

Intertemporal consumption with anticipating, remembering, and experiencing selves*

Sam Cosaert[†] and Tom Potoms[‡]

February 4, 2023

Abstract

We study intertemporal choice through a novel and flexible framework that accounts for savoring of future consumption and memories of past consumption. The model uses standard intertemporal budget constraints (Samuelson, 1937) but enriches preference structures with utility from anticipation, remembering, and experience (Kahneman et al., 1997). We also present an internal commitment mechanism that ensures dynamic consistency. We provide a revealed preference characterization of this model and apply it to quarterly consumption data from Spanish households. Utility from anticipation is important—and time inconsistency not strictly needed—to rationalize consumption patterns in the data.

Keywords: intertemporal consumption, instant utility, savoring, memory, revealed preferences, Encuesta Continua de Presupuestos Familiares.

JEL codes: C14, D11, D12, D91.

1 Introduction

Many consumption decisions have an intertemporal aspect: When to go on a city trip? When to go to a fancy restaurant? When to open an expensive bottle of champagne? The

*We thank Abi Adams-Prassl, Laura Blow, Laurens Cherchye, Ian Crawford, Thomas Demuyne, Bram De Rock, Paweł Dzielinski, Matthew Embrey, Hamish Low, Paola Manzini, Pietro Ortoleva, Michèle Tertilt, and seminar participants in the Applied Microeconomics seminar at Oxford University, Sussex University, the University of Antwerp, and the Luxembourg Household Economics and Labor workshop for useful comments. We gratefully acknowledge the Luxembourg National Research Fund for financial support of the *TIMING* project (C18/SC/12668748).

[†]University of Antwerp. Prinsstraat 13, 2000 Antwerp, Belgium. E-mail: Sam.Cosaert@uantwerpen.be

[‡]University of Sussex. Jubilee Building 9SL, Falmer Brighton, BN1 9SN United Kingdom. E-mail: T.Potoms@sussex.ac.uk

workhorse model to analyze intertemporal choice is the discounted utility (DU) framework.¹ Yet, there are interesting anomalies that DU fails to capture.

A first observation is the pleasurable deferral of desirable outcomes in behavioral experiments. [Loewenstein \(1987\)](#) asked respondents about their preferred timing to receive a kiss from their favorite movie star. Respondents stated their desire to postpone the kiss for a couple of days, to maximize the duration of ‘savoring’. These behavioral patterns arise not only in hypothetical experiments but also in practice ([Laajaj, 2017](#)).² This is in clear opposition to the standard lifecycle model, which assumes positive devaluing (impatience).³ A second observation is the disproportionate spending by young people on celebrations and ceremonies early in life. One explanation suggested by [Gilboa et al. \(2016\)](#) is that some goods generate happy memories long after consumption took place. The authors refer to these goods as ‘memorable’ consumption. [Hai et al. \(2020\)](#) recently showed that households optimally choose a *nonsmooth* consumption profile of memorable goods. Decreasing streams of physical consumption not necessarily imply declining welfare.

Both behavioral patterns have the following feature in common: agents care not just about physical outcomes but also about the mental image of past and future outcomes. There is a temporal dissociation between physical consumption and utility benefits, which can be enjoyed *before*, *during*, and *after* consumption. In addition, the level of anticipatory emotions and pleasant memories will generally depend on the good under consideration.

What we do and preview of results. To account for these behavioral patterns, we develop a model of intertemporal consumption where we enrich the structure of preferences, but preserve standard intertemporal budget constraints. Hence, our model is situated between the DU framework ([Samuelson, 1937](#)) and the theory of *total* utility ([Kahneman et al., 1997](#)). Like DU , it uses a linear budget constraint and can deal with multiple observations over time. However, consumers enjoy utility from anticipation of future consumption bundles, experience of a current bundle, and remembering of past bundles ([Elster and Loewenstein, 1992](#)). The underlying utility functions do not have to agree on the valuation of each bundle. This fits in a larger literature on multiple selves: individual decisions are driven by several rationales or selves. According to one stream of multi-self models, typically in intertemporal problems, there is *one ‘self’ per decision moment* and the decision makers interact in non-cooperative ways ([O’Donoghue and Rabin, 1999](#); [Kőszegi, 2009](#)). In view of

¹We refer to the seminal contributions of [Samuelson \(1937\)](#) and [Koopmans \(1960\)](#).

²Social gatherings (e.g., visits to cinema, restaurants, or bars) and holidays all induce a sense of savoring in the time leading up to the actual event. Furthermore, survey data collected by [Loewenstein and Sichertman \(1991\)](#) challenge the notion of discounted present-value maximization for wages; there is a widespread preference for *increasing* wage profiles.

³But see [Guo \(2020\)](#) for a recent exception to this positive devaluing.

our research objectives, we go one step further and admit that *three selves* (an anticipating, a remembering, and an experiencing one) *can be active at each decision moment*. Per period, the preferences of these selves are aggregated into a collective decision (Chiappori, 1988, 1992), following the seminal paper of May (1954) and recent contributions by Ambrus and Rozen (2014) and Jackson and Yariv (2014).⁴ The tools and techniques to fully characterize intertemporal consumption patterns generated by these preferences then depend on the *dynamic consistency* of the decisions.

Turning to a description of our results, we show in Proposition 1 that the model is generally not time consistent. First, the decision power of the different selves can change over time, which is a form of limited commitment. Second, Caplin and Leahy (2001) provided a discussion of the complex relationship between anticipation and time consistency. In our set-up, the duration of savoring decreases naturally as time moves forward. This creates incentives for agents to revise their anticipated consumption downwards; so called *reverse time inconsistency*. Reverse time inconsistency has hampered analyses of anticipation due to practical and theoretical reasons. From a practical perspective, Gilboa et al. (2016) explained that reverse time inconsistency invalidates the standard use of decision theory to study dynamic interactions with anticipating selves. More fundamentally, Loewenstein (1987) argued that systematic acts of time inconsistency raise concerns about self-credibility. If the agent is aware that she will postpone consumption indefinitely, how can she still enjoy benefits from anticipation at all? This suggests the existence of an internal mechanism to neutralize the reduced *duration* of anticipation over time. Proposition 2 introduces additional structure that ensures dynamic consistency. In particular, we posit that the decision weight of the anticipating self increases towards the end of the planning period. This structure is consistent with the view that in ‘shorter’ planning periods the consumer’s preferences for improvement are activated. Under this condition, one can resort to a decision theoretic approach to intertemporal consumption.

The rest of the paper then develops a time consistent version of the theory, which has clear testable implications for finite datasets of price and quantity observations per consumer. We name this *ICARES*, short for *intertemporal consumption with anticipating, remembering, and experiencing selves*. We bring *ICARES* to the data by means of revealed preference theory. The revealed preference approach, following Afriat (1967) and Varian (1982), offers an elegant way of assessing the empirical content of intertemporal models. Proposition 3 summarizes the data restrictions imposed by *ICARES*. We also present tests for relevant polar cases: intertemporal consumption with anticipating and experiencing selves (*ICAES*),

⁴Cherchye et al. (2020) recently represented food choices as the result of an efficient agreement between a healthy and an unhealthy self.

remembering and experiencing selves (*ICRES*), and an experiencing self only (*ICES*). The final characterization collapses to the nonparametric test of the lifecycle model (Browning, 1989). Two further remarks are in order. First, the empirical content of *ICARES* differs from that of static utility maximization (*GARP*). Second, characterizations *ICAES* and *ICRES* separately identify preferences for anticipation and preferences for recall. This feature distinguishes our framework from other nonparametric tests of intertemporal models. Crawford and Polisson (2014) demonstrated that the testable implications of rational habit formation are indistinguishable from the ones of rational anticipation.⁵ *ICARES*, by contrast, starts from a formal description of the asymmetric roles of temporal selves. It also has separability properties that make identification possible.

We then apply the empirical characterization to budget survey data. This demonstrates that the framework is practically relevant outside tailor-made experimental settings. Budget survey data are still the basis of much empirical work in demand theory and consumer analysis. To be more precise, we apply our revealed preference tests to a widely studied quarterly dataset from a Spanish panel of consumers (Encuesta Continua de Presupuestos Familiares; ECPF). For each respondent, we study expenditure on various nondurable commodities over the course of a year. We apply *ICARES*, *ICAES*, *ICRES*, and *ICES* separately to singles and couples. *ICARES* has a near perfect in-sample fit but lacks power, due to its flexibility. Special case *ICAES*, by contrast, is sufficiently powerful and still rationalizes about two thirds of the data. These results hold both for single consumers and for couples. We thus find strong evidence of savoring in these budget survey data. More generally, the results suggest that time *inconsistency* is not strictly needed to rationalize the consumption behavior of singles or couples. We then examine the ‘anticipatory’ nature of the goods in our data. We measure this by the increase in expenditure on each good towards the end of the sequence, for the subset of consumers who pass *ICAES*. The degree of anticipation varies considerably across goods. It is more pronounced for restaurant expenditures, expenditure on recreation services and cinema, long distance traveling, and a larger group of goods that are complementary to leisure activities and special celebrations. Overall, the results are intuitive and in line with earlier research into savoring, thus supporting the external validity of our approach. Finally, we also discuss the limitations of our budget survey data. In practice, consumers sometimes face binding liquidity constraints or unexpected income shocks. Both these are unobserved to the econometrician. For the econometrician, there is additional uncertainty in the form of measurement error. We therefore test the sensitivity of our results to these empirical challenges. Our investigations highlight that, the data limitations notwithstanding, richer

⁵However, the authors’ definition of anticipation is different from ‘savoring’. In their setting, anticipation implies that the consumer’s tastes in t are affected by her expected outcomes at future times $t + 1, t + 2, \dots$

models of temporal utility flows are empirically supported. These models can be fruitfully applied to recover aspects of the consumer’s, short-to-medium run, planning problem.

Overview of the paper. The rest of the paper unfolds as follows. Section 2 introduces the theory and section 3 presents a convenient revealed preference approach to implement the theory. Section 4 applies the model to quarterly consumption data. Section 5 presents related literature and section 6 concludes.

2 A theory of anticipating, remembering, and experiencing selves

2.1 Model setup

We start with an exposition of our notation. Agents choose a *consumption path* with expenditure on N commodities over the planning period $\mathcal{T} = \{1, \dots, T\}$, where $T \in \mathbb{N}_0$ indicates the length of the period.⁶ Let $(\mathbf{q}_t)_{t \geq 1}$ denote a consumption path $(\mathbf{q}_1, \dots, \mathbf{q}_T)$ where $\mathbf{q}_t \in \mathbb{R}_+^N$, and let $(\mathbf{p}_t)_{t \geq 1}$ denote the corresponding price vectors $(\mathbf{p}_1, \dots, \mathbf{p}_T)$ where $\mathbf{p}_t \in \mathbb{R}_{++}^N$. With r_t the applicable interest rate at time t , we can compute the relevant *discounted* prices as follows:

$$\boldsymbol{\rho}_t = \frac{\mathbf{p}_t}{\prod_{i=1}^{t-1} (1 + r_i)}.$$

In the empirical application, we use the average nominal interest rate on consumer loans. Furthermore, we use y to denote the *discounted* total resources of the consumer. For the moment, we will abstract from uncertainty regarding future prices and income. We thus adhere to perfect foresight, which is a restrictive yet common assumption in the revealed preference literature.⁷ We will relax this assumption as we implement the theory in section 3.3. Feasible consumption plans are contained in the intertemporal budget set,⁸

$$\mathcal{B}((\boldsymbol{\rho}_t)_{t \geq 1}, y) = \left\{ (\mathbf{q}_t)_{t \geq 1} \in \mathbb{R}_+^{N \times T} \mid \sum_{t \geq 1} \boldsymbol{\rho}_t \cdot \mathbf{q}_t \leq y \right\}. \quad (1)$$

⁶Our theoretical framework can easily accommodate infinite horizons but we opt to present the finite horizon version here, for ease of exposition and consistency with our empirical application. The finite horizon version is clearly more in line with our emphasis on short-to-medium run planning of nondurable consumption. We refer to Appendix B for details on the infinite horizon extension.

⁷For instance, Demuyneck and Verriest (2013), Adams et al. (2014) and Blow et al. (2021) also assume perfect foresight and apply their methods to the same dataset as in the present paper. We conduct an analysis of the predictability of prices in our dataset in Appendix C.1.

⁸We use $\mathbf{x} \cdot \mathbf{z}$ to denote the dot product between the vectors \mathbf{x} and \mathbf{z} .

A main feature of our study is that it allows a temporal dissociation between physical consumption and the utility gains from this consumption. To illustrate the utility flows in our framework, we first consider a single consumption bundle \mathbf{q}_t . Figure 1 situates this bundle graphically in the consumer’s timeline, where the decision moment is given by $\tau \geq 1$. We let functions $u^i(\mathbf{q}_t)$ capture the instantaneous trade-offs between different goods in




Figure 1: Timeline of the utility flows from anticipation, recall, and experience of \mathbf{q}_t

anticipation ($i = A$), experience ($i = E$), and recall ($i = R$). First, the mental image of future consumption \mathbf{q}_t produces utility benefits in the form of savoring (Loewenstein, 1987). Agents enjoy benefits $u^A(\mathbf{q}_t)$ from anticipation *just before* consumption of \mathbf{q}_t . Furthermore, these psychological benefits carry over to other periods before t , and may even accumulate through time. The cumulative anticipated utility, seen from τ , is given by $D^A(t - \tau) \times u^A(\mathbf{q}_t)$. Map D^A is an increasing function of the remaining duration of savoring, $t - \tau$.⁹ Cumulative anticipated utility from \mathbf{q}_t is therefore explicitly time-dependent. Second, the mental image of past consumption produces utility gains in the form of happy memories (Gilboa et al., 2016). Agents enjoy benefits $u^R(\mathbf{q}_t)$ from remembering \mathbf{q}_t *just after* its physical consumption. These psychological benefits carry over to other periods after t , and may also accumulate through time. The cumulative recollected utility from \mathbf{q}_t is $D^R(T - t) \times u^R(\mathbf{q}_t)$. Map D^R is an increasing function of the total duration of remembering, $T - t$. The cumulative utility flow from recall of \mathbf{q}_t is again explicitly time-dependent. Lastly, agents also derive value $u^E(\mathbf{q}_t)$ from physical consumption \mathbf{q}_t at time t . We assume that mappings $u^i : \mathbb{R}_+^N \rightarrow \mathbb{R}$ are concave and monotonically increasing.¹⁰

⁹In the special case where $D^A(t - \tau) = t - \tau + 1$, the ‘cumulative’ anticipated utility is simply the product of instant utility $u^A(\mathbf{q}_t)$ and the number of periods that precede it (Loewenstein and Prelec, 1993). Our generalization of this structure admits that anticipated utility changes *non-linearly* with the time distance to consumption. A similar remark holds for recollected utility.

¹⁰Given our focus on savoring, we rule out aversive anticipatory emotions such as anxiety (Caplin and Leahy, 2001). Past consumption may also induce a sense of ‘loss’. This is not taken into account by the

We now generalize this logic from an environment with a single consumption bundle to an environment with repeated consumption events. Consider an agent who plans \mathbf{q}_t for $t \in \{\tau, \dots, T\}$. The total utility from anticipation of a consumption stream $(\mathbf{q}_t)_{t \geq \tau}$ is the sum of cumulative anticipated utilities:

$$U^A((\mathbf{q}_t)_{t \geq \tau}) = \sum_{t \geq \tau} D^A(t - \tau) u^A(\mathbf{q}_t). \quad (2)$$

Next, the total utility from recall is the sum of cumulative recollected utilities:

$$U^R((\mathbf{q}_t)_{t \geq \tau}) = \sum_{t \geq \tau} D^R(T - t) u^R(\mathbf{q}_t). \quad (3)$$

Finally, the total utility from experience integrates over all (standard) event utilities $u^E(\mathbf{q}_t)$, similar to the objective function of the lifecycle model:

$$U^E((\mathbf{q}_t)_{t \geq \tau}) = \sum_{t \geq \tau} u^E(\mathbf{q}_t). \quad (4)$$

The total utilities from anticipation (2), recall (3), and experience (4) differ in two main ways. The first distinction lies in the asymmetric weighting of consumption over time. Utility functions from anticipation, respectively recall, attach more weight to consumption at the end (start) of the planning period. This asymmetry is reflected in the properties of the maps D^A (increasing in t) and D^R (decreasing in t). Functions $D^A(\cdot)$ and $D^R(\cdot)$ may also differ in form. The second distinction lies in the marginal rates of substitution between goods. Utility functions can differ in arguments and in shape. This accounts for heterogeneity in anticipatory emotions and enjoyable memories across commodities. The utility function from anticipation $u^A(\cdot)$ may for instance attach more weight to vacations and holidays and less weight to convenience goods from the supermarket. We do not a priori restrict commodities to be ‘anticipatory’, ‘memorable’, or ‘ordinary’. In fact, in our setting, each commodity can have elements of all the temporal motives.

The natural next question is how consumers aggregate their preferences for anticipation, recall, and experience. In the presence of anticipation and given the conflicting temporal motives, such an aggregation might not be trivial. We will assume that consumption is the outcome of a bargaining process between an anticipating self (with preferences U^A),

baseline model, which focuses on ‘pleasant’ memories. A small modification suffices to incorporate this. The function $-\tilde{u}^R(\mathbf{q}_t)$ can capture the negative impact of memories on utility. To keep the problem convex, assume that $\tilde{u}^R(\cdot)$ is subdifferentiable. The drawback is that u^A and this novel definition of \tilde{u}^R have almost identical testable implications. Anticipatory emotions and ‘spiteful’ memories push consumption in exactly the same direction (i.e., postpone consumption).

a remembering self (with preferences U^R), and an experiencing self (with preferences U^E) within the consumer. In that sense, at a given decision moment τ , the selves are assumed to behave cooperatively and select an efficient allocation of resources (Cherchye et al., 2020). We do admit that the respective decision weights of the selves depend on the time of decision $\tau \in \mathcal{T}$, thus allowing a form of limited commitment between the selves. The consumer's objective function in decision moment τ is then

$$V((\mathbf{q}_t)_{t \geq \tau}, \tau) = \omega^A(\tau) U^A((\mathbf{q}_t)_{t \geq \tau}) + \omega^R(\tau) U^R((\mathbf{q}_t)_{t \geq \tau}) + U^E((\mathbf{q}_t)_{t \geq \tau}) \quad (5)$$

where maps $\omega^i : \mathcal{T} \rightarrow \mathbb{R}$ determine the decision power of each self.

In sum, a consumer solves the following optimization problem in decision period τ :

$$\begin{aligned} \max_{(\mathbf{q}_t)_{t \geq \tau}} V((\mathbf{q}_t)_{t \geq \tau}, \tau), \text{ subject to} \\ (\mathbf{q}_t)_{t \geq \tau} \in \mathcal{B}((\boldsymbol{\rho}_t)_{t \geq \tau}, y_\tau), \end{aligned} \quad (6)$$

where $y_\tau = y - \sum_{t=1}^{\tau-1} \boldsymbol{\rho}_t \mathbf{q}_t$ captures the consumer's residual income in period τ , conditional on past expenditures. The associated first order conditions then read as follows:¹¹

$$\begin{aligned} \omega^A(\tau) D^A(t - \tau) \partial u^A(\mathbf{q}_t^*) + \omega^R(\tau) D^R(T - t) \partial u^R(\mathbf{q}_t^*) + \partial u^E(\mathbf{q}_t^*) = \lambda_\tau \boldsymbol{\rho}_t; \text{ and} \\ (\mathbf{q}_t^*)_{t \geq \tau} \in \mathcal{B}((\boldsymbol{\rho}_t)_{t \geq \tau}, y_\tau). \end{aligned} \quad (7)$$

The left hand side of (7) denotes the marginal utility of consumption at time t , seen from moment τ . This can be decomposed in three main effects: the marginal utility from savoring, memories, and experience. A necessary condition for \mathbf{q}_t^* to lie on the optimal path is that these marginal utilities (in money terms) sum up to the market price. Parameter λ_τ captures the effect of relaxing the budget constraint, i.e., the effect of a marginal change in the resources y_τ available for the remaining period $[\tau, T]$.

2.2 Reverse time inconsistency

The DU framework has strong mathematical appeal because its predicted choices are time consistent. As time moves forward, the consumer has no incentives to diverge from her original plan (e.g., chosen in $\tau = 1$). In this section, we show that time consistency generally does *not* hold for consumption with anticipating, remembering, and experiencing selves. In particular, compare the consumer's problem in decision moment τ , given by (6), with a

¹¹We use notation $\partial u^i(\mathbf{q}_t)$ to denote the superdifferential of u^i ($i = A, R, E$), following Rockafellar (1970). In case u^i is differentiable, the superdifferential equals the gradient of the function.

similar problem when the consumer would make consumption choices at time $\tau' > \tau$:

$$\begin{aligned} & \max_{(\mathbf{q}_t)_{t \geq \tau'}} V((\mathbf{q}_t)_{t \geq \tau'}, \tau'), \text{ subject to} \\ & (\mathbf{q}_t)_{t \geq \tau'} \in \mathcal{B}((\boldsymbol{\rho}_t)_{t \geq \tau'}, y_{\tau'}), \end{aligned} \tag{8}$$

We now first formalize the notion of dynamic consistency. In particular, let $(\mathbf{q}_t^*)_{t \geq \tau}$ be the solution to (6) and $(\hat{\mathbf{q}}_t)_{t \geq \tau'}$ the solution to (8). The consumer is then said to be dynamically consistent if $\mathbf{q}_t^* = \hat{\mathbf{q}}_t$ for all $t \in [\tau', T]$. Note that, if the consumer behaves in a time consistent way, we can find consumption choices by solving (6) for $\tau = 1$. There are however several differences between (6) and (8). The most important of these is that there is a change in the remaining duration of savoring associated with a consumption bundle in any future period. More specifically, $D^A(t - \tau') < D^A(t - \tau)$ because $\tau' > \tau$ and D^A is an increasing function. The consumer will enjoy *less* benefits from anticipation of consumption bundle \mathbf{q}_t because the duration of savoring has become shorter. She will therefore have an incentive to reduce (or further postpone) this consumption. Loewenstein (1987) called this ‘*reverse time inconsistency*’. Moreover, at the same time, there is a possible shift in the relative decision power of the anticipating self and the remembering self, $\omega^i(\tau') \neq \omega^i(\tau)$, for $i = A, R$. This can give rise to further time inconsistency problems on the part of the consumer. Proposition 1 formalizes this (general) time inconsistency property. We prove Proposition 1 by offering a counterexample to time consistency in Appendix A.

Proposition 1 (Time inconsistency) *Consider any $\tau, \tau' \in \mathcal{T}$, with $\tau < \tau'$. Let $(\mathbf{q}_t^*)_{t \geq \tau}$ be the solution to (6) and let $(\hat{\mathbf{q}}_t)_{t \geq \tau'}$ be the solution to (8), with $y_{\tau'} = y_\tau - \sum_{t=\tau}^{\tau'-1} \boldsymbol{\rho}_t \mathbf{q}_t^*$. Then we do not necessarily have that $\hat{\mathbf{q}}_t = \mathbf{q}_t^*$ for all $t \in [\tau', T]$.*

2.3 The increasing influence of the anticipating self

Researchers interested in anticipation face issues of reverse time inconsistency. First, anticipating agents will generally, when the time leading up to the actual consumption shrinks, want to lower this consumption and reshuffle plans accordingly. This has blocked the use of standard decision theoretic tools to analyze anticipation. Second, and more fundamentally, acts of reverse time inconsistency give rise to concerns about self-credibility. If agents delay consumption of a given commodity indefinitely—to maintain a sense of savoring—why would they still expect to consume this good at all? Loewenstein (1987) argued that acts of reverse time inconsistency ultimately come at a cost: they interfere with the agents’ ability to savor *any* form of future consumption. These issues may be responsible for the lack of attention for savoring in the literature.

One way to circumvent the inconsistency is to assume *external* commitment devices that bind consumers to their plans. In laboratory experiments, for instance, subjects commit to chosen plans by design (Andreoni and Sprenger, 2012; Echenique et al., 2020). In a social context, consumption often requires organization and reservation (e.g., going out for drinks or dinner, buying tickets for the theater) thereby involving slightly more effort and commitment. Revising holiday plans may entail large cancellation costs. However, limiting anticipation and savoring to circumstances with clear reinforcement mechanisms seems overly restrictive (i.e., it would rule out anticipation for most types of nondurable consumption).

Anticipation is prevalent in everyday life. This suggests that other *internal* mechanisms mitigate acts of reverse time inconsistency and secure self-credibility.¹² We consider such internal mechanism to mitigate the reverse time inconsistency problem. To that end, when we study (5) again in detail, we can see that the anticipated utilities are weighted by two different factors: (i) the temporal shape of preferences, as represented by $D^A(t - \tau)$, and (ii) the decision weight $\omega^A(\tau)$ of the anticipating self at time τ . The interplay of both these factors determines the consumer's incentives to revise planned consumption. Proposition 2 presents the conditions of an internal mechanism that guarantees time consistent behavior.

Proposition 2 (Time consistency) *Consider any $\tau, \tau' \in \mathcal{T}$, with $\tau < \tau'$. Let $(\mathbf{q}_t^*)_{t \geq \tau}$ be the solution to (6) and let $(\hat{\mathbf{q}}_t)_{t \geq \tau'}$ be the solution to (8), with $y_{\tau'} = y_\tau - \sum_{t=\tau}^{\tau'-1} \rho_t \mathbf{q}_t^*$. Then $\hat{\mathbf{q}}_t = \mathbf{q}_t^*$ for all $t \in [\tau', T]$ if*

$$\log \omega^A(\tau) = -\log D^A(t - \tau) + \alpha(t), \text{ for all } t \geq \tau, \text{ and } \tau \in \mathcal{T}, \quad (9)$$

for some mapping $\alpha : \mathbb{R}_+ \rightarrow \mathbb{R}$, and $\frac{\partial \omega^R(\tau)}{\partial \tau} = 0$, for all $\tau \in \mathcal{T}$.

The conditions in Proposition 2 are sufficient for consumers to behave in a dynamically consistent way. When (9) holds, the effect of a decrease in the duration of savoring is exactly neutralized by an increase in the decision power of the anticipating self. This perfectly balances the two main forces that give rise to dynamic inconsistency issues in our set-up: *reverse time inconsistency* associated with anticipation itself and *limited commitment* with respect to ‘bargaining’ between the selves.

One qualitative restriction (9) imposes on the planning process is that the anticipating self's decision weight is increasing over time, i.e., $\frac{\partial \omega^A(\tau)}{\partial \tau} > 0$. Such an increasing influence of anticipation on a consumer's decision making can be fruitfully compared to other, well-established behavioral patterns found and discussed in the literature. For example, notice

¹²In this respect, Benhabib and Bisin (2005) posit that the human brain consists of ‘automatic’ and ‘controlled’ processes. The latter can be used to exert internal commitment to earlier made consumption decisions, thereby overruling the temptation to alter consumption triggered by more automatic processes in the brain.

that an increase in τ shortens the remaining planning period $[\tau, T]$, thereby drawing more attention to the ‘*sequential nature*’ of the choice. This then activates the motives for anticipation (preferences for improvement), due to an increase in the perceived ‘integrity’ aspects of outcome sequences (Loewenstein and Prelec, 1993). Such a negative correlation between the duration of the planning period and a preference for improvement (anticipation) has also been shown elsewhere. For example, Chapman (1996) conducted a study where subjects were asked about their preferences over different sequences of outcomes, varied by domain (health or money) and duration (lifetime or shorter, e.g., 1 year). The results suggested that both mattered for anticipatory preferences. In particular, though over the lifetime subjects seem to be relatively coherent with the notion of impatience, for shorter periods subjects had a more pronounced preference for improvement, indicating perhaps a stronger influence of anticipation. In a similar vein, Castillo et al. (2022) have shown that a larger fraction of subjects indicated a preference for improving sequences when the array size (the length of the sequence) decreased.

There are several tangential behavioral motives underpinning such anticipatory preferences *in sequences*. One of these motives is *loss aversion*, which implies that improving sequences or anticipatory behavior lead to a continual stimulation of positive benefits, in contrast with a declining sequence of outcomes in which a consumer experiences continuous losses. The loss aversion motive is also closely related to the so-called ‘contrast effect’, where present outcomes are evaluated using a comparison with the past or future. Given that the strength of past comparisons is plausibly stronger (Prelec and Loewenstein, 1991), the contrast effect tends to stimulate preferences for anticipation (improvements over time). Another phenomenon underlying higher degrees of anticipation is the *recency effect*, in which evaluations of outcomes depend on the most recent experience. It is therefore beneficial to place desirable outcomes later in the planning period.

An important advantage of Proposition 2 is that it allows us to study consumption behavior by solving (6) for $\tau = 1$. In fact, given condition (9), the consumer’s overall utility (5) can also be written as:

$$\tilde{V}((\mathbf{q}_t)_{t \geq 1}) = V((\mathbf{q}_t)_{t \geq 1}, 1) = \sum_{t \geq 1} [a(t) u^A(\mathbf{q}_t) + b(t) u^R(\mathbf{q}_t) + u^E(\mathbf{q}_t)], \quad (10)$$

where $a(t) = \omega^A(\tau) D^A(t - \tau)$, and $b(t) = \omega^R(\tau) D^R(T - t)$. Given that D^A is increasing in the residual time $t - \tau$ between the consumption event and the decision moment, we have that a is increasing in t . Furthermore, given that D^R is increasing in the time distance $T - t$ between the end of the planning period and the consumption event, we have that b is decreasing in t .

An interesting feature of the model is that, in contrast to the discounted utility (DU) model, the overall utility of the consumer is non-stationary, as evidenced from the form in (10). This highlights the fact that stationarity of the objective function is not necessary for dynamic consistency. This point has also been made recently in the contribution by Drouhin (2020), who characterized all felicity functions that yield dynamically consistent choices. The form of \tilde{V} is consistent with this class of utility functions.¹³

3 Implementation of the theory

3.1 Empirical characterization

We now turn to an empirical characterization of the model. We have access to a (finite) dataset of prices, interest rates, and consumption choices across multiple commodities; $\mathcal{D} = \{\mathbf{p}_t, r_t, \mathbf{q}_t\}_{t \in \mathcal{T}}$. The goal is to test whether the observed choices, together with the prices and interest rates, are consistent with the implications of the model. The characterization is based on revealed preference theory. The revealed preference approach has two main advantages. First, it allows us to test the intertemporal behavior of each consumer separately, thus maximally accounting for preference heterogeneity. Second, it imposes no parametric specification on utility functions $u^i(\cdot)$ of the different selves. We can run our tests without specifying, *a priori*, which goods are partly anticipatory or partly memorable in nature.

In principle, beyond the restriction in Proposition 2, the time functions D^A and D^R can also remain unspecified. It is however common in the revealed preference approach to choose a functional form for time discounting. A functional specification is useful for practical purposes (i.e., to reduce the curse of dimensionality) and because it improves empirical tractability. We opt for an exponential form: the accumulation of anticipated utilities from \mathbf{q}_t over time is given by $D^A(t - \tau) = (\beta^A)^{t-\tau+1}$. Similarly, the compounding of recollected utilities from \mathbf{q}_t is captured by $D^R(T - t) = (\beta^R)^{T-t+1}$. Thus parameter $\beta^A \geq 1$ (respectively $\beta^R \geq 1$) determines the extent to which instantaneous utilities from anticipation (resp. remembering) accumulate and carry over to period τ (resp. period T). Note that, in the limiting case where $\beta^i = 1$, the position of \mathbf{q}_t in the sequence no longer matters. To satisfy condition (9) (needed for dynamic consistency) our parametrization of compounding psychological utility implies that the decision weight of the anticipating self takes the form $\omega^A(\tau) = \omega_0^A (\beta^A)^{\tau-1}$, while the decision weight of the remembering self is fixed at $\omega^R(\tau) = \omega_0^R$. Notice that ω_0^A and ω_0^R reflect the ‘initial’ decision weights at the start of

¹³Drouhin uses a continuous time, single consumption good, representative agent framework with a capital accumulation constraint. The class of time consistent utility functions is of the form $u(q, \tau, t) = d(\tau)u(q, t) + \gamma(t, \tau)$.

$\tau = 1$.

Consumption choices are made as a solution to maximizing (10), subject to the intertemporal budget constraint $(\mathbf{q}_t^*)_{t \geq 1} \in \mathcal{B}((\boldsymbol{\rho}_t)_{t \geq 1}, y)$. The associated first order conditions are then:

$$\begin{aligned} \omega_0^A (\beta^A)^t \partial u^A(\mathbf{q}_t^*) + \omega_0^R (\beta^R)^{T-t+1} \partial u^R(\mathbf{q}_t^*) + \partial u^E(\mathbf{q}_t^*) &= \lambda \boldsymbol{\rho}_t; \text{ and} \\ (\mathbf{q}_t^*)_{t \geq 1} &\in \mathcal{B}((\boldsymbol{\rho}_t)_{t \geq 1}, y). \end{aligned} \quad (11)$$

Clearly (11) are very similar to (7), with the difference that now we have specified D^A, D^R (and $\omega^A(\tau), \omega^R(\tau)$). Since the chosen specification satisfies dynamic consistency condition (9) by construction, any consumption plan $(\mathbf{q}_t^*)_{t \geq 1}$ that solves system (11) will also solve subsequent optimization problems for decision periods $\tau > 1$.

To ease reading, we simplify the notion of ‘intertemporal consumption with anticipating, remembering, and experiencing selves, and exponential time factors satisfying (9)’ to ‘rationalizability by *ICARES*’. Definition 1 formalizes our notion of rationalizability by *ICARES*. We also define three polar cases with stronger testable implications: intertemporal consumption with anticipating and experiencing selves (*ICAES*), remembering and experiencing selves (*ICRES*), and an experiencing self only (*ICES*).

Definition 1 (Rationalizability) Consider the dataset $\mathcal{D} = \{\mathbf{p}_t, r_t, \mathbf{q}_t\}_{t \in \mathcal{T}}$.

- A consumption plan $(\mathbf{q}_t)_{t \geq 1}$ is rationalizable by *ICARES* if there exist monotone and concave functions u^A, u^R, u^E , decision weights ω_0^A, ω_0^R , and time factors $\beta^A, \beta^R \geq 1$ so that condition (11) holds.
- A consumption plan $(\mathbf{q}_t)_{t \geq 1}$ is rationalizable by *ICAES* if there exist monotone and concave functions u^A, u^E , a decision weight ω_0^A , and a time factor $\beta^A \geq 1$ so that condition (11) holds with $\omega_0^R = 0$.
- A consumption plan $(\mathbf{q}_t)_{t \geq 1}$ is rationalizable by *ICRES* if there exist monotone and concave functions u^R, u^E , a decision weight ω_0^R , and a time factor $\beta^R \geq 1$ so that condition (11) holds with $\omega_0^A = 0$.
- A consumption plan $(\mathbf{q}_t)_{t \geq 1}$ is rationalizable by *ICES* if there exists a monotone and concave function u^E so that condition (11) holds with $\omega_0^A = \omega_0^R = 0$.

A few remarks are in order. First, following Browning (1989), Crawford (2010), De-muynck and Verriest (2013) and Adams et al. (2014), we define rationalizability in terms of consistency with the first order conditions of the associated model. Second, the revealed

preference approach allows researchers to test a model for each consumer separately. This is equally reflected in Definition 1, where the data $\mathcal{D} = \{\mathbf{p}_t, r_t, \mathbf{q}_t\}_{t \in \mathcal{T}}$ is a time series for one consumer. In principle, each consumer is characterized by a unique combination of selves (ω_0^A, ω_0^R) and a unique set of utility functions associated with experience, anticipation and/or remembering.

We define shadow price vectors $\tilde{\mathbf{p}}_t^i = \omega_0^i / \lambda \times \partial u^i(\mathbf{q}_t)$ for $i = A, R, E$. Using these prices, we formulate our main result. The proofs are in Appendix A.

Proposition 3 *Consider the dataset $\mathcal{D} = \{\mathbf{p}_t, r_t, \mathbf{q}_t\}_{t \in \mathcal{T}}$. A consumption plan $(\mathbf{q}_t)_{t \in \mathcal{T}}$ is rationalizable by ICARES if and only if there exist numbers u_t^A, u_t^R, u_t^E , nonnegative shadow prices $\tilde{\mathbf{p}}_t^A, \tilde{\mathbf{p}}_t^R, \tilde{\mathbf{p}}_t^E$, and time factors $\beta^A, \beta^R \geq 1$ such that for all $s, t \in \mathcal{T}$:*

$$u_s^A - u_t^A \leq \tilde{\mathbf{p}}_t^A \cdot (\mathbf{q}_s - \mathbf{q}_t); \quad (12)$$

$$u_s^R - u_t^R \leq \tilde{\mathbf{p}}_t^R \cdot (\mathbf{q}_s - \mathbf{q}_t); \quad (13)$$

$$u_s^E - u_t^E \leq \tilde{\mathbf{p}}_t^E \cdot (\mathbf{q}_s - \mathbf{q}_t); \quad (14)$$

$$(\beta^A)^t \times \tilde{\mathbf{p}}_t^A + (\beta^R)^{T-t+1} \times \tilde{\mathbf{p}}_t^R + \tilde{\mathbf{p}}_t^E = \boldsymbol{\rho}_t. \quad (15)$$

In addition, ICAES, respectively ICRES, requires that $\tilde{\mathbf{p}}_t^R = 0$ ($\tilde{\mathbf{p}}_t^A = 0$). Finally, ICES requires that $\tilde{\mathbf{p}}_t^A = \tilde{\mathbf{p}}_t^R = 0$.

Proposition 3 shows that our definitions of rationalizability have equivalent representations in terms of data consistency with systems of (in-)equalities. Conditions (12)–(14) are Afriat-style inequalities and stem from the properties of the utility functions (in particular concavity). Conditions (15) reflect the first order conditions (11). The additional structure imposed by ICAES, ICRES, and ICES comes in the form of straightforward restrictions on the shadow price vectors $\tilde{\mathbf{p}}_t^A$ and $\tilde{\mathbf{p}}_t^R$.

Let us now study the conditions of Proposition 3 in more detail. First, perhaps surprisingly, ICARES does not nest (repeated) static utility maximization. The empirical content of the utility maximization hypothesis is exhausted, in a revealed preference sense, by the generalized axiom of revealed preferences (GARP). More formally,

Corollary 1 *ICARES and GARP have independent testable implications.*

Second, the RP conditions in Proposition 3 resemble the RP characterization for additive separability by Varian (1983). Our additive structure stems from two modeling choices: (i) we ‘decompose’ the consumption problem by means of subutility functions from anticipation, recall, and experience, (ii) the total subutility from each temporal motive is the sum of the

corresponding (discounted) utility flows across consumption events. We acknowledge that the summation in (ii) is in principle restrictive, but our model is still a strict generalization of the lifecycle model which also sums up utilities from consumption events. [Varian \(1983\)](#) derived necessary and sufficient conditions for rationalizability of the data by a sum of utility functions. Our characterization has three distinguishing features: the marginal utility of wealth λ is constant (and hence absorbed in utilities u_t^i and shadow prices $\tilde{\mathbf{p}}_t^i$), we have a *weighted* sum of utilities with *time-dependent* weights $(\beta^A)^t$ and $(\beta^R)^{T-t+1}$, and the same bundle \mathbf{q}_t can enter as an argument in all three subutility functions. The latter property demonstrates that our framework does not impose weak (or even latent, [Blundell and Robin \(2000\)](#)) separability between anticipatory, memorable, and ordinary consumption goods.

The additive structure of our model facilitates implementation to budget survey data. The contribution of the (instantaneous) utility from anticipation, remembering, or experience of a *given* consumption event to overall welfare is independent of the consumption events before/after that. Thus the revealed preference results will be robust to other consumption events that may fall within the planning period but outside our window of observations; we refer to section 3.3 for more details. The additive structure also helps to keep the problem empirically tractable. Conditional on β^A and β^R , the system of inequalities (12)-(15) is linear in the unknowns (the utility numbers and shadow prices), which makes testing for rationalizability by *ICARES* computationally manageable. In practice, we run a grid search on β^A and β^R . A final attractive feature of our set-up is that preferences for anticipation and preferences for recall have separate testable implications. This is in sharp contrast to the more general ‘non-separable’ preference structures discussed in [Crawford and Polisson \(2014\)](#). The asymmetric roles of temporal selves, combined with the separability properties of *ICARES*, enable us to empirically distinguish anticipated from recollected utilities. Corollary 2 formalizes the independence of *ICAES* and *ICRES*.

Corollary 2 *ICAES and ICRES have independent testable implications.*

3.2 Implications for consumption patterns: Simulations

We now use our behavioral model to simulate intertemporal consumption patterns. The simulated behavior will illustrate key features of the model. Moreover, the simulations will offer insight into the typical data patterns needed to empirically disentangle *ICAES* from *ICRES*, *ICAES* or *ICRES* from *ICARES*, and *ICARES* from static models of utility maximization or models of hyperbolic discounting.

In all our simulations, we set $T = 4$ and we restrict attention to $N = 2$ commodities for

simplicity. We use utility functions of the (logarithmic) Cobb-Douglas form:

$$u^i(q_{1,t}, q_{2,t}) = \alpha^i \log(q_{1,t}) + (1 - \alpha^i) \log(q_{2,t}) \text{ where } i = A, R, E. \quad (16)$$

Finally, the intertemporal budget constraint is

$$\sum_{t=1}^4 [\rho_{1,t} q_{1,t} + \rho_{2,t} q_{2,t}] = 8.$$

We will present 3 scenarios. In scenario I we analyze the differential consumption patterns generated by *ICARES* and its special cases. Scenario II compares the consumption patterns from *ICARES* to those from static utility maximization and models of (quasi-)hyperbolic discounting. We highlight how we can empirically distinguish between all these models. Finally, scenario III analyzes the nestedness of *ICARES* specifications with different time factors β^A and β^R .

Scenario I. In this scenario, we set $\alpha^E = 0.2$ and $\alpha^A = \alpha^R = 1$. In words, good 1 produces little utility from experience, while good 2 yields no utility from anticipation or recall. The extent to which the anticipated (recollected) utilities of consumption accumulate through time and carry over to earlier (later) periods is given by a time factor $\beta^A = \beta^R = 1.2$. The consumer faces the following prices, where each row corresponds to a time period t :

$$\begin{pmatrix} \mathbf{p}_t \\ (1.025, 1) \\ (1, 1) \\ (1, 1) \\ (1.025, 1) \end{pmatrix}$$

We then use these prices and the overall resources available, $y = 8$, to simulate consumption behavior according to *ICES* ($\omega_0^A = \omega_0^R = 0$), *ICAES* ($\omega_0^A = 1, \omega_0^R = 0$), *ICRES* ($\omega_0^A = 0, \omega_0^R = 1$), and *ICARES* ($\omega_0^A = \omega_0^R = 1$). Table 3.1 reports the consumption patterns for good 1 (first column) and good 2 (second column) across different time periods (rows). Under *ICES*, consumption is more or less smoothed over time. By contrast, *ICAES*, respectively *ICRES*, predicts increasing (decreasing) consumption of good 1. *ICARES* finally produces a small increase in consumption of good 1 at the start *and* at the end of the sequence.

We now apply the tests of Proposition 3 to the simulations of Table 3.1. The smoothed consumption profile (generated with $\omega_0^A = \omega_0^R = 0$) satisfies the revealed preference charac-

Data generated by <i>ICES</i>	Data generated by <i>ICAES</i>	Data generated by <i>ICRES</i>	Data generated by <i>ICARES</i>
$\begin{pmatrix} \mathbf{q}_t \\ (0.3902, 1.6) \\ (0.4, 1.6) \\ (0.4, 1.6) \\ (0.3902, 1.6) \end{pmatrix}$	$\begin{pmatrix} \mathbf{q}_t \\ (1.0465, 0.6129) \\ (1.2565, 0.6129) \\ (1.4772, 0.6129) \\ (1.6995, 0.6129) \end{pmatrix}$	$\begin{pmatrix} \mathbf{q}_t \\ (1.6995, 0.6129) \\ (1.4772, 0.6129) \\ (1.2565, 0.6129) \\ (1.0465, 0.6129) \end{pmatrix}$	$\begin{pmatrix} \mathbf{q}_t \\ (1.6058, 0.3791) \\ (1.5959, 0.3791) \\ (1.5959, 0.3791) \\ (1.6058, 0.3791) \end{pmatrix}$

Table 3.1: Simulated consumption behavior under scenario I

terizations of *ICES*, *ICAES*, *ICRES*, and *ICARES*. This must necessarily hold because the data was generated by *ICES*, and *ICES* is a special case of all other models. Next, the increasing consumption pattern satisfies the conditions of *ICAES* but violates the conditions of *ICRES*, while the decreasing pattern satisfies *ICRES* but violates *ICAES*. This shows that *ICAES* and *ICRES* have separate testable implications, thereby confirming Corollary 2. The final consumption profile (generated with $\omega_0^A = \omega_0^R = 1$) passes the revealed preference conditions of *ICARES* but violates all other (nested) characterizations. Only *ICARES* can rationalize the consumption profile that peaks in the beginning *and* in the end of the sequence—the periods with the highest relative prices.

Scenario II. We also use simulations to investigate the differences between *ICARES* and two other classes of consumption models: static models of utility maximization and dynamic models with *hyperbolic* discounting. Hyperbolic discounting offers a time inconsistent description of intertemporal consumption. The static model is a useful benchmark because it nests many intertemporal models, such as exponential discounting (and hyperbolic discounting).¹⁴

Under the current scenario, we adjust the utility specification as follows. In particular, for the *ICARES* simulations, we will now assume that all anticipation comes from good 2 ($\alpha^A = 0$) whereas for the other models we assume a simple logarithmic form,

$$u(q_1) = \log(q_1)$$

We again assume that total resources over the entire time horizon are given by $y = 8$, and for the static utility maximization case we consider a sequence of incomes $\{1, 3, 1, 3\}$.

¹⁴As originally shown by [Browning \(1989\)](#), testing the lifecycle model under perfect foresight (and in the absence of borrowing constraints) is equivalent to testing a condition called cyclical monotonicity (CM), which is stronger than GARP. For quasi-hyperbolic discounting, it has been shown that its RP characterization implies intra-period consistency with GARP, see [Blow et al. \(2021\)](#).

Prices are given by:

$$\begin{pmatrix} \mathbf{p}_t \\ (1.025, 1) \\ (\sqrt{1.025} + 0.001, 1) \\ (1, \sqrt{1.025} + 0.001) \\ (1, 1.025) \end{pmatrix}$$

We then simulate consumption over $T = 4$ periods for *ICARES*, static utility maximization, and the quasi-hyperbolic discounting model. With regards to the latter, we assume that the hyperbolic planner is sophisticated. She is aware that her future self's preferences over consumption will be different; nonetheless, she cannot commit to any future plans. Consequently, a sophisticated quasi-hyperbolic discounter still behaves in a time inconsistent manner. A full revealed preference analysis of such time inconsistent model has been studied in [Blow et al. \(2021\)](#). We select a discount factor of $\delta = 0.97$ for this consumer. For *ICARES*, we assume a strong weighting of both anticipated and recollected utilities: we set $\omega_0^A = \omega_0^R = 4$. Table 3.2 reports the simulated consumption patterns for good 1 (first column) and good 2 (second column) under scenario II. *ICARES* predicts that consumption of the first—memorable—good decreases systematically while consumption of the second—anticipatory—good increases continuously. The hyperbolic model predicts that consumption of the first good peaks in period 2. Finally, according to the static model, consumption peaks in periods 2 and 4. This is simply because the consumer is more wealthy in these periods.

Data generated by <i>ICARES</i>	Data generated by Hyperbolic model	Data generated by Static utility maximization
$\begin{pmatrix} \mathbf{q}_t \\ (1.1938, 0.8067) \\ (1.0110, 0.9450) \\ (0.8586, 1.0963) \\ (0.7203, 1.2782) \end{pmatrix}$	$\begin{pmatrix} \mathbf{q}_t \\ (1.9961, 0) \\ (1.9983, 0) \\ (1.9943, 0) \\ (1.9345, 0) \end{pmatrix}$	$\begin{pmatrix} \mathbf{q}_t \\ (0.9756, 0) \\ (2.9603, 0) \\ (1, 0) \\ (3, 0) \end{pmatrix}$

Table 3.2: Simulated consumption behavior under scenario II

We now apply revealed preference tests to these simulated data. We first note that the data generated by *ICARES* violates *GARP* (and thereby immediately violates the RP characterization of quasi-hyperbolic discounting). The main reason for this violation is that, in the second and third periods, the consumer buys more of a commodity when it is relatively more expensive. We further note that the data generated by the static model and the hyperbolic model are inconsistent with the revealed preference characterization of *ICARES*.

This is because *ICARES* cannot explain why the consumer buys more in period 2—in the middle of a sequence—when consumption is expensive. For completeness, we also verify that the final increase in consumption in period 4, generated by the static model, violates the revealed preference conditions of the hyperbolic model.

Scenario III. In this scenario we study the effect of changes in the factors β^A and β^R . It can be easily observed that higher values of β^A and β^R allow heavier shifts in the first order conditions from one period to the next. Intuitively, this broadens the range of data patterns that can be rationalized. In this exercise, however, we show that higher values for β^A and β^R *not necessarily* explain the data better. We again select the parametric form specification as in scenario I. Furthermore, we will assume a new time series for prices,

$$\begin{pmatrix} \mathbf{p}_t \\ (1.175, 1) \\ (1.08, 1) \\ (1, 1) \\ (1, 1) \end{pmatrix}$$

We now simulate the *ICRES* model with a parameter value $\beta^R = 1.1$. The other preference parameters are fixed to the values of scenario I.

Data generated by *ICRES*

$$\begin{pmatrix} \mathbf{q}_t \\ (1.2444, 0.7029) \\ (1.2455, 0.7029) \\ (1.2389, 0.7029) \\ (1.1422, 0.7029) \end{pmatrix}$$

Table 3.3: Simulated consumption behavior under scenario III

Table 3.3 shows the results of the simulation. The consumption of good 1 first increases, reaches its peak in period 2, and then decreases until period 4. We first verify that the data indeed satisfy the revealed preference characterization of *ICRES* with $\beta^R = 1.1$. However, the data violate the revealed preference conditions with $\beta^R = 1.2$. Intuitively, the fact that the consumption of good 1 peaks in period 2 (and not in period 1) can only be rationalized by *moderate* levels of remembering.

3.3 Practical considerations

The simulations showed that the theory from section 2 can generate a wide variety of consumption patterns. In a choice experiment in the lab, one could easily bring the RP characterizations of *ICARES*, *ICAES*, and *ICRES* to the data without further adjustments. However, our main objective is to study anticipated and recollected consumption *outside* specially tailored lab environments.

In reality, the consumer will likely solve several successive planning problems over her lifetime. The (short) series of consumption events observed in budget survey data will generally be only a subset of one planning period.¹⁵ The start and end times of the underlying plan are not observed. A special property of our choice of D^A and D^R is that the revealed preference results will be independent of the total duration of the consumption plan. Notice for instance that one can simply redefine $\tilde{u}^R = \beta^T \times u^R$, which absorbs the effect of T . The true bounds of the planning period leave *changes* in the anticipated and recollected utility flows between subsequent consumption observations unaffected. So we can abstract from arbitrary definitions of the true bounds. The downside of this property is that we cannot use our revealed preference characterization to identify elements of the planning period beyond the observed sequence.

We discuss three further complications that may arise in an application to budget survey data. First, the presence of (binding) borrowing constraints impedes on the consumers' possibilities to smooth consumption and may thus lead to an over-rejection of standard exponentially discounted utility models (Dean and Sautmann, 2021). Second, the strong yet common assumption of perfect foresight with respect to prices, interest rates, and income levels may not hold in practice. Lastly, the data may suffer from measurement error. The latter can affect the empirical performance of economic models in RP analyses. For example, Aguiar and Kashaev (2021) have shown that there is a tendency to over-reject the hypothesis of static utility maximization (*GARP*) when one ignores mismeasurement. We will take several steps to test the robustness of our findings to these challenges. We now provide a brief overview of how we plan to tackle these issues.

¹⁵By contrast, for *very* short planning periods ($T = 2$ or $T = 3$ quarters) it is possible that our sequence of observations is longer than one full planning window. One can then test data consistency with *ICARES* for separate *subsets* of the observed sequence. These results do not alter the main conclusions of the following analyses.

Liquidity constraints. It is convenient in this context to reformulate the optimization problem associated with *ICARES* using spot prices. In particular, consider the following:

$$\begin{aligned} & \max_{(\mathbf{q}_t)_{t \geq 1}} \sum_{t=1}^T \left[\omega_0^A (\beta^A)^t u^A(\mathbf{q}_t) + \omega_0^R (\beta^R)^{T-t+1} u^R(\mathbf{q}_t) + u^E(\mathbf{q}_t) \right], \text{ subject to} \\ & s_t \geq -b_t, \text{ and} \\ & \mathbf{p}_t \cdot \mathbf{q}_t + s_t = \tilde{y}_t + (1 + r_{t-1}) s_{t-1}, \text{ for all } t \in \mathcal{T}. \end{aligned} \quad (17)$$

Problem (17) is very similar to the baseline *ICARES* optimization problem, with a few key differences. A first (obvious) difference is that the budget constraint is now expressed in sequential form, where s_t refers to *savings*. Second, the presence of a borrowing constraint, $s_t \geq -b_t$, puts a limit on the amount of debt a consumer can incur. Finally, $(\tilde{y}_t)_{t \in \mathcal{T}}$ is the sequence of income levels. The first order conditions are as follows:

$$\omega_0^A (\beta^A)^t \partial u^A(\mathbf{q}_t) + \omega_0^R (\beta^R)^{T-t+1} \partial u^R(\mathbf{q}_t) + \partial u^E(\mathbf{q}_t) = \tilde{\lambda}_t \mathbf{p}_t, \quad (18)$$

$$\tilde{\lambda}_t = (1 + r_t) \tilde{\lambda}_{t+1} + \tilde{\mu}_t. \quad (19)$$

These conditions are relatively standard. Condition (18) is clearly similar to the first order optimality conditions of the *ICARES* model *without* liquidity constraints. The main difference between (18) and (11) is the fact that shadow price $\tilde{\lambda}_t$ is now the Lagrange multiplier of the sequential budget constraint, and the price level refers to spot prices rather than discounted prices. The condition in (19) is a consumption Euler equation which takes into account the presence of the liquidity constraint with Lagrange multiplier $\tilde{\mu}_t$. In case the borrowing constraint is not binding, we would have $\tilde{\mu}_t = 0$ and (19) collapses to the standard consumption Euler equation. If the liquidity constraint does bind, $\tilde{\mu}_t > 0$. We can rewrite (19) more succinctly as follows:

$$\tilde{\lambda}_t \geq (1 + r_t) \tilde{\lambda}_{t+1}. \quad (20)$$

We adjust our definition of shadow prices slightly: $\tilde{\mathbf{p}}_t^i = \omega_0^i \times \partial u^i(\mathbf{q}_t)$ with $i = A, R, E$ and $t \in \mathcal{T}$. The RP characterization of problem (17) will be very similar to the RP conditions in Proposition 3, but the Lagrange multipliers $\tilde{\lambda}_t$ associated with the budget constraints will now explicitly enter the conditions to be verified for rationalizability. In particular, the sequence $(\tilde{\lambda}_t)_{t \in \mathcal{T}}$ will have to satisfy the monotonicity condition in (20). Since the interest rates are observable, the RP characterization remains computationally tractable in the sense that, conditional on the β s, the system of RP conditions is linear in its unknowns.

Uncertainty. Another concern with RP tests to budget survey data is the assumption of perfect foresight with regards to the economic environment. Without this assumption, or further stronger assumptions about the expectational process of agents, RP tests will lack empirical content. Specifically, the smoothing of marginal utilities from wealth ($\tilde{\lambda}_t = (1 + r_t) \tilde{\lambda}_{t+1}$) may no longer hold after a series of large unexpected shocks to the sequence of income levels $(\tilde{y}_t)_{t \in \mathcal{T}}$. Ex ante, the conditions could be adjusted in the following way,

$$\tilde{\lambda}_t = \mathbb{E} \left[(1 + r_t) \tilde{\lambda}_{t+1} \right], \quad (21)$$

but ex post, there may be strong variation of the marginal utilities of wealth $\tilde{\lambda}_t, \tilde{\lambda}_{t+1}$. We therefore adapt our conditions as follows,

$$\alpha (1 + r_t) \tilde{\lambda}_{t+1} \leq \tilde{\lambda}_t \leq \frac{1 + r_t}{\alpha} \tilde{\lambda}_{t+1}, \quad (22)$$

where $\alpha \in (0, 1]$. Note that the test becomes weaker in case $\alpha < 1$, which captures the possibility that $\tilde{\lambda}_t = (1 + r_t) \tilde{\lambda}_{t+1}$ is violated due to unobserved randomness (e.g., in income flows $\tilde{y}_t, t \in \mathcal{T}$). In practice, we will implement several tests of rationalizability for given values of α on a grid, thus varying the degree of uncertainty.

Measurement error. As a final extension, we admit that our data is affected by measurement error. Similar to the case of unobserved randomness due to a lack of foresight, the presence of measurement error implies that we need to weaken our ('sharp')¹⁶ tests for rationalizability by *ICARES*, *ICAES*, *ICRES*, *ICES*. This allows for a wider range of uncertainty from the econometrician's point of view. To explore the issue of measurement error in our context, we use a procedure similar to the one in [Adams et al. \(2014\)](#), which is itself based on an idea of [Varian \(1985\)](#). To be more precise, assume that the true prices are given by ρ_t^* , for $t \in \mathcal{T}$. Then suppose we quantify the error $\varepsilon_{n,t}$, for each good $n \in \{1, \dots, N\}$ and for each period $t \in \mathcal{T}$, as the divergence between true and observed prices:

$$\rho_{n,t}^* = \rho_{n,t} + \varepsilon_{n,t}.$$

Clearly, the sequence of errors $\varepsilon_{n,t}$ is not observed as we do not observe the true prices. Our procedure will therefore select values $\tilde{\varepsilon}_{n,t}$ from which we define the following altered prices,

$$\tilde{\rho}_{n,t} = \rho_{n,t} + \tilde{\varepsilon}_{n,t}.$$

¹⁶In the sense of [Demuyne and Verriest \(2013\)](#), [Adams et al. \(2014\)](#), and [Blow et al. \(2021\)](#).

The main concern behind measurement issues is the over-rejection of stringent models, such as the lifecycle model. We therefore choose the smallest perturbations $\tilde{\varepsilon}_{n,t}$ that can still rationalize a given fraction of (single) consumers in our dataset by the standard lifecycle model.¹⁷ We can then test rationalizability for *ICARES*, *ICAES*, and *ICRES* by applying the conditions in Proposition 3 to the adjusted dataset $\{\tilde{\rho}_t, \mathbf{q}_t\}_{t \in \mathcal{T}}$.

4 Empirical application

In this section we apply our models of intertemporal consumption with anticipating, remembering, and experiencing selves to budget survey data. We start with a discussion of the data from the Encuesta Continua de Presupuestos Familiares (ECPF). We then bring the time consistent versions *ICAES*, *ICRES*, and *ICARES* to the data by means of our revealed preference characterizations. We find strong empirical support for *ICAES*. This shows that dynamic inconsistency, as in Adams et al. (2014) or Blow et al. (2021), is not the only possible mechanism that can explain intertemporal consumption data. It suffices to allow for flexible forms of anticipation and remembering. Next, we identify which goods have a strong anticipatory element. We classify goods on the basis of their anticipatory nature. We finish this section with several generalizations of our analysis: including liquidity constraints, uncertainty, and measurement error. Agents may also be more likely to revise their consumption plans around New Year (Cherchye et al., 2020). We therefore study a specific subsample of consumers for whom the observed sequence starts precisely in January and ends in December.

4.1 Data description

For our application, we use data from the Encuesta Continua de Presupuestos Familiares (ECPF) collected by the Spanish Statistics Office (INE). We selected this data for two main reasons. First, consumption data is collected in every quarter. We focus on respondents who reported their expenditure in four *consecutive* quarters between 1985 and 1996.¹⁸ The quarterly nature of the data is important. After all, our objective is to describe consumption planning decisions in the short-to-medium run. The consumer’s entire lifecycle can be composed of several of these planning periods (i.e., preferences can change with age). The

¹⁷We opted for singles as our target group to construct the minimal perturbations because it has been shown that singles are more likely to satisfy standard Euler equations, see for example Mazzocco (2007). Similarly, in a revealed preference context, Adams et al. (2014) found that singles are more likely to pass the restrictions of the lifecycle model.

¹⁸Respondents are rotated at a rate of 12.5% each quarter. Only 25 respondents in total are observed for more than four successive quarters.

quarterly dimension also limits our assumptions of stationary utility u^i and perfect foresight to a relatively narrow time interval (i.e., one year). We study a relaxation of the perfect foresight assumption in section 4.4. Second, the ECPF is used frequently for nonparametric tests of intertemporal consumption models, see for instance Crawford (2010), Demuynck and Verriest (2013), Adams et al. (2014), and Blow et al. (2021). This allows us to benchmark the empirical performance of *ICARES* against the performance of other intertemporal models.

We implement the sample selection criteria discussed in Adams et al. (2014). In particular, we only keep consumers who completed all four interviews. Furthermore, the sample selection only keeps households that have a fixed number of children over the four interviews, to exclude ‘big shocks’ such as child birth. We also select households in which individuals have a stable employment status over the interviews, again excluding shocks such as job loss. This alleviates issues of nonseparability between time use and consumption. We only keep consumers with positive total expenditures over the four interviews. Seen together, the selection criteria give us a sample of 2,052 consumers (1,880 couples plus 172 singles). This strikes a balance between a sufficiently sizable dataset on one hand and a long enough panel to conduct revealed preference tests on the other hand.

We consider eight nondurable commodities: (i) food and drinks, (ii) clothing and footwear, (iii) household services including heating, water, and furniture repair, (iv) transport, (v) petrol, (vi) leisure including cinema, theatre, and sports, (vii) personal services, and (viii) restaurants and bars. At this point, we also would like to add that the level of aggregation of goods in our empirical application mirrors real-life situations where consumption events are inherently conglomerate in nature, thereby adding to the point that our framework can be useful in real-life datasets. However, we revisit the expenditure on specific subgroups of goods in section 4.3. Table 4.1 reports summary statistics of the expenditure shares of our eight commodities. This shows that 46% of the budget is spent on food, 16% on clothing, and 13% on restaurant costs.

	Sample mean	Sample stdev
Allfood	45.72	19.02
Clothing	16.39	13.46
Hhserv	4.58	6.61
Transport	4.62	7.60
Petrol	7.13	7.79
Leisure	5.67	7.38
Pserv	2.52	4.40
Foodout	13.38	13.09

Table 4.1: Descriptive statistics of expenditure shares (in percent) on eight nondurable commodities

The averages in Table 4.1 shed little light on the degree of improvement over the observed sequence. To help us distinguish special cases of *ICARES*, we now study the temporal profiles of consumption in the data. Expenditures vary *between consumers* and *over time*. Let $w_{ht}^n = p_{ht}^n q_{ht}^n$ denote expenditure on commodity n by consumer h at time t . We regress w_{ht}^n on total intertemporal budget y_h . We also add interactions between y_h and time dummies ($t = 2$), ($t = 3$), and ($t = 4$). We further include 44 dummies to absorb effects of quarterly variation in economic conditions between 1985 and 1996 (Browning and Collado, 2007). We thus exploit independent variation between the calendar date of the data point (i.e., 44 quarterly dummies) and the position of the data point in the sequence of observations (i.e., 3 positional dummies $t = 2, 3, 4$). We finally include a constant. We can then estimate

$$w_{ht}^n = \eta^n + \gamma^n \times y_h + \theta_t^n \times y_h + \sum_{k=5}^{48} \delta_k^n \times (\text{Calendar time}_{ht} = k) + \varepsilon_{ht}^n \quad (23)$$

Parameter γ^n captures income effects when $t = 1$ while coefficients δ_k^n capture the effects of yearly and seasonal variations. The main parameters of interest, θ_t^n , reflect the effect of the data point's position in the sequence of observations for that consumer. Each regression uses 8,208 consumer-observations. We cluster standard errors by consumer. Table 4.2 presents our estimates.

	Allfood	Clothing	Hhserv	Transport	Petrol	Leisure	Pserv	Foodout
allexp	0.079*** (0.003)	0.045*** (0.002)	0.018*** (0.001)	0.014*** (0.001)	0.017*** (0.001)	0.020*** (0.001)	0.009*** (0.001)	0.042*** (0.003)
t=2 × allexp	-0.002 (0.002)	0.000 (0.002)	0.001 (0.001)	0.001 (0.001)	0.000 (0.001)	0.001 (0.001)	0.000 (0.001)	0.003 (0.002)
t=3 × allexp	0.001 (0.002)	-0.001 (0.002)	0.001 (0.001)	0.002 (0.001)	0.001 (0.001)	0.001 (0.001)	0.001 (0.001)	0.002 (0.002)
t=4 × allexp	-0.000 (0.002)	0.000 (0.002)	0.001 (0.001)	0.002 (0.001)	0.001* (0.001)	0.001 (0.001)	-0.000 (0.001)	0.004** (0.002)
Observations	8208	8208	8208	8208	8208	8208	8208	8208

Table 4.2: Regressions of consumption expenditure on total intertemporal budget, interactions with time dummies, and quarterly controls. The sample consists of all consumers. Estimates of quarterly controls are suppressed for compactness.

The positive coefficients of total expenditure ($\gamma^n > 0$) show that all goods are normal. The coefficients θ_t^n associated with $t = 2, 3, 4$ are mostly nonnegative, reflecting a small increase in expenditure towards the end of the observed sequence. However, the differences in consumption are not statistically significant. There are two exceptions: the increase in

fuel in $t = 4$ is significant at the 10% level and the increase in food expenditure away from home in $t = 4$ is significant at the 5% level. The latter suggests that restaurant visits exhibit a strong anticipatory motive. But generally speaking, the consumption profiles are relatively smooth. One possible explanation is that the results in Table 4.2 are averages among consumers with different preferences types. It aggregates the temporal expenditure profiles of consumers with strong anticipating selves and the profiles of consumers with strong remembering selves. Increasing and decreasing profiles may cancel each other out. We therefore turn to our revealed preference method to identify different consumer types.

4.2 Revealed preference tests

We now assess the empirical performance of *ICARES* and its special cases *ICAES*, *ICRES*, and *ICES* through a revealed preference analysis. An important feature of the revealed preference approach is that it allows us to verify data consistency for each consumer separately. This accounts for general unobserved preference heterogeneity between consumers. For each consumer, we test whether their choices satisfy the revealed preference conditions of Proposition 3. The test has as inputs the prices and quantities of eight commodities over some observed *subset* of the planning period. As discussed in section 3.3, the revealed preference results are robust to the time distance between the actual start of the planning period and our first observation, and robust to the time distance between the last observation in our data and the end of the planning period. We finally specify a range of parameter values $\beta^A, \beta^R \in [1, 1.2]$, with stepsize 0.05, for our grid search. This allows $(\beta^A)^t$ and $(\beta^R)^{T-t+1}$ to accumulate to at most 2, representing a doubling of anticipated/recollected utilities for consumption at the very start/end of the sequence. Each test takes the form of a linear programming problem. If the program has a solution for *any* combination of β^A, β^R (e.g., $\beta^A = \beta^R = 1.2$)¹⁹ then we assign a value of one ('pass').

We first compute goodness-of-fit. Averaging over all consumers gives the mean in-sample fit of the model with the data ('pass rate'). Table 4.3 contains the pass rates of *ICES*, *ICAES*, *ICRES*, and *ICARES*. We report pass rates separately for singles and couples because household composition may also affect the decision-making structure of the consumer. Only 9% of the singles and 4% of the couples can be rationalized by *ICES*. This is not surprising: *ICES* is equivalent to the nonparametric test of the lifecycle model without discounting. [Browning \(1989\)](#) showed that this imposes strong restrictions on observed behavior. On the contrary, up to 49% of the singles and 46% of the couples can be rationalized by *ICAES*. This already supports our notion of intertemporal consumption with anticipating and experiencing

¹⁹Higher values of β^i generally improve goodness-of-fit, and the main gains in empirical fit are realized between $\beta^i = 1$ and $\beta^i = 1.2$.

selves. Finally, *ICRES* fits between 80% (couples) and 82% (singles) of the data while *ICARES* fits 95% (singles) to 96% (couples). We also explore the probability that consumers pass *ICAES* but not *ICRES*, and vice versa. We find that 8% of the data satisfy *ICAES* (not *ICRES*) while almost 42% satisfy *ICRES* (not *ICAES*). Corollary 2 is not just a theoretical curiosity: our method separates preferences for anticipation from preferences for recall on the basis of widely available data from budget surveys. In the next subsection, we will return to the subset of consumers consistent with *ICAES* to learn more about the (anticipatory) nature of commodities.

A limitation of assessing the empirical quality of a test merely by its goodness-of-fit is that the latter does not control for the fact that flexible models such as *ICAES*, *ICRES*, and especially *ICARES* more easily rationalize any kind of behavior. The revealed preference literature offers an elegant method to quantify the empirical power of nonparametric tests. Discriminatory power is the probability that a test rejects random behavior. In practice, it is one minus the pass rate of *simulated* random datasets. We follow the majority of papers in this literature and implement Bronars (1987)’s method of power measurement, adapted to an intertemporal setting. In particular, at each iteration b and for each consumer h , we simulate random behavior by drawing 4×8 random budget shares from a uniform distribution on the unit simplex. Using the actual (observed) prices, we can then compute consumption vectors $\left(\mathbf{q}_t^{h,b}\right)_{t \in \mathcal{T}}$ that exhaust the intertemporal budget constraint. This vector, together with prices and interest rates, then serves as a (simulated) dataset $\{\mathbf{p}_t, r_t, \mathbf{q}_t^{h,b}\}_{t \in \mathcal{T}}$. We repeat this procedure 205,200 times.²⁰ We then apply each model to the *same* set of random quantities. Table 4.3 reports our power estimates. As expected, *ICES* is the most powerful model. It rejects consistency for about 95% of the simulated data. *ICAES* is also still powerful, rejecting consistency for 77% of the simulated data. Power drops dramatically for *ICRES* (30%) and *ICARES* (12%).²¹

Given the trade-off between in-sample fit (i.e., psychological realism) and discriminatory power (i.e., empirical tractability), we combine both measures into a single metric of ‘predictive success’. We follow Selten (1991) and Beatty and Crawford (2011) and compute predictive success as the sum of pass rates and power minus one. Predictive success is always situated between -1 and 1 . Higher scores indicate better empirical performance. A predic-

²⁰We simulate $B = 100$ random datasets *per consumer*. We draw expenditure shares from a uniform distribution on the unit simplex.

²¹Given a decreasing path for discounted prices (Figure 4 of Appendix C.1) the most likely candidate to obtain *violations* of *ICES* is a decreasing consumption path. The latter cannot be rationalized by *ICAES* either, which is mainly responsible for generating increasing consumption paths. However, simulated decreasing consumption paths can easily be rationalized by *ICRES*, and this explains why *ICRES* has less discriminatory power (when prices decrease) than *ICAES*. Furthermore, given that the dataset mostly contains increasing consumption profiles, the pass rate of actual data with *ICAES* remains relatively high.

tive success of -1 represents the worst possible outcome: none of the observed behavior but all of the random data can be rationalized. A predictive success of 1 is the best possible result: all of the observed behavior but none of the random data can be rationalized. Finally, a predictive success of 0 suggests that the empirical performance of the model is similar to the performance of a ‘model’ of completely random behavior. The predictive success of the most restrictive model (*ICES*) is small overall, although it is strictly positive (5%) for singles. The most flexible model (*ICARES*) performs a bit better but predictive success is still limited to 7% for singles and 9% for couples. *ICRES* improves further on this empirical performance to 12% for singles and 10% for couples. We find the strongest empirical support for *ICAES*, with predictive success scores that exceed 20% for both singles and couples. This improves further on Selten’s index found for hyperbolic discounting models (around 7% in Blow et al. (2021)) and models in which habits form as durables (around 15% in Demuyneck and Verriest (2013)).²²

The last lines of Table 4.3 report 95%–confidence intervals around the mean predictive success. To construct these intervals, we follow an econometric approach put forward by Demuyneck (2015). The procedure uses as inputs the variance of observed ‘pass’ results among consumers, the variance of simulated power results among consumers, and finally the covariance of the consumers’ pass results and power estimates. For couples, the lower bound on predictive success of *ICAES* (20.5%) exceeds the upper bounds on predictive success of the other characterizations (maximum 11.3%). For singles, the lower bound on predictive success of *ICAES* (19%) is lower due to the small number of singles in the sample. Still, this exceeds the upper bounds on predictive success in all other models (maximum 18%).

	Singles				Couples			
	<i>ICES</i>	<i>ICAES</i>	<i>ICRES</i>	<i>ICARES</i>	<i>ICES</i>	<i>ICAES</i>	<i>ICRES</i>	<i>ICARES</i>
Pass rates	0.093	0.494	0.820	0.948	0.035	0.455	0.795	0.964
Power	0.955	0.770	0.301	0.121	0.953	0.773	0.300	0.121
Selten	0.048	0.265	0.121	0.068	-0.012	0.228	0.095	0.085
Selten lb	0.005	0.190	0.061	0.033	-0.020	0.205	0.077	0.077
Selten ub	0.092	0.339	0.180	0.104	-0.004	0.251	0.113	0.094

Table 4.3: Pass rates, power, and predictive success for different specifications of *ICARES*

To summarize, 49% of the singles and 46% of the couples behave exactly like predicted by *ICAES*. The natural next question is what distinguishes these *ICAES* consumer types from the rest. We conducted a probit regression of consistency with *ICAES* on the basis of household type (single or couple, with or without children), age of the household head,

²²Note that, similar to Demuyneck and Verriest (2013) and Blow et al. (2021), we focus on ‘sharp’ rationalizability tests; however, we can easily accommodate for alternative measures such as the Afriat index.

and the household’s total expenditure over the period of observation. Age data is available in broad intervals: younger than 26, 26 to 35 years old, 36 to 45 years old, 46 to 55 years old, and older than 55. The regression uses a total of 2,052 households as data points. Neither age nor relationship status help to explain consistency with *ICAES*. We do find that the likelihood of *ICAES* consistency declines with the number of (young) children in the household. These households may have experienced a recent fertility event, and this may distort their planning activities in the short-to-medium run. We refer to Appendix C.2 for the complete regression outputs.

4.3 Identification of ‘anticipatory’ goods

The empirical support for *ICAES* is consistent with behavioral patterns observed in lab experiments (Loewenstein, 1987). However, *ICAES* opens the door for more comprehensive analyses of anticipation based on budget survey data. These analyses can yield insight in the (anticipatory) nature of a wider range of commodities. It can also deal with nonseparabilities in anticipation across goods. Common temporal profiles may indicate a degree of complementarity between anticipatory consumption goods. Practical considerations typically limit the range of commodities that can be used in lab experiments. Moreover, the majority of experiments focus on one commodity in isolation, in a tightly controlled decision-making environment. We now want to further illustrate how standard budget survey data with expenditure information, on a range of goods and services, can be fruitfully combined with our revealed preference analysis to shed more light on the relative importance of anticipated utilities across these different commodities.

As a first step, in Table 4.4, we replicate Table 4.1 specifically for the subset of consumers who satisfy *ICAES*. The mean expenditure shares are similar overall, but the share of food at home decreases from 46% to 44% while the share of food away increases from 13% to 14%. Anticipating types spend more on restaurants and less on food and drinking at home. However, the commodities used in this first comparison are aggregates of a wide range of different goods. For instance, transport includes long distance transportation but also standard public transport. Leisure is a combination of books, newspapers, and magazines but also cinema, theater, football, and other entertainment services. Moreover, the results in Table 4.4 shed little light on the *temporal* consumption profile of these goods.

As a second step, we regress expenditure associated with separate subgroups (included in the eight commodities) on total expenditure, interactions of total expenditure and the observation’s position in the sequence, and quarterly dummies to absorb seasonal and yearly effects. We estimate specification (23) again but this time we limit the sample to all con-

	Sample mean	Sample stdev
Allfood	43.90	18.59
Clothing	16.45	13.14
Hhserv	4.82	6.55
Transport	4.81	7.88
Petrol	7.47	7.93
Leisure	5.89	7.19
Pserv	2.66	4.56
Foodout	14.01	13.34

Table 4.4: Descriptive statistics of expenditure shares (in percent) on eight nondurable commodities, for consumers consistent with *ICAES*

sumers consistent with *ICAES*. In this way, we can identify which goods enter as arguments in the utility function of an anticipating self. To recall, the anticipating self values consumption more when it lies in the future. So, increasing parameter estimates $\theta_4^n > \theta_3^n > \theta_2^n > 0$ associated with the expenditure on good n suggest that n is an argument of the anticipatory utility function. Table 4.5 presents the results. The first two columns contain the names of subgroups (goods) and the commodities to which they belong. Columns three and four show the mean and standard deviation of expenditure shares of each subgroup. The next column gives the estimates of income effects divided by standard error. The final three columns present the estimates of θ_2^n , θ_3^n , θ_4^n , again deflated by the respective standard errors. We rank goods from high θ_4^n to low θ_4^n .

Most estimates of ‘sequence’ effects θ_t^n are nonnegative. This is not surprising: *ICAES* typically generates increasing consumption profiles. Yet, we find considerable heterogeneity in θ^n between goods. In the first set of goods (*restaurants* to *womensouter*), consumption is *much higher* in the last observation of the sequence. This is statistically significant at the 5% level. In the second set (*cookoil* to *fruit*), the effect is still there but only significant at the 10% level. We do not find significant temporal profiles for goods between *fooddrink_remain* and *poultry*. A small set of goods at the bottom of Table 4.5 is characterized by declining profiles.

Restaurant expenditure is the consumption category that increases the most over the sequence of observations. Interestingly, this appears to validate the use of restaurant visits in hypothetical choice experiments related to anticipation. [Loewenstein and Prelec \(1993\)](#) reported preferences for improvement when respondents could choose between sequences of restaurant visits. We must also note that the expenditure share of *restaurants* is high. Food at home (and clothing) subgroups have diverse temporal profiles so they appear in all parts of the table. Some foods and drinks are characterized by strongly increasing patterns: beer and wine, nuts, pastry and sugar, cheese, processed and delicacy meat and lamb. Other

Item	Commodity	Mean	Sd	γ	θ_2	θ_3	θ_4
restaurants	Foodout	14.005	13.345	14.152	3.261	3.931	5.728
childfoot	Clothing	1.462	2.601	7.652	1.549	2.934	4.438
mensfoot	Clothing	0.959	2.230	3.996	2.291	1.969	3.956
motfuel	Petrol	7.474	7.929	13.707	3.227	2.937	3.872
womensunder	Clothing	0.581	1.572	5.391	1.088	2.653	3.681
other_trans	Transport	3.272	6.415	7.966	1.666	3.158	3.447
mensouter	Clothing	3.277	6.161	6.191	1.060	0.941	3.404
beer	Allfood	0.540	1.421	5.303	2.248	2.266	3.312
childouter	Clothing	2.294	4.800	7.546	-0.549	1.638	3.229
recservs	Leisure	1.469	3.375	8.971	2.060	1.622	3.162
pcaservs	Pserv	1.100	2.836	6.301	2.617	1.983	3.155
cinema	Leisure	0.655	2.361	4.591	1.687	2.359	2.773
pcarendur	Pserv	1.561	3.474	5.725	1.726	2.352	2.684
nuts	Allfood	0.566	1.051	6.749	-0.117	2.018	2.682
accessories	Clothing	0.619	1.877	5.806	1.257	2.336	2.674
cleaning	Hhserv	2.164	2.811	10.351	0.028	1.462	2.479
longdistance	Transport	0.835	4.110	2.676	2.525	2.174	2.450
pastry	Allfood	2.035	2.469	8.817	1.937	1.929	2.372
processed_meat	Allfood	0.751	1.533	7.665	0.124	2.035	2.297
lamb	Allfood	1.074	2.717	4.758	0.537	2.235	2.216
deli_meat	Allfood	3.234	3.966	8.639	0.838	1.961	2.212
cheese	Allfood	2.194	2.400	12.520	0.336	0.920	2.202
wine	Allfood	0.566	1.619	3.761	0.188	1.494	2.147
sugar	Allfood	0.249	0.601	2.367	1.655	2.161	2.055
womensouter	Clothing	4.057	7.108	8.608	1.731	1.181	2.031
cookoil	Allfood	1.338	2.803	5.529	0.739	1.945	1.978
prime_meat	Allfood	2.229	3.442	8.905	0.704	1.426	1.901
other_alc	Allfood	0.190	1.141	4.104	0.214	1.197	1.887
mensunder	Clothing	0.458	1.572	4.848	-0.334	1.325	1.873
recgoods	Leisure	1.816	4.317	6.213	0.481	2.246	1.824
fruit	Allfood	2.818	2.626	12.587	1.520	1.602	1.751
fooddrink_remain	Allfood	0.663	4.304	3.512	1.246	0.937	1.566
fresh_fish	Allfood	2.688	3.391	10.262	-0.168	1.990	1.550
domservs	Hhserv	1.931	5.843	6.087	0.750	1.110	1.523
chocolate	Allfood	0.523	1.064	6.437	-0.692	1.464	1.466
processed_fish	Allfood	0.861	1.583	6.709	0.556	0.422	1.338
hhservs	Hhserv	0.210	1.341	3.009	1.159	1.206	1.293
processed_veg	Allfood	0.528	1.122	6.111	1.631	1.466	1.269
pubtrans	Transport	0.703	2.075	5.634	2.429	1.864	1.269
spirits	Allfood	0.298	1.202	3.702	-0.364	1.010	1.249
foot_remain	Clothing	0.016	0.316	-2.064	1.064	1.383	1.235
newsbook	Leisure	1.947	3.221	9.616	1.312	1.717	1.131
other_meat	Allfood	0.559	1.512	4.499	-1.087	0.238	1.044
nonalcbev	Allfood	0.865	1.382	8.670	0.234	2.044	1.039
potatoes	Allfood	0.721	1.243	5.999	0.821	2.230	1.020
eggs	Allfood	0.769	0.949	7.708	0.155	0.304	0.979
fresh_veg	Allfood	1.645	1.767	9.766	-0.766	0.564	0.928
molluscs	Allfood	1.049	2.415	4.257	-1.131	0.442	0.847
rice	Allfood	0.183	0.410	3.049	0.496	1.576	0.836
nondur_article	Hhserv	0.511	1.091	5.855	-0.681	0.193	0.829
cereals	Allfood	0.065	0.258	2.949	0.576	1.091	0.778
dried_veg	Allfood	0.352	0.848	2.757	-0.081	0.798	0.690
bread	Allfood	2.467	2.107	11.017	1.071	1.414	0.670
other_food	Allfood	0.963	1.781	6.417	-0.521	0.442	0.633
womensfoot	Clothing	1.118	2.420	6.624	0.575	1.206	0.422
pork	Allfood	1.449	2.616	4.888	-0.825	-0.009	0.355
milk	Allfood	2.734	2.628	5.944	-1.183	0.169	0.297
tobacco	Allfood	2.935	4.014	7.072	-1.555	-0.667	0.239
cloth_remain	Clothing	0.158	1.659	2.255	0.699	-0.113	0.218
beef	Allfood	0.447	1.830	1.952	0.904	0.073	0.198
pasta	Allfood	0.665	1.493	4.369	0.236	1.870	0.167
poultry	Allfood	1.657	2.304	7.021	0.706	0.233	0.010
childunder	Clothing	1.388	3.105	7.426	-0.694	0.039	-0.743
butter	Allfood	0.125	0.332	5.435	0.200	1.084	-0.801
coffee	Allfood	0.588	1.301	5.718	-1.221	-0.105	-0.939
footrepair	Clothing	0.058	0.262	3.652	-0.766	-2.836	-2.228
preserved_milk	Allfood	0.318	1.518	4.434	0.638	-0.983	-2.707

Table 4.5: Mean and spread of budget shares (in percent) and estimates of income effects γ and sequence effects θ (coefficient divided by standard error; b/se), per subgroup of goods. The sample is restricted to consumers consistent with *ICAES*.

foods, including pork, milk, tobacco, beef, pasta, poultry, butter, coffee, and preserved milk have constant or declining profiles. The first group of goods appears to be complementary to leisure activities and special celebrations, while the second group reflects more habit purchases. Intuitively, it makes sense that the utility from psychological consumption—such as savoring—is associated with the *less* frequently purchased commodities. [Hai et al. \(2020\)](#) analogously used frequent zero purchases, and lumpy expenditure spikes, to operationalize memorable consumption goods. In the leisure commodity, the increase in expenditure on recreation services and cinema (theater) is more outspoken than the increase in recreation goods and books purchases. In the transport commodity, long distance traveling (and other transportation) increases more than public transportation towards the end of the sequence. Overall, the ranking of goods in [Table 4.5](#) is not inconsistent with the notion of ‘anticipatory’ goods. This validates our interpretation of *ICAES* in terms of utility flows from savoring.

4.4 Robustness

In this section, we address the empirical issues of liquidity constraints, uncertainty, and measurement error, thus generalizing the tests of our framework.

The first exercise relaxes the assumption that all consumers can borrow money without limits. More technically, we admit that the marginal utility from wealth in period t *exceeds* the marginal utility from wealth (multiplied by $1 + r_t$) in period $t + 1$. This relaxes the revealed preference conditions as in [\(18\)](#) and [\(20\)](#) of [section 3.3](#). The second row of [Table 4.6](#) presents the new predictive success results. These results are very similar to the baseline findings without liquidity constraints. The predictive success of *ICAES* does not change much. At the same time, the empirical performance of *ICRES* and *ICARES* goes down. Overall, the predictive success of *ICAES* clearly exceeds that of other characterizations. We report the corresponding pass rates in [Appendix C.3](#). [Table C.2](#) in [Appendix C.3](#) shows that liquidity constraints have only a small effect on pass rates. Even with the extension of liquidity constraints, *ICES* explains less than 20% of the choices in the sample.

The second exercise relaxes the assumption of perfect foresight. Foresight predicts that the discounted marginal utilities from wealth are constant between time periods. To relax this structure, we admit that $\alpha \leq \lambda_t / (\lambda_{t+1}(1 + r_t)) \leq 1/\alpha$ with $\alpha \leq 1$. This relaxes the revealed preference conditions as in [\(18\)](#) and [\(22\)](#) of [section 3.3](#). Rows $\alpha_{1/5}$ to α_1 of [Table 4.6](#) present the new predictive success results for various lower bounds on α . Note that $\alpha_{1/5}$ is the highest level of α that still rationalizes 20% of the singles with *ICES*, while α_1 is the highest α that can rationalize the behavior of *all* singles. The results indicate that larger deviations (i.e., smaller α) systematically reduce predictive success for all specifications

under consideration. At $\alpha_{1/3}$, the predictive success of *ICAES* is still 15% while that of other characterizations is less than 5%. At $\alpha_{1/2}$, *ICAES* is the only specification with a positive predictive success for both singles and couples. Finally, α_1 imposes so little structure that discriminatory power—and hence predictive success—becomes zero for all specifications.

The third exercise addresses measurement error in the price data. Price data in the application comes from the Spanish Statistical Office. It is not measured at the level of each consumer. The general price indices in the analysis are proxies for the consumers' true prices (which in reality may vary between consumers). We thus repeat our analyses with unobserved consumer-specific prices $\tilde{\rho}$ rather than the observed prices ρ . We capture the distance between both prices by the sum of absolute errors $\varepsilon = \sum_t \sum_n |\tilde{\rho}_{n,t} - \rho_{n,t}|$. Discounted prices change by 0.005 to 0.011 index points from one quarter to the next; thus errors can be expected to have a large impact. Rows $\varepsilon_{1/5}$ to ε_1 of Table 4.6 present the new predictive success results for various upper bounds on ε (i.e., $\varepsilon \leq \varepsilon_{1/5}, \varepsilon \leq \varepsilon_{1/4}, \dots$). Like before, $\varepsilon_{1/5}$ is the smallest distance that allows us to rationalize 20% of the singles with *ICES*, while ε_1 is the smallest distance that allows us to rationalize the behavior of *all* singles. The addition of limited levels of error improves the predictive success of both *ICES* and *ICAES*. The latter still receives the strongest empirical support but the gap between *ICAES* and *ICES* becomes smaller especially for singles. This confirms that the classical lifecycle model *with measurement error* performs reasonably well for singles. However, the predictive success of *ICAES* still dominates even if we add substantial measurement error to the data.

Selten	Singles				Couples			
	<i>ICES</i>	<i>ICAES</i>	<i>ICRES</i>	<i>ICARES</i>	<i>ICES</i>	<i>ICAES</i>	<i>ICRES</i>	<i>ICARES</i>
Baseline	0.048	0.265	0.121	0.068	-0.012	0.228	0.095	0.085
Liq Constr	0.092	0.276	-0.008	-0.004	0.003	0.234	-0.009	-0.004
Uncertainty								
$\alpha_{1/5} = 0.98$	0.048	0.228	0.020	0.003	-0.019	0.260	0.008	0.014
$\alpha_{1/4} = 0.974$	0.006	0.195	0.033	-0.004	-0.045	0.208	0.001	0.008
$\alpha_{1/3} = 0.97$	0.006	0.149	0.023	-0.007	-0.061	0.151	-0.004	0.004
$\alpha_{1/2} = 0.9535$	-0.054	0.033	0.003	-0.010	-0.129	0.038	-0.014	-0.003
$\alpha_1 = 0.01$	-0.008	0.000	0.000	0.000	-0.005	-0.001	-0.001	-0.001
Error								
$\varepsilon_{1/5} = 0.0255$	0.109	0.284	0.099	0.034	0.039	0.296	0.075	0.047
$\varepsilon_{1/4} = 0.0353$	0.139	0.289	0.078	0.027	0.065	0.311	0.063	0.037
$\varepsilon_{1/3} = 0.0520$	0.177	0.329	0.068	0.016	0.107	0.328	0.046	0.025
$\varepsilon_{1/2} = 0.0845$	0.246	0.315	0.040	0.011	0.188	0.334	0.026	0.013
$\varepsilon_1 = 0.5065$	0.015	0.011	0.000	0.000	0.011	0.009	0.000	0.000
First quarter	0.071	0.271	0.138	0.048	-0.004	0.26	0.08	0.091

Table 4.6: Predictive success for different specifications of *ICARES*, with liquidity constraints, uncertainty, measurement error, and for subsamples that start in the first quarter

In a final exercise, we further separate the effect of a consumption event's 'position in the

observed sequence’ from the ‘calendar date of the observation’. Our previous analyses already control for seasonal variation through quarterly dummies. Here we go one step further to rule out any correlation between seasonal effects and consumer characteristics. We repeat the analyses for the subgroup of consumers for whom the first (final) observation coincides with the first (final) quarter of the year. If yearly consumption plans were formed at the beginning of each new year, for instance, then the integrity of the sequence of observations should be strong in this subgroup. The restricted subsample consists of 609 consumers: 51 singles and 558 couples. The results are in the final row of Table 4.6. The empirical performance of the models is very similar to the baseline sample overall, but predictive success of *ICAES* among couples increases further from 22.8% to 26%.

5 Related literature

We now position the present paper in the larger literature. First, there have been other studies describing the complex psychological features pertaining to consumption events. [Morewedge \(2015\)](#), for instance, asked a sample of Americans to describe the contribution of anticipation, remembering, and experience to the total pleasure derived from various activities. The relative contribution of anticipation varied from 15% (exercise) to 24% (vacation); the relative contribution of memory from 13% (dinner) to 30% (wedding). The contributions in [Loewenstein \(1987\)](#) and [Loewenstein and Sicherman \(1991\)](#) have focused on the feature of anticipated utility, whereas [Gilboa et al. \(2016\)](#) and [Hai et al. \(2020\)](#) presented models of remembering and memorable goods. [Gilboa et al. \(2016\)](#) emphasized the complexity, due to reverse time inconsistency, of allowing for anticipation without additional assumptions on anticipatory preferences. In that regard, our paper provides such additional structure, which makes the use of a decision theoretic framework possible.

To the best of our knowledge, [Baucells and Bellezza \(2017\)](#) is the only other paper that models the temporal profile of ‘instant’ utilities from anticipation, remembering, and experience. Beside physical and psychological consumption, the authors also consider reference points. A distinguishing feature of their framework is that the carrier of utility is *effective* consumption: the difference between consumption and some reference point. The reference point is endogenous: savoring increases the target against which future consumption is valued (adaptation). The model of [Baucells and Bellezza \(2017\)](#) incorporates a wide range of insights from psychology. The main difference with respect to *ICARES* is that [Baucells and Bellezza \(2017\)](#) study the utility flows from a *single* consumption event. *ICARES*, by contrast, considers an environment with more than one good, and with consumption observations at multiple points in time. Given the wider range of commodities, *ICARES* also

allows the utility functions before, during, and after events to differ in arguments and in shape. Finally, we set our framework in discrete time, to tailor it to revealed preference testing and identification. Indeed, as we have shown, *ICARES* and its restricted versions produce straightforward testable implications outside specially tailored lab settings.

Second, our paper uses tools from the revealed preference literature to analyze the empirical content of *ICARES* and its special cases. Revealed preference theory was introduced early by Samuelson (1938) and Houthakker (1950). The seminal contributions by Afriat (1967) and Varian (1982) made revealed preference analysis operational and applicable in survey data on expenditures. By now, the technique has been used in a wide range of applications: household consumption choices, choices from non-linear budget sets, analyzing stable matching patterns, etc. We refer to Crawford and De Rock (2014) for more applications of (empirical) revealed preference methods. The revealed preference approach has several advantages. First, it is intrinsically nonparametric, and thus completely independent of the specific functional form of utility. Next, it allows econometricians to analyze each consumer separately, thus incorporating a large degree of individual heterogeneity. Revealed preference methods have already been applied fruitfully to the analysis of intertemporal models; notable examples are the study of rational habit formation (Crawford, 2010), rational addiction (Demuyneck and Verriest, 2013), intertemporal collective choice (Adams et al., 2014), discounted utility models (Dziwulski, 2018), the exponentially discounted utility model (Echenique et al., 2020), and models of quasi-hyperbolic discounting (Blow et al., 2021).

Finally, our paper adds to the large literature on behavioral deviations from (exponential) discounting; the *DU* framework. The *DU* model has excellent empirical tractability but imposes strong assumptions. We focus on two of these assumptions and refer to Frederick et al. (2002) for a comprehensive overview. *DU* typically assumes *positive devaluing*—with constant discount factors $\beta \leq 1$ —and *independence of discounting from consumption*. Utility from anticipation violates the first assumption; the dependency of anticipatory emotions and pleasant memories on the type of consumption good violates the second. Capturing these behavioral phenomena within the context of *DU* requires a flexible definition of discount factors β . A first amendment is to let discount factors vary over time: $\beta(t)$. This also permits negative devaluing (i.e., $\beta(t)$ increasing in t). A second amendment is to let discount factors vary between goods: $\beta(n)$. This addresses differences in temporal profiles between anticipatory, memorable, and ordinary consumption goods. However, even with these extensions, *DU* still has important limitations. First, *DU* with $\beta(t, n)$ suffers from a curse of dimensionality. The number of parameters grows multiplicatively with the number of observations *and* the number of goods. Second, the literature shows a large dispersion

of discount factors.²³ Estimates change dramatically from one experiment to the next, and this does not shed much light on the behavioral mechanisms underlying intertemporal consumption. Finally, [Manzini et al. \(2010\)](#) studied choices between time sequences of monetary outcomes. The authors found that standard models based on discounting could not explain the data, no matter how much variability in discount factors was allowed.

6 Conclusion

This paper is motivated by two observations of intertemporal behavior that violate the predictions of the lifecycle model. First, consumers sometimes postpone desirable outcomes in order to extend the duration of ‘savoring’ ([Baucells and Bellezza, 2017](#)). Second, consumers sometimes spend disproportionate amounts of income on holidays, celebrations, and ceremonies early in life to maximize the duration of ‘memories’ generated by this consumption ([Gilboa et al., 2016](#)). Both observations have one feature in common: a dissociation between physical consumption and the utility flows from this consumption.

We propose and test a new model in which consumers enjoy utility from savoring of future events, experience of a current event, and remembering of past events. We thus represent consumption as the outcome of a bargaining process between three temporal selves: an anticipating self, a remembering self, and an experiencing self. The selves can have different valuations of the same commodity.

The choices produced by our model are generally not time consistent. First, the bargaining between selves may suffer from commitment issues. This complicates the aggregation of different temporal motives. Second, and more fundamentally, the duration of savoring decreases naturally as time moves forward. This induces acts of *reverse* time inconsistency and undermines the very notion of anticipation through loss of self-credibility. We put forward an internal mechanism that can mitigate acts of reverse time inconsistency: the decreasing duration of savoring is offset by an increasing decision weight of the anticipating self. This imposes the qualitative condition that the decision weight of the anticipating self increases towards the end of each planning period. Such condition is in line with the view that shorter planning periods ‘activate’ preferences for improvement.

To bring the theory to the data, we specify time factors that satisfy the conditions of the internal commitment mechanism. We leave the utility functions and the initial decision weights unspecified. The corresponding model, *ICARES*, is a time consistent version of intertemporal consumption with anticipating, remembering, and experiencing selves. *ICARES* nests a number of interesting polar cases: consumption with anticipating and experiencing

²³[Frederick et al. \(2002\)](#) list a range of estimates from experimental and survey data in Table 1, Figure 2.

selves (*ICAES*), with remembering and experiencing selves (*ICRES*), and with an experiencing self alone (*ICES*). *ICARES* and its special cases have straightforward testable implications even outside specially tailored lab experiments.

We then derive a revealed preference characterization. We apply this characterization to a panel dataset of quarterly consumption by Spanish households (ECPF). *ICARES* rationalizes almost all observations, but lacks discriminatory power. The most successful specification is *ICAES*: it rationalizes close to two thirds of the data, and is still fairly powerful. We then investigate heterogeneity in the ‘anticipatory’ nature of consumption goods. For *ICAES* consumers, we find that restaurant expenditure, leisure services, and food expenditure complementary to these leisure activities increase more sharply over the planning period compared to other expenditure items. In line with experimental findings of [Loewenstein \(1987\)](#), our evidence from budget survey data confirms that anticipation matters for understanding consumption patterns. More generally, this is one of the first papers to provide a successful rationalization of consumption patterns for the full dataset (that is, both singles and couples) with a model that also satisfies time consistency.

A large literature has studied deviations from the discounted utility framework, but most of this work focused on violations of time consistency (present bias or myopia). Anticipatory emotions and memorable consumption have received less attention. Especially remarkable is the lack of evidence from budget survey data. The intertemporal framework proposed in this paper is situated between the theory of total utility ([Kahneman et al., 1997](#)) and the discounted utility model ([Samuelson, 1937](#)). The former enhances psychological realism by incorporating *all* the utility flows from savoring and memories; the latter maintains empirical tractability for consumption choices from standard intertemporal budget constraints.

References

- Adams, A., L. Cherchye, B. De Rock, and E. Verriest (2014). Consume now or later? Time inconsistency, collective choice, and revealed preference. *American Economic Review* 104, 4147–4183.
- Afriat, S. N. (1967). The construction of utility functions from expenditure data. *International Economic Review* 8, 67–77.
- Aguiar, V. H. and N. Kashaev (2021). Stochastic revealed preferences with measurement error. *The Review of Economic Studies* 88, 2042–2093.
- Ambrus, A. and K. Rozen (2014). Rationalising choice with multi-self models. *The Economic Journal* 125, 1136–1156.

- Andreoni, J. and C. Sprenger (2012). Risk preferences are not time preferences. *American Economic Review* 102, 3357–3376.
- Baucells, M. and S. Bellezza (2017). Temporal profiles of instant utility during anticipation, event, and recall. *Management Science* 63, 729–748.
- Beatty, T. and I. Crawford (2011). How demanding is the revealed preference approach to demand. *American Economic Review* 101, 2782–2795.
- Benhabib, J. and A. Bisin (2005). Modeling internal commitment mechanisms and self-control: A neuroeconomics approach to consumption-saving decisions. *Games and Economic Behavior* 52, 460–492.
- Blow, L., M. Browning, and I. Crawford (2021). Nonparametric analysis of time-inconsistent preferences. *The Review of Economic Studies* 88, 2687–2734.
- Blundell, R. and J.-M. Robin (2000). Latent separability: Grouping goods without weak separability. *Econometrica* 68, 53–84.
- Bronars, S. (1987). The power of nonparametric tests of preference maximization. *Econometrica* 55, 693–698.
- Browning, M. (1989). A nonparametric test of the life-cycle rational expectations hypothesis. *International Economic Review* 30, 979–992.
- Browning, M. and M. Collado (2007). Habits and heterogeneity in demands: a panel data analysis. *Journal of Applied Econometrics* 22, 625–640.
- Caplin, A. and J. Leahy (2001). Psychological expected utility theory and anticipatory feelings. *The Quarterly Journal of Economics* 116, 55–79.
- Castillo, M. I., S. Sun, M. A. Frank-Crawford, and J. C. Borrero (2022). Save the best for last i: Young adults demonstrate negative time preference—a replication and extension. *Behavior Analysis: Research and Practice* 22, 143–163.
- Chapman, G. B. (1996). Expectations and preferences for sequences of health and money. *Organizational Behavior and Human Decision Processes* 67, 59–75.
- Cherchye, L., B. De Rock, R. Griffith, M. O’Connell, K. Smith, and F. Vermeulen (2020). A new year, a new you? Within-individual variation in food purchases. *European Economic Review* 127, 1034–1078.

- Chiappori, P. A. (1988). Rational household labor supply. *Econometrica* 56, 63–89.
- Chiappori, P. A. (1992). Collective labor supply and welfare. *Journal of Political Economy* 100, 437–467.
- Crawford, I. (2010). Habits revealed. *The Review of Economic Studies* 77, 1382–1402.
- Crawford, I. and B. De Rock (2014). Empirical revealed preference. *Annual Review of Economics* 6, 503–524.
- Crawford, I. and M. Polisson (2014). Testing for intertemporal nonseparability. *Journal of Mathematical Economics* 52, 46–49.
- Dean, M. and A. Sautmann (2021). Credit constraints and the measurement of time preferences. *The Review of Economics and Statistics* 103, 119–135.
- Demuyck, T. (2015). Statistical inference for measures of predictive success. *Theory and Decision* 79, 689–699.
- Demuyck, T. and E. Verriest (2013). I’ll never forget my first cigarette: A revealed preference analysis of the habits as durables model. *International Economic Review* 54, 717–738.
- Drouhin, N. (2020). Non-stationary additive utility and time consistency. *Journal of Mathematical Economics* 86, 1–14.
- Dziewulski, P. (2018). Revealed time preference. *Games and Economic Behavior* 112, 67–77.
- Echenique, F., T. Imai, and K. Saito (2020). Testable implications of models of intertemporal choice: Exponential discounting and its generalizations. *American Economic Journal: Microeconomics* 12, 114–143.
- Elster, J. and G. Loewenstein (1992). *Choice over time*, Chapter Utility from memory and anticipation, pp. 213–234. Russell Sage Foundation.
- Frederick, S., G. Loewenstein, and T. O’Donoghue (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature* 40, 351–401.
- Gilboa, I., A. Postlewaite, and L. Samuelson (2016). Memorable consumption. *Journal of Economic Theory* 165, 414–455.
- Guo, L. (2020). Anticipatory consumptions. *Management Science* 66, 3717–3734.

- Hai, R., D. Krueger, and A. Postlewaite (2020). On the welfare cost of consumption fluctuations in the presence of memorable goods. *Quantitative Economics* 11, 1177–1214.
- Houthakker, H. S. (1950). Revealed preference and the utility function. *Economica* 17, 159–174.
- Jackson, M. O. and L. Yariv (2014). Present bias and collective dynamic choice in the lab. *American Economic Review* 104, 4184–4204.
- Kahneman, D., P. P. Wakker, and R. Sarin (1997). Back to Bentham? Explorations of experienced utility. *The Quarterly Journal of Economics* 112, 375–406.
- Koopmans, T. (1960). Stationary ordinal utility and impatience. *Econometrica* 28, 287–309.
- Kőszegi, B. (2009). Utility from anticipation and personal equilibrium. *Economic Theory* 44, 415–444.
- Laajaj, R. (2017). Endogenous time horizon and behavioral poverty trap: Theory and evidence from mozambique. *Journal of Development Economics* 127, 187–208.
- Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *The Economic Journal* 97, 666–684.
- Loewenstein, G. and D. Prelec (1993). Preferences for sequences of outcomes. *Psychological Review* 100, 91–108.
- Loewenstein, G. and N. Sicherman (1991). Do workers prefer increasing wage profiles? *Journal of Labor Economics* 9, 67–84.
- Manzini, P., M. Mariotti, and L. Mittone (2010). Choosing monetary sequences: theory and experimental evidence. *Theory and Decision* 69, 327–354.
- May, K. O. (1954). Intransitivity, utility, and the aggregation of preference patterns. *Econometrica* 22, 1.
- Mazzocco, M. (2007). Household intertemporal behaviour: A collective characterization and a test of commitment. *The Review of Economic Studies* 74, 857–895.
- Morewedge, C. K. (2015). Utility. In *The Wiley Blackwell Handbook of Judgment and Decision Making*, pp. 295–330. John Wiley & Sons, Ltd.
- O’Donoghue, T. and M. Rabin (1999). Doing it now or later. *American Economic Review* 89, 103–124.

- Prelec, D. and G. Loewenstein (1991). Decision making over time and under uncertainty: A common approach. *Management Science* 37, 770–786.
- Rockafellar, R. T. (1970). *Convex Analysis*. Princeton University Press.
- Samuelson, P. (1937). A note on measurement of utility. *The Review of Economic Studies* 4, 155–161.
- Samuelson, P. (1938). A note on the pure theory of consumer’s behavior. *Economica* 5, 61–71.
- Selten, R. (1991). Properties of a measure of predictive success. *Mathematical Social Sciences* 21, 153–167.
- Varian, H. R. (1982). The nonparametric approach to demand analysis. *Econometrica* 50, 945–974.
- Varian, H. R. (1983). Non-parametric tests of consumer behaviour. *The Review of Economic Studies* 50, 99.
- Varian, H. R. (1985). Nonparametric analysis of optimizing behavior with measurement error. *Journal of Econometrics* 30, 445–458.

A Proofs

A.1 Proof of Proposition 1

We prove Proposition 1 by offering a counterexample to dynamic consistency. Comparing the optimal solutions to (25) and (26) in the numerical example below shows that in general (i.e., without additional structure on the dynamics of intra-selves bargaining and/or the time functions D^A and D^R) our framework does not satisfy time consistency.

A numerical example. We start from a parametric specification of utility functions u^A , u^R , u^E , decision weight functions ω^A , ω^R , and an intertemporal budget constraint. We simulate a consumption plan $(\mathbf{q}_t^*)_{t \geq 1}$ at the start of $\tau = 1$. We then show that the optimal plan $(\hat{\mathbf{q}}_t)_{t \geq 2}$ changes at the start of $\tau = 2$. That is, $\hat{\mathbf{q}}_t \neq \mathbf{q}_t^*$ for some $t \geq 2$.

We restrict this exercise to $N = 2$ goods, so $\mathbf{q}_t = (q_{1,t}, q_{2,t})$. We choose the following parametric specification for the utility functions,

$$u^i(q_{1,t}, q_{2,t}) = \alpha^i \log(q_{1,t}) + (1 - \alpha^i) \log(q_{2,t}), \text{ where } i = A, R, E. \quad (24)$$

We let $\alpha^A = \alpha^R = 1$: the utility from anticipation and recall comes exclusively from good 1. For simplicity we assume that $\omega^R(\tau) = \omega_0^R$ and $\omega^A(\tau) = \omega_0^A$, where $\omega_0^A > \omega_0^R$. In words, the anticipating self has relatively more influence over the decision making process than the remembering self. With regards to the time functions D^A and D^R , we assume a simple exponential form: $D^A(t - \tau) = (\beta^A)^{t-\tau+1}$ and $D^R(T - t) = (\beta^R)^{T-t+1}$, where $\beta^A = \beta^R > 1$. We normalize all prices to unity. The optimization problem at $\tau = 1$ reads as follows:

$$\begin{aligned} \max_{(\mathbf{q}^t)_{t \geq 1}} \sum_{t=1}^T & \left[\left(\omega_0^A (\beta^A)^t + \omega_0^R (\beta^R)^{T-t+1} + \alpha^E \right) \log(q_{1,t}) + (1 - \alpha^E) \log(q_{2,t}) \right], \text{ subject to} \\ \sum_{t=1}^T & (q_{1,t} + q_{2,t}) = y. \end{aligned} \quad (25)$$

Solving the associated system of first order conditions then yields the following solutions for consumption levels,

$$\begin{aligned} q_{1,t}^* &= \frac{C(t) y}{\sum_{t=1}^T [1 + C(t)]}, \text{ and} \\ q_{2,t}^* &= \frac{y}{\sum_{t=1}^T [1 + C(t)]}, \text{ for all } t \geq 1, \end{aligned}$$

where $C(t) = \left(\omega_0^A (\beta^A)^t + \omega_0^R (\beta^R)^{T-t+1} + \alpha^E \right) (1 - \alpha^E)^{-1}$. Notice that the profile for consumption of the first good is increasing, given $\omega_0^A > \omega_0^R$ and $\beta^A = \beta^R > 1$. This is in line with the assumption in this example that the anticipating self is more influential in the decision process than the remembering self. Also noteworthy is the consumption smoothing with respect to good 2. This good produces no utility from anticipation or remembering.

Next, we study the consumption choices made by this consumer, but starting from the decision period $\tau = 2$. The equivalent of (25) can then be formulated as follows:

$$\begin{aligned} \max_{(\mathbf{q}^t)_{t \geq 2}} \sum_{t=2}^T & \left[\left(\omega_0^A (\beta^A)^{t-1} + \omega_0^R (\beta^R)^{T-t+1} + \alpha^E \right) \log(q_{1,t}) + (1 - \alpha^E) \log(q_{2,t}) \right], \text{ subject to} \\ \sum_{t=2}^T & (q_{1,t} + q_{2,t}) = \hat{y}, \end{aligned} \quad (26)$$

where $\hat{y} = y - q_{1,1} - q_{2,1}$ are the available resources for expenditures over the horizon $t \in \{2, \dots, T\}$. By again solving the associated system of first order conditions, we obtain the

optimal consumption choices:

$$\hat{q}_{1,t} = \frac{\hat{C}(t) \hat{y}}{\sum_{t=2}^T [1 + \hat{C}(t)]}, \text{ and}$$

$$\hat{q}_{2,t} = \frac{\hat{y}}{\sum_{t=2}^T [1 + \hat{C}(t)]}, \text{ for all } t \geq 2,$$

with $\hat{C}(t) = \left(\omega_0^A (\beta^A)^{t-1} + \omega_0^R (\beta^R)^{T-t+1} + \alpha^E \right) (1 - \alpha^E)^{-1}$. It can now be shown that, given $\omega_0^A \neq 0$, the optimal consumption choices for $t \geq 2$ are such that $\hat{\mathbf{q}}_t \neq \mathbf{q}_t^*$. This simple parametric example shows how reverse time inconsistency, driven by the anticipating self, can produce dynamically inconsistent behavior on the part of consumers.

A.2 Proof of Proposition 2

Consider any $\tau, \tau' \in \mathcal{T}$, with $\tau < \tau'$. Let $(\mathbf{q}_t^*)_{t \geq \tau}$ denote the solution to (6). Now suppose, towards a contradiction, that $(\mathbf{q}_t^*)_{t \geq \tau'}$ is *not* the solution to (8). Then there must exist a sequence $(\hat{\mathbf{q}}_t)_{t \geq \tau'} \in \mathcal{B}((\boldsymbol{\rho}_t)_{t \geq \tau'}, y_{\tau'})$ such that

$$\begin{aligned} & \sum_{t \geq \tau'} [\omega^A(\tau') D^A(t - \tau') u^A(\hat{\mathbf{q}}_t) + \omega^R(\tau') D^R(T - t) u^R(\hat{\mathbf{q}}_t) + u^E(\hat{\mathbf{q}}_t)] \\ & > \sum_{t \geq \tau'} [\omega^A(\tau') D^A(t - \tau') u^A(\mathbf{q}_t^*) + \omega^R(\tau') D^R(T - t) u^R(\mathbf{q}_t^*) + u^E(\mathbf{q}_t^*)]. \end{aligned} \quad (27)$$

1. Assume (9) holds. We thus have that, for all $\tau \in \mathcal{T}$ and $t \geq \tau$:

$$\omega^A(\tau) D^A(t - \tau) = \exp(\alpha(t)) \equiv a(t). \quad (28)$$

Taking derivatives on both sides yields the condition:

$$\frac{\partial (\omega^A(\tau) D^A(t - \tau))}{\partial \tau} = 0. \quad (29)$$

This implies independence of function $\omega^A(\tau) D^A(t - \tau) = a(t)$ from the decision moment. Similarly, $\frac{\partial \omega^R(\tau)}{\partial \tau} = 0$ immediately imposes independence of $b(t) = \omega^R(\tau) D^R(T - t)$

from the decision moment. Using this information, we can rewrite (27) as follows:

$$\begin{aligned} & \sum_{t \geq \tau'} [a(t) u^A(\hat{\mathbf{q}}_t) + b(t) u^R(\hat{\mathbf{q}}_t) + u^E(\hat{\mathbf{q}}_t)] + \sum_{t=\tau}^{\tau'-1} [a(t) u^A(\mathbf{q}_t^*) + b(t) u^R(\mathbf{q}_t^*) + u^E(\mathbf{q}_t^*)] \\ & > \sum_{t \geq \tau} [a(t) u^A(\mathbf{q}_t^*) + b(t) u^R(\mathbf{q}_t^*) + u^E(\mathbf{q}_t^*)]. \end{aligned} \quad (30)$$

2. We now show that $(\mathbf{q}_\tau^*, \mathbf{q}_{\tau+1}^*, \dots, \hat{\mathbf{q}}_{\tau'}, \dots, \hat{\mathbf{q}}_T)$ was also feasible at $t = \tau$. To that end, note that $(\hat{\mathbf{q}}_t)_{t \geq \tau'} \in \mathcal{B}(\boldsymbol{\rho}_t, y_{\tau'})$ implies:

$$\begin{aligned} \sum_{t \geq \tau'} \frac{\mathbf{p}_t}{\prod_{i=1}^{t-1} (1+r_i)} \cdot \hat{\mathbf{q}}_t & \leq y_{\tau'} \\ & = \left(y - \sum_{t \leq \tau'-1} \frac{\mathbf{p}_t}{\prod_{i=1}^{t-1} (1+r_i)} \cdot \mathbf{q}_t^* \right). \end{aligned} \quad (31)$$

Rearranging terms in (31), we obtain:

$$\begin{aligned} \sum_{t \geq \tau'} \frac{\mathbf{p}_t}{\prod_{i=1}^{t-1} (1+r_i)} \cdot \hat{\mathbf{q}}_t + \sum_{t \geq \tau}^{\tau'-1} \frac{\mathbf{p}_t}{\prod_{i=1}^{t-1} (1+r_i)} \cdot \mathbf{q}_t^* & \leq \left(y - \sum_{t \leq \tau-1} \frac{\mathbf{p}_t}{\prod_{i=1}^{t-1} (1+r_i)} \cdot \mathbf{q}_t^* \right) \\ \sum_{t \geq \tau}^{\tau'-1} \frac{\mathbf{p}_t}{\prod_{i=1}^{t-1} (1+r_i)} \cdot \mathbf{q}_t^* + \sum_{t \geq \tau'} \frac{\mathbf{p}_t}{\prod_{i=1}^{t-1} (1+r_i)} \cdot \hat{\mathbf{q}}_t & \leq y_\tau. \end{aligned} \quad (32)$$

3. We now put results (1) and (2) together. Condition (32) shows that $(\mathbf{q}_\tau^*, \mathbf{q}_{\tau+1}^*, \dots, \hat{\mathbf{q}}_{\tau'}, \dots, \hat{\mathbf{q}}_T) \in \mathcal{B}(\boldsymbol{\rho}_t, y_\tau)$ was feasible at decision period $t = \tau$. Condition (30) implies that $(\mathbf{q}_\tau^*, \mathbf{q}_{\tau+1}^*, \dots, \hat{\mathbf{q}}_{\tau'}, \dots, \hat{\mathbf{q}}_T)$ would have also yielded higher overall utility than $(\mathbf{q}_t^*)_{t \geq \tau}$. This contradicts our opening statement that $(\mathbf{q}_t^*)_{t \geq \tau}$ solves (6). We conclude that $(\mathbf{q}_t^*)_{t \geq \tau'}$ must solve (8), thereby confirming dynamic consistency of the consumer under condition (9).

A.3 Proof of Proposition 3

The proof consists of a necessary and a sufficient part.

Necessity: We first prove that *ICARES* implies the system of conditions in Proposition

3. From concavity of u^A , u^R , and u^E , we know that for all $s, t \in \mathcal{T}$:

$$\begin{aligned} u^A(\mathbf{q}_s) - u^A(\mathbf{q}_t) & \leq \partial u^A(\mathbf{q}_t) \cdot (\mathbf{q}_s - \mathbf{q}_t); \\ u^R(\mathbf{q}_s) - u^R(\mathbf{q}_t) & \leq \partial u^R(\mathbf{q}_t) \cdot (\mathbf{q}_s - \mathbf{q}_t); \\ u^E(\mathbf{q}_s) - u^E(\mathbf{q}_t) & \leq \partial u^E(\mathbf{q}_t) \cdot (\mathbf{q}_s - \mathbf{q}_t). \end{aligned}$$

Define utilities $u_t^A = \omega_0^A/\lambda \times u^A(\mathbf{q}_t)$ and marginal utilities $\tilde{\mathbf{p}}_t^A = \omega_0^A/\lambda \times \partial u^A(\mathbf{q}_t)$ in line with section 3, and similarly for u_t^R and u_t^E , and $\tilde{\mathbf{p}}_t^R$ and $\tilde{\mathbf{p}}_t^E$. This produces conditions (12)–(14). Condition (15) is a direct translation of the first order conditions.

Sufficiency: We subsequently prove that the system of conditions in Proposition 3 implies existence of utility functions u^A , u^R , and u^E , weights ω_0^A and ω_0^R , and parameters β^A and β^R so that (11) holds. Consider a subset of observations $\tilde{\mathcal{T}} \subseteq \mathcal{T}$ and sum conditions (14) over this subset. We obtain:

$$0 \leq \sum_{s,t \in \tilde{\mathcal{T}}} \tilde{\mathbf{p}}_t^E \cdot (\mathbf{q}_s - \mathbf{q}_t). \quad (33)$$

Condition (33) is referred to as *cyclical monotonicity* (Rockafellar, 1970) and implies existence of a concave utility function u^E so that

$$\partial u^E(\mathbf{q}_t) = \tilde{\mathbf{p}}_t^E.$$

We can repeat this argument and sum conditions (13) over a subset of observations in $\tilde{\mathcal{T}}$. We obtain:

$$0 \leq \sum_{s,t \in \tilde{\mathcal{T}}} \tilde{\mathbf{p}}_t^R \cdot (\mathbf{q}_s - \mathbf{q}_t) \quad (34)$$

and this implies existence of a concave map u^R such that:

$$\partial u^R(\mathbf{q}_t) = \tilde{\mathbf{p}}_t^R.$$

Likewise, summing conditions (12) over a subset of observations $\tilde{\mathcal{T}}$, we obtain:

$$0 \leq \sum_{s,t \in \tilde{\mathcal{T}}} \tilde{\mathbf{p}}_t^A \cdot (\mathbf{q}_s^A - \mathbf{q}_t^A). \quad (35)$$

Thus there exists a concave function u^A so that:

$$\partial u^A(\mathbf{q}_t) = \tilde{\mathbf{p}}_t^A.$$

Finally, we take $\omega_0^A = \omega_0^R = \lambda = 1$ without losing generality. Then (15), with $\partial u^i(\mathbf{q}_t) = \tilde{\mathbf{p}}_t^i$, yields first order conditions (11) for consistency with *ICARES*.

A.4 Proof of Corollary 1

It is sufficient to provide 2 datasets, one which satisfies *ICARES* but violates *GARP* and conversely. To that end, consider Table 3.2 of our simulation exercise. The dataset generated

by *ICARES* in the simulation exercise violates *GARP*. On the other hand, the dataset consistent with *GARP* violates the RP characterization of *ICARES*. This shows that the empirical content of *ICARES* is independent from the content of static utility maximization.

A.5 Proof of Corollary 2

We again prove this result by providing 2 datasets, one which satisfies the RP restrictions of *ICAES* but violates the restrictions of *ICRES*, and another which satisfies the restrictions of *ICRES* but violates the restrictions of *ICAES*. We can resort back to our simulation exercise, in particular the consumption time series provided in Table 3.1. The consumption time series generated by *ICAES* violates the RP conditions of the *ICRES* model, whereas the consumption data generated from *ICRES* violates the RP characterization of *ICAES*.

B Extension to infinite horizon

Throughout the main paper we have restricted attention to the case where the horizon \mathcal{T} is finite. However, we can easily generalize our theoretical framework to a case where $\mathcal{T} = \mathbb{N}_0$. To that end, we will readjust the consumer's overall utility, as seen from decision period $\tau \geq 1$, as follows:

$$\sum_{t \geq \tau} \delta^{t-\tau} [\omega^A(\tau) D^A(t-\tau) u^A(\mathbf{q}_t) + \omega^R(\tau) D^R(t) u^R(\mathbf{q}_t) + u^E(\mathbf{q}_t)]. \quad (36)$$

A few remarks are in order at this point. First note that in this case, the time function D^R is simply a function of t and not of the length of the planning period (as T is infinite). In accordance with the finite horizon setting, we merely assume that D^R is decreasing in t . Next, the main difference between (36) and the objective function in the main paper is the presence of the exponential discount factor δ , contained in the (open) unit interval $(0, 1)$. Although this factor is assumed to be uniform, time functions D^A and D^R remain heterogeneous across selves. We can then show that, under the same sufficient condition as in Proposition 2, the dynamic consistency result remains robust:

Proposition 4 *Let $(\mathbf{q}_t^*)_{t \geq \tau}$ be the solution to maximizing (36) subject to $\sum_{t=\tau}^{+\infty} \rho_t \cdot \mathbf{q}_t \leq y_\tau$ and let $(\hat{\mathbf{q}}_t)_{t \geq \tau'}$ be the solution to the same problem with the decision period shifted to $\tau' \neq \tau$. Then $\hat{\mathbf{q}}_t = \mathbf{q}_t^*$ for all $t \in [\tau', +\infty)$ if*

$$\log \omega^A(\tau) = -\log D^A(t-\tau) + \alpha(t), \text{ for all } t \geq \tau, \text{ and } \tau \in \mathcal{T}, \quad (37)$$

for some mapping $\alpha : \mathbb{R}_+ \rightarrow \mathbb{R}$, and $\frac{\partial \omega^R(\tau)}{\partial \tau} = 0$.

The proof of this Proposition is very similar to the proof of Proposition 2. For completeness, we will replicate it here. Assume, by way of contradiction, that the solution $(\mathbf{q}_t^*)_{t \geq \tau}$ to optimizing (36) subject to $\sum_{t \geq \tau} \boldsymbol{\rho}_t \cdot \mathbf{q}_t \leq y_\tau$ is no longer optimal from decision period τ' onwards. This implies there must exist another sequence $(\hat{\mathbf{q}}_t)_{t \geq \tau'}$ such that,

$$\begin{aligned} & \sum_{t \geq \tau'} \delta^{t-\tau'} [\omega^A(\tau') D^A(t-\tau') u^A(\hat{\mathbf{q}}_t) + \omega^R(\tau') D^R(t) u^R(\hat{\mathbf{q}}_t) + u^E(\hat{\mathbf{q}}_t)] \\ & > \sum_{t \geq \tau'} \delta^{t-\tau'} [\omega^A(\tau') D^A(t-\tau') u^A(\mathbf{q}_t^*) + \omega^R(\tau') D^R(t) u^R(\mathbf{q}_t^*) + u^E(\mathbf{q}_t^*)], \end{aligned} \quad (38)$$

and given assumption (37), we can write $a(t) = \omega^A(\tau') D^A(t-\tau')$ and $b(t) = \omega^R(\tau') D^R(t)$. Multiplying both sides of (38) by $\delta^{\tau'-\tau}$, we obtain the following,

$$\begin{aligned} & \sum_{t \geq \tau'} \delta^{t-\tau} [\omega^A(\tau') D^A(t-\tau') u^A(\hat{\mathbf{q}}_t) + \omega^R(\tau') D^R(t) u^R(\hat{\mathbf{q}}_t) + u^E(\hat{\mathbf{q}}_t)] \\ & > \sum_{t \geq \tau'} \delta^{t-\tau} [\omega^A(\tau') D^A(t-\tau') u^A(\mathbf{q}_t^*) + \omega^R(\tau') D^R(t) u^R(\mathbf{q}_t^*) + u^E(\mathbf{q}_t^*)]. \end{aligned} \quad (39)$$

Adding $\sum_{t=\tau}^{\tau'-1} \delta^{t-\tau} [\omega^A(\tau') D^A(t-\tau') u^A(\mathbf{q}_t^*) + \omega^R(\tau') D^R(t) u^R(\mathbf{q}_t^*) + u^E(\mathbf{q}_t^*)]$ to both sides of the inequality in (39), we obtain:

$$\begin{aligned} & \sum_{t \geq \tau} \delta^{t-\tau} [a(t) u^A(\hat{\mathbf{q}}_t) + b(t) u^R(\hat{\mathbf{q}}_t) + u^E(\hat{\mathbf{q}}_t)] \\ & > \sum_{t \geq \tau} \delta^{t-\tau} [a(t) u^A(\mathbf{q}_t^*) + b(t) u^R(\mathbf{q}_t^*) + u^E(\mathbf{q}_t^*)]. \end{aligned} \quad (40)$$

But (40) is in contradiction with the optimality of $(\mathbf{q}_t^*)_{t \geq \tau}$. We therefore conclude that $(\mathbf{q}_t^*)_{t \geq \tau'}$ is also optimal for τ' . Given that τ, τ' were chosen randomly, we have the desired result.

C Additional results

C.1 Predictability of prices

In this appendix we discuss the time series of spot prices (Figure 2), interest rates (Figure 3), and log discounted prices (Figure 4) for our sample, along the lines of Blow et al. (2021). We normalize all prices by using the corresponding mean prices in 1992 as the base. Figure

2 shows that all spot prices are increasing over the observation period. One exception is the price of petrol, which declines between 1985 and 1988. Figure 3 summarizes the evolution of the yearly interest rate on consumer loans. This interest rate varies between 10% and 20%. Yet it is worth noting that, due to compounding of interest across time, the discount rate is systematically increasing over time. Because of this, the discounted prices in Figure 4 are declining. The log discounted price curves are moreover fairly linear. This reflects a relatively uniform (negative) growth rate over time.

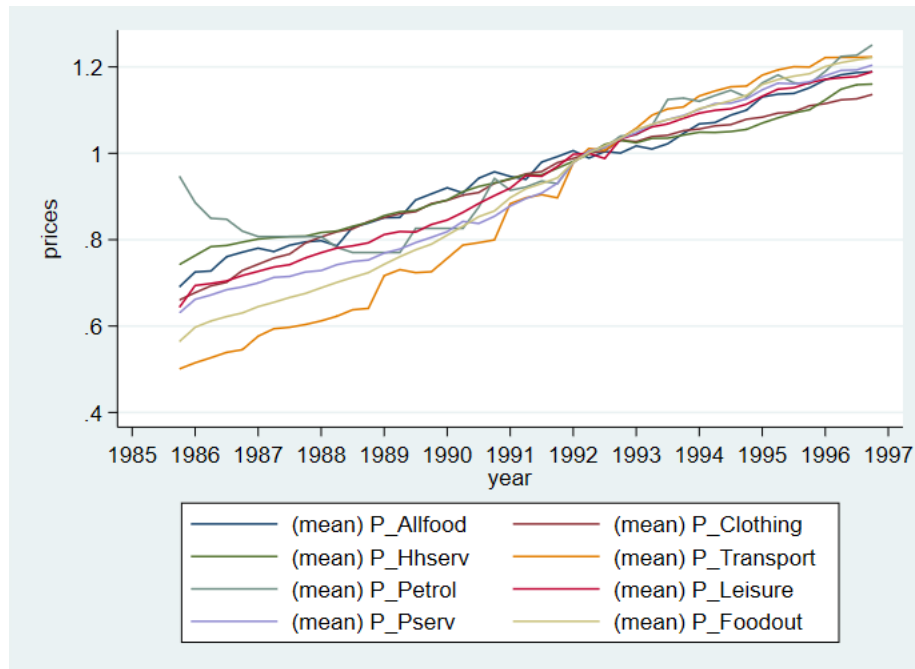


Figure 2: Time series of spot prices

Like Blow et al. (2021), we subsequently study predictability by regressing the (log) discounted prices of each commodity on a linear time trend. We do this for every commodity and for every consumer separately, because the sequence of observations differs by consumer. This leads to 16,416 linear regressions. The R-squared of each regression reflects the degree of (temporal) variation in prices that is captured by the time trend. Higher R-squared values imply better predictability. The histogram in Figure 5 shows the distribution of the R-squared among 16,416 consumer-commodity pairs. The median R-squared among all consumer-commodity regressions is as high as 93%.




Figure 3: Time series of interest rate on consumer loans




Figure 4: Time series of log discounted prices




Figure 5: Distribution of the proportion of variation in log discounted prices that is predictable from a linear time trend

C.2 Explaining consistency with *ICAES*

	<i>ICAES</i>
<i>ICAES</i>	
allexp	0.000 (0.000)
Between 26 and 35	0.142 (0.164)
Between 36 and 45	0.094 (0.169)
Between 46 and 55	-0.173 (0.189)
Older than 55	-0.240 (0.190)
Couple without children (under 14 y.o)	-0.123 (0.120)
Couple with a child	-0.177 (0.119)
Couple with two children	-0.311*** (0.116)
Couple with three or more children	-0.443*** (0.146)
Observations	2052

Table C.1: Regression of consistency with *ICAES* on total intertemporal budget, age, household composition, and year dummies. Age data is available in intervals 26 to 35 (dummy 2), 36 to 45 (dummy 3), 46 to 55 (dummy 4) and older than 55 (dummy 5). The sample consists of all consumers. Year dummies, all insignificant, are suppressed for compactness.

C.3 Extensions (Pass rates)

Pass rates	Singles				Couples			
	<i>ICES</i>	<i>ICAES</i>	<i>ICRES</i>	<i>ICARES</i>	<i>ICES</i>	<i>ICAES</i>	<i>ICRES</i>	<i>ICARES</i>
Baseline	0.093	0.494	0.820	0.948	0.035	0.455	0.795	0.964
Liq Constr	0.174	0.512	0.988	0.994	0.088	0.468	0.988	0.995
Uncertainty								
$\alpha_{1/5}$	0.203	0.651	0.930	0.977	0.137	0.680	0.913	0.984
$\alpha_{1/4}$	0.267	0.762	0.965	0.977	0.210	0.762	0.929	0.986
$\alpha_{1/3}$	0.349	0.814	0.965	0.977	0.268	0.799	0.935	0.987
$\alpha_{1/2}$	0.512	0.901	0.971	0.983	0.417	0.903	0.953	0.989
α_1	0.983	1.000	1.000	1.000	0.986	0.998	0.998	0.999
Error								
$\varepsilon_{1/5}$	0.198	0.564	0.924	0.971	0.130	0.575	0.901	0.980
$\varepsilon_{1/4}$	0.250	0.593	0.936	0.977	0.177	0.614	0.918	0.984
$\varepsilon_{1/3}$	0.331	0.674	0.965	0.983	0.261	0.675	0.940	0.987
$\varepsilon_{1/2}$	0.500	0.750	0.983	0.994	0.439	0.768	0.964	0.993
ε_1	1.000	1.000	1.000	1.000	0.996	0.997	1.000	1.000
First quarter	0.118	0.490	0.843	0.922	0.043	0.480	0.799	0.977

Table C.2: Pass rates for different specifications of *ICARES*, with liquidity constraints, uncertainty, measurement error, and for subsamples that start in the first quarter