Review of Economic Studies (2025) **0**, 1–0 doi:10.1093/restud/rdaa002

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

APRAJIT MAHAJAN
UC Berkeley, Dept. of Agricultural and Resource Economics, USA
CHRISTIAN MICHEL
Barcelona School of Economics, Spain
ALESSANDRO TAROZZI
Department of Economics, European University Institute (Italy), Universitat Pompeu Fabra (Spain), Barcelona School of Economics and CEPR

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.

Key words: Time Inconsistency, Identification of Types, Partial Identification, Mixture Models, Expectations, Bed nets, Dynamic Discrete Choice
JEL Codes: I1, I3, D9.

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

1. See 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.

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 malaria. 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 whether 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 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

2. See, 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.

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. (2023) 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. (2023) 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 results, 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. THE 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.

<u>STATE SPACE: $\mathcal{S}_t$</u>. 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 1 for a formal statement and the subsequent discussion of the exclusion restriction. We 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) with dimension equal to the number of actions relevant at t.

Action Space: $\mathcal{A}_t$. In period t, the agent takes one of a finite number K_t of actions $a_t \in \mathcal{A}_t$.

Transition Probabilities: $\mathbb{P}(s_t|s_{t-1},a_{t-1})$. We assume that agent beliefs about future states can be represented by a Markov transition probability: $\mathbb{P}(s_t|s_{t-1},a_{t-1})$ denotes the distribution function of the random vector $s_t \in \mathcal{S}_t$ conditional on (s_{t-1},a_{t-1}). The Markov assumption is standard in dynamic choice (see e.g. Rust, 1994; Aguirregabiria and Mira, 2010); 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. This rules out for instance direct feedback from current shocks to future state variables. One can deal with this limitation by including a sufficiently rich 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 $\mathcal{A}_t$ (and $\mathbf{d} \equiv \{d_t\}_{t=1}^3$). We do not allow for history-dependent decision rules (i.e. mappings from $\prod_{s=1}^{t-1}(\mathcal{S}_s,\mathcal{A}_s)\times\mathcal{S}_t \to \mathcal{A}_t$). Given the Markov property for the agent belief transition probabilities and the time-separability assumption 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, the payoff to a type τ agent from a given decision rule $\mathbf{d}$ at time t is :

$$\tilde{u}_t(s_t,d_t;\tau)+\beta_\tau \sum_{j=t+1}^{3} \delta_\tau^{j-t}\mathbb{E}_{\mathbf{d}}\left(\tilde{u}_j(s_j,d_j;\tau)\right), \tag{1}$$

where $\mathbb{E}_{\mathbf{d}}$ denotes the expectation induced by the decision rule $\mathbf{d}$. Broadly, we deal with three types of agents: time-consistent (denoted by τ_C or C), time-inconsistent naïve (τ_N or N) and time-inconsistent sophisticated (τ_S or S) although: (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

3. This 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.

accommodate partially sophisticated agents and we provide set identification results for this case.

Following O'Donoghue and Rabin (1999), time-consistent agents have $\beta_{\tau_C} \equiv \beta_C = 1$, as in the standard case of exponential discounting. Such agents will maximize eq. (1) using standard dynamic programming methods (backward induction in this finite horizon case). Both the other types are time-inconsistent, with hyperbolic parameter $\beta_\tau < 1$, and both are aware of their *current* present-bias, and solve the maximization problem using backward induction. However, while sophisticated agents ($\beta_{\tau_S} \equiv \beta_S < 1$) 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 *over time* can be separately attributed to evolving preferences over states and to time-preferences. 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 1. We have already discussed the first two (and they are standard in the dynamic discrete choice literature) and discuss the remaining three below.

Assumption ***1 (Basic Assumptions)***

Markov Property*:* $\mathbb{P}(s_t|s_{t-1},...,s_1,a_{t-1},...,a_1) = \mathbb{P}(s_t|s_{t-1},a_{t-1}), \forall\{a_1,s_1,\cdots,a_{t-1},s_{t-1}\}$
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. In addition, $\{\epsilon_t\}_{t=1}^T$ are independent random vectors.
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 per-period utility. Intuitively, z_t induces variation in the forward-looking part 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, as long as they are effectively exogenous (i.e. not determined by actions, preferences or other state variables). The elicitation and use of data on expectations and beliefs, as proposed forcefully by Manski (2004), has become 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). 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 conditions within Assumption 1 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 and 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 that act as type indicator or type proxy (collectively denoted by $r \in \mathcal{R}$). In the second case, we assume that r onlyimperfectly reveals an agent's type (for instance due to the 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

4. In addition to the state variables, per-period utility can also depend on time-invariant factors (e.g. schooling or sex) that we denote by v. Since these play no role in identification, we omit them as arguments in preferences although we do include such variables in the empirical application.

5. For an example on how to relax this assumption see the time-consistent model with utility known up to a constant shift in Abbring and Daljord (2020b, Appendix A).

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 $\mathcal{R}=\{r_{C_1},\ldots,r_{C_J},r_{S_1},\ldots,r_{S_K},r_{N_1},\ldots,r_{N_L}\}$. 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 in this sub-section 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 ***2.1 (Directly Observed Types and Exclusion Restriction)***
Agent type is a known deterministic function of r, therefore choice probabilities are directly observed for each type. For each 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 1 we can write the choice probability as

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

The decision in the terminal period is described by a standard static discrete choice model with a known error distribution. We can thus invert the choice probability to directly identify the period 3 utilities up to the normalization in Assumption 1.[7]

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

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

6. Since 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.

7. See Appendix C of Mahajan et al. (2025) for an alternative and self-contained argument or see Hotz and Miller (1993).

where $v^*_{\tau,3}(s_3)$ is defined in eq. (3), and $\mathrm{dF}(s_{t+1}|x_t,z_t,a)$ denotes the distribution of $s_{t+1}=(x_{t+1},z_{t+1},\epsilon_{t+1})$ conditional on (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,

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

where $\mathrm{dF}(\epsilon_{t+1})$ is known (by Assumption 1). We further assume that $\mathrm{dF}(x_{t+1},z_{t+1}|x_t,z_t,a)$ is identified so that eq. (2) is identified. However, how 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 $\mathrm{dF}(x_{t+1},z_{t+1}|x_t,z_t,a)$ (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

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

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. $\mathrm{dF}(z_{t+1}|z_t,x_t,a)=\mathrm{dF}(z_{t+1})$). 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 $\mathrm{dF}(x_{t+1}|x_t,z_t,a)=q(a,x_t,x_{t+1})$. 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 $\mathrm{dF}(x_{t+1}|x_t,z,a)$. These are directly identified since they are elicited from each agent.[8]

Finally, define

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

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:

$$g_{\tau,2,k}(x_2,z_2)\equiv u_2(x_2,k;\tau)-u_2(x_2,0;\tau)+\beta_\tau\delta_\tau\int h_{\tau,3}(x_3,z_3)\mathrm{dF}_{\Delta,k}(x_3,z_3|x_2,z_2), \tag{4}$$

8. In 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) to cut down survey length, and we do not allow beliefs to vary over time, see Section 4.1 for details. In Section 3 we use only within-period variation in beliefs for identification (so identification is feasible with time-invariant beliefs). We thank a referee for comments and suggestions on this point.

where $h_{\tau,3}(x_3,z_3) \equiv \int v^*_{\tau,3}(s_3)\mathrm{dF}(\epsilon_3)$, and where $\mathrm{dF}_{\Delta,k}(x_3,z_3|x_2,z_2) \equiv \mathrm{dF}(x_3,z_3|x_2,z_2,k) - \mathrm{dF}(x_3,z_3|x_2,z_2,0)$ 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 (RHS) of eq. (4)—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 $u_3(.)$ is identified and the distribution of ϵ_3 is known) and (b) $\mathrm{dF}(x_3,z_3|x_2,z_2,k)$ is directly identified from the data so that the signed measure $\mathrm{dF}_{\Delta,k}(\cdot)$ is identified. Next, z_2 only enters the last term in eq. (4) 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. (4), 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 *2.2 (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*

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

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,

$$\int h^B_S(x_2,z_2,\beta_S\delta_S)\left(\mathrm{dF}_{\Delta,k}(x_2,z_2|x_1,z'_1) - \mathrm{dF}_{\Delta,k_1}(x_2,z_2|x_1,z''_1)\right) \neq 0.$$

We define and describe $h_{\tau,2}(\cdot)$ and $h^B_S(\cdot)$ 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 $\mathrm{dF}_\Delta(\cdot)$ 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 1, 2.1 and 2.2. 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.*

9. When 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.

3. *The product of exponential and 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 utility functions in $t=1$. As before, the conditional choice probability of choosing action k in $t=1$ is given by

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

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

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

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 $A_\tau(\cdot)$, and where we use A to represent both the event and an indicator for the event.

We introduce the arguments $(\mathsf{d}_1,\mathsf{d}_2)$ since the identification results will involve evaluating the value function at different candidate values for these parameters. The argument 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.[10] The argument d_2 is the rate of time-preference that the period 1 self *believes* she will use to make choices in t=2. In standard dynamic programming problems—i.e for time-consistent agents, $\mathsf{d}_2=\delta$. Agents who are completely aware of their future present-bias will have $\mathsf{d}_2=\beta\delta$. In contrast, agents who are completely unaware of it are those with $\mathsf{d}_2=\delta$ and $\beta<1$. More generally, we can write $\mathsf{d}_2=\tilde{\beta}_\tau\delta_\tau$, 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

10. Note 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 $v_{\tau,2}(\cdot)$.

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 "naiv̈ete". For the main results in this paper we make the assumption that agents are either completely sophisticated or completely naiv̈e.

Assumption *2.3 (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, when $\tilde{\beta}_\tau \in [\beta_\tau, 1]$ is an additional parameter that is only set-identified. Intuitively, with partial sophistication the extra parameter $\tilde{\beta}_\tau$ drives choices in $t=1$ 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 $g_{\tau,1,k}(\cdot)$:

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

where
$$h_{\tau,2}(x_2,z_2) \equiv \int v^*_{\tau,2}(s_2)\,\mathrm{dF}(\epsilon_2) = \sum_{a\in\mathcal{A}_2}\int v_{\tau,2}(s_2,a,\delta_\tau)A_\tau(s_2,a,\tilde{\beta}_\tau\delta_\tau)\,\mathrm{dF}(\epsilon_2)$$

$$= \sum_{a\in\mathcal{A}_2}\int\Big(u_2(x_2,a;\tau)+\epsilon_2(a)+\delta_\tau\underbrace{\int v^*_{\tau,3}(s_3)\,\mathrm{dF}(s_3|x_2,z_2,a)}_{\equiv q_\tau(x_2,z_2,a)}\Big)A_\tau(s_2,a,\tilde{\beta}_\tau\delta_\tau)\,\mathrm{dF}(\epsilon_2)$$

$$= \underbrace{\sum_{a\in\mathcal{A}_2}\int(u_2(x_2,a;\tau)+\epsilon_2(a))A_\tau(s_2,a,\tilde{\beta}_\tau\delta_\tau)\,\mathrm{dF}(\epsilon_2)}_{\equiv\tilde{h}^A_\tau(x_2,z_2,\tilde{\beta}_\tau\delta_\tau)} + \delta_\tau\cdot\underbrace{\sum_{a\in\mathcal{A}_2}q_\tau(x_2,z_2,a)\int A_\tau(s_2,a,\tilde{\beta}_\tau\delta_\tau)\,\mathrm{dF}(\epsilon_2)}_{\equiv\tilde{h}^B_\tau(x_2,z_2,\tilde{\beta}_\tau\delta_\tau)}$$

$$= \tilde{h}^A_\tau(x_2,z_2,\tilde{\beta}_\tau\delta_\tau) + \delta_\tau\cdot\tilde{h}^B_\tau(x_2,z_2,\tilde{\beta}_\tau\delta_\tau). \quad (8)$$

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 further assumptions.

Second, for a sophisticated type, $\tilde{\beta}_S = \beta_S$ (by Assumption 2.3) 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:

$$g_{S,1,k}(x_1,z_1) = u_1(x_1,k;\tau_S) - u_1(x_1,0;\tau_S) + \beta_S\delta_S \int h_S^A(x_2,z_2,\beta_S\delta_S)\mathrm{dF}_{\Delta,k}(x_2,z_2|x_1,z_1)$$

$$+\delta_S(\beta_S\delta_S)\int h_S^B(x_2,z_2,\beta_S\delta_S)\mathrm{dF}_{\Delta,k}(x_2,z_2|x_1,z_1), \tag{9}$$

where the δ_S term (that multiplies $\beta_S\delta_S$ in the last term) is the only term that is not identified. We can again use variation in z_1 (guaranteed by the second rank condition of Assumption 2.2) 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 (1) at $t=1$ and suppose that the model satisfies Assumptions 1, 2.1 and 2.2. 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 1, 2.1, 2.2 and 2.3. 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) 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 $\tilde{\beta}_\tau\delta_\tau$ (since for these types $\tilde{\beta}_\tau\delta_\tau = \beta_\tau\delta_\tau$) but for naïve agents knowledge of $\beta_N\delta_N$ is *not* sufficient to identify $\tilde{\beta}_N\delta_N$ since the latter is equal to δ_N which is not identified. The unknown parameter δ_N enters eq. (7) non-linearly 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 such assumption is that δ_N is equal to δ_S or δ_C: since the latter 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 Section 1.1.1, 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 2.2) 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 more realistic (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 2.1 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. The broad intuition here is that types can be separately identified as long they behave sufficiently differently over a sufficiently rich state space. Let $\mathcal{T}$ denote the finite set of possible types. As in the previous section, under Assumption 2.3, 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.[11]

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

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

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 ***3.1 (Exclusion Restrictions)***

1. *Conditional upon type, the type proxy r is uninformative about choice:* $\mathbb{P}_\tau(a_1,\mathbf{x}_1|r,v) = \mathbb{P}_\tau(a_1,\mathbf{x}_1|v)\ \forall (a_1,\mathbf{x}_1,v)$, *and* $\mathbb{P}_\tau(a_t|\mathbf{x}_t,r,v) = \mathbb{P}_\tau(a_t|\mathbf{x}_t,v)\ \forall (a_t,\mathbf{x}_t,v)$ $t > 1$.
2. *Transition probabilities do not vary by type and are independent of r:*

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

Both parts of the assumption can be viewed as exclusion restrictions. The first part is reasonable if r is only informative about choices because it signals agent type but it will fail if, in contrast, r is also informative about other aspects of the decision process.[12] 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).

Next, we write the joint distribution of actions and states in two adjacent periods conditional on (r,v) as a mixture of the corresponding type-specific joint distributions

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

11. More generally, we can subsume all inconsistent agents under the rubric of partially sophisticated agents (i.e. those with $\tilde{\beta}_\tau \in [\beta_\tau,1]$) and within each type there might exist further subtypes with different preference parameters. Identification in this scenario is more complex, we provide set identification results below.

12. For instance, if r signals time-inconsistency but it also reflects low numeracy or other flaws in an agent's cognitive processes it may have an independent effect on choice, even after conditioning on type. One can mitigate this problem by including a rich set of observables v (e.g. schooling or measures of cognitive skills).

Using Assumption 3.1 and the Markov decision rule (see Assumption 1) 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

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

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 $(a_t,\mathbf{x}_t^1,\ldots,\mathbf{x}_t^{\underline{M}},a_{t+1},\mathbf{x}_{t+1}^1,\ldots,\mathbf{x}_{t+1}^{\underline{M}})$ define the $(\underline{M}+1)\times(\underline{M}+1)$ directly identified matrix using the expressions defined above in eq. (11):

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

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

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

where we use $\lambda_{\cdot}^{\cdot,\tau}$ as a short-hand notation for the type-specific choice probabilities and $\lambda_v^{a_t,\mathbf{x}_t,\tau}\equiv\mathbb{P}_\tau(a_t,x_t|v)$. To express the relationship between the directly identified object P and the type-specific choice probabilities, we define the two following $M_{r,v}\times(\underline{M}+1)$ matrices

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

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

Next, defining the diagonal matrix of type frequencies $\mathsf{V}_{r,v} \equiv \mathrm{Diag}(\pi_{\tau_1}(r,v),\dots,\pi_{\tau_{M_{r,v}}}(r,v))$, we can express the identified matrix P in terms of the unknown objects of interest:

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

We will use this relationship to determine the number of types (and subsequently identify the elements on the RHS as well). The dimension and rank of the RHS of eq. (16) is important because as we show below they are related to the number of types. Formally:

Assumption ***3.2 (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* $\mathsf{L}_v^{a_t,\mathbf{x}_t,(\underline{M}+1)}$ *and* $\mathsf{L}_{\mathbf{x}_{t+1},v}^{a_{t+1},(\underline{M}+1)}$ *have rank equal to* $M_{r,v}$.

The first part of the assumption ensures that the elements of P in eq. (12) are well defined and is, in principle, testable since it involves 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 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 L matrices are linearly independent) and at a given point in the state space (columns in the 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 3.1 and Assumption 3.2 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 3.1 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 $M_{r,v} \geq \mathrm{Rank}(\mathsf{P}_{r,v}^{\underline{M}})$ where the directly identified matrix $\mathsf{P}_{r,v}^{\underline{M}}$ is defined in eq. (12); (2) Suppose in addition that Assumption 3.2 holds; then, $M_{r,v} = \mathrm{Rank}(\mathsf{P}_{r,v}^{\underline{M}})$.

The first part provides a lower bound when the support of the state space (i.e. $\underline{M}$) is relatively small and links the richness of the state space to the number of types that can be identified. The second part shows that given a sufficiently rich state space and sufficient variation in type behavior across this state space, the lower bound is also the upper bound. Proposition 1 provides a result for the number of types at each point (r,v). The conditioning on time-invariant 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 $\mathsf{P}_{r,v}^{\underline{M}}$ as well as the exclusion restriction in Assumption 3.1. 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 $\pi_{\mathcal{T}}$ 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 $\pi_{\mathcal{T}}(r,v)=\pi_{\mathcal{T}}(r)$. 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 $\mathsf{L}_v^{a_t,\mathbf{x}_t,M_{r,v}}$ and $\mathsf{L}_{\mathbf{x}_{t+1},v}^{a_{t+1},M_{r,v}}$ 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 $\mathsf{L}_{t,r}$ and as $\mathsf{L}_{t+1,r}$. We collect both assumptions discussed above in the assumption below (since they are always invoked together).

Assumption ***3.3 (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 $(a_t,\mathbf{x}_t^1,\ldots,\mathbf{x}_t^{M_{r,v}-1},a_{t+1},\mathbf{x}_{t+1}^1,\ldots,\mathbf{x}_{t+1}^{M_{r,v}-1})$ such that*
   (a) $\mathbb{P}(\mathbf{x}_{t+1}^j|\mathbf{x}_t^k,a_t,v)\neq 0$ *for* $(j,k)\in\{1,\ldots,(M_{r,v}-1)\}^2$.
   (b) *The $M_{r,v}\times M_{r,v}$ matrices $\mathsf{L}_v^{a_t,\mathbf{x}_t,M_{r,v}}$ and $\mathsf{L}_{\mathbf{x}_{t+1},v}^{a_{t+1},M_{r,v}}$ are invertible.*

The second part above is a restatement of Assumption 3.2 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 3.1 and Assumption 3.3 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 1.2.1 with the main result being Lemma A4, 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 $\tilde{h}^A_\tau(x_2,z_2,\beta_\tau\delta_\tau)$ and $\tilde{h}^B_\tau(x_2,z_2,\beta_\tau\delta_\tau)$ defined in eq. (8). Recall that the function $\tilde{h}^A_\tau(x_2,z_2,\mathsf{d}_2)$ 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 d_2. Similarly, $\tilde{h}^B_\tau(x_2,z_2,\mathsf{d}_2)$ 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 d_2. Then, define the following terms for a pre-specified point (x_{20},z_{20}):

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

Functions with a Δ superscript can be interpreted as future expected utilities 'normalized' relative to a given point (x_{20},z_{20}). Lemma A1 of Section 1.1.1 shows that the following function is identified:

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

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):

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

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 two 'views of the future' are sufficiently different across the observed state space. To state this formally, consider the function $\tilde{h}_\tau(x_2,z_2,\mathsf{d}_1,\mathsf{d}_2) \equiv \tilde{h}^A_\tau(x_2,z_2,\mathsf{d}_2) + \mathsf{d}_1\tilde{h}^B_\tau(x_2,z_2,\mathsf{d}_2)$ (defined in eq. (10) in the Appendix). Normalizing it by subtracting $\tilde{h}_\tau(x_{20},z_{20},\mathsf{d}_1,\mathsf{d}_2)$ we can define the identified function $\tilde{h}^\Delta_\tau(x_2,z_2,\mathsf{d}_1,\mathsf{d}_2)$.

$$\tilde{h}^\Delta_\tau(x_2,z_2,\mathsf{d}_1,\mathsf{d}_2) \equiv \tilde{h}_\tau(x_2,z_2,\mathsf{d}_1,\mathsf{d}_2) - \tilde{h}_\tau(x_{20},z_{20},\mathsf{d}_1,\mathsf{d}_2). \tag{19}$$

The function $\tilde{h}^\Delta_N(x_2,z_2,\delta_N,\delta_N)$ is the (normalized) value of being in state (x_2,z_2) from the view-point of a naïve agent's period 1 self while $\tilde{h}^\Delta_N(x_2,z_2,\delta_N,\beta_N\delta_N)$ 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 3.3 and is stated in Section 1.2.2 in the interest of brevity). It seems reasonable to assume that the naïve and sophisticated calculations differ for if they did not, period 1 choice probabilities would be identical for naïve types and for sophisticated types who have the same preference parameters as their naïve counterparts. We can then state the formal result (details and intuition for the proof are in the Appendix):

Proposition 2 (Assigning Type-Identities). *Suppose that the type-specific choice probabilities are identified and Assumptions 1, 2.2.1 and 2.2.2 hold (so that Lemma 1 holds). Further, suppose that Assumption 3.3 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*

We now show that if we relax Assumption 2.3 so that agents are only *partially* aware of their future present-bias, then none of the time preference parameters are point identified without additional assumptions. The arguments are similar to those made for naïve types analyzed above. We start by relaxing Assumption 2.3 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, where the latter may have different values of $(\tilde{\beta}_\tau,\beta_\tau)$ as well as different exponential parameters and per period utilities. We maintain 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 A5). In particular, a type will be partially sophisticated if and only if the directly identified object $\hat{\delta}(x_2,z_2)$ (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 ***3.4 (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, $\tilde{h}^{\Delta,B}_\tau(x_2,z_2,\beta_\tau\delta_\tau) \neq 0$.
3. *For types* τ *such that* $\tilde{\beta}_\tau \neq \beta_\tau$, $\mathrm{Var}\left(\frac{\tilde{h}^{\Delta}_\tau(x_2,z_2,\delta_\tau,\tilde{\beta}_\tau\delta_\tau)-\tilde{h}^{\Delta}_\tau(x_2,z_2,\delta_\tau,\beta_\tau\delta_\tau)}{\tilde{h}^{\Delta,B}_\tau(x_2,z_2,\beta_\tau\delta_\tau)}\right) > 0$.

Recall that $\tilde{h}^{\Delta}_\tau(x_2,z_2,\delta_\tau,\tilde{\beta}_\tau\delta_\tau)$ is the (normalized) value at $t=1$ of being in state (x_2,z_2), defined in eq. (19) when one assumes that decisions in $t=2$ will be made using $\tilde{\beta}_\tau\delta_\tau$ to discount period 3 utility back to period 2. The function $\tilde{h}^{\Delta}_\tau(x_2,z_2,\delta_\tau,\beta_\tau\delta_\tau)$ 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 Section 1.3. As with Proposition 2, a key ingredient is the function $h^{\Delta}_\tau(x_2,z_2)$ 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 3.4 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 $\tilde{h}^{\Delta}_\tau(x_2,z_2,\mathsf{d}_1,\mathsf{d}_2)$. Recall that period 1 choices identify the function $h^{\Delta}_\tau(x_2,z_2)$ and we know that $h^{\Delta}_\tau(x_2,z_2) = \tilde{h}^{\Delta}_\tau(x_2,z_2,\delta_\tau,\tilde{\beta}_\tau\delta_\tau)$. This information, however, is not enough to identify the time-preference parameters since the identified function $\tilde{h}^{\Delta}_\tau(\cdot)$ is not one-to-one in $(\delta_\tau,\tilde{\beta}_\tau)$. Further, since types are partially sophisticated, we cannot impose any other restrictions on $\tilde{\beta}_\tau$ separately. This is in sharp contrast to the consistent or the completely sophisticated case where $\tilde{\beta}_\tau = \beta_\tau$, so that period 2 and 3 choices (which identify $\beta_\tau\delta_\tau$) identify $\tilde{\beta}_\tau\delta_\tau$. Equivalently, given the structure of the model, the product $\beta_\tau\delta_\tau$ is not sufficiently informative about either $\tilde{\beta}_\tau\delta_\tau$ or δ_τ. The following proposition states the most general result for partially sophisticated types, which only allows for set identification.

Proposition 4. *Suppose that the conditions for Proposition 3 hold. Then, the identified set for the parameters* $(\beta_\mathcal{T},\tilde{\beta}_\mathcal{T},\delta_\mathcal{T})$ *is given by* $\Theta_{\beta,\tilde{\beta},\delta} = \{(\mathsf{b},\tilde{\mathsf{b}},\mathsf{d}) \in (\beta_\mathcal{T}\delta_\mathcal{T},1]^2 \times [\beta_\mathcal{T}\delta_\mathcal{T},1] : \tilde{h}^\Delta_\mathcal{T}(x_2,z_2,\mathsf{d},\mathsf{bd}) = h^\Delta_\mathcal{T}(x_2,z_2)\ \forall(x_2,z_2), \mathsf{b} = (\beta_\mathcal{T}\delta_\mathcal{T})/\mathsf{d}, \mathsf{b} \leq \tilde{\mathsf{b}}\}$.

4. EMPIRICS: ITNS IN ODISHA, INDIA

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 adopt the general framework described in Section 3.2 with three unobserved types (time-consistent, hyperbolic naïve or sophisticated) and we abstract from the possibility of partial sophistication. All three types are assumed to exist with positive probability. This assumption is made in light of our sample size.[13] Agents have preferences as in eq. (1) and in our preferred specification we restrict the exponential discount factor to be common across all three types.[14] Agents choose whether to purchase an ITN and whether to retreat it periodically to ensure its insecticidal power. Given sample size concerns, we impose functional forms on the utility function so that the structural model is characterized by a small number of parameters; inference follows 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" 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.

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 ITNs and their proper use and maintenance. BISWA clients were then 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

13. With larger sample sizes 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.

14. We also estimate a model where the exponential discount factor differs for consistent and inconsistent types and find both very close to each other albeit imprecisely estimated.

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). The second contract (c henceforth) included the treated net plus two re-treatments, at a price of Rs. 203 (259), again to be repaid within one year. With this contract no additional cash payment was required for re-treatment as its cost 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, leaving a sample of 566 households. Panel A of Table 1 shows summary statistics at baseline. Mean monthly total expenditure per head was about twice the official poverty line for rural Odisha in 2004–5. Net ownership was not uncommon, with a mean of one bed net for every three persons, but net treatment was rare, with only 0.06 ITNs per head on average. Only 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.[15]

Respondents were aware of mosquitoes transmitting malaria, of the high cost of malaria episodes, and that nets reduced risk. The latter was reflected in subjective beliefs, measured asking respondents to hold up a number of fingers increasing in the perceived likelihood of an event, with none representing "no chance" and ten certain occurrence (we divide the number by ten to calculate probabilities).[16] 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.[17]

The histograms in Figure 1 show that about 3/4 of respondents believed that without nets one would certainly get malaria, or that regular use of ITNs would rule out risk.

15. Malaria 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).

16. Our data do not allow us to gauge the degree of uncertainty around the reports. In addition, since most respondents were unfamiliar with the formal concept of probability, the interviewer discussed first hypothetical examples of certain and uncertain events.

17. We also elicited perceived protective power of bed nets and treatment for children and pregnant women. Responses were almost identical, and so we only use information for adults.

FIGURE 1
Perceived Protective Power of Bednets

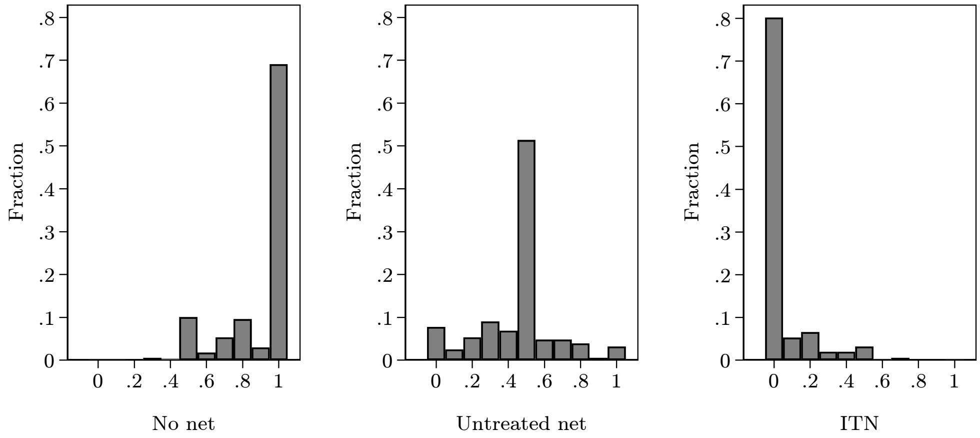

Notes: Histograms of subjective beliefs about the protective power of bed nets and treatment with insecticide from malaria risk. Data from March-April 2007 baseline survey, $n = 566$.

About half of respondents reported a 50% chance of malaria if an untreated net was used. Despite the spikes over the focal numbers 0, 5, and 10, there remains sufficient variation to be exploited in estimation.[18] 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 nets.

We also included questions to gauge respondents' intertemporal preferences and the extent of time inconsistency. First, 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, the same rewards described for the first group were offered, but with time horizons shifted by three months.[19] 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

18. Recall that the identification results required two points of support for z. The Monte Carlo results in Section 2 use two point distribution and show encouraging results.

19. Interviewers told respondents that one of the twelve chosen rewards, selected at random, would be paid by our micro-lender partner BISWA at the chosen time horizon. In practice, to avoid logistical difficulties, 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.

TABLE 1
Baseline Summary Statistics

	Mean	Median	S.d.	Obs.
(A) Summary statistics at baseline				
Household size	5.3	5	2.1	566
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
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 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.), if >0	1737	855	2729	228
(B) Intertemporal choices at baseline				
Prefers Rs. 10 in 1 month to Rs. 10 in 4 months	0.84			566
Prefers Rs. 10 in 1 month to Rs. 12 in 4 months	0.71			566
Prefers Rs. 10 in 1 month to Rs. 14 in 4 months	0.65			566
Prefers Rs. 10 in 1 month to Rs. 16 in 4 months	0.60			566
Prefers Rs. 10 in 1 month to Rs. 10 in 7 months	0.82			566
Prefers Rs. 10 in 1 month to Rs. 15 in 7 months	0.63			566
Prefers Rs. 10 in 1 month to Rs. 20 in 7 months	0.52			566
Prefers Rs. 10 in 1 month to Rs. 25 in 7 months	0.49			566
Prefers Rs. 10 in 4 months to Rs. 10 in 7 months	0.84			566
Prefers Rs. 10 in 4 months to Rs. 12 in 7 months	0.74			566
Prefers Rs. 10 in 4 months to Rs. 14 in 7 months	0.65			566
Prefers Rs. 10 in 4 months to Rs. 16 in 7 months	0.57			566
Always prefers earlier reward	0.27			566
At least one “hyperbolic” preference reversal	0.25			566
Mean no. of “hyperbolic” preference reversals (>0)	1.31			143

Notes: Data from March-April 2007 baseline survey. Un-weighted averages for n=566 households. Expenditure refers to usual consumption of 18 categories, including home production. Both actual and expected costs of malaria episodes elicited using an itemized list including doctor fees, drugs and tests, hospitalization, surgery, costs of lodging andtransportation (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, during the six months before the interview. All monetary values are in nominal Rs. (PPP exchange rate $\approx 16 Rs/USD$, World Bank, 2008). “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.

discounting.[20] In panel (B) of Table 1 we show that 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 1/4 of respondents exhibit at least one reversal and we denote agents with such reversals as having type signal $r=1$ ($r=0$ otherwise).[21]

Table 2 summarizes the results of the ITN sales and re-treatment. Slightly more than 50% of sample households purchased at least one net on credit (287 of 566). Of these, 141 purchased only ITNs (contract b), while 146 opted for the "commitment" product that also included two re-treatments (contract c). Re-treatment choices were strongly associated with the choice of contract. 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 of Mahajan et al. (2025) we provide additional descriptions of the association between actions (purchase and re-treatment) and a list of household characteristics.

TABLE 2
Summary of ITN purchase and re-treatment choices

Purchased at least one ITN	0.51
# ITNs purchased, any contract	1.03
# ITNs purchased, any contract, if >0	2.02
Purchased at least one ITN without commitment to retreat (contract 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 (contract 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
% Bednets re-treated after 6 months without commitment to retreat (b)	0.36
% Bednets re-treated after 6 months with commitment to retreat (c)	0.92
% Bednets re-treated after 12 months without commitment to retreat (b)	0.19
% Bednets re-treated after 12 months with commitment to retreat (c)	0.84

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

Our data do not include beliefs about the *joint* distribution of income and malaria, but they 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. We assume that gross income and malaria status are independent,

20. See Andreoni and Sprenger (2012) for an alternative view. Rubinstein (2003) discusses cases in which reversals arise despite preferences that are neither consistent nor time-inconsistent hyperbolic. We focus on 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.

21. We ignore "future bias" ($\beta>1$), rarely considered in the literature. In case of anti-hyperbolic behavior (when the respondent is 'patient now' but 'impatient later') we set $r=0$. Results are qualitatively similar if we drop these respondents. Also recall that the model is robust against imperfect type signals. Finally, in previous work we show that these reversals are important predictors (in a reduced form sense) of subsequent decisions about net re-treatment (Tarozzi et al., 2009).

although malaria is allowed to affect current utility by reducing consumption due to the monetary costs of illness episodes. Appendix E of Mahajan et al. (2025) provides 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 standard dynamic models is the presence of time-inconsistent agents and the further complication that types are unobserved. Furthermore, instead of the three periods required in our model, we use four because of the features of our intervention. We provide details on variable construction in Appendix E of Mahajan et al. (2025).

Preferences ($t=4$): In the last period, the state variables x_4 include income y_4 and health h_4, where $h_4=m$ if someone in the household has malaria and $=h$ ('healthy') otherwise. For simplicity we discretize income so that y_t is a dichotomous variable with 'high' or 'low' value, depending on whether household income is above or below the median. The time-invariant baseline household characteristics v that enter preferences include household size (v_{hhs}), a measure of households assets (v_{assets}), and indicators for risk aversion (v_{risk})[22] and ownership of untreated nets (v_{oldnet}). Previously owned nets also affect perceived malaria risk and we account for this in the estimation. We specify

$$u_4(x_4,v;\tau)=\mathcal{C}(x_4)+\boldsymbol{\phi}_\tau(v), \tag{20}$$

where $\mathcal{C}(x_4)$ is consumption in state x_4, and $\boldsymbol{\phi}$ 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 is the monetary cost of malaria, set to equal the median cost of an episode, including expenses for doctors and treatment as well as wages paid to labor hired to replace a sick worker. This choice is conservative in the sense that alternatives (such as using 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.

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

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

Preferences ($t=2,3$): now state variables also include the choice in $t=1$ (a_1), so

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

where $a_t=1$ if the net is re-treated 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 if b was chosen it is paid only if the net is re-treated. The multiplication of $\boldsymbol{\phi}(\cdot)$ by the action ensures that the ϕ coefficients are identified.[23]

22. We measure v_{risk} with an abbreviated version of the procedure in Holt and Laury (2002).

23. We lack reliable data on the timing of loan repayments. We assume repayment within six months, except for the cost of re-treatment. The assumption is conservative as present-bias is exacerbated

Preferences (Period 1): In $t=1$, preferences are given by

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

where $a_1\in\{n,b,c\}$ indicates the decision to purchase no contract (n), a baseline contract (b) or the commitment contract (c) and p_b is the price of a baseline contract b.[24] 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 specify the mixture probabilities using as building blocks a logit function ψ that depends on two parameters γ_1 and γ_2 and the observable type-proxy r:

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

We use a Hardy-Weinberg (Hardy, 1908; Weinberg, 1908) parameterization to define the type probabilities as $\pi_C(r)=\psi^2(\gamma,r)$, $\pi_N(r)=2\psi(\gamma,r)(1-\psi(\gamma,r))$, $\pi_S(r)=1-\pi_C(r)-\pi_N(r)$.

4.2.1. Solving the Model.

Given the finite horizon, 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). This is possible because while we do not observe types, the choice probabilities are identified using the earlier results. For clarity we sometimes suppress the dependence on observed characteristics v.

Period 3 Choice: the event that an agent (who has bought an ITN) re-treats the net is

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

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

The agent-specific variables z reflects beliefs about malaria risk as a function of bed net usage and transition probabilities of income. We emphasize the dependence on the

with later repayments while the estimates for type and relative adoption probabilities remain similar. The modeling assumption that the re-treatment component of contract c (≈13% of total price) is paid in later periods—as for optional re-treatment if b was chosen—is mostly due to technical considerations. Dropping this assumption does not qualitatively change estimates, but it worsens model fit. Fit worsens because high present-bias (i.e., small β_N and β_S) makes c less attractive than b due to the higher initial payment.

24. 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 =1.5. We set the price =1.5× (mean price of single and double nets, weighted by the respective purchase frequencies). Compared to assuming that demand is 1 or 2, 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.

hyperbolic parameter β_τ since it will be useful in the subsequent analysis. In our preferred specification we assume a common exponential discount parameter across types (i.e. $\delta_\tau = \delta$) although we also estimate models where δ differs between consistent and inconsistent types. Finally, we assume that $(\epsilon_3(0),\epsilon_3(1))/\sigma$ are *i.i.d.* standard GEV random variables, and to ease notation we set σ=1. The type-specific choice probability is then

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

<u>Period 2 Choice</u>: Under the GEV assumption on the errors, the choice probabilities are

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

The $v_{\tau,2}(\cdot)$ functions will provide insight into the time-inconsistency problem, and are defined as in $t=3$, although the calculation of the forward-looking component is more complex:

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

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 v_τ^* on the beliefs about the hyperbolic parameter $(\tilde\beta_\tau)$, 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

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

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 $\tilde\beta_S=\beta_S$. We can simplify eq. (27) further. First,

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

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

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

where $\mathbb{P}(A_{\tau,k}) = \exp(v_{\tau,3}(x_3,z,k,\tilde{\beta}_\tau\delta))/\sum_{j=0}^{1}\exp(v_{\tau,3}(x_3,z,j,\tilde{\beta}_\tau\delta))$ and γ_e is Euler's constant. 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.[25] Finally, this term is not equal to zero for sophisticated types since they are aware of their period 3 self's future bias and adjust their period 2 decisions accordingly.[26]

PERIOD 1 CHOICE: The argument is similar to the one above with the difference that there are now three possible actions and the choice probabilities for a type-τ agent are

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

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

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

As in the discussion of period 2,

$$h_\tau(x_2,z) = \sum_{k\in\{n,b,c\}} \mathbb{P}(A_{\tau,k}) \Big[v_{\tau,2}(x_2,z,k,\delta) - v_{\tau,2}(x_2,z,k,\tilde{\beta}_\tau\delta)\Big] + \gamma_e + \log\big(\sum_{j\in\{n,b,c\}} \exp(v_{\tau,2}(x_2,z,j,\tilde{\beta}_\tau\delta))\big),$$

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

25. Note, 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).

26. To 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.

4.2.2. Identification. The identification arguments and setting are broadly similar to those in the general section on identification but there are some differences that we highlight here: (a) for tractability we assume that all three types exist in the population so we do not need to first identify the total number of types. (b) 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. (c) 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.[27] (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. A discussion on identification accounting for these differences is provided in Section 3 and greater details are provided in Appendix B of Mahajan et al. (2025). Here we only discuss briefly the key variables required for identification and the assumptions imposed on them.

As a type proxy r we use the intertemporal ('MEL') choices described in Section 4.1. These responses are likely 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. In principle, the decision to not commit could also depend on low perceived benefits of re-treatment.[28] 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. Recall that r needs to satisfy two key (untestable) assumptions (see Assumption 3.1). 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. We further require that r be informative about types in a monotone likelihood ratio sense. 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—this is stated formally as Assumption 3.4, see Section 3.2.1 for details. This condition is weak in that it does not require the fraction of inconsistent agents to be larger in the sub-population with $r=1$ relative to $r=0$. Preference reversals are an imperfect proxy for time-inconsistency, and the main requirement is that they shift type probabilities in the sense above. Examining the key ratio $\frac{\pi_\tau(0)}{\pi_\tau(1)},\tau\in\{C,N,S\}$, the sufficiency criterion for type identification is met if $\frac{\pi_C(0)}{\pi_C(1)}>\frac{\pi_N(0)}{\pi_N(1)}>\frac{\pi_S(0)}{\pi_S(1)}$, which is not unreasonable in our context.

27. 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 in Appendix F of Mahajan et al., 2025).

28. However, we will show that the latter two are less of a concern to the extent that such perceptions are reflected in agents' elicited beliefs.

The second key variables for identification are subjective beliefs about malaria risk 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 of Mahajan et al. 2025 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 Assumption 1. While we cannot test this assumption directly, it seems plausible that conditional on current health status and income, beliefs about income and malaria in the next period do not affect the static payoff function.

4.2.3. Monte Carlo simulations. In Section 2 we illustrate the properties of our model through Monte Carlo (MC) simulations. For the observed types case, Table 2.1 shows that for moderate sample sizes (300 and above) both mean and median estimated time preference parameters are close to their true values. Table 2.2 show that when types are unobserved, the estimates (which now also include the type probabilities) remain close to their true values, albeit with more variability. The largest deviations 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.

In summary, the evidence from the Monte Carlo simulations is 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 (for details on the objective function see Appendix F of Mahajan et al., 2025). We parameterize the type probabilities conditional on the type signal r ($\mathbb{P}(\tau|r) \equiv \pi_\tau(r)$) as in eq. (23) and then use Bayes' rule to compute $\mathbb{P}(\tau|r,a_1) \equiv \pi_\tau(a_1,r)$.[29] The former estimate the distribution of types in the population conditional on an observed type signal, an object of interest not previously estimated in the literature. The latter estimate the type distribution conditional on both r and purchase decision. A comparison of the two evaluates the attractiveness of commitment products to sophisticated types (recall that types are unobserved and purchase decisions are not assumed to uniquely reveal type). In our preferred specification, we estimate three time preference parameters, *i.e.* the common discount factor δ 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.

29. We compute this as $\hat{\mathbb{P}}(\tau|a_1,r;\hat{\theta}) = \frac{\hat{\mathbb{P}}(a_1,\tau|r)}{\hat{\mathbb{P}}(a_1|r)} = \frac{\hat{\mathbb{P}}(a_1|\tau)\hat{\mathbb{P}}(\tau|r)}{\hat{\mathbb{P}}(a_1|r)}$, where the last equality follows from the exclusion restriction that the signal r is not informative on actions conditional on type, see Assumption 3.1.

We begin by discussing the estimated population type probabilities in col. 1 of Panel A of Table 3. We find that 36% of agents are time consistent, and the majority of time-inconsistent agents (about half of the total population or 75% of the 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.[30]

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.[31] 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 3 we present results for alternative specifications. Our preferred results (in col. 1) indicate that $\hat{\delta}\approx 1$, 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 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.

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

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

TABLE 3
Structural Estimates

		(1)	(2)	(3)	(4)	(5)
	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)
Panel A					$\pi_C(0,c)$	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					$\pi_C(1,c)$	0.663 (0.638-0.687)
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)
					$\pi_S(1,c)$	0.086 (0.074-0.098)
	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)
					$\pi_N(1,c)$	0.252 (0.239-0.264)
Panel B		Full model	No ϕ	$\beta_N=\beta_S$	Known types	One type
	δ	0.9989 (0.0438)	0.0323 (0.2456)	0.9989 (0.0461)	0.0822 (0.0237)	0.0100 (0.3549)
Preference	β_N	0.0580 (0.0613)	0.0102 (0.1261)	0.0829 (0.0266)	0.9988 (0.0311)	
Parameters	β_S	0.1550 (0.0708)	0.0100 (0.1581)		0.9986 (0.0302)	
	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)	
	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)
	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

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 29), and π_τ are unconditional type probabilities averaged using the empirical distribution of r (2.5^{th} and 97.5^{th} 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.

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 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 $\hat{\beta}_S = .05$.[32]

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 online Appendix Table 2.3. 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 mis-specified by falsely assuming the existence of time-inconsistent types, the estimates of β_S and β_N could plausibly have been large relative to δ while in fact we find that both estimates are very small.

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

4.4. *Comparisons with Existing Literature*

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

32. While 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.

experiments in rich countries with inter-temporal choices 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 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 (c).[33] 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 t=1, because 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 b to c, anticipating that present-bias problems will be exacerbated in future periods because of higher re-treatment costs. This latter effect is, however, moderated by the effect of the corresponding increase in the price of c. 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 4, 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 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 4 we re-evaluate the impact of price changes with a model where

33. 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).

TABLE 4
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 3).

all three types are time-consistent. 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 malaria risk relative to 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 Appendix G of Mahajan et al. (2025). First, we compute purchase and re-treatment probabilities using the parameters from our preferred model but setting $\beta_S = \beta_N = 1$ (i.e. assuming no present-bias) for all agents. Next, we use these probabilities to compute the expected costs of malaria for each agent.[34] 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 5 presents the results from using each measure. In both scenarios, we find that present-bias substantially increases expected costs from malaria. The median cost

34. The 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.”

TABLE 5
Median cost of malaria and days missed attributable to time-inconsistent preferences

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

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

To recover population type probabilities we require a signal 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).

Acknowledgments. We are grateful to Øystein Daljord, Hanming Fang, Paul Heidhues, Nathan Hendren, Han Hong, Gautam Rao, Charlie Sprenger, Petra Todd, exceptionally constructive anonymous referees and seminar participants at several conferences and institutions for comments and suggestions. Florens Odendahl provided excellent research assistance. We gratefully acknowledge financial support from the following sources for the completion of the field work that originated the data used in the paper: Centre for Micro Finance (Chennai, India); Stanford Presidential Fund for Innovation

35. In 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.

in International Studies; Stanford Center for International Development; Stanford OTL Research Incentive Fund; RAND Corporation; Duke Arts and Sciences Committee on Faculty Research; 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 opinions and views expressed in this article are those of the authors and do not necessarily reflect the positions of their affiliated organizations. Cross-references to the Online Appendix are not clickable in the published version; a combined PDF with working hyperlinks is available at https://www.nber.org/papers/w27198.

Data Availability
The data and code necessary to replicate the analysis in the paper is available on Zenodo at https://doi.org/10.5281/zenodo.15699365

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**, 565–571.

ABBRING, J. H. AND Ø. DALJORD (2020b), "Identifying the Discount Factor in Dynamic Discrete Choice Models", *Quantitative Economics*, **11**, 471–501.

ABBRING, J. H., Ø. DALJORD, AND F. ISKHAKOV (2019), "Identifying Present-Biased Discount Functions in Dynamic Discrete Choice Models" (Working Paper, Tilburg University).

AGUIRREGABIRIA, V. (2010), "Another Look at the Identification of Dynamic Discrete Decision Processes: An Application to Retirement Behavior", *Journal of Business & Economic Statistics*, **28**, 201–218.

AGUIRREGABIRIA, V. AND P. MIRA (2010), "Dynamic discrete choice structural models: A survey, *Journal of Econometrics*, **156**, 38–67.

ANDREONI, J., M. CALLEN, Y. KHAN, K. JAFFAR, AND C. SPRENGER (2023), "Using Preference Estimates to Customize Incentives: An Application to Polio Vaccination Drives in Pakistan", *Journal of the European Economic Association*, **21**, 1428–1477.

ANDREONI, J. AND C. SPRENGER (2012), "Estimating Time Preferences from Convex Budgets", *American Economic Review*, **102**, 3333–56.

ARCIDIACONO, P. AND P. B. ELLICKSON (2011), "Practical Methods for Estimation of Dynamic Discrete Choice Models", *Annu. Rev. Econ.*, **3**, 363–394.

ARCIDIACONO, P. AND R. A. MILLER (2020), "Identifying Dynamic Discrete Choice Models Off Short Panels", *Journal of Econometrics*, **215**, 473–485.

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**, 635–672.

AUGENBLICK, N., M. NIEDERLE, AND C. SPRENGER (2015), "Working Over Time: Dynamic Inconsistency in Real Effort Tasks", *Quarterly Journal of Economics*, **130**, 1067–1115.

AUGENBLICK, N. AND M. RABIN (2019), "An Experiment on Time Preference and Misprediction in Unpleasant Tasks", *Review of Economic Studies*, **86**, 941–975.

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**, 1–22.

BALAKRISHNAN, U., J. HAUSHOFER, AND P. JAKIELA (2020), “How soon is now? Evidence of present bias from convex time budget experiments”, *Experimental Economics*, **23**, 294–321.

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**, 1–9.

BANERJEE, A. AND S. MULLAINATHAN (2010), “The Shape Of Temptation: Implications For The Economic Lives Of The Poor” (NBER Working Paper w15973).

BERNHEIM, B. D., D. RAY, AND Ş. YELTEKIN (2015), “Poverty and Self-Control”, *Econometrica*, **83**, 1877–1911.

BISIN, A. AND K. HYNDMAN (2020), “Present-bias, procrastination and deadlines in a field experiment”, *Games and Economic Behavior*, **119**, 339–357.

BRUNNERMEIER, M. K. AND J. A. PARKER (2005), “Optimal Expectations”, *The American Economic Review*, **95**, 1092–1118.

BRYAN, G., D. KARLAN, AND S. NELSON (2010), “Commitment Devices”, *Annual Reviews of Economics*, **2**, 671–698.

CARRERA, M., H. ROYER, M. STEHR, J. SYDNOR, AND D. TAUBINSKY (2022), “Who Chooses Commitment? Evidence and Welfare Implications”, *Review of Economic Studies*, **89**, 1205–1244.

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**, 260–284.

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, NBER).

COHEN, J. AND P. DUPAS (2010), “Free Distribution Or Cost-Sharing?”, *Quarterly Journal of Economics*, **125**, 1–45.

COHEN, J., K. M. ERICSON, D. LAIBSON, AND J. M. WHITE (2020), “Measuring Time Preferences”, *Journal of Economic Literature*, **58**, 299–347.

DE GROOTE, O. AND F. VERBOVEN (2019), “Subsidies And Time Discounting In New Technology Adoption: Evidence From Solar Photovoltaic Systems”, *American Economic Review*, **109**, 2137–72.

DELAVANDE, A. (2014), “Probabilistic Expectations In Developing Countries”, *Annual Reviews of Economics*, **6**, 1–20.

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**, 151–163.

DELLAVIGNA, S. (2009), “Psychology And Economics: Evidence From The Field”, *Journal of Economic Literature*, **47**, 315–372.

DELLAVIGNA, S. (2018), “Structural behavioral economics, In B. D. Bernheim, S. DellaVigna, and D. Laibson (Eds.) *Handbook of Behavioral Economics - Foundations and Applications 1*, 613–723. (North Holland).

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**, 1768–1774.

DUFLO, E., M. KREMER, AND J. ROBINSON (2011), “Nudging Farmers To Use Fertilizer: Theory and Experimental Evidence from Kenya”, *American Economic*

Review, **101**, 2350–2390.

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*, 1–67 (North-Holland).

FALK, A., A. BECKER, T. DOHMEN, B. ENKE, D. HUFFMAN, AND U. SUNDE (2018), "Global Evidence on Economic Preferences", *Quarterly Journal of Economics*, **133**, 1645–1692.

FANG, H. AND D. SILVERMAN (2009), "Time-Inconsistency And Welfare Program Participation: Evidence From The NLSY", *International Economic Review*, **50**, 1043–1077.

FANG, H. AND Y. WANG (2015), "Estimating Dynamic Discrete Choice Models with Hyperbolic Discounting, with an Application to Mammography Decisions", *International Economic Review*, **56**, 565–596.

FREDERICK, S., G. LOEWENSTEIN, AND E. D. O'DONOGHUE (2002), "Time Discounting: A Critical Review", *Journal of Economic Literature*, **40**, 351–401.

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**, 1059–1078.

GINÉ, X., J. GOLDBERG, D. SILVERMAN, AND D. YANG (2018), "Revising Commitments: Field Evidence on the Adjustment of Prior Choices", *The Economic Journal*, **128**, 159–188.

GUISO, L., T. JAPPELLI, AND L. PISTAFERRI (2002), "An Empirical Analysis Of Earnings And Employment Risk", *Journal of Business & Economic Statistics*, **20**, 241–253.

GUL, F. AND PESENDORFER (2001), "Temptation and Self-Control", *Econometrica*, **69**, 1403–35.

GUL, F. AND W. PESENDORFER (2004), "Self-control and the theory of consumption", *Econometrica*, **72**, 119–158.

HARDY, G. H. (1908), "Mendelian Proportions in a Mixed Population", *Science*, **28**, 49–50.

HEIDHUES, P. AND P. STRACK (2021), "Identifying present-bias from the timing of choices", *American Economic Review*, **111**, 2594–2622.

HOLT, C. A. AND S. K. LAURY (2002), "Risk aversion and incentive effects", *American Economic Review*, **92**, 1644–1655.

HOTZ, V. AND R. MILLER (1993), "Conditional choice probabilities and the estimation of dynamic models", *The Review of Economic Studies*, **60**, 497–529.

KASAHARA, H. AND K. SHIMOTSU (2009), "Nonparametric identification of finite mixture models of dynamic discrete choices", *Econometrica*, **77**, 135–175.

KAUR, S., M. KREMER, AND S. MULLAINATHAN (2014), "Self-control at work", *Journal of Political Economy*, **123**, 1227–1277.

KŐSZEGI, B. (2010), "Utility From Anticipation And Personal Equilibrium", *Economic Theory*, **44**, 415–444.

KREMER, M. AND E. MIGUEL (2007), "The illusion of sustainability", *Quarterly Journal of Economics*, **122**, 1007–1065.

LAIBSON, D. I. (1997), "Golden Eggs And Hyperbolic Discounting", *Quarterly Journal of Economics*, **112**, 443–477.

LENGELER, C. (2009), "Insecticide-treated bed nets and curtains for preventing malaria", (Issue 2. Art. No.: CD000363. DOI: 10.1002/14651858.CD000363.pub2).

MAGNAC, T. AND D. THESMAR (2002), “Identifying dynamic discrete decision processes”, *Econometrica*, **70**, 801–816.

MAHAJAN, A., C. MICHEL, AND A. TAROZZI (2025), “Identification of time-inconsistent models: The case of insecticide treated nets” (Working Paper 27198, NBER).

MANSKI, C. F. (2004), “Measuring expectations”, *Econometrica*, **72**, 1329–1375.

MARTINEZ, S.-K., S. MEIERY, AND C. SPRENGER (2023), “Procrastination in the Field: Evidence from Tax Filing”, *Journal of the European Economic Association*, **21**, 1119–1153.

MULLAINATHAN, S. (2004), “Development economics through the lens of psychology”, (Working Paper 28974, World Bank).

NORETS, A. AND X. TANG (2013), “Semiparametric Inference In Dynamic Binary Choice Models”, *Review of Economic Studies*, **81**, 1229–1262.

O’DONOGHUE, T. AND M. RABIN (1999), “Doing it now or later”, *American Economic Review*, **89**, 103–124.

PASERMAN, M. (2008), “Job search and hyperbolic discounting: Structural estimation and policy evaluation”, *The Economic Journal*, **118**, 1418–1452.

ROTH, C. AND J. WOHLFART (2020), “How do expectations about the macroeconomy affect personal expectations and behavior?”, *Review of Economics and Statistics*, **102**, 731–748.

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

Rust, J. (1994), “Structural estimation of Markov decision processes, in R. Engle And D. McFadden (Eds.), *Handbook of Econometrics*, Vol. 4, 3081–3143 (Elsevier).

SCHILBACH, F. (2019), “Alcohol and Self-Control: A Field Experiment in India”, *American economic review*, **109**, 1290–1322.

SHLEIFER, A. (2019), “The Return of Survey Expectations”, (NBER Reporter, 2019).

SPRENGER, C. (2015), “Judging experimental evidence on dynamic inconsistency”, *American Economic Review*, **105**, 280–85.

STROTZ, R. H. (1955), “Myopia and inconsistency in dynamic utility maximization”, *The Review of Economic Studies*, **23**, 165–180.

TAROZZI, A., A. MAHAJAN, B. BLACKBURN, D. KOPF, L. KRISHNAN, AND J. YOONG (2014), “Micro-loans, bednets and malaria: Evidence from a randomized controlled trial in Orissa (India)”, *American Economic Review*, **104**, 1909–1941.

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**, 231–235.

TOUSSAERT, S. (2018), “Eliciting temptation and self-control through menu choices: A lab experiment”, *Econometrica*, **86**, 859–889.

UBFAL, D. (2016), “How General Are Time Preferences? Eliciting Good-Specific Discount Rates”, *Journal of Development Economics*, **118**, 150–170.

VAN DER KLAAUW, W. (2012), “On the use of expectations data in estimating structural dynamic choice models”, *Journal of Labor Economics*, **30**, 521 – 554.

VAN DER KLAAUW, W. AND K. WOLPIN (2008), “Social security and the retirement and savings behavior of low income households”, *Journal of Econometrics*, **145**, 21–42.

WEINBERG, W. (1908), *On The Demonstration Of Heredity In Man.*

WORLD BANK (2008), “Global purchasing power parities and real expenditures, 2005” (Technical report, International Comparison Program, Washington DC: World Bank).