsf{d} \cdot \tilde{h}_{\tau, 2}(x_{2}, z_{2}, \mathsf{b}), \end{split}$$
(8)

where $A_{\tau}(s_2, a, \mathbf{b}, 1)$ is defined in (5) and

$$q_{\tau}(x_2, z_2, a) \equiv \int v_{\tau,3}^*(s_3) \,\mathrm{dF}(s_3 | x_2, z_2, a).$$
(9)

The function $h_{\tau}(\cdot)$ is identified up to the arguments (\mathbf{b}, \mathbf{d}) since $u_2(\cdot), q_{\tau}(\cdot)$ and the indicator $A_{\tau}(s_2, a, \mathbf{b}, 1)$ are all identified up to (\mathbf{b}, \mathbf{d}) from previous arguments and the distribution of ϵ_2 is known.

3.1.1 Identification of Time Preferences for Consistent and Sophisticated Agents

Recall that for sophisticated agents, $\tilde{\beta}_{\tau_S} = \beta_{\tau_S}$ by Assumption 9. This implies that (a) $\tilde{\beta}_{\tau_S} \delta_{\tau_S} = \beta_{\tau_S} \delta_{\tau_S}$ which is identified and (b) δ_{τ_S} is the only unknown on the right hand side of (7) and enters the function linearly. Formally, consider the following transformation of equation (8) evaluated at $\mathbf{b} = \beta_{\tau_S} \delta_{\tau_S}$:

$$\int \tilde{h}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}, \mathsf{d}) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1') - \mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ = \int \left(\tilde{h}_{\tau, 1}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) + \mathsf{d}\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) \right) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1') - \mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ \leq \int \left(\tilde{h}_{\tau, 1}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) + \mathsf{d}\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) \right) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1') - \mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ \leq \int \left(\tilde{h}_{\tau, 1}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) + \mathsf{d}\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) \right) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') - \mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ \leq \int \left(\tilde{h}_{\tau, 1}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) + \mathsf{d}\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) \right) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') - \mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ \leq \int \left(\tilde{h}_{\tau, 1}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) + \mathsf{d}\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) \right) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') - \mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ \leq \int \left(\tilde{h}_{\tau, 1}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) + \mathsf{d}\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) \right) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') - \mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ \leq \int \left(\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) + \mathsf{d}\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) \right) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') - \mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ \leq \int \left(\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) + \mathsf{d}\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) \right) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ \leq \int \left(\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) + \mathsf{d}\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) \right) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ \leq \int \left(\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) \right) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ \leq \int \left(\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) \right) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ \leq \int \left(\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) \right) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ \leq \int \left(\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}) \right) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \\ \leq \int \left(\tilde{h}_{\tau, 2}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S})$$

Then, as long as the object multiplying d in the expression above is non-zero, we can identify δ_{τ_S} . We thus make the following assumption.

ASSUMPTION 11.

 $\sum_{a \in \mathcal{A}_2} \int q_{\tau_S}(x_2, z_2, a) A_{\tau_S}(s_2, a, \beta_{\tau_S} \delta_{\tau_S}, 1) \, \mathrm{dF}(\epsilon_2) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1') - \mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right) \neq 0 \text{ for the same action } k \text{ and points } (x_1, z_1', z_1'') \text{ as in Assumption } 10.$

As before, this rank condition is testable since it consists of directly identified quantities. Finally, we need a normalization for second-period utility.

ASSUMPTION 12. Utility for a base action (action 0) is known for all types for all states, i.e. $u_1(x_1, 0; \tau)$ is known for all $x_1 \in \mathcal{X}_1$ for all types.

LEMMA 2. Consider an agent of type τ_C solving the problem (2) at t = 1 and suppose that the model satisfies Assumptions 1- 8 and Assumption 12. Then, $u_1(x_1, a; \tau_C)$ is identified $\forall (a \in \mathcal{A}_1, x_1 \in \mathcal{X}_1)$. Next, consider an agent of type τ_S solving the problem (2) at t = 1 and suppose that the model satisfies Assumptions 1- 12. Then,

- 1. Period 1 utility $u_1(x_1, a; \tau_S)$ is identified $\forall (a \in \mathcal{A}_1, x_1 \in \mathcal{X}_1)$.
- 2. The exponential and hyperbolic parameters (δ_{τ_S} and β_{τ_S}) for sophisticated agents are identified.

3.1.2 Identification of Time Preferences for naïve Agents

For naïve agents, observed choices in periods 2 and 3 are informative about $\beta_{\tau_N} \delta_{\tau_N}$, but unlike for the other two types, this information is not enough to point-identify $\tilde{\beta}_{\tau_N} \delta_{\tau_N} (= \delta_{\tau_N})$. This introduces complications and point-identification requires much stronger assumptions than for the other types.

In order to obtain reasonably sharp partial identification results we place stronger assumptions on the transition probabilities that allow us to point identify the function $h_{\tau}(x_2, z_2)$ up to a normalization. Since these assumptions will play a key role in the identification argument when types are unknown we state them for a general type τ though in this subsection we only invoke them for naïve types.

3.1.3 Strengthening Variation in Transitions

In order to identify $h_{\tau}(x_2, z_2)$ we need stronger assumptions on the transition probabilities. In particular, we require that there exist at least two actions in period 1 such that the resulting transition probabilities (for period 2) are sufficiently different from each other. In addition, we require that a certain function of the differences in transition probabilities (across actions) is sufficiently variable in z_1 . These assumptions, although strong, are directly testable since they are placed on observable quantities.

To formalize these notions, we need to introduce some notation. Let $\{x_{s,j}, z_{s,j}\}$ be elements of $\mathcal{X}_s \times Z_s$ and define the probabilities, all of which are identified:

$$\mathrm{dF}_k(x_{s,j}, z_{s,j} | x_{s-1,j'}, z_{s-1,j'}) \equiv \mathbb{P}(x_s = x_{s,j}, z_s = z_{s,j} | x_{s-1} = x_{s-1,j'}, z_{s-1} = z_{s-1,j'}, a_{s-1} = k).$$

Let S denote the cardinality of $\mathcal{X}_2 \times \mathcal{Z}_2$. Define the matrix $\mathsf{dF}(k, z_1)$ as follows:

$$\mathsf{dF}(k, z_1) \equiv \begin{bmatrix} \mathrm{dF}_k(x_{2,1}, z_{2,1} | x_{1,1}, z_1) & \dots & \mathrm{dF}_k(x_{2,S-1}, z_{2,S-1} | x_{1,1}, z_1) \\ \vdots & \vdots & \vdots \\ \mathrm{dF}_k(x_{2,1}, z_{2,1} | x_{1,S-1}, z_1) & \dots & \mathrm{dF}_k(x_{2,S-1}, z_{2,S-1} | x_{1,S-1}, z_1) \end{bmatrix}.$$
(10)

The elements of this matrix are the transition probabilities for all (but one) possible values of the period 2 observed state variables conditional on S-1 possible values of the period 1 state variable x_1 , for period 1 action k. Next, define the matrix $dF_{\Delta}(k, z_1) \equiv dF(k, z_1) - dF(0, z_1)$ for the action pair (k, 0).

ASSUMPTION 13. \mathcal{X}_1 has at least S-1 points of support where S is the cardinality of $\mathcal{X}_2 \times \mathcal{Z}_2$. The $(S-1) \times (S-1)$ identified matrix $(\mathsf{dF}_\Delta(k, z'_1) - \mathsf{dF}_\Delta(k, z''_1))$ is invertible for some action k and two points $z'_1 \neq z''_1$.

This assumption is restrictive in that it requires that the number of possible states in period 1 is at least as large as the cardinality of $\mathcal{X}_2 \times \mathcal{Z}_2$. For instance, this assumption would fail if the support of x in period 2 is larger than in period 1 (or equal, if there are at least two points of support for z_2). It is possible, however, to relax this assumption by placing restrictions on the support of $\mathcal{A}_1 \times \mathcal{X}_1$ (this approach is presented in Appendix A as a modified version of Lemma 3 that follows below) and is useful when \mathcal{X}_1 is small. Assumption 13 can be viewed as a strengthening of Assumption 10 in that it imposes sufficient variation in the transition probabilities for many more points in the support of \mathcal{X}_1 than Assumption 10 and further that the vector of transitions across points in \mathcal{X}_1 be linearly independent in the precise sense specified above. With this additional assumption we can recover the function $h_{\tau}(\cdot)$ up to a location shift:

LEMMA 3. Consider an agent of type τ maximizing equation (2) and suppose that the model satisfies Assumptions 1–9 and Assumptions 12 and 13. Then

- 1. The function $h_{\tau}(x_2, z_2)$ is identified up to a constant k_{τ} for all types τ and $(x_2, z_2) \in \mathcal{X}_2 \times \mathcal{Z}_2$.
- 2. Period 1 utility $u_1(x_1, a; \tau)$ is identified $\forall (a \in \mathcal{A}_1, x_1 \in \mathcal{X}_1, \tau \in \mathcal{T}).$

We apply the previous lemma for naïve agents to identify the function $h_{\tau_N}(x_2, z_2) + k_{\tau_N}$. The proof shows that $k_{\tau_N} = -h_{\tau_N}(x_{20}, z_{20})$ for some pre-specified point $(x_{20}, z_{20}) \in \mathcal{X}_2 \times \mathcal{Z}_2$. Next, recalling the definitions of $h_{\tau}(x_2, z_2)$ in equation (7) and $h_{\tau}(x_2, z_2, \mathbf{b}, \mathbf{d})$ in (8) we define two functions and note that they are related as:

$$h_{\tau_N}^{\Delta}(x_2, z_2) \equiv h_{\tau_N}(x_2, z_2) - h_{\tau_N}(x_{20}, z_{20})$$
(11)

$$=\tilde{h}_{\tau_N}(x_2, z_2, \delta_{\tau_N}, \delta_{\tau_N}) - \tilde{h}_{\tau_N}(x_{20}, z_{20}, \delta_{\tau_N}, \delta_{\tau_N})$$

$$(12)$$

$$\equiv \tilde{h}^{\Delta}_{\tau_N}(x_2, z_2, \delta_{\tau_N}). \tag{13}$$

where the third argument in \tilde{h}_{τ_N} is δ_{τ_N} because for naïve individuals $\tilde{\beta}_{\tau_N} = 1$. If the function $\tilde{h}_{\tau_N}^{\Delta}(x_2, z_2, d)$ were one-to-one in d, then one could recover δ_{τ_N} by inverting the function at the point $h_{\tau_N}^{\Delta}(x_2, z_2)$ for a given (x_2, z_2) or by carrying out a minimum distance type strategy. Unfortunately, the function is not in general one-to-one in d and the corresponding minimum distance function will not be uniquely minimized at δ_{τ_N} so the parameter is not point-identified. We therefore begin by defining the identified set for δ_{τ_N} as all those values of d that are consistent with the identified function $h_{\tau_N}^{\Delta}(x_2, z_2)$ defined in (11) above:

$$\Theta_{\delta_N} \equiv \left\{ \mathsf{d} \in (\beta_{\tau_N} \delta_{\tau_N}, 1) : \tilde{h}^{\Delta}_{\tau_N}(x_2, z_2, \mathsf{d}) = h^{\Delta}_{\tau_N}(x_2, z_2) \quad \forall \ (x_2, z_2) \in \mathcal{X}_2 \times \mathcal{Z}_2 \right\}.$$

This leads to a corresponding identified set for β_{τ_N} :

$$\Theta_{\beta_N} \equiv \left\{ \frac{\beta_{\tau_N} \delta_{\tau_N}}{\mathsf{d}} : \mathsf{d} \in \Theta_{\delta_N} \right\}.$$

We can now state the identification results for naïve types.

LEMMA 4. Consider a time-inconsistent naïve agent solving the problem (2) at t = 1 and suppose that the model satisfies Assumptions 1– 9 and Assumptions 12 and 13. Then, the identified set for δ_{τ_N} is given by Θ_{δ_N} and the identified set for β_{τ_N} is given by Θ_{β_N} .

The proof also shows that the bounds above are sharp. General conditions for point identification are not always available. Here, we outline one relatively straightforward (and testable) assumption yielding point identification by requiring that $\tilde{h}_{\tau_N}^{\Delta}(x_2, z_2, \mathsf{d})$ is strictly monotone in d .

ASSUMPTION 14. There exists $(x_2, z_2) \in \mathcal{X}_2 \times \mathcal{Z}_2$ such that the following difference always has the same sign $\forall d \in \Theta_{\delta_N}$: $\sum_{a \in \mathcal{A}_2} (q_{\tau_N}(x_2, z_2, a) \int A_{\tau}(s_2, a, d, 1) dF(\epsilon_2) - q_{\tau_N}(x_{20}, z_{20}, a) \int A_{\tau}(s_2, a, d, 1) dF(\epsilon_2))$. Recall that $q_{\tau_N}(x_2, z_2, a)$ is the (not discounted) expected period 3 utility when the agent (a) contemplates being in state (x_2, z_2) in period 2, (b) takes action $a \in \mathcal{A}_2$ and (c) assumes he behaves optimally in the static period 3 problem. The integral over $A_{\tau}(\cdot)$ is the probability that action a is optimal in period 2 when the discount rate between period 2 and period 3 is d. The assumption above ensures that the derivative of $\tilde{h}_{\tau_N}^{\Delta}(x_2, z_2, d)$ with respect to d is strictly positive (or negative) everywhere in d at least for some value of the state variable. Roughly speaking, consider a situation where one treats (x_2, z_2) and (x_{20}, z_{20}) as being fixed and instead treats the action as a random variable with the probability of a occurring being the integral over $A_{\tau}(\cdot)$. Then, the equation above states that expected utility in state (x_2, z_2) is larger than in state (x_{20}, z_{20}) no matter what discount rate (in the identified set) the agent uses. With this additional assumption, we can separately identify both time preference parameters for naïve types. **LEMMA 5.** Consider a time-inconsistent naïve agent solving the problem (2) at t = 1 and suppose that the model satisfies Assumptions 1– 9 and Assumptions 12- 14. Then δ_{τ_N} and β_{τ_N} are identified.

Another assumption trivially allows point-identification by requiring that the exponential parameter for naïve type is the same as that for one of the other two (i.e. for the consistent or sophisticated) types:

ASSUMPTION 15. Time-inconsistent agents have the same exponential discount rate, $\delta_{\tau_S} = \delta_{\tau_N}$ or alternatively $\delta_{\tau_N} = \delta_{\tau_C}$

Since δ_{τ_S} and δ_{τ_C} are already identified, this assumption trivially guarantees identification of δ_{τ_N} . However, in order for this assumption to be substantive, both sophisticated and naïve types (or alternatively time-consistent and naïve types) have to exist. In other words, the time preferences of time-consistent and time-inconsistent sophisticated agents, respectively, can always be identified even if no naïve agents are present, while Assumption 15 states that when sophisticated and naïve time-inconsistent agents are both present and have the same per-period discount factor, this is sufficient to point identify also naïve agents' time preferences.

3.2 Unobserved Types

We next turn to the case where types are not directly observed. This is both a more realistic scenario (since observables typically do not completely reveal type) and a more general model because it nests the perfectly observed types model. The starting point is the joint distribution of $(\{a_t^*, x_t, z_t\}_{t=1}^T, w)$ but now without Assumption 5 so that we do not know which type an observation belongs to. Recall that w = (r, v) includes the type proxy r and other time-invariant characteristics v. There are now four steps involved in going from this observed joint distribution to the preference parameters for each type of agent:

- 1. Identify the total number of types.
- 2. Identify the type-specic choice probabilities, without assigning them to the types they refer to.
- 3. Assign the type-specific choice probabilities to the different types.
- 4. Identify the preference parameters for each type.

To illustrate, step 1 could determine that the population contains eight distinct types. Then step 2 identifies the eight sets of type-specific choice probabilities $\mathbb{P}_{\tau}(a_t|x_t, z_t, r, v)$, but leaving τ unknown. The specific type τ is next identified in step 3, while the preference parameters are identified in step 4. Note that a key implication of this more general approach is that while the *frequency of each type in the population* can be identified, the *type of each individual* cannot. This is in sharp contrast to the case discussed in Section 3.1 where the signal r directly identified the type for each individual.

We discuss each of the four steps above in a separate sub-section, although there is considerable overlap in the last two steps. We begin by introducing the additional elements needed to describe the unobserved types model starting from the mixture probabilities. The joint distribution of the observed data identifies the 'aggregate' choice probabilities $\mathbb{P}(a_t|x_t, z_t, r, v)$ which are mixtures over all of the typespecific choice probabilities. The mixture probabilities $\pi_{\tau}(r, v)$ denote the probabilities that an agent is of type τ conditional on a vector of exogenous variables v and the type proxy r. These probabilities have a substantive economic interpretation since they represent the relative sizes of the different types of agents in the population.

3.2.1 Identifying the Total Number of Types

Let \mathcal{T} denote the finite set of possible types. In the previous section under Assumption 9 we distinguished between completely sophisticated agents ($\tilde{\beta}_{\tau} = \beta_{\tau} < 1$), completely naïve agents ($\tilde{\beta}_{\tau} = 1, \beta_{\tau} < 1$) and time-consistent agents ($\tilde{\beta}_{\tau} = \beta_{\tau} = 1$). Within each type of agent we can allow for further subtypes—e.g. consistent agents with different preference parameters or sophisticated agents with different preference parameters—so the cardinality of \mathcal{T} can be larger than three. More generally, we can subsume all inconsistent agents under the rubric of partially sophisticated agents ($\tilde{\beta}_{\tau} \in [\beta_{\tau}, 1]$) within which type there might exist further sub-types with different preference parameters. We will show that whether all these types can be separately identified depends upon how different their behavior is both at a given point in the state space and across different points in the state space.

Define $M_{r,v}$ as the total number of types that exist at the support point (r, v):

$$M_{r,v} = \sum_{\tau \in \mathcal{T}} \mathbb{I}\{\pi_{\tau}(r,v) > 0\}$$
(14)

We first provide a lower bound on the total number of types that depends upon the size of the statespace. Under an additional (albeit unverifiable) assumption on the differential behavior of types across the state space this is also an upper bound so that we can identify the total number of types. To state these restrictions formally we begin by clarifying the link between observed choice probabilities and the underlying type-specific choice probabilities. For the purpose of identifying the type-specific choice probabilities $\mathbb{P}_{\tau}(a_t|x_t, z_t, r, v)$ there is no conceptual distinction between x_t and z_t so we denote their union by $\mathbf{x}_t \equiv (x_t, z_t)$. We place several restrictions on the distribution of states and actions:

ASSUMPTION 16. 1. Conditional upon type, r is uninformative about choice: $\mathbb{P}_{\tau}(a_1, \mathbf{x}_1 | r, v) = \mathbb{P}_{\tau}(a_1, \mathbf{x}_1 | v) \ \forall \ (a_1, \mathbf{x}_1, v), \ and \ \mathbb{P}_{\tau}(a_t | \mathbf{x}_t, r, v) = \mathbb{P}_{\tau}(a_t | \mathbf{x}_t, v) \ \forall \ (a_t, \mathbf{x}_t, v) \ t > 1.$

2. The transition probabilities for all points in the state space do not vary by type and are independent of $r: \mathbb{P}_{\tau}(\mathbf{x}_{t+1}|\mathbf{x}_t, a_t, r, v) = \mathbb{P}(\mathbf{x}_{t+1}|\mathbf{x}_t, a_t, v).$

The first part of the assumption is reasonable to the extent that r is only informative about choices through its predictive power for agent type. If, however, r provides information about other aspects of the decision process this assumption would fail. For instance, if r is not just a measure of time-inconsistency but also reflects a lack of numeracy or other flaws in an agent's cognitive processes it may have an independent effect on choice, even after conditioning on type. However, one can mitigate this problem by specifying a rich set of observables v (e.g. one could include a measure of literacy or cognitive skill if available in v). The second part of the assumption states that once we condition on action and current state, the evolution of future states is independent of type. This will be implausible if, for instance, different types take different unobserved actions that affect transition probabilities. One drawback of this assumption is that it is not testable since it imposes conditions on unobserved quantities (i.e., the type-specific choice and transition probabilities). Consider next the joint distribution of actions and states in two adjacent time periods $(a_t, a_{t+1}, \mathbf{x}_t, \mathbf{x}_{t+1})$ conditional upon (r, v) and express it as a mixture of the corresponding type distributions

$$\mathbb{P}(a_{t+1}, a_t, \mathbf{x}_{t+1}, \mathbf{x}_t | r, v) = \sum_{\tau \in \mathcal{T}} \pi_\tau(r, v) \mathbb{P}_\tau(a_{t+1}, a_t, \mathbf{x}_{t+1}, \mathbf{x}_t | r, v).$$

Using Assumption 16 and the Markov nature of the decision rule (see Section 2) we obtain

$$\mathbb{P}(a_{t+1}, a_t, \mathbf{x}_{t+1}, \mathbf{x}_t | r, v) = \sum_{\tau} \pi_{\tau}(r, v) \mathbb{P}_{\tau}(a_{t+1} | \mathbf{x}_{t+1}, v) \mathbb{P}(\mathbf{x}_{t+1} | \mathbf{x}_t, a_t, v) \mathbb{P}_{\tau}(a_t, \mathbf{x}_t | v).$$

Next, define (for $\mathbb{P}(\mathbf{x}_{t+1}|\mathbf{x}_t, a_t, v) \neq 0$) the directly identified quantity

$$\mathbf{F}_{\mathbf{x}_{t},\mathbf{x}_{t+1},r,v}^{a_{t},a_{t+1}} \equiv \frac{\mathbb{P}(a_{t},a_{t+1},\mathbf{x}_{t},\mathbf{x}_{t+1}|r,v)}{\mathbb{P}(\mathbf{x}_{t+1}|\mathbf{x}_{t},a_{t},v)}$$
$$= \sum_{\tau} \pi_{\tau}(r,v)\mathbb{P}_{\tau}(a_{t+1}|\mathbf{x}_{t+1},v)\mathbb{P}_{\tau}(a_{t},\mathbf{x}_{t}|v), \tag{15}$$

and, for an ordered pair (t, t + 1) also define the directly identified quantities

$$\mathbf{F}_{r,v}^{a_t,\mathbf{x}_t} \equiv \mathbb{P}(a_t, \mathbf{x}_t | r, v) = \sum_{\tau} \pi_{\tau}(r, v) \mathbb{P}_{\tau}(a_t, \mathbf{x}_t | v)$$
(16)

$$\mathbf{F}_{\mathbf{x}_{t+1},r,v}^{a_{t+1}} \equiv \mathbb{P}(a_{t+1}|\mathbf{x}_{t+1},r,v) = \sum_{\tau} \pi_{\tau}(r,v) \mathbb{P}_{\tau}(a_{t+1}|\mathbf{x}_{t+1},v).$$
(17)

Next, given a point (r, v), let $\underline{M}_{r,v}^t$ (\underline{M} for short) denote the cardinality of the smaller of the state spaces in two subsequent periods (min{ $\#\mathbf{X}_t, \#\mathbf{X}_{t+1}$ }). For given values $(a_t, \mathbf{x}_t^1, \dots, \mathbf{x}_t^M, a_{t+1}, \mathbf{x}_{t+1}^1, \dots, \mathbf{x}_{t+1}^M)$ define the ($\underline{M} + 1$) × ($\underline{M} + 1$) directly identified matrix

$$\mathsf{P}_{r,v}^{a_{t},a_{t+1},\underline{M}} \equiv \begin{pmatrix} 1 & \mathbf{F}_{r,v}^{a_{t+1},\mathbf{x}_{t+1}^{1}} & \cdots & \mathbf{F}_{r,v}^{a_{t+1},\mathbf{x}_{t+1}^{M}} \\ \mathbf{F}_{r,v}^{a_{t},\mathbf{x}_{t}^{1}} & \mathbf{F}_{\mathbf{x}_{t}^{1},\mathbf{x}_{t+1}^{1},r,v}^{a_{t},a_{t+1}} & \cdots & \mathbf{F}_{\mathbf{x}_{t}^{1},\mathbf{x}_{t+1}^{M},r,v}^{a_{t},a_{t+1}} \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{F}_{r,v}^{a_{t},\mathbf{x}_{t}^{M}} & \mathbf{F}_{\mathbf{x}_{t}^{a_{t},a_{t+1}}}^{a_{t},a_{t+1}} & \cdots & \mathbf{F}_{\mathbf{x}_{t}^{M},\mathbf{x}_{t+1}^{M},r,v}^{a_{t},a_{t+1}} \end{pmatrix}$$
(18)

and we will sometimes abbreviate this matrix as $\mathsf{P}_{r,v}^M$ for brevity. We also introduce the following shorthand notation for the unobserved type-specific choice probabilities:

$$\lambda_{\mathbf{x}_{t+1},v}^{a_{t+1},\tau} \equiv \mathbb{P}_{\tau}(a_{t+1}|\mathbf{x}_{t+1},v), \quad \lambda_{v}^{a_{t},\mathbf{x}_{t},\tau} \equiv \mathbb{P}_{\tau}(a_{t},\mathbf{x}_{t}|v).$$
(19)

Finally, define the $M_{r,v} \times \underline{M} + 1$ matrices

$$\mathsf{L}_{v}^{a_{t},\mathbf{x}_{t},\underline{M}+1} \equiv \begin{pmatrix} 1 & \lambda_{v}^{a_{t},\mathbf{x}_{t}^{1},\tau_{1}} & \dots & \lambda_{v}^{a_{t},\mathbf{x}_{t}^{M},\tau_{1}} \\ 1 & \lambda_{v}^{a_{t},\mathbf{x}_{t}^{1},\tau_{2}} & \dots & \lambda_{v}^{a_{t},\mathbf{x}_{t}^{M},\tau_{2}} \\ \vdots & \vdots & \dots & \vdots \\ 1 & \lambda_{v}^{a_{t},\mathbf{x}_{t}^{1},\tau_{M_{r,v}}} & \dots & \lambda_{v}^{a_{t},\mathbf{x}_{t}^{M},\tau_{M_{r,v}}} \end{pmatrix}$$
(20)

and

$$\mathsf{L}_{\mathbf{x}_{t+1},v}^{a_{t+1},\underline{M}+1} \equiv \begin{pmatrix} 1 & \lambda_{\mathbf{x}_{t+1}^{l},v}^{a_{t+1},\tau_{1}} & \dots & \lambda_{\mathbf{x}_{t+1}^{M},v}^{a_{t+1},\tau_{1}} \\ 1 & \lambda_{\mathbf{x}_{t+1}^{l},v}^{a_{t+1},\tau_{2}} & \dots & \lambda_{\mathbf{x}_{t+1}^{M},v}^{a_{t+1},\tau_{2}} \\ \vdots & \vdots & \dots & \vdots \\ 1 & \lambda_{\mathbf{x}_{t+1}^{l},v}^{a_{t+1},\tau_{M_{r},v}} & \dots & \lambda_{\mathbf{x}_{t+1}^{M},v}^{a_{t+1},\tau_{M_{r},v}} \end{pmatrix}.$$
(21)

The dimension of these matrices is important as we show below there is a relationship between the number of types that can be identified and the richness of the state space, as formalized in the assumption below:

ASSUMPTION 17. Given (r, v), there exist $(a_t, \mathbf{x}_t^1, \dots, \mathbf{x}_t^M, a_{t+1}, \mathbf{x}_{t+1}^1, \dots, \mathbf{x}_{t+1}^M)$ such that

- 1. $\mathbb{P}(\mathbf{x}_{t+1}^{j}|\mathbf{x}_{t}^{k}, a_{t}, v) \neq 0$ for $(j, k) \in \{1, \dots, M_{r,v} 1\}^{2};$
- 2. The matrices $L_v^{a_t,\mathbf{x}_t,\underline{M}+1}$ and $L_{\mathbf{x}_{t+1},v}^{a_{t+1},\underline{M}+1}$ have rank equal to $M_{r,v}$.

The assumption requires the existence of strictly more points in the state space than the number of types. It also requires that sufficient variation in the type-specific choice probabilities across the state space (i.e. the rows of the matrices are linearly independent). These assumptions formalize the intuition that type-specific choice probabilities are not identified if they do not vary sufficiently across types and state. This assumption although untestable (since it is placed on unobserved quantities) is reasonable here to the extent that the model is only interesting if it is true. This leads to our first Proposition.

PROPOSITION 1. Fix (r, v) and suppose that Assumption 16 holds and we can write

$$\mathbb{P}(a_t, a_{t+1}, \mathbf{x}_t, \mathbf{x}_{t+1} | r, v) = \sum_{\tau} \pi_{\tau}(r, v) \mathbb{P}_{\tau}(a_{t+1} | \mathbf{x}_{t+1}, v) \mathbb{P}(\mathbf{x}_{t+1} | \mathbf{x}_t, a_t, v) \mathbb{P}_{\tau}(a_t, \mathbf{x}_t | v).$$

Then

- 1. For a given point $(a_t, \mathbf{x}_t^1, \dots, \mathbf{x}_t^M, a_{t+1}, \mathbf{x}_{t+1}^1, \dots, \mathbf{x}_{t+1}^M)$, the total number of types $M_{r,v} \ge \operatorname{Rank}(\mathsf{P}_{r,v}^M)$ where the directly identified matrix $\mathsf{P}_{r,v}^M$ is defined in (18).
- 2. Suppose in addition that Assumption 17 holds. Then, $M_{r,v} = \operatorname{Rank}(\mathsf{P}^{\underline{M}}_{r,v})$.

The first part of the proposition only provides a lower bound on the number of types and is useful when the support of the state space is relatively small (i.e. \underline{M} is quite small). The second part shows that given a sufficiently rich state space and sufficient variation in type behavior across the space, the lower bound is also an upper bound. The proposition provides a result for the number of types at each point (r, v). The conditioning on time-invariant household characteristics v can be dropped (and will be dropped in the subsequent analysis) without affecting the identification results. However, we maintain the dependence on r since some of the identification arguments do depend upon the type indicator.

3.2.2 Identifying the Type-Specific Choice Probabilities

We now turn to identifying the type-specific choice probabilities, making use of the structure of the identified matrix $\mathsf{P}_{r,v}^M$ as well as the exclusion restriction in Assumption 16. We prove two different sets of identification results: the first set of results uses variation in the type proxy. These results are more flexible in allowing for variation across types and most useful in situations with limited state space transitions (and only consider transitions in adjacent periods). The second set of results use additional information from the markovian nature of the dynamic problem but requires a richer set of possible transitions beyond just adjacent periods. In what follows we allow the type probability π_{τ} to depend upon the exogenous variables v as usual. This is done in the interest of generality but nothing would be lost if we assumed (as one might for tractability reasons) that $\pi_{\tau}(r, v) = \pi_{\tau}(r)$. In what follows, we fix (r, v) and let $\mathcal{T}_{r,v}$ denote the set of $M_{r,v}$ types existing at (r, v). In the first approach, we assume that the following condition holds:

ASSUMPTION 18. Fix (r, v). There exists an $r' \neq r$ such that $\mathcal{T}_{r',v} = \mathcal{T}_{r,v}$.

This overlap condition therefore requires a priori knowledge about the existence of types at different values of (r, v). Note that a sufficient condition for Assumption 18 to hold is that all types exist at all values (r, v) and this is what we assume in the empirical application.

Finally, types must behave sufficiently differently at some point in the state space and in addition the behavior of a type across the state space must be sufficiently variable. Formally, this translates into an invertibility condition on the matrices $L_v^{a_t,\mathbf{x}_t,M_{r,v}}$ and $L_{\mathbf{x}_{t+1},v}^{a_{t+1},M_{r,v}}$ defined using (20) and (21) but replacing \underline{M} with $M_{r,v} - 1$ in the definition (so the dimensions now depend only upon the number of types $M_{r,v}$). For simplicity, we will abbreviate these two matrices as $L_{t,r}$ and as $L_{t+1,r}$. This condition is in essence a restatement of Assumption 17 for the square matrix case.

ASSUMPTION 19. Given (r, v), there exist $(a_t, \mathbf{x}_t^1, \dots, \mathbf{x}_t^{M_{r,v}-1}, a_{t+1}, \mathbf{x}_{t+1}^1, \dots, \mathbf{x}_{t+1}^{M_{r,v}-1})$ such that (a) $\mathbb{P}(\mathbf{x}_{t+1}^j | \mathbf{x}_t^k, a_t, v) \neq 0$ for $(j, k) \in \{1, \dots, M_{r,v} - 1\}^2$ and (b) the $M_{r,v} \times M_{r,v}$ matrices $\mathsf{L}_v^{a_t, \mathbf{x}_t, M_{r,v}}$ and $\mathsf{L}_{\mathbf{x}_{t+1}, v}^{a_{t+1}, M_{r,v}}$ (abbreviated as $\mathsf{L}_{t,r}$ and $\mathsf{L}_{t+1,r}$) are invertible.

We can then state the following lemma for point identification of the type-specific choice probabilities:

LEMMA 6. Fix (r, v) and suppose Assumption 16 holds. In addition, suppose that for all (a_t, a_{t+1}) there exist $(\mathbf{x}_t^1, \ldots, \mathbf{x}_t^{M_{r,v}-1}, \mathbf{x}_{t+1}^1, \ldots, \mathbf{x}_{t+1}^{M_{r,v}-1})$ such that Assumptions 18 and 19 hold and that the optimal decision process is Markov. Then, the type-specific choice probabilities $\{\mathbb{P}_{\tau}(a_t|\mathbf{x}_t, v)\}_{\tau \in \mathcal{T}_{r,v}; t \in \{1,2,3\}}$ for $(\mathbf{x}_t, v) \in \mathcal{X}_t \times \mathcal{V}$ are identified.

If Assumption 18 fails, identification results can be derived that make stronger assumptions about the nature of state transitions across three (as opposed to two) periods. In this case, a set of type-specific choice probabilities are recovered for each value of (r, v) without requiring any overlap. These results, which for brevity are stated in Appendix A.2.1, are helpful for instance if the type proxy creates a mutually exclusive partition of agent types. Note that this is not the same as in the special case of

directly observed types analyzed in Section 3.1. In this latter case, not only is there a single type in each partition, but it is also known which specific type is observed in each partition.

3.2.3 Assigning Identities to Choice Probabilities

The previous subsection identified the type-specific choice probabilities but did not identify the identities of the specific types. We now outline a procedure that infers the identifies of types from these typespecific choice probabilities. First, we apply Lemma 1 to these identified type-specific choice probabilities $\mathbb{P}_{\tau}(a_t|\mathbf{x}_t, v)$ to identify period 2 and 3 utilities $(\{u_t(x_t, a_t; \tau))\}_{t \in \{2,3\}, \tau \in \mathcal{T}})$ as well as the product $\beta_{\tau} \delta_{\tau}$.⁵ We use these to construct the function $\tilde{h}_{\tau}(x_2, z_2, \beta_{\tau} \delta_{\tau}, \mathsf{d})$ which using the notation from (8) we can write as $\tilde{h}_{\tau}(x_2, z_2, \beta_{\tau} \delta_{\tau}, \mathsf{d}) = \tilde{h}_{\tau,1}(x_2, z_2, \beta_{\tau} \delta_{\tau}) + \mathsf{d}\tilde{h}_{\tau,2}(x_2, z_2, \beta_{\tau} \delta_{\tau})$ and define the following terms for a pre-specified point (x_{20}, z_{20}) :

$$\tilde{h}_{\tau,j}^{\Delta}(x_2, z_2, \mathbf{b}) \equiv \tilde{h}_{\tau,j}(x_2, z_2, \mathbf{b}) - \tilde{h}_{\tau,j}(x_{20}, z_{20}, \mathbf{b}) \quad j \in \{1, 2\}
\tilde{h}_{\tau}^{\Delta}(x_2, z_2, \beta_{\tau} \delta_{\tau}, \mathbf{d}) \equiv \tilde{h}_{\tau}(x_2, z_2, \beta_{\tau} \delta_{\tau}, \mathbf{d}) - \tilde{h}_{\tau}(x_{20}, z_{20}, \beta_{\tau} \delta_{\tau}, \mathbf{d})
= \tilde{h}_{\tau,1}^{\Delta}(x_2, z_2, \beta_{\tau} \delta_{\tau}) + \mathsf{d}\tilde{h}_{\tau,2}^{\Delta}(x_2, z_2, \beta_{\tau} \delta_{\tau}).$$
(22)

As in Section 3.1, functions with a Δ superscript can be interpreted as 'normalized' future expected utilities from the viewpoint of period 1 for an agent with given perceived future discounting factors, and where the normalization is relative to a given point (x_{20}, z_{20}) in the state space. Next, using Lemma 3 we identify the function

$$h_{\tau}^{\Delta}(x_2, z_2) \equiv h_{\tau}(x_2, z_2) - h_{\tau}(x_{20}, z_{20}), \tag{23}$$

where $h_{\tau}(x_2, z_2)$ was defined in (7). Next, define the identified function

$$\hat{\delta}_{\tau}(x_2, z_2) \equiv \frac{h_{\tau}^{\Delta}(x_2, z_2) - \tilde{h}_{\tau, 1}^{\Delta}(x_2, z_2, \beta_{\tau} \delta_{\tau})}{\tilde{h}_{\tau, 2}^{\Delta}(x_2, z_2, \beta_{\tau} \delta_{\tau})}.$$
(24)

The variation of this function across the state space will be key to identifying the different types since (some of the) different types will exhibit different types of variation. First, however, in order for this to be well-defined we need the denominator to be non-zero:

ASSUMPTION 20. There exists a set $S \subset \mathcal{X}_2 \times \mathcal{Z}_2$ with positive measure such that for all types τ , $\tilde{h}_{\tau,2}^{\Delta}(x_2, z_2, \beta_{\tau} \delta_{\tau}) \neq 0$.

This is a significant strengthening of Assumption 11 and applies to all types (not just sophisticated types) which is necessary since type identities are not known (note that it is testable). The assumption ensures that δ_{τ} enters $h_{\tau}(x_2, z_2)$ separately for sophisticated types and guarantees that (24) is well-defined at least on S for all types.

The next assumption ensures that $\hat{\delta}$ defined above is non-constant across the state space for naïve agents. It does so by ensuring that two views of the future are sufficiently different across the state space.

⁵Because we do not know the identity of the type, we cannot conclude anything more than $\beta_{\tau}\delta_{\tau}$ is identified. Specifically, since we do not know whether the type is time-consistent we cannot set $\beta_{\tau} = 1$.

Roughly, the function $\tilde{h}_{\tau_N}^{\Delta}(x_2, z_2, \delta_{\tau_N}, \delta_{\tau_N})$ is the (normalized) value of being in state (x_2, z_2) from the view-point of period 1 for a naïve agent, while $\tilde{h}_{\tau_N}^{\Delta}(x_2, z_2, \beta_{\tau_N} \delta_{\tau_N}, \delta_{\tau_N})$ is the same value for the agent if instead he were sophisticated. The assumption below states that these two views of the future have to vary over the state space. It seems reasonable to assume that the naïve and sophisticated type differ over the future and if it failed, period 1 choice probabilities would be identical for naïve types and sophisticated types who have the same preference parameters as their naïve counterparts.

ASSUMPTION 21. Var
$$\left(\frac{\tilde{h}_{\tau_N}^{\Delta}(x_2, z_2, \delta_{\tau_N}, \delta_{\tau_N}) - \tilde{h}_{\tau_N}^{\Delta}(x_2, z_2, \beta_{\tau_N}, \delta_{\tau_N}, \delta_{\tau_N})}{\tilde{h}_{\tau_N, 2}^{\Delta}(x_2, z_2, \beta_{\tau_N}, \delta_{\tau_N})}\right) > 0.$$

We can now state the results for assigning type identities to the (type-specific) choice probabilities. We begin with a useful result and then state the general result immediately after.

LEMMA 7. Suppose that the type-specific choice probabilities $\{\mathbb{P}_{\tau}(a_t|\mathbf{x}_t, v)\}_{\tau \in \mathcal{T}_{r,v}; t \in \{1,2,3\}}$ are identified and the conditions for Lemma 3 and Assumptions 20 and 21 hold. Then, $\hat{\delta}_{\tau}(x_2, z_2)$ is a constant for all $(x_2, z_2) \in S \iff \tau \neq \tau_N$.

We next apply this result to finally assign type identities, which is our second proposition.

PROPOSITION 2. Suppose that the type-specific choice probabilities are identified and that Assumptions 20 and 21 hold. In addition, suppose that Assumptions 1 - 4, 6 - 13 hold. Then, type identities are identified.

The proof proceeds by first applying the results of Lemma 1 and the first two parts of Lemma 2, all of which do not require the type identity to be known. Using the quantities thus identified, we can construct the function $\hat{\delta}_{\tau}(x_2, z_2)$ and then apply Lemma 7.

3.2.4 Identifying Preferences for each Type

Note that most of the work in identifying preferences was already done in the previous section while identifying type identities. In particular, we identified per-period utilities $\{u_t(\cdot, \tau)\}_{t,\tau}$ for each period and the product of the time-preference parameters $\beta_{\tau}\delta_{\tau}$. In addition, for sophisticated and time-consistent agents the identified object $\hat{\delta}(x_2, z_2) = \delta_{\tau}$ so that for these two types the time-preference parameters are also separately identified. For naïve agents, we can use Lemma 4 and Lemma 5 to (set or point) identify the time-preference parameters.

3.3 Partial Sophistication

A key assumption in the previous analysis was that $\beta_{\tau} \in \{\beta_{\tau}, 1\}$ (Assumption 9). This restriction greatly simplified the analysis by reducing the number of time preference parameters. This came at the cost of making a strong assumption about behavior—that agents are either completely aware of the extent of their future present-bias (sophisticated agents) or completely unaware about it (naïve types). In this section, we discuss identification when agents are only partially aware of their future present-bias. Our main finding is that without further assumptions, none of the time preference parameters are point identified. We start the analysis by replacing Assumption 9 by

ASSUMPTION 22. Agents are partially sophisticated (or equivalently partially naïve): $\tilde{\beta}_{\tau} \in [\beta_{\tau}, 1]$.

Under this assumption, the only sharp distinction is between time-consistent agents ($\beta_{\tau} = 1$) on the one hand and partially sophisticated agents on the other hand. Partially sophisticated agents are not all identical since they may have different values of ($\tilde{\beta}_{\tau}, \beta_{\tau}$) as well as different exponential parameters and per period utilities. Given the finiteness assumption on the total number of types, however, we can only have a finite number of partially sophisticated agents.

We first discuss type identification. Starting with a given type-specific choice probability we show that one can decide whether the type associated with the probability is consistent (and if not consistent, the type must perforce be partially sophisticated). The arguments are very similar to those employed in identifying whether a given type is naïve in Lemma 7. In particular, under the assumption below, a type will be partially sophisticated if and only if the directly identified object $\hat{\delta}(x_2, z_2)$ varies over the state space.

ASSUMPTION 23. For types
$$\tau$$
 such that $\tilde{\beta}_{\tau} \neq \beta_{\tau}$, $\operatorname{Var}\left(\frac{\tilde{h}_{\tau}^{\Delta}(x_{2},z_{2},\tilde{\beta}_{\tau}\delta_{\tau},\delta_{\tau})-\tilde{h}_{\tau}^{\Delta}(x_{2},z_{2},\beta_{\tau}\delta_{\tau},\delta_{\tau})}{\tilde{h}_{\tau,2}^{\Delta}(x_{2},z_{2},\beta_{\tau}\delta_{\tau})}\right) > 0.$

The function $\tilde{h}_{\tau}^{\Delta}(x_2, z_2, \tilde{\beta}_{\tau} \delta_{\tau}, \delta_{\tau})$ is the normalized value of being in state (x_2, z_2) from the period 1 viewpoint, defined in (22), for a partially sophisticated agent. The function $\tilde{h}_{\tau}^{\Delta}(x_2, z_2, \beta_{\tau} \delta_{\tau}, \delta_{\tau})$ is the same value for the agent if instead he were fully sophisticated. The assumption above states that the difference between these two views of the future has to vary over the state space. In its absence period 1 choice probabilities would be identical for partially and completely sophisticated types. This assumption is a strengthening of Assumption 21 to apply to all types and we can state the analogous result to Lemma 7 in the following proposition:

PROPOSITION 3. Suppose that Assumptions 1 - 4, 6 - 8, 12, 13, 20, 22 and 23 hold. Further, suppose that the type-specific choice probabilities $\{\mathbb{P}_{\tau}(a_t|\mathbf{x}_t, v)\}_{\tau \in \mathcal{T}_{r,v}; t \in \{1,2,3\}}$ are identified. Then,

- 1. $\hat{\delta}_{\tau}(x_2, z_2)$ is a constant for all $(x_2, z_2) \in \mathcal{S} \iff \tilde{\beta}_{\tau} = \beta_{\tau}$
- 2. Time-consistent types ($\tilde{\beta}_{\tau} = \beta_{\tau} = 1$), completely sophisticated types ($\tilde{\beta}_{\tau} = \beta_{\tau} < 1$) and partially sophisticated types ($\tilde{\beta}_{\tau} \neq \beta_{\tau}$) are identified.

We next turn to the identification of the time-preference parameters for the partially sophisticated agents. The main result here is that without further assumptions, the three time-preference parameters for these agents (i.e., $\delta_{\tau}, \beta_{\tau}$, and $\tilde{\beta}_{\tau}$) are not point-identified although we show that if the exponential discount factor (δ_{τ}) is identified, then the remaining two parameters are also identified. The reasons for non-identification are similar to the reasons for non-identification of the time-preference parameters for the completely naïve agents in Section 3.1. As before, we can identify the per period utility functions $\{u_t(\cdot,\tau)\}_{t\in\{2,3\},\tau\in\mathcal{T}}$ and the product $\beta_{\tau}\delta_{\tau}$ using period 2 and 3 choices regardless of type τ . This information allows us to construct the function $\tilde{h}^{\Delta}_{\tau}(x_2, z_2, \mathbf{b}, \mathbf{d})$. Recall that period 1 choices identify the function $h^{\Delta}_{\tau}(x_2, z_2)$ and we know that $h^{\Delta}_{\tau}(x_2, z_2) = \tilde{h}^{\Delta}_{\tau}(x_2, z_2, \tilde{\beta}_{\tau}\delta_{\tau}, \delta_{\tau})$. This information, however, is not enough to identify the time-preference parameters since the identified function $\tilde{h}^{\Delta}_{\tau}(\cdot)$ is not one-to-one in $(\delta_{\tau}, \tilde{\beta}_{\tau})$. Further, since types are partially sophisticated, we cannot impose any other restrictions on $\tilde{\beta}_{\tau}$ separately. This is in sharp contrast to the consistent or the completely sophisticated case where $\tilde{\beta}_{\tau} = \beta_{\tau}$, so that period 2 and 3 choices (which identify $\beta_{\tau}\delta_{\tau}$) identify $\tilde{\beta}_{\tau}\delta_{\tau}$. Equivalently, given the structure of the model, the product $\beta_{\tau}\delta_{\tau}$ is not sufficiently informative about either $\tilde{\beta}_{\tau}\delta_{\tau}$ or δ_{τ} . The following proposition states the most general result for partially sophisticated types, which is a set identification result.

PROPOSITION 4. Consider a partially sophisticated agent maximizing (2) under assumptions 1- 4, 6-8, 12, 13, 22 and 23. Then, the identified set for the parameters $(\beta_{\tau}, \tilde{\beta}_{\tau}, \delta_{\tau})$ is given by

$$\Theta_{\beta,\tilde{\beta},\delta} = \left\{ (\mathsf{b},\tilde{\mathsf{b}},\mathsf{d}) \in (\beta_\tau \delta_\tau, 1]^3 : \tilde{h}_\tau^\Delta(x_2, z_2, \mathsf{bd}, \mathsf{d}) = h_\tau^\Delta(x_2, z_2) \;\; \forall (x_2, z_2), \mathsf{b} = (\beta_\tau \delta_\tau)/\mathsf{d}, \mathsf{b} < \tilde{\mathsf{b}} \right\}.$$

4 Empirical Application: Adoption and Retreatment of ITNs

In this section, we use the identification results developed above to examine the role of time-inconsistent preferences in explaining demand for and proper maintenance of insecticide-treated nets (ITNs) in a sample of households from rural Orissa, India. Bed nets are hung over sleeping areas to protect from mosquito-borne diseases such as malaria. In online Appendix B we discuss in detail the mappings from the general model to our application and provide the associated identification proofs.

We assume that agents have preferences as in (2), and are drawn from a population that includes time-consistent as well as hyperbolic naïve and sophisticated types. We adopt the more general framework described in Section 3.2 where types are not directly observed, although for simplicity we assume that the exponential discount factor δ is common to all agents and we abstract from the possibility of partial sophistication. We also assume that all three types exist with positive probability.⁶

Agents choose whether to purchase an ITN and whether to retreat it periodically to ensure that the net maintains its ability to kill mosquitoes that come into contact with it. For our structural estimation, given sample size concerns, we impose functional forms on the utility function so that the model reduces to one characterized by a vector-valued parameter, so that inference will follow from standard asymptotic arguments. Recall that for our identification strategy, the key variables required in the general model are the type indicator r and the excluded variables z. Prior to the ITN distribution, we elicited time-preferences by asking respondents to make a series of inter-temporal choices and this information forms part of the type indicator that we denote by \tilde{r} . The other element of the type indicator vector is the purchase decision of a product designed to appeal to sophisticated agents (see below), thereby being potentially informative about type. Finally, the excluded variables z are elicited subjective beliefs about ITN efficacy in preventing malaria and about future income.

4.1 Data

The data used in this paper were collected in the context of a randomized controlled trial (RCT) carried out in 2007–2009 in Orissa, the most malaria endemic state in India (Dhingra et al. 2010). The study evaluated impacts of alternative mechanisms of providing ITNs on health and socio-economic outcomes of potential users users, and was carried out in collaboration with our local partner, Bharat Integrated Social Welfare Agency (BISWA), a micro-lender with a large presence in Orissa and elsewhere in India, see Tarozzi et al. (2014) for details. We use data collected from a sample of 621 households in 47 villages

⁶This assumption is made for convenience, but one could apply the results from Section 3.2.1 to first identify the total number of types before proceeding to the analysis undertaken below.

randomly assigned to an experimental arm where BISWA offered all its clients the opportunity to purchase high quality ITNs on credit, with repayment over one year.

A baseline, pre-intervention survey was carried out in March–April 2007. In September–November, all villages were exposed to a brief community-based information campaign about the importance of ITN use and about their proper use and maintenance. BISWA clients were offered the opportunity to purchase ITNs. Purchases were completed 2–3 days later, to allow careful consideration of the offers. A second visit was scheduled approximately one month later, and nets were offered again with the same contracts (no further sales were made after the second visit). The first net re-treatment was completed approximately six months after the ITN sale, in March–April 2008, while the second and final re-treatment took place another six months later, in September–November 2008.

Two alternative contracts were offered to BISWA clients. With the first offer (referred to as b henceforth), single (double) nets were sold on credit for Rs. 173 (223), to be repaid within one year. For perspective, daily wages for agricultural labor in the area were around Rs. 50. Nets were immediately treated with insecticide, with a chemical concentration that made 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 c 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), again to be repaid within one year. 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. For both contracts the price was inclusive of 20% annual interest—the standard annual rate charged by BISWA in its micro-finance operations—but for simplicity in the sequel we do not explicitly model that nets were sold on credit.

Of the initial sample of 621 households, we exclude from the analysis 32 that could not be re-contacted at endline, 13 that purchased bednets with both contracts, 9 that purchased nets for cash, and one because the contract type was not recorded. We are thus left with a sample of 566 households.

Table 1 shows summary statistics at baseline. Mean monthly total expenditure per head was approximately twice as large as the official poverty line for rural Orissa in 2004–5. Net ownership was not uncommon, with a mean of one bed net for every three persons, although one third of households did not own any nets. However, net treatment was rare, with only 0.06 ITNs per head on average. On average, 16% of individuals slept under a net the night before the survey, and 3% under an ITN. Results from blood tests show high prevalence of malaria (11%) and anemia (46%), where the latter denotes hemoglobin (Hb) levels < 11 g/dl blood.⁷

Respondents were aware of the role of mosquitoes in transmitting malaria, of the high cost of malaria episodes, and that bednets reduced malaria risk. The latter was also reflected in subjective beliefs, that we elicited 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.⁸ Since most

⁷Malaria infection and Hb were measured via rapid diagnostic tests that only required fingerprick blood specimens, and where immediately communicated to individuals, see Appendix A.2 in Tarozzi et al. (2014) for additional details. At baseline, consent was requested to test pregnant women, children under five years (U5) as well as their mothers, and one randomly selected adult (age 15-60).

⁸We did not measure ranges of probability, so we cannot identify the degree of uncertainty around the reports.

respondents were illiterate and unfamiliar with the formal concept of probability, the interviewer discussed first hypothetical examples of certain and uncertain events to explain the procedure. Beliefs were elicited using wording such as the following: "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 an adult that does not sleep under a bed net will contract malaria in the next 1 year?" Perceived malaria risk was also recorded conditional on using an untreated bed net or an ITN.⁹

The histograms in Figure 1 show that bednets and re-treatment with insecticide were recognized as very effective at reducing malaria risk. About three quarters of respondents believed that without using nets one would certainly get malaria, and approximately the same fraction believed that regular use of ITNs would virtually rule out risk. According to about half of all respondents, there was a 50% chance of developing malaria if an untreated net was used. Despite the spikes over the focal numbers 0, 5, and 10, there remains a sufficient degree of variation to be exploited by the structural model outlined below.

The baseline survey also included 12 questions intended to gauge respondents' intertemporal preferences and the extent of time inconsistency. 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 16) 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 questions the same rewards described for the first group were offered, but with time horizons shifted by three months.¹⁰

Standard expected utility models imply that if a respondent prefers, say, Rs. 16 four months later to Rs. 10 paid a month from today, s/he should also prefer Rs. 16 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 reversed for the later rewards—to be correlated with a form of inconsistency in time preferences consistent with hyperbolic discounting.¹¹ 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 reversal.¹²

Panel A of Table 3 includes a summary of the results of the ITN sale. Slightly more than 50% of sample households purchased at least one net on credit (287 of 566). Of these, 141 chose to purchase

⁹Analogous beliefs were elicited about the protective power of bednets and treatment for children and pregnant women. Responses were almost identical across demographic groups, and so we only use information for adults.

¹⁰Interviewers told respondents that one of the twelve chosen rewards, selected at random, would be paid by our microlender 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 likely not critical because all respondents belonged to households with at least one BISWA client.

¹¹ See Andreoni and Sprenger (2012) for an alternative view. Rubinstein (2003) shows cases in which preference reversals arise despite preferences that are neither consistent not time-inconsistent hyperbolic. We focus on the identification of types with preferences that are compatible with time-consistency or hyperbolic discounting. We only impose that agents with choice reversals are potentially more likely to have time-inconsistent preferences than time-consistent preferences. Identifying different forms of time-inconsistent preferences, albeit important, is beyond the scope of our paper.

¹²We ignore the possibility of agents with "future bias" for whom $\beta > 1$. Such types are rarely considered in the literature, and our identification results do not extend to such fourth type. For households who exhibited anti-hyperbolic behavior at least once (that is, chose the later payoff with options closer in time, but then chose the earlier one with options farther away in time) we assign the signal r = 0. Our results are qualitatively similar if we drop these respondents. Also note that we do not require a one-to-one mapping from signals into types, such that the model is robust against imperfect type signals.

only ITNs (contract b), while 146 opted for the "commitment" product whose price also included the cost of two re-treatments (contract c).¹³ In panel B we show that the prevalence of re-treatment was strongly associated with the choice of contract. At the time of the first re-visit, about six months after the ITN sale, on overwhelming majority (92%) of the ITNs purchased with contract c were re-treated with insecticide. However, the fraction was only 36% for bednets who had been purchased with contract b, and for which re-treatment required a (small) cash payment to be paid on the spot. Six months later, re-treatment rates declined to 84% for contract c and dropped by almost half for contract b.

4.2 Empirical Model

We begin by specifying preferences and then discuss the transition probabilities and other key ingredients of the dynamic programming problem. The central difference from standard analysis of dynamic models in what follows is the presence of (the different types of) time-inconsistent agents and the further complication that these types are unobserved. These two deviations alter the standard results and we highlight these differences below. Furthermore, instead of the minimum necessary 3 periods in our model, we use 4 periods because of the specific features of our intervention.

PREFERENCES (PERIOD 4): In period 4, the state variables are income and health $(x_4 = (y_4, h_4))$, where h_4 is equal to m if someone in the household contracts malaria in period 4 and 0 otherwise, and y_4 indicates whether the household is in a high or low income state. The household characteristics that enter preferences are denoted by v. In the base specification v includes household size at baseline (v_{hhs}) , a measure of households assets (v_{assets}) , and an indicator of risk aversion (v_{risk}) . 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 online Appendix D for details. We specify

$$u_4(x_4, v; \tau) = \mathcal{C}(x_4) + \boldsymbol{\phi}_{\tau}(v),$$

where $C(x_4)$ is consumption in state x_4 , and ϕ captures other factors that can affect per-period utility. Consumption $C(\cdot)$ depends on the combination of both health and income. In any period we compute consumption as $C(x_t) = y_t - \mathbb{I}\{h_t = m\}\eta_m$, where η_m accounts for costs of malaria, including the direct cost of treatment and a lower income when being sick. We set η_m to the median cost of recent malaria episodes in our data.

We allow for heterogeneity in the per period utility both in terms of observable dimensions and across unobservable types (τ):

$$\boldsymbol{\phi}_{\tau}(v) = \phi_0 + \phi_1 \mathbb{I}(\tau = \text{sophisticated}) + \phi_2 \mathbb{I}(\tau = \text{naïve}) + \phi_3 v_{\text{hhs}} + \phi_4 v_{\text{assets}} + \phi_5 v_{\text{risk}}, \qquad (25)$$

The type-specific component in per-period utility does not depend on the action to be taken in that

¹³ Due to sample size considerations and for tractability, we do not model the decision to buy single *versus* double nets, or the number of nets purchased. In the data, buyers purchased on average close to two nets and in the model we assume demand, when positive, is always equal to 1.5 ITNs. For each contract type, we set the price equal to 1.5 times the weighted average of prices for single and double nets, weighted by the respective purchase frequencies. Compared to the cases in which each household buys only one net or 2 nets, the estimated type probabilities and the discount factor are virtually unchanged. A higher number of nets slightly decreases the degree of present bias, because the higher cost of purchasing more nets decreases demand.

period so that identification of the parameters here depends upon the forward-looking component of the value function (which contains these parameters and is a function of the action taken in that period).

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

$$u_t(x_t, a_t, v; \tau) = (\mathcal{C}(x_t) - p_r a_t \mathbb{I}\{a_1 = b\} - p_r \mathbb{I}\{a_1 = c\}) + \phi_\tau(v);$$

where p_r is the price of retreatment an agent has to pay in period t regardless of the choice to retreat or not in case he has purchased a commitment contract, and that he has to pay in case he chooses to retreat the net in period t after having chosen a baseline contract in period 1.¹⁴

PREFERENCES (PERIOD 1): In period 1, preferences are given by

$$u_1(x_1, a_1, v; \tau) = (\mathcal{C}(x_1) - p_b \mathbb{I}\{a_1 = b\} - p_c \mathbb{I}\{a_1 = c\}) + \phi_\tau(v),$$

where p_b is the price of a baseline contract comprised of the ITN alone, and p_c is the price of the commitment contract that includes retreatment in periods 2 and 3.

We parametrize the mixture probabilities $\pi_{\tau}(r)$ as a function of two parameters γ_1 and γ_2 using the parsimonious Hardy-Weinberg functional form (Hardy, 1908; Weinberg, 1908). In a first step, a logit function ψ is computed that depends on two estimable mixture parameters γ_1 and γ_2 and the observable value of the signal r, and is defined as

$$\psi(\gamma, r) = \frac{exp(\gamma_1 + \gamma_2 \cdot r)}{1 + exp(\gamma_1 + \gamma_2 \cdot r)}.$$
(26)

From this function one can derive the type probabilities conditional on the signal r as

$$\pi_{C}(r) = \psi^{2}(\gamma, r);$$

$$\pi_{N}(r) = 2\psi(\gamma, r)(1 - \psi(\gamma, r));$$

$$\pi_{S}(r) = 1 - \pi_{C}(r) - \pi_{N}(r) = 1 - \psi^{2}(\gamma, r) - 2\psi(\gamma, r)(1 - \psi(\gamma, r)).$$

We find similar estimated type distributions using alternative parametrizations.

¹⁴We do not have reliable data on the timing of debt repayments. We make the conservative assumption that debts are repaid within the first six months, except for the component attributable to (optional or commitment) retreatment. The assumption is conservative in the sense that present-bias is exacerbated if we consider repayments in future periods. Nevertheless, the results remain qualitatively similar in terms of estimated type probabilities and relative adoption probabilities when putting greater weight on later repayments. The fundamental difference is that with a commitment contract the agent has to pay the additional retreatment price regardless of whether he retreats or not, while with a standard contract he only pays the retreatment price in case he chooses to retreat. The modeling assumption that the retreatment component of the commitment contract (which amounts to about 13 percent of the total payment) is paid in later periods as is optional retreatment is mostly due to technical considerations. Dropping this assumption does not qualitatively change the present-bias nor type probability results nor the accuracy of Monte Carlo simulations, but it worsens the fit of the model and in some cases the sufficiency conditions regarding type identification can be violated. The results are available from the authors upon request. The worse fit arises because the high present-bias (i.e., low values of β_N and β_S) makes the commitment contract much less attractive than a baseline contract because of the higher initial payment. We believe that given our data constraints regarding the effective time of repayment, having both (mandatory) commitment and (optional) baseline contract payments deducted in the later periods is a reasonable choice.

4.2.1 Solving the Model

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 typespecific value functions (defined below). This is justified because even though we do not observe types, the type-specific choice probabilities are identified using the results in the previous sections.

PERIOD 3 CHOICE: At t = 3 an agent who has purchased an ITN will choose to re-treat his net if

$$u_3(x_3, 1, v; \tau) + \epsilon_3(1) + \beta_\tau \delta \int u_4(x_4, v; \tau) \, \mathrm{dF}(x_4 | x_3, 1; z) > u_3(x_3, 0, v; \tau) + \epsilon_3(0) + \beta_\tau \delta \int u_4(x_4, v; \tau) \, \mathrm{dF}(x_4 | x_3, 0; z)$$

which we rewrite as

$$\epsilon_3(0) - \epsilon_3(1) < v_{\tau,3}(x_3, 1, z, v, \beta_{\tau}, \delta) - v_{\tau,3}(x_3, 0, z, v, \beta_{\tau}, \delta),$$

where $a_3 = 1$ reflects retreatment and $a_3 = 0$ reflects no retreatment. Recall that $v_{\tau,3}(\cdot)$ is the choice specific value function for type τ :

$$v_{\tau,3}(x_3, a, z, v, \beta_{\tau}, \delta) \equiv u_3(x_3, a, v; \tau) + \beta_{\tau} \delta \int u_4(x_4, v; \tau) \,\mathrm{dF}(x_4 | x_3, a, z).$$
(27)

Here z reflects the household-specific vector of beliefs for the probability of contracting malaria when (a) sleeping without a net, (b) sleeping under an untreated net, and (c) sleeping under an ITN. 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 also set $\sigma = 1$ in what follows). Under the GEV assumption, the type-specific choice probability is given by

$$\mathbb{P}_{\tau}(a_3 = 1 | x_3; z, v) = \frac{\exp(v_{\tau,3}(x_3, 1, z, v, \beta_{\tau}, \delta))}{\sum_{j=0}^{1} \exp(v_{\tau,3}(x_3, j, z, v, \beta_{\tau}, \delta))}.$$
(28)

PERIOD 2 CHOICE: Similarly to period 3, the choice probabilities are given by

$$\mathbb{P}_{\tau}(a_2 = 1 | x_2; z, v) = \frac{\exp(v_{\tau,2}(x_2, 1, z, v, \beta_{\tau}, \delta))}{\sum_{j=0}^{1} \exp(v_{\tau,2}(x_2, j, z, v, \beta_{\tau}, \delta))}.$$
(29)

The $v_{\tau,2}(\cdot)$ functions—whose form will provide some insight into the time-inconsistency problem—are defined as in period 3, except that the calculation of the integral is more involved:

$$v_{\tau,2}(x_2, a, z, v, \beta_{\tau}, \delta) \equiv u_2(x_2, a, v; \tau) + \beta_{\tau} \delta \int v_{\tau}^*(s_3, z, v) \,\mathrm{dF}(s_3 | x_2, a, z), \tag{30}$$

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

$$A_k^{\tau} \equiv \{ v_{\tau,3}(x_3, k, z, v, \tilde{\beta}_{\tau}, \delta) + \epsilon_3(k) > v_{\tau,3}(x_3, s, z, v, \tilde{\beta}_{\tau}, \delta) + \epsilon_3(s) \quad \forall s \neq k \},$$
(31)

which is the event that action k is optimal in period 3 given a type- τ agent's future expected present-bias

of $\tilde{\beta}_{\tau}$. Next, if $\mathbb{I}_{A_k^{\tau}}$ is an indicator for the event A_k^{τ} , we have

$$v_{\tau}^{*}(s_{3}, z, v) \equiv \sum_{k \in \{0,1\}} \left(v_{\tau,3}(x_{3}, k, z, v, 1, \delta) + \epsilon_{3}(k) \right) \mathbb{I}_{A_{k}^{\tau}}.$$

This expression takes into account that (a) from the viewpoint of period 2, utility only depends on β_{τ} through the discounting of the integral in (30) above (note that $\beta_{\tau} = 1$ in $v_{\tau,3}(.)$); (b) the agent takes into account that in period 3 his choices (captured by A_k^{τ}) will be affected by the *perceived* discount factor $\tilde{\beta}_{\tau}$, whose specific value will depend on the degree of the agent's sophistication. For time-consistent agents $\tilde{\beta}_{\tau} = \beta_{\tau} = 1$ while for naïve agents $\tilde{\beta}_{\tau} = 1$ and for fully sophisticated agents $\tilde{\beta}_{\tau} = \beta_{\tau}$.

Integrating with respect to the GEV error distribution yields

$$\begin{aligned} v_{\tau}^{*}(x_{3}, z, v) &\equiv \int v_{\tau}^{*}(x_{3}, \epsilon_{3}, z, v) \,\mathrm{dF}(\epsilon_{3}) \\ &= \sum_{k \in \{0, 1\}} \mathbb{P}(A_{k}^{\tau}) \bigg(v_{\tau, 3}(x_{3}, k, z, v, 1, \delta) + \gamma_{\texttt{euler}} - \log \mathbb{P}(A_{k}^{\tau}) \bigg), \end{aligned}$$

where γ_{euler} is Euler's constant, and

$$\mathbb{P}(A_k^{\tau}) = \frac{\exp(v_{\tau,3}(x_3, k, z, v, \beta_{\tau}, \delta))}{\sum_{j=1}^J \exp(v_{\tau,3}(x_3, j, z, v, \tilde{\beta}_{\tau}, \delta))}$$

so that

$$\begin{split} v_{\tau}^{*}(x_{3},z,v) &= \sum_{k \in \{0,1\}} \mathbb{P}(A_{k}^{\tau}) \\ \times \left[v_{\tau,3}(x_{3},k,z,v,1,\delta) - v_{\tau,3}(x_{3},k,z,v,\tilde{\beta}_{\tau},\delta) + \gamma_{\texttt{euler}} + \log\left(\sum_{j \in \{0,1\}} \exp(v_{\tau,3}(x_{3},j,z,v,\tilde{\beta}_{\tau},\delta))\right) \right]. \end{split}$$

The first two terms in brackets give the difference between the net value (in period 2) of taking action k assuming no present-bias (in period 3) and the corresponding net value of the same action assuming a present-bias of $\tilde{\beta}_{\tau}$ (in period 3). 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 contrast, sophisticated inconsistent agents recognize their future inconsistency and this expression is non-zero for them.¹⁵

PERIOD 1 CHOICE: The argument is again similar 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

$$\mathbb{P}_{\tau}(a_1 = a | x_1; z, v) = \frac{\exp(v_{\tau,1}(x_1, a, z, v, \beta_{\tau}, \delta))}{\sum_{j \in \{n, b, c\}} \exp(v_{\tau, 1}(x_1, j, z, v, \beta_{\tau}, \delta))},$$
(32)

 $^{^{15}}$ Our model does not allow for additional learning about the usefulness of retreatment between periods 1 and 3.

where b reflects the purchase of a baseline contract, c reflects a purchase of a commitment contract, and n reflects no contract purchase. For period 1, the $v_{\tau,1}(\cdot)$ function is

$$v_{\tau,1}(x_1, s, z, v, \beta_{\tau}, \delta) \equiv u_1(x_1, s, v; \tau) + \beta_{\tau} \delta \int v_{\tau}^*(s_2, z, v) \,\mathrm{dF}(s_2 | x_1, s, z), \tag{33}$$

where $v_{\tau}^*(\cdot)$ is the value function for type τ in period 2 and following a by now familiar argument

$$\begin{split} v_{\tau}^*(x_2, z, v) &= \sum_{k \in \{0,1\}} \mathbb{P}(A_k^{\tau}) \\ \times \left[v_{\tau,2}(x_2, k, z, v, 1, \delta) - v_{\tau,2}(x_2, k, z, v, \tilde{\beta}_{\tau}, \delta) + \gamma_{\texttt{euler}} + \log \left(\sum_{j \in \{0,1\}} \exp(v_{\tau,2}(x_2, j, z, v, \tilde{\beta}_{\tau}, \delta)) \right) \right] \end{split}$$

where

$$\mathbb{P}(A_k^{\tau}) = \frac{\exp(v_{\tau,2}(x_2, k, z, v, \beta_{\tau}, \delta))}{\sum_{j \in \{0,1\}} \exp(v_{\tau,2}(x_2, j, z, v, \tilde{\beta}_{\tau}, \delta))}.$$

We also account for the presence of untreated bednets in the households prior to the intervention.¹⁶ This is because owning an untreated bednet in period 1 affects the utility in case of not purchasing an ITN as well as the transition probabilities in any period because for almost all households it is perceived to lower the probability of getting sick compared to the case of not having any net at all.

Monte Carlo simulations In online Appendix G we illustrate the properties of our model with a set of Monte Carlo simulations. For the observed types case, Tables G.1 and G.2 show that for moderate sample sizes (300 and above) both the mean and the median estimated time preference parameters are close to their true values. Tables G.3 and G.4 show that when types are unobserved, the estimates (which now also include the type probabilities) continue to be close to their true values, albeit with more variability. The largest differences between the estimates and the true parameter values in small samples occur if $\beta_S \neq \beta_N$. This appears to be related to the additional uncertainty that is introduced by the need to estimate the type probabilities. In fact, when we assume that types are unknown but type probabilities are known, the time preferences as well as the per-period parameters are again very close to their true values. In summary, we consider the evidence from the Monte Carlo simulations to be encouraging enough to conduct a meaningful empirical analysis for the case with three unobserved types, unknown population type probabilities, and two distinct present bias parameters, but we also present the results under the more restrictive assumption of $\beta_S = \beta_N$.

4.3 Structural Estimation Results

We estimate the model outlined in Section 4.2 using Maximum Likelihood, and relegate the derivation of the objective function to online Appendix E. We estimate the type probabilities conditional on the type signal r, { $\mathbb{P}(\tau|r) : \tau \in \{\tau_C, \tau_N, \tau_S\}$ } as well as conditional upon both the type signal and first period

¹⁶This has an effect on the value functions when not buying any contract and when choosing not to retreat after having bought a contract in subsequent periods. We suppress the dependence of the value functions on first-period ownership of an old net to ease exposition.

choice, $\{\mathbb{P}(\tau|r, a_1) : \tau \in \{\tau_C, \tau_N, \tau_S\}\}$. The former provide an estimate of the distribution of types in the population conditional on an observed type predictor, objects that have not previously been estimated in the literature and that are of direct interest. The latter probabilities estimate the distribution of types conditional on purchase; a comparison of the two provides a measure of the attractiveness of commitment products to sophisticated types (recall that agent types are unobserved and purchase decisions are not assumed to deterministically reveal type).

In our main specification, we estimate three time preference parameters, i.e. the discount factor, δ , and two type-specific present-bias parameters β_N and β_S . We also include type indicators in the per-period utility functions as well as household size, households assets, and a measure of household risk aversion. The type-specific dummies allow take-up and retreatment decisions to vary by type for reasons unrelated to differences in time-preferences.

We adopt a sufficiency criterion for the identification of population type probabilities that is strictly weaker than the assumptions required for the known types case (i.e. assumptions that map from (r, a_1) directly to types). In particular, we require that the proxy r is informative about types in a monotone likelihood ratio sense. This condition 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 relative to r = 0. Formally, we require that for some $r \neq r'$, the three ratios $\left\{\frac{\pi_{\tau_C}(r)}{\pi_{\tau_C}(r')}, \frac{\pi_{\tau_N}(r)}{\pi_{\tau_N}(r')}, \frac{\pi_{\tau_S}(r)}{\pi_{\tau_S}(r')}\right\}$ can be strictly ordered ex-ante, see online Appendix B.3 for more details. Recall that in the context of the model, preference reversals are an imperfect proxy for time-inconsistency and are potentially affected by measurement and cognitive issues as well as factors such as seasonality and other constraints. The main requirement is that they shift type probabilities. Examining the key ratio $\frac{\pi_{\tau}(r=0,b)}{\pi_{\tau}(r=1,c)}, \tau \in \{C, N, S\}$, and the difference in the type probabilities depending on r, the sufficiency criterion for type identification is clearly met for our baseline model.

We begin by discussing the estimated population type probabilities presented in the top panel of Table 4. We estimate that 21% of agents are time consistent and that the majority of the time-inconsistent agents (about half of the total population) are naïve. The fraction of time-consistent agents is higher for the subpopulation that did not exhibit preference reversals (i.e. r = 0), while the fraction of sophisticated agents is higher for the subpopulation for which r = 1 as shown in the second panel of Table 4.

We next examine the informativeness of the standard and commitment contracts by estimating type probabilities conditional on first period choice and r which are presented in the bottom panel of Table 4. 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 two indicators. For instance, among the sub-population of agents who exhibit preference reversals (i.e., had r = 1) and purchase either b or c, almost 40 percent are estimated to be time-consistent. In contrast, recall that the directly observed types model assumes that $\pi_C(0, \cdot) = 1$, $\pi_N(1, b) = 1$, and $\pi_S(1, c) = 1$.

Overall, across all combinations of (r, a_1) and conditional upon any net purchase, time-consistent agents account for about 45 percent of all purchases despite only comprising about one fifth of the total population. Naïve inconsistent agents account for 42 percent of total purchases while accounting for 49 percent of the total population. The remaining 13 percent of purchases are made by sophisticated agents who are approximately 30 percent of the total population.

We next explore differences in the type-specific adoption probabilities by focusing on the estimated preference parameters. Table 5 shows the estimation results for the baseline model and for several alternative specifications. The baseline results in the first column indicate that the exponential discount factor is equal to one, which implies that time-consistent households do not significantly discount future periods relative to the present (for the time horizons relevant for our study). However, the two timeinconsistent types dramatically discount the future relative to the present: the estimated present-bias coefficients are $\beta_N = 0.159$ for naïve time-inconsistent households and $\beta_S = 0.077$ for sophisticated timeinconsistent households. Thus, the high present-bias of a large part of the population can rationalize the low adoption of ITNs despite the substantially higher expected cost of malaria when not using an ITN.

The second column of Table 5 shows that the results remain almost identical when we do not estimate any per-period utility parameters suggesting that differences across types in the per-period utility functions are unimportant in explaining adoption rates relative to the differences in time-preference parameters. When we impose $\beta_N = \beta_S$ (column 3), the results remain qualitatively similar, with the estimated present-bias parameter lying in between the two estimated present-bias parameters in the baseline specification. Columns 4 and 5 present results assuming that types are observed (based on a deterministic mapping from (a_1, r) to types) or that there is a single time-consistent type for the whole population, respectively. The results are quite different from our baseline specification: in the known type case, the estimated discount factor is 0.330, and in the single type case it is 0.336. Furthermore, in the known type case the present-bias parameters are much higher than in the baseline case, with $\hat{\beta}_S$ being almost indistinguishable from 1. Recall that 80% of respondents report r = 0 so that in the known types case they are labeled as time-consistent by assumption. Thus, in both of these cases, the vast majority of households are time-consistent by construction. It thus seems reasonable that the discount factor has to be low enough to rationalize the low ITN adoption rate given the high expected costs of malaria.

As a further robustness check, we estimate a version of our baseline model in which we allow the discount factor to differ for the time-consistent type (δ_C) and for the time-inconsistent types (δ_{NS}) . Following our identification proofs, in this specification all the time-preference parameters are still point-identified. This specification nests models in which there are only two time-consistent types (i.e. if $\delta_C \neq \delta_{NS}$ and $\beta_S = \beta_N = 1$), or at most two time-consistent types and one present-biased type (i.e. if either β_N or β_S is equal to 1). However, when estimating the model allowing for two distinct discount factors, the results are virtually identical to our baseline model, with $\delta_C = \delta_{NS} = 1$ and all other parameters being virtually unchanged. This suggest that our three-type model is preferred to a model with fewer types and to an alternative three-type model that includes multiple time-consistent types.¹⁷

To shed further light on the discrepancies between our results and those of the single-type or known type case, we also conduct a set of "placebo" exercises that present results from simulations that estimate a misspecified model. We examine two forms of misspecification: (a) the data are generated from a population with three types but the researcher maximizes a likelihood assuming only one (consistent) type; (b) the data are generated from a population with one time-consistent type but the researcher maximizes a likelihood that assumes the three types from our baseline model (i.e., consistent, naïve, and sophisticated). The results are presented in Table G.5 in online Appendix G. Under scenario (a), we estimate the discount factor δ to be 0.36, which is close to our empirical results when we impose one time-consistent type. Under scenario (b), the estimated discount factor is somewhat higher than the

¹⁷While in principle one can also test for the total number of types in the population as we have outlined above, we do not pursue this here because of our relatively small sample size.

true underlying discount factor while both the present-bias parameters are strictly less than 1 (albeit imprecisely estimated). The model estimates nearly three-fifths of agents to be consistent. This suggests that if our baseline model had been mis-specified by falsely assuming the existence of time-inconsistent types, the estimates of β_S and β_N could plausibly have been large relative to δ while in fact we show that both estimates are very small.

These results suggest that if the one type model was true, our baseline model should yield significantly different results. On the other hand, if the baseline model is true, the misspecified one-type model in our simulations yields an estimated discount factor similar to that produced by estimating our application with a single time-consistent type. We interpret this as further support for our model relative to the alternatives. These results also highlight the importance of separately identifying the population type distributions and time preference parameters.

4.4 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 consequences of timeinconsistency measured in terms of additional costs from sickness.

Changing Retreatment Prices: We first discuss the consequences of doubling the price of retreatment, balanced by a corresponding increase in the price of the commitment contract.¹⁸ Intuitively, the price change has several effects. First, the increase will reduce contemporaneous demand for retreatment through a substitution and income effect. Second, the price increase may reduce overall ITN adoption in the first period, because the dynamic nature of the problem implies that agents predict that the cost of maintaining the protective power of the net with the treatment has increased. Third, a sophisticated agent who cares about retreatment may switch from the standard to the commitment contract, anticipating that present-bias problems will be exacerbated in future periods because of the higher cost of re-treatment. 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.

Averaging across types, demographics and states, we find that after a doubling of the retreatment price from an initial retreatment price of Rs. 16.5 to Rs 33 per bednet, retreatment rates under contract b decline by about 8.8% (see Table 6). We find no effect on retreatment decisions under contract c, since retreatment price increases have no effect on those who have committed to retreatment. Demand for c declines by 3.8% while demand for b decreases by 1.5%. This suggests a substitution from buying c to either buying b or not buying at all, and a substitution from buying b to not buying at all.

We further examine (also in Table 6) changes in take-up and retreatment when the price of retreatment is halved (to Rs 8.25 per bednet). First, we find that retreatment rates for buyers of contract b increase by 4.4%. Second, we find that overall purchase of the standard contract slightly increases by 0.1% while demand for the commitment contract increases by 1.9%.

¹⁸Counterfactuals without the corresponding increase in the price of c imply that demand for c increases unambiguously. We omit these results here. For other recent examples of identification of counterfactual policy interventions in dynamic discrete choice models see Aguirregabiria (2010), Norets and Tang (2013), and Arcidiacono and Miller (2019).

Finally, we estimate the effects of removing the retreatment option for the b contract. In this case, 66% of households that had previously bought a baseline contract switch to a different option, and of those who switch, 54% adopt a commitment contract instead.

Quantifying the effect of time-inconsistency on health costs: Present-bias reduces the present value of purchasing an ITN and thus reduces demand. This in turn increases the probability that time-inconsistent agents contract malaria – relative to the probability for otherwise identical but time-consistent agents. A natural next step is thus to use a counterfactual exercise to estimate the resulting increase in health costs due to medical treatment and lost wages.

We provide a broad outline here and relegate the details of the construction to online Appendix F. First, we compute purchase and retreatment probabilities using the parameters from our baseline model but setting $\beta_S = \beta_N = 1$ (i.e. assuming no present-bias) for each agent in the sample. Next, we use these probabilities to compute the expected costs of malaria for each agent.¹⁹ We then compare this expected cost to the actual expected cost for each agent (i.e. using all the estimated parameters from our baseline model) starting with period 2 (i.e. the first period in which period 1 actions affect health) and summing across periods without discounting.

Table 7 presents the results from using each measure. Across both sets of results, we find that presentbias substantially increases expected costs from malaria. The median cost associated with presentbias is Rs. 1216 (using the numbers from the meta-analysis in (Lengeler, 2009)) or Rs. 1624 (using household beliefs on net efficacy). Overall, present-bias leads to a median reduction of between 6.3 and 8.8 workdays per malaria episode. Even though these costs are high relative to the cost of an ITN, the estimated preference parameters make it clear that investing in ITNs is not a very attractive option for the median present-biased household (naïve or sophisticated) relative to a time-consistent household. This provides concrete empirical evidence of an important dichotomy raised in theoretical treatments of time-inconsistency: a long-run self and a social planner with sufficiently high discount rates will have a strong preference for encouraging ITN adoption to reduce long-run health costs and increase productivity. However, time-inconsistent households do not find ITN purchases particularly attractive. This also suggests that small subsidies may not significantly increase ITN adoption.

5 Conclusions

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. Importantly, we are also able to identify type distributions i.e. the fraction of time-consistent, naïve, and sophisticated agents. We further extend the identification results to any finite set of types in the population. Our Monte Carlo simulations suggest that the

¹⁹The key ingredients in the expected cost calculation are the probability of contracting malaria when sleeping under an ITN relative to an untreated net or no net, and the expected number of workdays lost due to malaria elicited during the baseline survey. We use two alternative measures of the probability of contracting malaria. First, we use the household-specific elicited beliefs about the efficacy of ITNs, untreated nets and sleeping without a net. Second, we use the meta-analysis in Lengeler (2009) that concludes that "in areas with stable malaria, ITNs reduced the incidence of uncomplicated malarial episodes in areas of stable malaria by 50% compared to no nets, and 39% compared to untreated nets." While the latter is clearly an extrapolation, it provides an alternative measure of the efficacy of ITNs relative to our survey measures.
individual time-preference parameters of interest and the associated population type probabilities can indeed be estimated with reasonable precision.

We estimate the model on a specifically collected dataset containing detailed information on beliefs combined with a field intervention. Our empirical results suggest that time-inconsistency is a strong predictor of investment in a preventive health technology. We estimate that time-inconsistent agents account for almost 80 percent of the population, with about 50 percent of the population being naïve timeinconsistent. While the standard exponential time-preference parameter is close to 1, time-inconsistent types are substantially present-biased, with estimated present-bias parameters of 0.159 (for naïve types) and 0.077 (for sophisticated types).²⁰ We find that present-bias among sophisticated households is so pronounced that our specifically designed commitment products are not particularly appealing to them (the purchase of these products is in fact higher among naïve households).

Estimating models with a single time-consistent type or pre-determined types (as standard in earlier work) leads to significantly different results, in particular to a low exponential discount factor. We provide further evidence for our preferred specification from a set of placebo simulations. Overall, our results highlight the importance of separately identifying the type distribution, time preferences, and the other utility parameters.

Our identification strategy can also be applied to other contexts. Key variables in the identification strategy are the "excluded" variables z that affect future, but not current utility. Besides directly eliciting beliefs by asking households, as is increasingly common in surveys (Manski, 2004; Delavande et al., 2010; Delavande, 2014), one could use other available data that similarly indicate the future value of an action to generate exclusion restrictions. One example could be firms' disclosed expectations regarding the return on a specific investment when it is being announced. In such a context, De Groote and Verboven (2019) use an alternative restriction in a model with only time-consistent agents by assuming that the discount factor for adopting an investment is the same as the one relevant for weighing investment costs against future benefits.

To recover population type probabilities we require a signal r that is correlated with time-inconsistency but uninformative about type once conditioned on. In our case, we use a specifically designed intervention to generate such observables, but a similar strategy can be followed in other contexts where researchers are designing their own interventions or data-collection. In addition, in other contexts there may be other data that could plausibly be informative about self-control problems (e.g., data on binge-watching). If there is evidence that certain consumption patterns are associated with agents having less self-control, then such information can also be used (provided they do not affect utility directly).

Our estimates suggest that the magnitude of the present-bias is large enough (both in terms of the estimated present-bias parameters as well as in terms of the large fraction of inconsistent agents in the population) to affect the adoption of ITNs, despite their proven ability to reduce malaria. Small or partial subsidies may thus only have limited effects on adoption, consistent with recent research that argues that, in poor areas, free provision may be the only way to ensure universal coverage for important health-related products (Kremer and Miguel, 2007; Cohen and Dupas, 2010).

 $^{^{20}}$ Bisin and Hyndman (2018) also find that present-bias is more pronounced among sophisticated individuals relative to naïve ones in an experiment among U.S. students. Unlike us, however, they find that the hyperbolic discount factor among naïve individuals is close to one.

References

- Abbring, J. H. and Ø. Daljord (2018). Identifying The Discount Factor In Dynamic Discrete Choice Models. Quantitative Economics (forthcoming). [4, 5, 8, 9]
- Abbring, J. H. and Ø. Daljord (2019). A Comment on Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting by Hanming Fang and Yang Wang. International Economic Review (forthcoming). [5]
- Abbring, J. H., Ø. Daljord, and F. Iskhakov (2019). Identifying Present-Biased Discount Functions in Dynamic Discrete Choice Models. *mimeo*. [4, 5, 8]
- Aguirregabiria, V. (2010). Another look at the identification of dynamic discrete decision processes: An application to retirement behavior. Journal of Business & Economic Statistics 28(2), 201–218. [34]
- Aguirregabiria, V. and P. Mira (2010). Dynamic discrete choice structural models: A survey. Journal of Econometrics 156(1), 38–67. [4, 6, 30]
- Andreoni, J., M. Callen, Y. Khan, K. Jaffar, and C. Sprenger (2016). Using Preference Estimates to Customize Incentives: An Application to Polio Vaccination Drives in Pakistan. Technical report, National Bureau of Economic Research. [5]
- Andreoni, J. and C. Sprenger (2012). Estimating time preferences from convex budgets. American Economic Review 102(7), 3333–56. [2, 26]
- Arcidiacono, P. and P. B. Ellickson (2011). Practical Methods for Estimation of Dynamic Discrete Choice Models. Annu. Rev. Econ. 3(1), 363–394. [4]
- Arcidiacono, P. and R. A. Miller (2019). Identifying dynamic discrete choice models off short panels. Journal of Econometrics. [34]
- Ashraf, N., D. Karlan, and W. Yin (2006). Tying Odysseus to the mast: evidence from a commitment savings product in the Philippines. *Quarterly Journal of Economics* 121(2), 635–672. [2, 5]
- Augenblick, N., M. Niederle, and C. Sprenger (2015). Working Over Time: Dynamic Inconsistency in Real Effort Tasks. Quarterly Journal of Economics 130(3), 1067–1115. [2, 5]
- Bai, L., B. Handel, E. Miguel, and G. Rao (2017). Self-control and demand for preventive health: Evidence from hypertension in india. Technical report, National Bureau of Economic Research. [5]
- Banerjee, A., E. Duflo, R. Glennerster, and D. Kothari (2010). Improving immunization coverage in rural India: A clustered randomized controlled evaluation of immunization campaigns with and without incentives. British Medical Journal 340(c2220). [2]
- Banerjee, A. and S. Mullainathan (2010). The Shape of Temptation: Implications for the Economic lives of the Poor. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1601724. NBER Working Paper w15973. [2]
- Bauer, M., J. Chytilová, and J. Morduch (2012). Behavioral foundations of microcredit: Experimental and survey evidence from rural india. American Economic Review 102(2), 1118–39. []
- Bernheim, B. D., D. Ray, and S. Yeltekin (2015). Poverty and Self-Control. Econometrica 83(5), 1877–1911. [2]
- Bisin, A. and K. Hyndman (2018). Present-bias, procrastination and deadlines in a field experiment. National Bureau Economic Research Working Paper 19874, http://www.nber.org/papers/w19874. [36]
- Bryan, G., D. Karlan, and S. Nelson (2010). Commitment devices. Annual Reviews of Economics 2, 671–698. [4]
- Carvalho, L. S., S. Meier, and S. W. Wang (2016). Poverty and Economic Decision-Making: Evidence from Changes in Financial Resources at Payday. *American Economic Review* 106(2), 260–284. [2]
- Cohen, J. and P. Dupas (2010). Free distribution or cost-sharing? Evidence from a randomized malaria prevention experiment. *Quarterly Journal of Economics* 125(1), 1–45. [36]
- De Groote, O. and F. Verboven (2019). Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems. *American Economic Review* 109(6), 2137–72. [36]
- Delavande, A. (2014). Probabilistic expectations in developing countries. Annual Reviews of Economics 6, 1–20. [8, 36]
- Delavande, A., X. Giné, and D. McKenzie (2010). Measuring subjective expectations in developing countries: A critical review and new evidence. *Journal of Development Economics* 94(2), 151–163. [8, 36]
- DellaVigna, S. (2009). Psychology and economics: Evidence from the field. Journal of Economic Literature 47(2), 315-372. [2]
- DellaVigna, S. (2017). Structural behavioral economics. Forthcoming in Handbook of Behavioral Economics, eds. David Laibson, Douglas Bernheim, and Stefano DellaVigna. [4]
- Dhingra, N., P. Jha, V. Sharma, A. Cohen, R. Jotkar, P. Rodriguez, D. Bassani, W. Suraweera, R. Laxminarayan, and R. Peto (2010). Adult and child malaria mortality in India: a nationally representative mortality survey. *Lancet* 376(9754), 1768–1774. [24]

- Duflo, E., M. Kremer, and J. Robinson (2011). Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya. American Economic Review 101(6), 2350–2390. [2]
- Ericson, K. M. and D. Laibson (2018, December). Intertemporal choice. NBER Working Number 25358 http: //www.nber.org/papers/w25358. [2]
- Fang, H. and D. Silverman (2009). Time-inconsistency and welfare program participation: evidence from the NLSY. International Economic Review 50(4), 1043–1077. [4]
- Fang, H. and Y. Wang (2015). Estimating dynamic discrete choice models with hyperbolic discounting, with an application to mammography decisions. *International Economic Review* 56(2), 565–596. [4, 5]
- Frederick, S., G. Loewenstein, and E. D. O'Donoghue (2002). Time discounting: A critical review. Journal of Economic Literature 40(2), 351–401. [2]
- Giné, X., J. Goldberg, D. Silverman, and D. Yang (2018). Revising Commitments: Field Evidence on the Adjustment of Prior Choices. *The Economic Journal* 128(608), 159–188. [7]
- Gul, F. and W. Pesendorfer (2001). Temptation and self-control. *Econometrica* 69(6), 1403–1435. [7]
- Gul, F. and W. Pesendorfer (2004). Self-control and the theory of consumption. Econometrica 72(1), 119–158. [7]
- Hardy, G. H. (1908). Mendelian proportions in a mixed population. Classic papers in genetics. Prentice-Hall, Inc.: Englewood Cliffs, NJ, 60–62. [28]
- Heidhues, P. and P. Strack (2019). Identifying Present-Bias from the Timing of Choices. mimeo. [5]
- Holt, C. A. and S. K. Laury (2002, December). Risk Aversion and Incentive Effects. American Economic Review 92(5), 1644–1655. [27, 12]
- Hotz, V. and R. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. The Review of Economic Studies 60(3), 497–529. [4, 9, 48, 11]
- Kasahara, H. and K. Shimotsu (2009, 01). Nonparametric identification of finite mixture models of dynamic discrete choices. *Econometrica* 77(1), 135–175. [5, 53, 6]
- Kaur, S., M. Kremer, and S. Mullainathan (2014). Self-control at work. Journal of Political Economy. [2]
- Kremer, M. and E. Miguel (2007). The illusion of sustainability. *Quarterly Journal of Economics* 122(3), 1007–1065. [36]
- Laibson, D. I. (1997). Golden eggs and hyperbolic discounting. Quarterly Journal of Economics 112(2), 443–477.
- Lengeler, C. (2009). Insecticide-treated bed nets and curtains for preventing malaria. Cochrane database of systematic reviews (Online). Issue 2. Art. No.: CD000363. DOI: 10.1002/14651858.CD000363.pub2. [4, 35, 15]
- Magnac, T. and D. Thesmar (2002). Identifying Dynamic Discrete Decision Processes. *Econometrica* 70(2), 801–816. [2, 4, 8, 10]
- Manski, C. (2004, September). Measuring Expectations. *Econometrica* 72(5), 1329–1375. [8, 36]
- Martinez, S.-K., S. Meier, and C. Sprenger (2019). Procrastination in the field: Evidence from tax filing. *mimeo*. [5]
- Mullainathan, S. (2004). Development economics through the lens of psychology. Available at http://www.economics.harvard.edu/admin/person_pages/edit/html/files/LENS%200F%20PSYCH0L0GY.pdf. [2]
- Norets, A. and X. Tang (2013). Semiparametric inference in dynamic binary choice models. *Review of Economic Studies* 81(3), 1229–1262. [34]
- O'Donoghue, T. and M. Rabin (1999). Doing it now or later. American Economic Review 89(1), 103–124. [7]
- Paserman, M. (2008, August). Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation. The Economic Journal 118(531), 1418–1452. [4]
- Roth, C. and J. Wohlfart (2018). How do expectations about the macroeconomy affect personal expectations and behavior? [8]
- Rubinstein, A. (2003). "Economics and psychology"? The case of hyperbolic discounting. International Economic Review 44(4), 1207–1216. [26]
- Rust, J. (1994). Structural estimation of markov decision processes. In R. Engle and D. McFadden (Eds.), Handbook of Econometrics, pp. 3081–3143. Elsevier Science. [2, 7]
- Schilbach, F. (2019). Alcohol and Self-Control: A Field Experiment in India. American economic review 109(4), 1290–1322. [2]
- Shleifer, A. (2019). The return of survey expectations. NBER Reporter 2019(1), 14. [8]
- Sprenger, C. (2015, May). Judging experimental evidence on dynamic inconsistency. American Economic Review 105(5), 280–85. [2]
- Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. The Review of Economic Studies 23(3), 165–180. [7]
- Tarozzi, A., A. Mahajan, B. Blackburn, D. Kopf, L. Krishnan, and J. Yoong (2014). Micro-loans, bednets and

malaria: Evidence from a randomized controlled trial in Orissa (India). American Economic Review 104(7), 1909–1941. [2, 24, 25]

- Tarozzi, A., A. Mahajan, J. Yoong, and B. Blackburn (2009). Commitment mechanisms and compliance with health-protecting behavior: Preliminary evidence from Orissa, India. American Economic Review (Papers and Proceedings) 99(2), 231–235. [2]
- Toussaert, S. (2018). Eliciting temptation and self-control through menu choices: A lab experiment. *Econometrica* 86(3), 859–889. [7]
- Ubfal, D. (2016). How General Are Time Preferences? Eliciting Good-Specific Discount Rates. Journal of Development Economics 118(C), 150–170. [2]
- van der Klaauw, W. (2012). On the use of expectations data in estimating structural dynamic choice models. Journal of Labor Economics 30(3), 521 – 554. [5]
- van der Klaauw, W. and K. Wolpin (2008). Social security and the retirement and savings behavior of low income households. *Journal of Econometrics* 145(1-2), 21–42. [5]
- Weinberg, W. (1908). On the demonstration of heredity in man. (1963) Papers on human genetics. [28]
- World Bank (2008). Global purchasing power parities and real expenditures, 2005. Technical report, International Comparison Program, Washington DC: World Bank. Available at www.worldbank.org/data/icp. [41]



Notes: Histograms of subjective beliefs about the protective power of bednets and treatment with insecticide from malaria risk. Data from March-April 2007 baseline survey, n = 566.

	Mean	Median	S.d.	Obs.
Household size	5.3	5	2.1	566
no. children under 5	0.47	0	0.7	566
Head is male	0.94	1	0.24	566
Head age	45	45	12	566
Head at least secondary school	0.11	0	0.32	554
Head any schooling	0.71	1	0.45	560
Total monthly expenditure per head	753	607	574	566
Bednets per head	0.32	0.25	0.31	562
ITNs per head	0.059	0	0.19	561
At least one bednet	0.68	1	0.47	562
Fraction of member slept under bednet last night	0.16	0	0.32	566
Fraction of member slept under ITN last night	0.032	0	0.16	564
Fraction of member sleeps under net in peak malaria season	0.56	0.79	0.46	566
Fraction of members +ve to malaria	0.11	0	0.29	522
Fraction of members anemic $(\text{Hb} < 11g/dl)$	0.46	0.5	0.46	514
Aware mosquito bites can cause malaria	0.96	1	0.19	566
Aware bednets can protect against malaria	0.96	1	0.19	566
Expected cost of a malaria episode (working man) (Rs.)	2919	2330	2383	566
Expected cost of a malaria episode (non-working) (Rs.)	1753	1400	1537	566
Cost of recent (actual) malaria episodes (Rs.)	700	0	1928	566
Cost of recent (actual) malaria episodes (Rs.), if > 0	1737	855	2729	228

 Table 1: Baseline Summary Statistics

Notes: Data from March-April 2007 baseline survey. Data from 566 households. All means as un-weighted averages across sample households. The varying sample size for different variables is explained by missing values. Mean expenditure per head was measured asking about usual consumption of 18 item categories, including home production of foodstuff. Both the actual and expected costs of malaria episodes were elicited using an itemized list including doctor fees, drugs and tests, hospitalization, surgery, costs of lodging and transportation (including those for any caretaker), lost earnings from days of lost work, and cost of non-household members hired to replace the sick at work. Costs of recent malaria episodes refer to all health episodes in the household reported as malaria by the respondents, during the six months before the interview. All monetary values are in nominal Rs. (PPP exchange rate $\approx 16Rs/USD$, World Bank, 2008).

Table 2: Baseline Time Preferences

Prefers Rs. 10 in 1 month to Rs. 10 in 4 months	0.84
Prefers Rs. 10 in 1 month to Rs. 12 in 4 months	0.71
Prefers Rs. 10 in 1 month to Rs. 14 in 4 months	0.65
Prefers Rs. 10 in 1 month to Rs. 16 in 4 months	0.60
Prefers Rs. 10 in 1 month to Rs. 10 in 7 months	0.82
Prefers Rs. 10 in 1 month to Rs. 15 in 7 months	0.63
Prefers Rs. 10 in 1 month to Rs. 20 in 7 months	0.52
Prefers Rs. 10 in 1 month to Rs. 25 in 7 months	0.49
Prefers Rs. 10 in 4 months to Rs. 10 in 7 months	0.84
Prefers Rs. 10 in 4 months to Rs. 12 in 7 months	0.74
Prefers Rs. 10 in 4 months to Rs. 14 in 7 months	0.65
Prefers Rs. 10 in 4 months to Rs. 16 in 7 months	0.57
Always prefers earlier reward	0.27
At least one "hyperbolic" preference reversal	0.25
Mean no. of "hyperbolic" preference reversals (> 0)	1.31

Notes: Data from March-April 2007 survey. n = 566. "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 mean in the last row is calculated including only the 147 respondents who displayed at least one hyperbolic preference reversal.

	Mean
(A) Purchase	
Purchased at least one ITN	0.51
# ITNs purchased, any contract	1.03
# ITNs purchased, any contract, if > 0	2.03
Purchased at least one ITN without 'commitment' to retreat (b)	0.25
# ITNs purchased, without 'commitment' to re-treatments (b)	0.44
# ITNs purchased, without 'commitment' to re-treatments (b), if > 0	1.76
Purchased at least one ITN with commitment' to 2 re-treatments (c)	0.26
# ITNs purchased, with 'commitment' to re-treatments (c)	0.59
# ITNs purchased, with 'commitment' to re-treatments (c), if > 0	2.29
(B) Re-treatment	
% Bednets re-treated after 6 months	
without 'commitment' to retreat (b)	0.36
with 'commitment' to retreat (c)	0.92
% Bednets re-treated after 12 months	
without 'commitment' to retreat (b)	0.19
with 'commitment' to retreat (c)	0.84

Table 3: Summary of purchases

Notes: Data from September-November 2007.n=566.

$\pi_{ au}$	Estimate	2.5	97.5
π_C	0.206	0.194	0.219
π_N	0.494	0.490	0.499
π_S	0.299	0.283	0.316
$\pi_{ au}(r)$	Estimate	2.5	97.5
$\pi_C(0)$	0.217	0.206	0.228
$\pi_N(0)$	0.498	0.496	0.499
$\pi_S(0)$	0.285	0.273	0.297
$\pi_C(1)$	0.158	0.140	0.177
$\pi_N(1)$	0.479	0.470	0.489
$\pi_S(1)$	0.363	0.335	0.391
$\pi_{\tau}(r,a_1)$	Estimate	2.5	97.5
$\pi_C(0,1)$	0.430	0.412	0.447
$\pi_N(0,1)$	0.440	0.383	0.496
$\pi_S(0,1)$	0.131	0.085	0.176
$\pi_C(1,1)$	0.305	0.281	0.329
$\pi_N(1,1)$	0.506	0.432	0.580
$\pi_S(1,1)$	0.189	0.122	0.256
$\pi_C(0,2)$	0.510	0.495	0.526
$\pi_N(0,2)$	0.379	0.321	0.438
$\pi_S(0,2)$	0.110	0.066	0.154
$\pi_C(1,2)$	0.390	0.357	0.423
$\pi_N(1,2)$	0.444	0.365	0.522
$\pi_{S}(1,2)$	0.166	0.098	0.234

 Table 4: Type Probabilities

Notes: π_{τ} is the unconditional probability that an agent in the population is of type τ , where *C* refers to a timeconsistent agent, *N* refers to a time-inconsistent naïve agent, and *S* refers to a time-inconsistent sophisticated agent. $\pi_{\tau}(r)$ is the probability that an agent is of type τ given their response to the time-inconsistency question *r*. $\pi_{\tau}(r, a)$ 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 . The second and third columns are the 2.5th and 97.5th percentiles of the distribution of the type-probabilities computed using the delta method.

	(1)	(2)	(3)	(4)	(5)
Variable	Full model	No ϕ	$\beta_N = \beta_S$	Known types	One type
δ	1.000(0.031)	$1.000 \ (0.059)$	1.000(0.064)	0.330(0.004)	$0.336\ (0.000)$
β_N	0.159(0.019)	$0.159\ (0.007)$	$0.129\ (0.013)$	$0.650\ (0.112)$	
β_S	0.077(0.024)	$0.077 \ (0.030)$		$0.999\ (0.039)$	
ϕ_0	0.093(0.049)		0.600(0.216)	$1.313\ (0.056)$	$0.001 \ (0.012)$
$\phi_{ m Naiv}$	$0.037\ (0.083)$		$0.022 \ (0.092)$	$1.522 \ (0.020)$	
$\phi_{ m Soph}$	$0.030\ (0.083)$		$0.093\ (0.606)$	$0.030\ (0.112)$	
$\phi_{ m HHS}$	$0.056\ (0.033)$		$0.037 \ (0.149)$	$0.056\ (0.163)$	$0.006\ (0.045)$
$\phi_{\rm Assets}$	$0.095\ (0.079)$		$0.030\ (0.087)$	$0.095\ (0.377)$	$0.027 \ (0.014)$
ϕ_{Risk}	$0.060\ (0.028)$		$0.056\ (0.049)$	$0.060\ (0.050)$	$0.003\ (0.036)$

Table 5: Time Preferences (Unobserved Types)

Notes: The first column shows the estimated preference parameters for the baseline model with three unobserved types including one common discount factor, two present-bias parameters, and several per-period utility parameters. The second column presents the estimated preference parameters for the estimation without per-period utility parameters. The third column presents the results when we impose $\beta_S = \beta_N$. The fourth column presents the results when we impose $\beta_S = \beta_N$. The fourth column presents the results when types are uniquely identified by (r, a_1) , implying that r = 0 reflects a time-consistent type, $(r = 1, a_1 = c)$ reflects a time-inconsistent sophisticated type, and $(r = 1, a_1 = n)$, $(r = 1, a_1 = b)$ both reflect time-inconsistent naïve type. Column 5 presents the results under the assumption of a single time-consistent type. Standard errors are in parentheses.

Table 6: Counterfactuals: Change in Take up and Retreatment Rates

Outcome	Double retreatment price	Half retreatment price
% Change b Take up	-1.482	0.817
	(0.021)	(0.012)
% Change c Take up	-3.838	1.865
	(0.056)	(0.026)
% Change Retreatment b	-8.822	4.360
	(0.033)	(0.012)

Notes: All changes are relative to the retreatment price of Rs.16.5 per bednet. 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. Standard errors are in parentheses.

	(1)	(2)	(3)	(4)	(5)	(6)
(A)	$\Delta \text{Cost } t = 2$	$\Delta \text{Cost } t = 3$	$\Delta \text{Cost } t = 4$	ΔC_{total}	$\Delta C_{total} \cdot \beta_N$	$\Delta C_{total} \cdot \beta_S$
Cost ind. beliefs	616.000	485.462	522.194	1623.655	257.950	124.426
	(5.970)	(8.943)	(28.400)	(43.313)	(6.881)	(3.319)
Cost population prob.	479.058	354.397	382.511	1215.966	193.180	93.183
	(19.001)	(13.516)	(24.743)	(57.260)	(9.097)	(4.388)
(B)	$\Delta \text{Days } t = 2$	$\Delta \text{Days } t = 3$	$\Delta \text{Days } t = 4$	$\Delta Days_{total}$	$\Delta Days_{total} \cdot \beta_N$	$\Delta Days_{total} \cdot \beta_S$
Missed days ind. beliefs	3.411	2.677	2.748	8.837	1.404	0.677
	(0.067)	(0.048)	(0.126)	(0.242)	(0.038)	(0.019)
Missed days pop. prob.	2.526	1.803	1.991	6.319	1.004	0.484
	(0.133)	(0.037)	(0.252)	(0.423)	(0.067)	(0.032)

Table 7: Median cost of malaria and days missed attributable to time-inconsistent preferences

Notes: Panel A presents the expected additional costs of malaria that is attributable to the lower investment into ITNs and retreatment because of present-bias. The first specification is computed based on the stated beliefs of malaria risk, while the second specification uses an aggregate estimate from the literature. The first 3 columns present the median changes in expected cost for periods 2-4 in the population. Column 4 presents the total median cost changes for all three periods. Columns 5 and 6 present the total median cost changes discounted with the different estimated present-bias parameters. Panel B presents the same statistics for the median expected days missed at work instead of the median expected cost of malaria. Standard errors are in parentheses.

Appendix

A Proofs

A.1 Proofs for Identification of Directly Observed Types

Proof of Lemma 1

Proof. The probability that an agent in period 3 chooses action k conditional upon state x_3 is given by

$$\mathbb{P}_{\tau}(a_3^* = k | x_3) = \mathbb{P}\left(k = \operatorname*{argmax}_{a \in \mathcal{A}_3} \left\{ u_3(x_3, a; \tau) + \epsilon_3(a) \right\} | x_3 \right).$$

The decision in the terminal period is a standard static discrete choice model and with a known error distribution we can invert the relationship (see Hotz and Miller (1993) or see online Appendix C for a self-contained argument) to directly identify the functions $u_3(x_3, k; \tau) - u_3(x_3, 0; \tau)$. The normalization for period 3 utility (Assumption 6) ensures that $u_3(x_3, k; \tau)$ is identified.

Next, note that because period 3 utility is identified and the error distribution is assumed to be known, the expected value function $\int v_{\tau,3}^*(s_3) dF(s_3|x_2, k, z_2)$ is also identified. Turning now to period 2, the probability that an agent of type τ will choose action k given state variables x_2 and beliefs z_2 is given by

$$\mathbb{P}_{\tau}(a_{2}^{*}=k|x_{2},z_{2}) = \mathbb{P}\left(k = \operatorname*{argmax}_{a \in \mathcal{A}_{2}} \left\{u_{2}(x_{2},a;\tau) + \epsilon_{2}(a) + \beta_{\tau}\delta_{\tau}\int v_{\tau,3}^{*}(s_{3})\,\mathrm{dF}(s_{3}|x_{2},a,z_{2})\right\} \middle| x_{2},z_{2}\right)$$

Inverting the type-specific conditional choice probabilities as before (cf. Hotz and Miller (1993)) we can identify the function

$$g_{\tau,2,k}(x_2, z_2) = u_2(x_2, k, \tau) - u_2(x_2, 0, \tau) + \beta_\tau \delta_\tau \int v_\tau^*(s_3) \,\mathrm{dF}_{\Delta,k}(s_3 | x_2, z_2) \tag{34}$$

for all (x_2, z_2, k) . Next, Assumption 7 allows us to express (for at least one action k and two points (z'_2, z''_2) and all $x_2 \in \mathcal{X}_2$)

$$\beta_{\tau}\delta_{\tau} = \frac{g_{\tau,2,k}(x_2, z_2') - g_{\tau,2,k}(x_2, z_2'')}{\int v_{\tau}^*(s_3)(\mathrm{dF}_{\Delta,k}(s_3|x_2, z_2') - \mathrm{dF}_{\Delta,k}(s_3|x_2, z_2''))}$$
(35)

so that $\beta_{\tau} \delta_{\tau}$ is identified for all τ .

Next, (a) substituting (35) into (34), (b) using the fact that $\int v_{\tau,3}^*(s_3) dF(s_3|x_2,k,z_2)$ is identified and (c) Assumption 8 we conclude that the period 2 utility function $u_2(x_2,k;\tau)$ is identified for all $k \in \mathcal{A}_2$.

Proof of Lemma 2

Proof. The (conditional) probability that an agent chooses action k in period 1 is given by

$$\mathbb{P}_{\tau}(a_{1}^{*}=k|x_{1},z_{1}) = \mathbb{P}\left(k = \operatorname*{argmax}_{a \in \mathcal{A}_{1}} \left\{u_{1}(x_{1},a;\tau) + \epsilon_{1}(a) + \beta_{\tau}\delta_{\tau}\int v_{\tau,2}^{*}(s_{2})\,\mathrm{dF}(s_{2}|x_{1},a,z_{1})\right\} \middle| x_{1},z_{1}\right)$$

Inverting the type-specific conditional choice probabilities we next identify the function $g_{\tau,1,k}(\cdot)$:

$$g_{\tau,1,k}(x_1, z_1) = u_1(x_1, k; \tau) - u_1(x_1, 0; \tau) + \beta_\tau \delta_\tau \int v_{\tau,2}^*(s_2) \,\mathrm{dF}_{\Delta,k}(s_2 | x_1, z_1) \tag{36}$$

where we have

$$v_{\tau,2}^*(s_2) \equiv \sum_{a \in \mathcal{A}_2} v_{\tau,2}(s_2, a, 1, \delta_{\tau}) A_{\tau}(s_2, a, \tilde{\beta}_{\tau}, \delta_{\tau})$$

and $v_{\tau,2}(\cdot)$, $A_{\tau}(\cdot)$ are defined in equation (5). Recall that $v_{\tau,2}^*(s_2)$ is the continuation value from period 2 onwards, from the standpoint of period 1, assuming that the event that action a will be chosen in period 2 is given by $A_{\tau}(s_2, a, \tilde{\beta}_{\tau}, \delta_{\tau}) = 1$. The parameter $\tilde{\beta}_{\tau}$ is interpreted as the amount of present-bias that the agent in period 1 thinks his period 2 self will be subject to.

We begin by noting that for consistent agents the last term on the right hand side of (36) is identified so that period 1 preferences are then identified.

Next, we show identification for sophisticated agents. We note that under Assumption 11 the function

$$\int \tilde{h}_{\tau_S}(x_2, z_2, \beta_{\tau_S} \delta_{\tau_S}, \mathsf{d}) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1') - \mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right)$$

is strictly monotone in d and we know that at $d = \delta_{\tau_S}$ the function is equal to the directly identified object

$$\int h_{\tau_S}(x_2, z_2) \left(\mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1') - \mathrm{dF}_{\Delta, k}(x_2, z_2 | x_1, z_1'') \right).$$

Therefore, we solve for the value of d that sets the two functions above equal to each other and thus point identify δ_{τ_s} . Following this we can identify β_{τ_s} since $\beta_{\tau_s} \delta_{\tau_s}$ was identified by Lemma 1.

Proof of Lemma 3

Proof. We begin by first isolating the last expression in (36). Under Assumption (13) we can identify the difference

$$g_{\tau,1,k,\Delta}(x_1) \equiv g_{\tau,1,k}(x_1, z_1') - g_{\tau,1,k}(x_1, z_1'') = \beta_\tau \delta_\tau \left(\int v_{\tau,2}^*(s_2) \bigg(dF_{\Delta,k}(s_2|x_1, z_1') - dF_{\Delta,k}(s_2|x_1, z_1'') \bigg) \bigg)$$

Rewriting this difference using the previously introduced notation and using the fact that $\beta_{\tau}\delta_{\tau}$ is identified we can identify

$$\frac{g_{\tau,1,k,\Delta}(x_1)}{\beta_{\tau}\delta_{\tau}} = \sum_{(x_2,z_2)\in\mathcal{X}_2\times\mathcal{Z}_2} \int v_{\tau,2}^*(s_2) \left(dF(\epsilon_2) dF_{\Delta,k}(x_2,z_2|x_1,z_1') - dF(\epsilon_2) dF_{\Delta,k}(x_2,z_2|x_1,z_1'') \right)
= \sum_{(x_2,z_2)\in\mathcal{X}_2\times\mathcal{Z}_2} h_{\tau}(x_2,z_2) \left(dF_{\Delta,k}(x_2,z_2|x_1,z_1') - dF_{\Delta,k}(x_2,z_2|x_1,z_1'') \right)
= \sum_{(x_2,z_2)\in\mathcal{X}_2\times\mathcal{Z}_2\setminus(x_{20},z_{20})} (h_{\tau}(x_2,z_2) - h_{\tau}(x_{20},z_{20})) \left(dF_{\Delta,k}(x_2,z_2|x_1,z_1') - dF_{\Delta,k}(x_2,z_2|x_1,z_1'') \right).$$
(37)

(38)

We have added $h_{\tau}(x_{20}, z_{20})$ where (x_{20}, z_{20}) is a fixed point in $\mathcal{X}_2 \times \mathcal{Z}_2$ to incorporate the constraint that

$$\sum_{x_2, z_2} \mathrm{dF}_k(x_2, z_2 | x_1, z_1) = 1.$$

Without incorporating this restriction the corresponding $S \times S$ matrix needed in Assumption 13 would not be invertible. Recall that S-1 denotes the cardinality of the (finite) set $\mathcal{X}_2 \times \mathcal{Z}_2 \setminus (x_{20}, z_{20})$. Define the S-1 column vectors $g_{\tau,\Delta}(k)$ and h_{τ} :

$$\mathbf{g}_{\tau,\Delta}(k) \equiv \frac{1}{\beta_{\tau}\delta_{\tau}} \begin{bmatrix} g_{\tau,1,k,\Delta}(x_{1,1}) \\ \vdots \\ g_{\tau,1,k,\Delta}(x_{1,S-1}) \end{bmatrix} ; \mathbf{h}_{\tau} \equiv \begin{bmatrix} h_{\tau}(x_{2,1},z_{2,1}) - h_{\tau}(x_{20},z_{20}) \\ \vdots \\ h_{\tau}(x_{2,S-1},z_{2,S-1}) - h_{\tau}(x_{20},z_{20}) \end{bmatrix}.$$
(39)

Using the notation above we can rewrite (37) in matrix form as

$$\mathsf{g}_{ au,\Delta}(k) = (\mathsf{dF}_\Delta(k,z_1') - \mathsf{dF}_\Delta(k,z_1''))\mathsf{h}_{ au}$$

where $g_{\tau,\Delta}(k)$ is identified by (37) and the matrices $dF_{\Delta}(k, z'_1)$ and $dF_{\Delta}(k, z''_1)$ (defined below (10)) are identified since they are constructed from observed beliefs. Under the (testable) invertibility assumption (Assumption 13) it follows that

$$h_{\tau} = \left(\mathsf{dF}_{\Delta}(k, z_1') - \mathsf{dF}_{\Delta}(k, z_1'')\right)^{-1} \mathsf{g}_{\tau, \Delta}(k), \tag{40}$$

so that $h_{\tau}(x_2, z_2) - h_{\tau}(x_{20}, z_{20})$ is identified. To simplify notation in the statement of the lemma we have defined

$$\mathbf{k}_{\tau} \equiv -h_{\tau}(x_{20}, z_{20})$$

Next, we see that the last expression on the right hand side of the equation below is identified:

$$g_{\tau,1,k}(x_1, z_1) = u_1(x_1, k; \tau) - u_1(x_1, 0; \tau) + \beta_\tau \delta_\tau \int (h_\tau(x_2, z_2) - h_\tau(x_{20}, z_{20})) \,\mathrm{dF}_{\Delta,k}(x_2, z_2 | x_1, z_1),$$

so that first period preferences are identified for all types τ . Note that $\sum_{x_2,z_2} h_{\tau}(x_{20}, z_{20}) dF_{\Delta,k}(x_2, z_2|x_1, z_1) = 0$ since $dF_{\Delta,k}$ is a signed measure with total measure equal to zero.

Alternative Assumptions for Lemma 3

Note that Assumption 13 places a support condition on \mathcal{X}_1 which may be too strong. We now outline an alternative assumption that does not require this assumption. However, since the assumption is relatively cumbersome to state we have not included it in the main text. To begin, recall that K_1 denotes the cardinality of \mathcal{A}_1 and let S_1 denote the cardinality of \mathcal{X}_1 . Define the $((K_1 - 1) \times S_1) \times (S - 1)$ identified matrix (recall that S denotes the cardinality of $\mathcal{X}_2 \times \mathcal{Z}_2$)

ASSUMPTION 24. The matrix $dF_{\Delta}(z'_1) - dF_{\Delta}(z''_1)$ has rank S - 1.

This assumption requires that the number of points in the support of $\mathcal{A}_1 \times \mathcal{X}_1$ be at least as large as S and so relaxes the requirement on \mathcal{X}_1 . We can then state the modified version of Lemma 3.

LEMMA 8. Consider an agent of type τ maximizing equation (2) and suppose that the model satisfies Assumptions 1–9 and 12 and 24 (the last of which replaces Assumption 13). Then

- 1. We can identify the function $h_{\tau}(x_2, z_2)$ up to a constant k_{τ} for all types τ and $(x_2, z_2) \in \mathcal{X}_2 \times \mathcal{Z}_2$.
- 2. Period 1 utility $u_1(x_1, a; \tau)$ is identified $\forall (a \in \mathcal{A}_1, x_1 \in \mathcal{X}_1, \tau \in \mathcal{T}).$

Proof. The proof is essentially identical to the proof of Lemma 3 with the only modification being how the function $h_{\tau}(x_2, z_2)$ is identified. To this end, Collect the S-1 linearly independent rows of $\mathsf{dF}_{\Delta}(z_1') - \mathsf{dF}_{\Delta}(z_1'')$ into a matrix and denote it by H. Recall that each row of H corresponds to a unique value of (a_1, x_1) and denote the collection of these values by $(x_{1,j}, a_{1,j})_{j=1}^{S-1}$ and recall that for each value $(x_{1,j}, a_{1,j})$

$$\frac{g_{\tau,1,a_{1,j},\Delta}(x_{1,j})}{\beta_{\tau}\delta_{\tau}} = \sum_{(x_{2},z_{2})\in\mathcal{X}_{2}\times\mathcal{Z}_{2}} \int v_{\tau,2}^{*}(s_{2}) \bigg(dF(\epsilon_{2}) dF_{\Delta,a_{1,j}}(x_{2},z_{2}|x_{1,j},z_{1}') - dF(\epsilon_{2}) dF_{\Delta,a_{1,j}}(x_{2},z_{2}|x_{1,j},z_{1}'') \bigg)
= \sum_{(x_{2},z_{2})\in\mathcal{X}_{2}\times\mathcal{Z}_{2}} h_{\tau}(x_{2},z_{2}) \bigg(dF_{\Delta,a_{1,j}}(x_{2},z_{2}|x_{1,j},z_{1}') - dF_{\Delta,a_{1,j}}(x_{2},z_{2}|x_{1,j},z_{1}'') \bigg).$$

Denote the vector

$$\mathbf{g}_{\tau,\Delta} \equiv \frac{1}{\beta_{\tau}\delta_{\tau}} \begin{bmatrix} g_{\tau,1,a_{1,1},\Delta}(x_{1,1}) \\ \vdots \\ g_{\tau,1,a_{1,S-1},\Delta}(x_{1,S-1}), \end{bmatrix}$$

so that we obtain the equation

 $g_{\tau,\Delta} = Hh_{\tau}.$

where h_{τ} is defined in (39). Consequently, under Assumption 24 we recover h_{τ} . The remainder of the proof is identical to the one above and is omitted.

Proof of Lemma 4

Proof. The argument here is straight-forward. First consider any (\mathbf{b}, \mathbf{d}) belonging to the identified set. Then, by definition we must have $\beta_{\tau_N} \delta_{\tau_N} = \mathbf{b} \mathbf{d}$ and $h_{\tau_N}^{\Delta}(x_2, z_2) = \tilde{h}_{\tau_N}^{\Delta}(x_2, z_2, \mathbf{d})$. Moreover, this choice of time-preference parameters is consistent with the remaining identified preference parameters (the per-period utility functions) so we can use this set of parameters to generate the same joint observed distributions as the original set of parameters. Moreover, this shows that the bounds are sharp.

Proof of Lemma 5

Proof. The result follows directly from computing the derivative of the function

$$\tilde{h}_{\tau}(x_2, z_2, \mathsf{d}, \mathsf{d}) = \sum_{a \in \mathcal{A}_2} \int \left(u_2(x_2, a, \tau) + \epsilon_2(a) + \mathsf{d}q_{\tau}(x_2, z_2, a) \right) A_{\tau}(s_2, a, \mathsf{d}, 1) \, \mathrm{dF}(\epsilon_2).$$

In fact, this function is convex in d (proof available on request). To keep the exposition straight-forward we demonstrate the result for the case with 3 possible actions in period 2. In the following, we shall use repeatedly the fact that ϵ_2 has a strictly positive density everywhere on its domain.

$$\frac{\partial \dot{h}(x_2, z_2, \mathsf{d}, \mathsf{d})}{\partial \mathsf{d}} = \sum_{a \in \mathcal{A}_2} \frac{\partial \ddot{h}^a_{\tau}(x_2, z_2, \mathsf{d})}{\partial \mathsf{d}},$$

where

$$\begin{aligned} \frac{\partial \tilde{h}_{\tau}^{a}(x_{2}, z_{2}, \mathsf{d})}{\partial \mathsf{d}} &= \frac{\partial}{\partial \mathsf{d}} \int \left(u_{2}(x_{2}, a, \tau) + \epsilon_{2}(a) + \mathsf{d}q_{\tau}(x_{2}, z_{2}, a) \right) A_{\tau}(s_{2}, a, \mathsf{d}, 1) \, \mathrm{dF}(\epsilon_{2}) \\ &= \frac{\partial}{\partial \mathsf{d}} \int \left(m(x_{2}, z_{2}, a, \tau, \mathsf{d}) + \epsilon_{2}(a) \right) A_{\tau}(s_{2}, a, \mathsf{d}, 1) \, \mathrm{dF}(\epsilon_{2}), \end{aligned}$$

where

$$m_a(x_2, z_2, \tau, \mathsf{d}) = u_2(x_2, a, \tau) + \mathsf{d}q_\tau(x_2, z_2, a).$$

In what follows we will refer to this simply as m_a and its derivative with respect to d as m'_a for brevity. In addition, we will use ϵ_{2a} to refer to $\epsilon_2(a)$ Applying Leibniz's rule and the Dominated Convergence Theorem repeatedly we can show

$$\begin{split} \frac{\partial \tilde{h}_{\tau}^{1}(x_{2}, z_{2}, \mathsf{d})}{\partial \mathsf{d}} &= (m_{1}' - m_{0}') \int_{\epsilon_{20}} \int_{\epsilon_{22} < m_{0} - m_{2} + \epsilon_{20}} (m_{0} + \epsilon_{20}) f(\epsilon_{20}, m_{0} - m_{1} + \epsilon_{20}, \epsilon_{22}) \mathrm{d}\epsilon_{22} \, \mathrm{d}\epsilon_{20} \\ &+ m_{1}' \int_{\epsilon_{20}} \int_{\epsilon_{22} < m_{0} - m_{2} + \epsilon_{20}} \int_{\epsilon_{21} > m_{0} - m_{1} + \epsilon_{20}} f(\epsilon_{20}, \epsilon_{21}, \epsilon_{22}) \, \mathrm{d}\epsilon_{21} \, \mathrm{d}\epsilon_{22} \, \mathrm{d}\epsilon_{20} \\ &+ (m_{1}' - m_{2}') \int_{\epsilon_{20}} \int_{\epsilon_{22} > m_{0} - m_{2} + \epsilon_{0}} (m_{2} + \epsilon_{2}) f(\epsilon_{0}, m_{2} - m_{1} + \epsilon_{2}, \epsilon_{2}) \, \mathrm{d}\epsilon_{22} \, \mathrm{d}\epsilon_{20} \\ &+ m_{1}' \int_{\epsilon_{20}} \int_{\epsilon_{22} > m_{0} - m_{2} + \epsilon_{0}} \int_{\epsilon_{21} > m_{2} - m_{1} + \epsilon_{22}} f(\epsilon_{20}, \epsilon_{21}, \epsilon_{22}) \, \mathrm{d}\epsilon_{21} \, \mathrm{d}\epsilon_{22} \, \mathrm{d}\epsilon_{20} \end{split}$$

$$\begin{aligned} \frac{\partial h_{\tau}^2(x_2, z_2, \mathbf{d})}{\partial \mathbf{d}} &= (m_2' - m_0') \int_{\epsilon_{20}} \int_{\epsilon_{22} < m_0 - m_2 + \epsilon_{20}} (m_0 + \epsilon_{20}) f(\epsilon_{20}, \epsilon_{21}, m_0 - m_2 + \epsilon_{20}) \mathrm{d}\epsilon_{21} \, \mathrm{d}\epsilon_{20} \\ &+ m_2' \int_{\epsilon_{20}} \int_{\epsilon_{21} < m_0 - m_1 + \epsilon_{20}} \int_{\epsilon_{22} > m_0 - m_2 + \epsilon_{20}} f(\epsilon_{20}, \epsilon_{21}, \epsilon_{22}) \, \mathrm{d}\epsilon_{22} \, \mathrm{d}\epsilon_{21} \, \mathrm{d}\epsilon_{20} \\ &+ (m_2' - m_1') \int_{\epsilon_{20}} \int_{\epsilon_{21} > m_0 - m_1 + \epsilon_{20}} (m_1 + \epsilon_1) f(\epsilon_0, \epsilon_1, m_1 - m_2 + \epsilon_1) \, \mathrm{d}\epsilon_{21} \, \mathrm{d}\epsilon_{20} \\ &+ m_2' \int_{\epsilon_{20}} \int_{\epsilon_{21} > m_0 - m_1 + \epsilon_{20}} \int_{\epsilon_{22} > m_0 - m_1 + \epsilon_{20}} f(\epsilon_{20}, \epsilon_{21}, \epsilon_{22}) \, \mathrm{d}\epsilon_{22} \, \mathrm{d}\epsilon_{21} \, \mathrm{d}\epsilon_{20} \end{aligned}$$

$$\begin{split} \frac{\partial \tilde{h}_{\tau}^{0}(x_{2}, z_{2}, \mathsf{d})}{\partial \mathsf{d}} &= (m_{0}' - m_{1}') \int_{\epsilon_{22}} \int_{\epsilon_{21} > m_{2} - m_{1} + \epsilon_{22}} (m_{1} + \epsilon_{21}) f(m_{1} - m_{0} + \epsilon_{21}, \epsilon_{21}, \epsilon_{22}) \mathrm{d}\epsilon_{21} \, \mathrm{d}\epsilon_{22} \\ &+ m_{0}' \int_{\epsilon_{22}} \int_{\epsilon_{21} > m_{2} - m_{1} + \epsilon_{22}} \int_{\epsilon_{20} > m_{1} - m_{0} + \epsilon_{21}} f(\epsilon_{20}, \epsilon_{21}, \epsilon_{22}) \, \mathrm{d}\epsilon_{20} \, \mathrm{d}\epsilon_{21} \, \mathrm{d}\epsilon_{22} \\ &+ (m_{0}' - m_{2}') \int_{\epsilon_{22}} \int_{\epsilon_{21} < m_{2} - m_{1} + \epsilon_{22}} (m_{2} + \epsilon_{22}) f(m_{2} - m_{0} + \epsilon_{22}, \epsilon_{21}, \epsilon_{22}) \, \mathrm{d}\epsilon_{21} \, \mathrm{d}\epsilon_{22} \\ &+ m_{0}' \int_{\epsilon_{22}} \int_{\epsilon_{21} < m_{2} - m_{1} + \epsilon_{22}} \int_{\epsilon_{20} > m_{2} - m_{0} + \epsilon_{22}} f(\epsilon_{20}, \epsilon_{21}, \epsilon_{22}) \, \mathrm{d}\epsilon_{20} \, \mathrm{d}\epsilon_{21} \, \mathrm{d}\epsilon_{22}. \end{split}$$

Adding the three terms and simplifying suitably, we obtain

$$\frac{\partial \tilde{h}_{\tau}(x_2, z_2, \mathsf{d}, \mathsf{d})}{\partial \mathsf{d}} = \sum_{a \in \mathcal{A}_2} q_{\tau}(x_2, z_2, a) \int_{\epsilon_2} A_{\tau}(s_2, a, \mathsf{d}, 1) \, \mathrm{dF}(\epsilon_2), \tag{41}$$

so that $\frac{\partial \tilde{h}_{\tau}^{\Delta}(x_2, z_2, \mathsf{d}, \mathsf{d})}{\partial \mathsf{d}} = \sum_{a \in \mathcal{A}_2} \left(q_{\tau_N}(x_2, z_2, a) \int A_{\tau}(s_2, a, \mathsf{d}, 1) \, \mathrm{dF}(\epsilon_2) - q_{\tau_N}(x_{20}, z_{20}, a) \int A_{\tau}(s_2, a, \mathsf{d}, 1) \, \mathrm{dF}(\epsilon_2) \right), \text{ which is exactly the expression in Assumption 14.} \Box$

A.2 Proofs for Identification for Unobserved Types

Proof of Proposition 1

Proof. The proof follows closely the arguments in Proposition 3 of Kasahara and Shimotsu (2009). Given a value (r, v) define the $M_{r,v} \times M_{r,v}$ diagonal matrix $\bigvee_{r,v}^{M_{r,v}}$ that contains the strictly positive type probabilities $\pi_{\tau}(r, v)$ as its diagonal elements. Next, note that

$$\mathsf{P}^{a_2,a_3,\underline{M}}_{r,v} = (\mathsf{L}^{a_2,\mathbf{x}_2,\underline{M}+1}_v)'\mathsf{V}^{M_{r,v}}_{r,v}\mathsf{L}^{a_3,\underline{M}+1}_{\mathbf{x}_3,v},$$

from which it follows that $\operatorname{Rank}(\mathsf{P}_{r,v}^{a_2,a_3,\underline{M}}) \leq \min\left\{\operatorname{Rank}(\mathsf{L}_v^{a_2,\mathbf{x}_2,\underline{M}+1}), \operatorname{Rank}(\mathsf{L}_{\mathbf{x}_3,v}^{a_3,\underline{M}+1}), \operatorname{Rank}(\mathsf{V}_{r,v}^{M_{r,v}})\right\}$. Since the rank of $\mathsf{V}_{r,v}^{M_{r,v}} = M_{r,v}$ we have that

$$M_{r,v} \ge \operatorname{Rank}(\mathsf{P}_{r,v}^{a_2,a_3,\underline{M}}). \tag{42}$$

Next, suppose that Assumption 17 holds. We will show that in that case the reverse inequality holds. First note that then $L_{\mathbf{x}_3,v}^{a_3,\underline{M}+1}(L_{\mathbf{x}_3,v}^{a_3,\underline{M}+1})'$ is invertible. Post-multiplying both sides by $\left(L_{\mathbf{x}_3,v}^{a_3,\underline{M}+1}\right)'\left(L_{\mathbf{x}_3,v}^{a_3,\underline{M}+1}\left(L_{\mathbf{x}_3,v}^{a_3,\underline{M}+1}\right)'\right)^{-1}$,

$$\mathsf{P}_{r,v}^{a_{2},a_{3},\underline{M}}\left(\mathsf{L}_{\mathbf{x}_{3},v}^{a_{3},\underline{M}+1}\right)'\left(\mathsf{L}_{\mathbf{x}_{3},v}^{a_{3},\underline{M}+1}\left(\mathsf{L}_{\mathbf{x}_{3},v}^{a_{3},\underline{M}+1}\right)'\right)^{-1} = (\mathsf{L}_{v}^{a_{2},\mathbf{x}_{2},\underline{M}+1})'\mathsf{V}_{r,v}^{M_{r,v}}$$

Since $\mathsf{L}_{v}^{a_{2},\mathbf{x}_{2},\underline{M}+1}$ has rank $M_{r,v}$ and $V_{r,v}^{M_{r,v}}$ has strictly positive diagonal elements, it must be the case that the rank of $\mathsf{P}_{r,v}^{a_{2},a_{3},\underline{M}}\left(\mathsf{L}_{\mathbf{x}_{3},v}^{a_{3},\underline{M}+1}\left(\mathsf{L}_{\mathbf{x}_{3},v}^{a_{3},\underline{M}+1}\left(\mathsf{L}_{\mathbf{x}_{3},v}^{a_{3},\underline{M}+1}\right)'\right)^{-1}$ is $M_{r,v}$ and it follows that

$$M_{r,v} \leq \min \left\{ \operatorname{Rank}(\mathsf{P}_{r,v}^{a_2,a_3,\underline{M}}), \operatorname{Rank}\left(\left(\mathsf{L}_{\mathbf{x}_3,v}^{a_3,\underline{M}+1}\right)' \left(\mathsf{L}_{\mathbf{x}_3,v}^{a_3,\underline{M}+1} \left(\mathsf{L}_{\mathbf{x}_3,v}^{a_3,\underline{M}+1}\right)'\right)^{-1} \right) \right\},$$

so that $\operatorname{Rank}(\mathsf{P}_{r,v}^{a_2,a_3,\underline{M}}) \geq M_{r,v}$. Combining this with (42) we conclude that $\operatorname{Rank}(\mathsf{P}_{r,v}^{a_2,a_3,\underline{M}}) = M_{r,v}$. Therefore, the rank of the directly identified matrix $\mathsf{P}_{2,r,v}^{a_2,a_3,\underline{M}}$ gives the total number of types in the population. \Box

A.2.1 Proofs for Identification of Type Specific Choice Probabilities

Proof of Lemma 6

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 (Assumption 16) to generate identification instead of using observations that are more than one period apart. As on p.18 we begin by using Assumption 16 to simplify the joint distribution of states and actions conditional on (r, v).

$$\begin{aligned} & \mathbb{P}(a_t, a_{t+1}, \mathbf{x}_t, \mathbf{x}_{t+1} | r, v) \\ &= \sum_{\tau \in \mathcal{T}} \pi_{\tau}(r, v) \mathbb{P}_{\tau}(a_{t+1}, a_t, \mathbf{x}_{t+1}, \mathbf{x}_t | r, v) \\ &= \sum_{\tau} \pi_{\tau}(r, v) \mathbb{P}_{\tau}(a_{t+1} | a_t, \mathbf{x}_{t+1}, \mathbf{x}_t, r, v) \mathbb{P}_{\tau}(\mathbf{x}_{t+1} | \mathbf{x}_t, a_t, r, v) \mathbb{P}_{\tau}(a_t, \mathbf{x}_t | r, v) \\ &= \sum_{\tau} \pi_{\tau}(r, v) \mathbb{P}_{\tau}(a_{t+1} | \mathbf{x}_{t+1}, v) \mathbb{P}(\mathbf{x}_{t+1} | \mathbf{x}_t, a_t, v) \mathbb{P}_{\tau}(a_t, \mathbf{x}_t | v), \end{aligned}$$

where in the last line we have used the Markov nature of the optimal choices (see the discussion on p.7) and Assumption 16. Next, recall that we defined (in (15))

$$\mathbf{F}_{\mathbf{x}_{t},\mathbf{x}_{t+1},r,v}^{a_{t},a_{t+1}} = \sum_{\tau} \pi_{\tau}(r,v) \mathbb{P}_{\tau}(a_{t+1}|\mathbf{x}_{t+1},v) \mathbb{P}_{\tau}(a_{t},\mathbf{x}_{t}|v), \tag{43}$$

and (in (16))

$$\mathbf{F}_{r,v}^{a_t,\mathbf{x}_t} = \sum_{\tau} \pi_{\tau}(r,v) \mathbb{P}_{\tau}(a_t,\mathbf{x}_t,v)$$
$$\mathbf{F}_{\mathbf{x}_{t+1},r,v}^{a_{t+1}} = \sum_{\tau} \pi_{\tau}(r,v) \mathbb{P}_{\tau}(a_{t+1}|\mathbf{x}_{t+1},v).$$

Next, for given values $(\mathbf{x}_t^1, \dots, \mathbf{x}_t^{M_{r,v}-1}, \mathbf{x}_{t+1}^1, \dots, \mathbf{x}_{t+1}^{M_{r,v}-1})$ define the $M_{r,v} \times M_{r,v}$ directly identified matrix

$$\mathsf{P}_{r,v}^{a_{t},a_{t+1},M_{r,v}} = \begin{pmatrix} 1 & \mathbf{F}_{\mathbf{x}_{t+1}^{a_{t+1}},r,v}^{a_{t+1}} & \cdots & \mathbf{F}_{\mathbf{x}_{t+1}^{d_{t+1}},r,v}^{a_{t+1}} \\ \mathbf{F}_{\mathbf{x}_{t}^{1},r,v}^{a_{t}} & \mathbf{F}_{\mathbf{x}_{t}^{1},\mathbf{x}_{t+1}^{1},r,v}^{a_{t},a_{t+1}} & \cdots & \mathbf{F}_{\mathbf{x}_{t}^{1},\mathbf{x}_{t+1}^{1},r,v}^{a_{t},a_{t+1},r,v} \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{F}_{\mathbf{x}_{t}^{M_{r,v}-1},r,v}^{a_{t}} & \mathbf{F}_{\mathbf{x}_{t}^{M_{r,v}-1},\mathbf{x}_{t+1}^{1},r,v}^{a_{t},a_{t+1}} & \cdots & \mathbf{F}_{\mathbf{x}_{t}^{M_{r,v}-1},\mathbf{x}_{t+1}^{M_{r,v}-1},r,v} \end{pmatrix},$$
(44)

which is just the matrix defined in (18) but with \underline{M} replaced by $M_{r,v} - 1$. We will abbreviate $\mathsf{P}_{r,v}^{a_t,a_{t+1},M_{r,v}}$ to P_r in the sequel to economize on notation. Recall that the type-specific choice probabilities are defined in (19). Next, we define the matrices $\mathsf{L}_v^{a_t,\mathbf{x}_t,M_{r,v}}$ and $\mathsf{L}_{\mathbf{x}_{t+1},v}^{a_{t+1},M_{r,v}}$ using (20) and (21) but replacing \underline{M} with $M_{r,v} - 1$ (so the dimensions now depend only upon the number of types $M_{r,v}$). Thus,

$$\mathsf{L}_{v}^{a_{t},\mathbf{x}_{t},M_{r,v}} \equiv \begin{pmatrix} 1 & \lambda_{v}^{a_{t},\mathbf{x}_{t}^{1},\tau_{1}} & \dots & \lambda_{v}^{a_{t},\mathbf{x}_{t}^{M_{r,v}-1},\tau_{1}} \\ 1 & \lambda_{v}^{a_{t},\mathbf{x}_{t}^{1},\tau_{2}} & \dots & \lambda_{v}^{a_{t},\mathbf{x}_{t}^{M_{r,v}-1},\tau_{2}} \\ \vdots & \vdots & \dots & \vdots \\ 1 & \lambda_{v}^{a_{t},\mathbf{x}_{t}^{1},\tau_{M_{r,v}}} & \dots & \lambda_{v}^{a_{t},\mathbf{x}_{t}^{M_{r,v}-1},\tau_{M_{r,v}}} \end{pmatrix},$$
(45)

which we will abbreviate as $L_{t,r}$. Similarly,

$$\mathsf{L}_{\mathbf{x}_{t+1},v}^{a_{t+1},M_{r,v}} \equiv \begin{pmatrix} 1 & \lambda_{\mathbf{x}_{t+1},\tau_{1}}^{a_{t+1},\tau_{1}} & \dots & \lambda_{\mathbf{x}_{t+1},v-1}^{a_{t+1},\tau_{1}} \\ 1 & \lambda_{\mathbf{x}_{t+1},v}^{a_{t+1},\tau_{2}} & \dots & \lambda_{\mathbf{x}_{t+1},\tau_{2}}^{a_{t+1},\tau_{2}} \\ 1 & \lambda_{\mathbf{x}_{t+1},v}^{a_{t+1},\tau_{2}} & \dots & \lambda_{\mathbf{x}_{t+1},v}^{a_{t+1},\tau_{2}} \\ \vdots & \vdots & \dots & \vdots \\ 1 & \lambda_{\mathbf{x}_{t+1}^{a_{t+1},\tau_{M_{r,v}}}}^{a_{t+1},\tau_{M_{r,v}}} & \dots & \lambda_{\mathbf{x}_{t+1}^{M_{r,v-1}},v}^{a_{t+1},\tau_{M_{r,v}}} \end{pmatrix}.$$
(46)

which we will abbreviate as $L_{t+1,r}$.

Assumption 16 and Assumption 19 guarantee that $L_{t,r} = L_{t,r'}$ and $L_{t+1,r} = L_{t+1,r'}$. Finally, define the $M_{r,v} \times M_{r,v}$ matrix $V_r = \text{Diag}(\pi_{\tau_1}(r, v), \ldots, \pi_{\tau_{M_{r,v}-1}}(r, v))$ which we will abbreviate to V_r . It is easy to show the following factorization holds:

$$\mathsf{P}_r = \mathsf{L}'_{t,r} \mathsf{V}_r \mathsf{L}_{t+1,r},\tag{47}$$

and by assumption each term on the right hand side is invertible. Next, for $r \neq r'$ consider the directly identified object A defined by

$$A \equiv \mathsf{P}_{r}^{-1}\mathsf{P}_{r'} = \mathsf{L}_{t+1,r}^{-1}\mathsf{V}_{r}^{-1}\mathsf{V}_{r',v}\mathsf{L}_{t+1,r},$$
(48)

so that

$$\mathsf{L}_{t+1,r}\mathsf{A} = \widehat{V}_{r,r'}\mathsf{L}_{t+1,r},$$

where $\hat{\mathsf{V}}_{r,r'} \equiv \mathsf{V}_r^{-1}\mathsf{V}_{r'}$ is a diagonal matrix. The expression above asserts that the diagonal matrix $\hat{\mathsf{V}}_{r,r',v}$ contains the eigenvalues of A and that the rows of $\mathsf{L}_{t+1,v}$ 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 $\mathsf{E} \equiv \mathsf{DL}_{t+1,r}$ where D is a diagonal matrix (and we have $\mathsf{L}_{t+1,r} = \mathsf{D}^{-1}\mathsf{E}$). Next,

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

Since the first row of $L'_{t,r}$ consists of ones, the first row of the identified matrix $P_r E^{-1}$ identifies the elements of the diagonal matrix $V_r D^{-1}$. Define $F \equiv V_r D^{-1}$ to be the identified matrix from this analysis. Next,

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

where all the terms on the right hand side are identified, so that $L_{t,r}$ is identified. Next,

$$V_r L_{t+1,r} = (L'_{t,r})^{-1} \mathsf{P}_r = \mathsf{P}_r \mathsf{E}^{-1} \mathsf{F}^{-1} \mathsf{P}_r.$$
(49)

where the right hand side is directly identified. The first column on the left hand side consists of the diagonal elements of the matrix $V_{r,v}$. Therefore $V_{r,v}$ is identified. Denote by $G(\equiv V_r)$ the diagonal matrix obtained by this argument. Then,

$$\mathbf{L}_{t+1,r} = \mathbf{G}^{-1} \mathbf{P}_r \mathbf{E}^{-1} \mathbf{F}^{-1} \mathbf{P}_r, \tag{50}$$

where the matrix G is invertible since by assumption all its diagonal entries are non-zero. Finally, note that since V_r is identified, then $V_{r'} = G\widehat{V}_{r,r'}$ and so $V_{r'}$ is also identified since both G and \widehat{V} are identified.

We first apply this result for $(t, t+1) \in \{(2, 3), (1, 2)\}$ to identify the type-specific choice probabilities $\mathbb{P}_{\tau}(a_{t+1}|\mathbf{x}_{t+1}, v)$ and $\mathbb{P}_{\tau}(a_t, \mathbf{x}_t|v)$ for each period. Note that the model is actually overidentified in a sense since we can recover period 2 choice probabilities from both applications of the argument. In principle, one could use this to propose a specification test (i.e. that the period 2 choice probabilities obtained by two applications of the argument should be the same). Finally, we also note that we have identified the type probabilities $\pi_{\tau}(r, v)$ so we have identified the relative sizes of the different types of agent in the population.

We next state a result that does not require Assumption 18 (the overlap of types across (r, v)). In its stead we require the existence of a set of state-variables across three periods that satisfy a stronger version of Assumption 19:

ASSUMPTION 25. Given (r, v), there exist $(\mathbf{x}_1, \mathbf{x}_2^1, \dots, \mathbf{x}_2^{M_{r,v}-1}, \mathbf{x}_3^1, \dots, \mathbf{x}_3^{M_{r,v}-1})$ such that (a) $\mathbb{P}(\mathbf{x}_3^j | \mathbf{x}_2^k, r, v) \mathbb{P}(\mathbf{x}_2^k | \mathbf{x}_1, r, v) \neq 0$ for $(j, k) \in \{1, \dots, M_{r,v} - 1\}$ and (b) the $M_{r,v} \times M_{r,v}$ matrices L_v and $\mathsf{L}_{t+1,v}$ are invertible. In addition, the matrix $\mathsf{P}^{\mathbf{x}_1}$ defined below is invertible.

We can, however, substantially relax Assumption 16 to:

ASSUMPTION 26. The transition probabilities do not vary by type: $\mathbb{P}_{\tau}(\mathbf{x}_{t+1}|\mathbf{x}_t, a_t, r, v) = \mathbb{P}(\mathbf{x}_{t+1}|\mathbf{x}_t, a_t, r, v)$.

In particular, type-specific choice probabilities need not be independent of the type proxy r so that the recovered preference parameters will also be indexed by r (i.e. $u_{\tau,r}(x_t, v)$ and $(\beta_{\tau,r}, \delta_{\tau,r})$).

LEMMA 9. Fix (r, v) and suppose that Assumptions 25, 18 and 26 hold and that the optimal decision process is Markov. Then, the type-specific choice probabilities $\{\mathbb{P}_{\tau}(a_t|\mathbf{x}_t, r, v)\}_{\tau \in \mathcal{T}_{r,v}; t \in \{1,2,3\}}$ for $\mathbf{x}_t \in \mathcal{X}_t$. are identified

Proof. The proof is very similar to the proof of Lemma 6 with the main addition being that we now examine events in three consecutive time-periods (as opposed to two periods earlier). First, define the identified quantities

$$\begin{aligned} \mathbf{F}_{\mathbf{x}_{1},\mathbf{x}_{2},\mathbf{x}_{3},r,v} &\equiv \sum_{\tau \in \mathcal{T}_{r,v}} \pi_{\tau}(r,v) \mathbb{P}_{\tau}(a_{3}|\mathbf{x}_{3},r,v) \mathbb{P}_{\tau}(a_{2}|\mathbf{x}_{2},r,v) \mathbb{P}_{\tau}(a_{1},\mathbf{x}_{1}|r,v) \\ &= \sum_{\tau \in \mathcal{T}_{r,v}} \pi_{\tau}(r,v) \lambda_{\mathbf{x}_{3},v}^{a_{3},\tau} \lambda_{\mathbf{x}_{2},v}^{a_{2},\tau} \lambda_{v}^{a_{1},\mathbf{x}_{1},\tau} \end{aligned}$$

$$\begin{aligned} \mathbf{F}_{\mathbf{x}_{1},\mathbf{x}_{2},r,v} &= \sum_{\tau \in \mathcal{T}_{r,v}} \pi_{\tau}(r,v) \mathbb{P}_{\tau}(a_{1},\mathbf{x}_{1}|r,v) \mathbb{P}_{\tau}(a_{2}|\mathbf{x}_{2},r,v) = \sum_{\tau \in \mathcal{T}_{r,v}} \pi_{\tau}(r,v) \lambda_{\mathbf{x}_{2},v}^{a_{2},\tau} \lambda_{v}^{a_{1},\mathbf{x}_{1},\tau} \\ \mathbf{F}_{\mathbf{x}_{1},\mathbf{x}_{3},r,v} &= \sum_{\tau \in \mathcal{T}_{r,v}} \pi_{\tau}(r,v) \mathbb{P}_{\tau}(a_{3}|\mathbf{x}_{3},r,v) \mathbb{P}_{\tau}(a_{1}\mathbf{x}_{1}|r,v) = \sum_{\tau \in \mathcal{T}_{r,v}} \pi_{\tau}(r,v) \lambda_{\mathbf{x}_{3},v}^{a_{3},\tau} \lambda_{v}^{a_{1},\mathbf{x}_{1},\tau} \\ \mathbf{F}_{\mathbf{x}_{1},r,v} &= \sum_{\tau \in \mathcal{T}_{r,v}} \pi_{\tau}(r,v) \mathbb{P}_{\tau}(a_{1}\mathbf{x}_{1}|r,v) = \sum_{\tau \in \mathcal{T}_{r,v}} \pi_{\tau}(r,v) \lambda_{v}^{a_{1},\mathbf{x}_{1},\tau} \end{aligned}$$

Next, for given values of $\mathbf{x}_1, \mathbf{x}_2^1, \dots, \mathbf{x}_2^{M_{r,v}-1}, \mathbf{x}_3^1, \dots, \mathbf{x}_3^{M_{r,v}-1}$, define the $M_{r,v} \times M_{r,v}$ square matrix

$$\mathsf{P}_{r,v}^{\mathbf{x}_{1}} = \begin{pmatrix} \mathbf{F}_{\mathbf{x}_{1}} & \mathbf{F}_{\mathbf{x}_{1},\mathbf{x}_{3}^{1}} & \cdots & \mathbf{F}_{\mathbf{x}_{1},\mathbf{x}_{3}^{M_{r,v}-1}} \\ \mathbf{F}_{\mathbf{x}_{1},\mathbf{x}_{2}^{1}} & \mathbf{F}_{\mathbf{x}_{1},\mathbf{x}_{2}^{1},\mathbf{x}_{3}^{1}} & \cdots & \mathbf{F}_{\mathbf{x}_{1},\mathbf{x}_{2}^{1},\mathbf{x}_{3}^{M_{r,v}-1}} \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{F}_{\mathbf{x}_{1},\mathbf{x}_{2}^{M_{r,v}-1}} & \cdots & \mathbf{F}_{\mathbf{x}_{1},\mathbf{x}_{2}^{M_{r,v}-1},\mathbf{x}_{3}^{M_{r,v}-1}} \end{pmatrix},$$
(51)

Define the matrix $\mathsf{D}_{\mathbf{x}_1,r} \equiv \operatorname{Diag}(\lambda_v^{a_1,\mathbf{x}_1,\tau_1},\ldots,\lambda_v^{a_1,\mathbf{x}_1,\tau_{M_{r,v}-1}})$ and as before $\mathsf{V}_r = \operatorname{Diag}(\pi_{\tau_1}(r,v),\ldots,\pi_{\tau_{M_{r,v}-1}}(r,v))$. Then, the following factorizations hold – suppressing the dependence on v and r since identification does not dependence upon variation in these variables:

$$\mathsf{P}^{\mathbf{x}_1} = \mathsf{L}_2' \mathsf{V} \mathsf{D}_{\mathbf{x}_1} \mathsf{L}_3; \qquad \qquad \mathsf{P} = \mathsf{L}_2' \mathsf{V} \mathsf{L}_3,$$

where the matrix $P = P_{r,v}^{a_2,a_3,M_{r,v}}$ is defined in (44) and the matrices (L₂, L₃) are defined in (45). The argument from here onwards follows the same broad outlines as the previous lemma but using the period ahead decompositions (rather than the variation in r). Consider the directly identified object A defined by

$$\mathbf{A} \equiv \mathbf{P}^{-1} \mathbf{P}^{\mathbf{x}_1} = \mathbf{L}_3^{-1} \mathbf{D}_{\mathbf{x}_1} \mathbf{L}_3,\tag{52}$$

so that

 $\mathsf{L}_3\mathsf{A}=\mathsf{D}_{\mathbf{x}_1}\mathsf{L}_3.$

The eigenvalues of A determine $D_{\mathbf{x}_1}$ and the rows of L_3 are the left eigenvectors of A. 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 $E \equiv HL_3$ where H is a diagonal matrix with non-zero diagonal entries (and we have $L_3 = H^{-1}E$). Next,

$$\mathsf{P}\mathsf{E}^{-1} = \mathsf{L}_2'\mathsf{V}\mathsf{H}^{-1}.$$

Since the first row of L'_2 consists of ones, the first row of the identified matrix PE^{-1} identifies the elements of the diagonal matrix VH^{-1} . Define $F \equiv VH^{-1}$ to be this identified matrix. Next,

$$\mathsf{L}_2'=\mathsf{P}\mathsf{L}_3^{-1}\mathsf{V}^{-1}=\mathsf{P}\mathsf{E}^{-1}\mathsf{H}\mathsf{V}^{-1}=\mathsf{P}\mathsf{E}^{-1}\mathsf{F}^{-1},$$

where all the terms on the right hand side are identified, so that L_2 is identified. Next,

$$VL_3 = (L_2')^{-1}P = PE^{-1}F^{-1}P.$$
(53)

The first column on the left hand side consists of the diagonal elements of the matrix V_r . Therefore V is identified since all the matrices on the right hand side in (53) are identified. Denote by $G(\equiv V)$ the diagonal matrix obtained by this argument. Then,

$$L_3 = G^{-1} P E^{-1} F^{-1} P, \tag{54}$$

where the matrix G is invertible since by assumption all its diagonal entries are non-zero. Finally, note that since V is identified, then $\mathsf{D}_{x_1}=(\mathsf{L}_2'\mathsf{V})^{-1}\mathsf{P}^{\mathbf{x}_1}\mathsf{L}_3^{-1}$ is also identified. \Box

A.2.2 Proofs for Identification of Type Identities

Proof of Lemma 7

Proof. " \Rightarrow ": Suppose that $\hat{\delta}_{\tau}$ is a constant but $\tau = \tau_N$. First, noting that

$$\begin{split} h^{\Delta}_{\tau_N}(x_2, z_2) &= \tilde{h}^{\Delta}_{\tau}(x_2, z_2, \delta_{\tau_N}, \delta_{\tau_N}) \\ \tilde{h}^{\Delta}_{\tau_N}(x_2, z_2, \beta_{\tau_N}\delta_{\tau_N}, \delta_{\tau_N}) &= h^{\Delta}_{\tau_N, 1}(x_2, z_2, \beta_{\tau_N}\delta_{\tau_N}, \delta_{\tau_N}) + \delta_{\tau_N} h^{\Delta}_{\tau_N, 2}(x_2, z_2, \beta_{\tau_N}\delta_{\tau_N}) \end{split}$$

we can rewrite

$$\begin{split} \hat{\delta}_{\tau_N}(x_2, z_2) &= \frac{h_{\tau_N}^{\Delta}(x_2, z_2) - h_{\tau_N, 1}^{\Delta}(x_2, z_2, \beta_{\tau_N} \delta_{\tau_N})}{\tilde{h}_{\tau_N, 2}^{\Delta}(x_2, z_2, \beta_{\tau_N} \delta_{\tau_N})} \\ &= \frac{\tilde{h}_{\tau}^{\Delta}(x_2, z_2, \delta_{\tau_N}, \delta_{\tau_N}) - \tilde{h}_{\tau_N}^{\Delta}(x_2, z_2, \beta_{\tau_N} \delta_{\tau_N}, \delta_{\tau_N})}{\tilde{h}_{\tau_N, 2}^{\Delta}(x_2, z_2, \beta_{\tau_N} \delta_{\tau_N})} - \delta_{\tau_N} \end{split}$$

and the first term is non-constant by Assumption 21 so we have contradiction. The " \Leftarrow " follows by simply observing that for $\tau \in \{\tau_S, \tau_N\}$, $\hat{\delta}_{\tau}(x_2, z_2) = \delta_{\tau}$.

Proof of Proposition 2

Proof. First, note that the results of Lemma 1 and Lemma 3 do not require the type identity to be known (i.e. they apply to all types). Therefore, starting with a given set of type-specific choice probabilities $\{\mathbb{P}_{\tau}(a_t|x_t, z_t)\}_t$ we can identify the per period utilities for all three periods (without knowing the type τ) as well as the product $\beta_{\tau}\delta_{\tau}$ and the function $h_{\tau}(x_2, z_2)$ defined in (7). Using these, we can construct the known function $\hat{\delta}_{\tau}(x_2, z_2)$ defined in (24): First, suppose that $\hat{\delta}_{\tau}(x_2, z_2)$ is a constant (in (x_2, z_2)). Then, by Lemma 7 the type must be either sophisticated or consistent and and that $\hat{\delta}(x_2, z_2) = \delta_{\tau}$. Next, if the ratio of the directly identified objects $\frac{\beta_{\tau}\delta_t}{\delta_{\tau}} = 1$ then we can conclude the type must be consistent while if the ratio is strictly less than one, the type must be sophisticated inconsistent. If, however, $\hat{\delta}_{\tau}(x_2, z_2)$ is not a constant, then by Lemma 7 the type must be naïve inconsistent. \Box

A.3 Partially Sophisticated Agents

Proof of Proposition 3

Proof. We first prove part 1 starting with the " \Leftarrow " implication by noting that if $\tilde{\beta}_{\tau} = \beta_{\tau}$ then $h_{\tau}^{\Delta}(x_2, z_2) = \tilde{h}_{\tau}^{\Delta}(x_2, z_2, \beta_{\tau}\delta_{\tau}, \delta_{\tau})$ so that $\hat{\delta}_{\tau}(x_2, z_2) = \delta_{\tau}$ which is constant. We prove the " \Rightarrow " implication using a proof by contradiction. Suppose $\tilde{\beta}_{\tau} \neq \beta_{\tau}$ but $\hat{\delta}_{\tau}(x_2, z_2)$ is a constant. First, as in the proof of Lemma 7 rewrite $\hat{\delta}_{\tau}(x_2, z_2)$ as

$$\delta_{\tau}(x_2, z_2) = \frac{\tilde{h}_{\tau}^{\Delta}(x_2, z_2, \tilde{\beta}_{\tau} \, \delta_{\tau}, \delta_{\tau}) - \tilde{h}_{\tau}^{\Delta}(x_2, z_2, \beta_{\tau} \delta_{\tau}, \delta_{\tau})}{h_{\tau,2}^{\Delta}(x_2, z_2, \beta_{\tau} \delta_{\tau})} - \delta_{\tau}.$$

By Assumption 23 if $\tilde{\beta}_{\tau} \neq \beta_{\tau}$ then the first term is not-constant (in (x_2, z_2)) and we have a contradiction.

The argument for the second part of the lemma is essentially identical to the arguments for the proof of Proposition 2. As before, using the results of Lemma 1 and Lemma 3 we identify the per period utilities for all three periods for any given type τ , the product $\beta_{\tau}\delta_{\tau}$ and the function $h_{\tau}(x_2, z_2)$ defined in (7). Using these, we can construct the known function $\hat{\delta}_{\tau}(x_2, z_2)$ defined in (24).

First, suppose that $\hat{\delta}_{\tau}(x_2, z_2)$ is a constant (in (x_2, z_2)). Then, by part 1 of this lemma $\hat{\beta}_{\tau} = \beta_{\tau}$ so that agents are either completely sophisticated or time-consistent and in either case $\hat{\delta}_{\tau} = \delta_{\tau}$. Next, if the ratio of the two identified quantities $\beta_{\tau} \delta_t / \hat{\delta}_{\tau} = 1$ then we can conclude the type must be consistent while if the ratio is strictly less than one, the type must be completely sophisticated.

Suppose instead that $\hat{\delta}_{\tau}(x_2, z_2)$ is not a constant. Then again by the first part of this lemma we must have $\tilde{\beta}_{\tau} \neq \beta_{\tau}$ so that the type under consideration must be partially sophisticated. Note, however at this point we cannot further categorize partially sophisticated agents on the basic of the values of $\tilde{\beta}_{\tau}$ (e.g. into completely naïve agents) since the latter is not identified.

Proof of Proposition 4

Proof. Under the assumptions, the function $h_{\tau}^{\Delta}(x_2, z_2)$ is identified (part 1 of Lemma 3) and is equal to $\tilde{h}_{\tau}^{\Delta}(x_2, z_2, \tilde{\beta}_{\tau} \delta_{\tau}, \delta_{\tau})$. The definition of the identified set follows. Note that the identified set is sharp in the sense that all values in Θ are candidate values for the true parameter $(\tilde{\beta}_{\tau}, \delta_{\tau})$.

Online Appendix for Identification of Time-Inconsistent Models: The Case of Insecticide Treated Nets

B Identification Proofs for the Empirical Application

In this appendix we discuss model identification when applying our general results from Section 3 to our empirical application of ITN choice. We begin by defining the state space and the action space and then discuss the fidelities to and departures from the original set of assumptions of the general model.

State Space: S_t

In the empirical work and formal identification results we allow for a rich observable state space (including income and other characteristics), but for expositional purposes we simplify the state space to the bare minimum. Online Appendix D contains details on the construction of the state space for the empirical implementation.

To fix ideas, 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 the purchase decision of the agent. The agent can either purchase no net (n) or purchase one with one of two contracts. Contract b involves the purchase of an ITN that is repaid over the next 12 months. With contract c ('commitment'), the agent purchases both an ITN and a set of two retreatments with insecticide. Buyers who choose b can still purchase retreatments for cash, but contract c may be appealing for sophisticated agents who wish to commit to the ITN maintenance schedule at the time of purchase. However, there may be non-time preference reasons for agents choice of c and that is an important reason for us to consider the unobserved types model. The second letter captures whether anyone in the agent's household suffered from malaria in the past six months, with m denoting someone had malaria and h ('healthy') denoting that no-one contracted malaria. 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. As before, we can allow for a richer (albeit still discrete) state space and in the application we also include income as a state variable.

Action Space: A_t

The action space in period 1 (\mathcal{A}_1) has three elements denoted by (n, b, c), which are defined as above. In periods 2 and 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 do so. 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. Beliefs

Beliefs about the likelihood of contracting malaria (when using ITNs, untreated nets and no nets) were elicited from households during the baseline as described in the Data section. We use these elicited beliefs – now denoted without a time-subscript since they are collected at baseline – z as the excluded variable in Assumption 3. We are thus abstracting away from learning issues in general, and more specifically ruling out learning about the efficacy of ITNs as a motive for ITN purchase. Extending the model to allow for such motives is an important topic for future research but we ignore it here.

With the definition of the state and action space in hand we can now discuss the substantive content of the assumptions made in the previous section. First, Assumption 1 holds. This rules out (in the simple state space example above) for instance the possibility that the probability of malaria infection in period 3 depends on malaria status in the first period given malaria status and the retreatment decisions from the second period. Incorporating such dependencies, however, is straightforward by suitably redefining the state space at t to contain the entire malaria history up to t.

Next, Assumption 2 holds as well but with the modification that z_t is no longer a state variable but rather a conditioning variable so that the modified assumption is

ASSUMPTION 27. (Modification of Assumption 2)

$$\mathbb{P}(x_t, \epsilon_t | x_{t-1}, \epsilon_{t-1}, a_{t-1}, z) = \mathbb{P}(x_t | x_{t-1}, a_{t-1}, z) \mathbb{P}(\epsilon_t),$$
(55)

where the distribution of the vector ϵ_t is absolutely continuous on \mathbb{R}^{K_t} w.r.t. Lebesgue measure and independently distributed across t.

Assumption 3 holds (with $z_t = z$) and the substantive import of the assumption is that elicited beliefs about the efficacy of ITNs (or income) do not enter the static payoff function. Similarly, additive separability given by Assumption 4 also holds. Before discussing the modifications to Assumption 5 we describe the type indicators in greater detail.

Type Indicators

We use two important pieces of information from the intervention to learn about type: 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 $\tilde{r} = 1$ and agents who exhibit no preference reversals have $\tilde{r} = 0$. Second, as described previously, we designed a contract that should appeal to sophisticated inconsistent agents, thereby providing a signal about their type.

DIRECTLY IDENTIFIED TYPES

In the directly observed types case, we use both these pieces of information to characterize three types of agent. Agents with $\tilde{r} = 0$ are classified as time-consistent and agents with $\tilde{r} = 1$ are classified as time-inconsistent. Further, agents with $\tilde{r} = 1$ who purchase the commitment product $(a_1 = c)$ are classified as time-inconsistent sophisticated types and agents with $\tilde{r} = 1$ and who purchase the standard product $(a_1 = b)$ 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 below and this is one important departure from the general results for directly observed types in the previous section. UNOBSERVED TYPES

In this case, the researcher does not directly observe the type of any individual. We assume instead that the variables (\tilde{r}, a_1) are only imperfect proxies, as is likely the case. For instance, an agent may choose $\tilde{r} = 1$ 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.

As before, we first discuss identification for the directly observed types case so Assumption 5 holds but only for t > 1 and for $r = (\tilde{r}, a_1)$. For clarity, we restate that assumption with the requisite modifications here

ASSUMPTION 28. (Modification of Assumption 5) Choice probabilities for types that purchase a product are directly observed. In particular, for a time-consistent agent

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

For a naïve time-inconsistent agent

$$\mathbb{P}_{\tau_N}(a_t^* = a | x_t, z) = \mathbb{P}(a_t^* = a | x_t, z, \tilde{r} = 1, a_1 = b) \text{ for } t > 1$$

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

$$\mathbb{P}_{\tau_S}(a_t^* = a | x_t, z) = \mathbb{P}(a_t^* = a | x_t, z, \tilde{r} = 1, a_1 = c) \text{ for } t > 1.$$

As before, we start the backward induction from the last decision made by the agent, which is the decision to retreat the net in period 3. Since the decision to retreat is based on an expectation about the future, we introduce a terminal period (period 4) where no action is taken but over which expectations are formed in period 3 and which affect the period 3 decision to retreat. To be specific, utility for type τ in period 4 is $u_4(x_4; \tau)$ and the agent's choice is period 3 is

$$a_3 = 1 \iff u_3(x_3, 1; \tau) - u_3(x_3, 0; \tau) + \beta_\tau \delta_\tau \int u_4(x_4; \tau) \, \mathrm{dF}_\Delta(x_4 | x_3, z) + \epsilon_3(1) - \epsilon_3(0) > 0.$$

As before, using the Hotz-Miller inversion argument we can identify the function

$$g_{\tau,3,1}(x_3,z) = u_3(x_3,1;\tau) - u_3(x_3,0;\tau) + \beta_\tau \delta_\tau \int u(x_4;\tau) \,\mathrm{dF}_\Delta(x_4|x_3,z),$$

and as before, using variation in z we can isolate the last term in the expression above. Next, we seek to identify its

constituent components using the variation in x_3 . Intuitively, as long as x_3 has sufficiently rich support (as large as that of x_4) we can generate a system of linear equations using

$$g_{\tau,3,\Delta}(x_3) \equiv g_{\tau,3,1}(x_3,z) - g_{\tau,3,1}(x_3,z') = \int_{x_4 \in \mathcal{X}_4} \beta_\tau \delta_\tau u_4(x_4;\tau) \left(\mathrm{dF}_\Delta(x_4|x_3,z) - \mathrm{dF}_\Delta(x_4|x_3,z') \right).$$
(56)

Formally, define the matrix

$$\mathsf{dF}(k,z) \equiv \begin{bmatrix} \mathrm{dF}_k(x_{4,1}|x_{3,1},z) & \dots & \mathrm{dF}_k(x_{4,S_4-1}|x_{3,1},z) \\ \vdots & \vdots & \vdots \\ \mathrm{dF}_k(x_{4,1}|x_{3,S_4-1},z) & \dots & \mathrm{dF}_k(x_{4,S_4-1}|x_{3,S_4-1},z) \end{bmatrix}.$$

and define the matrix $dF_{\Delta,1}(z) \equiv dF(1,z) - dF(0,z)$. We can then state the remaining assumptions required for the first result.

ASSUMPTION 29. (a) \mathcal{X}_3 has at least $S_4 - 1$ points of support (where S_4 is the cardinality of \mathcal{X}_4). (b) The distribution of z conditional on $x_3 \in \mathcal{X}_3$ has at least two points of support. (c) The $(S_4 - 1) \times (S_4 - 1)$ identified matrix $\mathsf{dF}_{\Delta,1}(z) - \mathsf{dF}_{\Delta,1}(z')$ is invertible for $z \neq z'$.

The value of z can differ across columns but we omit that generality in the notation to avoid cumbersome notation. Finally, we need the appropriate normalization assumptions.

ASSUMPTION 30. Utility for a base action (action 0) is known for all types and for all states, i.e. $u_3(x_3, 0; \tau)$ is known for all $x_3 \in \mathcal{X}_3$ for all types. In addition, $u_4(x_{4,0}; \tau)$ is known for all types τ for a given point $x_{4,0} \in \mathcal{X}_4$.

Finally, we can state the identification result for preferences in period 3 and 4:

LEMMA 10. Consider an agent solving 2 and suppose that Assumption 1, 27 (which replaces Assumption 2), 3, 4, 28 (which replaces Assumption 5), 29 and 30 hold. Then,

- 1. $u_3(x_3, a; \tau)$ is identified for all types τ and $x_3 \in \mathcal{X}_3$.
- 2. $\beta_{\tau}\delta_{\tau}(u_4(x_4;\tau)-u_4(x_{40};\tau))$ is identified for all types τ and $x_4 \in \mathcal{X}_4$ and a fixed given $x_{4,0} \in \mathcal{X}_4$.

Proof of Lemma 10

Proof. The argument of the proof is similar to the argument in the proof of Lemma 3. We begin by defining the $S_4 - 1$ column vectors $g_{\tau,3,\Delta}(k)$ and $h_{\tau,4}$

$$\mathbf{g}_{\tau,3,\Delta} \equiv \begin{bmatrix} g_{\tau,\Delta}(x_{3,1}) \\ \vdots \\ g_{\tau,\Delta}(x_{3,S_4-1}) \end{bmatrix} \mathbf{h}_{\tau,4} \equiv \beta_{\tau} \delta_{\tau} \begin{bmatrix} u_4(x_{4,1};\tau) - u_4(x_{4,0};\tau) \\ \vdots \\ u_4(x_{4,S_4-1};\tau) - u_4(x_{4,0};\tau) \end{bmatrix}.$$
(57)

We have subtracted $u(x_{40}; \tau)$ where x_{40} is a fixed point in \mathcal{X}_4 to incorporate the constraint that

$$\sum_{x_4 \in \mathcal{X}_4} \mathrm{dF}_k(x_4 | x_3, z) = 1$$

Without incorporating this restriction the corresponding matrix needed in Assumption 29 would not be invertible. Using the notation above we can rewrite (56) in matrix form as

$$\mathsf{g}_{\tau,3,\Delta} = (\mathsf{dF}_{\Delta,1}(z) - \mathsf{dF}_{\Delta,1}(z'))\mathsf{h}_{\tau,4},$$

where $g_{\tau,3,\Delta}$ is identified from the argument culminating in (56) and the matrices $dF_{\Delta,1}(z)$ and $dF_{\Delta,1}(z')$ are identified since they are constructed from observed beliefs. Under the testable invertibility assumption (Assumption 29) it follows that we can identify

$$h_{\tau,4} = \left(\mathsf{dF}_{\Delta,1}(z) - \mathsf{dF}_{\Delta,1}(z')\right)^{-1} \mathsf{g}_{\tau,3,\Delta}.$$
(58)

Therefore, the object $\beta_{\tau}\delta_{\tau}(u(x_4;\tau) - u(x_{40};\tau))$ is identified for all $x_4 \in \mathcal{X}_4$. Next, note that

$$g_{\tau,3,1}(x_3,z) = u_3(x_3,1;\tau) - u_3(x_3,0;\tau) + \beta_\tau \delta_\tau \int \left(u(x_4;\tau) - u(x_{40};\tau) \right) \mathrm{dF}_\Delta(x_4|x_3,z), \tag{59}$$

since

$$\int u(x_{40};\tau)dF_{\Delta}(x_4|x_3,z)=0.$$

Since the second term in (59) is now identified and $u_3(x_3, 0; \tau)$ is known, we can identify $u_3(x_3, 0; \tau)$ for all $x_3 \in \mathcal{X}_3$.

Note that we are only able to identify the product of the time-preference parameters and the period four payoff function. We next examine the choices made in period 2 to isolate the time-preference parameters from the other preference parameters. As before, we apply the Hotz-Miller inversion to directly identify

$$g_{\tau,2,1}(x_2,z) \equiv u_2(x_2,1;\tau) - u_2(x_2,0;\tau) + \beta_\tau \delta_\tau \int h_\tau(x_3,z) \,\mathrm{dF}_\Delta(x_3|x_2,z),\tag{60}$$

which is conceptually analogous to the expression (6). Analogous the notation in that section, we define

$$h_{\tau}(x_3, z) \equiv \int v_{\tau}^*(x_3, \epsilon_3, z) \,\mathrm{dF}(\epsilon_3) \tag{61}$$

$$v_{\tau}^{*}(s_{3}, z) \equiv \sum_{a \in \mathcal{A}_{3}} \left(v_{\tau,3}(s_{3}, z, a, 1, \delta_{\tau}) \right) A_{\tau}(s_{3}, z, a, \tilde{\beta}_{\tau}, \delta_{\tau})$$
(62)

$$v_{\tau,3}(s_3, z, a, \mathbf{b}, \mathbf{d}) \equiv u_3(x_3, a; \tau) + \epsilon_3(a) + \mathbf{bd} \int u(x_4; \tau) \, \mathrm{dF}(x_4 | x_3, z, a)$$
(63)

$$A_{\tau}(s_3, z, a, \tilde{\beta}_{\tau}, \delta_{\tau}) \equiv \mathbb{I}\bigg\{a = \operatorname*{argmax}_{j \in \mathcal{A}_3} v_{\tau,3}(s_3, z, j, \tilde{\beta}_{\tau}, \delta_{\tau})\bigg\}.$$
(64)

As before, we use variation in beliefs z to isolate the quantity

$$g_{\tau,2,\Delta}(x_2) \equiv g_{\tau,2,1}(x_2,z) - g_{\tau,2,1}(x_2,z'), \tag{65}$$

where we have written $g_{\tau,2,\Delta}(x_2,z,z')$ as $g_{\tau,2,\Delta}(x_2)$ for brevity. Next, we rewrite this as

$$g_{\tau,2,\Delta}(x_2) = \beta_\tau \delta_\tau G_{\tau,\Delta}(x_2, \tilde{\beta}_\tau \delta_\tau) + \delta_\tau H_{\tau,\Delta}(x_2, \tilde{\beta}_\tau \delta_\tau), \tag{66}$$

where

$$G_{\tau,\Delta}(x_2, \mathsf{b}) \equiv G_{\tau}(x_2, z, \mathsf{b}) - G_{\tau}(x_2, z', \mathsf{b})$$
$$H_{\tau,\Delta}(x_2, \mathsf{b}) \equiv H_{\tau}(x_2, z, \mathsf{b}) - H_{\tau}(x_2, z', \mathsf{b}),$$

and as before we have omitted (z, z') as arguments for brevity, and where

$$G_{\tau}(x_{2}, z, \mathbf{b}) \equiv \sum_{a \in \mathcal{A}_{3}} \int \left(u_{3}(x_{3}, a; \tau) + \epsilon_{3}(a) \right) A_{\tau}(s_{3}, a, z, \mathbf{b}, 1) \, \mathrm{dF}_{\Delta}(s_{3} | x_{2}, a, z)$$

$$H_{\tau}(x_{2}, z, \mathbf{b}) \equiv \sum_{a \in \mathcal{A}_{3}} \int \int \beta_{\tau} \delta_{\tau} \left(u_{4}(x_{4}; \tau) - u(x_{40}; \tau) \right) \, \mathrm{dF}(x_{4} | x_{3}, a, z) A_{\tau}(s_{3}, a, z, \mathbf{b}, 1) \, \mathrm{dF}_{\Delta}(s_{3} | x_{2}, a, z)$$
(67)

Note that in general, the functions G and H are completely known functions of **b**. Next, note that for sophisticated and consistent agents $G_{\tau}(x_2, z, \tilde{\beta}_{\tau} \delta_{\tau}) = G_{\tau}(x_2, z, \beta_{\tau} \delta_{\tau})$ and $H_{\tau}(x_2, z, \tilde{\beta}_{\tau} \delta_{\tau}) = H_{\tau}(x_2, z, \beta_{\tau} \delta_{\tau})$ and these are both identified since $u_3(\cdot)$ and the distribution of ϵ_3 are known, $\tilde{\beta}_{\tau} \delta_{\tau} = \beta_{\tau} \delta_{\tau}$ and this term only enters as $\beta_{\tau} \delta_{\tau} (u_4(x_4; \tau) - u_4(x_{40}; \tau))$ which in turn is identified by part 2 of Lemma 10. This argument, however, breaks down for naïve agents since for them $\tilde{\beta}_{\tau} \delta_{\tau} = \delta_{\tau}$ and the latter is not identified either in isolation or multiplicatively. Note that for naïve agents part 2 of Lemma 11 identifies $\beta_{\tau_N} \delta_{\tau_N} (u_4(x_4; \tau_N) - u_4(x_{40}; \tau_N))$ which is not enough to identify $H(x_2, z, \tilde{\beta}_{\tau_N} \delta_{\tau_N})$. Therefore for naïve agents $G_{\tau}(x_2, z, \beta_{\tau} \delta_{\tau})$ and $H_{\tau}(x_2, z, \tilde{\beta}_{\tau} \delta_{\tau})$ are not identified. For consistent agents (for whom $\beta_{\tau} = 1$) the exponential parameter $\delta_{\tau C}$ can be obtained directly by solving (66). For sophisticated agents, instead, the expression in (66) for a given value of x_2 is a linear equation in two unknowns ($\beta_{\tau}\delta_{\tau}, \delta_{\tau}$) and an appropriate invertibility condition is required for its solution. First, define The 2 × 2 matrix

$$\mathsf{K}_{\tau}(\beta_{\tau}\delta_{\tau}, x_{2}, x_{2}', z, z') \equiv \begin{bmatrix} \mathsf{G}_{\tau,\Delta}(x_{2}, \beta_{\tau}\delta_{\tau}) & \mathsf{H}_{\tau,\Delta}(x_{2}, \beta_{\tau}\delta_{\tau}) \\ \mathsf{G}_{\tau,\Delta}(x_{2}', \beta_{\tau}\delta_{\tau}) & \mathsf{H}_{\tau,\Delta}(x_{2}', \beta_{\tau}\delta_{\tau}) \end{bmatrix}.$$
(69)

ASSUMPTION 31. There exist $x_2 \neq x'_2, z \neq z'$ such that the matrix $\mathsf{K}_{\tau_S}(\beta_{\tau_S}\delta_{\tau_S}, x_2, x'_2, z, z')$ is invertible

Finally, we need the standard normalization assumption for period 2 preferences.

ASSUMPTION 32. Utility for a base action (action 0) is known for all types and for all states, i.e. $u_2(x_2, 0; \tau)$ is known for all $x_2 \in \mathcal{X}_2$ for all types.

We can then state the second identification result. Under the maintained assumptions we can only identify the individual time-preference parameters for the sophisticated and the consistent types as in the general section, and for the same reasons, the time-preference parameters for the naïve types are not point identified.

LEMMA 11. Consider an agent solving 2 and suppose that Assumption 1, 27 (which replaces Assumption 2),3, 4, 28 (which replaces Assumption 5), 29, 30, 31 and 32 hold. Then,

- 1. $u_2(x_2, a; \tau)$ is identified for time-consistent and sophisticated types for $x_3 \in \mathcal{X}_2$.
- 2. The time-preference parameters for sophisticated agents, $(\beta_{\tau_S}, \delta_{\tau_S})$ are identified.
- 3. The exponential parameter for time-consistent agents δ_{τ_C} is identified.

Proof of Lemma 11

Proof. For sophisticated types $\tilde{\beta}_{\tau} = \beta_{\tau}$ so that upon examining the function

$$g_{\tau,\Delta,2}(x_2) = \beta_\tau \delta_\tau G_{\tau,\Delta}(x_2, \beta_\tau \delta_\tau) + \delta_\tau H_{\tau,\Delta}(x_2, \beta_\tau \delta_\tau),$$

we have a linear equation in two unknowns $(\beta_{\tau_S}\delta_{\tau_S}, \delta_{\tau_S})$ for each value x_2 . Under Assumption 31 we can solve the system of linear equations obtained by evaluating the expression above at two different values of x_2 . Therefore, both β_{τ_S} and δ_{τ_S} are separately identified. Note that for time-consistent agents, the equation above directly identifies δ_{τ_C} since $\beta_{\tau_C} = 1$. On the other hand, for naïve agents, the term $G_{\tau_N,\Delta}(x_2, \tilde{\beta}_{\tau_N} \delta_{\tau_N}) = G_{\tau_N,\Delta}(x_2, \delta_{\tau_N})$ is not identified so we cannot use the same argument as above.

Once we have identified the time-preference parameters the last term on the right hand side of the expression (60) is identified and consequently second period preferences are also identified under the normalization assumption. To verify the first statement, note that the last term for consistent or sophisticated agents is given by

$$\beta_{\tau}\delta_{\tau}\int h_{\tau}(x_{3},z)\,\mathrm{dF}_{\Delta}(x_{3}|x_{2},z) = \beta_{\tau}\delta_{\tau}G_{\tau}(x_{2},z,\tilde{\beta}_{\tau}\delta_{\tau}) + \delta_{\tau}H(x_{2},z,\tilde{\beta}_{\tau}\delta_{\tau})$$
$$= \beta_{\tau}\delta_{\tau}G_{\tau}(x_{2},z,\beta_{\tau}\delta_{\tau}) + \delta_{\tau}H(x_{2},z,\beta_{\tau}\delta_{\tau}),$$

and all terms in the second line are identified.

The proof relies on variation in both beliefs as well as in the state variables to achieve identification and can loosely be thought of as a "double-differencing" argument with the first difference exploiting variation in beliefs to eliminate one part of the model and the second difference exploiting variation in the state variable x_2 to separately identify the exponential and the hyperbolic parameters.

This argument however does not extend to naïve agents because for such agents, the expression (66) for a given value of x_2 is a non-linear equation of the unknowns $(\beta_{\tau_N} \delta_{\tau_N}, \delta_{\tau_N})$ and in general the equation will not uniquely determine both time preference parameters. In particular,

$$g_{\tau_N,2,\Delta}(x_2) = \beta_{\tau_N} \delta_{\tau_N} G_{\tau_N,\Delta}(x_2, z, z', \delta_{\tau_N}) + \delta_{\tau_N} H_{\tau_N}(x_2, z, z', \delta_{\tau_N}),$$

where $G_{\tau_N}(\cdot)$ and $H_{\tau_N}(\cdot)$ are identified up to the time preference parameter δ_{τ_N} . We define the identified set as

$$\Theta_N \equiv \bigg\{ (\mathsf{db}) \in]0,1[^2: g_{\tau_N,\Delta}(x_2, z, z') = \mathsf{bd}G_{\tau_N,\Delta}(x_2, z, z', \mathsf{d}) + \mathsf{d}H_{\tau_N}(x_2, z, z', \mathsf{d}) \ \forall (x_2, z, z') \bigg\}.$$

			Т
			I
			I
_	_	_	л

In general, this set will not be a singleton since it is possible to show that the right hand side is not 1-1 in (c, d). In the interests of space, we do not present the set identification results here but they are available upon request from the authors.

The definition of the identified set also makes clear that if δ_{τ_N} were identified, then so would β_{τ_N} . This would happen for instance if we assumed that sophisticated and naïve agents had the same exponential parameter.

ASSUMPTION 33. Time-inconsistent agents have identical exponential parameters $\delta_{\tau_S} = \delta_{\tau_N}$.

There are potentially alternative assumptions for achieving identification. We note, however, that assuming the hyperbolic parameters are identical does not solve the identification problem completely. The reason is that even if β_{τ_N} is known, (the right hand side of) $g_{\tau_N,2,\Delta}$ is not 1-1 in d and consequently δ_{τ_N} need not be identified.

LEMMA 12. Consider an agent solving (2) and suppose that Assumptions 1, 27 (which replaces Assumption 2),3, 4, 28 (Observability of type-specific choice probabilities and which replaces Assumption 5),29 (Invertibility), 30 (Normalization), 31 (needed to identify δ_{τ_S}), 32 (Normalization) and 33 (exponential discount rates across inconsistent types are equal) hold. Then,

- 1. The hyperbolic parameter β_{τ_N} is identified.
- 2. Period 2 preferences $u_2(x_2, a; \tau_N)$ are identified.

Proof of Lemma 12

Proof. By assumption, $\delta_{\tau_N} = \delta_{\tau_S}$ and δ_{τ_S} is identified from the previous Lemma. Note that we need Assumption 31 in addition to the assumptions stated in the Lemma to ensure that δ_{τ_S} is identified. The function

$$h_{\tau_N}(x_3, z) = \sum_{a \in \mathcal{A}_3} \int \left(u_3(x_3, a; \tau_N) + \epsilon_3(a) + \delta_{\tau_N} \int u_4(x_4; \tau_N) \,\mathrm{dF}(x_4 | x_3, a, z) \right) A_{\tau_N}(s_3, z, a, \delta_{\tau_N}, 1) \,\mathrm{dF}(\epsilon_3)$$

is identified since all the terms on the right hand side are identified. This in turn implies that

$$\beta_{\tau_N} = g_{\tau_N,\Delta}(x_2) \bigg(\delta_{\tau_N} \int \bigg(h_{\tau_N}(x_3, z) \,\mathrm{dF}_{\Delta}(x_3 | x_2, z) - h_{\tau_N}(x_3, z') \,\mathrm{dF}_{\Delta}(x_3 | x_2, z') \bigg)^{-1},$$

as long as $g_{\tau_N,\Delta}(x_2) \neq 0$ (which is observable). Once both time-preference parameters are identified we can recover period 2 preferences as per usual from examing

$$g_{\tau_N,2}(x_2,z) \equiv u_2(x_2,1;\tau_N) - u_2(x_2,0;\tau_N) + \beta_{\tau_N}\delta_{\tau_N} \int v_{\tau_N}^*(s_3,z) \,\mathrm{dF}_{\Delta}(s_3|x_2,z),$$

and noting that the last term is identified and applying the normalization (Assumption 32).

B.1 Period One Preferences

There is a sharp distinction in period 1 relative to the later periods regarding direct type identification for individual agents. In particular, in period 1 types are not observed for all agents so we are in effect in a model with unobserved types. Specifically, 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, because 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 typespecific choice probability $\mathbb{P}_{\tau}(a_1|x_1, z)$. For t > 1 we identified $\mathbb{P}_{\tau}(a_t|x_t, z)$ for agents who purchased a product since the agent's choice of product revealed his type perfectly (as ensured by Assumption 28).

As in the section on unobserved types we adapt the insights from Kasahara and Shimotsu (2009) by imposing a set of exclusion restrictions. As in Section 3.2.2 we will use the structure imposed by the markov assumptions and the exclusion restrictions on the identified matrix $P_{r,v}$ for t = 1 to identify the type-specific choice probabilities. To ease notation we omit dependence of the objects below on household time-invariant characteristics v. This is entirely without loss of generality since these variables are not used for identification. We do, however, discuss these at length in the empirical discussion that follows.

First, we adapt the results for the case where $\tilde{r} = 1$, that is to the sub-populations that expressed timeinconsistent preferences as measured in the baseline survey. As noted before, we assume that both naïve and sophisticated types exist so the total number of types is $M_{\tilde{r}=1} = 2$. In order to proceed, we first restate Assumption 16 in terms of the variables in the empirical application and only for periods 1 and 2 (the plausibility of these assumptions is examined in greater detail in the next section):

ASSUMPTION 34. 1. Conditional upon type, the survey response \tilde{r} is uninformative about choice $\mathbb{P}_{\tau}(a_t|x_t, z, \tilde{r}) = \mathbb{P}_{\tau}(a_t|x_t, z)$ for t = 2 and $\mathbb{P}_{\tau}(a_1, x_1|\tilde{r}, z) = \mathbb{P}_{\tau}(a_1, x_1|z)$.

2. The transition probabilities do not vary by type and are independent of \tilde{r} : $\mathbb{P}_{\tau}(x_{t+1}|x_t, a_t, z, \tilde{r}) = \mathbb{P}(x_{t+1}|x_t, a_t, z)$ for t = 1.

The assumption above is only used for period 1 and 2 in the directly observed types case but we will need it to hold for t > 2 when types are completely unobserved as was the case in Section 3.2.2. We then use the Markov property and the assumption above to obtain

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

and as in equation (16) we define the quantities $\mathbf{F}_{x_1,x_2,\tilde{r}}^{a_1,a_2}, \mathbf{F}_{x_1,\tilde{r}}^{a_1}$ and $\mathbf{F}_{x_2,\tilde{r}}^{a_2}$ and then use them to define the identified matrix $\mathsf{P}_{1,\tilde{r}}^{a_1,a_2}$ which in this case is a 3×3 matrix since there are three types of agent (note that this means we only need two points in the state space for identification arguments):

$$\mathsf{P}_{1,\vec{r}}^{a_1,a_2} = \begin{pmatrix} 1 & \mathbf{F}_{x_2^1,\vec{r}}^{a_2} & \mathbf{F}_{x_2^2,\vec{r}}^{a_2} \\ \mathbf{F}_{x_1^1,\vec{r}}^{a_1} & \mathbf{F}_{x_1^1,x_2^1,\vec{r}}^{a_1,a_2} & \mathbf{F}_{x_1^1,x_2^2,\vec{r}}^{1,2} \\ \mathbf{F}_{x_1^2,\vec{r}}^{a_1} & \mathbf{F}_{x_1^2,x_2^2,\vec{r}}^{a_1,a_2} & \mathbf{F}_{x_1^2,x_2^2,\vec{r}}^{a_1,a_2} \end{pmatrix}.$$

while the type-specific choice probabilities are

$$\lambda_{x_2}^{a_2,\tau} \equiv \mathbb{P}_{\tau}(a_2|x_2, z, v), \quad \lambda_{x_1}^{a_1,\tau} \equiv \mathbb{P}_{\tau}(a_1, x_1|z, v),$$

(where we have suppressed dependence on (z, v) for convenience) and which are used to define the matrices (for $t \in \{1, 3\}$)

$$\mathsf{L}_{t}^{a_{t}} = \begin{pmatrix} 1 & \lambda_{x_{t}^{1}}^{a_{t},\tau_{C}} & \lambda_{x_{t}^{2}}^{a_{t},\tau_{C}} \\ 1 & \lambda_{x_{t}^{1},\tau_{N}}^{a_{t},\tau_{N}} & \lambda_{x_{t}^{2}}^{a_{t},\tau_{N}} \\ 1 & \lambda_{x_{t}^{1}}^{a_{t},\tau_{N}} & \lambda_{x_{t}^{2}}^{a_{t},\tau_{N}} \end{pmatrix}.$$
(70)

Note that the difference in notation across sections reflects the fact that now (a) we have time-varying proxies and we are only using the first period proxy \tilde{r} ; (b) we are suppressing dependence on v and (c) since z is time-invariant we do not need to incorporate it into the state space and so the state space here is just x_t and not $\mathbf{x}_t \equiv (x_t, z_t)$ as before. $\mathsf{L}_2^{a_t}$ is directly identified since we observe type-specific choice probabilities from period 2 onwards, which simplifies the identification argument in this sub-section considerably. We can state the remaining assumptions required for identification.

ASSUMPTION 35. The matrix L_2^1 defined in (70) is invertible for two pairs

- 1. (x_2, x'_2) such that $\mathbb{P}(x_2|x_1, a_1 = b) > 0, \mathbb{P}(x'_2|x_1, a_1 = b) > 0.$
- 2. (x_2^*, x_2^{**}) such that $\mathbb{P}(x_2^*|x_1, a_1 = c) > 0, \mathbb{P}(x_2^{**}|x_1, a_1 = c) > 0.$

The condition above is the usual assertion (familiar from Assumption 19) that requires that types behave sufficiently differently over the state space (and at a given point in the state space as well). Since the first period action a_1 is a part of the second period state variable and we require a separate identification argument for each first period choice, we need the assumption to hold at two pairs of states. Finally, we need the normalization for period 1 preferences

ASSUMPTION 36. $u_1(x_1, 0; \tau)$ is known for $x_1 \in \mathcal{X}_1$ for all types τ .

We can now state the result for identification of first-period choice probabilities and preferences as well.

LEMMA 13. Consider an agent solving 2 and suppose that Assumption 1, 27 (which replaces Assumption 2),3, 4, 28 (observability of type-specific choice probabilities),29 (invertibility), 30 (normalization), 32 (normalization), 33 (exponential discount rates across inconsistent types are equal), 35 (invertibility of L_2) and 36 (normalization) hold. Then,

- 1. First period preferences $u_1(x_1, a; \tau)$ are identified for $x_1 \in \mathcal{X}_1$, $a \in \mathcal{A}_1$ for all types τ
- 2. The type probabilities (conditional on the survey response) $\pi_{\tau}(\tilde{r})$ are identified.

Proof of Lemma 13

Proof. Given the notation introduced we can write

$$\mathsf{P}_{1,1}^{b,1} = (\mathsf{L}_1^b)' \mathsf{V}_1 \mathsf{L}_2^1,$$

where we have set $(a_1, a_2, \tilde{r}) = (b, 1, 1)$ so we are examining the sub-population that expressed inconsistency in the survey, purchased a net and retreated it in period 2. Note that since the purchase decision in period 1 is part of the state variable in period 2, we only consider those states in period 2 such that can be transitioned into from the net purchase decision. Concretely, it means that we cannot examine the behavior of sophisticated agents since they definitionally do not purchase contract b.

Next, evaluating $V_{\tilde{r}}$ at $\tilde{r} = 1$: $V_1 = \text{Diag}(0, \pi_{\tau_N}(1), \pi_{\tau_S}(1))$. Note that the first entry is zero since by definition (in this section) consistent agents are those who have $\tilde{r} = 0$. Next, note that given our assumptions about how ITN purchase reveals type, it is clear that $\lambda_{b,x_1}^{\tau_S} = 0$ since sophisticated types do not purchase the standard contract bby construction (again, only for this section) so we do not learn their choice probabilities when $\tilde{r} = 1$ and $a_1 = b$ so that $\lambda_{a_2,x_2}^{\tau_S} = 0$. for all x_2 where $b \in x_2$ (recall that the first period action a_1 is an element of x_t for t > 1) With this in mind, note that

$$(\mathsf{L}_2^1)^{-1}\mathsf{P}_{1,1}^{b,1} = (L_1^b)'\mathsf{V}_1,$$

where the left-hand side is identified and the right hand side is equal to

$$= \begin{pmatrix} 1 & \pi_{\tau_N}(1) & \pi_{\tau_S}(1) \\ 1 & \pi_{\tau_N}(1)\lambda_{b,\pi_1}^{\tau_N} & 0 \\ 1 & \pi_{\tau_N}(1)\lambda_{b,\pi_1'}^{\tau_N} & 0 \end{pmatrix},$$

so that we have identified $\tau_{\tau_N}(1)$ and $\pi_{\tau_S}(1)$. We also have identified $\mathbb{P}_{\tau_N}(a_1 = b|x_1, z)$ and also consequently $\mathbb{P}_{\tau_N}(a_1 = n|x_1, z)$ since naïve agents can only choose among (b, n).

Redoing the arguments with $(a_1 = c)$ will yield identification of $\mathbb{P}_{\tau_S}(a_1 = c | x_1, z)$ as well as $\mathbb{P}_{\tau_S}(a_1 = n | x_1, z)$.

Once first-period choice probabilities are identified and since all time-preference parameters are identified we can recover first-period preferences straightforwardly. In particular, we use the Hotz-Miller inversion to identify

$$g_{\tau,1,k}(x_1,z) = u_1(x_1,k;\tau) - u_1(x_1,n;\tau) + \beta_\tau \delta_\tau \int v_\tau^*(s_2) \,\mathrm{dF}_\Delta(s_2|x_1,z),$$

and then since the entire last term in the expression above is identified, we can identify first period payoff functions. \Box

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 computing counterfactuals. 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 on choice. 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, heterogeneity in take-up, *ceteris paribus*, provides us with a measure of how attractive the commitment contract is for the different types of agents. We explore each of these issues in the estimation section below.

B.2 Unobserved Types

Although survey responses and contract choice are informative about agents' time preferences it is not clear that they are definitively so in the presence of other factors that may affect these responses that are unrelated to time preference. For instance, an agent may choose $\tilde{r} = 1$ due to an imperfect understanding of the choices offered rather than 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 decision not to commit could also depend on low perceived benefits of retreatment. However, this is not a concern for identification to the extent that such perceptions are reflected in agents' perceived beliefs.

The arguments for identification here are identical to those in Section 3.2.2 once we have accounted for the relevant differences between the empirical application and the general model. For this section, there are three important differences worth highlighting: (a) In the application the excluded variable z_t is time-invariant and consequently is not part of the state-space. Consequently we condition all probabilities on z and choices are denoted by $\mathbb{P}_{\tau}(a_t|x_t, z, v)$ instead of $\mathbb{P}(a_t|\mathbf{x}_t, v)$, where $\mathbf{x}_t = (x_t, z_t)$. Similarly, transition probabilities are written as $\mathbb{P}(x_{t+1}|x_t, z, v)$ instead of $\mathbb{P}(\mathbf{x}_{t+1}|\mathbf{x}_t, v)$. (b) The type proxy is type-varying in the application. For t > 1 the type proxy $r = (\tilde{r}, a_1)$ while for period 1 $r = \tilde{r}$. This does not create any complications for estimation though it does make the notation more cumbersome. (c) Finally, for tractability we assume that all three types exist in the population so we do not need to first identify the total number of types. With these alterations to the model, the identification arguments are identical to those in Section 3.2.2.

LEMMA 14. Let t > 1 and fix (r, v). Suppose Assumptions 16 holds with the modifications in (a) above, Assumption 18 holds with $M_{r,v} = 3 \forall (r, v)$ and Assumption 19 holds and that the optimal decision process is Markov. Then, the type-specific choice probabilities $\{\mathbb{P}_{\tau}(a_t|\mathbf{x}_t, v)\}_{\tau \in \mathcal{T}_{r,v}; t \in \{1,2,3\}}$ for $(\mathbf{x}_t, v) \in \mathcal{X}_t \times \mathcal{V}$ are identified for t > 1. For t = 1, fix \tilde{r} and assume that the previous conditions hold. Then the type-specific choice probabilities for period 1 are also identified.

Proof of Lemma 14

Proof. The proof is a direct application of Lemma 6. Note that the alternative approach of using a longer panel is not feasible here.

B.3 Identifying Type Identities

Having identified the type-specific choice probabilities, the next step is to identify the identities of the different types (i.e. classify a given choice probability as belonging to a consistent, naïve or sophisticated type). We adopt two alternative approaches towards identification. The first is relatively straightforward but involves placing stronger assumptions on the type proxy r. The second method avoids the stronger assumptions on the type-proxy and relies more strongly on different types behaving sufficiently differently. We discuss each strategy in turn.

The second strategy assumes that the type proxy is informative about types in a monotone likelihood ratio sense which is achieved by imposing a monotonicity restriction on $\frac{\pi_{\tau}(r)}{\pi_{\tau}(r')}$. To motivate the weakest condition, we start with a stronger set of sufficient conditions. Suppose that 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: $\left\{\frac{\pi_{\tau_C}(r)}{\pi_{\tau_C}(r')} \ge \frac{\pi_{\tau_N}(r)}{\pi_{\tau_N}(r')} \ge \frac{\pi_{\tau_S}(r)}{\pi_{\tau_S}(r')}\right\}$ for r = (0, b) and r' = (1, c). This particular ordering then guarantees the identification of type-identities. While the direct assumptions on the probabilities themselves appear reasonable in our empirical framework, we only need the following weaker condition to hold (which in fact allows us to test the previous set of conditions) for the ratios of the probabilities:

ASSUMPTION 37. For some $r \neq r'$, the three ratios $\left\{\frac{\pi_{\tau_{C}}(r)}{\pi_{\tau_{C}}(r')}, \frac{\pi_{\tau_{N}}(r)}{\pi_{\tau_{N}}(r')}, \frac{\pi_{\tau_{S}}(r)}{\pi_{\tau_{S}}(r')}\right\}$ can be strictly ordered ex-ante.

Under this additional assumption the type-identities are identified.

PROPOSITION 5. Suppose that Lemma 14 holds and that Assumption 37 (Monotonicity) holds. Then, type identities are identified.

Note that in principle one could use the strategy and assumptions outlined in Section 3.2.3 and Proposition 2. However, sample size concerns militated against such an approach - i.e. estimating the identified function 24 non-parametrically is not a feasible approach for the empirical application.

Proof of Proposition 5

Proof. The monotonicity assumption allows the ordering of the type-specific choice probabilities associated with each type-probability so the proof is trivial since the ordering is strict. \Box

This result is useful when we have greater confidence in the ability of the proxy to distinguish between different types of agent. We next discuss a set of assumptions that instead place greater confidence in the sufficiently different behavior of the different types of agent. The argument follows the broad outlines of the discussion in Section 3.2.3 but is more complicated because (a) no action is taken in the terminal state and so $\beta_{\tau}\delta_{\tau}$ is not separately identified within the first two steps of the backward induction (unlike the general case), (b) the constancy of z across time introduces more cumbersome notation and arguments.

The line of argument is very similar to that in the general section. In particular, we start with identified type-specific choice probabilities $P(a_t|x_t, z)$ for period t = 2, 3 but for which the type identity (i.e. whether the agent is consistent or naïve or sophisticated). We then note that Lemma 10 can be applied without knowledge of the type identity to identify period 3 utilities $u(x_3, 0; \tau)$ and the product $\beta_{\tau \delta_{\tau}} (u(x_4; \tau) - u(x_{40}; \tau))$.

Next, using the Hotz-Miller inversion and the differencing argument as earlier on p.4 we can identify the function $g_{\tau,\Delta,2}(x_2)$ (viz. the left hand side of (66)). For brevity, define the 2×1 vector $\mathbf{g}_{\tau,\Delta,2}(x_2, x'_2, z, z') \equiv (g_{\tau,\Delta,2}(x_2), (g_{\tau,\Delta,2}(x'_2))')$ and define the identified quantity

$$\widehat{\mathsf{d}} \equiv (\mathsf{K}(\beta_{\tau}\delta_{\tau}, x_2, x_2', z, z'))^{-1}\mathsf{g}_{\tau, \Delta, 2}(x_2, x_2', z, z'),$$

assuming the inverse exists (see assumption below). We then identify types by examining the solution $\hat{\mathbf{d}}$ for different pairs of points in S. For naïve and sophisticated types, \hat{d} will always be a constant equal to $(\beta_{\tau}\delta_{\tau}, \delta_{\tau})$. This will not be true for the naïve types if the following is true: (a) their view the trade-off between period 2 and period 3 outcomes is different depending upon whether they use δ_{τ_N} or $\beta_{\tau_N}\delta_{\tau_N}$ as the discount rate (which is reasonable) and (b) these differential views of the future vary across the state space. More formally, we need the following conditions to hold:

ASSUMPTION 38. The matrix $K_{\tau}(\beta_{\tau}\delta_{\tau}, x_2, x'_2, z, z')$ defined in (69) is invertible for all types and for values in a subset $S_{\tau} \subset \mathcal{X}_2^2 \times \mathcal{Z}^2$ where $x_2 \neq x'_2$ and $z \neq z'$.

ASSUMPTION 39. Var
$$\left(\operatorname{Vec} \left(\mathsf{K}_{\tau_N}(\beta_\tau \delta_\tau, x_2, x_2', z, z') \left(\mathsf{K}_{\tau_N}(\tilde{\beta}_\tau \delta_\tau, x_2, x_2', z, z'))^{-1} \right) \right) \right) > 0$$
 over the subset \mathcal{S}_{τ_N} .

Note that for sophisticated and consistent agents, since $\tilde{\beta}_{\tau} = \beta_{\tau}$, the product of the two matrices in the assumption above will be the identity matrix. We use this important distinction between types to identify type identities in the following Proposition.

PROPOSITION 6. Suppose that Lemma 14 holds and that Assumptions 38 and 39 hold. Then, type identities are identified.

Proof of Proposition 6

Proof. The proof follows by examining the behavior of \hat{d}_{τ} over the space $\mathcal{X}_2^2 \times \mathcal{Z}^2$. If \hat{d}_{τ} is a constant over the entire space, by Assumption 39 the associated type must be either sophisticated or consistent. In both these cases,

$$\vec{d} = (\beta_\tau \delta_\tau, \delta_\tau)',$$

and so type-identities will be identified for these two types by examining \hat{d}_{τ} . On the other hand, if \hat{d}_{τ} is not a constant over $\mathcal{X}_2^2 \times \mathcal{Z}^2$, the associated type must be naïve.

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. 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. 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 analogously. We maintain Assumptions 1-4 for the argument. The probability that an agent chooses action 0 is

$$\mathbb{P}(a_{2} = 0|x_{2}) = \mathbb{P}_{x_{2}} \left(\begin{array}{c} u(x_{2}, 0) + \epsilon(0) + \beta\delta \int v^{*}(s_{3}) \, \mathrm{dF}(s_{3}|s_{2}, 0) \geq \\ u(x_{2}, 1) + \epsilon(1) + \beta\delta \int v^{*}(s_{3}) \, \mathrm{dF}(s_{3}|s_{2}, 1), \\ u(x_{2}, 0) + \epsilon(0) + \beta\delta \int v^{*}(s_{3}) \, \mathrm{dF}(s_{3}|s_{2}, 0) \geq \\ u(x_{2}, 2) + \epsilon(2) + \beta\delta \int v^{*}(s_{3}) \, \mathrm{dF}(s_{3}|s_{2}, 2) \end{array} \right).$$

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

$$\mathbb{P}(a_{2} = 1|x_{2}) = \mathbb{P}_{x_{2}} \left(\begin{array}{c} u(x_{2}, 0) + \epsilon(0) + \beta\delta \int v^{*}(s_{3}) \, \mathrm{dF}(s_{3}|s_{2}, 0) \leq \\ u(x_{2}, 1) + \epsilon(1) + \beta\delta \int v^{*}(s_{3}) \, \mathrm{dF}(s_{3}|s_{2}, 1), \\ u(x_{2}, 1) + \epsilon(1) + \beta\delta \int v^{*}(s_{3}) \, \mathrm{dF}(s_{3}|s_{2}, 1) \geq \\ u(x_{2}, 2) + \epsilon(2) + \beta\delta \int v^{*}(s_{3}) \, \mathrm{dF}(s_{3}|s_{2}, 2) \end{array} \right)$$

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

$$\mathbb{P}(a_2 = 0 | x_2) = \mathbb{P}(\epsilon(0) - \hat{u}_1 \ge \epsilon(1), \epsilon(0) - \hat{u}_2 \ge \epsilon(2) | x_2)$$

$$\mathbb{P}(a_2 = 1 | x_2) = \mathbb{P}(\epsilon(0) - \hat{u}_1 \le \epsilon(1), \epsilon(1) + (\hat{u}_1 - \hat{u}_2) \ge \epsilon(2) | x_2).$$

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

$$\mathbb{P}(\epsilon(0) - \hat{u}_1 \ge \epsilon(1), \epsilon(0) - \hat{u}_2 \ge \epsilon(2) | x_2) - \mathbb{P}(\epsilon(0) - u_1^* \ge \epsilon(1), \epsilon(0) - u_2^* \ge \epsilon(2) | x_2) = 0$$
(71)

$$\mathbb{P}\left(\epsilon(0) - \hat{u}_1 \le \epsilon(1), \epsilon(1) + (\hat{u}_1 - \hat{u}_2) \ge \epsilon(2) | x_2\right) - \mathbb{P}\left(\epsilon(0) - u_1^* \le \epsilon(1), \epsilon(1) + (u_1^* - u_2^*) \ge \epsilon(2) | x_2\right) = 0.$$
(72)

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 $\hat{u}_2 \ge u_2^*$ then the set

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

and as long as $dF(\epsilon|x_2)$ had strictly positive measure on all of \mathbb{R}^3 , the equality in (71) cannot hold. 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 (72) cannot hold because

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

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

D Variable Description

State Space

<u>Malaria</u>: In period 1, 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 2, 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]) and the parametric distribution assumption imply that the C.D.F. of income is

$$F(y) = \mathbb{I}\left\{y \le \frac{l+u}{2}\right\} \left(\frac{4q}{(u-l)^2}(y-l)^2\right) + \\ \mathbb{I}\left\{y \ge \frac{l+u}{2}\right\} \left(\frac{4(1-q)}{(u-l)^2}(u-y)^2\right).$$

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

$$u \le q \Rightarrow y = l + \left(\frac{u-l}{2}\right)\sqrt{\frac{u}{q}}$$
$$u \ge q \Rightarrow y = u - \left(\frac{u-l}{2}\right)\sqrt{\frac{1-u}{1-q}},$$

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 choose option A in response to the choice problems.

<u>Household Assets</u>: 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.

<u>Beliefs</u> The beliefs data is discussed in detail in Section 4.1.

E Maximum Likelihood Estimation

In this appendix, we show how the model's joint probability distribution can be rewritten in a form that yields an estimable equation for a Conditional Maximum Likelihood Estimation (CMLE).

Given the exclusion restrictions, the likelihood function of the joint model can be written as the sum of the type-specific conditional choice probabilities weighted with the respective type probabilities conditional on the signal r:

$$P(a_1, a_2, a_3, x_1, x_2, x_3, r; z) = \prod_{t=1}^{2} P(x_{t+1}|x_t; z) P(x_1; z, r) P(r) \sum_{\tau} \left(\prod_{t=1}^{3} P(a_t|x_t; z, \tau) P(\tau|r) \right).$$
(73)

Taking logs and removing the parts that do not depend on the estimable parameters, this results in the objective function

$$\sum_{i=1}^{n} \log \left(\sum_{\tau \in \mathcal{T}} P(\tau | r_i) \left(\prod_{t=1}^{3} P(a_{it} | x_{it}; z_i, \tau) \mathbb{I}\{a_{i1} \neq 0\} + P(a_{i1} | x_{i1}, z_i, \tau) \mathbb{I}(a_{i1} = 0) \right) \right).$$
(74)

Denote the vector of model parameters by $\theta \equiv (\delta, \beta_{\tau_N}, \beta_{\tau_S}, \phi, \gamma)$, where δ is the usual exponential discounting parameter, $(\beta_{\tau_N}, \beta_{\tau_S})$ are the hyperbolic parameters for the naïve and sophisticated agents, respectively, ϕ are within-period parameters inside the utility function, and γ are the parameters that explain the population type distribution, as discussed before.

For each draw of θ , the different parts of equation (74) will have to be computed: the population type probabilities that depend on γ , and the type-specific choice probabilities. The latter can be identified using the methods outlined in Lemma 14 (which is an application of Lemma 6). For any candidate of θ , the choice probabilities can be calculated by starting with the value functions for the last period and then working backwards using equations (27), (30) and (33). Using these value functions one can compute the model choice probabilities using the right hand side of equations (28), (29) and (32) for any given set of parameter values. In order to compute these value functions we also need estimates of the transition probabilities $dF(x_{t+1}|x_t, a_t)$ used by agents in solving the problem. We obtain these using elicited beliefs about the two stochastic components of this distribution (income and health) along with information on the monetary costs of illness. The time-invariant variables comprise $v = (v_{hhs}, v_{assets}, v_{risk}, \pi, \delta, \gamma)$, type-specific fixed effects, and elicited subjective beliefs (z) about the reduction of malaria risk from using untreated bednets and ITNs. We choose our estimate of θ_0 to be the value of θ that maximizes the empirical likelihood function implied by equation (74). For considerations of space, we derive only the case of unobserved population types below. The derivation of the known type case is available from the authors upon request. For our CMLE estimations we used a constrained minimization using the Nelder Mead algorithm. For this algorithm, we used a function value convergence criterion of 10^{-7} and used multiple random starting points.

Maximum Likelihood Estimation with Unobserved Types When type is unobserved we need to integrate over the distribution of the unknown types to arrive at the likelihood and the parameters of this distribution become additional parameters to be estimated. We note that for this case we require the "signal" variable r_i to be able to condition out the population type probabilities. In this case, the likelihood is given by

$$P(a_1, a_2, a_3, x_1, x_2, x_3, r; z) = \sum_{\tau \in \mathcal{T}} P(a_1, a_2, a_3, x_1, x_2, x_3, r, \tau; z)$$

=
$$\sum_{\tau \in \mathcal{T}} P(a_1, a_2, a_3, x_1, x_2, x_3 | r, \tau; z) P(\tau | r) P(r),$$

where we have imposed that z does not enter the conditional distribution of types given r or the marginal distribution of r. Next, we use the same decomposition as in the directly observed types case to get

$$= \sum_{\tau} \left(\prod_{t=1}^{3} P(a_t | x_t; z, \tau) P(\tau | r) \right) \prod_{t=1}^{2} P(x_{t+1} | x_t; z) P(x_1; z, r) P(r)$$

$$= \prod_{t=1}^{2} P(x_{t+1} | x_t; z) P(x_1; z, r) P(r) \sum_{\tau} \left(\prod_{t=1}^{3} P(a_t | x_t; z, \tau) P(\tau | r) \right),$$

so that talking logs and dropping the term before the parenthesis, we obtain

$$\log(P(a_1, a_2, a_3, x_1, x_2, x_3, r; z))t \propto \log\left(\sum_{\tau \in \mathcal{T}} \prod_{t=1}^3 P(a_t | x_t; z, \tau) P(\tau | r)\right)$$

=
$$\log\left(\sum_{\tau \in \mathcal{T}} P(\tau | r) \left(\prod_{t=1}^3 P(a_t | x_t; z, \tau) \mathbb{I}\{a_1 \neq 0\} + P(a_1 | x_1, z, \tau) \mathbb{I}(a_1 = 0)\right)\right).$$

Summing over the entire sample the objective function will be given by

$$\sum_{i=1}^{n} \log \left(\sum_{\tau \in \mathcal{T}} P(\tau | r_i) \left(\prod_{t=1}^{3} (P(a_{it} | x_{it}; z_i, \tau) \mathbb{I}\{a_{i1} \neq 0\} + P(a_{i1} | x_{i1}, z_i, \tau) \mathbb{I}(a_{i1} = 0) \right) \right).$$

Where we will compute the probability $P(a_{it}|x_{it}, z_i, \theta)$ as before and the parameters in $P(\tau|r_i)$ are obtained from estimating additional parameters in the model. The last equation above is the one that is implemented for our structural estimations.

F Computation of Effect of Time-Inconsistency on Health Costs and Loss of Workdays

Denote by $P_{\tau,t}(a_t = a | x_i; z_i, \theta)$ the probability that an agent *i* of type τ with observed states x_i and z_i and a vector of preference parameters θ (including a subparameter γ that explains the population type probabilities) chooses action *a* in period *t*. Denote the probability of consumer *i* choosing to adopt any contract (in period 1) or choosing to retreat an ITN (in periods 2 and 3) given the parameter vector θ by $\Psi_{t,i}(\theta)$

The probability of agent i with signal r_i choosing to adopt any contract in period 1 can be written as

$$\Psi_{1,i}(\theta) = \sum_{\tau \in \{C,N,S\}} \pi_{\tau}(r_i,\gamma) [P_{\tau,1}(a_1 = 1 | x_i; z_i, \theta) + P_{\tau,1}(a_1 = 2 | x_i; z_i, \theta)].$$

Similarly, the probability of agent i choosing to retreat an ITN in period $t \in \{2,3\}$ can be written as

$$\Psi_{t,i}(\theta) = \sum_{\tau \in \{C,N,S\}} \pi_{\tau}(r_i,\gamma) \{ P_{\tau,1}(a_1 = 1 | x_i; z_i, \theta) P_{\tau,t}(a_t = 1 | x_i; z_i, \theta) + P_{\tau,1}(a_1 = 2 | x_i; z_i, \theta) P_{\tau,t}(a_t = 1 | x_i; z_i, \theta) \}.$$

Next denote by $\hat{\theta}$ the vector of estimated parameters. Further denote by θ^C the parameter vector that is identical to $\hat{\theta}$ except for having $\beta_N = \beta_S = 1$ instead of the estimated present bias parameters. Because both vectors include the same discount factor δ , this implies that differences in the probability of agent *i* adopting a contract or retreating in periods 2–3 depending on the two parameter vectors is entirely attributable to how the change in the present bias parameters β_N and β_S affect choices of the naïve and sophisticated subtypes. The difference in the probability of agent *i* adopting any contract or retreating an ITN in period $t \in \{1, 2, 3\}$ can thus be written as

$$\Delta \Psi_{t,i}(\hat{\theta}, \theta^C) \equiv \Psi_{t,i}(\hat{\theta}) - \Psi_{t,i}(\theta^C).$$

For period 1, this can be rewritten as

$$\Delta \Psi_{1,i}(\hat{\theta}) = \sum_{\tau \in \{N,S\}} \pi_{\tau}(r_i, \gamma) [\Delta P_{\tau,1}(a_1 = 1 | x_i; z_i, \hat{\theta}, \theta^C) + \Delta P_{\tau,1}(a_1 = 2 | x_i; z_i, \hat{\theta}, \theta^C)],$$

where

$$\Delta P_{\tau,t}(a_t = a | x_i; z_i, \hat{\theta}, \theta^C) \equiv P_{\tau,t}(a_t = a | x_i; z_i, \hat{\theta}) - P_{\tau,t}(a_t = a | x_i; z_i, \theta^C)$$

We next focus on how the change in the probability of adopting a contract or retreating an ITN depending on experiencing present-bias or not leads to a change in the probability of having malaria in the subsequent period.
To do so, define the change in the probability of getting malaria when sleeping under no bednet relative to sleeping under an ITN as $D_{ITN,0}$, the difference in the same probability when sleeping under no bednet relative to sleeping under an untreated net as $D_{untr,0}$, and the difference in the same probability when sleeping under an untreated net relative to an ITN as $D_{ITN,untr}$. Denote by $\mathbb{1}_{i,nonet}$ a dummy variable that is equal to 1 if agent *i* owns an untreated net already before period 1 and that is equal to 0 otherwise.

As values for $D_{ITN,0}$, $D_{ITN,untr}$, and $D_{untr,0}$, respectively, we use two different measures: First, we use the individual elicited beliefs. Second, we use the values indicated in Lengeler (2009), i.e. $D_{ITN,0} = 0.5$, $D_{ITN,untr} = 0.39$, and $D_{untr,0} = 0.11$.

To estimate the difference in the probability of agent *i* having malaria depending on having present-bias relative to not having present-bias, we compute how the adoption probability changes depending on having present-bias as estimated under the parameter vector $\hat{\theta}$ as opposed to the hypothetical parameter vector θ^{C} that includes $\beta_{N} = \beta_{S} = 1$. For period 2, this amounts to the difference in the adoption probabilities times the difference in the in the probability of getting sick between having an ITN and the alternative option. This alternative option is equal to not having a net if the household does not own any net prior to period 1, and equal to having an untreated net if an old net is available in the household.

For period 2, the change in the probability of getting malaria can thus be written as

$$\Delta \Psi_{1,i}(\theta, \theta^C) [D_{ITN,0}(1 - \mathbb{1}_{i,untr}) + D_{ITN,untr} \mathbb{1}_{i,untr}].$$

For periods $t + 1 \in \{3, 4\}$, the change in the probability of getting malaria consists of several parts: First, the difference in the probability that agent *i* does not choose an ITN under present-bias in period 1, but does so under present-bias, has to be multiplied with how this changes the probability of having malaria in period *t* depending on whether the agent retreats the net in period *t*. Second for the probabilities with which the agent already purchases a standard or commitment contract in period 1, respectively, one has to compute how the difference in retreating the nets depending on whether the agent has present-bias or not affects the probability of having malaria. Thus, the probability of getting malaria can be written as

$$\sum_{\tau \in \{N,S\}} \pi_{\tau}(r_{i},\gamma) \{ \sum_{j \in \{1,2\}} \Delta P_{\tau,1}(a_{1} = j | x_{i}; z_{i}, \hat{\theta}, \theta^{C}) [P_{\tau,t}(a_{t} = 0 | x_{i}; z_{i}, \theta^{C}) \mathbb{1}_{i,untr} D_{untr,0} + P_{\tau,t}(a_{t} = j | x_{i}; z_{i}, \theta^{C}) [(1 - \mathbb{1}_{i,untr}) D_{ITN,0} + \mathbb{1}_{i,untr} D_{ITN,untr}]]] + \sum_{\tau \in \{N,S\}} \pi_{\tau}(P_{\tau,1}(a_{1} = 1 | x_{i}; z_{i}, \hat{\theta}) [\Delta P_{\tau,t}(a_{t} = 1 | x_{i}; z_{i}, \hat{\theta}, \theta^{C}) \mathbb{1}_{i,untr} D_{ITN,untr}] + P_{\tau,1}(a_{1} = 2 | x_{i}; z_{i}, \hat{\theta}) [\Delta P_{\tau,t}(a_{t} = 1 | x_{i}; z_{i}, \hat{\theta}, \theta^{C}) \mathbb{1}_{i,untr} D_{ITN,untr}]) \}.$$

In a last step, we (a) multiply the computed changes in the probability of having malaria in each period t with the expected cost of having a malaria episode and with the expected numbers of days lost for each observation in the data and (b) we sum these values across the three time periods and (c) we take the median value of this sum and of each individual period to compute the figures in Table 7.

G Monte Carlo Simulations

To focus attention on the accurate estimation of the time preference parameters, we provide a parsimonious model parametrization for per-period utilities, imposing that they are common across types. We begin by specifying utility in each period as a function of the state variables and actions taken in the last period.

• **Period 4**: $x_4 \in \{h, m\}$

$$u(x_4) = y - \mathbb{I}\{x_4 = m\}\eta_m + \theta_2$$

where h refers to being healthy, m refers to having malaria, y is an agent's income, θ is a utility parameter, and η_m accounts for the costs of malaria.

• **Periods** $t = 2, 3: x_t \in \{b, c, n\} \times \{h, m\} \equiv \{bh, bm, ch, cm, nh, nm\}$

$$u(x_t, a_t) = y - \mathbb{I}\{x_t \in \{bm, cm, nm\}\}\eta_m - p_r \mathbb{I}\{x_t \in \{bh, bm\}\}\mathbb{I}\{a_t = 1\} - p_r \mathbb{I}\{x_t \in \{ch, cm\}\} + \theta.$$

where p_r is the price of retreatment, $a_t = 1$ if the net is re-treated in period t, and $a_t = 0$ otherwise.

• **Period 1**: $x_1 \in \{h, m\}$ and $a_1 \in \{b, c, n\}$

$$u(x_1, a_1) = y - \mathbb{I}\{x_1 = m\}\eta_m - p_b\mathbb{I}\{a_1 = b\} - p_c\mathbb{I}\{a_1 = c\} + \theta$$

where p_b is the price of the standard 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

$$\mathbb{P}_{\tau}(a_t = a | x_t, z) = \frac{\exp(v_{\tau}(x_t, a, z))}{\sum_{s \in \mathcal{A}_t} \exp(v_{\tau}(x_t, s, z))},$$

where the $v_{\tau}(\cdot)$ functions are constructed using backward induction.

We estimate θ along with the time preference parameters $(\delta, \beta_N, \beta_S)$ and (in case of unobserved types) the type probabilities. We use 200 simulations for each model. For the belief variables z, we use a distribution that is close to the empirical distribution in the data. We use the following distributions to draw the subjective probabilities of individual *i* contracting malaria when not using any net $(p_{i,nonet})$, an untreated net $(p_{i,netunt})$, and an ITN $(p_{i,itn})$: $p_{i,nonet} = 0.8 + 0.2u_{i,nonet}$; $p_{i,netunt} = 0.2 \cdot 1\{b_{i,netunt} = 0\} + 0.7 \cdot 1\{b_{i,netunt} = 1\} + 0.05u_{i,netunt}$; $p_{i,itn} =$ $0.05 \cdot 1\{b_{i,itn} = 0\} + 0.3 \cdot 1\{b_{i,itn} = 0\} + 0.05u_{i,itn}$; where $u_{i,itn}, u_{i,netunt}$, and $u_{i,itn}$ are all uniformly distributed over (0,1), $b_{i,netunt}$ follows a Bernoulli distribution with mean 0.6, $b_{i,netunt}$ follows a Bernoulli distribution with mean 0.5. For the signal r we use a binomial distribution with mean r = 0.4. For y we use a value of 9600, which is close to the median income in our data, while η_m is set to 660. Note that because y is constant, the transition probabilities are fully pinned down by the beliefs z.

Tables G.1 and G.2 show the results for the observed types case for one and two separate present bias parameters, respectively. Tables G.3 and G.4 show associated results when the types are unobserved.

	Mean	Median	Std.Dev	True
N=300				
δ	0.7117	0.7033	0.0697	0.7
β	0.3969	0.3903	0.0672	0.4
θ	1.0683	1.2885	0.4856	1.0
N=600				
δ	0.7073	0.6998	0.0529	0.7
β	0.3966	0.3970	0.0525	0.4
θ	1.0856	1.2961	0.4732	1.0
N=300				
δ	0.8938	0.9019	0.0796	0.9
β	0.3098	0.3007	0.0573	0.3
θ	1.0177	1.2978	0.5343	1.0
N=600				
δ	0.9046	0.9070	0.0623	0.9
β	0.3032	0.3045	0.0414	0.3
θ	1.0481	1.2333	0.5057	1.0

Table G.1: Monte Carlo Results: Directly Observed Types, $\beta_S = \beta_N = \beta$

Notes: Each model was simulated 200 times.

	Mean	Median	Std.Dev	True
N=300				
δ	0.9087	0.9072	0.0538	0.9
β_N	0.2958	0.2952	0.0376	0.3
β_S	0.5967	0.5930	0.0712	0.6
θ	1.0967	1.3574	0.4801	1.0
N=600				
δ	0.9080	0.9080	0.0621	0.9
β_N	0.2990	0.2967	0.0455	0.3
β_S	0.6002	0.6003	0.0816	0.6
θ	1.0595	1.2847	0.5043	1.0
N=300				
δ	0.7085	0.7051	0.0304	0.7
β_N	0.3923	0.3902	0.0369	0.4
β_S	0.0966	0.0966	0.0315	0.1
θ	1.0998	1.2641	0.4535	1.0
N=600				
δ	0.7065	0.6986	0.0483	0.7
β_N	0.3977	0.3963	0.0524	0.4
β_S	0.0944	0.0970	0.0442	0.1
θ	1.0667	1.2847	0.5011	1.0

Table G.2: Monte Carlo Results: Directly Observed Types, $\beta_N \neq \beta_S$

Notes: Each model was simulated 200 times.

	Mean	Median	Std.Dev	True
N=300				
δ	0.6952	0.6581	0.1603	0.7
β	0.4293	0.4244	0.1582	0.4
θ	1.0067	1.1291	0.4899	1.0
π_C	0.2958	0.2684	0.1435	0.2214
π_N	0.3858	0.3889	0.0649	0.4446
π_S	0.3214	0.2990	0.1613	0.3340
N=600				
δ	0.7085	0.6900	0.1226	0.7
eta	0.4194	0.3967	0.1484	0.4
θ	0.9537	1.0847	0.5599	1.0
π_C	0.2305	0.2355	0.0588	0.2190
π_N	0.4218	0.4321	0.0484	0.4444
π_S	0.3477	0.3436	0.0823	0.3366
N=1200				
δ	0.7163	0.7091	0.1098	0.7
β	0.4027	0.3927	0.1076	0.4
θ	1.0111	1.1356	0.4997	1.0
π_C	0.2321	0.2234	0.0655	0.2212
π_N	0.4260	0.4351	0.0440	0.4447
π_S	0.3419	0.3321	0.0859	0.3341
N=300				
δ	0.9086	0.9250	0.0787	0.9
β	0.3805	0.3827	0.1094	0.3
θ	1.0948	1.0562	0.2403	1.0
π_C	0.2254		0.1468	0.1947
π_N	0.4144		0.1016	0.4875
π_S	0.3602		0.2242	0.3178
N=600				
δ	0.8677	0.8943	0.1228	0.9
β	0.3362	0.3025	0.1037	0.3
θ	1.0771	1.2614	0.4816	1.0
π_C	0.2339	0.2230	0.0507	0.2213
π_N	0.4184	0.4343	0.0507	0.4447
π_S	0.3477	0.3347	0.1012	0.3340
N=1200				
δ	0.8950	0.9096	0.0896	0.9
β	0.3111	0.2944	0.0696	0.3
θ	1.075	1.2913	0.4886	1.0
π_C	0.2245	0.2202	0.0464	0.2212
π_N	0.4344	0.4414	0.0336	0.4447
π_S	0.3411	0.3355	0.0671	0.3341

Table G.3: Monte Carlo Results: Unobserved Types, $\beta_S = \beta_N = \beta$

Notes: Each model was simulated 200 times.

	Mean	Median	Std.Dev	True
N=300				
δ	0.6613	0.6013	0.1795	0.7
β_N	0.3773	0.2870	0.2790	0.4
β_S	0.2500	0.1781	0.2699	0.1
θ	1.0714	1.4649	0.5245	1.0
π_C	0.2780	0.2543	0.1475	0.2215
π_N	0.3837	0.3794	0.0490	0.4447
π_S	0.3383	0.3086	0.1582	0.3339
N=600				
δ	0.6805	0.6239	0.1670	0.7
β_N	0.3813	0.3476	0.2359	0.4
β_S	0.2297	0.1725	0.2358	0.1
θ	1.1217	1.4057	0.4615	1.0
π_C	0.2516	0.2395	0.1283	0.2213
π_N	0.3965	0.4064	0.0470	0.4448
π_S	0.3519	0.3380	0.1471	0.3339
N=1200				
δ	0.6825	0.6492	0.1411	0.7
β_N	0.3996	0.3451	0.2251	0.4
β_S	0.2165	0.1424	0.2563	0.1
θ	1.040	1.2328	0.5022	1.0
π_C	0.2251	0.2066	0.1079	0.2213
π_N	0.4095	0.4261	0.0453	0.4449
π_S	0.3654	0.3488	0.1338	0.3335
N=300				
δ	0.8291	0.8358	0.1478	0.9
β_N	0.4491	0.4124	0.2061	0.3
β_S	0.6177	0.5963	0.2599	0.6
θ	1.0717	1.4364	0.5242	1.0
π_C	0.2991	0.2882	0.1689	0.2215
π_N	0.3927	04031	0.0699	0.4446
π_S	0.3082	0.2771	0.1676	0.3339
N=600				
δ	0.8361	0.8399	0.1365	0.9
β_N	0.4014	0.3709	0.1661	0.3
β_S	0.6615	0.6505	0.2300	0.6
θ	1.0224	1.1752	0.5124	1.0
π_C	0.2781	0.2614	0.1296	0.2218
π_N	0.4158	0.4288	0.0584	0.4449
π_S	0.3062	0.3092	0.1265	0.3333
 N=1200				
δ	0.8522	0.8598	0.1102	0.9
β_N	0.3759	0.3535	0.1266	0.3
β_S	0.6329	0.6057	0.2085	0.6
θ	1.0728	1.2702	0.4887	1.0
π_C	0.2376	0.2347	0.0805	0.2218
π_N	0.4362	0.4453	0.0413	0.4449
πc	0.3261	0.3244	0.1004	0.3333

Table G.4: Monte Carlo Results: Unobserved Types, $\beta_N \neq \beta_S$

Notes: Each model was simulated 200 times.

_

	Mean	Median	Std . Dev	True
N=300				
δ	0.3594	0.3598	0.0141	0.9
β_N	1	1		0.2
β_S	1	1		0.1
θ	1.3474	1.3846	0.7131	1.0
π_C	1			0.2215
π_N	0			0.4447
π_S	0			0.3338
N=600				
δ	0.3598	0.3598	0.0091	0.9
β_N	1	1		0.2
β_S	1	1		0.1
θ	1.2034	1.1637	0.7562	1.0
π_C	1			0.2215
π_N	0			0.4447
π_S	0			0.3338
N=300				
δ	0.4472	0.3885	0.1482	0.36
β_N	0.5391	0.5641	0.3776	1.0
β_S	0.6049	0.7694	0.3765	1.0
θ	1.2867	1.3637	0.7457	1.0
π_C	0.5642	0.6279	0.3137	1.0
π_N	0.1797	0.1279	0.1421	0
π_S	0.2561	0.1823	0.2541	0
N = 600				
δ	0.4173	0.3798	0.0958	0.36
β_N	0.6636	0.7130	0.3320	1.0
β_S	0.6117	0.7662	0.3729	1.0
θ	1.3364	1.4409	0.7272	1.0
π_C	0.5658	0.6258	0.3242	1.0
π_N	0.1841	0.1297	0.1467	0
π_S	0.2502	0.1568	0.2580	0

Table G.5: Monte Carlo Results: 3 Types vs 1 Type Comparison

Notes: The top two panels provide placebo simulations by estimating a single time-consistent type when the data is generated from three distinct types. Tho bottom two panels provide the analogous simulations by estimating three type models when the data is generated from a single type. Each model was simulated 200 times.