

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, Revised December 2025

We are grateful to Øystein Daljord, Hanming Fang, Paul Heidhues, Nathan Hendren, Han Hong, Gautam Rao, 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 usual disclaimers apply. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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, Revised December 2025
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 Odisha, India, we find that about two-thirds of our sample comprises time-inconsistent agents and that both 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 four times the price of an ITN.

Aprajit Mahajan
University of California, Berkeley
and NBER
aprajit@gmail.com

Alessandro Tarozzi
Universitat Pompeu Fabra Barcelona
Department of Economics
alessandro.tarozzi@upf.edu

Christian Michel
Universitat Pompeu Fabra Barcelona
Department of Economics and Business
christian.michel@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 low costs and high expected returns (Frederick et al. 2002, DellaVigna 2009, Sprenger 2015, Ericson and Laibson 2019). 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), so that, 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 a reward at current date t over a larger one at date $t + k$, but instead choose the larger reward if both payoff dates are shifted forward by the same length of time s (i.e. to $t + s$ and $t + s + k$).¹ Such choices are not consistent with standard models of inter-temporal choice.

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’ agents who are aware of having such time-inconsistent preferences and ‘naïve’ ones 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 key parameters β and δ is non-trivial. In

¹See Andreoni and Sprenger (2012) for an alternative explanation for these findings and Augenblick et al. (2015) for a similar finding when choices are over effort rather than money.

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 *a fortiori* to both β and δ in the hyperbolic “ β – δ ” formulations of time-inconsistent preferences which dominate empirical work on time-inconsistency. In addition, it seems desirable and important to allow for heterogeneity in time preferences, especially in applied work—for instance allowing for both time-consistent and inconsistent agents. However, it is rare for agent type to be directly observed by the researcher so that a model with *unobserved* types seems more appropriate.

We make 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 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 Odisha, 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 variables z that only affect current utility via the perceived value of future states. The second is the presence of variables r that act as (imperfect) signals of agent type but which, conditional on agent type and observables, provide 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.

In the general version of the model we allow for an unknown (but finite) number of types with possibly time-inconsistent preferences. We first identify the total number of types in the population and the nature of the time-preferences for each type, classifying each 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 existence of 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, and we show point identification under a further set of additional (but commonly assumed) conditions.

Next, we introduce our empirical application. Malaria is an enormous global health burden and is endemic in our study region, where survey respondents indeed report high expected costs of malaria as well as strong beliefs in the efficacy of ITNs in preventing it. Despite this, nearly half of our sample households do not purchase ITNs offered through a micro-credit intervention. To rationalize these choices, we explore counterfactual pricing policies by estimating a structural

dynamic choice model of ITN purchase and re-treatment that allows for time-inconsistent agents. Monte Carlo simulations suggest that time-preference parameters are well estimated with sample sizes similar to those in the application. We then estimate the model and find that approximately one-third of respondents are time-consistent while about one-half are naïve inconsistent and the remaining one-sixth are sophisticated inconsistent. The discount rate for consistent agents is close to one. Further, we find that naïve and sophisticated agents are considerably present-biased with our preferred estimates of β being 0.06 (for naïve agents) and 0.16 (for sophisticated agents). The estimates of both the population distribution of types and of the separate β parameters are, to our knowledge, new to the literature. Finally, we find that the per-period utilities do vary across agent types and that ignoring this exacerbates the already considerable high levels of present-bias.

Next, we evaluate the extent to which present-biased but sophisticated agents are more likely to choose specially designed “commitment” products (Bryan et al. 2010). The ITNs in our context require regular re-treatment with insecticide in order to remain effective against mosquitoes. We offered a choice between a standard contract (with the option to purchase re-treatment at a later time) and a “commitment contract” which also included a bundle of two consecutive re-treatments. The commitment contract was designed to mitigate the time-inconsistency problem associated with re-treatment. 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—commonly assumed—deterministic mapping whereby the choice of commitment products reveals an agent to be sophisticated. 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.

Finally, 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 re-treatments. Since ITNs reduce the risk of malaria, fewer ITN purchases and re-treatments increase the likelihood of contracting malaria. We find that the median (un-discounted) additional expected total cost of malaria during our study period exceeds the price of a treated net by a factor of around four. 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 rationalizes 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 large literature on time-inconsistency and structural estimation with unknown types we focus on work closest to our approach.² 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 (2020b), Abbring et al. (2019) and Fang and Wang (2015), with important differences. First, we allow for multiple unobserved types while these papers consider the case of a single observed type. Second, the distribution of these multiple types

²See, e.g., Aguirregabiria and Mira (2010) or Arcidiacono and Ellickson (2011) for a survey on dynamic discrete choice structural models, and DellaVigna (2018) on structural models in behavioral economics.

is of intrinsic interest, and we can assess the time-inconsistency problem in terms of both the type frequency in the population and the type-specific magnitude of time-inconsistency. Third, our model is motivated by our specific setting in which purposely collected data (beliefs about future state evolution) provide a natural candidate for the exclusion restrictions. This is a key 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 \(2002\)](#)—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 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 2020a](#) 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 impose 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, and 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 \(2008\)](#) who use information about expected future choices to improve precision in the context of a structural dynamic model, while we use expectations about state transitions 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 lab experiment with real effort choices to identify potentially heterogeneous time-preference parameters for agents who may be partially sophisticated. [Bai et al. \(2021\)](#) 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, attributing this to partial naïveté. [Heidhues and Strack \(2021\)](#) provide identification results with partial naïveté in a stopping problem when data on both the stopping probabilities and the continuation value are available. [Martinez et al. \(2021\)](#) adapt their model in the context of filing tax returns and find non-negligible present-bias (assuming a

per-period discount factor $\delta = 1$). Our paper is also related to Andreoni et al. (2016) who estimate individual level time-preference parameters and use 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 dynamic discrete choice model with a finite action and state space. The model has three periods, the minimum required to identify the time-preference parameters. We begin by defining and placing assumptions on the state and action spaces, the transition probabilities, the class of acceptable decision rules and finally the preferences and objective function maximized by the agent.

STATE SPACE: \mathcal{S}_t . The state space $\mathcal{S}_t \equiv \mathcal{X}_t \times \mathcal{Z}_t \times \mathcal{E}_t$ 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 is 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 $z_t \in \mathcal{Z}_t$ are excluded. In the empirical application z_t comprises subjective beliefs elicited from the agent about elements of the distribution of x_{t+1} and these are plausibly excludable from the static payoff function (conditional on the observed state) – see Assumption B for a formal statement and the subsequent discussion of the exclusion restriction.

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 K_t of actions $a_t \in \mathcal{A}_t$.

TRANSITION PROBABILITIES: $\mathbb{P}(s_t|s_{t-1}, a_{t-1})$.

Let $\mathbb{P}(s_t|s_{t-1}, a_{t-1})$ denote the distribution function of the random vector $s_t \in \mathcal{S}_t$ conditional on (s_{t-1}, a_{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 the conditional distribution of s_t given the entire state and action history through period $t-1$ only depends on last period's state and action, that is, (s_{t-1}, a_{t-1}) . Incorporating dependencies across longer horizons requires redefining the state variable to include sufficient lags.

ERROR TERMS ϵ_t .

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 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.

DECISION RULES: d_t .

The decision rule in period t , d_t , is a mapping from \mathcal{S}_t to \mathcal{A}_t . We do not allow for history-dependent decision rules which thus cannot be mappings from $\prod_{s=1}^{t-1}(\mathcal{S}_s, \mathcal{A}_s) \times \mathcal{S}_t$ to the action space. 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 common in empirical work, we assume that preferences are additively time-separable, and parameterize time inconsistency using the tractable (β, δ) formulation described in [Strotz \(1955\)](#).³ Then, for a given sequence of actions $\{a_t\}_{t=1}^3$, the utility of an agent of type τ is:

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

Broadly, we deal with three types of agents: time-consistent agents (denoted by τ_C or C), time-inconsistent naïve agents (τ_N or N) and time-inconsistent sophisticated agents (τ_S or 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) the theory can accommodate partially sophisticated agents and we provide set identification results for this case.

Following [O'Donoghue and Rabin \(1999\)](#), time-consistent agents ($\tau = \tau_c$) have $\beta_{\tau_c} \equiv \beta_C = 1$, which corresponds to the standard case of exponential discounting. Such agents will maximize eq. (1) 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

³This 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.

solve the maximization problem using backward induction. However, while sophisticated agents ($\tau = \tau_S, \beta_{\tau_S} \equiv \beta_S < 1$) also recognize their *future* present-bias, naïve agents ($\beta_{\tau_N} \equiv \beta_N < 1$) do not. For the econometrician, this generates differences in predicted behavior that can be exploited for identification, as we show below.

The formulation in eq. (1) allows for type-varying exponential (δ_τ) and hyperbolic (β_τ) parameters. Previous empirical work assumes that $\beta_N = \beta_S$ and that $\delta_N = \delta_S = \delta_C$. We relax these restrictions while still retaining point-identification for all parameters except the time-preference parameters for naïve agents. The formulation also allows for time-varying type-specific per-period utilities $\tilde{u}_t(\cdot; \tau)$. This flexibility is important since it allows us to 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-preferences, without confounding their relative role. Second, one can examine how much of the difference in choices *between types* is driven by differing preferences over states versus different time preferences. Third, the time- and type-varying formulation provides a mechanism for flexibly accounting for serially correlated unobserved heterogeneity. Our formulation nests the model where types only differ in the degree of present-bias so that we can evaluate the role of present bias relative to those of other differences in preferences in explaining behavior.

We now have sufficient notation in place to state the first set of basic assumptions. These assumptions are always invoked together and we will refer to them jointly as **Assumption B** (for “Basic” assumptions). We have already discussed the first two (and they are standard in the dynamic discrete choice literature) and discuss the remaining three below.

ASSUMPTION B (Basic Assumptions).

Markov Property:

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

Independent Errors with Known 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 .

Exclusion Restriction: The variable z_t does not enter the per-period utility function, that is, $\tilde{u}_t(x_t, z_t, \epsilon_t, a_t; \tau) = \tilde{u}_t(x_t, \epsilon_t, a_t; \tau)$.

Additive Separability: For each type $\tau \in \mathcal{T}$ $\tilde{u}_t(x_t, \epsilon_t, a_t; \tau) = u_t(x_t, a_t; \tau) + \epsilon_t(a_t)$.

Normalization: Utility in period t for a base action $a_t = 0$ is known for all types and for all states, i.e. $u_t(x_t, 0; \tau)$ is known for all $(x_t, \tau) \in \mathcal{X}_t \times \mathcal{T}$.

The exclusion restriction requires that there exist z_t that does not enter the per-period utility function. Intuitively, this variable will provide the basis for inducing variation in the forward-looking component of the value function while keeping current period utility constant. This strategy builds on the ideas (though not the precise assumption) in [Magnac and Thesmar \(2002\)](#) and is also used by [Abbring and Daljord \(2020b\)](#) and [Abbring et al. \(2019\)](#). As [Abbring and Daljord \(2020b\)](#) point out, the assumption in [Magnac and Thesmar](#) (in their Section 4.2) imposes conditions on the value function (rather than the per-period utility) and is therefore not straightforward to interpret.⁴

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 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, 2020](#)). The assumption does, however, rule out models where beliefs about the future affect current-period utility directly (e.g. [Brunnermeier and Parker, 2005; Kőszegi, 2010](#)). Beliefs that are effectively exogenous (i.e. are not determined by actions, preferences or other state variables) are potential candidates for the exclusion restriction. By the same token, endogenous beliefs or beliefs based on endogenous information acquisition (as in e.g. [Fuster et al., 2022](#)) may be incompatible with the exclusion restriction.

Beyond beliefs, any variable that does not affect current period pay-offs but does affect the forward-looking component of the pay-off function is a potential candidate for z . Another example could be a current measure of a future pay-off, such as marketing tools that promise a future pay-off based on current period action (e.g. a free coffee after 10 purchases), or variables that lead some agents to be better informed about the likelihood of future payoff-relevant events.

The last two assumptions within [Assumption B](#) are standard in the dynamic choice literature. First, we maintain the additive separability of utility in the unobserved state time-varying variables ϵ_t . Second, we assume that payoffs from a base action are known in each state in each period. Such normalizations are standard although recent work has emphasized that counterfactual analyses can be sensitive to them.⁵

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

⁴In 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 denote v . Since these play no role for identification, we will omit them as arguments in preferences for the most part, although we do include such variables in the empirical application.

⁵One example of how to relax this assumption is given in Appendix A of [Abbring and Daljord \(2020b\)](#) who consider a time-consistent model where the utility normalization is known only up to a constant shift.

agents' imperfect understanding of the choice problem, 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 $\mathcal{R} = \{r_C, r_S, r_N\}$ where each element corresponds to a unique type, but we can also allow for sub-types within a particular class of time-inconsistent preferences, in which case $\mathcal{R} = \{r_{C_1}, \dots, r_{C_J}, r_{S_1}, \dots, r_{S_K}, r_{N_1}, \dots, r_{N_L}\}$. This allows e.g. for multiple types of time-consistent agents who may differ in their time preferences or per-period utility functions. Here, types are directly observed so we could equivalently have stated these restrictions as a condition on the set of possible types, but we prefer this formulation because it provides a natural generalization to the unobserved types case.

The key starting point for identification are the type-specific choice probabilities $\mathbb{P}_\tau(a_t^* = a|x_t, z_t)$, which are directly observed since here we assume that agent type is a known function of the observed type proxy. In addition, we assume that conditional on type, the proxy is uninformative about choice.

ASSUMPTION D1 (Directly Observed Types and Exclusion Restriction). *Agent type is a known deterministic function of r and therefore choice probabilities are directly observed for each type. For an agent of type τ , $\mathbb{P}(a_t^* = a|x_t, z_t, r = r_\tau) = \mathbb{P}_\tau(a_t^* = a|x_t, z_t, r = r_\tau) = \mathbb{P}_\tau(a_t^* = a|x_t, z_t)$.*

Implicit in the formulation above is that type-observability is equivalent to knowledge of type-identity (i.e. whether a type is consistent, naïve or sophisticated). However, for the first set of results (collected in [Lemma 1](#)) we do not need to know the type identity – i.e. we do not need to know whether the identified type-specific choice probability corresponds to a consistent, naïve, or sophisticated type; this added generality will prove useful when we turn to the unobserved types case.

We now turn to identification of the preference parameters. Since this is a finite-horizon dynamic choice problem, we can use backward induction and we start from the terminal period, when the agent chooses action k if and only if $\tilde{u}_3(s_3, k; \tau) > \tilde{u}_3(s_3, a; \tau) \forall a \neq k$ (we do not index actions by time unless there is ambiguity). Under [Assumption B](#) we can write the choice probability as

$$\mathbb{P}_\tau(a_3^* = k|x_3, z_3) = \mathbb{P}\left(k = \operatorname{argmax}_{a \in \mathcal{A}_3} \left\{ u_3(x_3, a; \tau) + \epsilon_3(a) \right\} \middle| x_3, z_3\right).^6$$

The decision in the terminal period is described by a standard static discrete choice model with a

⁶Since the terminal period does not have a forward looking component and we do not incorporate learning for future periods, we do not require the existence of z_3 for identification.

known error distribution. We can thus invert the choice probability to directly identify the period 3 utilities up to the normalization in [Assumption B](#).⁷

Next, in period 2 the conditional probability that an agent chooses action k is given by

$$\mathbb{P}_\tau(a_2^* = k|x_2, z_2) = \mathbb{P}\left(k = \operatorname{argmax}_{a \in \mathcal{A}_2} \left\{ u_2(x_2, a; \tau) + \epsilon_2(a) + \beta_\tau \delta_\tau \int v_{\tau,3}^*(s_3) dF(s_3|x_2, z_2, a)\right\} \middle| x_2, z_2\right),$$

where $dF(s_{t+1}|x_t, z_t, a)$ is our notation for the distribution of the vector $s_{t+1} = (x_{t+1}, z_{t+1}, \epsilon_{t+1})$ conditional on the vector (x_t, z_t, a) that is used by the agent when making choices in period t . Given the independence between the unobserved and observed state variables,

$$dF(s_{t+1}|x_t, z_t, a) = dF(x_{t+1}, z_{t+1}|x_t, z_t, a) dF(\epsilon_{t+1}) \quad (2)$$

where $dF(\epsilon_{t+1})$ is known (by [Assumption B](#)). We further assume that $dF(x_{t+1}, z_{t+1}|x_t, z_t, a)$ is identified so that eq. (2) is identified. However, the precise manner in which this is achieved will depend upon the nature of the data generating process and the nature of the excluded variables z_t . One conventional approach in dynamic choice models is to impose rational expectations. Combined with knowledge of the joint distribution of $\{a_t^*, x_t, z_t\}_{t=1}^T$, this implies that $dF(x_{t+1}, z_{t+1}|x_t, z_t, a)$ (and hence eq. (2)) is identified. In some contexts, such an assumption may be infeasible or unreasonable without further modification. For instance, if z_t are elicited beliefs about the likelihood of future states then further assumptions are needed to ensure that such elicitations do not impose onerous data collection requirements for identification of eq. (2). One set of assumptions is to restrict beliefs z_t to be solely about x_{t+1} and require

$$dF(x_{t+1}, z_{t+1}|x_t, z_t, a) = dF(x_{t+1}|x_t, z_t, a) dF(z_{t+1}).$$

so that (a) next-period beliefs (z_{t+1}) and next-period states (x_{t+1}) are conditionally independent given (x_t, z_t, a) and that (b) the distribution of next-period beliefs does not depend upon current beliefs, state or action (i.e. $dF(z_{t+1}|z_t, x_t, a) = dF(z_{t+1})$). With these restrictions, beliefs are only about static-payoff relevant state variables and evolve independently of states and actions. This ensures, for instance, that we do not need to elicit beliefs in period t about beliefs in period $t+1$ and so the data collection requirements are not as onerous. It does, however, rule out learning or belief-updating as a function of past states and actions. Formally, define $z_t = \{Q_{t,t+1}(a) : a \in \mathcal{A}_t\}$ where $Q_{t,t+1}(a)$ is a matrix of elicited beliefs of dimension $\#\mathcal{X}_t \times \#\mathcal{X}_{t+1}$ with an element $q(a, x', x'')$ denoting the agent's elicited belief of being in state x'' in period $t+1$ conditional on being in state x' in period t and taking action a in period t . We then assume that $dF(x_{t+1}|x_t, z_t, a) = q(a, x_t, x_{t+1})$. In the application, elicited beliefs (our candidate excluded variable) are only observed at one point in time (so $z_t = z$) so the transition probabilities reduce to $dF(x_{t+1}|x_t, z, a)$. These are directly

⁷See online [Appendix C](#) for an alternative and self-contained argument or see [Hotz and Miller \(1993\)](#).

identified since they are elicited from each agent.⁸ Next, define

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

We can then use the standard Hotz-Miller inversion of the type-specific conditional choice probabilities to directly identify the left-hand side of the expression below:

$$g_{\tau,2,k}(x_2, z_2) \equiv u_2(x_2, k; \tau) - u_2(x_2, 0; \tau) + \beta_\tau \delta_\tau \int h_{\tau,3}(x_3, z_3) dF_{\Delta,k}(x_3, z_3 | x_2, z_2), \quad (3)$$

where

$$h_{\tau,3}(x_3, z_3) \equiv \int v_{\tau,3}^*(s_3) dF(\epsilon_3), \quad (4)$$

and where $dF_{\Delta,k}(x_3, z_3 | x_2, z_2) \equiv (dF(x_3, z_3 | x_2, z_2, k) - dF(x_3, z_3 | x_2, z_2, 0))$ is our short-hand notation for the signed measure which is the difference in the conditional probabilities of (x_3, z_3) given (x_2, z_2) when action k is taken and when action 0 is taken.

We next explore which of the unknown elements on the right hand side of eq. (3)—the utility functions and the discount rates—can be identified. First, the integral is directly identified since (a) $h_{\tau,3}(\cdot)$ is identified (because $u_3(\cdot)$ is identified and the distribution of ϵ_3 is known) and (b) $dF(x_3, z_3 | x_2, z_2, k)$ is directly identified from the data so that the signed measure $dF_{\Delta,k}(\cdot)$ is identified. Next, z_2 only enters the last term in eq. (3) so we can use variation in z_2 (conditional on x_2) to isolate this last term. This requires that the variation in z_2 translates into variation in the integral of the period 3 value function (where integrals are taken using the signed measure defined above).⁹ This variation allows us to isolate the forward-looking component of the value function and, along with the previously identified terms in eq. (3), to identify the product $\beta_\tau \delta_\tau$. While we do not prove that such variation is necessary for identification, the non-identification of discount parameters in standard dynamic choice models can be traced to the lack of variation of this kind. In fact, this variation is a version of the rank condition in Proposition 4 of [Magnac and Thesmar \(2002\)](#) adapted to the context of our model. Since we require such an assumption for all three periods, we state it here for all periods for brevity.

ASSUMPTION D2 (Rank Condition). *For $t \in \{2, 3\}$ the distribution of z_{t-1} conditional on x_{t-1} has at least two points of support (z'_{t-1}, z''_{t-1}) and there exists at least one action k_{t-1} and one point in the support of \mathcal{X}_{t-1} such that*

⁸In our empirical application we use beliefs about malaria risk under various ITN usage scenarios (i.e. different values of a) as well as about transition probabilities for income. We did not condition on current state (x_t) in the elicitation to cut down survey length, and we do not allow beliefs to vary over time, see [Section 4.1](#) for details. In [Appendix B](#) we show that identification is achieved also with time-invariant beliefs, given that we will only require within-period variation in z . We thank a referee for comments and suggestions on this point.

⁹When the z_t are beliefs about x_{t+1} , the assumption requires that the variation in beliefs across the population induces sufficient variation in the expectations of the forward-looking component of the value function.

$$\int h_{\tau,t}(x_t, z_t) (\mathrm{d}F_{\Delta, k_{t-1}}(x_t, z_t | x_{t-1}, z'_{t-1}) - \mathrm{d}F_{\Delta, k_{t-1}}(x_t, z_t | x_{t-1}, z''_{t-1})) \neq 0.$$

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_1 and one point in the support of \mathcal{X}_1 such that for sophisticated agents,

$$\int h_S^B(x_2, z_2, \beta_S \delta_S) (\mathrm{d}F_{\Delta, k}(x_2, z_2 | x_1, z'_1) - \mathrm{d}F_{\Delta, k_1}(x_2, z_2 | x_1, z''_1)) \neq 0.$$

We define and describe $h_{\tau,2}(\cdot)$ and $h_S^B(\cdot)$ below in eq. (8). For $t = 3$, this assumption is in principle testable, since the function $h_{\tau,3}(\cdot)$ is identified for all types and $\mathrm{d}F_{\Delta}(\cdot)$ is known. For $t = 2$, this is not the case since $h_{\tau,2}(\cdot)$ is not identified for all types. With the rank condition in place, we can separately identify the per-period preferences (for $t = 2, 3$) and the product of the time-preference parameters.

LEMMA 1 (Identification for Periods 3 and 2). *Consider an agent maximizing (1) and suppose that the model satisfies Assumptions B, D1 and D2. Then*

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

All proofs are relegated to the appendix. The intuition for the result is that, following the Hotz-Miller inversion, the exclusion restriction provides variation in the agents' future expected utilities which is used to identify $\beta_\tau \delta_\tau$. With that in hand, we can then recover the period 3 payoff functions. Two points are worth keeping in mind for the next section with unobserved types. First, the proof reveals that we only need the type-specific choice probabilities to be identified (so it is not necessary to observe each agent's type). Second, knowledge of the type-identities is not required (i.e. we do not need to know whether a given type-specific choice probability belongs to a time-consistent or inconsistent type).

Next, we turn to identification of β_τ and δ_τ separately, and of the period 1 utility functions. As before, the conditional choice probability that an agent chooses action k in period 1 is given by

$$\mathbb{P}_\tau(a_1^* = k | x_1, z_1) = \mathbb{P} \left(k = \operatorname{argmax}_{a \in \mathcal{A}_1} \left\{ u_1(x_1, a; \tau) + \epsilon_1(a) + \beta_\tau \delta_\tau \int v_{\tau,2}^*(s_2) \mathrm{d}F(s_2 | x_1, z_1, a) \right\} \middle| x_1, z_1 \right). \quad (5)$$

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

$$\begin{aligned}
v_{\tau,2}^*(s_2) &= \sum_{a \in \mathcal{A}_2} v_{\tau,2}(s_2, a, \delta_\tau) A_\tau(s_2, a, \tilde{\beta}_\tau \delta_\tau), \text{ where} \\
v_{\tau,2}(s_2, a, \mathbf{d}_1) &= u_2(x_2, a; \tau) + \epsilon_2(a) + \mathbf{d}_1 \int v_{\tau,3}^*(s_3) dF(s_3|x_2, z_2, a), \text{ and} \\
A_\tau(s_2, a, \mathbf{d}_2) &= \mathbb{I} \left\{ a = \underset{j \in \mathcal{A}_2}{\operatorname{argmax}} v_{\tau,2}(s_2, j, \mathbf{d}_2) \right\}.
\end{aligned} \tag{6}$$

Here, $v_{\tau,2}^*(s_2)$ is the continuation value from period 2 onwards *from the standpoint of period 1* and is defined in terms of the utility measure $v_{\tau,2}(\cdot)$ and the choice indicator $A_\tau(\cdot)$.¹⁰ We introduce the arguments $(\mathbf{d}_1, \mathbf{d}_2)$ since the identification results will involve evaluating the value function at different candidate values for these parameters. The argument \mathbf{d}_1 in $v_{\tau,2}(\cdot)$ governs utility trade-offs between periods 2 and 3 from the view-point of $t=1$ and is equal to δ_τ for all types.¹¹ The argument \mathbf{d}_2 is the rate of time-preference that the period 1 self *believes* she will use to make choices in period 2. In standard dynamic programming problems—i.e for time-consistent agents, $\mathbf{d}_2 = \delta$. Agents who are completely aware of their future present-bias will have $\mathbf{d}_2 = \beta\delta$. In contrast, agents who are completely unaware of their future present bias are those with $\mathbf{d}_2 = \delta$ and $\beta < 1$. More generally, we can write $\mathbf{d}_2 = \tilde{\beta}_\tau \delta_\tau$, where the parameter $\tilde{\beta}_\tau$ is interpretable as the extent of present-bias that the agent in $t = 1$ thinks her 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 “naiïete” (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 naiïe.

ASSUMPTION D3 (Three Types). *The parameter $\tilde{\beta}_\tau$ is equal to 1 for consistent and naïve agents and is equal to β_S for sophisticated agents.*

In Section 3.3 we explore the weaker assumption of partial sophistication, in which case $\tilde{\beta}_\tau \in [\beta_\tau, 1]$ and $\tilde{\beta}_\tau$ is an additional parameter that is only set-identified. Intuitively, with partial sophistication the extra parameter $\tilde{\beta}_\tau$ drives choices in period 1 but observed choices in periods 2 and 3 are uninformative about it, because choice in the later periods involve β_τ and not $\tilde{\beta}_\tau$.

Identification for $t = 1$ follows the same general strategy as for $t = 2$. We first invert the type-specific conditional choice probabilities to directly identify the function $g_{\tau,1,k}(\cdot)$:

$$g_{\tau,1,k}(x_1, z_1) = u_1(x_1, k; \tau) - u_1(x_1, 0; \tau) + \beta_\tau \delta_\tau \int h_{\tau,2}(x_2, z_2) dF_{\Delta,k}(x_2, z_2|x_1, z_1), \tag{7}$$

¹⁰We use A to represent both the event and an indicator for the event.

¹¹Note that $\beta_\tau \delta_\tau$ multiplies utility at both $t=2$ and $=3$ (see eq. (5)), so the trade-off between the two periods from the viewpoint of $t = 1$ is only governed by the discount factor in $v_{\tau,2}(\cdot)$.

where

$$\begin{aligned}
h_{\tau,2}(x_2, z_2) &\equiv \int v_{\tau,2}^*(s_2) dF(\epsilon_2) = \sum_{a \in \mathcal{A}_2} \int v_{\tau,2}(s_2, a, \delta_\tau) A_\tau(s_2, a, \tilde{\beta}_\tau \delta_\tau) dF(\epsilon_2) \\
&= \sum_{a \in \mathcal{A}_2} \int \left(u_2(x_2, a; \tau) + \epsilon_2(a) + \delta_\tau \underbrace{\int v_{\tau,3}^*(s_3) dF(s_3|x_2, z_2, a)}_{\equiv q_\tau(x_2, z_2, a)} \right) A_\tau(s_2, a, \tilde{\beta}_\tau \delta_\tau) dF(\epsilon_2) \\
&= \underbrace{\sum_{a \in \mathcal{A}_2} \int (u_2(x_2, a; \tau) + \epsilon_2(a)) A_\tau(s_2, a, \tilde{\beta}_\tau \delta_\tau) dF(\epsilon_2)}_{\equiv \tilde{h}_\tau^A(x_2, z_2, \tilde{\beta}_\tau \delta_\tau)} + \underbrace{\delta_\tau \cdot \sum_{a \in \mathcal{A}_2} q_\tau(x_2, z_2, a) \int A_\tau(s_2, a, \tilde{\beta}_\tau \delta_\tau) dF(\epsilon_2)}_{\equiv \tilde{h}_\tau^B(x_2, z_2, \tilde{\beta}_\tau \delta_\tau)} \\
&= \tilde{h}_\tau^A(x_2, z_2, \tilde{\beta}_\tau \delta_\tau) + \delta_\tau \cdot \tilde{h}_\tau^B(x_2, z_2, \tilde{\beta}_\tau \delta_\tau). \tag{8}
\end{aligned}$$

Note that in the expression above the expected value of $v_{\tau,3}^*(s_3)$ (i.e. $q_\tau(x_2, z_2, a)$) is multiplied by the discount factor δ_τ and not $\beta_\tau \delta_\tau$, because the hyperbolic parameter β_τ does not directly enter into the intertemporal decision problem 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 $h_{\tau,2}(x_2, z_2)$ represents how much an agent at $t = 1$ values being in state (x_2, z_2) at $t = 2$ after (a) incorporating her own perceived future behavior in periods 2 and 3 and (b) taking expectations over the unobserved state variables in period 2 and over all variables in period 3. As the right-hand side of eq. (8) makes explicit, the only unknowns in $h_\tau(\cdot)$ are the pair $(\tilde{\beta}_\tau \delta_\tau, \delta_\tau)$. Further, observe that the product $\beta_\tau \delta_\tau$ is identified (by [Lemma 1](#)). How informative $\beta_\tau \delta_\tau$ is for the unknown parameters $(\tilde{\beta}_\tau \delta_\tau, \delta_\tau)$ varies by type and so we discuss identification separately by type below.

3.1.1 Identification for Consistent and Sophisticated Agents

First, for consistent agents, $\tilde{\beta}_C = \beta_C = 1$ and $\beta_C \delta_C = \delta_C$ which is identified by the previous lemma. Therefore $(\tilde{\beta}_C \delta_C, \delta_C) = (\delta_C, \delta_C)$ is identified which in turn implies that $h_{C,2}(x_2, z_2)$ is identified. Thus, all the elements in the last term in eq. (7) are identified. Therefore, period 1 preferences (i.e. the first two terms in eq. (7)) are identified (without any further assumptions).

Second, for a sophisticated type, $\tilde{\beta}_S = \beta_S$ (by [Assumption D3](#)) so that $\tilde{\beta}_S \delta_S = \beta_S \delta_S$ and the latter is identified by the previous Lemma. Therefore, the only unknown parameter in eq. (7) is δ_S which enters it linearly. Rewriting eq. (7) for sophisticated types:

$$\begin{aligned}
g_{S,1,k}(x_1, z_1) &= u_1(x_1, k; \tau_S) - u_1(x_1, 0; \tau_S) + \beta_S \delta_S \int h_S^A(x_2, z_2, \beta_S \delta_S) dF_{\Delta,k}(x_2, z_2|x_1, z_1) \\
&\quad + \delta_S (\beta_S \delta_S) \int h_S^B(x_2, z_2, \beta_S \delta_S) dF_{\Delta,k}(x_2, z_2|x_1, z_1), \tag{9}
\end{aligned}$$

where the δ_S term that multiplies $\beta_S \delta_S$ in the last term is not identified. As before, we can use

variation in z_1 (guaranteed by the second rank condition of [Assumption D2](#)) to identify δ_S and then in the next step identify the first period payoffs. We collect the identification results for all remaining parameters for consistent and sophisticated types here.

LEMMA 2 (Period 1: Identification for Consistent and Sophisticated types).

Consider an agent of type τ_C solving the problem (1) at $t = 1$ and suppose that the model satisfies [Assumptions B](#), [D1](#) and [D2](#). 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 (1) at $t = 1$ and suppose that the model satisfies [Assumptions B](#), [D1](#), [D2](#) and [D3](#). 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_S and β_S) for sophisticated agents are identified.

3.1.2 Identification for Naïve Agents

For both consistent and sophisticated agents knowledge of $\beta_\tau \delta_\tau$ was sufficient to identify $\tilde{\beta}_\tau \delta_\tau$ but for naïve agents knowledge of $\beta_N \delta_N$ is *not* sufficient to identify $\tilde{\beta}_N \delta_N = \delta_N$. The unknown parameter δ_N enters eq. (7) in a non-linear fashion through the functions $(\tilde{h}_\tau^A(\cdot), \tilde{h}_\tau^B(\cdot))$ which are now not identified (unlike for consistent and sophisticated types), so that identification will require stronger assumptions.

One simple assumption that ensures point-identification is that the exponential parameter for the naïve type is the same as for the consistent or sophisticated types: since δ_S and δ_C are already identified, this trivially guarantees identification of δ_N , and this in turn implies identification of period 1 payoff functions for the naïve type. 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 be identified even if no naïve agents are present, while the equal discount rate assumption is only informative when sophisticated (or consistent) agents are also present in addition to naïve types.

In [Appendix A.2](#), we provide an alternative set of conditions to generate bounds on the time-preference parameters for the naïve type. The argument has two steps: (a) first, in [Lemma A1](#) we place additional structure on the transition probabilities to identify the function $h_{\tau,2}(\cdot)$ defined in eq. (8) above (up to a normalization). This allows us to identify the first-period payoff function using eq. (7). Next, we construct an alternative measure of $h_{\tau,2}(\cdot)$ using period 2 and period 3 choices only. A comparison of these two functions provides a measure of the difference between an agent's beliefs at $t = 1$ about her subsequent behavior in periods 2 and 3 and her actual choices in those periods. This comparison allows us to place bounds on δ_N (and consequently on β_N) which we do in [Lemma A2](#). We also discuss an (untestable) monotonicity restriction ([Assumption DA2](#)) that yields point identification for (δ_N, β_N) in [Lemma A3](#). In the empirical application we will assume that all types share the same exponential discount parameter, thereby circumventing the identification problem.

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 D1](#) so that we do not observe the type for each observation. 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-specific choice probabilities, without assigning them to their respective types.
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 type-specific choice probabilities $\{\mathbb{P}_\tau(a_t|x_t, z_t, r, v)\}_t$, leaving the type-identity of τ unknown. The identity of each type τ is then identified in step 3, and the preference parameters 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 any given 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 additional elements needed, 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 type-specific choice probabilities $\mathbb{P}_\tau(a_t|x_t, z_t, r, v)$. The mixture probabilities $\pi_\tau(r, v)$ denote the probability 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. As in the previous section under [Assumption D3](#) we distinguish 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) :

$$M_{r,v} = \sum_{\tau \in \mathcal{T}} \mathbb{I}\{\pi_\tau(r, v) > 0\}. \quad (10)$$

We first provide a lower bound for the total number of types that depends upon the size of the state-space. 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 unobserved 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 two restrictions on the distribution of states and actions:

ASSUMPTION U1 (Exclusion Restrictions).

1. *Conditional upon type, the type proxy r is uninformative about choice:*

$$\mathbb{P}_\tau(a_1, \mathbf{x}_1|r, v) = \mathbb{P}_\tau(a_1, \mathbf{x}_1|v) \quad \forall (a_1, \mathbf{x}_1, v), \text{ and } \mathbb{P}_\tau(a_t|\mathbf{x}_t, r, v) = \mathbb{P}_\tau(a_t|\mathbf{x}_t, v) \quad \forall (a_t, \mathbf{x}_t, v) \quad t > 1.$$

2. *Transition probabilities do not vary by type and are independent of r :*

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

Both parts of the assumption can be viewed as exclusion restrictions. The first part 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 conditional upon (r, v) and express it as a mixture of the corresponding type-specific joint 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 U1](#) and the Markov nature of the decision rule (see [Assumption B](#)) we can write

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

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

$$\begin{aligned} \mathbf{F}_{r,v}^{a_t, \mathbf{x}_t} &\equiv \mathbb{P}(a_t, \mathbf{x}_t | r, v), \\ \mathbf{F}_{\mathbf{x}_{t+1}, r, v}^{a_{t+1}} &\equiv \mathbb{P}(a_{t+1} | \mathbf{x}_{t+1}, r, v), \\ \mathbf{F}_{\mathbf{x}_t, \mathbf{x}_{t+1}, r, v}^{a_t, a_{t+1}} &\equiv \frac{\mathbb{P}(a_t, a_{t+1}, \mathbf{x}_t, \mathbf{x}_{t+1} | r, v)}{\mathbb{P}(\mathbf{x}_{t+1} | \mathbf{x}_t, a_t, v)}, \end{aligned} \quad (11)$$

and we will often suppress the dependence on v for brevity. Next, let \underline{M}^t (\underline{M} for short) denote the cardinality of the smaller of the state spaces in the two adjacent periods ($\min\{\#\mathbf{X}_t, \#\mathbf{X}_{t+1}\}$). For given values $(a_t, \mathbf{x}_t^1, \dots, \mathbf{x}_t^{\underline{M}}, a_{t+1}, \mathbf{x}_{t+1}^1, \dots, \mathbf{x}_{t+1}^{\underline{M}})$ define the $(\underline{M} + 1) \times (\underline{M} + 1)$ directly identified matrix using the expressions defined above in eq. (11):

$$\mathbf{P}_{r,v}^{a_t, a_{t+1}, \underline{M}} \equiv \begin{pmatrix} 1 & \mathbf{F}_{r,v}^{a_{t+1}, \mathbf{x}_{t+1}^1} & \dots & \mathbf{F}_{r,v}^{a_{t+1}, \mathbf{x}_{t+1}^{\underline{M}}} \\ \mathbf{F}_{r,v}^{a_t, \mathbf{x}_t^1} & \mathbf{F}_{\mathbf{x}_t^1, \mathbf{x}_{t+1}^1, r, v}^{a_t, a_{t+1}} & \dots & \mathbf{F}_{\mathbf{x}_t^1, \mathbf{x}_{t+1}^{\underline{M}}, r, v}^{a_t, a_{t+1}} \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{F}_{r,v}^{a_t, \mathbf{x}_t^{\underline{M}}} & \mathbf{F}_{\mathbf{x}_t^{\underline{M}}, \mathbf{x}_{t+1}^1, r, v}^{a_t, a_{t+1}} & \dots & \mathbf{F}_{\mathbf{x}_t^{\underline{M}}, \mathbf{x}_{t+1}^{\underline{M}}, r, v}^{a_t, a_{t+1}} \end{pmatrix}, \quad (12)$$

where we will sometimes abbreviate this matrix as $\mathbf{P}_r^{\underline{M}}$ or \mathbf{P} for brevity. Next, we express each element of this matrix in terms of the corresponding unknown type-specific probabilities

$$\begin{aligned} \mathbf{F}_{r,v}^{a_t, \mathbf{x}_t} &= \sum_{\tau} \pi_{\tau}(r, v) \mathbb{P}_{\tau}(a_t, \mathbf{x}_t | v) \equiv \sum_{\tau} \pi_{\tau}(r, v) \lambda_v^{a_t, \mathbf{x}_t, \tau}, \\ \mathbf{F}_{\mathbf{x}_{t+1}, r, v}^{a_{t+1}} &= \sum_{\tau} \pi_{\tau}(r, v) \mathbb{P}_{\tau}(a_{t+1} | \mathbf{x}_{t+1}, v) \equiv \sum_{\tau} \pi_{\tau}(r, v) \lambda_{\mathbf{x}_{t+1}, v}^{a_{t+1}, \tau}, \\ \mathbf{F}_{\mathbf{x}_t, \mathbf{x}_{t+1}, r, v}^{a_t, a_{t+1}} &= \sum_{\tau} \pi_{\tau}(r, v) \mathbb{P}_{\tau}(a_t, \mathbf{x}_t | v) \mathbb{P}_{\tau}(a_{t+1} | \mathbf{x}_{t+1}, v) \equiv \sum_{\tau} \pi_{\tau}(r, v) \lambda_v^{a_t, \mathbf{x}_t, \tau} \lambda_{\mathbf{x}_{t+1}, v}^{a_{t+1}, \tau}, \end{aligned} \quad (13)$$

where we have introduced $\lambda_{\cdot, \cdot}^{\cdot, \tau}$ as a short-hand notation for the type-specific choice probabilities. To express the relationship between the directly identified object \mathbf{P} and the type-specific choice probabilities, we define the two $M_{r,v} \times (\underline{M} + 1)$ matrices

$$\mathbf{L}_v^{a_t, \mathbf{x}_t, (\underline{M}+1)} \equiv \begin{pmatrix} 1 & \lambda_v^{a_t, \mathbf{x}_t^1, \tau_1} & \dots & \lambda_v^{a_t, \mathbf{x}_t^{\underline{M}}, \tau_1} \\ 1 & \lambda_v^{a_t, \mathbf{x}_t^1, \tau_2} & \dots & \lambda_v^{a_t, \mathbf{x}_t^{\underline{M}}, \tau_2} \\ \vdots & \vdots & \dots & \vdots \\ 1 & \lambda_v^{a_t, \mathbf{x}_t^1, \tau_{M_{r,v}}} & \dots & \lambda_v^{a_t, \mathbf{x}_t^{\underline{M}}, \tau_{M_{r,v}}} \end{pmatrix}, \quad (14)$$

and

$$\mathbf{L}_{\mathbf{x}_{t+1},v}^{a_{t+1},(\underline{M}+1)} \equiv \begin{pmatrix} 1 & \lambda_{\mathbf{x}_{t+1}^1,v}^{a_{t+1},\tau_1} & \dots & \lambda_{\mathbf{x}_{t+1}^{\underline{M}},v}^{a_{t+1},\tau_1} \\ 1 & \lambda_{\mathbf{x}_{t+1}^1,v}^{a_{t+1},\tau_2} & \dots & \lambda_{\mathbf{x}_{t+1}^{\underline{M}},v}^{a_{t+1},\tau_2} \\ \vdots & \vdots & \dots & \vdots \\ 1 & \lambda_{\mathbf{x}_{t+1}^1,v}^{a_{t+1},\tau_{M_{r,v}}} & \dots & \lambda_{\mathbf{x}_{t+1}^{\underline{M}},v}^{a_{t+1},\tau_{M_{r,v}}} \end{pmatrix}. \quad (15)$$

Next, define a diagonal matrix containing the type frequencies $\mathbf{V}_{r,v} \equiv \text{Diag}(\pi_{\tau_1}(r,v), \dots, \pi_{\tau_{M_{r,v}}}(r,v))$. With this notation in hand, we can now express the identified matrix \mathbf{P} in terms of the unknown objects of interest:

$$\mathbf{P}_{r,v}^{a_t,a_{t+1},\underline{M}} = (\mathbf{L}_v^{a_t,\mathbf{x}_t,(\underline{M}+1)})' \mathbf{V}_{r,v}^{M_{r,v}} \mathbf{L}_{\mathbf{x}_{t+1},v}^{a_{t+1},(\underline{M}+1)}. \quad (16)$$

We will use this relationship to determine the number of types (and subsequently identify the elements on the right hand side as well).

The dimension and rank of the right-hand side of eq. (16) is important because as we show below there is a relationship between them and the number of types. This is formalized as:

ASSUMPTION U2 (Existence and Rank Condition).

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

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

The first part of the assumption ensures that the elements of \mathbf{P} in eq. (12) are well defined and is, in principle, testable since it is placed on observed quantities. The second part requires the existence of at least as many points in the state space (\underline{M}) as the number of types ($M_{r,v}$) – and can be interpreted as an order condition. It further also imposes a rank condition – that there be sufficient variation in the type-specific choice probabilities across the state space (i.e. that the rows of the \mathbf{L} matrices are linearly independent) and at a given point in the state space (columns in the \mathbf{L} matrices cannot consist of identical entries). This assumption formalizes the intuition that type-specific choice probabilities are not identified if they do not vary sufficiently across types (i.e. across rows) so that the state space must be sufficiently rich to distinguish between them. Although untestable (since it involves unobserved quantities), this assumption is reasonable here to the extent that the model is only interesting – in the sense that types behave sufficiently differently – if it is true. With [Assumption U1](#) and [Assumption U2](#) in hand, we can now identify the total number of types.

PROPOSITION 1 (Identifying the Total Number of Types). Fix (r,v) and suppose that [Assumption U1](#) 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 $(a_t, \mathbf{x}_t^1, \dots, \mathbf{x}_t^{\underline{M}}, a_{t+1}, \mathbf{x}_{t+1}^1, \dots, \mathbf{x}_{t+1}^{\underline{M}})$, the total number of types $M_{r,v} \geq \text{Rank}(\mathbf{P}_{r,v}^{\underline{M}})$ where the directly identified matrix $\mathbf{P}_{r,v}^{\underline{M}}$ is defined in eq. (12).

2. Suppose in addition that [Assumption U2](#) holds. Then, $M_{r,v} = \text{Rank}(\mathbf{P}_{r,v}^{\underline{M}})$.

The first part of the proposition provides a lower bound on the number of types and is useful when the support of the state space is restricted (i.e. \underline{M} is relatively small). The second part shows that given a sufficiently rich state space and sufficient variation in type behavior across the state space, the lower bound is also an upper bound. [Proposition 1](#) 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 below 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 $\mathbf{P}_{r,v}^{\underline{M}}$ as well as the exclusion restriction in [Assumption U1](#). We prove two different sets of identification results: the first set of results uses variation in the type proxy r . These results are most useful in situations with limited state space transitions and only consider transitions between adjacent periods. The second set of results use additional information from the Markovian nature of the dynamic problem but requires a richer set of transitions beyond just adjacent periods.

In what follows we allow the type probability π_τ to depend upon the exogenous variables v . This is done in the interest of generality but nothing would be lost if we assumed (as one might for tractability reasons as we do in the empirical section) that $\pi_\tau(r, v) = \pi_\tau(r)$. 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 there is a common set of types that exist for at least two values of the type proxy, that is $\mathcal{T}_{r,v} = \mathcal{T}_{r',v}$ for $r' \neq r$. This requires *a priori* knowledge about the existence of types at different values of (r, v) and is therefore untestable. A simple and sufficient condition 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 across the state space in the sense outlined above. Formally, this translates into an invertibility condition on the matrices $\mathbf{L}_v^{a_t, \mathbf{x}_t, M_{r,v}}$ and $\mathbf{L}_{\mathbf{x}_{t+1}, v}^{a_{t+1}, M_{r,v}}$ defined using [eqs. \(14\)](#) and [\(15\)](#) but replacing \underline{M} with $M_{r,v} - 1$ in the definitions (so the dimensions now depend only upon the number of types $M_{r,v}$). For simplicity, we will abbreviate these two matrices as $\mathbf{L}_{t,r}$ and as $\mathbf{L}_{t+1,r}$. We collect both assumptions discussed above in the assumption below (since they are always invoked together).

ASSUMPTION U2' (Overlap Condition; Modified Existence and Rank Condition).

1. (Overlap) Fix (r, v) . There exists an $r' \neq r$ such that $\mathcal{T}_{r',v} = \mathcal{T}_{r,v}$.
2. 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 $(j, k) \in \{1, \dots, (M_{r,v} - 1)\}^2$.
 - (b) The $M_{r,v} \times M_{r,v}$ matrices $\mathbf{L}_v^{a_t, \mathbf{x}_t, M_{r,v}}$ and $\mathbf{L}_{\mathbf{x}_{t+1}, v}^{a_{t+1}, M_{r,v}}$ are invertible.

The second part above is a restatement of [Assumption U2](#) for the square matrix case (see the discussion following that assumption) and the first part is the overlap condition discussed immediately above. We can then state the following lemma for point identification of the type-specific choice probabilities:

LEMMA 3 (Identifying Type-Specific Choice Probabilities). *Fix (r, v) and suppose [Assumption U1](#) and [Assumption U2](#) hold and that the agent's optimal decision process is Markovian. 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 the overlap condition fails, identification results can follow from 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.3.1](#) with the main result being [Lemma A5](#), 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 not the identities of the specific types. We now outline a procedure to assign type identities to the identified type choice probabilities.

First, we apply [Lemma 1](#) to the identified type-specific choice probabilities $\mathbb{P}_\tau(a_t | \mathbf{x}_t, v)$ to identify period 2 and 3 utilities ($\{u_t(x_t, a_t; \tau)\}_{t \in \{2,3\}, \tau \in \mathcal{T}}$) and the products $\beta_\tau \delta_\tau$. Next, we use these objects to construct the functions $\tilde{h}_\tau^A(x_2, z_2, \beta_\tau \delta_\tau)$ and $\tilde{h}_\tau^B(x_2, z_2, \beta_\tau \delta_\tau)$ defined in eq. (8). Recall that the function $\tilde{h}_\tau^A(x_2, z_2, \mathbf{d}_2)$ represents the period 1 self's evaluation of (un-discounted) utility in period 2 (in state (x_2, z_2)) when actions in period 2 are taken assuming the discount factor between periods 2 and 3 is \mathbf{d}_2 . Similarly, $\tilde{h}_\tau^B(x_2, z_2, \mathbf{d}_2)$ is the period 1 self's expected period 3 (un-discounted) utility assuming that the discount factor between period 2 and 3 when choosing period 2 actions is given by \mathbf{d}_2 . Then, define the following terms for a pre-specified point (x_{20}, z_{20}) :

$$\tilde{h}_\tau^{\Delta, j}(x_2, z_2, \mathbf{d}_2) \equiv \tilde{h}_\tau^j(x_2, z_2, \mathbf{d}_2) - \tilde{h}_\tau^j(x_{20}, z_{20}, \mathbf{d}_2) \quad j \in \{A, B\}. \quad (17)$$

Functions with a Δ superscript can be interpreted as 'normalized' future expected utilities relative to a given point (x_{20}, z_{20}) . [Lemma A1](#) of [Appendix A.2](#) shows that the following function is identified:

$$h_\tau^\Delta(x_2, z_2) \equiv h_{\tau,2}(x_2, z_2) - h_{\tau,2}(x_{20}, z_{20}),$$

where $h_{\tau,2}(x_2, z_2)$ was defined in eq. (8). Next, define the identified function (which is well-defined

as long as the denominator is not zero):

$$\hat{\delta}_\tau(x_2, z_2) \equiv \frac{h_\tau^\Delta(x_2, z_2) - \tilde{h}_\tau^{\Delta, A}(x_2, z_2, \beta_\tau \delta_\tau)}{\tilde{h}_\tau^{\Delta, B}(x_2, z_2, \beta_\tau \delta_\tau)}. \quad (18)$$

The variation of this function across the state space will be key to distinguish naïve from consistent and sophisticated types. This will be true as long as, roughly speaking, two views of the future are sufficiently different across the observed state space. To state this formally, consider the function $\tilde{h}_\tau(x_2, z_2, d_1, d_2) \equiv \tilde{h}_\tau^A(x_2, z_2, d_2) + d_1 \tilde{h}_\tau^B(x_2, z_2, d_2)$ (defined in eq. (40) below). Normalizing it by subtracting $\tilde{h}_\tau(x_{20}, z_{20}, d_1, d_2)$ we can define the identified function $\tilde{h}_\tau^\Delta(x_2, z_2, d_1, d_2)$.

$$\tilde{h}_\tau^\Delta(x_2, z_2, d_1, d_2) \equiv \tilde{h}_\tau(x_2, z_2, d_1, d_2) - \tilde{h}_\tau(x_{20}, z_{20}, d_1, d_2). \quad (19)$$

The function $\tilde{h}_N^\Delta(x_2, z_2, \delta_N, \delta_N)$ is the (normalized) value of being in state (x_2, z_2) from the viewpoint of a naïve agent's period 1 self while $\tilde{h}_N^\Delta(x_2, z_2, \delta_N, \beta_N \delta_N)$ is the same value for the agent if instead she assumed that she would use $\beta_N \delta_N$ to discount utility between periods 2 and 3 when making decisions in period 2 (i.e. behaved as if she were completely sophisticated). We will require that these two views of the future have to vary over the state space (formally, this is [Assumption UA3](#) and is stated in [Appendix A.3.2](#) in the interest of brevity). It seems reasonable to assume that the naïve and sophisticated calculations differ over the future for if they did not, period 1 choice probabilities would be identical for naïve types and sophisticated types who have the same preference parameters as their naïve counterparts. We can then state the formal result (details and some intuition for the proof are relegated to the appendix):

PROPOSITION 2 (Assigning Type-Identities). *Suppose that the type-specific choice probabilities are identified and Assumptions B, D1 and D2 hold (so that [Lemma 1](#) holds). Further, suppose that [Assumption UA3](#) holds. Then, type identities are identified.*

3.2.4 Identifying Preferences for each Type

Note that most of the work in identifying preferences was already done in the previous sub-sections 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 A2](#) and [Lemma A3](#) to (set or point) identify the time-preference parameters.

3.3 Partial Sophistication

[Assumption D3](#) imposes that sophisticated agents are fully aware of their future present-bias, while naïve agents are fully unaware of it. This restriction greatly simplifies the analysis by reducing the number of time preference parameters, but it may be undesirable. 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. The arguments for why this is the case is similar to those made for naïve types analyzed above.

We start the analysis by relaxing [Assumption D3](#) and only require that $\tilde{\beta}_\tau \in [\beta_\tau, 1]$, so the only sharp distinction is between time-consistent ($\beta_\tau = 1$) and partially sophisticated agents. 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. As always, we assume that the total number of types is finite.

We first discuss the identification of types. Starting with a given type-specific choice probability we show that one can determine whether the type associated with the probability is consistent (or partially sophisticated). The reasoning is very similar to that employed in identifying whether a given type is naïve (formally worked out in [Lemma A6](#)). In particular, a type will be partially sophisticated if and only if the directly identified object $\hat{d}(x_2, z_2)$ (defined in eq. (18) above) varies over the state space. We collect the assumptions required to identify type-identity in the partially sophisticated case below.

ASSUMPTION U3 (Restriction for Partially Sophisticated Model).

1. *Agents are partially sophisticated (or equivalently partially naïve):* $\tilde{\beta}_\tau \in [\beta_\tau, 1]$.
2. *There exists a set $\mathcal{S} \subset \mathcal{X}_2 \times \mathcal{Z}_2$ with positive measure such that for all types τ ,* $\tilde{h}_\tau^{\Delta, B}(x_2, z_2, \beta_\tau \delta_\tau) \neq 0$.
3. *For types τ such that $\tilde{\beta}_\tau \neq \beta_\tau$,* $\text{Var} \left(\frac{\tilde{h}_\tau^{\Delta}(x_2, z_2, \delta_\tau, \tilde{\beta}_\tau \delta_\tau) - \tilde{h}_\tau^{\Delta}(x_2, z_2, \delta_\tau, \beta_\tau \delta_\tau)}{\tilde{h}_\tau^{\Delta, B}(x_2, z_2, \beta_\tau \delta_\tau)} \right) > 0$.

Recall that $\tilde{h}_\tau^{\Delta}(x_2, z_2, \delta_\tau, \tilde{\beta}_\tau \delta_\tau)$ is the period 1 self's (normalized) value of being in state (x_2, z_2) , defined in eq. (19) when they assume that decisions in period 2 will be made using $\tilde{\beta}_\tau \delta_\tau$ to discount period 3 utility back to period 2. The function $\tilde{h}_\tau^{\Delta}(x_2, z_2, \delta_\tau, \beta_\tau \delta_\tau)$ is the same value for the agent if instead they assume that decisions in period 2 will be made using $\beta_\tau \delta_\tau$ to discount period 3 utility back to period 2 (i.e. as if the agent 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 fully sophisticated types (who share the same remaining preference parameters) – so that it would not be possible to distinguish between them on the basis of the observed distributions. The argument behind the proof of the result below is very similar to that employed in [Proposition 2](#) and the proof is relegated to [Appendix A.4](#). As with [Proposition 2](#), a key ingredient is the function $h_\tau^{\Delta}(x_2, z_2)$ which is identified in [Lemma A1](#).

PROPOSITION 3 (Assigning Type Identities). *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 that the conditions for [Lemma A1](#) hold. Further, suppose that [Assumption U3](#) holds. 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 is that without further assumptions, the three parameters for

these agents (i.e., δ_τ, β_τ , and $\tilde{\beta}_\tau$) are not point-identified although if the exponential discount factor δ_τ is identified, then the remaining two parameters are also identified. 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 $\tilde{h}_\tau^\Delta(x_2, z_2, d_1, d_2)$. Recall that period 1 choices identify the function $h_\tau^\Delta(x_2, z_2)$ and we know that $h_\tau^\Delta(x_2, z_2) = \tilde{h}_\tau^\Delta(x_2, z_2, \delta_\tau, \tilde{\beta}_\tau \delta_\tau)$. This information, however, is not enough to identify the time-preference parameters since the identified function $\tilde{h}_\tau^\Delta(\cdot)$ 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 only allows for set identification.

PROPOSITION 4. *Suppose that the conditions for Proposition 3 hold. Then, the identified set for the parameters $(\beta_\tau, \tilde{\beta}_\tau, \delta_\tau)$ is given by*

$$\Theta_{\beta, \tilde{\beta}, \delta} = \left\{ (\mathbf{b}, \tilde{\mathbf{b}}, \mathbf{d}) \in (\beta_\tau \delta_\tau, 1]^2 \times [\beta_\tau \delta_\tau, 1] : \tilde{h}_\tau^\Delta(x_2, z_2, \mathbf{d}, \mathbf{b}\mathbf{d}) = h_\tau^\Delta(x_2, z_2) \quad \forall (x_2, z_2), \mathbf{b} = (\beta_\tau \delta_\tau)/\mathbf{d}, \mathbf{b} \leq \tilde{\mathbf{b}} \right\}.$$

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 Odisha, India.

We assume that agents have preferences as in eq. (1), 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 discussed in Section 3.3). We also assume that each type exists with positive probability.¹²

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. Given sample size concerns, we impose functional forms on the utility function so that the structural model reduces to one characterized by a finite vector-valued parameter; inference will therefore follow from standard asymptotic arguments.

Recall that the identification strategy requires two key variables: the type proxy (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 (commonly known as “Money Earlier or Later”

¹²This assumption is made for convenience particularly given our sample size, but one could in principle apply the results from Section 3.2.1 to first identify the total number of types before proceeding to the analysis undertaken below.

or MEL questions, see e.g. [Cohen et al., 2020](#)) and this information forms the type indicator r . Finally, the excluded variables z are elicited subjective beliefs about ITN and untreated net efficacy in preventing malaria.

We begin by providing some basic context for the study in [Section 4.1](#) and then in [Section 4.2](#) we set up the structural model and evaluate the performance of the estimator in a set of Monte Carlo simulations. In the interest of space and because the identification arguments are broadly similar to those in the general section (subject to a few key differences we outline below) we relegate details on identification to [Appendix B](#). [Section 4.3](#) presents the estimation results, [Section 4.4](#) provides a comparison of the time-preference parameter estimates to those in the literature and [Section 4.5](#) describes results from a set of counterfactual exercises.

4.1 Data

Our data were collected in the context of a randomized controlled trial (RCT) carried out in 2007–2009 in Orissa (now Odisha), the most malaria-endemic state in India ([Dhingra et al. 2010](#)). The study evaluated the impacts of alternative mechanisms of providing ITNs on the health and socio-economic outcomes of potential users, and was carried out in collaboration with a local partner, Bharat Integrated Social Welfare Agency (BISWA), a micro-lender with a large presence in Odisha, see [Tarozzi et al. \(2014\)](#) for details. We use data from a sample of 621 households in 47 villages 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 option (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 re-treatment for Rs. 15 (single) or Rs. 18 (double). With the second option (c henceforth), the household purchased the treated net plus 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 additional cash payment was required for re-treatment as the price of the chemical 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. This choice will, if anything, lead to an under-estimation of the extent of present-bias so that relaxing it would only further amplify the substantial present-bias we document below.

Of the initial 621 households, we exclude 32 that could not be re-contacted at endline, 13 that purchased bed nets 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.¹³

Respondents were aware of the role of mosquitoes in transmitting malaria, of the high cost of malaria episodes, and that bed nets 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 that the event would occur with certainty. We then estimated subjective probabilities by dividing the number by ten.¹⁴ Since most respondents were 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.¹⁵

The histograms in **Figure 1** show that about three quarters of respondents believed that without using nets one would certainly get malaria, or that regular use of ITNs would virtually rule out risk. About half of respondents reported 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. Note that these beliefs show that both bed nets and re-treatment with insecticide were recognized as very effective at reducing malaria risk. Coupled with the high monetary cost of malaria episodes, this is *prima facie* evidence that present bias may help explaining the low rates of ownership and especially re-treatment of bed

¹³Malaria infection and Hb were measured via rapid diagnostic tests that only required fingerprick blood specimens, and were 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).

¹⁴We did not measure ranges of probability, so we cannot identify the degree of uncertainty around the reports.

¹⁵Analogous beliefs were elicited about the protective power of bed nets and treatment for children and pregnant women. Responses were almost identical across demographic groups, and so we only use information for adults.

nets.

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 receiving Rs. 10 one month later and an equal or larger sum (Rs. 10, 12, 14 or 16) 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.¹⁶ Standard expected utility models imply that if a respondent prefers to receive, say, Rs. 10 a month from today to Rs. 16 paid four months from today, s/he should also prefer Rs. 10 paid four months in the future to Rs. 16 paid seven months in the future. We interpret preference "reversals"—whereby the former is true but the choice is reversed for the latter—to be correlated with a form of inconsistency in time preferences consistent with hyperbolic discounting.¹⁷ In [Table 2](#) we summarize the findings. As expected, in each set of four questions, the fraction of individuals choosing 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 and we denote agents with such reversals as having type signal $r = 1$ ($r = 0$ otherwise).¹⁸

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 only ITNs (contract b), while 146 opted for the "commitment" product whose price also included the cost of two re-treatments (contract c). 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, an overwhelming majority (92%) of the ITNs purchased with contract c were re-treated with insecticide. However, the fraction was only 36% for bed nets 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 . In [Appendix D](#) we provide additional descriptions of the association between actions (purchase and re-treatment) and a list of household characteristics.

Our data do not include beliefs about the *joint* distribution of income and malaria, but they

¹⁶Interviewers told respondents that one of the twelve chosen rewards, selected at random, would be paid by our micro-lender partner BISWA at the chosen time horizon. In practice, to avoid logistical difficulties, the selected reward was paid 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 rewards were to be paid at least one month later. This was done so that choices would not depend on issues of trust, although such issues were unlikely given that all sample households included at least one BISWA client.

¹⁷See [Andreoni and Sprenger \(2012\)](#) for an alternative view. [Rubinstein \(2003\)](#) shows cases in which preference reversals arise despite preferences that are neither consistent nor time-inconsistent hyperbolic. We focus on the identification of types with preferences that are compatible with time-consistency or hyperbolic discounting. Identifying different forms of time-inconsistent preferences, albeit important, is beyond the scope of our paper.

¹⁸We 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 them. 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 recall that we do not require a one-to-one mapping from signals into types, such that the model is robust against imperfect type signals.

do include beliefs about future income. These were recorded at baseline using an approach similar to [Guiso et al. \(2002\)](#), eliciting a plausible range for future income, followed by a simple question about the probability that income will be below the mid-point of the range. In the model we assume that gross income and malaria status are independent, although malaria is allowed to affect current utility by reducing consumption due to the monetary costs due to illness episodes. [Appendix E](#) includes details on the measurement of income, subjective transition probabilities for income, and malaria costs.

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 the standard analysis of dynamic models in what follows is the presence of time-inconsistent agents and the further complication that types are unobserved. This alters the standard results as we highlight below. Furthermore, instead of the minimum necessary three periods in our model, we use four periods because of the specific features of our intervention. For brevity we relegate the details of the variable construction to [Appendix E](#).

PREFERENCES (PERIOD 4): At $t = 4$, 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 has malaria and h ('healthy') otherwise, and y_4 measures income. For simplicity we discretize income so that y_t is a dichotomous variable that can take either a 'high' or a 'low' value, depending on whether household income is above or below the median. The time-invariant household characteristics v that enter preferences include household size at baseline (v_{hhs}), a measure of households assets (v_{assets}), an indicator of risk aversion (v_{risk}) and an indicator of untreated net ownership at baseline v_{oldnet} . The survey-based measure of attitudes towards risk is obtained by using an abbreviated version of the procedure proposed in [Holt and Laury \(2002\)](#). We specify

$$u_4(x_4, v; \tau) = \mathcal{C}(x_4) + \phi_\tau(v), \quad (20)$$

where $\mathcal{C}(x_4)$ is consumption in state x_4 , and ϕ captures other factors that can affect per-period utility. Consumption depends on both health and income and is calculated as $\mathcal{C}(x_t) = y_t - \mathbb{I}\{h_t = m\}\eta_m$, where η_m accounts for the monetary cost of malaria. We set this value equal to the median cost of a malaria episode, taking into account both expenses for doctor visits and treatment as well as any wages paid to labor hired to replace a sick worker. This choice is conservative in the sense that the use of alternative measures of malaria costs (such as the expected costs of a malaria episode elicited in our survey, or the inclusion of estimates of lost earnings due to illness) lead to greater estimated present bias.

The per period utility can vary along both observed (v) and unobserved (τ) dimensions:

$$\phi_\tau(v) = \phi_0 + \phi_1 \mathbb{I}_{\{\tau=\tau_S\}} + \phi_2 \mathbb{I}_{\{\tau=\tau_N\}} + \phi_3 v_{\text{hhs}} + \phi_4 v_{\text{assets}} + \phi_5 v_{\text{risk}} + \phi_6 v_{\text{oldnet}}. \quad (21)$$

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

$$u_t(x_t, a_t, v; \tau) = (\mathcal{C}(x_t) - p_r a_t \mathbb{I}\{a_1 = b\} - p_r \mathbb{I}\{a_1 = c\}) + a_t \phi_\tau(v); \quad (22)$$

where $a_t = 1$ if the net is re-treated in period t and $= 0$ otherwise, and p_r is the price of re-treatment. The price is paid in period t regardless of the choice to re-treat if the household purchased a commitment contract, while it is paid only if the household chose the baseline contract b in period 1 and the net is re-treated.¹⁹ The multiplication of $\phi(\cdot)$ by the action ensures that the ϕ coefficients are identified.

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

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

where p_b is the price of a baseline contract b .²⁰ Buyers choosing contract c pay a higher price that also includes the cost of re-treatment, but in the model we assume that such price is paid at the time of re-treatment, consistent with eq. (22).

We parameterize the mixture probabilities $\pi_\tau(r)$ using the parsimonious Hardy-Weinberg functional form (Hardy, 1908; Weinberg, 1908). In a first step, we estimate the following logit function ψ that depends on two mixture parameters γ_1 and γ_2 and the observable value of the MEL type-proxy r :

$$\psi(\gamma, r) = \frac{\exp(\gamma_1 + \gamma_2 r)}{1 + \exp(\gamma_1 + \gamma_2 r)}. \quad (23)$$

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

$$\begin{aligned} \pi_C(r) &= \psi^2(\gamma, r); \\ \pi_N(r) &= 2\psi(\gamma, r)(1 - \psi(\gamma, r)); \\ \pi_S(r) &= 1 - \pi_C(r) - \pi_N(r) = 1 - \psi^2(\gamma, r) - 2\psi(\gamma, r)(1 - \psi(\gamma, r)). \end{aligned}$$

¹⁹We do not have reliable data on the timing of debt repayments. We assume that debts are repaid within the first six months, except for the component attributable to re-treatment. The assumption is conservative in the sense that present-bias is exacerbated if we consider repayments in future periods. The results remain qualitatively similar in terms of estimated type probabilities and relative adoption probabilities when putting greater weight on later repayments. The modeling assumption that the re-treatment component of the commitment contract c (equal to about 13% of total price) is paid in later periods—as is optional re-treatment for households choosing b —is mostly due to technical considerations. Dropping this assumption does not qualitatively change the present-bias or 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 is violated. The results are available from the authors upon request. The worse fit arises because high present-bias (i.e., low values of β_N and β_S) makes c less attractive than b because of the higher initial payment.

²⁰Due to sample size considerations and for tractability, we do not model the decision to buy single *versus* double nets, nor the number of nets purchased. In the data, buyers purchased on average close to two nets and in the model we assume that demand, when positive, is always equal to 1.5 ITNs. For each contract type, we set the price equal to 1.5 times the average of prices for single and double nets, weighted by the respective purchase frequencies. Compared to assuming that each household buys one net, or two nets, the estimated type probabilities and discount factor are virtually unchanged. A higher number of nets slightly decreases the degree of present bias, because the higher cost decreases demand.

4.2.1 Solving the Model

Given the finite horizon model, we solve for the optimal decision rule using backward induction. We solve and estimate the model using the mapping between type-specific choice probabilities and type-specific value functions (defined below), possible because while we do not observe types, the type-specific choice probabilities are identified using the results in the previous sections. For clarity we will sometimes suppress the dependence of the functions below on v (which only enters the period utility functions).

PERIOD 3 CHOICE: At $t = 3$ the event that an agent (who has purchased an ITN) will retreat his net is

$$\begin{aligned}\{a_3 = 1\} &\equiv \left\{ 1 = \operatorname{argmax}_{a \in \{0,1\}} \left\{ u_3(x_3, a, \tau, v) + \epsilon_3(a) + \beta_\tau \delta \int u_4(x_4; \tau, v) dF(x_4|x_3, a; z) \right\} \right\} \\ &= \left\{ 1 = \operatorname{argmax}_{a \in \{0,1\}} \left\{ v_{\tau,3}(x_3, z, a, \beta_\tau \delta) + \epsilon_3(a) \right\} \right\},\end{aligned}$$

where

$$v_{\tau,3}(x_3, z, a, \beta_\tau \delta) = u_3(x_3, a; \tau) + \beta_\tau \delta \underbrace{\int u_4(x_4; \tau) dF(x_4|x_3, a; z)}_{q_\tau(x_3, z, a)}. \quad (24)$$

Here z reflects the household-specific vector of beliefs about malaria risk as a function of the chosen bed net usage as well as household-specific transition probabilities of income. We emphasize the dependence on the hyperbolic parameter β_τ since it will be useful in the subsequent analysis.

We assume that $(\epsilon_3(0), \epsilon_3(1))/\sigma$ are *i.i.d.* standard GEV random variables, and to ease notation we set $\sigma = 1$ in what follows. Under this assumption, the type-specific choice probability is given by

$$\mathbb{P}_\tau(a_3 = 1|x_3; z) = \frac{\exp(v_{\tau,3}(x_3, z, 1, \beta_\tau \delta))}{\sum_{j=0}^1 \exp(v_{\tau,3}(x_3, z, j, \beta_\tau \delta))}. \quad (25)$$

PERIOD 2 CHOICE: Under the GEV assumption on the errors, the choice probabilities are

$$\mathbb{P}_\tau(a_2 = 1|x_2; z) = \frac{\exp(v_{\tau,2}(x_2, z, 1, \beta_\tau \delta))}{\sum_{j=0}^1 \exp(v_{\tau,2}(x_2, z, j, \beta_\tau \delta))}. \quad (26)$$

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

$$v_{\tau,2}(x_2, z, a, \beta_\tau \delta) \equiv u_2(x_2, a; \tau) + \beta_\tau \delta \int v_\tau^*(s_3, z) dF(s_3|x_2, a, z), \quad (27)$$

where v_τ^* represents the optimized utility (from the point of view of the period 2 self) in state $s_3 = (x_3, \epsilon_3)$. To simplify notation and make clear the dependence of the $v_\tau^*(\cdot)$ function on the

beliefs about the future hyperbolic parameter ($\tilde{\beta}$), we define the choice indicator (as in eq. (6))

$$A_\tau(s_3, k, \tilde{\beta}_\tau \delta) \equiv \mathbb{I} \left\{ v_{\tau,3}(x_3, z, k, \tilde{\beta}_\tau \delta) + \epsilon_3(k) > v_{\tau,3}(x_3, z, s, \tilde{\beta}_\tau \delta) + \epsilon_3(s) \quad \forall s \neq k \right\},$$

which is an indicator for the event that action k is optimal in state s_3 given a type- τ agent's expected future present-bias of $\tilde{\beta}_\tau$. To ease exposition, we will sometimes shorten $A_\tau(s_3, k, \tilde{\beta}_\tau \delta)$ to $A_{\tau,k}$. With this notation, we can re-write

$$v_\tau^*(s_3, z) \equiv \sum_{k \in \{0,1\}} (v_{\tau,3}(x_3, z, k, \delta) + \epsilon_3(k)) A_\tau(s_3, k, \tilde{\beta}_\tau \delta).$$

Substituting this expression into eq. (27) makes clear that from the perspective of period 2 (a) period 3 utility is discounted back to period 2 using $\beta_\tau \delta$ (b) the period 2 self believes that his period 3 self will discount utility between periods 3 and 4 (captured by $A_{\tau,k}$) by the factor $\tilde{\beta}_\tau \delta$. For time-consistent agents $\tilde{\beta}_C = \beta_C = 1$ while for naïve agents $\tilde{\beta}_N = 1 \neq \beta_N$ and for fully sophisticated agents $\tilde{\beta}_S = \beta_S$.

In this section, we can simplify the expression in eq. (27) further. First,

$$\begin{aligned} h_\tau(x_3, z) &\equiv \int v_\tau^*(x_3, \epsilon_3, z) dF(\epsilon_3) \\ &= \sum_{k \in \{0,1\}} \left\{ \int A_\tau(x_3, \epsilon_3, k, \tilde{\beta}_\tau \delta) [u_3(x_3, k, \tau) + \epsilon_3(k)] dF(\epsilon_3) + \delta q_\tau(x_3, z, k) \int A_k^\tau(x_3, \epsilon_3, \tilde{\beta}_\tau \delta) dF(\epsilon_3) \right\} \end{aligned}$$

which is analogous to the expression derived in eq. (8). Next, using the GEV distribution of ϵ_3 :

$$h_\tau(x_3, z) = \sum_{k \in \{0,1\}} \mathbb{P}(A_{\tau,k}) [v_{\tau,3}(x_3, z, k, \delta) - v_{\tau,3}(x_3, z, k, \tilde{\beta}_\tau \delta)] + \gamma_{\text{euler}} + \log \left(\sum_{j=0}^1 \exp(v_{\tau,3}(x_3, z, j, \tilde{\beta}_\tau \delta)) \right), \quad (28)$$

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

$$\mathbb{P}(A_{\tau,k}) = \frac{\exp(v_{\tau,3}(x_3, z, k, \tilde{\beta}_\tau \delta))}{\sum_{j=0}^1 \exp(v_{\tau,3}(x_3, z, j, \tilde{\beta}_\tau \delta))}.$$

The term in square parentheses in eq. (28) captures the key differences between the three types of agent in the dynamic programming problem. It can be viewed as the adjustment made by the period 2 self to account for the *perceived* future present-bias of the period 3 self. For consistent agents, no such adjustment is needed, $\tilde{\beta}_C \delta = \delta$ so this term is zero and the expression reduces to the standard one in dynamic choice problems (see e.g. eq. 12 in [Aguirregabiria and Mira, 2010](#)). For naïve agents this term is also zero ($\tilde{\beta}_N \delta = \delta$) since such agents (incorrectly) do not perceive their period 3 self to be present-biased, consequently they do not adjust their period 2 value function to account for future present-bias.²¹ Finally, this term is not equal to zero for sophisticated types since

²¹Note, however, that the problem for naïve agents does not reduce to the standard problem since in period 3 such agents will use a different discount rate ($\beta_\tau \delta$) than the one they anticipated (that is, δ).

they are aware of their period 3 self's future bias and adjust their period 2 decisions accordingly.²²

PERIOD 1 CHOICE: The argument is similar to the one above with the only substantive difference that there are now three possible actions and the choice probabilities for an agent of type τ are given by

$$\mathbb{P}_\tau(a_1 = a|x_1; z, v) = \frac{\exp(v_{\tau,1}(x_1, z, a, \beta_\tau \delta))}{\sum_{j \in \{n, b, c\}} \exp(v_{\tau,1}(x_1, z, j, \beta_\tau \delta))}, \quad (29)$$

where b represents the purchase of a baseline contract, c the purchase of a commitment contract, and n no contract purchase (on occasion we will use 0 to indicate no purchase, 1 to indicate the standard contract and 2 to indicate the commitment contract). For period 1, the $v_{\tau,1}(\cdot)$ function is

$$v_{\tau,1}(x_1, z, k, \beta_\tau \delta) \equiv u_1(x_1, k; \tau) + \beta_\tau \delta \int h_\tau(x_2, z) dF(x_2|x_1, k, z), \quad (30)$$

As in the discussion of period 2,

$$h_\tau(x_2, z) = \sum_{k \in \{n, b, c\}} \mathbb{P}(A_{\tau, k}) \left[v_{\tau,2}(x_2, z, k, \delta) - v_{\tau,2}(x_2, z, k, \tilde{\beta}_\tau \delta) \right] + \gamma_{\text{euler}} + \log \left(\sum_{j \in \{n, b, c\}} \exp(v_{\tau,2}(x_2, z, j, \tilde{\beta}_\tau \delta)) \right).$$

where

$$\mathbb{P}(A_{\tau, k}) = \frac{\exp(v_{\tau,2}(x_2, z, k, \tilde{\beta}_\tau \delta))}{\sum_{j \in \{n, b, c\}} \exp(v_{\tau,2}(x_2, z, j, \tilde{\beta}_\tau \delta))}.$$

In the estimation we also account for the presence of untreated bed nets owned prior to the intervention.²³ This is because owning an untreated bed net in period 1 affects the utility in case of not purchasing an ITN as well as perceived malaria risk, which affects expected utility through the transition probabilities.

4.2.2 Identification

The identification arguments are broadly similar to those in the general section on identification but there are some differences that we highlight here: (a) In the application the excluded variable z_t is time-invariant and 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 $\mathbb{P}_\tau(a_t|\mathbf{x}_t, v)$, 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 $\mathbb{P}(\mathbf{x}_{t+1}|\mathbf{x}_t, v)$. Since we do not exploit any time-series variation in z for identification, this deviation does not impose any new difficulties. (b) The type proxies vary depending upon the model. For the observed types model, the type indicator is (r, a_1) once agents have made first-period choices, while it only includes r prior to that (i.e. at $t = 1$). In the unobserved types case the only proxy is r . (c) For tractability we assume that all three types exist in the population so we do not need to first identify the total

²²To connect this with the argument in Section 3.1.1 note $v_{\tau,3}(x_3, z, a, \delta) - v_{\tau,3}(x_3, z, a, \tilde{\beta}\delta) = \delta(1 - \tilde{\beta}_\tau)q_\tau(x_3, z, k)$ which is linear in the unknown parameter δ for sophisticated types when $\beta_S\delta$ is identified.

²³This 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.

number of types. (d) Finally, our setting requires an additional period ($t = 4$) to rationalize choices in period 3, given that such choices involve an expectation. This additional period (when no action is taken) adds a complication since terminal period utilities are not identified by the standard arguments as they were in [Section 3](#). In the interest of space, we relegate the formal discussion on identification accounting for these differences to [Appendix B](#). Here we only discuss briefly the key additional variables required for identification and the assumptions imposed on them.

As a type proxy r we use responses to standard MEL experiments that have become quite common in field surveys. Recall that r needs to satisfy two key (untestable) assumptions (see [Assumption U1](#)). First, conditional upon type the proxy must be uninformative about choice. Second, transition probabilities (assumed independent of type) do not depend upon the proxy either. In the empirical application we add a third condition, that is, that the proxy needs to be sufficiently informative about types that researchers can order specific likelihood ratios that we define below and formalize in [Assumption UE1](#) in the appendix.

The second key variables for identification are subjective beliefs about the likelihood of contracting malaria conditional on the choice of sleeping regularly under a bed net (either an ITN or an untreated net) or of not using a net. In addition, we also use household-specific transition probabilities for income, constructed from expectations on future income measured at baseline, see [Appendix E](#) for details. We assume that these beliefs are time-invariant and that they do not enter directly the per-period payoff function and only affect the forward-looking component of the value function, as outlined in the [Basic Assumptions](#).

4.2.3 Monte Carlo simulations

In online [Appendix H](#) we illustrate the properties of our model with a set of Monte Carlo simulations. For the observed types case, [Table H.1](#) shows that for moderate sample sizes (300 and above) both the mean and the median estimated time preference parameters are close to their true values. [Table H.2](#) shows 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 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 (results available upon request).

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 (we relegate the derivation of the objective function to online [Appendix F](#)). We parameterize and estimate the type probabilities conditional on the type signal r ($\mathbb{P}(\tau|r) \equiv \pi_\tau(r)$) and then use Bayes' rule to compute the probabilities conditional on both the type signal and first period choice ($\mathbb{P}(\tau|r, a_1) \equiv \pi_\tau(a_1, r)$).²⁴ The former estimate the distribution of types in the population conditional on an observed type predictor, an object that has not been previously estimated in the literature, and that is of direct interest. The latter estimate the type distribution conditional on both the type predictor and the purchase decision. A comparison of the two provides a measure of attractiveness of commitment products to sophisticated types (recall that agent types are unobserved and purchase decisions are not assumed to uniquely reveal type).

In our main specification, we estimate three time preference parameters, i.e. the discount factor δ (assumed to be constant across types), and two type-specific present-bias parameters β_N and β_S . As specified in [eq. \(21\)](#) we include type indicators, household size, assets, a measure of risk aversion and an indicator for old net ownership in the per-period utility function. The type-specific dummies allow take-up and re-treatment 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 that maps r directly into 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 $\left\{ \frac{\pi_C(r)}{\pi_C(r')}, \frac{\pi_N(r)}{\pi_N(r')}, \frac{\pi_S(r)}{\pi_S(r')} \right\}$ 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 $\frac{\pi_\tau(r=0)}{\pi_\tau(r=1)}$, $\tau \in \{C, N, S\}$, the sufficiency criterion for type identification would be met if $\frac{\pi_C(r=0)}{\pi_C(r=1)} > \frac{\pi_N(r=0)}{\pi_N(r=1)} > \frac{\pi_S(r=0)}{\pi_S(r=1)}$, which is not unreasonable in our context.

We begin by discussing the estimated population type probabilities presented in the top panel of [Table 4](#). We estimate that 36% of agents are time consistent and that the majority of the time-inconsistent agents (about half of the total population or 75% of the population of time-inconsistent agents) are naïve. The fraction of time-consistent agents is about the same regardless of the value of r (col. 3) and indeed the same is true for the two inconsistent types. The estimates satisfy the monotonicity condition described above for identification, a result that does not depend on constraints imposed during the estimation.²⁵

²⁴We compute this as $\hat{\mathbb{P}}(\tau|a_1, r; \hat{\theta}) = \frac{\hat{\mathbb{P}}(a_1, \tau|r)}{\hat{\mathbb{P}}(a_1|r)} = \frac{\hat{\mathbb{P}}(a_1|r)\hat{\mathbb{P}}(\tau|r)}{\hat{\mathbb{P}}(a_1|r)}$, where the last equality follows from the exclusion restriction that the signal r is not informative on actions conditional on type, see [Assumption U1](#).

²⁵In particular, $\frac{\hat{\pi}_C(r=0)}{\hat{\pi}_C(r=1)} = \frac{.36}{.35} > \frac{\hat{\pi}_N(r=0)}{\hat{\pi}_N(r=1)} = \frac{.480}{.483} > \frac{\hat{\pi}_S(r=0)}{\hat{\pi}_S(r=1)} = \frac{.160}{.167}$.

We next examine the informativeness of contract choice by computing type probabilities conditional on first period choice and r (col. 5). These two indicators do not perfectly predict type and, in fact, all three types exist for *every* value of the indicators. In contrast, recall that the directly observed types model assumes that $\pi_C(0, \cdot) = 1$, $\pi_N(1, b) = 1$, and $\pi_S(1, c) = 1$. Perhaps most strikingly, the results indicate that conditional on r , the probability of being sophisticated does not change substantively with the purchase of the commitment product.²⁶ This finding is consistent with recent work (see e.g. [Carrera et al., 2022](#)) that also finds commitment products to be of limited use in predicting time-inconsistency.

Overall, across all values of r and conditional upon any net purchase, time-consistent agents account for 67% of all purchases while comprising about a little over one-third of the total population. Inconsistent naïve (sophisticated) agents account for 25% (8%) of total purchases while accounting for 48%(16%) of the total population.

In Panel B of [Table 4](#) we present results for alternative specifications. Our preferred results (in col. 1) indicate that the exponential discount factor is ≈ 1 , so that for the time horizons relevant for our study time-consistent households do not significantly discount future utility. However, the two time-inconsistent types dramatically discount the future relative to the present, with $\hat{\beta}_N = 0.06$ and $\hat{\beta}_S = 0.16$. Thus, the high present-bias of a large part of the population can rationalize the low adoption of ITNs despite the much higher expected cost of malaria when not using an ITN. Col. 2 shows that all the estimated time preference parameters are considerably lower than in col 1 when we remove household characteristics and type intercepts. This exacerbation of the present-bias problem and increased impatience (relative to col. 1) suggest that ignoring heterogeneity in the per-period utility function results in even greater present-bias than the considerable amount already present in col 1. When we impose $\beta_N = \beta_S$ (col. 3), the results remain quantitatively similar with the common β lying between the two estimated parameters in col. 1.

In col. 4 we assume that types are observed (based on a deterministic mapping from (a_1, r) to types), while in col. 5 there is a single time-consistent type for the whole population. The results are now quite different: in the known type case, the estimated discount factor $\hat{\delta}$ is 0.08, and in the single type case it is even lower at 0.01. Furthermore, the present-bias parameters in col. 4 are much higher than in our preferred model being almost indistinguishable from 1. To better understand these results, recall that for about 3/4 of respondents $r = 0$, and in the known types case these are assumed to be time-consistent. Thus, in col. 4 (as well as 5, where all agents are assumed time-consistent) the vast majority of agents are time-consistent by construction. It is thus reasonable that under these scenarios the discount factor δ must be low enough to rationalize the overall low ITN adoption rate given the high expected costs of malaria and the high perceived protective power of bed nets.

As a further robustness check, we estimate a version of our model in which we allow the discount factor to differ for the time-consistent type (δ_C) and for the time-inconsistent types, although we assume it remains common to both these latter types (δ_{NS}). Recall that our arguments allow

²⁶In fact, demand for commitment is higher for N than for S types, i.e. $\mathbb{P}(a_1 = c | \tau_N, \hat{\theta}) > \mathbb{P}(a_1 = c | \tau_S; \hat{\theta})$.

point-identification of all the time-preference parameters, and it nests models with 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 similar to our preferred model, with $\hat{\delta}_C = \hat{\delta}_{NS} = .99$, $\hat{\beta}_N = .10$ and $\hat{\beta}_S = .05$.²⁷

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” simulations where we estimate a misspecified model. We examine two forms of misspecification: (a) there are three types in the population but we impose only one consistent type; (b) the population comprises only one time-consistent type but the researcher assumes the existence of consistent, naïve, and sophisticated agents. The results are presented in [Table H.3](#) in online [Appendix H](#). Under scenario (a), the common $\hat{\delta}$ is considerably less than the true value (.99) similar to when we impose one time-consistent type. Under scenario (b), $\hat{\delta}$ is somewhat higher than the true value while both the present-bias parameters are strictly less than 1 (albeit imprecisely estimated). This suggests that if our primary model had been misspecified 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 find 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 model 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. appeal

4.4 Comparisons of Estimated Time-Preference Parameters

There is growing evidence of substantial heterogeneity in preferences, including discounting, both across and within countries (see for instance [Falk et al., 2018](#)). Most estimates come from high-income countries, although evidence from low-income countries is also growing. To facilitate comparisons we discuss geometric discount rates δ using a six-month horizon (consistent with our empirical application). In contrast, the parameter β multiplies utility in any future period (with no difference in how far in the future the payoff is as long as it is in the future), and so estimates from different studies should be directly comparable. It is important to note that most of the estimates arise in models with just one type of agent.

[Balakrishnan et al. \(2020\)](#) uses a lab experiment in Kenya, assuming a single agent type, and estimate β and δ using inter-temporal choices in a convex time budget experiment (as in [Andreoni and Sprenger, 2012](#)). They estimate β in the 0.90-0.92 range, and a high degree of impatience in δ , in the 0.21-0.48 range. Using data from lab experiments in rich countries with inter-temporal choices

²⁷While in principle one can also test for the number of types in the population as we have outlined above, we do not pursue this here because of our relatively small sample size.

on effort, rather than money, Augenblick et al. (2015) and Augenblick and Rabin (2019) estimate β in the 0.83-0.90 range on average. Carrera et al. (2022) use a model of partial sophistication and offer contracts for gym attendance with a commitment component in a US city and estimate $\beta = 0.55$. A similar value (0.67) was found in Chaloupka et al. (2019) who study partial sophistication and demand for commitment products for smoking cessation. Using a job search model with hyperbolic discounting, Paserman (2008) estimates $\beta = 0.40$ and $\delta = 0.998$ among low-wage US workers. In a context closer to ours, Bai et al. (2021) study demand for commitment contracts for health care to prevent hypertension in rural Punjab, India. They estimate $\beta = 0.365$ on average; however, in their model agents are partially naïve and on average their perceived β ($\tilde{\beta}$) is more than twice as large. Their estimates for δ range from 0.234 to 0.780 although they note that while δ is technically identified, in practice it is not robustly estimated.

Overall, our estimates suggest a relatively large geometric discounting factor (δ), while we find a high degree of present bias (i.e. small estimates of β) relative to the literature.

4.5 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 additional costs from sickness associated to low purchase and re-treatment rates of ITNs due to present-bias.

Changing re-treatment Prices: We first discuss the consequences of doubling the price of re-treatment, balanced by a corresponding increase in the price of the commitment contract.²⁸ Intuitively, the price change has several effects. First, the increase will reduce contemporaneous demand for re-treatment 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 re-treatment 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 re-treatment price from Rs. 16.5 to 33 per bed net, re-treatment rates under contract b decline by 1% (see Table 5, panel A). We find no effect on re-treatment under contract c , which commits buyers to re-treatment. Demand for contract c declines by .8% while demand for contract b does not change substantially. This suggests a substitution from buying c to making no purchases. We

²⁸Counterfactuals 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 (2020).

further examine changes in take-up and re-treatment when the price of re-treatment is halved (to Rs 8.25 per bed net). First, we find that re-treatment rates for buyers of the standard contract b increase by .6%. Second, demand for the commitment contract increases by .39%.

Quantifying the effect of time-inconsistency on price responses: Next, we gauge to what extent price responses are a function of time-preferences. To this end, in Panel B of [Table 5](#) we re-evaluate the impact of price changes with a model where all three types are time-consistent with the common value of $\hat{\delta}$ estimated in our preferred specification and no type-specific intercepts for inconsistent types. In general we find greater responsiveness to price changes relative to Panel A. Doubling re-treatment price p_r reduces overall purchases in period 1 by 4% for a consistent population, compared to a decline of .4% in the model with inconsistent agents. Similarly, halving p_r increases purchases by 2% for a consistent population, compared to an increase of .2% in the model with inconsistent agents.

Quantifying the effect of time-inconsistency on health and health costs: Present-bias reduces the present value of purchasing an ITN and thus reduces demand. This in turn increases the probability that present-biased agents contract malaria relative to the probability for otherwise identical but time-consistent agents. A natural next step is thus to conduct 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 to online [Appendix G](#). First, we compute purchase and re-treatment probabilities using the parameters from our preferred model but assuming away time-inconsistency for all agents as above. Next, we use these probabilities to compute the expected costs of malaria for each agent.²⁹ While the latter is clearly an extrapolation, it provides an alternative measure of the efficacy of ITNs relative to our survey measures. We then compare these expected costs to the actual ones for each agent (i.e. using the estimated parameters from our preferred model) starting with $t=2$ (i.e. the first period in which period 1 actions affect health) and summing across periods without discounting.

[Table 6](#) presents the results from using each measure. In both scenarios, we find that present-bias substantially increases expected costs from malaria. The median cost associated with present-bias is Rs. 488 (using the numbers from the meta-analysis in [Lengeler, 2009](#)) or Rs. 812 (using elicited beliefs on net efficacy). Overall, present-bias leads to a median reduction of 3-5 workdays per malaria episode. Even though these costs are high relative to that of an ITN, the estimated present-bias is such that investing in ITNs is not an attractive option for the median present-biased agent (naïve or sophisticated) relative to a time-consistent one (cols 5 and 6). This provides concrete

²⁹The key ingredients in the expected cost calculation are (1) the probability of contracting malaria when sleeping under an ITN relative to an untreated net or no net, and (2) 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.”

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 prefer to encourage ITN adoption to reduce long-run health costs and increase productivity. However, time-inconsistent households do not find ITN purchases particularly attractive. The results, combined with those for the price elasticities, also suggest 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. Monte Carlo simulations suggest that both the time-preference parameters and the population type probabilities can 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 about two thirds of the population, with almost half of the population being naïve time-inconsistent. While the standard exponential time-preference parameter is close to 1, time-inconsistent types are substantially present-biased, with estimated present-bias parameters of 0.06 (for naïve types) and 0.16 (for sophisticated types).³⁰ 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, 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

³⁰In contrast, Bisin and Hyndman (2020) find that present-bias is more pronounced among sophisticated individuals relative to naïve ones in an experiment among U.S. students, while the hyperbolic discount factor among naïve individuals is close to one.

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 that is correlated with time-inconsistency but uninformative about choice conditional on type. In our case, we use inter-temporal choices that are fairly commonly included in household surveys. 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 of streaming programs). 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 degree of present-bias is large enough (in terms of both the present-bias parameters and the prevalence of time-inconsistency 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).

References

Abbring, J. H. and Ø. Daljord (2020a). A Comment on Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting by Hanming Fang and Yang Wang. *International Economic Review* 61(2), 565–571. [5]

Abbring, J. H. and Ø. Daljord (2020b). Identifying The Discount Factor In Dynamic Discrete Choice Models. *Quantitative Economics* 11(2), 471–501. [4, 9]

Abbring, J. H., Ø. Daljord, and F. Iskhakov (2019). Identifying Present-Biased Discount Functions in Dynamic Discrete Choice Models. Working Paper. [4, 5, 9]

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. [38]

Aguirregabiria, V. and P. Mira (2010). Dynamic discrete choice structural models: A survey. *Journal of Econometrics* 156(1), 38–67. [4, 6, 32]

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutstrom (2008). Eliciting risk and time preferences. *Econometrica* 76(3), 583–618. [OA-3]

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. [6]

Andreoni, J. and C. Sprenger (2012). Estimating Time Preferences from Convex Budgets. *American Economic Review* 102(7), 3333–56. [2, 28, 37]

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 (2020). Identifying Dynamic Discrete Choice Models Off Short Panels. *Journal of Econometrics* 215(2), 473–485. [38]

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, OA-3]

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, 38]

Augenblick, N. and M. Rabin (2019). An Experiment on Time Preference and Misprediction in Unpleasant Tasks. *Review of Economic Studies* 86, 941–975. [38]

Bai, L., B. Handel, E. Miguel, and G. Rao (2021). Self-Control and Demand for Preventive Health: Evidence from Hypertension in India. *The Review of Economics and Statistics* 103(5), 1–22. [5, 38]

Balakrishnan, U., J. Haushofer, and P. Jakiela (2020). How soon is now? Evidence of present bias from convex time budget experiments. *Experimental Economics* 23(2), 294–321. [37]

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), 1–9. [2]

Banerjee, A. and S. Mullainathan (2010). The Shape Of Temptation: Implications For The Economic Lives Of The Poor. NBER Working Paper w15973. [2]

Bernheim, B. D., D. Ray, and S. Yeltekin (2015). Poverty and Self-Control. *Econometrica* 83(5), 1877–1911. [2]

Bisin, A. and K. Hyndman (2020, January). Present-bias, procrastination and deadlines in a field experiment. *Games and Economic Behavior* 119, 339–357. [40]

Brunnermeier, M. K. and J. A. Parker (2005). Optimal Expectations. *The American Economic Review* 95, 1092–1118. A Structural Model Of "Optimal" Belief Distortions Due To Anticipatory Utility. [9]

Bryan, G., D. Karlan, and S. Nelson (2010). Commitment Devices. *Annual Reviews of Economics* 2, 671–698. [4]

Carrera, M., H. Royer, M. Stehr, J. Sydnor, and D. Taubinsky (2022). Who Chooses Commitment? Evidence and Welfare Implications. *Review of Economic Studies* 89(3), 1205–1244. [36, 38, OA-3]

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]

Chaloupka, Frank J. I., M. R. Levy, and J. S. White (2019). Estimating biases in smoking cessation: Evidence from a field experiment. Working Paper 26522, National Bureau of Economic Research. [38]

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. [41]

Cohen, J., K. M. Ericson, D. Laibson, and J. M. White (2020, June). Measuring Time Preferences. *Journal of Economic Literature* 58(2), 299–347. [26, OA-2]

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. [40]

Delavande, A. (2014). Probabilistic Expectations In Developing Countries. *Annual Reviews of Economics* 6, 1–20. [9, 40]

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. [9, 40]

DellaVigna, S. (2009). Psychology And Economics: Evidence From The Field. *Journal of Economic Literature* 47(2), 315–372. [2]

DellaVigna, S. (2018). Chapter 7 - structural behavioral economics. In B. D. Bernheim, S. DellaVigna, and D. Laibson (Eds.), *Handbook of Behavioral Economics - Foundations and Applications* 1, Volume 1 of *Handbook of Behavioral Economics: Applications and Foundations* 1, pp. 613 – 723. North-Holland. [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. [26]

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 (2019). Intertemporal choice. In B. D. Bernheim, S. DellaVigna, and D. Laibson (Eds.), *Handbook of Behavioral Economics - Foundations and Applications* 2, Volume 2, Chapter 1, pp. 1–67. North-Holland. [2]

Falk, A., A. Becker, T. Dohmen, B. Enke, D. Huffman, and U. Sunde (2018). Global Evidence on Economic Preferences. *Quarterly Journal of Economics* 133(4), 1645–1692. [37]

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]

Fuster, A., R. Perez-Truglia, M. Wiederholt, and B. Zafar (2022). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *The Review of Economics and Statistics* 104(5), 1059–1078. [9]

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]

Guiso, L., T. Jappelli, and L. Pistaferri (2002). An Empirical Analysis Of Earnings And Employment Risk. *Journal of Business & Economic Statistics* 20(2), 241–253. [29, OA-17]

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. *Science* 28(706), 49–50. [30]

Heidhues, P. and P. Strack (2021). Identifying present-bias from the timing of choices. *American Economic Review* 111(8), 2594–2622. [5]

Holt, C. A. and S. K. Laury (2002). Risk aversion and incentive effects. *American Economic Review* 92(5), 1644–1655. [29, OA-19]

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, 11, 52, OA-13]

Kasahara, H. and K. Shimotsu (2009). Nonparametric identification of finite mixture models of dynamic discrete choices. *Econometrica* 77(1), 135–175. [5, 60, 61, OA-7]

Kaur, S., M. Kremer, and S. Mullainathan (2014). Self-control at work. *Journal of Political Economy* 123(6), 1227–1277. [2]

Kőszegi, B. (2010). Utility From Anticipation And Personal Equilibrium. *Economic Theory* 44(3), 415–444. [9]

Kremer, M. and E. Miguel (2007). The illusion of sustainability. *Quarterly Journal of Economics* 122(3), 1007–1065. [41]

Laibson, D. I. (1997). Golden Eggs And Hyperbolic Discounting. *Quarterly Journal of Economics* 112(2), 443–477. [2]

Lengeler, C. (2009). Insecticide-treated bed nets and curtains for preventing malaria. Issue 2. Art. No.: CD000363. DOI: 10.1002/14651858.CD000363.pub2. [4, 39, 51, OA-21]

Magnac, T. and D. Thesmar (2002). Identifying dynamic discrete decision processes. *Econometrica* 70(2), 801–816. [3, 5, 9, 12]

Manski, C. F. (2004). Measuring expectations. *Econometrica* 72(5), 1329–1375. [9, 40]

Martinez, S.-K., S. Meiry, and C. Sprenger (2021). Procrastination in the Field: Evidence from Tax Filing. Working Paper. [5]

Mullainathan, S. (2004). Development economics through the lens of psychology. World Bank Working Paper, Report No. 28974. [2]

Norets, A. and X. Tang (2013). Semiparametric Inference In Dynamic Binary Choice Models. *Review of Economic Studies* 81(3), 1229–1262. [38]

O'Donoghue, T. and M. Rabin (1999). Doing it now or later. *American Economic Review* 89(1), 103–124. [7]

Paserman, M. (2008). Job search and hyperbolic discounting: Structural estimation and policy evaluation. *The Economic Journal* 118(531), 1418–1452. [4, 38]

Roth, C. and J. Wohlfart (2020). How Do Expectations About The Macroeconomy Affect Personal Expectations And Behavior? *Review of Economics and Statistics* 102(4), 731–748. [9]

Rubinstein, A. (2003). “Economics and psychology”? The case of hyperbolic discounting. *International Economic Review* 44(4), 1207–1216. [28]

Rust, J. (1994). Structural estimation of Markov decision processes. In R. Engle and D. McFadden (Eds.), *Handbook of Econometrics*, pp. 3081–3143. Elsevier Science. [3, 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. [9]

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, bed-nets and malaria: Evidence from a randomized controlled trial in Orissa (India). *American Economic Review* 104(7), 1909–1941. [2, 26, 27, OA-16]

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, OA-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]

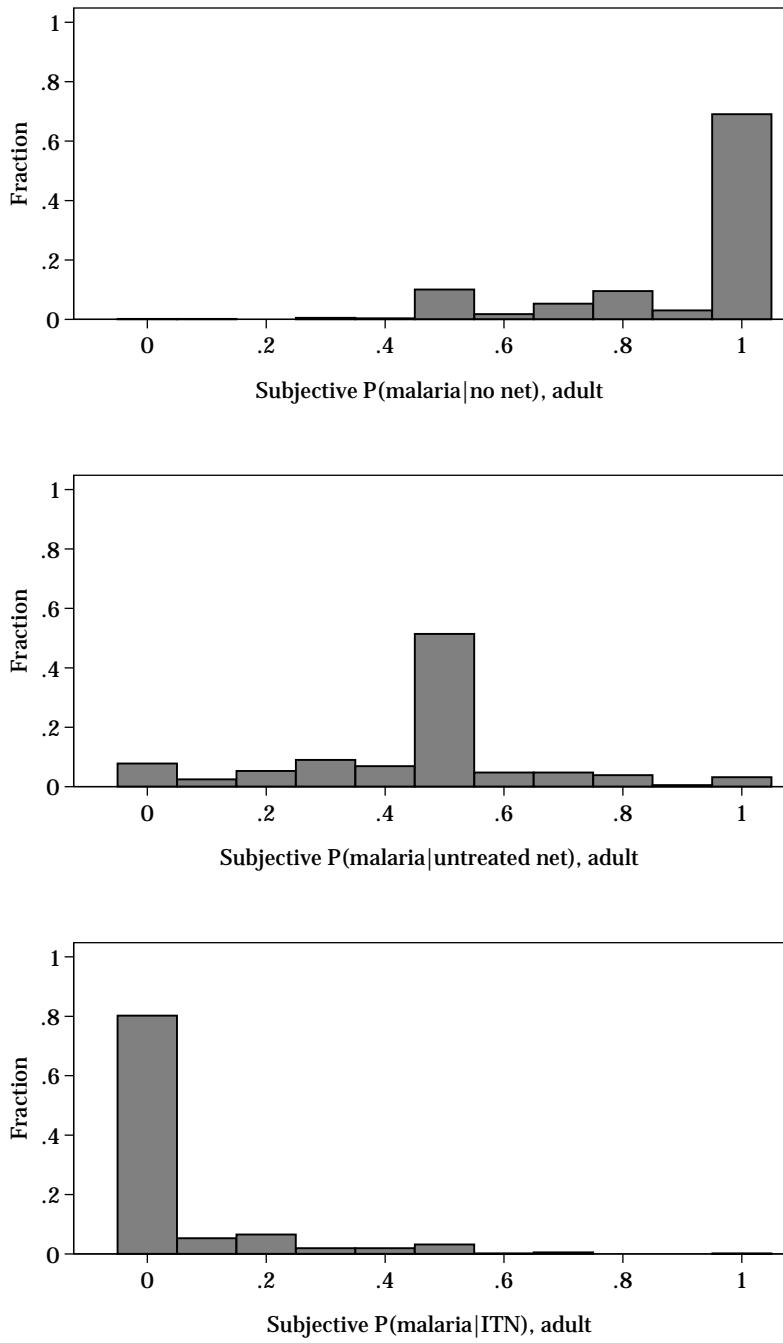
van der Klaauw, W. (2012). On the use of expectations data in estimating structural dynamic choice models. *Journal of Labor Economics* 30(3), 521 – 554. [5]

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. [30]

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. [46]

Figure 1: Perceived Protective Power of Bednets



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, $n = 566$.

Table 1: Baseline Summary Statistics

	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 ($Hb < 11g/dl$)	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

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.

Table 3: Summary of purchases

	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

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

Table 4: Structural Estimates

		(1)	(2)	(3)	(4)	(5)
Panel A	π_C	0.357 (0.342-0.372)	$\pi_C(0)$	0.360 (0.346-0.374)	$\pi_C(0, b)$	0.675 (0.662-0.688)
					$\pi_C(0, c)$	0.664 (0.639-0.688)
			$\pi_C(1)$	0.350 (0.325-0.376)	$\pi_C(1, b)$	0.674 (0.661-0.687)
Type Probabilities	π_S	0.162 (0.152-0.172)	$\pi_S(0)$	0.160 (0.151-0.169)	$\pi_S(0, b)$	0.081 (0.075-0.087)
					$\pi_S(0, c)$	0.086 (0.074-0.098)
			$\pi_S(1)$	0.167 (0.149-0.184)	$\pi_S(1, b)$	0.081 (0.075-0.087)
Panel B	π_N	0.481 (0.476-0.486)	$\pi_N(0)$	0.480 (0.475-0.485)	$\pi_N(0, b)$	0.244 (0.237-0.251)
					$\pi_N(0, c)$	0.250 (0.238-0.263)
			$\pi_N(1)$	0.483 (0.475-0.491)	$\pi_N(1, b)$	0.245 (0.238-0.252)
Preference Parameters	δ	0.9989 (0.0438)	No ϕ	$\beta_N = \beta_S$	Known types	One type
	β_N	0.0580 (0.0613)	0.0323 (0.2456)	0.9989 (0.0461)	0.0822 (0.0237)	0.0100 (0.3549)
	β_S	0.1550 (0.0708)	0.0100 (0.1581)		0.9988 (0.0311)	
Parameters	ϕ_0	1.1132 (0.0054)	0.1080 (0.0087)	1.1128 (0.0044)	0.0230 (0.0055)	0.1127 (0.0064)
	ϕ_{Naive}	-1.0800 (0.0028)		-1.0912 (0.0062)	-0.8782 (0.0070)	
	ϕ_{Soph}	-1.1200 (0.0027)		-1.0887 (0.0047)	1.9948 (0.0028)	
Parameters	ϕ_{HHS}	-0.1354 (0.0055)		-0.1354 (0.0023)	-0.1251 (0.0047)	-0.0472 (0.0241)
	ϕ_{Assets}	0.1066 (0.0019)		0.1066 (0.0044)	0.0475 (0.0029)	0.0631 (0.0529)
	ϕ_{Risk}	-0.1959 (0.0061)		-0.1959 (0.0067)	-0.1069 (0.0022)	-0.0516 (0.0139)
Parameters	ϕ_{OldNet}	0.1666 (0.0023)		0.1668 (0.0056)	0.1086 (0.0035)	-0.0140 (0.0797)
	1e6×Log-Likelihood	-1651716.1540	-1745755.9478	-1651717.0368	-1597132.8423	-1741713.9648

6†

Notes: In Panel A, $\pi_\tau(r)$ are type probabilities conditional on signal r , parameterized using eq. (23), $\pi_\tau(r, a)$ also condition on contract choice a_1 (see [Footnote 24](#)), and π_τ are unconditional type probabilities averaged using the empirical distribution of r (2.5th and 97.5th percentiles of the distributions computed using the delta method in parentheses). All type probabilities are computed using the preferred model with three unobserved types and preferences described in [Section 4.2](#). Col. (1) of Panel B shows estimated preference parameters for our preferred model with three unobserved types and utility function specified as in [Section 4.2](#). In Col. 2 we impose that per-period utility parameters $\phi_v = 0$. In Col. 3 we impose $\beta_S = \beta_N$. In Col. 4 types are uniquely identified by (r, a_1) , with $\tau = \tau_C$ if $r = 0$, $\tau = \tau_S$ if $(r = 1, a_1 = c)$, and $\tau = \tau_N$ if $(r = 1, a_1 \in \{n, b\})$. In Col. 5 we assume $\tau = \tau_C$ for all agents. In cols. 4-5 types are known so the likelihoods are not mixtures and are thus not directly comparable to those in cols. 1-3. The identification for all estimated models follows from the arguments in [Section 4.2.2](#).

Table 5: Counterfactual Choices with Changes in re-treatment Price

	Double p_r		Half p_r	
Panel A: Preferred model				
% Change No Purchase (n)	0.398	(0.047)	-0.194	(0.024)
% Change Take up standard contract (b)	-0.015	(0.020)	0.004	(0.010)
% Change Take up commitment product (c)	-0.795	(0.078)	0.389	(0.040)
% Change Re-treatment (b)	-1.153	(0.021)	0.579	(0.010)
Panel B: Model with no time-inconsistency				
% Change No Purchase (n)	4.013	(0.107)	-1.978	(0.024)
% Change Take up standard contract (b)	-0.020	(0.015)	0.001	(0.010)
% Change Take up commitment product (c)	-2.101	(0.033)	1.045	(0.040)
% Change Re-treatment (b)	-1.006	(0.009)	0.499	(0.010)

Notes: All changes are relative to the re-treatment price (p_r) of Rs.16.5 per bed net. All figures are averages over the empirical distribution of demographics, beliefs and types. Standard errors in parentheses estimated using the delta method. We use estimates from our preferred model (col. 1 in Panel B of [Table 4](#)).

Table 6: Median cost of malaria and days missed attributable to time-inconsistent preferences

		(1) $t = 2$	(2) $t = 3$	(3) $t = 4$	(4) Total	(5) $\text{Total} \times \beta_N$	(6) $\text{Total} \times \beta_S$
(A) Monetary costs	Elicited beliefs	307 (5.0)	307 (5.0)	198 (2.7)	812 (12.8)	47 (49.2)	126 (56.3)
	Lengeler (2004)	215 (3.2)	215 (3.2)	58 (2.1)	488 (8.3)	28 (29.6)	76 (33.9)
(B) Missed days	Elicited beliefs	1.79 (0.03)	1.79 (0.03)	1.15 (0.03)	4.72 (0.08)	0.28 (0.29)	0.73 (0.33)
	Lengeler (2004)	1.30 (0.03)	1.30 (0.03)	0.33 (0.01)	2.93 (0.06)	0.17 (0.18)	0.45 (0.20)

Notes: Panel (A) shows the additional costs of malaria attributable to the lower investment into ITNs and re-treatment due to present-bias, using either the survey-elicited beliefs about malaria risk, or estimates on the protective power of ITNs from the meta-analysis in [Lengeler \(2009\)](#). We report per-period median changes in expected costs (cols. 1-3), as well as their sum, either raw (col. 4) or discounted by the estimated present-bias parameters (cols 5-6). Panel (B) presents the same statistics for the median expected days missed at work. Standard errors in parentheses are estimated using the delta method.

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

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

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 B](#)) 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 x_2 and z_2 is given by

$$\mathbb{P}_\tau(a_2^* = k|x_2, z_2) = \mathbb{P}\left(k = \operatorname{argmax}_{a \in \mathcal{A}_2} \left\{u_2(x_2, a; \tau) + \epsilon_2(a) + \beta_\tau \delta_\tau \int v_{\tau,3}^*(s_3) dF(s_3|x_2, a, z_2)\right\} \middle| x_2, z_2\right).$$

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

$$g_{\tau,2,k}(x_2, z_2) = u_2(x_2, k, \tau) - u_2(x_2, 0, \tau) + \beta_\tau \delta_\tau \int v_{\tau,3}^*(s_3) dF_{\Delta,k}(s_3|x_2, z_2) \quad (31)$$

for all (x_2, z_2, k) . Next, [Assumption D2](#) (the [Rank Condition](#)) allows us to express (for at least one action k and two points (z'_2, z''_2) and all $x_2 \in \mathcal{X}_2$)

$$\beta_\tau \delta_\tau = \frac{g_{\tau,2,k}(x_2, z'_2) - g_{\tau,2,k}(x_2, z''_2)}{\int v_{\tau,3}^*(s_3) (dF_{\Delta,k}(s_3|x_2, z'_2) - dF_{\Delta,k}(s_3|x_2, z''_2))} \quad (32)$$

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

Next, substituting eq. (32) into eq. (31), (b) using the fact that $\int v_{\tau,3}^*(s_3) dF(s_3|x_2, k, z_2)$ is identified and invoking the normalization in [Assumption B](#) we conclude that the period 2 utility function $u_2(x_2, k; \tau)$ is identified for all $k \in \mathcal{A}_2$ for all types.

Proof of Lemma 2

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

$$\mathbb{P}_\tau(a_1^* = k|x_1, z_1) = \mathbb{P}\left(k = \operatorname{argmax}_{a \in \mathcal{A}_1} \left\{u_1(x_1, a; \tau) + \epsilon_1(a) + \beta_\tau \delta_\tau \int v_{\tau,2}^*(s_2) dF(s_2|x_1, a, z_1)\right\} \middle| x_1, z_1\right).$$

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

$$g_{\tau,1,k}(x_1, z_1) = u_1(x_1, k; \tau) - u_1(x_1, 0; \tau) + \beta_\tau \delta_\tau \int v_{\tau,2}^*(s_2) dF_{\Delta,k}(s_2 | x_1, z_1) \quad (33)$$

where

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

and $v_{\tau,2}(\cdot), A_\tau(\cdot)$ are defined in eq. (6). 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 the indicator $A_\tau(s_2, a, \tilde{\beta}_\tau \delta_\tau)$ being equal to one. 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 eq. (33) is identified so that period 1 preferences are then identified.

Next, we show identification for sophisticated agents. We begin by first isolating the last expression in eq. (33). Under [Assumption D2](#) we can identify the difference

$$\begin{aligned} g_{\tau,1,k,\Delta}(x_1) &\equiv g_{\tau,1,k}(x_1, z'_1) - g_{\tau,1,k}(x_1, z''_1) \\ &= \beta_\tau \delta_\tau \left(\int v_{\tau,2}^*(s_2) \left(dF_{\Delta,k}(s_2 | x_1, z'_1) - dF_{\Delta,k}(s_2 | x_1, z''_1) \right) \right) \\ &= \beta_\tau \delta_\tau \left(\int v_{\tau,2}^*(s_2) dF(\epsilon_2) \left(dF_{\Delta,k}(x_2, z_2 | x_1, z'_1) - dF_{\Delta,k}(x_2, z_2 | x_1, z''_1) \right) \right) \\ &= \beta_\tau \delta_\tau \left(\int h_{\tau,2}(x_2, z_2) \left(dF_{\Delta,k}(x_2, z_2 | x_1, z'_1) - dF_{\Delta,k}(x_2, z_2 | x_1, z''_1) \right) \right) \\ &= \beta_\tau \delta_\tau \left(\int \tilde{h}_\tau^A(x_2, z_2, \tilde{\beta}_\tau \delta_\tau) + \delta_\tau \cdot \tilde{h}_\tau^B(x_2, z_2, \tilde{\beta}_\tau \delta_\tau) \left(dF_{\Delta,k}(x_2, z_2 | x_1, z'_1) - dF_{\Delta,k}(x_2, z_2 | x_1, z''_1) \right) \right). \quad (34) \end{aligned}$$

Thus, for sophisticated types

$$g_{S,1,k,\Delta}(x_1) = \beta_S \delta_S \left(\int \tilde{h}_S^A(x_2, z_2, \beta_S \delta_S) + \delta_S \cdot \tilde{h}_S^B(x_2, z_2, \beta_S \delta_S) \left(dF_{\Delta,k}(x_2, z_2 | x_1, z'_1) - dF_{\Delta,k}(x_2, z_2 | x_1, z''_1) \right) \right),$$

where the $\tilde{h}_S^A(\cdot)$ functions are defined in eq. (8) and are identified. We can then directly identify δ_S as

$$\delta_S = \frac{g_{S,1,k,\Delta}(x_1) - \beta_S \delta_S \int \tilde{h}_S^A(x_2, z_2, \beta_S \delta_S) \left(dF_{\Delta,k}(x_2, z_2 | x_1, z'_1) - dF_{\Delta,k}(x_2, z_2 | x_1, z''_1) \right)}{\beta_S \delta_S \int \tilde{h}_S^B(x_2, z_2, \beta_S \delta_S) \left(dF_{\Delta,k}(x_2, z_2 | x_1, z'_1) - dF_{\Delta,k}(x_2, z_2 | x_1, z''_1) \right)}$$

where the integral in the denominator is guaranteed to be non-zero by the rank condition ([Assumption D2](#)) and $\beta_S \delta_S$ and all the remaining objects on the right hand side are also identified by [Lemma 1](#). \blacksquare

A.2 Identification Results for Naïve Types

To obtain informative partial identification results we place stronger assumptions on the transition probabilities that allow us to point identify the function $h_{\tau,2}(x_2, z_2)$ – defined in eq. (8) – up to a normalization. Since these assumptions and the function $h_{\tau,2}(\cdot)$ 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.

A.2.1 Strengthening Variation in Transitions

We place stronger assumptions on the transition probabilities and this allows us to identify $h_{\tau,2}(\cdot)$. 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 $\mathcal{X}_s \times \mathcal{Z}_s$ and define the probabilities, all of which are identified:

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

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

$$dF(k, z_1) = \begin{bmatrix} dF_k(x_{2,1}, z_{2,1} | x_{1,1}, z_1) & \dots & dF_k(x_{2,S-1}, z_{2,S-1} | x_{1,1}, z_1) \\ \vdots & \vdots & \vdots \\ dF_k(x_{2,1}, z_{2,1} | x_{1,S-1}, z_1) & \dots & dF_k(x_{2,S-1}, z_{2,S-1} | x_{1,S-1}, z_1) \end{bmatrix}. \quad (35)$$

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 . We require that the vector of period 2 transition probabilities display sufficient variation as the period 1 state varies. In particular, define the matrix

$$dF_\Delta(k, z_1) \equiv dF(k, z_1) - dF(0, z_1) \quad (36)$$

for the action pair $(k, 0)$. Then the formal statement is

ASSUMPTION DA1 (Invertibility). \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 $(dF_\Delta(k, z'_1) - dF_\Delta(k, z''_1))$ is invertible for some action k and two points $z'_1 \neq z''_1$.

$dF_\Delta(\cdot)$ is the difference in the transition probabilities when action k is taken in period 1 relative to a base action. [Assumption DA1](#) can be interpreted as requiring that the excluded variable z_1 induces sufficient variability in these probability differences. If we interpret the z_1 as beliefs about x_2 , then we can interpret the assumption as stating that these beliefs must induce sufficient variability in the transition probabilities (see the empirical application section for a discussion on this). [Assumption DA1](#) is restrictive in that the “order” condition – that the number of possible states in period 1 to be at least as large as the cardinality of $\mathcal{X}_2 \times \mathcal{Z}_2$ – may be quite onerous. This assumption would fail if, for instance, the support of the x state variable in period 2 is larger than its support 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 instead placing restrictions on the support of $\mathcal{A}_1 \times \mathcal{X}_1$ which are often more palatable and leads to order conditions that are more likely to be satisfied in practice. This approach is presented below in [Lemma A4](#) and is useful when x_1 has limited support. [Assumption DA1](#) is related to the rank conditions in [Assumption D2](#) in that it imposes sufficient variation in the transition probabilities arising from the variation in z_1 . Further, it requires that this variation is sufficiently independent as x_1 varies – i.e. the vectors of the differences of transition probabilities (indexed by points in \mathcal{X}_1) must be linearly independent in the sense specified above. With this additional assumption we can recover the function $h_{\tau,2}(\cdot)$ up to a location shift.

We state the result in slightly more generality that it is needed for this section. In particular, we will show identification of $h_{\tau,2}(\cdot)$ for all types τ and regardless of whether the particular identity of τ is known (i.e. without knowing whether the type is consistent or partially sophisticated). This is because this result will be useful in the subsequent section when we examine unobserved types.

LEMMA A1 (Identification of $h_{\tau,2}(\cdot)$ and First Period Payoffs). Consider an agent of type τ maximizing equation eq. (1). Suppose that period 2 and 3 utilities $\{u_t(\cdot; \tau)\}_{t \in \{2,3\}, \tau \in \mathcal{T}}$ and the product $\{\beta_\tau \delta_\tau\}_{\tau \in \mathcal{T}}$

are identified (though type identities are not necessarily known).³¹ In the directly observed types cases (Section 3.1) this follows from Assumptions *B*, *D1* and *D2*. In the unobserved types case, this follows from applying Lemma 1 (for which assumptions *B*, *D1* and *D2* are needed) once the type-specific choice probabilities are identified. Suppose that Assumption DA1 holds. Then,

1. The function $h_{\tau,2}(x_2, z_2)$ is identified up to an additive constant \mathbf{k}_τ for all types τ and $(x_2, z_2) \in \mathcal{X}_2 \times \mathcal{Z}_2$. We denote this identified function as $h_\tau^\Delta(x_2, z_2) = h_{\tau,2}(x_2, z_2) - h_{\tau,2}(x_{20}, z_{20})$
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. We begin by first isolating the last expression in eq. (34):

$$\begin{aligned} g_{\tau,1,k,\Delta}(x_1) &= \beta_\tau \delta_\tau \left(\int v_{\tau,2}^*(s_2) \left(dF_{\Delta,k}(s_2|x_1, z'_1) - dF_{\Delta,k}(s_2|x_1, z''_1) \right) \right) \\ &= \beta_\tau \delta_\tau \left(\int h_{\tau,2}(x_2, z_2) \left(dF_{\Delta,k}(x_2, z_2|x_1, z'_1) - dF_{\Delta,k}(x_2, z_2|x_1, z''_1) \right) \right) \end{aligned}$$

Rewriting the integral as a summation and using the fact that $\beta_\tau \delta_\tau$ is identified we can identify

$$\begin{aligned} \frac{g_{\tau,1,k,\Delta}(x_1)}{\beta_\tau \delta_\tau} &= \sum_{(x_2, z_2) \in \mathcal{X}_2 \times \mathcal{Z}_2} h_{\tau,2}(x_2, z_2) \left(dF_{\Delta,k}(x_2, z_2|x_1, z'_1) - dF_{\Delta,k}(x_2, z_2|x_1, z''_1) \right) \\ &= \sum_{(x_2, z_2) \in (\mathcal{X}_2 \times \mathcal{Z}_2)^-} (h_{\tau,2}(x_2, z_2) - h_{\tau,2}(x_{20}, z_{20})) \left(dF_{\Delta,k}(x_2, z_2|x_1, z'_1) - dF_{\Delta,k}(x_2, z_2|x_1, z''_1) \right). \end{aligned} \quad (37)$$

where $(\mathcal{X}_2 \times \mathcal{Z}_2)^- \equiv (\mathcal{X}_2 \times \mathcal{Z}_2) \setminus (x_{20}, z_{20})$ which has cardinality $S - 1$. We add $-h_\tau(x_{20}, z_{20})$ where (x_{20}, z_{20}) is a fixed point in $\mathcal{X}_2 \times \mathcal{Z}_2$ to incorporate the constraint

$$\sum_{(x_2, z_2) \in \mathcal{X}_2 \times \mathcal{Z}_2} dF_k(x_2, z_2|x_1, z_1) = 1.$$

Without incorporating this restriction the matrix needed in Assumption DA1 below would not be invertible. Next, define the two $(S-1)$ column vectors $\mathbf{g}_{\tau,\Delta}(k)$ and \mathbf{h}_τ (where we suppress dependence on state variables):

$$\mathbf{g}_{\tau,\Delta}(k) \equiv \frac{1}{\beta_\tau \delta_\tau} \begin{bmatrix} g_{\tau,1,k,\Delta}(x_{1,1}) \\ \vdots \\ g_{\tau,1,k,\Delta}(x_{1,S-1}) \end{bmatrix}; \mathbf{h}_\tau \equiv \begin{bmatrix} h_{\tau,2}(x_{2,1}, z_{2,1}) - h_{\tau,2}(x_{20}, z_{20}) \\ \vdots \\ h_{\tau,2}(x_{2,S-1}, z_{2,S-1}) - h_{\tau,2}(x_{20}, z_{20}) \end{bmatrix}. \quad (38)$$

Using the notation above we can rewrite eq. (37) (for each candidate value $x_{1,1} \dots x_{1,S-1}$) in matrix form as

$$\mathbf{g}_{\tau,\Delta}(k) = (dF_\Delta(k, z'_1) - dF_\Delta(k, z''_1)) \mathbf{h}_\tau,$$

where $\mathbf{g}_{\tau,\Delta}(k)$ is identified by eq. (37) and the matrices $dF_\Delta(k, z'_1)$ and $dF_\Delta(k, z''_1)$ (defined in eq. (36)) are identified since they are constructed only from observed variables. Under the (testable) invertibility assumption (Assumption DA1) it follows that

$$h_\tau = \left(dF_\Delta(k, z'_1) - dF_\Delta(k, z''_1) \right)^{-1} \mathbf{g}_{\tau,\Delta}(k), \quad (39)$$

so that $h_{\tau,2}(x_2, z_2) - h_{\tau,2}(x_{20}, z_{20})$ is identified. To simplify notation in the statement of the lemma we have

³¹That is, it is not known if a given type τ corresponds a consistent, naïve or completely sophisticated type or more generally whether τ is a consistent or partially sophisticated type.

defined

$$k_\tau \equiv -h_{\tau,2}(x_{20}, z_{20})$$

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

$$g_{\tau,1,k}(x_1, z_1) = u_1(x_1, k; \tau) - u_1(x_1, 0; \tau) + \beta_\tau \delta_\tau \int (h_{\tau,2}(x_2, z_2) - h_{\tau,2}(x_{20}, z_{20})) dF_{\Delta,k}(x_2, z_2 | x_1, z_1),$$

so that first period preferences are identified for all types τ as

$$u_1(x_1, k; \tau) - u_1(x_1, 0; \tau) = g_{\tau,1,k}(x_1, z_1) - \beta_\tau \delta_\tau \int (h_{\tau,2}(x_2, z_2) - h_{\tau,2}(x_{20}, z_{20})) dF_{\Delta,k}(x_2, z_2 | x_1, z_1),$$

Note that $\sum_{x_2, z_2} h_{\tau,2}(x_2, z_2) dF_{\Delta,k}(x_2, z_2 | x_1, z_1) = 0$ since $dF_{\Delta,k}$ is a signed measure with total measure equal to zero. \blacksquare

A.2.2 Application of Lemma A1 for Naïve Types and Partial Identification Results

Next, recalling the definition of $\tilde{h}_\tau^j(x_2, z_2, d)$ for $j \in \{A, B\}$ in eq. (8), define the function.

$$\tilde{h}_\tau(x_2, z_2, d_1, d_2) \equiv \tilde{h}_\tau^A(x_2, z_2, d_2) + d_1 \tilde{h}_\tau^B(x_2, z_2, d_2) \quad (40)$$

The functions $\tilde{h}_\tau^j(\cdot)$ are identified upto d_2 and therefore the function \tilde{h}_τ is identified upto (d_1, d_2) .³² Next, using the definition of $h_{\tau,2}(x_2, z_2)$ in eq. (8) we see that

$$h_{\tau,2}(x_2, z_2) = \tilde{h}_\tau(x_2, z_2, \delta_\tau, \tilde{\beta}_\tau \delta_\tau).$$

So that

$$\underbrace{h_{\tau,2}(x_2, z_2) - h_{\tau,2}(x_{20}, z_{20})}_{\equiv h_\tau^\Delta(x_2, z_2)} = \underbrace{\tilde{h}_\tau(x_2, z_2, \delta_\tau, \tilde{\beta}_\tau \delta_\tau) - \tilde{h}_\tau(x_{20}, z_{20}, \delta_\tau, \tilde{\beta}_\tau \delta_\tau)}_{\equiv \tilde{h}_\tau^\Delta(x_2, z_2, \delta_\tau, \tilde{\beta}_\tau \delta_\tau)}. \quad (41)$$

Lemma A1 identifies the object $h_\tau^\Delta(x_2, z_2)$ for all types $\tau \in \mathcal{T}$ and the right hand side of eq. (41) is known upto the time preference parameters $(\delta_\tau, \tilde{\beta}_\tau \delta_\tau)$.

Applying this argument to naïve types, notice that the right hand side of eq. (41) is $\tilde{h}_{\tau_N}^\Delta(x_2, z_2, \delta_N, \delta_N)$ which is known upto δ_N . If the function $\tilde{h}_{\tau_N}^\Delta(x_2, z_2, d, d)$ were one-to-one in d , then one could recover δ_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_N as all those values of d that are consistent with the identified function $h_{\tau_N}^\Delta(x_2, z_2)$ and the identified object $\beta_N \delta_N$. To simplify exposition define

$$\tilde{h}_{\tau_N}^\Delta(x_2, z_2, d) \equiv \tilde{h}_{\tau_N}^\Delta(x_2, z_2, d, d) \quad (42)$$

Then, define the identified set for δ_N as

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

This leads to a corresponding identified set for β_N :

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

We state this result formally below

³²To see this, notice that the $\tilde{h}_\tau^j(\cdot)$ are functions of the period two and period three payoff functions which are both identified by previous arguments and the distribution of ϵ_t which is assumed known.

LEMMA A2 (Identified set for δ_N and β_N). Consider a naïve agent solving the problem (1) at $t = 1$ and suppose that the model satisfies Assumptions B, D1, D2, D3 and DA1. Then, the identified set for δ_N is given by Θ_{δ_N} and the identified set for β_N is given by Θ_{β_N} .

Proof. The argument here is straight-forward. First consider any $\mathbf{d} \in \Theta_{\delta_N}$ and the corresponding $\mathbf{b} = (\beta_{\tau_N} \delta_{\tau_N}) / \mathbf{d}$. Then, we have that $\beta_{\tau_N} \delta_{\tau_N} = \mathbf{b} \mathbf{d}$ and since \mathbf{d} lies in the identified set for δ_N , $h_{\tau_N}^{\Delta}(x_2, z_2) = \tilde{h}_{\tau_N}^{\Delta}(x_2, z_2, \mathbf{d})$. 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 observed distribution as the original set of parameters. Moreover, this shows that the bounds are sharp. \blacksquare

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

ASSUMPTION DA2. There exists $(x_2, z_2) \in \mathcal{X}_2 \times \mathcal{Z}_2$ such that the following difference is non-zero and has the same sign $\forall \mathbf{d} \in \Theta_{\delta_N}$:

$$\sum_{a \in \mathcal{A}_2} \left(q_{\tau_N}(x_2, z_2, a) \int A_{\tau_N}(s_2, a, \mathbf{d}) dF(\epsilon_2) - q_{\tau_N}(x_2, z_2, a) \int A_{\tau_N}(s_2, a, \mathbf{d}) dF(\epsilon_2) \right).$$

where $s_2 = (x_2, z_2, \epsilon_2)$.

Recall that $q_{\tau_N}(x_2, z_2, a)$, defined in eq. (8) is the un-discounted expected period 3 utility when the period 1 agent (a) contemplates being in state (x_2, z_2) in period 2, (b) takes action $a \in \mathcal{A}_2$ and (c) assumes she behaves optimally in the static period 3 problem. $\int A_{\tau}(x_2, z_2, \epsilon_2, a, \mathbf{d}) dF(\epsilon_2)$ is the conditional probability that action a is optimal in state (x_2, z_2) when the period 1 agent believes that \mathbf{d} is used to discount 3 utility back to period 2 when making choices in period 2.

Roughly speaking, consider (x_2, z_2) (and (x_2, z_2)) as being fixed and view the action as a random variable with the probability of a occurring being $\mathbb{P}(A_{\tau_N}(x_2, z_2, a)) \equiv \int A_{\tau_N}(s_2, a, \mathbf{d}) dF(\epsilon_2)$. Then the summation $\sum_a q(x_2, z_2, a) \mathbb{P}(A_{\tau_N}(x_2, z_2, a))$ is the un-discounted period 3 expected utility from state (x_2, z_2) in period 2 from the viewpoint of the period 1 self who believes they will use \mathbf{d} to discount utility between periods 2 and 3. The assumption above states that are at least two points in $\mathcal{X}_2 \times \mathcal{Z}_2$ such that this expected utility is always strictly greater (or smaller) at one point (x_2, z_2) than at the other point (x_2, z_2) for all choices of the discount rate \mathbf{d} in the identified set for δ , Θ_{δ_N} – i.e. that the agent is always strictly better off (in period 3) at one point relative to the other point (in period 2), regardless of the choice of the exponential discount rate δ . The plausibility of the assumption clearly depends upon context and is likely to hold when there are certain states in period 2 that unambiguously lead to better outcomes in period 3 (relative other states).

Formally speaking, the assumption above ensures that the derivative of $\tilde{h}_{\tau_N}^{\Delta}(x_2, z_2, \mathbf{d})$ with respect to \mathbf{d} is strictly positive (or negative) everywhere in the identified set at least for some value of the state variable. With this additional assumption, we can separately identify both time preference parameters for naïve types.

LEMMA A3 (Point Identification for δ_N and β_N). Consider a time-inconsistent naïve agent solving the problem (1) at $t = 1$ and suppose that the model satisfies Assumptions B, D1, D2, D3, DA1 and DA2. Then δ_{τ_N} and β_{τ_N} are identified.

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

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

(where we defined $\tilde{h}_{\tau}(\cdot)$ in eq. (40)). In fact, this function is convex in \mathbf{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.

$$\frac{\partial \tilde{h}(x_2, z_2, d, d)}{\partial d} = \sum_{a \in \mathcal{A}_2} \frac{\partial \tilde{h}_\tau^a(x_2, z_2, d)}{\partial d},$$

where

$$\begin{aligned} \frac{\partial \tilde{h}_\tau^a(x_2, z_2, d)}{\partial d} &= \frac{\partial}{\partial d} \int (u_2(x_2, a, \tau) + \epsilon_2(a) + d q_\tau(x_2, z_2, a)) A_\tau(s_2, a, d, 1) dF(\epsilon_2) \\ &= \frac{\partial}{\partial d} \int (m(x_2, z_2, a, \tau, d) + \epsilon_2(a)) A_\tau(s_2, a, d, 1) dF(\epsilon_2), \end{aligned}$$

and where $m_a(x_2, z_2, \tau, d) = u_2(x_2, a, \tau) + d q_\tau(x_2, z_2, a)$.

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

$$\begin{aligned} \frac{\partial \tilde{h}_\tau^1(x_2, z_2, d)}{\partial d} &= (m'_1 - m'_0) \int_{\epsilon_{20}} \int_{\epsilon_{22} < m_0 - m_2 + \epsilon_{20}} (m_0 + \epsilon_{20}) f(\epsilon_{20}, m_0 - m_1 + \epsilon_{20}, \epsilon_{22}) d\epsilon_{22} d\epsilon_{20} \\ &\quad + m'_1 \int_{\epsilon_{20}} \int_{\epsilon_{22} < m_0 - m_2 + \epsilon_{20}} \int_{\epsilon_{21} > m_0 - m_1 + \epsilon_{20}} f(\epsilon_{20}, \epsilon_{21}, \epsilon_{22}) d\epsilon_{21} d\epsilon_{22} d\epsilon_{20} \\ &\quad + (m'_1 - m'_2) \int_{\epsilon_{20}} \int_{\epsilon_{22} > m_0 - m_2 + \epsilon_0} (m_2 + \epsilon_2) f(\epsilon_0, m_2 - m_1 + \epsilon_2, \epsilon_2) d\epsilon_{22} d\epsilon_{20} \\ &\quad + m'_1 \int_{\epsilon_{20}} \int_{\epsilon_{22} > m_0 - m_2 + \epsilon_{20}} \int_{\epsilon_{21} > m_2 - m_1 + \epsilon_{22}} f(\epsilon_{20}, \epsilon_{21}, \epsilon_{22}) d\epsilon_{21} d\epsilon_{22} d\epsilon_{20}. \end{aligned}$$

$$\begin{aligned} \frac{\partial \tilde{h}_\tau^2(x_2, z_2, d)}{\partial d} &= (m'_2 - m'_0) \int_{\epsilon_{20}} \int_{\epsilon_{22} < m_0 - m_2 + \epsilon_{20}} (m_0 + \epsilon_{20}) f(\epsilon_{20}, \epsilon_{21}, m_0 - m_2 + \epsilon_{20}) d\epsilon_{21} d\epsilon_{20} \\ &\quad + m'_2 \int_{\epsilon_{20}} \int_{\epsilon_{21} < m_0 - m_1 + \epsilon_{20}} \int_{\epsilon_{22} > m_0 - m_2 + \epsilon_{20}} f(\epsilon_{20}, \epsilon_{21}, \epsilon_{22}) d\epsilon_{22} d\epsilon_{21} d\epsilon_{20} \\ &\quad + (m'_2 - m'_1) \int_{\epsilon_{20}} \int_{\epsilon_{21} > m_0 - m_1 + \epsilon_{20}} (m_1 + \epsilon_1) f(\epsilon_0, \epsilon_1, m_1 - m_2 + \epsilon_1) d\epsilon_{21} d\epsilon_{20} \\ &\quad + m'_2 \int_{\epsilon_{20}} \int_{\epsilon_{21} > m_0 - m_1 + \epsilon_{20}} \int_{\epsilon_{22} > m_1 - m_2 + \epsilon_{21}} f(\epsilon_{20}, \epsilon_{21}, \epsilon_{22}) d\epsilon_{22} d\epsilon_{21} d\epsilon_{20}. \end{aligned}$$

$$\begin{aligned} \frac{\partial \tilde{h}_\tau^0(x_2, z_2, d)}{\partial d} &= (m'_0 - m'_1) \int_{\epsilon_{22}} \int_{\epsilon_{21} > m_2 - m_1 + \epsilon_{22}} (m_1 + \epsilon_{21}) f(m_1 - m_0 + \epsilon_{21}, \epsilon_{21}, \epsilon_{22}) d\epsilon_{21} d\epsilon_{22} \\ &\quad + m'_0 \int_{\epsilon_{22}} \int_{\epsilon_{21} > m_2 - m_1 + \epsilon_{22}} \int_{\epsilon_{20} > m_1 - m_0 + \epsilon_{21}} f(\epsilon_{20}, \epsilon_{21}, \epsilon_{22}) d\epsilon_{20} d\epsilon_{21} d\epsilon_{22} \\ &\quad + (m'_0 - m'_2) \int_{\epsilon_{22}} \int_{\epsilon_{21} < m_2 - m_1 + \epsilon_{22}} (m_2 + \epsilon_{22}) f(m_2 - m_0 + \epsilon_{22}, \epsilon_{21}, \epsilon_{22}) d\epsilon_{21} d\epsilon_{22} \\ &\quad + m'_0 \int_{\epsilon_{22}} \int_{\epsilon_{21} < m_2 - m_1 + \epsilon_{22}} \int_{\epsilon_{20} > m_2 - m_0 + \epsilon_{22}} f(\epsilon_{20}, \epsilon_{21}, \epsilon_{22}) d\epsilon_{20} d\epsilon_{21} d\epsilon_{22}. \end{aligned}$$

Adding the three terms and simplifying suitably, we obtain

$$\frac{\partial \tilde{h}_\tau(x_2, z_2, d, d)}{\partial d} = \sum_{a \in \mathcal{A}_2} q_\tau(x_2, z_2, a) \int_{\epsilon_2} A_\tau(s_2, a, d, 1) dF(\epsilon_2), \quad (43)$$

so that

$$\frac{\partial \tilde{h}_\tau^\Delta(x_2, z_2, d)}{\partial d} = \sum_{a \in \mathcal{A}_2} \left(q_{\tau_N}(x_2, z_2, a) \int A_\tau(s_2, a, d, 1) dF(\epsilon_2) - q_{\tau_N}(x_{20}, z_{20}, a) \int A_\tau(s_{20}, a, d, 1) dF(\epsilon_2) \right),$$

where $s_{20} = (x_{20}, z_{20}, \epsilon_2)$. The expression above is the expression in [Assumption DA2](#) from which the conclusion follows. \blacksquare

Point Identification Results for Naïve Types

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 DA3. *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 DA3](#) 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.

Alternative Assumptions for [Lemma A1](#)

Note that [Assumption DA1](#) 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. 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 $\mathcal{X}_2 \times \mathcal{Z}_2$)

$$dF_\Delta(z_1) \equiv \begin{bmatrix} dF_{\Delta, a_{1,1}}(x_{2,1}, z_{2,1} | x_{1,1}, z_1) & \dots & dF_{\Delta, a_{1,1}}(x_{2,S-1}, z_{2,S-1} | x_{1,1}, z_1) \\ dF_{\Delta, a_{1,2}}(x_{2,1}, z_{2,1} | x_{1,1}, z_1) & \dots & dF_{\Delta, a_{1,2}}(x_{2,S-1}, z_{2,S-1} | x_{1,1}, z_1) \\ \vdots & \vdots & \vdots \\ dF_{\Delta, a_{1,K_1-1}}(x_{2,1}, z_{2,1} | x_{1,1}, z_1) & \dots & dF_{\Delta, a_{1,K_1}}(x_{2,S-1}, z_{2,S-1} | x_{1,1}, z_1) \\ dF_{\Delta, a_{1,1}}(x_{2,1}, z_{2,1} | x_{1,2}, z_1) & \dots & dF_{\Delta, a_{1,1}}(x_{2,S-1}, z_{2,S-1} | x_{1,2}, z_1) \\ dF_{\Delta, a_{1,2}}(x_{2,1}, z_{2,1} | x_{1,2}, z_1) & \dots & dF_{\Delta, a_{1,2}}(x_{2,S-1}, z_{2,S-1} | x_{1,2}, z_1) \\ \vdots & \vdots & \vdots \\ dF_{\Delta, a_{1,K_1-1}}(x_{2,1}, z_{2,1} | x_{1,2}, z_1) & \dots & dF_{\Delta, a_{1,K_1}}(x_{2,S-1}, z_{2,S-1} | x_{1,2}, z_1) \\ \vdots & \vdots & \vdots \\ dF_{\Delta, a_{1,1}}(x_{2,1}, z_{2,1} | x_{1,S_1}, z_1) & \dots & dF_{\Delta, a_{1,1}}(x_{2,S-1}, z_{2,S-1} | x_{1,S_1}, z_1) \\ dF_{\Delta, a_{1,2}}(x_{2,1}, z_{2,1} | x_{1,S_1}, z_1) & \dots & dF_{\Delta, a_{1,2}}(x_{2,S-1}, z_{2,S-1} | x_{1,S_1}, z_1) \\ \vdots & \vdots & \vdots \\ dF_{\Delta, a_{1,K_1-1}}(x_{2,1}, z_{2,1} | x_{1,S_1}, z_1) & \dots & dF_{\Delta, a_{1,K_1}}(x_{2,S-1}, z_{2,S-1} | x_{1,S_1}, z_1) \end{bmatrix}.$$

ASSUMPTION DA4 (Alternative for [Assumption DA1](#)). *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 (the “order” condition) and so relaxes the requirement on \mathcal{X}_1 . We can then state the modified version of [Lemma A1](#).

LEMMA A4 (Alternative Period 1 Identification). *Consider an agent of type τ maximizing equation (1) and suppose that the model satisfies Assumptions B, D1, D2, D3 and DA4 (the last of which replaces Assumption DA1). Then*

1. *The function $h_{\tau,2}(x_2, z_2)$ 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 $\forall (a \in \mathcal{A}_1, x_1 \in \mathcal{X}_1, \tau \in \mathcal{T})$.*

Proof. The proof is essentially identical to the proof of [Lemma A1](#) with the only modification being how the function $h_\tau(x_2, z_2)$ is identified. To this end, first define $\mathbf{H} = \mathbf{dF}_\Delta(z'_1) - \mathbf{dF}_\Delta(z''_1)$. Then by [Assumption DA4](#) the matrix $\mathbf{H}^T \mathbf{H}$ has rank $S - 1$ and is invertible.

Consider the points $(a_{1,i}, x_{1,j})_{i=1 \dots K_1, j=1 \dots S_1}$ and consider the expression (derived in eq. (37))

$$\frac{g_{\tau, 1, a_{1,i}, \Delta}(x_{1,j})}{\beta_\tau \delta_\tau} = \sum_{(x_2, z_2) \in \mathcal{X}_2 \times \mathcal{Z}_2} h_{\tau, 2}(x_2, z_2) \left(\mathbf{dF}_{\Delta, a_{1,i}}(x_2, z_2 | x_{1,j}, z'_1) - \mathbf{dF}_{\Delta, a_{1,i}}(x_2, z_2 | x_{1,j}, z''_1) \right).$$

stacking the equations and defining the left-hand side $(K_1 \times S_1) \times 1$ vector as $\bar{\mathbf{g}}_{\tau, \Delta}$ we obtain

$$\bar{\mathbf{g}}_{\tau, \Delta} = \mathbf{H} \mathbf{h}_\tau.$$

where \mathbf{h}_τ is defined in eq. (38). Consequently, under [Assumption DA4](#) we recover \mathbf{h}_τ as $(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \bar{\mathbf{g}}_{\tau, \Delta}$. The remainder of the proof is identical to the one above and is omitted. \blacksquare

A.3 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)

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

where the matrices above are defined in [eqs. \(12\), \(14\)](#) and [\(15\)](#). It follows that

$$\text{Rank}(\mathbf{P}_{r,v}^{a_2, a_3, \underline{M}}) \leq \min \{ \text{Rank}(\mathbf{L}_v^{a_2, \mathbf{x}_2, \underline{M}+1}), \text{Rank}(\mathbf{L}_{\mathbf{x}_3, v}^{a_3, \underline{M}+1}), \text{Rank}(\mathbf{V}_{r,v}^{M_{r,v}}) \}.$$

Since the rank of $\mathbf{V}_{r,v}^{M_{r,v}} = M_{r,v}$ we have that

$$M_{r,v} \geq \text{Rank}(\mathbf{P}_{r,v}^{a_2, a_3, \underline{M}}). \quad (44)$$

For the second part, suppose that in addition [Assumption U2](#) holds. We will show that in that case the reverse inequality holds. First note that then $\mathbf{L}_{\mathbf{x}_3, v}^{a_3, \underline{M}+1} (\mathbf{L}_{\mathbf{x}_3, v}^{a_3, \underline{M}+1})'$ is invertible. Post-multiplying both sides by $(\mathbf{L}_{\mathbf{x}_3, v}^{a_3, \underline{M}+1})' \left(\mathbf{L}_{\mathbf{x}_3, v}^{a_3, \underline{M}+1} (\mathbf{L}_{\mathbf{x}_3, v}^{a_3, \underline{M}+1})' \right)^{-1}$,

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

Since $\mathbf{L}_v^{a_2, \mathbf{x}_2, \underline{M}+1}$ has rank $M_{r,v}$ and $\mathbf{V}_{r,v}^{M_{r,v}}$ has strictly positive diagonal elements, it must be the case that the rank of $\mathbf{P}_{r,v}^{a_2, a_3, \underline{M}} (\mathbf{L}_{\mathbf{x}_3, v}^{a_3, \underline{M}+1})' \left(\mathbf{L}_{\mathbf{x}_3, v}^{a_3, \underline{M}+1} (\mathbf{L}_{\mathbf{x}_3, v}^{a_3, \underline{M}+1})' \right)^{-1}$ is $M_{r,v}$ and it follows that

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

so that $\text{Rank}(\mathbf{P}_{r,v}^{a_2, a_3, \underline{M}}) \geq M_{r,v}$. Combining this with eq. (44) we conclude that $\text{Rank}(\mathbf{P}_{r,v}^{a_2, a_3, \underline{M}}) = M_{r,v}$. Therefore, the rank of the directly identified matrix $\mathbf{P}_{2,r,v}^{a_2, a_3, \underline{M}}$ gives the total number of types in the population. \blacksquare

A.3.1 Proofs for Identification of Type Specific Choice Probabilities

Proof of Lemma 3

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 U1](#)) to generate identification instead of using observations that are more than one period apart. As on p.18 we begin by using [Assumption U1](#) 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 U1](#). 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

$$\mathbf{P}_{r,v}^{a_t, a_{t+1}, M_{r,v}} = \begin{pmatrix} 1 & \mathbf{F}_{\mathbf{x}_{t+1}^1, r, v}^{a_{t+1}} & \dots & \mathbf{F}_{\mathbf{x}_{t+1}^{M_{r,v}-1}, r, v}^{a_{t+1}} \\ \mathbf{F}_{\mathbf{x}_t^1, r, v}^{a_t} & \mathbf{F}_{\mathbf{x}_t^1, \mathbf{x}_{t+1}^1, r, v}^{a_t, a_{t+1}} & \dots & \mathbf{F}_{\mathbf{x}_t^1, \mathbf{x}_{t+1}^{M_{r,v}-1}, r, v}^{a_t, a_{t+1}} \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{F}_{\mathbf{x}_t^{M_{r,v}-1}, r, v}^{a_t} & \mathbf{F}_{\mathbf{x}_t^{M_{r,v}-1}, \mathbf{x}_{t+1}^1, r, v}^{a_t, a_{t+1}} & \dots & \mathbf{F}_{\mathbf{x}_t^{M_{r,v}-1}, \mathbf{x}_{t+1}^{M_{r,v}-1}, r, v}^{a_t, a_{t+1}} \end{pmatrix}, \quad (45)$$

which is just the matrix defined in eq. (12) but with \underline{M} replaced by $M_{r,v} - 1$. We will abbreviate $\mathbf{P}_{r,v}^{a_t, a_{t+1}, M_{r,v}}$ to \mathbf{P}_r in the sequel to economize on notation.

Next, we define the matrices $\mathbf{L}_v^{a_t, \mathbf{x}_t, M_{r,v}}$ and $\mathbf{L}_{\mathbf{x}_{t+1}, v}^{a_{t+1}, M_{r,v}}$ using [eqs. \(14\)](#) and [\(15\)](#) but replacing \underline{M} with $(M_{r,v} - 1)$ (so the dimensions now depend only upon the number of types $M_{r,v}$). Thus,

$$\mathbf{L}_v^{a_t, \mathbf{x}_t, M_{r,v}} \equiv \begin{pmatrix} 1 & \lambda_v^{a_t, \mathbf{x}_t^1, \tau_1} & \dots & \lambda_v^{a_t, \mathbf{x}_t^{M_{r,v}-1}, \tau_1} \\ 1 & \lambda_v^{a_t, \mathbf{x}_t^1, \tau_2} & \dots & \lambda_v^{a_t, \mathbf{x}_t^{M_{r,v}-1}, \tau_2} \\ \vdots & \vdots & \dots & \vdots \\ 1 & \lambda_v^{a_t, \mathbf{x}_t^1, \tau_{M_{r,v}}} & \dots & \lambda_v^{a_t, \mathbf{x}_t^{M_{r,v}-1}, \tau_{M_{r,v}}} \end{pmatrix}, \quad (46)$$

which we will abbreviate as $\mathbf{L}_{t,r}$. Similarly,

$$\mathbf{L}_{\mathbf{x}_{t+1}, v}^{a_{t+1}, M_{r,v}} \equiv \begin{pmatrix} 1 & \lambda_{\mathbf{x}_{t+1}^1, v}^{a_{t+1}, \tau_1} & \dots & \lambda_{\mathbf{x}_{t+1}^{M_{r,v}-1}, v}^{a_{t+1}, \tau_1} \\ 1 & \lambda_{\mathbf{x}_{t+1}^1, v}^{a_{t+1}, \tau_2} & \dots & \lambda_{\mathbf{x}_{t+1}^{M_{r,v}-1}, v}^{a_{t+1}, \tau_2} \\ \vdots & \vdots & \dots & \vdots \\ 1 & \lambda_{\mathbf{x}_{t+1}^1, v}^{a_{t+1}, \tau_{M_{r,v}}} & \dots & \lambda_{\mathbf{x}_{t+1}^{M_{r,v}-1}, v}^{a_{t+1}, \tau_{M_{r,v}}} \end{pmatrix}. \quad (47)$$

which we will abbreviate as $\mathbf{L}_{t+1,r}$.

[Assumption U1](#) and the overlap condition (in [Assumption U2'](#)) guarantee that $\mathbf{L}_{t,r} = \mathbf{L}_{t,r'}$ and $\mathbf{L}_{t+1,r} = \mathbf{L}_{t+1,r'}$. Finally, define the $M_{r,v} \times M_{r,v}$ matrix $\mathbf{V}_r = \text{Diag}(\pi_{\tau_1}(r, v), \dots, \pi_{\tau_{M_{r,v}-1}}(r, v))$ which we will abbreviate to \mathbf{V}_r . It is easy to show the following factorization holds:

$$\mathbf{P}_r = \mathbf{L}_{t,r}' \mathbf{V}_r \mathbf{L}_{t+1,r}, \quad (48)$$

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

$$\mathbf{A} \equiv \mathbf{P}_r^{-1} \mathbf{P}_{r'} = \mathbf{L}_{t+1,r}^{-1} \mathbf{V}_r^{-1} \mathbf{V}_{r',v} \mathbf{L}_{t+1,r}, \quad (49)$$

so that

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

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

Next,

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

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

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

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

$$\mathbf{V}_r \mathbf{L}_{t+1,r} = (\mathbf{L}'_{t,r})^{-1} \mathbf{P}_r = \mathbf{P}_r \mathbf{E}^{-1} \mathbf{F}^{-1} \mathbf{P}_r. \quad (50)$$

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

$$\mathbf{L}_{t+1,r} = \mathbf{G}^{-1} \mathbf{P}_r \mathbf{E}^{-1} \mathbf{F}^{-1} \mathbf{P}_r, \quad (51)$$

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

We first apply this result to $(t, t+1) \in \{(2, 3), (1, 2)\}$ in order 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. \blacksquare

We next state a result that does not require the overlap condition ([Assumption U2'](#) – 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 U2'](#). We can, however, substantially relax [Assumption U1](#) to:

ASSUMPTION UA1. *The transition probabilities do not vary by type: $\mathbb{P}_\tau(\mathbf{x}_{t+1} | \mathbf{x}_t, a_t, r, v) = \mathbb{P}(\mathbf{x}_{t+1} | \mathbf{x}_t, a_t, r, v)$.*

So, type-specific choice probabilities need not be independent of the type proxy r – the recovered preference parameters will then also be indexed by r (i.e. $u_{\tau,r}(x_t, v)$ and $(\beta_{\tau,r}, \delta_{\tau,r})$). [Assumption U2'](#) is, however, strengthened to

ASSUMPTION UA2. *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 $(j, k) \in \{1, \dots, M_{r,v}\}$ and
- (b) the $M_{r,v} \times M_{r,v}$ matrices \mathbf{L}_v and $\mathbf{L}_{t+1,v}$ are invertible. In addition, the matrix $\mathbf{P}^{\mathbf{x}_1}$ defined below is invertible.

LEMMA A5 (Alternative Result for Identifying Type-Specific Choice Probabilities). *Fix (r, v) and suppose that [Assumption UA1](#), and [Assumption UA2](#) 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 3](#) with the main addition being that we now examine events in three consecutive time-periods (as opposed to two periods earlier). To ease exposition we suppress

the dependence on v throughout.

First, define the identified quantities

$$\begin{aligned}\mathbf{F}_{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, r} &\equiv \sum_{\tau \in \mathcal{T}_r} \pi_\tau(r) \mathbb{P}_\tau(a_3 | \mathbf{x}_3, r) \mathbb{P}_\tau(a_2 | \mathbf{x}_2, r) \mathbb{P}_\tau(a_1, \mathbf{x}_1 | r) \\ &= \sum_{\tau \in \mathcal{T}_r} \pi_\tau(r) \lambda_{\mathbf{x}_3, r}^{a_3, \tau} \lambda_{\mathbf{x}_2, r}^{a_2, \tau} \lambda_r^{a_1, \mathbf{x}_1, \tau}\end{aligned}$$

$$\begin{aligned}\mathbf{F}_{\mathbf{x}_1, \mathbf{x}_2, r} &= \sum_{\tau \in \mathcal{T}_r} \pi_\tau(r) \mathbb{P}_\tau(a_1, \mathbf{x}_1 | r) \mathbb{P}_\tau(a_2 | \mathbf{x}_2, r) = \sum_{\tau \in \mathcal{T}_r} \pi_\tau(r) \lambda_{\mathbf{x}_2, r}^{a_2, \tau} \lambda_r^{a_1, \mathbf{x}_1, \tau} \\ \mathbf{F}_{\mathbf{x}_1, \mathbf{x}_3, r} &= \sum_{\tau \in \mathcal{T}_r} \pi_\tau(r) \mathbb{P}_\tau(a_3 | \mathbf{x}_3, r) \mathbb{P}_\tau(a_1 | \mathbf{x}_1 | r) = \sum_{\tau \in \mathcal{T}_r} \pi_\tau(r) \lambda_{\mathbf{x}_3, r}^{a_3, \tau} \lambda_r^{a_1, \mathbf{x}_1, \tau} \\ \mathbf{F}_{\mathbf{x}_1, r} &= \sum_{\tau \in \mathcal{T}_r} \pi_\tau(r) \mathbb{P}_\tau(a_1 | \mathbf{x}_1 | r) = \sum_{\tau \in \mathcal{T}_r} \pi_\tau(r) \lambda_r^{a_1, \mathbf{x}_1, \tau}\end{aligned}$$

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

$$\mathbf{P}_{r, v}^{\mathbf{x}_1} = \begin{pmatrix} \mathbf{F}_{\mathbf{x}_1} & \mathbf{F}_{\mathbf{x}_1, \mathbf{x}_3^1} & \cdots & \mathbf{F}_{\mathbf{x}_1, \mathbf{x}_3^{M_r, v-1}} \\ \mathbf{F}_{\mathbf{x}_1, \mathbf{x}_2^1} & \mathbf{F}_{\mathbf{x}_1, \mathbf{x}_2^1, \mathbf{x}_3^1} & \cdots & \mathbf{F}_{\mathbf{x}_1, \mathbf{x}_2^1, \mathbf{x}_3^{M_r, v-1}} \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{F}_{\mathbf{x}_1, \mathbf{x}_2^{M_r, v-1}} & \cdots & \cdots & \mathbf{F}_{\mathbf{x}_1, \mathbf{x}_2^{M_r, v-1}, \mathbf{x}_3^{M_r, v-1}} \end{pmatrix}, \quad (52)$$

Define the matrix $\mathbf{D}_{\mathbf{x}_1, r} \equiv \text{Diag}(\lambda_r^{a_1, \mathbf{x}_1, \tau_1}, \dots, \lambda_r^{a_1, \mathbf{x}_1, \tau_{M_r, v-1}})$ and as before $\mathbf{V}_r = \text{Diag}(\pi_{\tau_1}(r), \dots, \pi_{\tau_{M_r-1}}(r))$. Then, the following factorizations hold – suppressing the dependence on r since identification does not depend upon variation in r

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

where the matrix $\mathbf{P} = \mathbf{P}_{r, v}^{a_2, a_3, M_r, v}$ is defined in eq. (45) and the matrices $(\mathbf{L}_2, \mathbf{L}_3)$ are defined in eq. (46). 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 \mathbf{A} defined by

$$\mathbf{A} \equiv \mathbf{P}^{-1} \mathbf{P}^{\mathbf{x}_1} = \mathbf{L}_3^{-1} \mathbf{D}_{\mathbf{x}_1} \mathbf{L}_3, \quad (53)$$

so that

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

The eigenvalues of \mathbf{A} determine $\mathbf{D}_{\mathbf{x}_1}$ and the rows of \mathbf{L}_3 are the left eigenvectors of \mathbf{A} . Therefore, these objects are identified by carrying out an eigenvalue decomposition of the identified matrix \mathbf{A} . Note that the eigenvectors are only identified up to scale, so that we can identify the matrix $\mathbf{E} \equiv \mathbf{H} \mathbf{L}_3$ where \mathbf{H} is a diagonal matrix with non-zero diagonal entries (and we have $\mathbf{L}_3 = \mathbf{H}^{-1} \mathbf{E}$). Next,

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

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

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

where all the terms on the right hand side are identified, so that \mathbf{L}_2 is identified. Next,

$$\mathbf{V} \mathbf{L}_3 = (\mathbf{L}_2')^{-1} \mathbf{P} = \mathbf{P} \mathbf{E}^{-1} \mathbf{F}^{-1} \mathbf{P}. \quad (54)$$

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 eq. (54) are identified. Denote by $G(\equiv V)$ the diagonal matrix obtained by this argument. Then,

$$L_3 = G^{-1}P E^{-1} F^{-1} P, \quad (55)$$

where the matrix G is invertible since by assumption all its diagonal entries are non-zero. Finally, note that since V is identified, then $D_{x_1} = (L_2' V)^{-1} P^{x_1} L_3^{-1}$ is also identified. \blacksquare

A.3.2 Proofs for Identification of Type Identities

ASSUMPTION UA3.

For $\tilde{h}_\tau^\Delta(x_2, z_2, d_1, d_2)$ is defined in eq. (41) and

$$\tilde{h}_\tau^{\Delta, j}(x_2, z_2, d) \equiv \tilde{h}_\tau^{\Delta, j}(x_2, z_2, d) - \tilde{h}_\tau^{\Delta, j}(x_{20}, z_{20}, d) \quad j \in \{A, B\}$$

1. There exists a set $\mathcal{S} \subset \mathcal{X}_2 \times \mathcal{Z}_2$ with positive measure such that for all types τ , $\tilde{h}_\tau^{\Delta, B}(x_2, z_2, \beta_\tau \delta_\tau) \neq 0$.
2. $\text{Var} \left(\frac{\tilde{h}_N^\Delta(x_2, z_2, \delta_N, \delta_N) - \tilde{h}_N^\Delta(x_2, z_2, \delta_N, \beta_N \delta_N)}{\tilde{h}_N^{\Delta, B}(x_2, z_2, \beta_N \delta_N)} \right) > 0$.

The first part of the assumption is a significant strengthening of the second part of [Assumption D2](#) and applies to all types (not just sophisticated types) which is necessary since type identities are not known. Note, however, that it is testable since it is imposed on an identified object. The assumption ensures that δ_τ enters $h_\tau^\Delta(x_2, z_2)$ linearly for sophisticated types and guarantees that eq. (18) is well-defined at least on \mathcal{S} for all types.³³

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 A6. 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 A1](#) hold. Further, suppose that [Assumption UA3](#) holds. Then,

$$\hat{\delta}_\tau(x_2, z_2) \text{ is a constant for all } (x_2, z_2) \in \mathcal{S} \iff \tau \neq \tau_N.$$

Proof. “ \Rightarrow ”: Suppose that $\hat{\delta}_\tau$ is a constant but $\tau = N$. First, note that

$$\begin{aligned} h_N^\Delta(x_2, z_2) &= \tilde{h}_N^\Delta(x_2, z_2, \delta_N, \tilde{\beta}_N \delta_N) = \tilde{h}_N^\Delta(x_2, z_2, \delta_N, \delta_N) \\ &= \tilde{h}_N^{\Delta, A}(x_2, z_2, \delta_N) + \delta_N \tilde{h}_N^{\Delta, B}(x_2, z_2, \delta_N) \end{aligned}$$

If instead, naïve types used $\beta_N \delta_N$ (instead of δ_N) to discount period three utility to period two (from the viewpoint of the period one self), then the corresponding version of the $\tilde{h}_N^\Delta(\cdot)$ function would be given by

$$\tilde{h}_N^\Delta(x_2, z_2, \delta_N, \beta_N \delta_N) = \tilde{h}_N^{\Delta, A}(x_2, z_2, \beta_N \delta_N) + \delta_N \tilde{h}_N^{\Delta, B}(x_2, z_2, \beta_N \delta_N)$$

Next, rewrite

$$\hat{\delta}_N(x_2, z_2) = \frac{h_N^\Delta(x_2, z_2) - \tilde{h}_N^{\Delta, A}(x_2, z_2, \beta_N \delta_N)}{\tilde{h}_N^{\Delta, B}(x_2, z_2, \beta_N \delta_N)}$$

by adding and subtracting $\delta_N \tilde{h}_N^{\Delta, B}(x_2, z_2, \beta_N \delta_N)$ to the numerator we obtain

$$\hat{\delta}_N(x_2, z_2) = \frac{\tilde{h}_N^\Delta(x_2, z_2, \delta_N, \delta_N) - \tilde{h}_N^\Delta(x_2, z_2, \delta_N, \beta_N \delta_N)}{\tilde{h}_N^{\Delta, B}(x_2, z_2, \beta_N \delta_N)} - \delta_N.$$

The first term in the expression above is non-constant by [Assumption UA3](#) so we have a contradiction. The “ \Leftarrow ” follows by simply observing that for $\tau \in \{\tau_S, \tau_C\}$, $\hat{\delta}_\tau(x_2, z_2) = \delta_\tau$. \blacksquare

³³To see this, note that $h_\tau^\Delta(x_2, z_2) \equiv \tilde{h}_\tau^{\Delta, A}(x_2, z_2, \tilde{\beta}_\tau \delta_\tau) + \delta_\tau \tilde{h}_\tau^{\Delta, B}(x_2, z_2, \tilde{\beta}_\tau \delta_\tau)$.

The proof for [Proposition 2](#) proceeds by applying [Lemma 1](#) to identify period 2 and 3 payoffs as well as the product $\beta_\tau \delta_\tau$ (note that [Lemma 1](#) does not require knowledge of the type identities). Using these objects, we then apply [Lemma A1](#) to identify $h_\tau^\Delta(x_2, z_2)$. [Lemma A1](#) requires an invertibility condition ([Assumption DA1](#)) on an appropriately differenced transition probability matrix. Such an assumption is likely to hold as long as the excluded variable z_1 induces sufficient variation in the transition probabilities (we discuss the formal statement and its plausibility in [Appendix A.2.1](#)).³⁴ We then use the objects identified in [Lemma A1](#) to construct the function $\hat{\delta}_{\tau,2}(x_2, z_2)$. Finally, [Lemma A6](#) shows that $\hat{\delta}_{\tau,2}(x_2, z_2)$ is a constant for all (x_2, z_2) if and only if the type τ is not naïve. When $\hat{\delta}_{\tau,2}(x_2, z_2)$ is a constant, we can then distinguish between consistent and sophisticated types by comparing the identified quantities $\beta_\tau \delta_\tau$ and $\hat{\delta}_{\tau,2}(x_2, z_2)$.

Proof of [Proposition 2](#)

Proof. First, note that the results of [Lemma 1](#) and [Lemma A1](#) 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 normalized function $h_\tau^\Delta(x_2, z_2)$ defined in eq. (8). Using these, we can construct the known function $\hat{\delta}_\tau(x_2, z_2)$ defined in eq. (18). First, suppose that $\hat{\delta}_\tau(x_2, z_2)$ is a constant (in (x_2, z_2)). Then, by [Lemma A6](#) the type must be either sophisticated or consistent and that $\hat{\delta}(x_2, z_2) = \delta_\tau$. Next, if the ratio of the directly identified objects $\frac{\beta_\tau \delta_\tau}{\hat{\delta}_\tau} = 1$ then we can conclude the type must be consistent while if the ratio is strictly less than one, the type must be inconsistent and sophisticated. If, however, $\hat{\delta}_\tau(x_2, z_2)$ is not a constant, then by [Lemma A6](#) the type must be naïve. ■

A.4 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 $h_\tau^\Delta(x_2, z_2) = \tilde{h}_\tau^\Delta(x_2, z_2, \delta_\tau, \beta_\tau \delta_\tau)$ (where the latter is defined on p.23 in eq. (19)). Then, $\hat{\delta}_\tau(x_2, z_2) = \delta_\tau$ 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 A6](#) rewrite $\hat{\delta}_\tau(x_2, z_2)$ as

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

By [Assumption U3](#) 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 A1](#) 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 eq. (8). Using these, we can construct the known function $\hat{\delta}_\tau(x_2, z_2)$ defined in eq. (18).³⁵

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 $\tilde{\beta}_\tau = \beta_\tau$ 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 $\beta_\tau \delta_\tau / \hat{\delta}_\tau = 1$ 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 basis of the values of $\tilde{\beta}_\tau$ (e.g. into completely naïve agents) since the latter is not identified. ■

³⁴In particular, [Lemma A4](#) shows identification of h_τ^Δ using [Assumption DA4](#) (as an alternative to [Assumption DA1](#)) which places fewer restrictions on the support of \mathcal{X}_1 .

³⁵Note that the assumptions for [Lemma A1](#) imply that [Lemma 1](#) holds.

Proof of [Proposition 4](#)

Proof. Under the assumptions, the function $h_\tau^\Delta(x_2, z_2)$ is identified (part 1 of [Lemma A1](#)) and is equal to the function

$\tilde{h}_\tau^\Delta(x_2, z_2, \delta_\tau, \tilde{\beta}_\tau \delta_\tau)$. 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 $(\tilde{\beta}_\tau, \delta_\tau)$. \blacksquare

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. The main reason for this discussion is to account for the deviations in our setting from that of the general setting in [Section 3](#). There are four 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 $\mathbb{P}(\mathbf{x}_{t+1}|\mathbf{x}_t, v)$. Since we do not exploit any time-series variation in z for identification, this deviation does not impose any new difficulties. (b) In the directly observed types case, not all types are directly identified in the first-period since we need first-period choices to determine types. We derive the relevant identification results for this period below. (c) For tractability we assume that exactly three types exist in the population so we do not need to first identify the total number of types. (d) Finally, our setting requires an additional period (period four) to rationalize choices in period three (since they involve an expectational component). This additional period (where no action is taken) adds a complication since terminal period utilities are not identified by the standard arguments as in [Section 3](#). We discuss the modifications needed for the identification arguments to go through below.

As earlier, 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 for the general model.

Observable State Space: \mathcal{X}_t

In the empirical work and formal identification results we allow for a rich observable state space (including income and other characteristics), but for ease of exposition we simplify the state space to the bare minimum required for identification. [Online Appendix E](#) 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 and zero otherwise. In periods 2 and 3, $x_t \in \mathcal{X}_t = (nm, nh, bm, bh, cm, ch)$, where the first letter in each pair records the purchase decision of the agent, while the second indicates malaria status. The agent can either purchase no net (n) or purchase one with one of two contracts. Contract b (“base”) 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 re-treatments for cash. Contract c may be appealing to sophisticated agents who wish to commit to the ITN maintenance schedule at the time of purchase. However, demand for c may also depend on factors different from time preferences. 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”) if no-one did. The state space can be easily extended so that agents keep track of their entire history of malaria.

Observables: \mathcal{Z} (Beliefs)

Beliefs form a key ingredient for identification in the empirical section. In particular, during the baseline we elicited beliefs about: (a) the likelihood of contracting malaria (when using ITNs, untreated nets and no nets) and (b) income expectations (see [Appendix E](#) for more details).

We use the elicited malaria baseline beliefs, denoted by z , as the excluded variable in [Assumption B](#). There is one significant difference in the use of beliefs in our empirical application and the z_t variable in the general identification results: our beliefs are measured at baseline and are time-invariant (and thus are best viewed as conditioning variables rather than state variables). This means that we are abstracting away from belief evolution, ruling out, for instance, learning about ITN efficacy as a motive for purchase.

Action Space: \mathcal{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 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.

With the definition of the state and action space in hand we can now discuss the substantive content of the assumptions made in [Section 3](#). We begin by imposing [Assumption B](#), with the modification that $z_t \equiv z$ and discuss its plausibility and implications below.³⁶ The *markov property* in [Assumption B](#) rules out a complicated state dependence structure, or requires a suitable re-definition of the state space if these dependencies are important. For instance, it rules out the possibility that the probability that malaria infection in period three conditional on malaria status and retreatment in period two, depends on malaria status in period one. Incorporating such a dependency will require suitably redefining the state space (in this instance we would need to re-define the state space at t to contain the entire malaria history up to t). The *exclusion restriction* in [Assumption B](#) requires that household beliefs—about the likelihood of malaria infection in the next period—do not enter the current period utility function. While we cannot test this assumption directly, it seems plausible that conditional on current health status and income, beliefs about malaria in the next period do not affect utility. Before discussing the next assumption on directly observed types, we outline our type indicator.

Type Indicator

We collect information about whether individuals exhibit preference reversals in a series of questions designed to gauge the extent of consistency in time preferences (these are known as “Money Earlier or Later” (MEL) questions) In previous work we show that these reversals are important predictors (in a reduced form sense) of subsequent decisions about net 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 $\tilde{r} = 0$. We use this as a type signal in both the observed types case as well as the unobserved types case.

We also designed a contract that should appeal to sophisticated inconsistent agents and we use this as a indicator of type in the directly observed types case (for $t > 1$). We do not use product choice as a type signal in the unobserved types case because the likelihood function generated by doing so had significant disadvantages (we discuss this below in [Appendix F](#)). Instead, we use only the MEL responses as type signals and then ex-post evaluate the ability of the commitment contract to predict type.

Directly Identified Types Case

In the directly observed types case, we use both the MEL response as well as the choice of contract to characterize the type of agent. Agents with $r = 0$ are classified as time-consistent and agents with $r = 1$ are classified as time-inconsistent. We can further classify agents with $(r = 1, a_1 = c)$ as sophisticated and $(r = 1, a_1 = n)$ as naïve and thereby provide a complete classification of agents into types and when we estimate the model assuming types are known, this is indeed the mapping we use. However, for completeness, we also consider identification under the possibility that agents who do not purchase a net ($a_1 = n$) and have $r = 1$ 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.

Unobserved Types Case

In this case, the researcher does not directly observe the type of any individual. We assume instead that the MEL variable r is only an imperfect indicator for type. For instance, an agent may choose $r = 1$ due to an imperfect understanding of the choices offered rather than genuine time-inconsistency. Alternatively, an agent who expects sufficiently high income at the time of re-treatment may not choose the commitment product regardless of time-inconsistency.³⁷ Finally, MEL responses may also reflect rates of return to investments (see [Cohen et al., 2020](#), for an overview of the debates around the relationship between MEL responses and time-preferences). For these reasons, we will assume that r does not map deterministically into types. Instead, we will impose a weaker requirement explicated below.

³⁶As a reminder, [Assumption B](#) comprises the markov property, the knowledge of the error distribution, the exclusion restriction, additive separability and normalization. The knowledge of the error distribution, the additive separability and normalization assumptions are all standard in dynamic choice models and so we do not discuss them here.

³⁷In principle, the decision to not 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.

B.1 Directly Observed Types

As before, we first discuss identification for the directly observed types case. The main difference is that we need to modify [Assumption D1](#) to reflect the fact that given our type-proxy, not all types are observed.

ASSUMPTION DE1 (Modification of [Assumption D1](#)). *Choice probabilities for types that purchase a product are directly observed. In particular, for a time-consistent agent*

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

For a naïve time-inconsistent agent

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

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

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

The assumptions that (a) information about preference reversals in MEL choice questions is informative about time preferences and (b) the purchase of a commitment product reveals agents as being sophisticated is common in the empirical literature (see e.g. [Ashraf et al., 2006](#); [Andersen et al., 2008](#)). [Carrera et al. \(2022\)](#) note 33 studies that examine various commitment contracts and offer a critical review of the plausibility of such contracts identifying sophisticated inconsistent agents (indeed, our unobserved types approach is an attempt weaken both (a) and (b)).

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. This adds a complication relative to the identification argument in [Section 3.1](#) since terminal period utilities are not identified by standard arguments as they would be if actions were taken in the terminal period.

To be specific, utility for type τ in period 4 is $u_4(x_4; \tau)$ and the agent's choice in period 3 is

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

where we define the signed measure

$$dF_{\Delta, k}(x_3|x_2, z) \equiv (dF(x_3|x_2, z, k) - dF(x_3|x_2, z, 0)), \quad (56)$$

and we will on occasion abbreviate $dF_{\Delta, k} = dF_\Delta$ if no ambiguity results. Next, using the Hotz-Miller inversion argument we can identify the function

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

and as before, using variation in z we can isolate the last term in the expression above.

$$g_{\tau, 3, \Delta}(x_3) \equiv g_{\tau, 3, 1}(x_3, z) - g_{\tau, 3, 1}(x_3, z') = \int_{x_4 \in \mathcal{X}_4} \beta_\tau \delta_\tau u_4(x_4; \tau) (dF_\Delta(x_4|x_3, z) - dF_\Delta(x_4|x_3, z')). \quad (57)$$

Next, we seek to identify its constituent components using the variation in x_3 . Intuitively, as long as x_3 has sufficiently rich support (specifically, as large as that of x_4) we can generate a system of linear equations using the support of x_3 as follows: let S_4 denote the number of elements in \mathcal{X}_4 . Define the square matrix

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

and define the matrix $\mathbf{dF}_{\Delta,1}(z) \equiv \mathbf{dF}(1, z) - \mathbf{dF}(0, z)$. As long as this matrix is invertible we can identify period three preferences as well as the product of the time preferences and period four utility (see e.g. [Assumption DA1](#) for a similar requirement in the general case):

ASSUMPTION DE2 (Invertibility, Modification of [Assumption DA1](#)). *Suppose*

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

In the empirical application, period three has a much richer state-space than of period 4 and the beliefs conditional on x_3 have at least two points of support, suggesting that this assumption is not onerous. The invertibility assumption requires that beliefs about the likelihood of period 4 states be sufficiently different across individuals (i.e. possess at least two points of support with at which future state likelihoods differ) and our belief data, though it has marked modes, does possess this level of variability.

LEMMA B1 (Identification of Period 3 and (Scaled) Period 4 Preferences). *Consider an agent solving eq. (1) and suppose that we assume that all the [Basic Assumptions](#) hold with the modification that $z_t \equiv z$ which is no longer a state variable but instead a conditioning variable. In addition, suppose that [Assumption DE1](#) and [Assumption DE2](#) hold.³⁸ 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 $x_{4,0} \in \mathcal{X}_4$.

It is useful to note that this lemma does not require knowledge of the type-identity (i.e. whether a type τ is consistent, naïve or sophisticated), which will prove useful in the sequel (specifically for [Proposition 6](#) below).

Proof. The argument of the proof is similar to the argument in the proof of [Lemma A1](#). We begin by defining the $S_4 - 1$ column vectors $\mathbf{g}_{\tau,3,\Delta}(k)$ and $\mathbf{h}_{\tau,4}$,

$$\mathbf{g}_{\tau,3,\Delta} \equiv \begin{bmatrix} g_{\tau,\Delta}(x_{3,1}) \\ \vdots \\ g_{\tau,\Delta}(x_{3,S_4-1}) \end{bmatrix} \quad \mathbf{h}_{\tau,4} \equiv \beta_\tau \delta_\tau \begin{bmatrix} u_4(x_{4,1}; \tau) - u_4(x_{4,0}; \tau) \\ \vdots \\ u_4(x_{4,S_4-1}; \tau) - u_4(x_{4,0}; \tau) \end{bmatrix}. \quad (58)$$

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

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

Without incorporating this restriction the corresponding matrix needed in [Assumption DE2](#) would not be invertible.

Using the notation above we can rewrite eq. (57) in matrix form as

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

where $\mathbf{g}_{\tau,3,\Delta}$ is identified from the argument culminating in eq. (57) and the matrices $\mathbf{dF}_{\Delta,1}(z)$ and $\mathbf{dF}_{\Delta,1}(z')$ are identified since they are constructed from observed beliefs. Under [Assumption DE2](#) it follows that we can identify $\mathbf{h}_{\tau,4}$ as:

$$\mathbf{h}_{\tau,4} = \left(\mathbf{dF}_{\Delta,1}(z) - \mathbf{dF}_{\Delta,1}(z') \right)^{-1} \mathbf{g}_{\tau,3,\Delta}. \quad (59)$$

Therefore, the object $\beta_\tau \delta_\tau (u_4(x_4; \tau) - u_4(x_{40}; \tau))$ is identified for all $x_4 \in \mathcal{X}_4$. Next, note that

$$g_{\tau,3,1}(x_3, z) = u_3(x_3, 1; \tau) - u_3(x_3, 0; \tau) + \beta_\tau \delta_\tau \int (u_4(x_4; \tau) - u_4(x_{40}; \tau)) \mathbf{dF}_\Delta(x_4 | x_3, z), \quad (60)$$

³⁸These replace [Assumption D1](#) and [Assumption DA1](#) respectively.

since

$$\int u_4(x_{40}; \tau) dF_{\Delta}(x_4|x_3, z) = 0.$$

Since the last term in eq. (60) is now identified and $u_3(x_3, 0; \tau)$ is known, we can identify $u_3(x_3, 1; \tau)$ for all $x_3 \in \mathcal{X}_3$. ■

In general, we are only able to identify the product of the time-preference parameters and the period four payoff function $\beta_{\tau}\delta_{\tau}(u_4(x_4; \tau) - u_4(x_{40}; \tau))$. However, for the functional forms used in the application (where period 4 utility is given by eq. (20)) we can go further. In particular, for the application $(u_4(x_4; \tau) - u_4(x_{40}; \tau)) = \mathcal{C}(x_4) - \mathcal{C}(x_{40})$ which is identified (since $\mathcal{C}(\cdot)$ is known). Then, it follows that in the application $\beta_{\tau}\delta_{\tau}$ is separately identified from period three choices alone.

However, for completeness we consider identification without using the period 4 payoff specification in eq. (20). To do so, we examine period 2 choices. We start by applying the Hotz-Miller inversion to directly identify the function $g_{\tau,2,1}(\cdot)$:

$$g_{\tau,2,1}(x_2, z) \equiv u_2(x_2, 1; \tau) - u_2(x_2, 0; \tau) + \beta_{\tau}\delta_{\tau} \int h_{\tau}(x_3, z) dF_{\Delta}(x_3|x_2, z), \quad (61)$$

which is conceptually analogous to eq. (7) in the [Directly Observed Types](#) section. Analogous to the notation in that section, define

$$\begin{aligned} h_{\tau}(x_3, z) &\equiv \int v_{\tau}^*(x_3, \epsilon_3, z) dF(\epsilon_3) = \sum_{a \in \mathcal{A}_3} \int v_{\tau,3}(s_3, z, a, \delta_{\tau}) A_{\tau}(s_3, z, a, \tilde{\beta}_{\tau}\delta_{\tau}) dF(\epsilon_3) \\ &= \sum_{a \in \mathcal{A}_3} \underbrace{\int (u_3(x_3, a; \tau) + \epsilon_3(a)) A_{\tau}(s_3, z, a, \tilde{\beta}_{\tau}\delta_{\tau}) dF(\epsilon_3)}_{\tilde{h}_{\tau,1}(x_3, z, \tilde{\beta}_{\tau}\delta_{\tau})} + \sum_{a \in \mathcal{A}_3} \delta_{\tau} \int \int u_4(x_4; \tau) dF(x_4|x_3, z, a) A_{\tau}(s_3, z, a, \tilde{\beta}_{\tau}\delta_{\tau}) dF(\epsilon_3) \\ &= \tilde{h}_{\tau,1}(x_3, z, \tilde{\beta}_{\tau}\delta_{\tau}) + \beta_{\tau}^{-1} \sum_{a \in \mathcal{A}_3} \underbrace{\int \int \beta_{\tau}\delta_{\tau} (u_4(x_4; \tau) - u_4(x_{40}; \tau)) dF(x_4|x_3, z, a) A_{\tau}(s_3, z, a, \tilde{\beta}_{\tau}\delta_{\tau}) dF(\epsilon_3)}_{\tilde{h}_{\tau,2}(x_3, z, \tilde{\beta}_{\tau}\delta_{\tau})} + \delta_{\tau} u_4(x_{40}; \tau) \\ &= \tilde{h}_{\tau,1}(x_3, z, \tilde{\beta}_{\tau}\delta_{\tau}) + \beta_{\tau}^{-1} \tilde{h}_{\tau,2}(x_3, z, \tilde{\beta}_{\tau}\delta_{\tau}) + \delta_{\tau} u_4(x_{40}; \tau) \end{aligned}$$

where³⁹

$$\begin{aligned} v_{\tau,3}(s_3, z, a, \mathbf{d}_1) &\equiv u_3(x_3, a; \tau) + \epsilon_3(a) + \mathbf{d}_1 \int u(x_4; \tau) dF(x_4|x_3, z, a) \\ A_{\tau}(s_3, z, a, \mathbf{d}_2) &\equiv \mathbb{I} \left\{ a = \underset{j \in \mathcal{A}_3}{\operatorname{argmax}} v_{\tau,3}(s_3, z, j, \mathbf{d}_2) \right\}. \end{aligned}$$

To ease notation in what follows, define

$$\bar{h}_{\tau,\Delta,1}(x_2, \mathbf{d}) \equiv \int \left(\tilde{h}_{\tau,1}(x_3, z, \mathbf{d}) dF_{\Delta}(x_3|x_2, z) - \tilde{h}_{\tau,1}(x_3, z', \mathbf{d}) dF_{\Delta}(x_3|x_2, z') \right) \quad (62)$$

$$\bar{h}_{\tau,\Delta,2}(x_2, \mathbf{d}) \equiv \int \left(\tilde{h}_{\tau,2}(x_3, z, \mathbf{d}) dF_{\Delta}(x_3|x_2, z) - \tilde{h}_{\tau,2}(x_3, z', \mathbf{d}) dF_{\Delta}(x_3|x_2, z') \right) \quad (63)$$

(where we have suppressed the dependence on (z, z') for convenience). As before, we use variation in beliefs z to identify objects of interest. Using the notation defined above and taking differences of the identified

³⁹Note that $\sum_a \int \int \delta_{\tau} u_4(x_{40}; \tau) dF(x_4|x_3, z, a) A(s_3, z, a, \tilde{\beta}_{\tau}\delta_{\tau}) dF(\epsilon_3) = \delta_{\tau} u_4(x_{40}; \tau)$

function $g_{\tau,2}(\cdot)$ at two different points (z, z') :

$$\begin{aligned}
g_{\tau,2,\Delta}(x_2) &\equiv g_{\tau,2,1}(x_2, z) - g_{\tau,2,1}(x_2, z') \\
&= \beta_\tau \delta_\tau \int \left(h_\tau(x_3, z) dF_\Delta(x_3|x_2, z) - h_\tau(x_3, z') dF_\Delta(x_3|x_2, z') \right) \\
&= \beta_\tau \delta_\tau \left(\bar{h}_{\tau,\Delta,1}(x_2, \tilde{\beta}_\tau \delta_\tau) + \beta_\tau^{-1} \bar{h}_{\tau,\Delta,2}(x_2, \tilde{\beta}_\tau \delta_\tau) \right) \\
&= \beta_\tau \delta_\tau \bar{h}_{\tau,\Delta,1}(x_2, \tilde{\beta}_\tau \delta_\tau) + \delta_\tau \bar{h}_{\tau,\Delta,2}(x_2, \tilde{\beta}_\tau \delta_\tau)
\end{aligned} \tag{64}$$

The subsequent arguments now depend upon the type of agent. Suppose that agents are either sophisticated or consistent (so that $\tilde{\beta}_\tau \delta_\tau = \beta_\tau \delta_\tau$). This simplifies matters because then the $\bar{h}_{\tau,\Delta,j}(x_2, \beta_\tau \delta_\tau)$ functions are identified. This is because $\bar{h}_{\tau,\Delta,j}(x_2, \beta_\tau \delta_\tau)$ is a function of (a) $u_3(\cdot; \tau)$ and $\beta_\tau \delta_\tau (u_4(x_4; \tau) - u_4(x_4; \tau))$ which are both identified by [Lemma B1](#); (b) $\int u_4(x_{40}) dF_\Delta(x_3|x_2, z) = 0$ since $dF_\Delta(\cdot)$ is a signed measure integrating to zero; (c) $A_\tau(s_3, z, a, \tilde{\beta}_\tau \delta_\tau) = A_\tau(s_3, z, a, \beta_\tau \delta_\tau)$ and the latter is identified.⁴⁰

For consistent agents (for whom $\beta_C = 1$) the exponential parameter δ_C can be obtained directly by solving eq. (64) as long as $\bar{h}_{C,\Delta,1}(x_2, z, \delta_C) + \bar{h}_{C,\Delta,2}(x_2, z, \delta_C) \neq 0$. For sophisticated agents, the expression in eq. (64) is a linear equation in two unknowns $(\beta_S \delta_S, \delta_S)$. One can then separately identify β_S and δ_S by imposing an appropriate invertibility condition as below.

Specifically, we would require the existence of two points (x_2, x'_2) such that the 2×2 matrix below is invertible. We record this assumption below.

ASSUMPTION DE3 (Invertibility). *There exist (x_2, x'_2) such that the matrix*

$$\begin{bmatrix} \bar{h}_{S,\Delta,1}(x_2, \beta_S \delta_S) & \bar{h}_{S,\Delta,2}(x_2, \beta_S \delta_S) \\ \bar{h}_{S,\Delta,1}(x'_2, \beta_S \delta_S) & \bar{h}_{S,\Delta,2}(x'_2, \beta_S \delta_S) \end{bmatrix}.$$

is invertible.

eq. (64) yields two equations in two unknowns $(\beta_S \delta_S, \delta_S)$ that can be solved for uniquely under [Assumption DE3](#). However, this argument however does not extend to naïve agents because for such agents, eq. (64) is a non-linear equation of the unknowns $(\beta_{\tau_N} \delta_{\tau_N}, \delta_{\tau_N})$ and in general it is difficult to write down an easily interpretable condition that guarantees a unique solution. One possible way forward, as long as both naïve and sophisticated types exist is to assume that sophisticated and naïve agents have the same exponential parameter.

ASSUMPTION DE4. *Time-inconsistent agents have identical exponential parameters $\delta_S = \delta_N (\equiv \delta_I)$.*⁴¹

LEMMA B2 (Identification of Time Preferences and Period 2 Preferences). *Consider an agent solving eq. (1) and suppose that the [Basic Assumptions](#) hold with the modification that $z_t \equiv z$ which is no longer a state variable but instead a conditioning variable. Suppose that [Assumption DE1](#), [Assumption DE2](#), [Assumption DE3](#) and [Assumption DE4](#) hold. In addition, assume that*

1. **(Identification of δ_C)** $\bar{h}_{C,\Delta,1}(x_2, \delta_C) + \bar{h}_{C,\Delta,2}(x_2, \delta_C) \neq 0$
2. **(Identification of β_N)** $\bar{h}_{N,\Delta,1}(x_2, \delta_I) \neq 0$.

Then,

1. (δ_C, δ_I) are identified.
2. The time-preference parameters for sophisticated and naïve agents, (β_S, β_N) are identified.
3. $u_2(x_2, a; \tau)$ is identified for all types τ for $x_2 \in \mathcal{X}_2$.

⁴⁰Note that the argument for identification for $A_\tau(x_3, z, j, d)$ requires $d((u_4(x_4; \tau) - u_4(x_{40}; \tau))$ to be identified. This is only true for $d = \beta_\tau \delta_\tau$ from [Lemma B1](#). Thus part (c) is not identified for naïve agents.

⁴¹There 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) $g_{N,2,\Delta}$ is not 1-1 in δ_N and consequently δ_N need not be identified.

Proof. Under the assumptions above, δ_C is identified as

$$\delta_C = \frac{g_{C,2,\Delta}(x_2)}{\bar{h}_{C,\Delta,1}(x_2, \delta_C) + \bar{h}_{C,\Delta,2}(x_2, \delta_C)}.$$

Note that δ_C enters the functions $\bar{h}_{C,\Delta,j}(x_2, \delta_C)$ as the product $\delta_C(u_4(x_4; \tau_C) - u_4(x_{40}; \tau_C))$ which term is identified from [Lemma B1](#) – therefore all terms on the right hand side of the expression above are identified. Under the invertibility condition [Assumption DE 3](#) we can identify β_S and δ_I . Then,

$$\beta_N = \frac{g_{N,2,\Delta}(x_2) - \delta_I \bar{h}_{N,\Delta,2}(x_2, \delta_I)}{\delta_I \bar{h}_{N,\Delta,1}(x_2, \delta_I)}$$

Once the time-preference parameters are identified, the second term on the right hand side of eq. (61) is identified so we can identify the (normalized) period two utility functions for all types as

$$u_2(x_2, 1; \tau) - u_2(x_2, 0; \tau) = g_{\tau,2,1}(x_2, z) - \beta_\tau \delta_\tau \int h_\tau(x_3, z) dF_\Delta(x_3|x_2, z), \quad (65)$$

■

Since we restrict $\delta_\tau = \delta$ in the empirical application (in order to be parsimonious with parameters given our sample size), we can use this common discount factor assumption to identify β_S and β_N without imposing the invertibility condition [Assumption DE3](#). In particular, having identified the common parameter δ as outlined above from the choices of consistent agents, we can solve for β_S and β_N as

$$\beta_S = \frac{g_{S,2,\Delta}(x_2) - \delta \bar{h}_{S,\Delta,2}(x_2, \beta_S \delta)}{\delta \bar{h}_{S,\Delta,1}(x_2, \beta_S \delta)} \quad (66)$$

$$\beta_N = \frac{g_{N,2,\Delta}(x_2) - \delta \bar{h}_{N,\Delta,2}(x_2, \delta)}{\delta \bar{h}_{N,\Delta,1}(x_2, \delta)} \quad (67)$$

as long as the denominators are not equal to zero.

In addition, for the empirical application, we can exploit the specification of the static utility functions and the common δ parameter to simplify and extend the identification argument above. In particular, recall from the penultimate lemma that for our empirical application $\beta_\tau \delta_\tau$ and $u(x_4; \tau) - u(x_{40}; \tau)$ are separately identified. Further, from the specifications in [eqs. \(21\)](#) and [\(22\)](#) we see that $u(x_2, 1; \tau) - u(x_2, 0; \tau) = -p_r \mathbb{I}\{a_1 = b\} + \phi_\tau(v)$ which is identified (note $\phi_\tau(\cdot)$ is identified from the period 3 utility identification argument). Therefore, it only remains to separately identify δ_τ and β_τ . Since we assume that the $\delta_\tau = \delta$ and that type identities are known, then we can recover β_τ directly as $\beta_\tau \delta / \delta$. Once the time-preference parameters are identified we can recover period 2 utility functions as above (since all terms on the right hand side of eq. (65) are identified).

B.1.1 Identification of Period One Preferences

The only remaining unidentified objects are now the period one payoff functions. There is a sharp distinction in period one (relative to the later periods) regarding direct type identification for individual agents. In particular, types are not observed for all agents in period one 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 uninformative of their type. To compound the problem, these agents make no further decisions.

We approach this problem by first noting that the key object required for the inversion argument is the type-specific choice probability $\mathbb{P}_\tau(a_1|x_1, z)$.⁴² As in [Section 3.2.2](#) for unobserved types we adapt the insights from [Kasahara and Shimotsu \(2009\)](#) by imposing a set of exclusion restrictions. We use the structure imposed by the markov assumptions and the exclusion restrictions on the identified matrix of choice probabilities ($\mathsf{P}_{r,v}$,

⁴²For $t > 1$ 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 DE1](#)).

defined below) 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 without loss of generality since these variables are not used for identification).

We consider here the case where $r = 1$, that is sub-population that expressed time-inconsistent preferences in the baseline survey. As noted before, we assume that both naïve and sophisticated types (as well as consistent types) exist. In order to proceed, we first restate [Assumption U1](#) in terms of the variables in the empirical application and only for periods one and two.

ASSUMPTION DE5 (Exclusion Restrictions). *1. Conditional upon type, the MEL survey response r is uninformative about choice $\mathbb{P}_\tau(a_t|x_t, z, r) = \mathbb{P}_\tau(a_t|x_t, z)$ for $t = 2$ and $\mathbb{P}_\tau(a_1, x_1|r, z) = \mathbb{P}_\tau(a_1, x_1|z)$. 2. The transition probabilities do not vary by type and are independent of r : $\mathbb{P}_\tau(x_{t+1}|x_t, a_t, z, r) = \mathbb{P}(x_{t+1}|x_t, a_t, z)$ 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 first use the Markov property and the assumption above to obtain

$$\mathbb{P}(a_1, a_2, x_1, x_2|r, z) = \sum_{\tau \in \mathcal{T}_r} \pi_\tau(r) \mathbb{P}_\tau(a_1, x_1|z) \mathbb{P}_\tau(a_2|x_2, z) \mathbb{P}(x_2|x_1, a_1, z), \quad (68)$$

and as in eq. (11) we define the quantities $\mathbf{F}_{x_1, x_2, r}^{a_1, a_2}$, $\mathbf{F}_r^{a_1, x_1}$ and $\mathbf{F}_{x_2, r}^{a_2}$ as functions of the type-specific choice probabilities (the objects of interest) – $\lambda_{x_2}^{a_2, \tau} \equiv \mathbb{P}_\tau(a_2|x_2, z, v)$, and $\lambda^{a_1, x_1, \tau} \equiv \mathbb{P}_\tau(a_1, x_1|z, v)$.

We then use eq. (68) to express the identified matrix $\mathbf{P}_{1,r}^{a_1, a_2}$ in terms of the objects of interest. In this case, since there are only three types we only need a 3×3 matrix and two points in the state space (in each period) for the identification arguments. Let (x_t^1, x_t^2) denote these elements for $t \in \{1, 2\}$

$$\begin{aligned} \mathbf{P}_{1,r}^{a_1, a_2} &= \begin{pmatrix} 1 & \mathbf{F}_{x_2, r}^{a_2} & \mathbf{F}_{x_2, r}^{a_2} \\ \mathbf{F}_r^{a_1, x_1} & \mathbf{F}_{x_1, x_2, r}^{a_1, a_2} & \mathbf{F}_{x_1, x_2, r}^{a_1, a_2} \\ \mathbf{F}_r^{a_1, x_2} & \mathbf{F}_{x_2, x_1, r}^{a_1, a_2} & \mathbf{F}_{x_2, x_1, r}^{a_1, a_2} \end{pmatrix}, \\ \mathbf{L}_1^{a_1} &= \begin{pmatrix} 1 & \lambda^{a_1, x_1, \tau_C} & \lambda^{a_1, x_2, \tau_C} \\ 1 & \lambda^{a_1, x_1, \tau_N} & \lambda^{a_1, x_2, \tau_N} \\ 1 & \lambda^{a_1, x_1, \tau_S} & \lambda^{a_1, x_2, \tau_S} \end{pmatrix} \quad \mathbf{L}_2^{a_2} = \begin{pmatrix} 1 & \lambda_{x_2}^{a_2, \tau_C} & \lambda_{x_2}^{a_2, \tau_C} \\ 1 & \lambda_{x_1}^{a_2, \tau_N} & \lambda_{x_2}^{a_2, \tau_N} \\ 1 & \lambda_{x_1}^{a_2, \tau_S} & \lambda_{x_2}^{a_2, \tau_S} \end{pmatrix}. \end{aligned} \quad (69)$$

Note that the difference in notation across sections reflects the fact that (a) we are suppressing dependence on v and (b) 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. $\mathbf{L}_2^{a_2}$ 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 DE6 (Invertibility). *The matrix \mathbf{L}_2^1 defined in eq. (69) is invertible for two pairs (x_2', x_2'') and (x_2^*, x_2^{**}) . Further, we assume that for $x_1 \in \{x_1^1, x_1^2\}$: (a) (x_2', x_2'') such that $\mathbb{P}(x_2'|x_1, a_1 = b, z) > 0, \mathbb{P}(x_2''|x_1, a_1 = b, z) > 0$; (b) (x_2^*, x_2^{**}) such that $\mathbb{P}(x_2^*|x_1, a_1 = c, z) > 0, \mathbb{P}(x_2^{**}|x_1, a_1 = c, z) > 0$.*

The invertibility assumption formalizes the precise sense in which different types must behave sufficiently differently at two points in the state space. The second condition ensures that the (directly identified) matrix $\mathbf{P}_{1,r}^{a_1, a_2}$ is well defined. Since the 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. We can now state the result for identification of first-period choice probabilities and preferences as well.

LEMMA B3 (Identification of Period One Preferences and Type Distribution). *Consider an agent solving eq. (1) and suppose that the conditions in [Lemma B2](#) hold. In addition, suppose that [Assumption DE5](#) and [Assumption DE6](#) hold. Then,*

1. *First period preferences $u_1(x_1, a; \tau)$ are identified $\forall x_1 \in \mathcal{X}_1 \ \forall a \in \mathcal{A}_1 \ \forall \tau \in \mathcal{T}$*
2. *The type probabilities (conditional on the MEL response r) $\pi_\tau(r)$ are identified.*

Proof. Given the notation introduced we can write

$$P_{1,1}^{\{b,c\},1} = (L_1^b)' V_1 L_2^1,$$

where we have set $(r, a_1, a_2) = (1, \{b, c\}, 1)$ so we are examining the sub-population that (a) expressed inconsistency in the survey and (b) purchased an ITN and retreated it in period 2. The decision to combine contracts (b, c) in the first-period action for the proof is to ensure that we can examine both naïve and sophisticated types in period two.⁴³

Next, evaluating V_r at $r = 1$: $V_1 = \text{Diag}(0, \pi_{\tau_N}(1), \pi_{\tau_S}(1))$. Note that the first entry is zero since by definition (for this section on directly observed types) consistent agents cannot have $r = 1$. Using the invertibility of L_2^1 (from [Assumption DE6](#) above)

$$(L_2^1)^{-1} P_{1,1}^{\{b,c\},1} = (L_1^b)' V_1,$$

where the left-hand side is identified (because the elements of L_2^1 are period-two type-specific choice probabilities which are identified by assumption and $P_{1,1}^{\{b,c\},1}$ is directly identified). The right hand side is equal to

$$\begin{bmatrix} 0 & \pi_N(1) & \pi_S(1) \\ 0 & \pi_N(1)\lambda_N^{a_1 x_1^1} & \pi_S(1)\lambda_S^{a_1 x_1^1} \\ 0 & \pi_N(1)\lambda_N^{a_1 x_1^2} & \pi_S(1)\lambda_S^{a_1 x_1^2} \end{bmatrix},$$

so that the type-frequencies $\{\pi_\tau(1) : \tau \in \{N, S\}\}$ are identified as well as the first-period type-specific choice probabilities $\mathbb{P}_{\tau_N}(a_1 = b|x_1, z)$ (and also consequently $\mathbb{P}_{\tau_N}(a_1 = n|x_1, z)$) since naïve agents can only choose among (b, n) and $\mathbb{P}_{\tau_S}(a_1 = c|x_1, z)$ (and likewise $\mathbb{P}_{\tau_S}(a_1 = n|x_1, z)$).

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

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

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

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.

B.2 Unobserved Types

As noted above, although survey responses 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 but are unrelated to time preference (see [Appendix page OA-2](#) for a longer discussion). For this reason we consider a model with unobserved types. The arguments for identification here are identical to those in [Section 3.2.2](#) once we have

⁴³Recall that naïve types are only choosing between contract b and no-purchase and similarly sophisticated types are only choosing between c and no-purchase (by assumption and only for this section). Therefore, $\mathbb{P}_N(a_1 \in \{b, c\}|x_1) = \mathbb{P}_N(a_1 = b|x_1)$ and $\mathbb{P}_S(a_1 \in \{b, c\}|x_1) = \mathbb{P}_S(a_1 = c|x_1)$.

accounted for the relevant differences between the empirical application and the general model (i.e. those mentioned at the start of the section).

LEMMA B4. *Let $t = 2$ and fix (r, v) . Suppose [Assumption U1](#) holds with the modifications that $z_t \equiv z$; [Assumption U2](#) holds with $M_{r,v} = 3$ and the optimal decision process is Markov. Then, the type-specific choice probabilities $\{\mathbb{P}_\tau(a_t|x_t, z, v)\}_{\tau \in \mathcal{T}; t \in \{1, 2, 3\}}$ for $(x_t, v) \in \mathcal{X}_t \times \mathcal{Z} \times \mathcal{V}$ are identified for $t > 1$. For $t = 1$, fix r and assume that the previous conditions hold. Then the type-specific choice probabilities for period 1 are also identified.*

Proof. The proof is a direct application of [Lemma 3](#). 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 method avoids strong assumptions on the type-proxy and relies instead on different types behaving sufficiently differently. The second method is relatively straightforward but involves placing stronger assumptions on the type proxy r . We discuss each strategy in turn.

The first 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 $\pi_\tau(r)/\pi_\tau(r')$. To motivate the weakest condition, we start with a stronger set of sufficient conditions. Suppose that the set of agents with responses $r = 1$ are most likely to be sophisticated inconsistent agents and least likely to be time-consistent agents. Second, the set of agents with $(r = 0)$ are most likely to be time-consistent agents and least likely to be sophisticated inconsistent agents. This implies an ordering on the ratios: $\{\pi_C(r)/\pi_C(r') \geq \pi_N(r)/\pi_N(r') \geq \pi_S(r)/\pi_S(r')\}$ for $r = 0$ and $r' = 1$. This ordering then guarantees the identification of type-identities (i.e. the labelling of types). While the direct assumptions on the probabilities themselves may 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 UE1 (Monotone Likelihood Ratio Like Property). *For some $r \neq r'$, the three ratios*

$$\left\{ \frac{\pi_C(r)}{\pi_C(r')}, \frac{\pi_N(r)}{\pi_N(r')}, \frac{\pi_S(r)}{\pi_S(r')} \right\}$$

can be strictly ordered ex-ante.

Under this additional assumption the type-identities are identified.

PROPOSITION 5. *Suppose that [Lemma B4](#) holds and that [Assumption UE1](#) holds. Then, type identities are identified.*

Proof of Proposition 5

Proof. [Lemma B4](#) identifies the type probabilities $\pi_\tau(r)$. We can then form $\{\pi_\tau(r)/\pi_\tau(r')\}_{\tau \in \mathcal{T}}$. By The monotonicity assumption, the strict ordering of these ratios allows us to identify the type-identity for each τ . ■

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 eq. (18) non-parametrically is a tall order in our empirical application.

B.3.1 Alternative Theorem for Type Identities

The result above is useful when we have sufficient confidence in the ability of the proxy to distinguish between different types of agent. We next discuss assumptions that instead rely on sufficiently different

behavior across different types. The argument follows the broad outlines of the discussion in [Section 3.2.3](#) but needs to account for (a) no action being taken in the terminal state – so that $\beta_\tau\delta$ is not identified in the first two steps of the backward induction (unlike the argument in [Section 3.2.3](#)), (b) the arguments need to account for the constancy of z across time.

Our starting point is the identified type-specific choice probabilities $\mathbb{P}_\tau(a_t|x_t, z)$ for $t = 2, 3$ but for which the type identity itself (i.e. whether the agent is consistent or naïve or sophisticated) is unknown. [Lemma B1](#) can be applied without knowledge of the type identity to identify period 3 utilities $u(x_3, a; \tau)$ and the product $\beta_\tau\delta(u(x_4; \tau) - u(x_{40}; \tau))$.

Next, using the Hotz-Miller inversion and the differencing argument as earlier on p.[OA-6](#) we can identify the function $g_{\tau, \Delta, 2}(x_2)$ defined in eq. (64). Define the 2×1 vector $\mathbf{g}_{\tau, \Delta, 2}(x_2, x'_2) \equiv [g_{\tau, \Delta, 2}(x_2), (g_{\tau, \Delta, 2}(x'_2))']'$ for two points $(x_2, x'_2) \in \mathcal{X}_2 \times \mathcal{X}_2$. Next, define the 2×2 identified matrix

$$\mathbf{K}_\tau(x_2, x'_2) = \begin{pmatrix} \bar{h}_{\tau, \Delta, 1}(x_2, \beta_\tau\delta) & \bar{h}_{\tau, \Delta, 2}(x_2, \beta_\tau\delta) \\ \bar{h}_{\tau, \Delta, 1}(x'_2, \beta_\tau\delta) & \bar{h}_{\tau, \Delta, 2}(x'_2, \beta_\tau\delta) \end{pmatrix} \quad (70)$$

where we have suppressed dependence on (z, z') for readability and the functions $\bar{h}_{\tau, \Delta, j}(\cdot)$ are defined in eq. (62).⁴⁴ If we assume this matrix is invertible, we can identify the 2×1 vector

$$\hat{\mathbf{d}}_\tau(x_2, x'_2) \equiv [\mathbf{K}_\tau(x_2, x'_2)]^{-1} \mathbf{g}_{\tau, \Delta, 2}(x_2, x'_2) \quad (71)$$

We then identify types by examining $\hat{\mathbf{d}}$ for different pairs of points (x_2, x'_2) . For consistent and sophisticated types, $\hat{\mathbf{d}}$ will always be a constant – equal to $(\beta_\tau\delta, \delta)$ – for all pairs (x_2, x'_2) . This will not be the case for the naïve types if the following is true: (a) their (period two) view of the trade-off between period three and period four differs depending upon whether they use δ_N or $\beta_N\delta_N$ 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 condition to hold:

ASSUMPTION UE2 (Invertibility and Variation over State Space). *Define the matrix*

$$\tilde{\mathbf{K}}_\tau(x_2, x'_2) \equiv \begin{pmatrix} \bar{h}_{\tau, \Delta, 1}(x_2, \tilde{\beta}_\tau\delta) & \bar{h}_{\tau, \Delta, 2}(x_2, \tilde{\beta}_\tau\delta) \\ \bar{h}_{\tau, \Delta, 1}(x'_2, \tilde{\beta}_\tau\delta) & \bar{h}_{\tau, \Delta, 2}(x'_2, \tilde{\beta}_\tau\delta) \end{pmatrix}. \quad (72)$$

There exist distinct points $(x_2, x'_2, x''_2) \in \mathcal{X}_2 \times \mathcal{X}_2 \times \mathcal{X}_2$ and $(z, z') \in \mathcal{Z} \times \mathcal{Z}$ such that the inverses in the display below exist and

$$\left[(\mathbf{K}_N(x_2, x'_2))^{-1} \tilde{\mathbf{K}}_N(x_2, x'_2) - (\mathbf{K}_N(x_2, x''_2))^{-1} \tilde{\mathbf{K}}_N(x_2, x''_2) \right] \begin{bmatrix} \beta_N\delta \\ \delta \end{bmatrix} \neq 0 \quad (73)$$

PROPOSITION 6 (Alternative Type Identification Result). *Suppose that the conditions for [Lemma B4](#) hold and that [Assumption UE2](#) holds. Then, type identities are identified.*

Proof. The proof follows by examining the behavior of the identified objects $\hat{d}_\tau(x_2, x'_2)$ and $\hat{d}_\tau(x_2, x''_2)$ as (x_2, x'_2, x''_2) range over \mathcal{X}_2^3 . If these objects are the same regardless of the choice of triplet, then $\tau \in \{S, C\}$ and

$$\hat{d}_\tau(x_2, x'_2) = \hat{d}_\tau(x_2, x''_2) = (\beta_\tau\delta, \delta)',$$

and we can distinguish between sophisticated and consistent types by examining whether the two elements of the vector are equal. Next, observe that for naïve types by [Assumption UE2](#) there exist points (x_2, x'_2, x''_2)

⁴⁴Note that even though $\beta_\tau\delta$ is not currently identified, it only enters the $\bar{h}_{\tau, \Delta, j}(\cdot)$ functions as the product $\beta_\tau\delta(u(x_4; \tau) - u(x_{40}; \tau))$ and this object identified (by [Lemma B1](#)). Thus the $\bar{h}_{\tau, \Delta, j}(\cdot)$ functions are identified.

such that $\widehat{d}_\tau(x_2, x'_2) \neq \widehat{d}_\tau(x_2, x''_2)$. To see this,

$$\begin{aligned}\widehat{d}_\tau(x_2, x'_2) &= (\mathsf{K}_\tau(x_2, x'_2))^{-1} \widetilde{\mathsf{K}}_\tau(x_2, x'_2) \left(\widetilde{\mathsf{K}}_\tau(x_2, x'_2) \right)^{-1} \mathsf{g}_{\tau, \Delta, 2}(x_2, x'_2) \\ &= (\mathsf{K}_\tau(x_2, x'_2))^{-1} \widetilde{\mathsf{K}}_\tau(x_2, x'_2) \begin{bmatrix} \beta_\tau \delta \\ \delta \end{bmatrix}.\end{aligned}$$

■

C Inversion Argument

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 Assumption [B](#) for the argument (but do not need the exclusion restriction). The probability that an agent chooses action 0 is

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

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

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

Next, define

$$\begin{aligned} \hat{u}_1 &\equiv u(x_2, 1) - u(x_2, 0) + \beta \delta \int v^*(s_3) dF_{\Delta, 1}(s_3|s_2), \\ \hat{u}_2 &\equiv u(x_2, 2) - u(x_2, 0) + \beta \delta \int v^*(s_3) dF_{\Delta, 2}(s_3|s_2), \end{aligned}$$

and as usual, the signed measure is defined as

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

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

$$\begin{aligned} \mathbb{P}(a_2 = 0|x_2) &= \mathbb{P}(\epsilon(0) - \hat{u}_1 \geq \epsilon(1), \epsilon(0) - \hat{u}_2 \geq \epsilon(2)|x_2) \\ \mathbb{P}(a_2 = 1|x_2) &= \mathbb{P}(\epsilon(0) - \hat{u}_1 \leq \epsilon(1), \epsilon(1) + (\hat{u}_1 - \hat{u}_2) \geq \epsilon(2)|x_2). \end{aligned}$$

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

$$\mathbb{P}(\epsilon(0) - \hat{u}_1 \geq \epsilon(1), \epsilon(0) - \hat{u}_2 \geq \epsilon(2)|x_2) - \mathbb{P}(\epsilon(0) - u_1^* \geq \epsilon(1), \epsilon(0) - u_2^* \geq \epsilon(2)|x_2) = 0 \quad (74)$$

$$\mathbb{P}(\epsilon(0) - \hat{u}_1 \leq \epsilon(1), \epsilon(1) + (\hat{u}_1 - \hat{u}_2) \geq \epsilon(2)|x_2) - \mathbb{P}(\epsilon(0) - u_1^* \leq \epsilon(1), \epsilon(1) + (u_1^* - u_2^*) \geq \epsilon(2)|x_2) = 0. \quad (75)$$

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

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

and as long as $dF(\cdot|x_2)$ had strictly positive measure on all of \mathbb{R}^3 , the equality in eq. (74) 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 (eq. 75) cannot hold because

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

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

D Correlates of Actions

In this Appendix we illustrate reduced-form patterns in the data, with a special attention to predictors of purchase and re-treatment. In Table OA-1 below we look at correlations between purchase and re-treatment decisions and a list of predictors that include proxies for household socio-economic status, beliefs about the protective power of nets and re-treatment, recent malaria episodes, and responses to the inter-temporal choice questions. In column 1 we show that the only strong (and statistically significant) predictor of purchase is malaria exposure: take up increases by 21 percentage points if either someone tested positive to malaria based on the RDTs conducted by our study team, or the respondent reported at least one malaria episode in the six months that preceded the interview. An increase in the perceived protective power of bed nets (relative to no nets) or ITNs (relative to untreated nets) predicts a *decrease* in demand, although neither estimate is significant at standard levels and the latter is close to zero. Note that this finding does not imply that beliefs are measured with error, or that they do not matter for demand. In fact, demand also depends on discount factors, and indeed our model rationalizes low demand despite high perceived benefits as indication of impatience (especially low β).

	(1) Any net purchased $t = 1$	(2) Choose net + retr. bundle $t = 1$	(3) First re-treatment $t = 2$	(4) Second re-treatment $t = 3$
Purchased ITN + 2 retreatments bundle			0.461*** (0.068)	0.643*** (0.090)
Pr(Malaria no net)–Pr(Malaria untreated net)	-0.105 (0.101)	0.123 (0.156)	0.124 (0.112)	-0.018 (0.115)
Pr(Malaria untreated net)–Pr(Malaria ITN)	-0.037 (0.118)	0.102 (0.181)	0.220 (0.147)	0.274*** (0.096)
ln(monthly income per person)	-0.055 (0.035)	0.000 (0.051)	-0.001 (0.031)	-0.059 (0.039)
Asset Index (First Principal Component)	-0.017 (0.015)	0.007 (0.026)	0.014 (0.020)	0.009 (0.015)
Any malaria episode last 6mts (reported or RDT)	0.209*** (0.036)	-0.020 (0.070)	0.004 (0.035)	0.023 (0.035)
Nets owned at baseline	0.044 (0.058)	-0.088 (0.076)	0.017 (0.053)	0.049 (0.047)
PC Costs malaria episodes last 6mts>Rs 500	-0.032 (0.073)	-0.157 (0.100)	0.061 (0.093)	-0.037 (0.103)
HH. head had any formal schooling	-0.022 (0.053)	0.101 (0.070)	0.078* (0.041)	0.004 (0.067)
Intertemporal choices: any preference reversal	0.037 (0.051)	0.060 (0.073)	-0.298*** (0.092)	-0.056 (0.082)
Intertemporal choices: always chooses earlier payoff	-0.024 (0.045)	0.059 (0.075)	-0.068 (0.059)	-0.007 (0.051)
Any preference reversal \times Purchased ITN+2R bundle			0.267*** (0.097)	0.097 (0.096)
Constant	0.910*** (0.305)	0.403 (0.504)	0.339 (0.260)	0.582* (0.296)
Observations	549	280	270	275
R-squared	0.068	0.029	0.402	0.485
Clusters	47	42	42	42

Table OA-1: Predictors of purchase and re-treatment

In column 2 we show that none of the regressors is significant at standard levels when we predict the choice of the bundle contract (net + two re-treatment, or contract c), conditional on purchase. Here we only note that beliefs about risk reduction from nets and ITNs predict about a 10 percentage point increase in the probability of choosing the bundle, but both are estimated very imprecisely.

When we look at predictors of re-treatment after six months ($t = 2$ in our model) and twelve months ($t = 3$), most coefficients are again not significant at standard levels, but some interesting patterns emerge. First, consistent with the results in Table 3 of the paper, re-treatment rates are substantially higher among buyers of the bundle product c —46 and 64 percentage points higher in $t = 2$ and $t = 3$, respectively (both p-values > 0.01 when we test the null of equality). Second, households are significantly more likely to re-treat when their perceived risk reduction from re-treatment is higher: in both time periods the slope is very large (0.22 and 0.27), although only in the second period it is significant (p-value < 0.01) at conventional levels. Note that this is not a result of learning after purchasing from our program, given that beliefs were measured at baseline, before the sales were conducted. Third, at $t = 2$ (although not at $t = 3$) the event that the inter-temporal choices included at least one preference reversal (the standard indicator of hyperbolic preferences that we include in our type signal r) predicts a very large and significant decrease in the probability of re-treatment ($\hat{\beta} = -0.298$, $p < 0.01$), among households who did not choose the ‘commitment’ bundle (c).⁴⁵ This is broadly consistent with present-bias playing a relevant role in re-treatment decisions, although once again this reduced form approach does not allow us to clearly disentangle the role of beliefs and household characteristics from that of (unobserved) time preference parameters.

⁴⁵ Among household who purchased the bundle, the predicted change in the probability of re-treatment is $-0.298 + 0.267 = -0.031$, which is close to zero and not significant at standard levels ($p = 0.4723$).

E Variable Description

We start by summarizing the timing of the data collection, which guides several of our choices in the construction of the variables used in the estimation. Recall that we denote the key time periods as $t = 1, 2, 3, 4$, where $t = 1$ corresponds to the time of ITN sales, $t = 2, 3$ to the two re-treatments, and $t = 4$ to the endline survey. However, in practice several variables were observed at baseline ($t = 0$).

The following figure illustrates the timeline. In what follows we assume that each two subsequent periods are separated by six months, which is approximately correct.

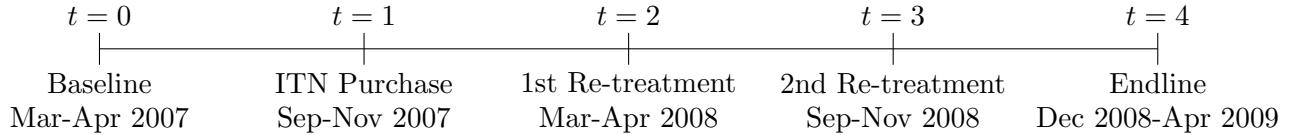


Figure OA-1: Timeline

Malaria: We define malaria cases at the household level, as a function of the information available to the household in each time period. Below we describe which information was collected in each period, and how we reconstruct the malaria indicators h_t used in the model for each period.

At baseline ($t = 0$), rapid diagnostic blood tests (RDTs) were used to measure malaria *prevalence*, that is, the fraction of individuals with ongoing malaria episodes at the time of the measurement, regardless of severity. Individuals targeted for blood tests included all pregnant women, children under the age of five (U5) and their mothers, and one randomly selected adult (age 15-60). For every household member we also recorded malaria *incidence* (that is, the number of cases) during the previous six months. Unlike prevalence, incidence was not measured by our research team with RDTs, and only relied on respondents' reports. In [Tarozzi et al. \(2014\)](#) we show that although respondent's reports of recent cases and RDTs were strongly correlated, a large majority of infections detected by RDTs were not reported. This suggests that most cases were asymptomatic, and thus not severe enough to cause loss of income. In [Tarozzi et al. \(2014\)](#) we show that both RDT-based prevalence and respondent-reported incidence were predictive of ITN purchase. Incidence (self-reported) and prevalence (from RDTs) were also recorded at endline ($t = 4$), although at this time *all* members were targeted for testing. Finally, at the time of the first re-treatment ($t = 2$) RDTs were not used and we only recorded how many members had malaria after the ITN sales, that is, between $t = 1$ and $t = 2$. Malaria status, one of the key state variables in the model, was thus constructed in each period as follows:

$t = 1$ Data on malaria cases were not collected at this time, and so we use, as malaria indicator, information from the baseline survey ($t = 0$). We construct a binary malaria indicator ($h_1 = 1$) if either someone in the household was found to be positive, or if someone was reported as having had malaria in the six months before the baseline survey.

$t = 2$ At this time we only have malaria incidence as reported by the respondent, so $h_2 = 1$ if any individual was reported as having been sick with malaria between $t = 1$ and $t = 2$, and $= 0$ otherwise.

$t = 3$ No information on malaria incidence or prevalence was collected at this time, and so we use data from the endline survey at $t = 4$ and we set $h_3 = 1$ if either someone in the household was found to be positive, or if someone was reported as having had malaria in the six months before the survey, and $= 0$ otherwise.

Expected Cost of Malaria Episodes: Forward-looking agents consider the expected cost of future malaria episodes when taking decisions about bed net purchases and re-treatment. In the empirical application we use the median monetary cost of a malaria episode reported by the respondent at baseline, equal to Rs. 386. This choice is conservative in the sense that the use of alternative measures of malaria costs (such as the expected costs of a malaria episode elicited in our survey, or the inclusion of estimates of lost earnings due to illness) lead to greater estimated present bias. Monetary costs take into account both expenses for doctor's visits and treatment as well as any wages paid to labor hired to replace a sick worker.

Income: Annual total household income was recorded both at baseline ($t = 0$) and endline ($t = 4$). At baseline the respondent was asked the following question:

Now please think about the income of everybody in your household from last year this time. Think about income from wages, sales, business, or any other source from each member of your household. Also include the value of any in-kind earnings. Now think about the period of time from today until last year this time. Can you please tell us the total income that your household has been able to earn during this period?

At endline, the question was phrased similarly, but respondents were allowed to indicate a range—instead of a single figure—in case they were not certain. More than half of respondents (57%) indicated a single figure, and even among those who indicated a range, its width was usually not very wide.⁴⁶ For these observations we assume that income was equal to the mid-point of the range.

Because income was not observed at $t = 1, 2, 3$, we impute it using simple interpolations between $t = 0$ and $t = 4$, with weights proportional to the distance in time between the two endpoints. As indicated in Figure OA-1, there are about two years between baseline and endline surveys, and approximately six months between each two periods. We also need to convert yearly into six-month household income. Letting Y_t denote yearly income in our data, six-month income y_t at times $t = 1, 2, 3$ is thus imputed as

$$y_t = \frac{(1 - t/4)Y_0 + (t/4)Y_4}{2}.$$

Note that respondents were asked to report *actual* total household income, that is, *net* of any earnings lost because of malaria episodes. Our data do not include beliefs about the joint distribution of earnings and malaria incidence, and so for simplicity we assume that malaria status and *gross* income are stochastically independent. For this reason, and taking into account that the cost of malaria episodes was recorded using a six-month recall, we set the two values Y_0 and Y_4 to be equal to total household yearly income as reported by the respondent *plus* twice the six-month income loss due to malaria episodes. We recorded earning losses due to malaria episodes due to reduced labor supply for both the sick individual and any care-taker.

For tractability, we assume that income in each period can only take two values ('High' or 'Low'). Let y_t^H and y_t^L denote median income conditional on y_t being above or below the overall time-specific median, respectively. For each household i , we then replace income y_t , $t = 0, 1, 2, 3, 4$ as constructed above with either y_t^H or y_t^L , depending on whether household income is above or below the median.

Subjective Transition Probabilities for Income In addition to the subjective beliefs about the protective power of nets, at baseline we also recorded beliefs about yearly income in the 12 months following the interview. The question was as follows:

Now please think about the income of everybody in your household. Think about income from wages, sales, business, or any other source from each member of your household. Also include the value of any in-kind earnings. Now think about the next agricultural year (April 2007 to March 2008): in your opinion, in the best possible situation, what is the largest amount of total income that your household may be able to earn during the next agricultural year?

A similar question was then asked about the smallest possible value, and then the surveyor would ask about the perceived probability that income would be below or above the midpoint between largest and smallest value. We used this information to generate a household-specific distribution, as described below.

Let the lower and upper bound of the reported range be denoted by l and u , respectively, and let q denote the reported probability that realized income will be smaller than the average of the lower and upper bounds. We follow Guiso et al. (2002, Fig. 1) in assuming that the distribution is 'triangular', so that the density function will take a shape like in the example shown in Figure OA-2, where we have assumed that $q < 1/2$.

⁴⁶Among the respondents who reported a range, the median ratio between the upper and lower bounds of the range was 1.2, while the 90th percentile was 1.5.

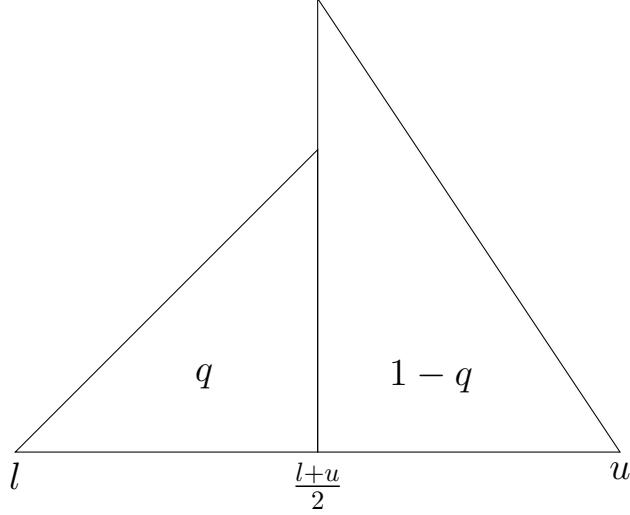


Figure OA-2: Triangular Distribution for Expected Income

One can show that the density function is described by the following expression

$$f(y) = \mathbb{I}\left\{y \leq \frac{l+u}{2}\right\} \left[8q \frac{y-l}{(u-l)^2} \right] + \mathbb{I}\left\{y > \frac{l+u}{2}\right\} \left[8(1-q) \frac{u-y}{(u-l)^2} \right],$$

while the cumulative distribution function is:

$$\begin{aligned} F(y) &= \mathbb{I}\left\{y \leq \frac{l+u}{2}\right\} \left[\frac{4q}{(u-l)^2} (y-l)^2 \right] \\ &+ \mathbb{I}\left\{y > \frac{l+u}{2}\right\} \left\{ q + \frac{8(1-q)}{(u-l)^2} \left[uy - \frac{y^2}{2} - \frac{u(u+l)}{2} + \frac{1}{2} \left(\frac{u+l}{2} \right)^2 \right] \right\}. \end{aligned} \quad (76)$$

We assume that the household-specific distribution of income described above remains constant over time. Although strong, this assumption is reasonable given our data, where we find that income at baseline is strongly correlated with income at endline ($\rho = 0.48$) and especially the mid-point of the distribution of one-year ahead income ($(l+u)/2$ in equation (76), $\rho = 0.84$). In addition, the cross-sectional distributions at baseline (y_0) and endline (y_4) appear overall quite similar, although mean and median are about 10% larger at endline.

We maintain the assumption that income and malaria are independent. Recall that this does *not* mean that malaria episodes have no monetary costs. In the model, agents interpret l and u as the upper and lower bound of predicted income, and then consider that any malaria episodes—whose likelihood depend on bed net ownership and re-treatment—will lead to monetary costs that will reduce consumption.

Recall that in the estimation we dichotomize income, with income in each period set to be equal to median income below or above the period-specific median, depending on whether actual reported income is below or above the median, respectively. Consistent with this, for each household we calculate subjective transition probabilities of income defined over two values. So, if y_0^H and y_0^L denote median income at baseline conditional on it being above or below the baseline median y_0^m , for household i we calculate the (stationary) transition probabilities as $\Pr_i(y_{i,t+1} = y_0^L) = \Pr_i(y_{i,t+1} \leq y_0^m) = F_i(y_0^m)$; $t = 1, 2, 3$, where $F_i(\cdot)$ is the respondent-specific CDF in equation (76). Lastly, this also implies that $\Pr_i(y_{i,t+1} = y_0^H) = 1 - F_i(y_0^m)$.

Malaria Beliefs. We use the baseline elicited beliefs about the likelihood of malaria summarized in Section 4.1. Since the beliefs were elicited for a horizon of one year we modify them to be reflective of a six-month horizon, so that for a subjective one-year probability of p_{12} we calculate the six-month probability as $p_6 = 1 - \sqrt{1 - p_{12}}$, where we have assumed for simplicity that malaria risk is equal and independent in the two halves of the year. In order to avoid imputing $p_6 = 1$ (certainty of malaria within six months) when the respondent

chooses p_{12} (certainty over one year), we interpret p_{12} as ‘almost certainty’ of malaria, replacing $p_{12} = 1$ with $p_{12} = 0.95$, so that $p_6 = 1 - \sqrt{0.05} = 0.76$. Using the unadjusted beliefs resulted in estimating even higher rates of present-bias.

Time-invariant Household-specific controls. The utility function in equation (21) also accounts for time-invariant controls z recorded at baseline that include a measure of risk aversion, an asset index, household size, and whether the household already owned bed nets before our sales program. The first three variable are standardized by subtracting the mean and dividing by the standard deviation, while bed net ownership is binary, as described below. Attitudes Towards Risk: This is measured 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 (standardized) number of times the household chose option A in response to the choice problems. Household Assets: This is the (standardized) first principal component of the following baseline binary asset indicators, equal to one if the household owned the asset: dwelling, motorbike, bicycle, radio, clock, car, television, fan, poultry, livestock (small and large), land. Household size: This is the number of household members, censored at nine and standardized. Bed net ownership: This is a binary variable = 1 if the household owned at least one bed net at baseline, and zero otherwise.

F 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 Conditional Maximum Likelihood Estimation (CMLE). We start with the joint distribution of the observed variables for a single agent and express them as a mixture over the type distributions:

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

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 exclusion restrictions and the Markov property to simplify the right hand side as:

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

so that (and dropping inessential quantities)

$$P(a_1, a_2, a_3, x_1, x_2, x_3, r; z) \propto \prod_{t=1}^2 P(x_{t+1} | x_t; z) P(x_1; z, r) P(r) \sum_{\tau} \left(\prod_{t=1}^3 P(a_t | x_t; z, \tau) P(\tau | r) \right). \quad (77)$$

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

$$\sum_{i=1}^n \log \left(\sum_{\tau \in \mathcal{T}} P(\tau | r_i) \left(\prod_{t=1}^3 P(a_{it} | x_{it}; z_i, \tau) \mathbb{I}\{a_{i1} \neq 0\} + P(a_{i1} | x_{i1}, z_i, \tau) \mathbb{I}\{a_{i1} = 0\} \right) \right). \quad (78)$$

Denote the vector of model parameters by $\theta \equiv (\delta, \beta_N, \beta_S, \phi, \gamma)$, where δ is the usual exponential discounting

parameter, (β_N, β_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 choice of θ , the different parts of equation (eq. 78) 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 B4](#) (which is an application of [Lemma 3](#)). 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 [eqs. \(24\), \(27\)](#) and [\(30\)](#). Using these value functions one can compute the model choice probabilities using the right hand side of [eqs. \(25\), \(26\)](#) and [\(29\)](#) for any given set of parameter values. In order to compute these value functions we also need estimates of the transition probabilities $dF(x_{t+1}|x_t, a_t; z)$ 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 $v = (v_{\text{hhs}}, v_{\text{assets}}, v_{\text{risk}}, v_{\text{oldnet}})$, type-specific fixed effects, and elicited subjective beliefs (z) about the reduction of malaria risk from using untreated bednets and ITNs. We also use household beliefs about income transitions to compute this transition probability (but suppress dependence of this variable in the notation, see p.[OA-17](#) for details). We choose our estimate of θ_0 to be the value of θ that maximizes the sample likelihood function implied by eq. (78). For considerations of space, we have derived only the case of unobserved population types. The derivation of the known type case is available from the authors upon request. For our CMLE estimations we used a constrained minimization routine in MATLAB (using the sqp and interior-point algorithms) with multiple random starting points and a function value convergence tolerance of 10^{-16} . We minimized the (negative of the) log-likelihood using 40 random starting points and for each point carried out a global minimization (using the sqp algorithm) followed by another minimization using the interior-point algorithm. We then chose the value that minimized the objective function over the range of starting points. For additional details on the computational analysis please see the Computational Reproducibility Analysis on [Zenodo](#).⁴⁷

Note that if we were to use a_1 as a type indicator, the decomposition used to obtain eq. (68) would be different in that we would decompose (suppressing x_1) $P(a_1, \tau; r)$ as $P(\tau|a_1, r)P(a_1|r)$. Using this in the likelihood function involves taking a stand on how to treat first-period choices $P(a_1|r)$. One option is to leave them un-modeled in which case they would not contribute to the likelihood function. This has the obvious disadvantage that first-period choices are not used in estimation. The second option would be to decompose $P(a_1|r)$ as $\sum_{\tau} P(a_1|r, \tau)$ and use the structural model implied type-specific probabilities. However, this approach has the disadvantage that it assumes types can vary over time (i.e. agents can be one type in period 1 and another type in periods 2 and 3). Since neither approach is attractive, we do not use a_1 as a type indicator for estimation. Instead, we derive $P(\tau|a_1, r)$ as derived from the model using Bayes rule (i.e. $P(a_1|\tau, r)P(\tau|r)/P(a_1|r)$) where estimates of all the objects on the right-hand are available after the structural estimation.

G 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 of choosing action a in period t for an agent i of type τ with observed states x_i , beliefs z_i , and a vector of preference parameters θ which also includes the sub-parameter γ that describes the population type probabilities. Denote the probability of agent i buying a bed net with either contract b or c (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 purchase a net in period 1 can be written as

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

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

⁴⁷ <https://zenodo.org/records/15699365>

(recall that only nets sold through our program could be re-treated, and so re-treatment was only possible if at least one net was purchased)

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

Next, denote by $\hat{\theta}$ the vector of estimated parameters, and by θ^C a parameter vector 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 , differences in the probability of bednet purchase or re-treatment in periods 2–3 depending on the two parameter vectors can be interpreted as the ‘effect’ of present bias on the choices of the naïve and sophisticated types. The impact of present bias on the choices of agent i in periods $t \in \{1, 2, 3\}$ 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

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

where for $a \in \{b, c\}$ we have

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

We next focus on how the change in the probabilities of buying or re-treating a bednet affect the probability of having malaria in the subsequent periods. To do so, define the difference in the probability of getting malaria when not using any bednet relative to sleeping under an ITN as $D_{ITN,0}$, the difference relative to sleeping under an untreated net as $D_{untr,0}$, and the difference in the probability between sleeping under an untreated net and an ITN as $D_{ITN,untr}$. Denote by $\mathbb{1}_{i,untr}$ a binary variable = 1 if agent i already owned any untreated nets before period 1. To determine these quantities we use either the subjective beliefs measured in the survey, or we use values from Lengeler (2009), i.e. $D_{ITN,0} = 0.5$, $D_{ITN,untr} = 0.39$, and $D_{untr,0} = 0.11$. We maintain the assumption that both bednets purchased through our intervention and previously owned bednets remain available for usage during the study period, while (consistent with the actual rule followed during the study) re-treatment is only possible for nets purchased from us.

The impact of present-bias on malaria risk for household i at $t = 2$ (that is, in the time interval between $t = 1$ and $t = 2$) can be calculated as the product between the impact on the purchase rate and the difference in malaria risk between using an ITN and an alternative that is equal to either not using a net (if none was owned prior to $t = 1$) or using an untreated net (if at least one was). The impact of present bias on malaria risk (MR) can thus be written as

$$\Delta MR_{2,i} = \Delta\Psi_{2,i}(\hat{\theta}, \theta^C) [D_{ITN,0}(1 - \mathbb{1}_{i,untr}) + D_{ITN,untr}\mathbb{1}_{i,untr}].$$

In periods $t \in \{3, 4\}$, the change in the household-specific malaria risk needs to take into account both the impact of present bias on the decision to purchase a bednet with either contract, and the decision to re-treat the nets if any are indeed purchased. Thus, the change in the probability of malaria for household i can be written as

$$\begin{aligned}
\Delta MR_{t,i} = & \sum_{\tau \in \{N,S\}} \pi_\tau(r_i, \hat{\gamma}) \left\{ \sum_{j \in \{b,c\}} \Delta P_{\tau,1}(a_1 = j \mid x_i; z_i, \hat{\theta}, \theta^C) \left[P_{\tau,t}(a_t = 0 \mid x_i; z_i, \theta^C)(1 - D_{untr,0}) \right. \right. \\
& + P_{\tau,t}(a_t = 1 \mid x_i; z_i, \theta^C)[(1 - \mathbb{1}_{i,untr})D_{ITN,0} + \mathbb{1}_{i,untr}D_{ITN,untr}] \left. \right] \left. \right\} \\
& + \sum_{\tau \in \{N,S\}} \pi_\tau(r_i, \hat{\gamma}) \left\{ P_{\tau,1}(a_1 = b \mid x_i; z_i, \hat{\theta}) [\Delta P_{\tau,t}(a_t = 1 \mid x_i; z_i, \hat{\theta}, \theta^C)D_{ITN,untr}] \right. \\
& \left. \left. + P_{\tau,1}(a_1 = c \mid x_i; z_i, \hat{\theta}) [\Delta P_{\tau,t}(a_t = 1 \mid x_i; z_i, \hat{\theta}, \theta^C)D_{ITN,untr}] \right. \right\}.
\end{aligned}$$

Next, we use these estimates to recover the household-specific ‘impact’ of present bias on malaria costs, evaluated using either the expected monetary cost of a malaria episode or the expected numbers of days lost, both measured in our survey. We show the median costs in [Table 6](#).

H 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:** $x_4 \in \{h, m\}$

$$u(x_4) = y - \mathbb{1}\{x_4 = m\}\eta_m + \theta,$$

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

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

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 $a_1 \in \{b, c, n\}$

$$u(x_1, a_1) = y - \mathbb{1}\{x_1 = m\}\eta_m - p_b \mathbb{1}\{a_1 = b\} - p_c \mathbb{1}\{a_1 = c\} + \theta,$$

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

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

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 ($p_{i,itn}$): $p_{i,nonet} = 0.8 + 0.2u_{i,nonet}$; $p_{i,netunt} = 0.2 \cdot 1\{b_{i,netunt} = 0\} + 0.7 \cdot 1\{b_{i,netunt} = 1\} + 0.05u_{i,netunt}$; $p_{i,itn} = 0.05 \cdot 1\{b_{i,itn} = 0\} + 0.3 \cdot 1\{b_{i,itn} = 1\} + 0.05u_{i,itn}$; where $u_{i,itn}, u_{i,netunt}$, and $u_{i,nonet}$

Table H.1: Monte Carlo Results: Directly Observed Types

		$\beta_S = \beta_N = \beta$				$\beta_S \neq \beta_N$			
		Mean	Median	Std.Dev	True	Mean	Median	Std.Dev	True
N=300	δ	0.897	0.917	0.094	0.9	0.918	0.922	0.070	0.9
	β	0.315	0.309	0.071	0.3	0.295	0.288	0.055	0.3
	β_S					0.596	0.582	0.109	0.6
	θ	1.002	1.014	0.111	1	1.016	1.011	0.128	1
N=600	δ	0.911	0.910	0.074	0.9	0.907	0.901	0.061	0.9
	β	0.300	0.294	0.052	0.3	0.301	0.297	0.046	0.3
	β_S					0.606	0.604	0.092	0.6
	θ	1.001	1.002	0.080	1	1.001	1.001	0.086	1

Notes: Each model was simulated 200 times.

are all uniformly distributed over (0,1), $b_{i,\text{netunt}}$ follows a Bernoulli distribution with mean 0.6, $b_{i,\text{netunt}}$ 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 .

Table H.1 show the results for the observed types case for one and two separate present bias parameters, respectively. Table H.2 show associated results when the types are unobserved.

Table H.2: Monte Carlo Results: Unobserved Types

		$\beta_S = \beta_N = \beta$				$\beta_S \neq \beta_N$			
		Mean	Median	Std.Dev	True	Mean	Median	Std.Dev	True
N=300	δ	0.752	0.740	0.181	0.700	0.733	0.713	0.195	0.700
	β	0.383	0.336	0.188	0.400	0.374	0.304	0.253	0.400
	β_S					0.283	0.225	0.223	0.200
	θ	0.984	0.979	0.139	1.000	0.981	0.976	0.140	1.000
N=600	π_C	0.282	0.238	0.182	0.200	0.262	0.231	0.164	0.200
	π_N	0.411	0.427	0.074	0.478	0.409	0.425	0.059	0.478
	π_S	0.307	0.283	0.186	0.322	0.329	0.317	0.183	0.322
	δ	0.747	0.736	0.156	0.700	0.750	0.725	0.176	0.700
N=1200	β	0.383	0.348	0.156	0.400	0.362	0.307	0.205	0.400
	β_S					0.278	0.244	0.172	0.200
	θ	1.001	0.993	0.099	1.000	0.995	1.000	0.100	1.000
	π_C	0.244	0.232	0.136	0.201	0.219	0.208	0.123	0.200
	π_N	0.431	0.446	0.065	0.478	0.423	0.435	0.056	0.478
	π_S	0.325	0.296	0.168	0.321	0.358	0.339	0.153	0.322
	δ	0.712	0.702	0.119	0.700	0.718	0.697	0.147	0.700
	β	0.402	0.393	0.118	0.400	0.384	0.340	0.162	0.400
	β_S					0.270	0.240	0.156	0.200
	θ	1.003	1.001	0.066	1.000	1.002	1.007	0.076	1.000
	π_C	0.230	0.221	0.120	0.201	0.215	0.209	0.103	0.201
	π_N	0.440	0.462	0.059	0.478	0.438	0.441	0.044	0.478
	π_S	0.330	0.306	0.151	0.321	0.347	0.328	0.129	0.321

Notes: Each model was simulated 200 times.

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

		Assuming 1 Type				Assuming 3 Types			
		Mean	Median	Std.Dev	True	Mean	Median	Std.Dev	True
N=300	δ	0.308	0.308	0.019	0.99	0.560	0.513	0.132	0.46
	β_N	1	1	0	0.06	0.701	0.738	0.296	1
	β_S	1	1	0	0.16	0.556	0.639	0.337	1
	θ	0.953	0.990	0.064	1	1.017	1.011	0.129	1
	π_C	1	1	0	0.222	0.489	0.504	0.283	1
	π_N	0	0	0	0.445	0.217	0.147	0.148	0
	π_S	0	0	0	0.333	0.294	0.188	0.286	0
N=600	δ	0.306	0.306	0.014	0.99	0.529	0.501	0.083	0.46
	β_N	1	1	0	0.06	0.788	0.816	0.218	1
	β_S	1	1	0	0.16	0.599	0.694	0.341	1
	θ	0.965	0.990	0.046	1	1.001	1.000	0.099	1
	π_C	1	1	0	0.221	0.516	0.555	0.290	1
	π_N	0	0	0	0.445	0.225	0.155	0.148	0
	π_S	0	0	0	0.334	0.259	0.172	0.257	0

Notes: The left panel shows placebo simulations when assuming a single time-consistent type when, in fact, the data is generated from three distinct types. In the right panel we estimate a 3-type model when, in fact, there is a single type. Each model was simulated 200 times.