



KATHOLIEKE UNIVERSITEIT
LEUVEN

DISCUSSION PAPER SERIES
12.12

AUGUST 2012



Consume now or later? Time inconsistency, collective choice and revealed preference

*Abi ADAMS, Laurens CHERCHYE,
Bram DE ROCK and Ewout VERRIEST*

Public Economics

Faculty of Economics
And Business



Consume Now or Later?

Time Inconsistency, Collective Choice and Revealed Preference*

Abi Adams[†] Laurens Cherchye[‡] Bram De Rock[§] Ewout Verriest[¶]

August 28, 2012

Abstract

This paper develops a revealed preference methodology for exploring whether time inconsistencies in household choice are the product of nonstationarities at the individual level or the result of individual heterogeneity and renegotiation within the collective unit. An empirical application to household-level microdata highlights that an explicit recognition of the collective nature of choice allows the vast majority of household behaviour to be rationalised by theory that assumes preference stationarity at the individual level. For our particular short panel data set, simply permitting limited intrahousehold heterogeneity in time preferences allows the choices of 98.4% of the sample to be rationalised by a model that assumes exponential discounting at the individual level. We also find that couples characterized by lower divergence in spousal discount rates are older, more likely to have children and wealthier, which we take as indications of experiencing higher match quality.

JEL Classification: D11, D12, D13, C14.

Keywords: time consistency, collective choice, full efficiency, renegotiation, revealed preference.

1 Introduction

Understanding the dynamics of intertemporal choice is a valuable, and increasingly salient, goal for modern society. Life expectancy, and life expectancy upon retirement, are on an upward trajectory, pushing issues of pension provision and retirement planning to the forefront. These trends are especially important given the rapid transition to defined-contribution pension schemes in developed countries, under which the characteristics of dynamic choice and individual self-control issues have more relevance (Thaler and Benartzi 2004). Traditional economic models of intertemporal choice are increasingly criticised for adding little to the analysis of issues surrounding these trends, lacking normative, descriptive and predictive value.

The discounted utility benchmark. Samuelson’s (1937) canonical “discounted utility” (DU) model is the standard framework by which economists conceptualise intertemporal choice. Consider a consumer who faces a consumption stream $\mathbf{C}_{ij} = \{\mathbf{q}_t\}_{t=i,\dots,j}$, with \mathbf{q}_t the consumption quantities at time t . Under the DU model, consumer preferences are represented by the following functional form:

$$U(\mathbf{C}_{ij}) = \sum_{t=i}^j \beta^{t-i} u(\mathbf{q}_t),$$

*We thank Richard Blundell, Ian Crawford, Vincent Crawford and David Gill for useful discussions.

[†]Department of Economics, University of Oxford, Manor Road Building, Manor Road, Oxford OX1 3UQ, United Kingdom. E-mail: abi.adams@economics.ox.ac.uk.

[‡]Center for Economic Studies, University of Leuven. E. Sabbelaan 53, B-8500 Kortrijk, Belgium. E-mail: laurens.cherchye@kuleuven-kulak.be. Laurens Cherchye gratefully acknowledges financial support from the Research Fund K.U.Leuven through the grant STRT1/08/004.

[§]ECARES-ECORE, Université Libre de Bruxelles. Avenue F. D. Roosevelt 50, CP 114, B-1050 Brussels, Belgium. E-mail: bderock@ulb.ac.be. Bram De Rock gratefully acknowledges the European Research Council (ERC) for his Starting Grant.

[¶]Center for Economic Studies, University of Leuven. E. Sabbelaan 53, B-8500 Kortrijk, Belgium. E-mail: Ewout.Verriest@kuleuven-kortrijk.be. Ewout Verriest gratefully acknowledges the Fund for Scientific Research - Flanders (FWO-Vlaanderen) for financial support.

where u is a felicity function and $\beta = 1/(1 + \sigma)$, with $\sigma \in [0, \infty)$ denoting the consumer’s discount rate. These preferences are *time consistent*: the choice between alternatives does not depend on when in time that choice occurs. Thus, if receiving X at t is preferred to receiving Y at $t + d$, the decision maker will always prefer X at τ to Y at $\tau + d$. Results from Koopmans (1960) provide an axiomatic foundation to the DU model and highlight that preferences are time consistent only if they can be expressed in the above format.

However, it is its virtues of simplicity and tractability, rather than claims concerning its normative validity or predictive accuracy, that explains the widespread application of the DU model. Samuelson himself expressed skepticism over the descriptive worth of the approach, commenting that “it is completely arbitrary to assume that the individual behaves so as to maximise an integral of the form envisaged in [the discounted-utility model]”. Yet, ease and elegance won out and the framework quickly became the standard way of representing intertemporal preferences in dynamic optimisation frameworks.

Despite its theoretical virtues, casual observation and empirical studies consistently cast doubt on the validity of the DU model. Decision makers systematically behave in a time inconsistent manner, acting impatiently today while planning to act patiently in the future. For example, \$100 may be preferred today to \$110 tomorrow, but \$110 in 31 days to \$100 in 30 days. Such “preference reversals” are well documented in the psychology (Green, Fristoe and Myerson 1994; Kirby and Herrnstein 1995) and economics literature (for a survey, see Frederick *et al.* 2002). Good intentions are often expressed regarding saving and responsible consumption, but the implementation of such plans tends to fall by the wayside. For example, in Choi *et al.*’s (2001) survey of 401(k) participants, 86% of the self-reported undersavers who expressed an intention to increase their savings rate, had made no changes to their behaviour four months later.

Collective choice and time inconsistency. Research effort has largely focused on modelling the sources of time inconsistency at the individual level. Standard methods of modelling discounting have been a prime target of criticism, with the behavioural economics literature increasingly favouring frameworks that assume hyperbolic discount functions. Hyperbolic discount functions are characterised by a relatively high discount rate over short horizons and a relatively low rate over long horizons. This lack of constancy in the discount rate introduces a conflict between today’s and future preferences and a “present bias” to decision making. Hyperbolic discounting has been offered as an explanation for many stylised facts, from under-saving and excess co-movement of income and consumption (Laibson 1997, 1998; Angeletos *et al.* 2001) to procrastination, addiction and lack of exercise (O’Donoghue and Rabin 1999, 2001; Gruber and Koszegi 2001; della Vigna and Malmendier 2006).

This paper takes a different approach. We consider how acknowledging the *collective* nature of choice can rationalise time inconsistencies in revealed preferences. The preference structure associated with the DU approach is often applied to model group behaviour without modification. Under this “unitary” approach, one assumes the collective acts as a single decision making unit, and therefore can be treated as if a rational individual. No explicit allowance is made for the separateness of persons and preference heterogeneity within a group.

Yet, many dimensions of intertemporal choice are better modelled as the outcome of group, rather than individual, decision making. Savings decisions are typically made at the household level, and the allocation of budgets over time within firms and committees are the product of multi-party deliberation. Even in the context of individual choice, one can consider the existence of multiple selves with distinct personalities rather than a single decision making unit. Thaler and Shefrin (1981) contrast the long-sighted “planner” within us to the short-sighted “doer”, while Metcalfe and Mischel (1999) contrast our “hot” and “cool” systems. There is scientific evidence in support of the multiple selves hypothesis. Psychological studies find that different decision making systems interact within the brain during the evaluation of intertemporal prospects, and that these systems respond differently to the temporal dimensions of reward (Ainslie *et al.* 2005; McClure *et al.* 2004, 2007). Such evidence supports the application of collective choice models to characterise the behaviour of individuals.

Acknowledgement of the collective nature of choice can help to rationalise the apparent time inconsistencies in household choice behaviour. In a group context, inconsistencies can arise simply from the aggregation of heterogeneous preferences. Both variation in individual discount rates and innovations in the Pareto weight summarising the household decision making process represent relevant considerations in this regard. First, consider the effect of discount rate heterogeneity. Deriving a time independent discount rate from the underlying preferences of a heterogeneous population has long been recognised as problematic (Marglin

1963; Feldstein 1964). When time preferences within a group differ, the collective preference is typically time inconsistent, even when the underlying population has perfectly time consistent preferences (Jackson and Yariv 2011; Zuber 2010). A present bias can be introduced to household choice through the aggregation of heterogeneous preferences. As time passes, the preferences of more impatient individuals are weighted less in the group’s overall maximisation problem, causing the effective aggregate discount rate to increase over time. In fact, Jackson and Yariv (2011) show that, for a uniform distribution of discount rates in an otherwise homogeneous population, group utility maximisation generates aggregate behaviour that corresponds to hyperbolic discounting.

Further, renegotiations of the household choice rule can generate nonstationarities in family behaviour. Relative decision making power within the collective unit can vary, and differences in time preferences can prompt periodic innovations in the intrahousehold preference weighting. Other things equal, it is optimal to relatively favour impatient group members in early periods and patient members in later periods. However, as time passes and impatient members start to receive lower shares of the group surplus, there is an incentive for them to demand a renegotiation of allocations in their favour or threaten to leave the group. Renegotiations prompt changes in the intra-household preference weighting, generating nonstationarities in the collective preference.

Although it is true that nonstationarities at the individual level will translate into a failure of time consistency at the collective level, understanding whether the primary locus of inconsistent behaviour is at the individual self or group level is important from both a methodological and policy perspective. The DU preference structure is tractable and parsimonious. Thus, if it cannot be rejected on the basis of choice behaviour, an appeal to Occam’s razor provides a compelling analytic reason for its retention. Further, policy design should be influenced according to whether time inconsistent behaviour is the product of individual nonstationarities or collective aggregation issues.

Our contribution. This paper puts forward a nonparametric characterisation of household intertemporal choice and develops a revealed preference methodology for analysing the sources of collective time inconsistency. The approach taken follows in the spirit of Afriat (1967), Diewert (1973), and Varian (1982) and incorporates insights gained from the extension of the revealed preference methodology to an intertemporal setting by Browning (1989) and Crawford (2010), and to the collective model by Cherchye, De Rock and Vermeulen (2007). Our empirical strategy allows for the recovery of theory-consistent spousal discount rates and an assessment of the degree of intrahousehold commitment. The framework presented allows us to explore whether time inconsistencies in household choice can be rationalised by preference heterogeneity and renegotiation within the collective unit rather than individual nonstationarities. We demonstrate the value of our approach via a practical application to household microdata.

Our methodology is novel in this context and has clear advantages over existing empirical tests of time consistency and household intertemporal behaviour. Current tests of dynamic collective choice models and time discounting are parametric and tend to reject the assumptions of constant discounting (see Frederick *et al.* 2002) and a time-independent intrahousehold preference weighting (Lich-Tyler 2004; Mazzocco 2007). However, this may be due to restrictive parametric assumptions, rather than a failure of the underlying, functional form free, theoretical framework.

Empirical studies of discounting and time consistency are sensitive to the parametric specification employed. The common assumption of linear consumption utility imparts an upward bias to discount rate estimates and is thought to contribute to the unrealistically high discount rates observed in the literature. Although recent developments have seen the linear-utility specification somewhat relaxed (Anderson *et al.* 2008; Andreoni and Sprenger 2012), estimates are still dependent on the set of functional form assumptions made concerning the form of the utility function. The experimental nature of existing time discounting studies can also be critiqued. Dohman *et al.* (2012) highlight that elicited preferences are not procedurally independent and discount rate estimates are hugely sensitive to the experimental design.

Turning to studies of intrahousehold commitment and renegotiation, Mazzocco’s (2007) tests of intertemporal decision making and the constancy of the intrahousehold preference weighting centre on how “distribution factors” (Browning *et al.* 1994) enter log-linearised household Euler equations.¹ The risk of misspecification is high in this context. Carroll (2001), using simulated data, finds that log-linearisation introduces

¹Distribution factors are factors that influence the relative power of family members but are independent of their preferences. See also Section 2.1 for a formal discussion.

significant bias to the estimation of preference parameters, concluding that empirical estimation of approximated Euler equations should be abandoned. Furthermore, this strategy requires the observation of all relevant distribution factors. This requirement is hugely demanding, especially given the limited information available in household survey data.

The methodology and empirical application presented in this paper avoids such criticisms. Our revealed preference approach is wholly nonparametric and, thus, our results are not contingent on any particular specification of family member utility functions. Rather than directly estimate the preference parameters that best “fit” with some assumed functional specification, we ask whether there exists a non-empty feasible set to the system of inequalities implied by maximising household behaviour within the framework imposed by economic theory. The existence of a non-empty feasible set to these inequalities is then a necessary and sufficient condition for household behaviour and the theory in question to be consistent. Using this approach, we are able to determine whether time inconsistencies in the revealed household preference can be explained as a result of discount rate heterogeneity and/or imperfect commitment within the collective unit, or whether one must additionally allow for nonstationarities at the individual level. In fact, if our nonparametric tests support a particular theoretical model, then this may provide also provide a strong motivation for parametric studies using exactly this model. In this sense, our nonparametric assessment here may also be seen as a complement to existing parametric studies.

Our tests are also explicitly designed for use with household consumption data, although they can also be profitably applied to an experimental setting. Our empirical application is thus one of the few in recent years to be fully grounded in “real world” household behaviour, rather than make use of choice data that has been elicited in an artificially constructed environment. This allows us to avoid many of the procedural nuances that plague experimental studies.

We find that accounting for the collective nature of choice allows us to rationalise time inconsistencies in aggregate household behaviour without positing nonstationarities in individual preferences. Simply allowing for some limited heterogeneity in familial discount rates allows the behaviour of 98.4% of households in our sample to be rationalised by standard models of household intertemporal behaviour.

Outline. This paper proceeds as follows. Section 2 defines time consistency of the collective preference and outlines the associated revealed preference restrictions for establishing the time consistency of group choice. Section 3 uses these revealed preference conditions to evaluate the empirical validity of a time consistent household consumption for a Spanish panel of household microdata. We find that the hypothesis of time consistent household choice is heavily rejected for our data set, even using nonparametric revealed preference restrictions. Given this result, Section 4 explores how a recognition of the collective nature of household choice can rationalise nonstationarities in the revealed collective preference and derives a nonparametric methodology for testing hypotheses on the sources of time inconsistent household behaviour. Section 5 continues our empirical application and provides strong empirical support for a collective rationalisation of observed time inconsistencies. This section also considers the relationship between intrahousehold time preference heterogeneity and observable household characteristics. Section 6 concludes this paper.

2 Time consistency and collective choice

The aim of this paper is to provide a framework for exploring the sources of time inconsistencies in household choice. Specifically, we wish to determine whether explicitly accounting for the collective nature of choice can allow one to rationalise patterns of household behaviour without positing nonstationarities at the individual level. This section formally defines the concept of time consistency tested in this paper and derives simple revealed preference conditions that can be used to determine the time consistency of observed household choice.

2.1 The “collective” preference

Collective models explicitly recognise the separateness of persons within the household, allowing for complete heterogeneity in family member felicity functions and discount rates. For notational simplicity, we focus on

a two-member household, constituted of members $m \in \{A, B\}$. The extension of results to an M -member ($M > 2$) household is straightforward.

Individual preferences are represented by a time-additive discounted utility function, defined over private and public consumption. We assume N private goods and K public goods. At a given time t , private quantities $\mathbf{q}_t^m \in \mathbb{R}_+^N$, with associated (discounted) prices $\mathbf{p}_t \in \mathbb{R}_{++}^N$, are consumed non-jointly, while public quantities $\mathbf{Q}_t \in \mathbb{R}_+^K$, with associated (discounted) prices $\mathbf{P}_t \in \mathbb{R}_{++}^K$, are consumed jointly and non-exclusively. Thus, associated with each member m is a concave and strictly increasing felicity function u^m and discount rate $\sigma_m \in [0, \infty)$, such that a stream of public and private consumption $\mathbf{C}_{ij}^m = \{\mathbf{q}_t^m, \mathbf{Q}_t\}_{t=i, \dots, j}$ is evaluated as

$$U^m(\mathbf{C}_{ij}^m) = \sum_{t=i}^j \beta_m^{t-1} u^m(\mathbf{q}_t^m, \mathbf{Q}_t),$$

where $\beta_m = 1/(1 + \sigma_m)$.

Collective models do not assume a priori that individual preferences can be aggregated into a single time-independent household felicity function and the framework does not specify a single intrahousehold bargaining process. Rather, it simply assumes that some cooperative decision making process exists and that this process leads to Pareto efficient outcomes over the affordable budget set.² Given these assumptions, one can define the Pareto weight ω_t to summarise the bargaining process in period t . Consider a time horizon $|T|$, $T = \{1, \dots, |T|\}$. For household H , the ‘‘collective’’ preference over some lifecycle consumption profile $\mathbf{C}^H = \{\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t\}_{t \in T}$ is then given by

$$U^H(\mathbf{C}^H) = \sum_{t \in T} \{\beta_A^{t-1} u^A(\mathbf{q}_t^A, \mathbf{Q}_t) + \omega_t \beta_B^{t-1} u^B(\mathbf{q}_t^B, \mathbf{Q}_t)\},$$

with $\omega_t = f(\mathbf{Z}_t)$, where \mathbf{Z}_t denotes the set of relevant ‘‘distribution factors’’ at time t . The standard theory places no restrictions on what variables count as relevant distribution factors, beyond requiring that they are independent of individual preferences (Browning *et al.* 1994). This lack of structure makes our nonparametric framework especially attractive as, unlike parametric tests of intertemporal behaviour, our methodology does not require a formal specification of the factors that jointly determine ω_t .

2.2 Time consistency

Given a household’s time series of consumption choices, the first question of interest is whether there exists a time consistent household utility function that could have generated this choice pattern.

What does it mean for the collective preference to be time consistent? Following conditions given by Koopmans (1960), time consistency is seen primarily as a stationarity restriction: the passing of time should not have an effect on the preference ordering over fixed consumption streams. Formally, consider a consumption stream $\mathbf{C}_{ij}^H = \{\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t\}_{t=i, \dots, j}$. The collective preference over this consumption stream is defined as

$$U^H(\mathbf{C}_{ij}^H) = \sum_{t=i}^j \{\beta_A^{t-1} u^A(\mathbf{q}_t^A, \mathbf{Q}_t) + \omega_t \beta_B^{t-1} u^B(\mathbf{q}_t^B, \mathbf{Q}_t)\}.$$

Let $\mathbf{C}_{i'j'}^H$ represent an ‘‘updated’’ counterpart of \mathbf{C}_{ij}^H ($i, j \in T$, $i < j$), which denotes the consumption streams \mathbf{C}_{ij}^H shifted forward into the future by some amount $0 \leq \tau \leq |T| - j$, and let $\{\mathbf{C}_{ij}^H, \mathbf{C}_{kl}^H\}$ the combination of two (non-overlapping) consumption streams \mathbf{C}_{ij}^H and \mathbf{C}_{kl}^H ($i, j, k, l \in T$, $i < j$, $k < l$ and $j < k$ or $l < i$). Then, we can formally define time consistency of the collective preference.

Definition 1 *The collective preference is time consistent if the following two conditions are satisfied:*

1. For any $i, j \in T$ with $i < j$,

$$U^H(\mathbf{C}_{ij}^H) > U^H(\tilde{\mathbf{C}}_{ij}^H) \text{ if and only if } U^H(\mathbf{C}_{i'j'}^H) > U^H(\tilde{\mathbf{C}}_{i'j'}^H).$$

²See, for example, Chiappori (1988, 1992) and Browning and Chiappori (1998) for detailed discussions of the Pareto efficiency assumption in collective household models. Mazzocco (2007) discusses this framework in an intertemporal context.

2. For any $i, j, k, l \in T$ with $i < j$, $k < l$ and, in addition, $j < k$ or $l < i$,

$$U^H(\{\mathbf{C}_{ij}^H, \mathbf{C}_{kl}^H\}) > U^H(\{\tilde{\mathbf{C}}_{ij}^H, \mathbf{C}_{kl}^H\}) \text{ if and only if } U^H(\{\mathbf{C}_{ij}^H, \bar{\mathbf{C}}_{kl}^H\}) > U^H(\{\tilde{\mathbf{C}}_{ij}^H, \bar{\mathbf{C}}_{kl}^H\}).$$

The first condition in Definition 1 imposes stationarity on the collective preference; the ranking of consumption streams should not depend on when in time those streams are situated. The second condition requires the ranking of consumption streams to be independent of periods with identical consumption bundles.

Conditions for time consistency. Results originally given in Koopmans (1960), and applied to a collective setting by Jackson and Yariv (2012), imply that time consistency of the collective preference requires the ability to re-express the household preference in the following format:

$$U^H(\mathbf{C}_{ij}^H) = \sum_{t=i}^j \beta^{t-1} u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t),$$

for $\beta \in (0, 1]$. Thus, time consistency in our setting requires the existence of a single, constant household discount rate and a stationary household felicity function.

For the collective preference to be recast in the above format, two conditions need to hold. First, household members must discount the future at the same rate, $\beta_A = \beta_B = \beta$. Second, the intrahousehold decision making mechanism must give rise to a constant Pareto weight across the lifetime of the household, $\omega_t = \omega$ for all t . Under these conditions, for the lifecycle consumption profile \mathbf{C}^H we get

$$\begin{aligned} U^H(\mathbf{C}^H) &= \sum_{t \in T} \{\beta_A^{t-1} u^A(\mathbf{q}_t^A, \mathbf{Q}_t) + \omega_t \beta_B^{t-1} u^B(\mathbf{q}_t^B, \mathbf{Q}_t)\} \\ &= \sum_{t \in T} \beta^{t-1} u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t), \end{aligned}$$

where

$$u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t) = u^A(\mathbf{q}_t^A, \mathbf{Q}_t) + \omega u^B(\mathbf{q}_t^B, \mathbf{Q}_t).$$

Thus, only if the two conditions hold, the collective preference is time-consistent. The household can be modelled as a time-consistent representative agent with a latently separable, time-independent felicity function and constant discount rate.

2.3 Revealed preference conditions

Following the above, we assume $|T|$ observed consumption choices for household H , $T = \{1, \dots, |T|\}$. For each observation t , we know the privately consumed quantities, \mathbf{q}_t , and the publicly consumed quantities, \mathbf{Q}_t , as well as the corresponding (discounted) prices, \mathbf{p}_t and \mathbf{P}_t . This defines a set of observations $S = \{\mathbf{q}_t, \mathbf{Q}_t; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T}$. Note that we only assume that aggregate private quantities \mathbf{q}_t , but not the individual private quantities \mathbf{q}_t^A and \mathbf{q}_t^B (with $\mathbf{q}_t^A + \mathbf{q}_t^B = \mathbf{q}_t$), are observed. This assumption is motivated by the fact that, in most household surveys (including the one we use in our own application), information on “who gets what” is limited and the decomposition of private consumption into that consumed by household members is generally unobserved.

The revealed preference approach to establishing the time consistency of household behaviour asks whether one can find necessary and sufficient conditions under which observed choices can be rationalised by a stationary collective preference subject to the lifecycle budget constraint.

Definition 2 *The set of observations $S = \{\mathbf{q}_t, \mathbf{Q}_t; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T}$ can be rationalised by the time consistency model if there exist, for all $t \in T$, private quantities $\mathbf{q}_t^A, \mathbf{q}_t^B \in \mathbb{R}_+^N$ (with $\mathbf{q}_t^A + \mathbf{q}_t^B = \mathbf{q}_t$) and, in addition, a concave, strictly increasing felicity function u^H and discount factor $\beta \in (0, 1]$ such that*

$$\sum_{t \in T} \beta^{t-1} u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t) \geq \sum_{t \in T} \beta^{t-1} u^H(\zeta_t^A, \zeta_t^B, \zeta_t^H)$$

for all affordable consumption plans $\{\zeta_t^A, \zeta_t^B, \zeta_t^H\}_{t \in T}$ (with $\zeta_t^A, \zeta_t^B \in \mathbb{R}_+^N$ and $\zeta_t^H \in \mathbb{R}_+^K$) that satisfy

$$\sum_{t \in T} [\mathbf{p}'_t(\mathbf{q}_t^A + \mathbf{q}_t^B) + \mathbf{P}'_t \mathbf{Q}_t] \geq \sum_{t \in T} [\mathbf{p}'_t(\zeta_t^A + \zeta_t^B) + \mathbf{P}'_t \zeta_t^H].$$

In words, the data can be rationalised by a time consistent household preference if observed choices maximise discounted lifetime household utility out of affordable lifetime consumption plans for a coherent, stationary collective preference. We recognise that this definition of rationalisable behavior assumes perfect foresight and perfect capital markets. However, we consider how the robustness of our empirical results to the relaxation of these assumptions in Section 3.

The next theorem states the revealed preference conditions for a data rationalisation as defined above. We refer the reader to Appendix A for proofs of all our main results.

Theorem 1 *The set of observations $S = \{\mathbf{q}_t, \mathbf{Q}_t; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T}$ can be rationalised by the time consistency model if and only if there exist, for all $t \in T$, a utility number $u_t^H \in \mathbb{R}$ and a positive constant $\beta \in (0, 1]$ that, for any $s, t \in T$, satisfy*

$$u_s^H - u_t^H \leq \frac{1}{\beta^{t-1}} [\mathbf{p}'_t(\mathbf{q}_s - \mathbf{q}_t) + \mathbf{P}'_t(\mathbf{Q}_s - \mathbf{Q}_t)].$$

Theorem 1 is an equivalence result. In words, if there exists a household discount rate β and constants $\{u_t^H\}_{t \in T}$ such that the stated inequalities hold, then there exists a stationary household felicity function and discount rate that provide a perfect within-sample rationalisation of the data. The existence of a non-empty feasible set to the inequalities implies that one cannot reject the hypothesis of discount rate homogeneity and a stable intertemporal weighting of individual preferences within the household. Determining the actual existence of such a non-empty feasible set is easily done empirically. Conditioning on β , the inequalities defined by Theorem 1 are linear in unknowns and thus can be verified using standard linear programming techniques, which we use in our following empirical application.

3 An empirical test of time consistency

In this section we test the time consistency conditions in Theorem 1 for household panel data, taken from the Spanish Continuous Family Expenditure Survey (the *Encuesta Continua de Presupuestos Familiares*, ECPF). We strongly reject the hypothesis of a time consistent collective preference, as defined above. We also examine the robustness of this result, evaluating the possibility that actual behaviour is effectively time consistent but some of the (admittedly strong) assumptions underlying our revealed preference conditions are violated. Specifically, we consider measurement error, imperfect capital markets and imperfect foresight. However, we find that none of these alternative explanations convincingly “rationalises” the observed rejections of time consistency. This result motivates our focus on individual heterogeneity within households as a source of time inconsistency in the revealed household preference in Sections 4 and 5.

3.1 The data

The ECPF is a quarterly budget survey in which households are randomly rotated at a rate of 12.5% each quarter. Participating households are surveyed in the same week of each successive quarter, with each adult family member completing an expenditure diary in which they record their spending during the survey week. The data used are drawn from the period 1985-1997 for a sub-sample of couples in which the husband is in full time employment and the wife is out of the labour force in all periods. The requirement of stable employment status is to allow the potential nonseparability of leisure and consumption to be ignored for the time being, as our theoretical framework does not presently consider the household labour supply decision. For matters of comparability, we only consider those households that participated in the survey for eight consecutive quarters ($|T| = 8$), which is the maximum number of observations in the ECPF. This still leaves a sizeable sample of 1585 households.

We consider household choice over a commodity bundle of 15 nondurable goods. Each good is classified as either ‘private’ or ‘public’ consumption. Our bundle of private consumption consists of 11 ($= N$) goods:

(1) Food and non-alcoholic drinks, (2) Alcohol, (3) Tobacco, (4) Clothing and footwear, (5) Nondurable medicines, (6) Medical services, (7) Transportation, (8) Petrol, (9) Leisure (cinema, theatre, clubs for sport), (10) Personal services, (11) Restaurants and bars. Our bundle of public consumption consists of 4 ($= K$) goods: (1) Rent, (2) Energy, (3) Home Services (heating, water and furniture repair) and (4) Nondurables at home (cleaning products). Prices are calculated from published prices aggregated to correspond to the listed expenditure categories, discounted by the average interest rate on consumer loans.³ We refer to Browning and Collado (2001) and Crawford (2010) for more details on the ECPF data.

3.2 Revealed preference tests

Revealed preference tests of a given model are defined by hypotheses of the following form:

\mathbf{H}^0 : Household behaviour can be rationalised by the model.

\mathbf{H}^1 : Household behaviour cannot be rationalised by the model.

Our tests thus yield ‘yes/no’ answers; either household behaviour is consistent with the model in question or it is not. A “yes” result implies that the model cannot be rejected on the basis of observed behaviour. However, it does not necessarily imply that the model is “the truth”. Popper (1959) points to the logical asymmetry between verification and falsification. No number of observed passes of model X allows one to derive the universal statement: “All households can be rationalised by model X ”. However, failure of a revealed preference tests allows us to logically derive the conclusion: “The household cannot be rationalised by model X ”.

As explained in Section 2, testing whether household consumption choices can be rationalised by a time-consistent household utility function boils down to checking the linear conditions defined by Theorem 1 for a given value of the discount factor β . Therefore, in our empirical application we conduct a grid search on β . In particular we report results for a grid search of individual discount factors on $[0.9, 1]$ with a spacing of 0.005.⁴ At this point it is worth noting that all our following results (reported in Tables 1-9 and Figures 1-4) are robust to alternative grid search specifications (including a grid search across the full interval $(0, 1]$). These robustness checks can be obtained from the authors on request.

We also remark here that our tests allow for unrestricted preference heterogeneity across households, as captured by the individual-specific felicity functions. The theory-consistency of each household’s behaviour is tested independently and the data is not pooled at any stage. Therefore, when we refer to “the data”, we in fact refer to “the data for some household h ”.

3.3 Test results

Table 1 reports the test results obtained from applying the revealed preference conditions in Theorem 1 to the ECPF data. Only two households in our sample of 1585 households can be rationalised by a time-consistent collective preference. Putting it differently, the pass rate is an exceptionally low 0.13%, representing a decisive failure of the time consistent model. This result is consistent with Crawford (2010), who reports a pass rate of less than 5% for a closely similar model and his particular sample of ECPF households. Thus, we can safely conclude that the data does not endorse the applicability of the time consistent model as it stands to explain household intertemporal choice behaviour.

Table 1: Time consistency

	<i>Number of households</i>	<i>Percentage (%) of households</i>
<i>Pass rate</i>	2	0.13

The following subsections report additional results that lay claim to the robustness of our finding. Further, in Section 5, we evaluate the time consistent model in terms of two performance measures that are frequently used in revealed preference analysis such as ours: discriminatory power and predictive success. The results reported in Section 5 provide further empirical evidence against the time consistent model for explaining observed household behaviour.

³The *actual* (quarterly) interest rate on savings and/or loans faced by a particular household (i_t^*) may differ from the publicly listed (quarterly) interest rate on consumer loans (i_t). However, the results prove robust if we account for some deviations, e.g. by assuming $i_t^* = i_t \pm \epsilon$, where $\epsilon = \pm 0.01$ or ± 0.02 .

⁴This corresponds to a search for discount rate on $[0, 0.11]$, which given the quarterly periodicity of our data is not unreasonable.

3.4 Robustness checks

Our revealed preference tests are based upon a number of very exacting conditions and make some implicit assumptions that do not specifically relate to the time consistency of household behaviour. It is plausible that measurement error or a failure of perfect foresight and perfect capital market assumptions, rather than any time inconsistencies in household choice, explain the decisive rejection of the time consistent model. In the following sections, we explore the impact of relaxing these background assumptions on the pass rate of the time consistent model and end by concluding that none of these factors can fully account for the empirical failings of the model.

3.4.1 Measurement error

The strong rejection of the time consistent model could be the result of errors in listed expenditure data, rather than deviations of actual choice from model’s prescriptions. Revealed preference tests are “sharp” in that household behaviour is either consistent with the model in question or it is not. Thus, small deviations in observed quantities away from the “truth” could have a large impact on pass rates. We apply Varian’s (1985) procedure to explore the possibility that measurement error is behind the rejection of our revealed preference tests. We restrict ourselves here to an outline of the main intuition behind our measurement error procedure; a formal presentation is given in Appendix B.

Consider a data set S that does not meet the conditions defined by Theorem 1. Our measurement error procedure first estimates the minimal adjustments of the quantity data required to obtain consistency with the revealed preference conditions, and then calculates the standard deviation on measurement error in quantities that would be necessary for the null hypothesis of theory rationalisable behaviour to be accepted at some significance level α .

For a household observation t , let $q_{t,n}$ represent the observed quantity of the n -th private good and $Q_{t,k}$ the observed quantity of the k -th public good. Due to measurement error, these quantities may deviate from $q_{t,n}^*$ and $Q_{t,k}^*$, the true (but unobserved) values of the private and public quantities. The divergence between listed and true quantities are quantified by the “relative quantity errors”

$$\epsilon_{t,n} = \frac{q_{t,n}^* - q_{t,n}}{q_{t,n}} \text{ and } \varepsilon_{t,k} = \frac{Q_{k,n}^* - Q_{k,n}}{Q_{k,n}},$$

As we cannot observe the true values, $q_{t,n}^*$ and $Q_{t,k}^*$, we approximate the necessary measurement error required for listed expenditure data to pass the revealed preference conditions by calculating the minimal adjustments of the quantity data required to obtain theory consistency. Under the assumption that relative quantity errors are independently normally distributed with zero mean and constant variance σ^2 , we then calculate the “critical” standard deviation σ_α on measurement error needed such that the null hypothesis of theory-rationalisable behaviour cannot be rejected at some significance level α . We then reject rationalisability at a significance level α if this σ_α exceeds our prior beliefs about the “true” standard deviation σ .

For each household and any significance level α , we can compute a critical standard deviation σ_α for household behaviour to be rationalisable by the time consistent model. Figure 1 shows the kernel distribution that corresponds to a significance level $\alpha = 0.05$, and Table 2 shows summary statistics on the distribution of σ_α for various significance levels. Our results are relatively insensitive to the chosen rejection level α . We observe considerable inter-household variation in σ_α , ranging from zero for the two households that were consistent with the time consistent model in the absence of measurement error, to approximately 17%. Thus, the entire sample can be rationalised by the time consistent model if we believed that the standard deviation of relative quantity errors ranged to 17%. To rationalise three-quarters of the households in our sample, we still require a standard deviation on relative quantity errors of approximately 8%. These results indicate that substantial measurement error in observed quantities is required if we are to rationalise a majority of households by the time consistent model. Given the unlikelihood that the error in the ECPF data is that large, we reject the hypothesis that measurement error explains the decisive rejection of the time consistent model.

Figure 1: Kernel plot of σ_α distribution




Table 2: Summary statistics of σ_α distribution

	<i>Mean</i>	<i>Min</i>	<i>Q1</i>	<i>Median</i>	<i>Q3</i>	<i>Max</i>
$\alpha = 0.01$	0.0614	0	0.0420	0.0592	0.0786	0.1610
$\alpha = 0.05$	0.0640	0	0.0437	0.0616	0.0819	0.1677
$\alpha = 0.10$	0.0654	0	0.0447	0.0630	0.0837	0.1715

3.4.2 Imperfect capital markets and imperfect foresight

Our theoretical framework assumes perfect capital markets and perfect foresight. These assumptions are clearly very strong even though we only consider household behaviour over a relatively short, two-year period. We here explore the impact that relaxing these assumptions could have on the pass rate of the time consistent model.

Imperfect capital markets. Under the assumption of perfect capital markets, households are able to borrow and save at the same interest rate. To relax this assumption, we use a method analogous to Demuynck and Verriest (2012). Again, we refer the interested reader to Appendix B for a formal presentation. Intuitively, we allow households to face some (unobserved) borrowing constraint that hinders their ability to smooth consumption across time periods. If this additional constraint never binds, the revealed preference conditions for the time consistent model are unchanged from the ones in Theorem 1. However, a binding borrowing constraint causes household consumption to deviate from its optimal path: a binding borrowing constraint in period t implies that total household consumption in t is less than it would be under unconstrained borrowing.⁵

Table 3 presents our results. Accounting for a binding borrowing constraint allows the pass rate of the time consistent model to increase to approximately 24%. Although this represents a sizeable improvement, it is still very modest. In Section 5 we show that an allowance for individual heterogeneity has a substantially more beneficial impact on our rationalisability results.

Table 3: Perfect capital markets versus imperfect capital markets

	<i>Perfect capital markets</i>	<i>Imperfect capital markets</i>
<i>Pass rate (%)</i>	0.13	23.66

Imperfect foresight. Finally, we relax the assumption that the household has perfect foresight with respect to all relevant information, such as future prices, incomes and interest rates. In many settings, this assumption is utterly implausible. However, given the short nature of our panel, the assumption may not

⁵This generally entails weaker revealed preference conditions (see Corollary 1 in Appendix B).

be overly restrictive. In the absence of market imperfections, assuming perfect foresight is equivalent to assuming a time-constant marginal utility of wealth. We relax the assumption of perfect foresight over the entire time horizon in a simple way, by reducing the set of observations for which a time-constant marginal utility of wealth is imposed. In particular, we sequentially reduce the length of a time series on a household from $|T| = 8$ to $|T'|$, $2 \leq |T'| \leq 7$. By construction, the pass rate will decrease monotonically as the set of observations shrinks.

Table 4 reports our test results. The time consistent model is still heavily rejected if perfect foresight is imposed for longer periods, suggesting that households may not be able to accurately forecast prices and incomes that are far into the future. However, in Section 5 we will see that perfect foresight over the entire time horizon of eight quarters cannot be rejected for the vast majority of households as soon as we allow for intrahousehold heterogeneity in discount rates. Therefore, we argue that the heavy rejection of the time consistent model cannot be solely attributed to the failure of our perfect foresight assumption.

Table 4: Imperfect foresight

<i>Time horizon</i> $ T' $	2	3	4	5	6	7	8
<i>Pass rate (%)</i>	90.54	51.80	20.19	4.86	1.20	0.25	0.13

In summary, this section has sought to establish the empirical validity of a time consistent model of household consumption behaviour. We have decisively rejected the hypothesis that this model provides an adequate framework for explaining observed patterns of family choice in our sample. This result is robust to relaxing the assumptions of no measurement error, perfect capital markets and perfect foresight upon which our testing framework is prefixed. Given this negative result, in the following sections we explore whether an explicit recognition of the collective nature of household choice can allow for the rationalisation of consumption patterns.

4 Collective choice and time inconsistency

The collective preference cannot be recast in the required format for time consistency in the presence of either heterogeneity in individual time preferences or innovations in the Pareto weight. The failure of these assumptions manifest themselves in different ways in observed choice behaviour. This section explores these sources of time inconsistency in greater depth and utilises Mazzocco's (2007) theoretical framework to develop an empirical strategy for distinguishing between the different sources of time inconsistent household behaviour.

4.1 Individual heterogeneity

Within the collective unit, individual heterogeneity and innovations in the Pareto weight, ω_t , can both introduce time inconsistencies to household consumption patterns, even if individuals within the family have perfectly time consistent preferences.

With intrahousehold discount rate heterogeneity, one cannot collapse individual discount rates into a single household rate of time preference, $\beta_A^t + \beta_B^t \neq (\beta_A + \beta_B)^t$. This implies that members' preferences are weighted differently in the household allocation problem at different points in time, even if the Pareto weight remains constant. Other things equal, the preferences of the more patient member become relatively more important in future time periods, introducing a time inconsistency to the collective preference.

Jackson and Yariv (2011) prove that for a group of otherwise homogeneous individuals choosing a common consumption stream, any heterogeneity in time preferences necessitates a present biased collective preference, and that with a uniform distribution of discount rates in a population, the collective utility function is hyperbolic. Thus, time inconsistencies in group behaviour need not be derivative of nonstationarities at the individual level. Rather, present biases in household choice can arise from the aggregation of heterogeneous time preferences.

Individual heterogeneity is the only source of time inconsistency in Mazzocco's (2007) "full efficiency" model. The model assumes the existence of a perfect commitment mechanism, removing the possibility of intrahousehold renegotiation. This implies the existence of a single, fixed Pareto weight to summarise the household decision making process.

For a given data set $S = \{\mathbf{q}_t, \mathbf{Q}_t; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T}$ the full efficiency model corresponds to the following rationalisation condition.⁶

Definition 3 *The set of observations $S = \{\mathbf{q}_t, \mathbf{Q}_t; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T}$ can be rationalised by the full efficiency model if there exist, for all $t \in T$, private quantities $\mathbf{q}_t^A, \mathbf{q}_t^B \in \mathbb{R}_+^N$ (with $\mathbf{q}_t^A + \mathbf{q}_t^B = \mathbf{q}_t$) and, in addition, concave, strictly increasing felicity functions u^A and u^B , discount factors $\beta_A, \beta_B \in (0, 1]$ and a Pareto weight $\omega > 0$ such that*

$$\sum_{t \in T} \{ \beta_A^{t-1} u^A(\mathbf{q}_t^A, \mathbf{Q}_t) + \omega \beta_B^{t-1} u^B(\mathbf{q}_t^B, \mathbf{Q}_t) \} \geq \sum_{t \in T} \{ \beta_A^{t-1} u^A(\zeta_t^A, \zeta_t^H) + \omega \beta_B^{t-1} u^B(\zeta_t^B, \zeta_t^H) \}$$

for all affordable consumption plans $\{\zeta_t^A, \zeta_t^B, \zeta_t^H\}_{t \in T}$ (with $\zeta_t^A, \zeta_t^B \in \mathbb{R}_+^N$ and $\zeta_t^H \in \mathbb{R}_+^K$), which satisfy

$$\sum_{t \in T} [\mathbf{p}'_t(\mathbf{q}_t^A + \mathbf{q}_t^B) + \mathbf{P}'_t \mathbf{Q}_t] \geq \sum_{t \in T} [\mathbf{p}'_t(\zeta_t^A + \zeta_t^B) + \mathbf{P}'_t \zeta_t^H].$$

As for the time consistent model, the constant Pareto weight ω incorporates the combined impact of all changes in distribution factors over time; it can be considered as the average relative power of family members across the lifetime of the household. Thus, with the existence of a perfect commitment mechanism, the only source of time inconsistent aggregate behaviour is discount rate heterogeneity. If $\beta_A = \beta_B = \beta$ then the collective preference could be recast in a time consistent format with a stationary household felicity function.

Revealed preference conditions. How can we test for the importance of individual heterogeneity as a source of time inconsistency? Discount rate heterogeneity negates the possibility of representing the collective preference in representative-consumer format. Given this, the composition of household consumption and its distribution between family members plays a central role in revealed preference tests of the full efficiency model. This has two important implications for the revealed preference conditions associated with the full efficiency model. First, for privately consumed goods, the information on \mathbf{q}_t^A and \mathbf{q}_t^B is relevant. Second, for publicly consumed goods, the relevant “prices” for an individual family member will be so-called Lindahl prices, \mathbf{P}_t^A and \mathbf{P}_t^B . These prices coincide with a family member’s marginal willingness to pay and, given the maintained assumption of cooperative decision making, sum to observed prices, $\mathbf{P}_t^A + \mathbf{P}_t^B = \mathbf{P}_t$.

Theorem 2 gives the condition under which household choice can be rationalised by the full efficiency model. If choices can be rationalised by the model, we cannot reject the hypothesis that time inconsistencies in aggregate behaviour are the result of discount rate heterogeneity within the family.

Theorem 2 *The set of observations $S = \{\mathbf{q}_t, \mathbf{Q}_t; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T}$ can be rationalised by the full efficiency model if and only if there exist, for all $t \in T$, private quantities $\mathbf{q}_t^A, \mathbf{q}_t^B \in \mathbb{R}_+^N$ (with $\mathbf{q}_t^A + \mathbf{q}_t^B = \mathbf{q}_t$), Lindahl prices $\mathbf{P}_t^A, \mathbf{P}_t^B \in \mathbb{R}_+^N$ (with $\mathbf{P}_t^A + \mathbf{P}_t^B = \mathbf{P}_t$), utility numbers $u_t^A, u_t^B \in \mathbb{R}$ and constants $\beta_A, \beta_B \in (0, 1]$ that, for any $s, t \in T$, satisfy*

$$\begin{aligned} u_s^A - u_t^A &\leq \frac{1}{\beta_A^{t-1}} [\mathbf{p}'_t(\mathbf{q}_s^A - \mathbf{q}_t^A) + \mathbf{P}_t^{A'}(\mathbf{Q}_s - \mathbf{Q}_t)], \\ u_s^B - u_t^B &\leq \frac{1}{\beta_B^{t-1}} [\mathbf{p}'_t(\mathbf{q}_s^B - \mathbf{q}_t^B) + \mathbf{P}_t^{B'}(\mathbf{Q}_s - \mathbf{Q}_t)]. \end{aligned}$$

In words, if we can find a discount factor β_m and constants $\{u_t^m\}_{t \in T}$ for each household member $m \in \{A, B\}$, along with feasible private quantities and Lindahl prices, such that the inequalities defined by Theorem 2 hold, then there exists a pair of felicity functions and constant discount rates that provide a perfect within-sample rationalisation of the household data. Conversely, if we cannot find values of all

⁶One important difference between our theoretical framework and that of Mazzocco’s (2007) “full efficiency” model is the assumption of perfect foresight. Revealed preference tests of martingale processes lack content as, without a specification of the expectation process, one can always posit an unexpected shock to rationalise behaviour. Following Mazzocco, our framework also assumes perfect capital markets. We recognise that these assumptions are very strong; see also our Section 3.3. Still, in our empirical application we will find that nearly all observed household behavior in our sample can be rationalised even when maintaining these assumptions. See Section 5 for more discussion.

relevant variables such that these inequalities hold, then there does not exist a theory-consistent specification of household member preferences and a constant Pareto weight that rationalise the observed consumption stream. This allows us to test whether time inconsistencies in choice can be explained by appeal to discount rate heterogeneity alone. If a non-empty feasible set is associated with the full efficiency constraints, one cannot reject the hypothesis that time inconsistency in household choice is simply the product of individual heterogeneity within the collective unit. One does not necessarily require nonstationarities in individual preferences or renegotiations of the household choice rule over time to rationalise the observed behaviour.

4.2 Renegotiation

If household behaviour is inconsistent with the full efficiency model, an appeal to more than just discount rate heterogeneity is required. The second condition for time consistency of the collective preference is the existence of a constant Pareto weight across the full lifetime of the household. Whether this is necessarily attained depends upon the existence of a perfect intrahousehold commitment mechanism. Without a perfect commitment device, the Pareto weight can vary over time to reflect renegotiations of the household choice rule. These renegotiations thus open up an additional mechanism for time inconsistent behaviour.

Mazzocco’s (2007) “no-commitment” model weakens the assumption of perfect intrahousehold commitment. Instead, the household solves the lifetime bargaining problem subject to additional incentive compatibility constraints. Mazzocco (2007) classes a consumption stream as incentive compatible if it does not provide an incentive for any family member to quit the household at some point to take their “outside option”. An individual’s outside option is defined as the utility they could derive from divorcing and continuing in the world alone.⁷

Using the method developed by Marcat and Marimon (1998), the no-commitment model can be formulated as a recursive saddle point problem and theory consistent behaviour corresponds to the following rationalisation condition.⁸

Definition 4 *The set of observations $S = \{\mathbf{q}_t, \mathbf{Q}_t; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T}$ can be rationalised by the no-commitment model if there exist, for all $t \in T$, private quantities $\mathbf{q}_t^A, \mathbf{q}_t^B \in \mathbb{R}_+^N$ (with $\mathbf{q}_t^A + \mathbf{q}_t^B = \mathbf{q}_t$) and, in addition, concave, strictly increasing felicity functions u^A and u^B , discount factors $\beta_A, \beta_B \in (0, 1]$, Pareto weights $\omega_t^A, \omega_t^B > 0$, multipliers φ_t^A and φ_t^B and outside utilities \underline{u}_t^A and \underline{u}_t^B such that*

$$\sum_{t \in T} \left\{ \left[\omega_t^A \beta_A^{t-1} u^A(\mathbf{q}_t^A, \mathbf{Q}_t) - \varphi_t^A \underline{u}_t^A \right] + \left[\omega_t^B \beta_B^{t-1} u^B(\mathbf{q}_t^B, \mathbf{Q}_t) - \varphi_t^B \underline{u}_t^B \right] \right\} \geq \sum_{t \in T} \left\{ \left[\omega_t^A \beta_A^{t-1} u^A(\zeta_t^A, \zeta_t^H) - \varphi_t^A \underline{u}_t^A \right] + \left[\omega_t^B \beta_B^{t-1} u^B(\zeta_t^A, \zeta_t^H) - \varphi_t^B \underline{u}_t^B \right] \right\}$$

for all affordable consumption plans $\{\zeta_t^A, \zeta_t^B, \zeta_t^H\}_{t \in T}$ (with $\zeta_t^A, \zeta_t^B \in \mathbb{R}_+^N$ and $\zeta_t^H \in \mathbb{R}_+^K$), which satisfy

$$\sum_{t \in T} [\mathbf{p}'_t(\mathbf{q}_t^A + \mathbf{q}_t^B) + \mathbf{P}'_t \mathbf{Q}_t] \geq \sum_{t \in T} [\mathbf{p}'_t(\zeta_t^A + \zeta_t^B) + \mathbf{P}'_t \zeta_t^H].$$

In this definition, φ_t^m represents the Lagrange multiplier on m ’s incentive compatibility constraint, and \underline{u}_t^m gives the utility associated with m ’s outside option, that is, the stream of utility that m could receive when leaving the household at period t . The household choice rule, summarised by the set of intrahousehold preference weights, ω_t^m , is sequentially renegotiated to reflect changes in the slackness of the incentive compatibility constraints: $\omega_1^A = 1$, $\omega_t^A = \omega_{t-1}^A + \varphi_t^A$ and $\omega_1^B = \omega$, $\omega_t^B = \omega_{t-1}^B + \varphi_t^B$. Assuming positive gains to marriage continuation for at least one spouse in every time period, there will always be at least

⁷This definition is akin to Shaked and Sutton’s (1984) formulation of outside options.

⁸This formulation implicitly embodies the requirement that a particular consumption profile must be incentive compatible. i.e. for each member $m = \{A, B\}$

$$\sum_{s=0}^{|T|-t} \beta_m^s u^m(\mathbf{q}_{t+s}^m, \mathbf{Q}_{t+s}) \geq \underline{u}_t^m.$$

In words, by remaining within the household each household member must achieve welfare at least as great as when exiting via divorce.

one individual who is strictly better off if the marriage continues rather than dissolving through divorce. Thus, the incentive compatibility condition can only bind for one family member at any point in time, i.e. if $\varphi_t^A \neq 0$ then $\varphi_t^B = 0$, and vice versa. In periods where the incentive compatibility constraint binds for some member, the weight assigned to her preferences is increased until she is indifferent between taking their outside option and staying within the household. This new weighting of family member preferences then prevails in subsequent time periods until an incentive constraint again binds and another reweighting of preferences is implemented.

Stationarity of the household per-period felicity function can now also be undermined by renegotiations of the household choice rule, which take place whenever an incentive compatibility constraint binds. That allocations are sensitive to outside option heterogeneity is clear from the family maximisation problem; u_t^m appears explicitly in the household objective function. However, the interaction between discount rate heterogeneity and incentive compatibility is more subtle. Consider a couple who are identical in every respect except for their patience, $\beta_A < \beta_B$. The optimal lifecycle plan would see a more present-weighted consumption profile for A than B . Thus in early periods, A would receive a greater relative share of per-period expenditure, and the opposite in later periods. However, without a commitment mechanism, this plan may be infeasible. In some period, as her per-period resource share drops, A could conceivably do better by quitting the household, especially given the low weight she attaches to future marital surpluses. The Pareto weight will then be renegotiated to re-emphasise A 's preferences in the household allocation problem to prevent her from dissolving the household.

Revealed preference conditions. The no-commitment model implies the existence of a set of mutually exclusive subsets within there is no renegotiation and thus, the same Pareto weight is applied.⁹ To introduce the potential for renegotiation into our revealed preference set-up, consider a partition of the set T into Υ mutually exclusive subsets T_τ of the form

$$\mathbf{T} = \{T_1, \dots, T_\Upsilon\},$$

with

$$\mathbf{T} = \cup_{\tau=1}^{\Upsilon} T_\tau \quad \text{and} \quad T_{\tau_s} \cap T_{\tau_t} = \emptyset \quad \text{if} \quad \tau_s \neq \tau_t,$$

such that

$$\tau_1 < \tau_2 \text{ implies } t_1 < t_2 \text{ for all } t_1 \in T_{\tau_1} \text{ and } t_2 \in T_{\tau_2}.$$

Each subset represents a distinct 'Pareto weight regime', thus $\omega_s^m = \omega_t^m$ for all $s, t \in T_\tau$. Let the Pareto weight in subset T_τ thus be denoted ω_τ^m . We then have the following testability result.

Theorem 3 *For a partition \mathbf{T} , the set of observations $S = \{\mathbf{q}_t, \mathbf{Q}_t; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T}$ can be rationalised by the no-commitment model only if there exist, for all $t \in T$, private quantities $\mathbf{q}_t^A, \mathbf{q}_t^B \in \mathbb{R}_+^N$ (with $\mathbf{q}_t^A + \mathbf{q}_t^B = \mathbf{q}_t$), Lindahl prices $\mathbf{P}_t^A, \mathbf{P}_t^B \in \mathbb{R}_+^N$ (with $\mathbf{P}_t^A + \mathbf{P}_t^B = \mathbf{P}_t$), utility numbers $u_t^A, u_t^B \in \mathbb{R}$ and constants $\beta_A, \beta_B \in (0, 1]$ that, for any $s, t \in T_\tau$ ($\tau \in \{1, \dots, \Upsilon\}$), satisfy*

$$\begin{aligned} u_s^A - u_t^A &\leq \frac{1}{\beta_A^{t-1}} [\mathbf{p}'_t(\mathbf{q}_s^A - \mathbf{q}_t^A) + \mathbf{P}_t^{A'}(\mathbf{Q}_s - \mathbf{Q}_t)], \\ u_s^B - u_t^B &\leq \frac{1}{\beta_B^{t-1}} [\mathbf{p}'_t(\mathbf{q}_s^B - \mathbf{q}_t^B) + \mathbf{P}_t^{B'}(\mathbf{Q}_s - \mathbf{Q}_t)]. \end{aligned}$$

The interpretation is similar to before. In this particular case, innovations in the Pareto weight define the partitions of the set T . Thus, within sub-periods T_τ , the Pareto weight is constant and choices must satisfy the revealed preference inequalities associated with the full efficiency model (in Theorem 2).

⁹For rich enough data sets we can define these subsets by using information on outside options. In our own application (in Section 5), however, we do not have such information and, therefore, we consider all possible partitions \mathbf{T} .

5 Rationalising observed time inconsistency

In this section, we resume our empirical application to the ECPF data. We find that simply allowing for limited intrahousehold heterogeneity in the discount rate allows the behaviour of 98.4% of families to be rationalised without recourse to hyperbolic discounting at the individual level. Given this positive result, we conduct a detailed investigation of the theory-consistent differences in spousal discount rates.

Furthermore, although the vast majority of household behaviour can be explained without any mention of intrahousehold renegotiation, we provide results for a strengthened version of the conditions in Theorem 3, which assumes equal discount factors for the individual household members A and B , $\beta_A = \beta_B$. We consider the success of this strengthened version of Theorem 3 to allow for an assessment of the relative importance of time preference heterogeneity and renegotiation in generating observed time inconsistency. Our results suggest that, for our short panel, discount rate heterogeneity is the more relevant channel for explaining patterns of household choice.

Similar to Section 3, we conduct “yes/no” tests for the models that account for individual heterogeneity. Thus, the interpretation of the test results reported below is directly analogous to before: the pass rate gives the proportion of households that can be rationalised by a given a model. Our testing procedure is slightly modified from that of Section 3 to account for the collective nature of choice. We test the conditions in Theorems 2 and 3 using a two-dimensional grid search over individual discount factors on $[0.9, 1]^2$, with a spacing of 0.005. Further, to test the strengthened version of Theorem 3, we consider alternative scenarios defined by the maximum number of renegotiations that are permissible in the two-year period that a household is observed: this maximum can range from 0 (i.e. time consistent behaviour) to 7 (i.e. the Pareto weight changes in each different consumption quarter).

5.1 Test results

Table 5 presents our results. Let us first consider the full-efficiency collective model, which allows for β -heterogeneity within the household but imposes a single choice rule for the period of consideration. The results associated with the full efficiency model stand in stark contrast to those reported in Section 3 for the time consistent model. We can explain the behaviour of the overwhelming majority of households using the framework of this extremely simple intertemporal collective model, under which the only source of time inconsistency in household revealed preferences is variation in the time preferences of family members. No further recourse to nonstationarities at the individual level is required. We find this result surprising given the strong assumptions of perfect foresight, perfect commitment and perfect capital markets that the theoretical framework incorporates. However, we cannot reject the hypothesis that these assumptions are valid in the short term, or at least for two years. For 98.4% of households we are able to find a well-behaved felicity function and a constant discount rate for each family member that provide a perfect within-sample rationalisation of their choice behaviour.¹⁰

We next consider the pass rate associated with the strengthened version of Theorem 3 that admits renegotiation but assumes β -homogeneity. Unsurprisingly, allowing for more frequent renegotiations of the Pareto weight is associated with an increase in the pass rate.¹¹ For the extreme scenario that allows innovations in the Pareto weight between any two consecutive periods (quarters in our case) the pass rate amounts to 92.75%. However, as soon as we require a stable Pareto weight over two periods or more, the pass rate drops quite dramatically.

Table 5: Individual heterogeneity

<i>Full efficiency and β-heterogeneity</i>								
<i>Pass rate (%)</i>	98.36							
<i>Renegotiation and β-homogeneity</i>								
<i>Renegotiations (max.)</i>	0	1	2	3	4	5	6	7
<i>Pass rate (%)</i>	0.13	1.08	7.07	22.34	50.29	74.83	89.28	92.75

¹⁰Here it is worth to remark that Mazzocco (2007) actually rejected the full efficiency model in his empirical application. Our findings suggest that this rejection could be the result of biases introduced by misspecification, omitted relevant distribution factors or the synthetic nature of the panel used, rather than a failure of commitment itself.

¹¹This is unsurprising since increasing the number of (possible) renegotiations generally obtains less stringent rationalisability conditions.

We conclude that accounting for the collective nature of household choice allows the intertemporal behaviour of families in our sample to be explained using simple models that assume constant discounting at the individual level. For the given data set, our results provide particularly strong empirical support for a model which locates the primary source of time inconsistent family behaviour with intrahousehold β -heterogeneity. This full-efficiency model seems plausible given the short time span of our sample.

5.2 Discriminatory power and predictive success

Thus far, we have only considered pass rates of the various intertemporal models presented in this paper. However, alongside the pass rate, empirical applications of revealed preference analysis often consider two additional performance metrics: discriminatory power and predictive success. In this section, we compute these metrics for the three behavioural models under study, to complement the pass rates already reported.¹² As we will explain, predictive success gives a holistic measure of the empirical performance of a behavioural model by simultaneously accounting for the pass rate and discriminatory power. As such, it is particularly interesting to compare models by using this metric. We find that this comparison provides additional empirical support for the full efficiency model as framework for explaining the choice behaviour of households in our sample.

Discriminatory power. Following Bronars (1987), we define the discriminatory power of a revealed preference test for a particular behavioural model as the probability of detecting behaviour that is not rationalisable by the model.

Bronars suggests an iterative procedure to compute his power metric, which we apply to each household in our sample. At every iteration, the procedure simulates random behaviour (i.e. behaviour that is not generated by any optimising model) by drawing $|T| \times (N + K)$ random budget shares from the uniform distribution. For a given household, these budget shares then define a new random consumption stream $\{\mathbf{q}_t^R, \mathbf{Q}_t^R\}_{t \in T}$ that exhausts total wealth.¹³ We then test the revealed preference conditions of the model under evaluation on the correspondingly defined set $\{\mathbf{q}_t^R, \mathbf{Q}_t^R; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T}$. In our application, we iterate this procedure 1000 times and calculate the proportion of the randomly generated consumption streams that fail the revealed preference restrictions of the behavioural model.

We use the proportion of randomly generated consumption streams that fail the revealed preference restrictions to calculate a household-specific measure for discriminatory power. This proportion proxies the true probability that random household behaviour will fail the restrictions of the behavioural model for observed prices and total household expenditure. For example, if 50% of all randomly generated consumption streams fail to meet the requirements of a revealed preference test, then there is approximately a 50% chance that our tests will correctly reject random choice behaviour. Generally, high power signals a restrictive model and there is a high probability that revealed preference tests will detect irrational/random behaviour.

At this point, recall that we conducted a grid search on the discount factor (β for the time consistent model and (β_A, β_B) for the collective models) to compute pass rates for the different models. Computing power requires an analogous grid search at each iteration. We define a household specific grid size, conditional on whether the observed household choices are rationalisable by the model under study. For the collective models, if observed behaviour is not rationalisable, then at each iteration we define the same grid size as before: the interval $[0.9, 1]^2$ with a spacing of 0.005. However, if observed behaviour *can* be rationalised by the model in question, we use this information to define a finer grid. Specifically, we search only over (β_A, β_B) for which the difference $(\beta_A - \beta_B)$ is not greater than the minimum difference under which the observed behaviour is rationalisable.¹⁴ This adjusted grid search substantially limits the computational burden of our power assessment, while accounting for the information on individual time preferences as revealed by the observed behaviour. We proceed similarly for the time consistent model. For nonrationalisable behaviour, we consider β in the interval $[0.9, 1]$, with a spacing of 0.005, and, for nonrationalisable behaviour, our finer grid contains all β no lower than the maximum value under which rationalisability is obtained.

¹²For the collective model that allows for renegotiation, the following results all pertain to the specification that allows innovations in the Pareto weight between any two consecutive periods, of which the pass rate amounts to 92.75% (see Table 5).

¹³This simulated random behaviour corresponds to Becker’s (1962) notion of irrational behaviour as behaviour that randomly exhausts the available budget.

¹⁴In Section 5.3 we explain our procedure to recover this minimum difference $(\beta_A - \beta_B)$ that is consistent with rationalisability.

Results. Figure 2 shows the kernel distribution of discriminatory power for the three models under study, and Table 6 gives corresponding summary statistics. Some notable features emerge. First, the time consistent model has very high discriminatory power. For over half of the households, discriminatory power equals 100%, and the average power is no less than 99.97%. This is not too surprising given the extremely low pass rate (0.13%) that we obtained before. In contrast, the collective model that allows for renegotiation within the household but assumes homogenous discount rates has generally low power: the average value is only 8.99%, and the maximum value amounts to no more than 13.30%. Finally, the model that admits intrahousehold β -heterogeneity, while assuming a constant Pareto weight, takes an intermediate position: average power is 47.73%, which is high for revealed preference tests of the collective model (see Cherchye, De Rock and Vermuelen, 2009). There is also significant heterogeneity in discriminatory power across households: minimum power equals 0% and maximum power, 100%. The full efficiency model performs relatively well in terms of discriminatory power.

Figure 2: Kernel plot of power distribution

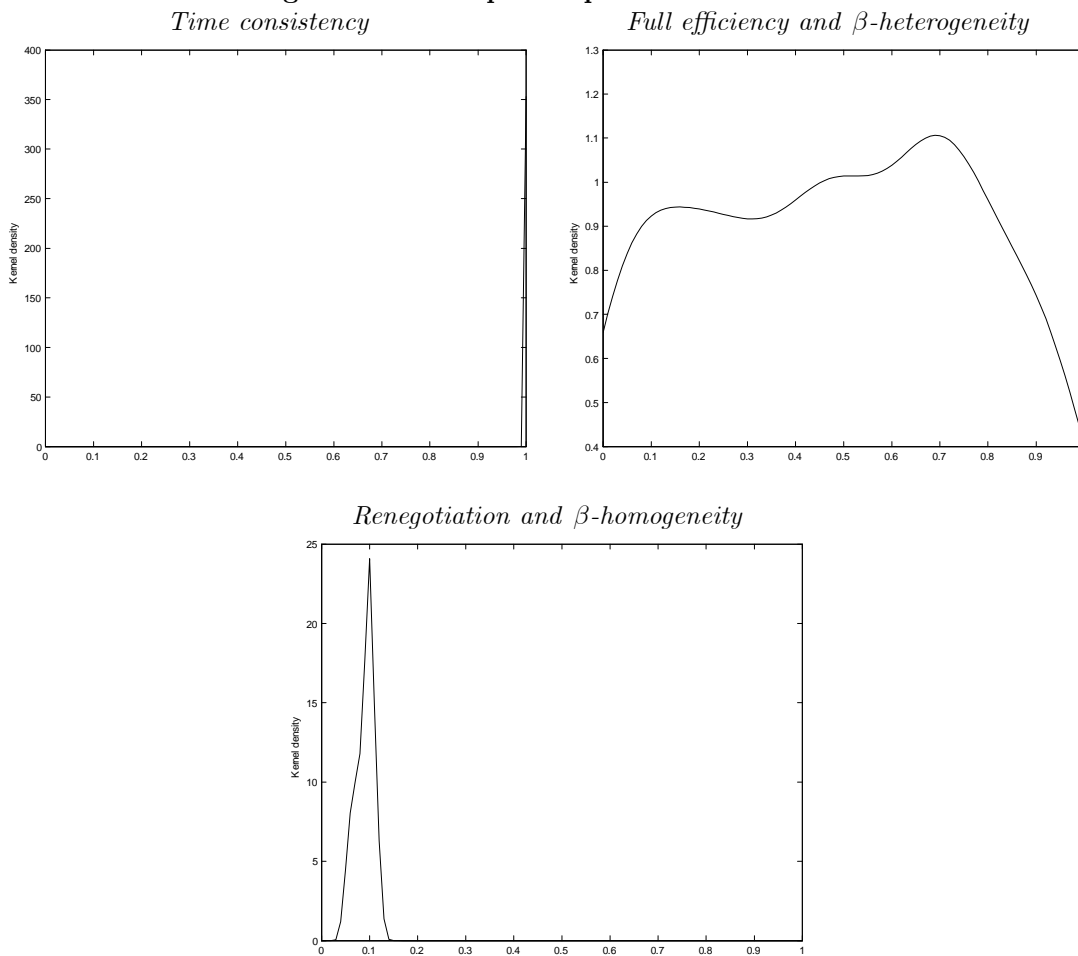


Table 6: Summary statistics of power distribution (in %)

	<i>Mean</i>	<i>Min</i>	<i>Q1</i>	<i>Median</i>	<i>Q3</i>	<i>Max</i>
<i>Time consistency</i>	99.97	99.60	99.90	100.00	100.00	100.00
<i>Full efficiency and β-heterogeneity</i>	47.73	0.00	22.00	48.70	72.02	100.00
<i>Renegotiation and β-homogeneity</i>	8.99	3.80	7.60	9.40	10.40	13.30

Predictive success. We also compute a predictive success metric to compare the different models that we study. This metric was recently axiomatised by Beatty and Crawford (2011) and is based on an original proposal of Selten (1991). It combines the pass rate and power of a particular behavioural model into a single metric: for each household, it subtracts 1 minus the power measure from the pass measure (1 or 0). As such, the predictive success measure can be interpreted as a power-adjusted pass rate. The measure is always situated between -1 and 1.

The higher the average predictive success measure, the better the empirical performance of the behavioural model under evaluation. A predictive success value in the neighbourhood of -1 pertains to a household that fails the rationalisability conditions, implying a pass measure equal to 0, even though the power of the test is low and thus relatively easy to pass (i.e. discriminatory power is close to 0). Conversely, a predictive success value in the neighbourhood of 1 indicates a household that passes the model restrictions in a situation where the model has high power. This represents the ‘ideal’ scenario if you will. Finally, a predictive success value equal to zero suggests that the model is not informative for the household at hand: the model does not outperform the uninformative assumption that households exhibit random consumption behaviour, for which the power is 0 and the pass measure equals 1, by construction.

Figure 3 and Table 7 summarise the predictive success results for the three models under study. These results suggest that both the time consistent model and the no commitment model with homogenous discount factors are uninformative for our sample of households. For many households, they cannot distinguish between “rationalisable” and “random” behaviour, although for different reasons. The very stringent time consistency model rejects almost *any* behaviour, which effectively makes it a neutral predictor. By contrast, the no commitment model with homogenous discount factors is not very restrictive for our Spanish data set and thus suffers from low power (see Figure 2). As a result, although the model achieves a high pass rate, it barely outperforms the neutral time consistent model in terms of predictive success.

Both these models are outperformed by the full efficiency model with heterogeneous discount factors, as is readily apparent from Figure 3 and Table 7. We can even establish a (partial) stochastic dominance result in terms of predictive success. If we disregard the small set of households that are not rationalizable by the full efficiency model, which amounts to only 1.6% of the sample, the whole predictive success distribution pertaining to the remaining sample is situated to the right of the distributions pertaining to the other two models. These results suggest that the full efficiency model is a particularly useful one for describing the short term dynamics of the consumption behaviour of the households studied here.

Figure 3: Kernel plot of predictive success distribution

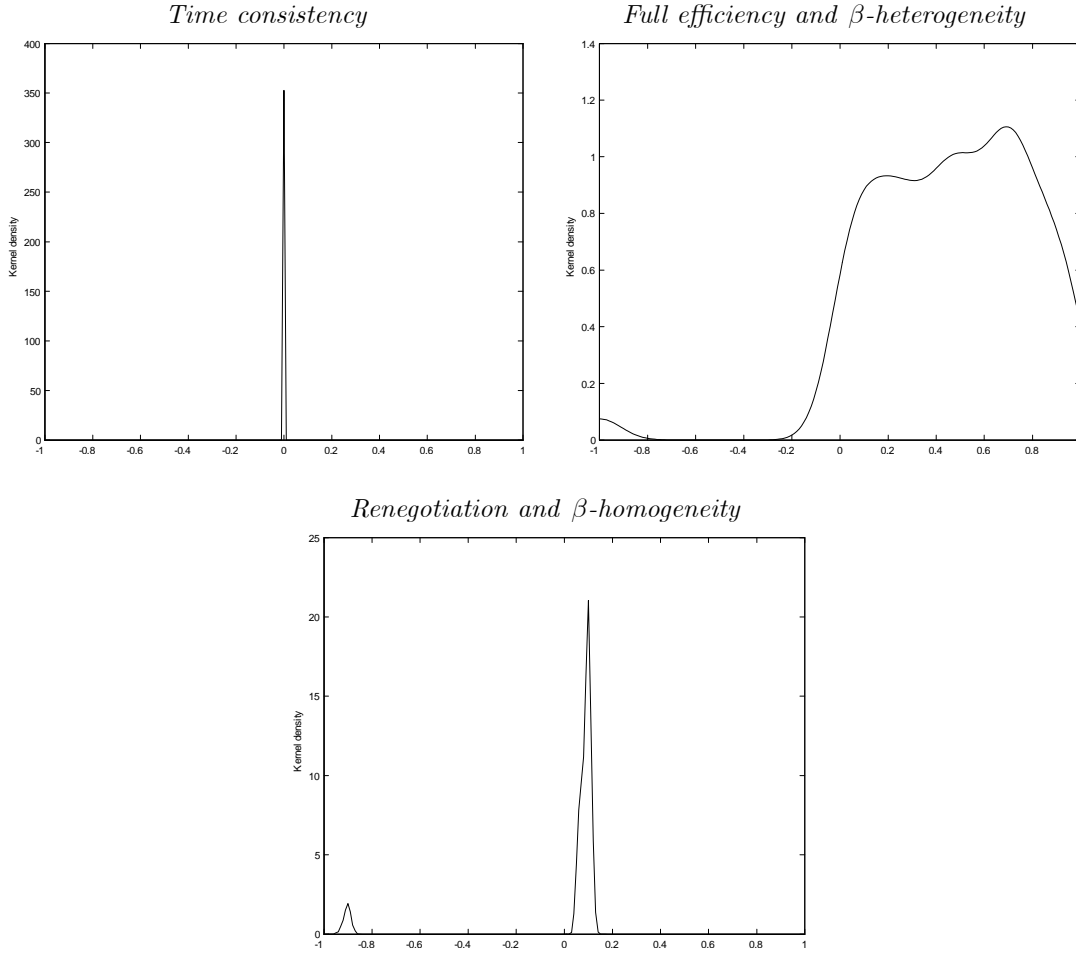


Table 7: Summary statistics of predictive success distribution (in %)

	<i>Mean</i>	<i>Min</i>	<i>Q1</i>	<i>Median</i>	<i>Q3</i>	<i>Max</i>
<i>Time consistency</i>	0.09	-0.40	-0.10	0.00	0.00	100
<i>Full efficiency and β-heterogeneity</i>	46.09	-100.00	22.00	48.70	72.02	100
<i>Renegotiation and β-homogeneity</i>	1.73	-94.80	7.00	9.10	10.30	13.30

5.3 Time preference recovery

The analysis above suggests that the full efficiency model, allowing for time preference heterogeneity, performs well for the data at hand. Given this implied importance of discount rate heterogeneity in accounting for household consumption behaviour, we now explore the nature of the theory-consistent set of household discount rates, and consider whether the necessary degree of unobservable preference heterogeneity correlates with observable family characteristics.

Minimum heterogeneity. For each household that can be rationalised by the full efficiency model, we recover the discount rates that make observed consumption behaviour consistent with the rationalisability conditions in Theorem 2. One drawback of our revealed preference methodology is that the identification of these discount rates is necessarily weakened by the lack of structure our framework imposes on individual preferences and the household choice problem. Our recovery problem is thus underdetermined and preferences

are only set identified in the sense of Manski (2007). We refer to this set of potential time preferences as the set of "theory consistent discount rates".

We deal with this non-uniqueness by reporting results for the *minimal* amount of discount rate heterogeneity necessary to rationalise the household consumption stream. To determine the minimum difference ($\beta_A - \beta_B$) in the discount rate grid, we iterate the testing procedure by considering the allowable grid points in a specific order. We initially set $\beta_A - \beta_B = 0$, which corresponds to the time consistent model, and thus is rejected for all but two households. For the remaining households, we ask if behaviour can be rationalised for $\beta_A - \beta_B = 0.005$, by testing for a non-empty set to the inequalities defined by Theorem 2 with $\beta_A \in [1, 0.995, \dots, 0.905]$ and $\beta_B = \beta_A - 0.005$. For the remaining non-rationalisable households, the difference is then set at 0.01, and so on until the maximum difference is reached when $\beta_A = 1$ and $\beta_B = 0.90$.

We are able to assume that $\beta_A \geq \beta_B$ without loss of generality because the data set does not contain assignable goods, whose consumption can be tied to a particular household member. Thus, the Afriat inequalities for $m = \{A, B\}$ are fully symmetric. Therefore, we make no assumptions on, and cannot test for, whether the husband or wife is the more patient household member. If assignable information were available, analysis using our methodology could be easily extended to address gender-related questions.


The distribution of the minimum heterogeneity level over all 1559 rationalisable households is summarised in Table 8, and shown with more detail in Figure 4.

Table 8: β -heterogeneity

	<i>Mean</i>	<i>Min</i>	<i>Q1</i>	<i>Median</i>	<i>Q3</i>	<i>Max</i>
<i>Min</i> ($\beta_A^* - \beta_B^*$)	0.0199	0	0.0100	0.0150	0.0250	0.0950

Only limited heterogeneity is required to rationalise the behaviour of most households. Almost three quarters of the households that are consistent with Theorem 2 only require $\beta_A - \beta_B \leq 0.025$ to rationalise their behaviour. The kernel plot in Figure 4 confirms that the distribution of minimal heterogeneity shows a high and relatively narrow peak for limited amounts of within-household heterogeneity in time preferences. We also note that the maximum allowable heterogeneity level of 0.10 is not observed for any of the rationalisable households.

Figure 4: β -heterogeneity



Appeal to observed heterogeneity. We investigate whether patterns in minimum unobservable time preference heterogeneity are related to observable heterogeneity via a simple regression of the log of minimum discount rate distances on recorded household characteristics. The ECPF contains information on a number of household characteristics including: age of each spouse, schooling and occupation of the head, the number of children in the household and housing tenure.¹⁵ We also consider total consumption expenditures and the

¹⁵We created dummy variables for housing tenure and job position. The variable "high level job" equals 1 if the household head is wage earner with a university degree, and 0 if he is a specialized worker without a university degree, or an unspecialized worker.

average budget share of public goods as proxies for wealth and marital surpluses respectively.

Table 7 shows the regression results. We find that intrahousehold heterogeneity in discount factors correlates with many household characteristics and other observables: average spousal age, the presence of children and the level and variance of household expenditures.

First let us consider possible rationalisations of the inverse relationship between average spousal age and time preference heterogeneity. We hypothesise that a positive association between age and match quality generates this correlation. Match quality is closely aligned to similarity in time preferences in the literature. For example, Schaner (2012) classes couples as either “well” or “badly” matched purely on the basis of differences in their elicited discount rates. We see two alternative explanations for the positive association between the average age of a couple and match quality. First, the older people are, the more likely it is for them to have met and married someone similar to them. Second, only well-matched couples stay together in the long term. Before considering these explanations in more detail, we must first comment on the place of divorce in our framework. Our empirical test assumes an invariant intrahousehold decision making rule for the two years a family is observed. However, this is not to say that renegotiation and divorce cannot occur in the long run. Clearly, some couples do divorce in reality but the probability of divorce is declining in match quality (Becker *et al* 1977; Weiss and Willis 1997). The higher the quality of a match, the more marital surplus is available to be shared by a couple and the less likely it is for marriage dissolution to dominate for either spouse in any time period.

Now, on the first proffered explanation for the association between match quality and average age, consider the search process leading to marriage. If this process is costly, individuals will accept imperfect matches even if it is known that differences in time preferences will create inefficiencies in the new household (see Burdett and Coles (1999) for a formal framework). Over time, if some search continues during marriage, individuals will continue to acquire new information on other potential matches. If a spouse meets a high enough quality match, it can be worth dissolving an existing marriage to take up a new opportunity. The older someone is, the longer they will have been a participant in the marriage market and therefore, the more likely it is that they will have met someone similar to themselves and thus be a member of a ‘high quality’ match. In this way, costly search can create a negative association between age and discount rate heterogeneity.

Alternatively, on the second explanation, it is not unreasonable to assume that older couples have been married longer. Some model match quality as an experience good.(Nelson 1970; Jovanovic 1979; Weiss and Willis 1997; Chiappori and Weiss 2006). In contrast to the above, these explanations assume that one cannot perfectly assess the suitability of a potential mate until marriage occurs and a match is experienced. Given that young couples have lower marriage experience, one can expect a greater variation in match quality (i.e. heterogeneity in discount rates) amongst this group, as these couples are still assessing the degree to which they are suited and not all poor matches will have been terminated. However, only well suited couples will remain married for a long period of time, again creating a negative association between the age of a couple and discount rate heterogeneity.

The presence of children and higher total expenditure are also associated with smaller differences in time preferences. These facts can also be rationalised by appeals to match quality. Children raise the cost of divorce and thus one would not want to bring children into a household unless one were sure that they are part of a high quality match. On the relationship between total expenditure and heterogeneity, more wealthy individuals are able to bear search costs for longer before entering into a match, which again raises the average quality of match that they can expect to make.

Another variable of interest is the variation in household spending. We find that a higher variation is generally associated with more time preference heterogeneity. We do not directly see a behavioural interpretation of this phenomenon. Actually this finding may simply be an artefact of our focus on the minimum β -heterogeneity needed for rationalisability in terms of the full efficiency model. Because this model excludes that (large) income changes cause renegotiation within the household, we may generally need a higher β -difference to obtain consistency with the conditions in Theorem 2 for households facing substantial income variation.

As a final note, we remark that our results suggest three further effects that are significant: university degree, home ownership and high level job all seem to bear a positive relation to intrahousehold β -heterogeneity. We could not directly see a “collective” interpretation of these effects. Follow-up research, potentially using richer data sets, may focus on better explaining these patterns. For example, as indicated above, data sets

that include information on assignable goods can allow for identifying gender-specific effects, which may be relevant here.

Table 9: β -heterogeneity and household characteristics

	<i>coefficient</i>	<i>standard error</i>
<i>Mean age</i>	-0.00725***	0.00213
<i>Age difference</i>	0.00153	0.00589
<i>University degree</i>	0.146*	0.0704
<i>High school degree</i>	-0.0542	0.0379
<i>Children (dummy)</i>	-0.111*	0.0432
<i>Log expenditures (logexp)</i>	-0.322***	0.0417
<i>Var(logexp)</i>	2.01***	0.240
<i>Pct. public cons.</i>	0.0124	0.159
<i>Home owner</i>	0.103*	0.0415
<i>High level job</i>	0.184**	0.0669
<i>R-squared</i>	0.0925	
<i>No. of observations</i>	1557	

* = significant at 5%-level, ** = significant at 1%-level, *** = significant at 0.001%-level.

6 Conclusion

We have provided a revealed preference analysis of time inconsistencies in household consumption. Adopting a collective perspective, we focused on rationalising these inconsistencies as the product of nonstationarities at the individual level or the result of individual heterogeneity and renegotiation within the collective unit. An empirical application to a Spanish consumption panel highlights that an explicit recognition of the collective nature of choice allows us to rationalise the vast majority of time inconsistent household behaviour. Almost all observed household behaviour turns out to be consistent with a simple model that assumes perfect intrahousehold commitment (i.e. no renegotiation) in the two-years period under consideration together with exponential discounting at the individual level, but allowing for heterogeneity in time preferences within the household.

We have also shown that revealed preference restrictions can be fruitfully applied to recover individual time preferences. One can then relate these results to specific individual or household characteristics. For example, for our application we found that the intrahousehold heterogeneity in discount factors correlates with average spousal age, the presence of children and the level and variance of household expenditures. This application demonstrates the potential usefulness of revealed preference methodology to address this type of questions in an effective manner. It is our belief that richer household data sets (e.g. including assignable goods and more information on observable characteristics) may yield additional and more refined insights.

References

- [1] Sydney Afriat (1967), The Construction of Utility Functions from Expenditure Data, *International Economic Review*, 8.1 p67-77.
- [2] George Ainslie, Shane Frederick, Xiupeng Li, John Lynch, Page Moreau, Andrew Mitchell, Daniel Read, Alan Sawyer, Yaacov Trope, Klaus Wertenbroch, Gal Zauberman (2005), The Psychology of Intertemporal Discounting: Why are Distant Events Valued Differently from Proximal Ones?, *Marketing Letters*, 16.3.4 p347-360.
- [3] George-Mario Angeletos, David Laibson, Andrea Repetto, Jeremy Tobacman and Stephen Weinberg (2001), The Hyperbolic Consumption Model: Calibration, Simulation and Empirical Evaluation, *Journal of Economic Perspectives*, 15.3 p47-68.
- [4] James Andreoni and Charles Sprenger (2012), Estimating Time Preferences from Convex Budget Sets, *American Economic Review*, Forthcoming 2012.

- [5] Timothy K. Beatty and Ian Crawford (2011), How Demanding is the Revealed Preference Approach to Demand?, *American Economic Review*, 101, p2782-2795.
- [6] Gary Becker (1962), Irrational Behavior and Economic Theory, *Journal of Political Economy*, 70, p1-13.
- [7] Gary Becker, Elizabeth Landes, Robert Michael (1977), An Economic Analysis of Marital Instability, *Journal of Political Economy*, 85, p1141-1187.
- [8] Stephen G. Bronars (1987), The Power of Nonparametric Tests of Preference Maximization, *Econometrica*, 55, p693-698.
- [9] Martin Browning (1989), A Nonparametric Test of the Life-Cycle Rational Expectations Hypothesis, *International Economic Review*, 30.4 p979-992.
- [10] Martin Browning, Francois Bourguignon, Pierre-Andre Chiappori, Valerie Lechene (1994). Income and Outcomes: A Structural Model of Intrahousehold Allocation, *Journal of Political Economy*, 102, p1067-96.
- [11] Martin Browning and Pierre Andre Chiappori (1998), Efficient Intra-household Allocations: A General Characterisation and Empirical Tests, *Econometrica*, 66.6 p1241-1278.
- [12] Martin Browning and M. Dolores Collado (2001), The Response of Expenditures to Anticipated Income Changes: Panel Data Estimates. *American Economic Review*, 91, 681-692.
- [13] Kenneth Burdett and Melvyn Coles (1999), Long Term Partnership Formation: Marriage and Employment, *The Economic Journal*, 109, p307-334.
- [14] Christopher Carroll (2001), Death to the Log-Linearized Consumption Euler Equation (And Very Poor Health to the Second-Order Approximation), *The B.E. Journal of Macroeconomics*, 1.1 p6.
- [15] Laurens Cherchye, Bram De Rock and Frederic Vermeulen (2007), The Collective Model of Household Consumption: A Nonparametric Characterization, *Econometrica*, 75.2 p553-574.
- [16] Laurens Cherchye, Bram De Rock and Frederic Vermeulen (2009), Opening the Black Box of Intra-household Decision-making: Theory and Nonparametric Empirical Tests of General Collective Consumption Models, *Journal of Political Economy*, 117, p1074-1104.
- [17] Pierre-André Chiappori (1988), Rational Household Labor Supply, *Econometrica*, 56, p63-89.
- [18] Pierre-André Chiappori (1992), Collective Labor Supply and Welfare, *Journal of Political Economy*, 100, p437-67.
- [19] Pierre-Andre Chiappori and Yoram Weiss (2006), Divorce, Remarriage and Welfare: A General Equilibrium Approach, *Journal of the European Economic Association* 4.2/3, Papers and Proceedings of the Twentieth Annual Congress of the European Economic Association p415-426.
- [20] James Choi, David Laibson, Brigitte Madrian and Andrew Metrick (2001), For Better or for Worse: Default Effects and 401(k) Savings Behavior, In *Perspectives on the Economics of Ageing*, edited by David Wise, University of Chicago Press for NBER.
- [21] Ian Crawford (2010), Habits Revealed, *Review of Economic Studies*, 77.4 p1382-1402.
- [22] Stefano DellaVigna and Ulrike Malmendier (2006), Paying Not To Go To The Gym, *The American Economic Review*, 96, p694-719.
- [23] Thomas Demuynek and Ewout Verriest (2012), I'll Never Forget My First Cigarette: A Revealed Preference Analysis of the 'Habits as Durables' Model, *International Economic Review*, Forthcoming.
- [24] Thomas Dohman, Armin Falk, David Huffman, Uwe Sunde (2012), Interpreting Time Horizon Effects in Inter-Temporal Choice, CESifo Working Paper No. 3750.

- [25] W. Erwin Diewert (1973), Afriat and Revealed Preference Theory, *The Review of Economic Studies*, 40.3 p419-425.
- [26] Martin S. Feldstein (1964), The Social Time Preference Discount Rate in Cost Benefit Analysis, *The Economic Journal*, 74, p360-379.
- [27] Shane Frederick, George Loewenstein, Ted O'Donoghue (2002), Time Discounting and Time Preference: A Critical Review, *Journal of Economic Literature*, Vol 40, No. 2, 351-401.
- [28] Leonard Green, Nathanael Fristoe and Joel Myerson (1994), Temporal Discounting and Preference Reversals in Choice Between Delayed Outcomes, *Psychonomic Bulletin Review*, Vol 1: 3 p383-89.
- [29] Jonathan Gruber and Botond Koszegi (2000), Is Addiction 'Rational?' Theory and Evidence, *Quarterly Journal of Economics*, 116, p1261-1305.
- [30] Glenn Harrison, Morten Lau and Melonie Williams (2002), Estimating individual discount rates in Denmark: A field experiment, *American Economic Review*, 92.5 p1606-1617.
- [31] Andrew Hertzberg (2010), Exponential Individuals, Hyperbolic Households, Mimeo.
- [32] Matthew Jackson and Leeat Yariv (2011), Collective Dynamic Choice: The Necessity of Time Inconsistency, Mimeo.
- [33] Boyan Jovanovic (1979), Job Matching and the Theory of Turnover, *Journal of Political Economy*, 87.5 Part 1 p972-990.
- [34] Tjalling Koopmans (1960), Stationary Ordinal Utility and Impatience, *Econometrica*, 28, p287-309.
- [35] Kris Kirby and Richard Herrnstein (1995), Preference Reversals due to Myopic Discounting of Delayed Reward, *Psychological Science*, 6.2 p83-89.
- [36] David Laibson (1997), Golden Eggs and Hyperbolic Discounting, *Quarterly Journal of Economics*, 112, p443-477.
- [37] David Laibson (1998), Life-Cycle Consumption and Hyperbolic Discount Functions, *European Economic Review*, 42, p861-871.
- [38] Stephen Lich-Tyler (2004), Negotiations and Love Songs: The Dynamics of Bargained Household Decisions, Manuscript: University of Michigan.
- [39] Maurizio Mazzocco (2007), Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment, *Review of Economic Studies*, 74.3 p857-895.
- [40] Samuel McClure, David Laibson, George Lowenstein and Jonathan Cohen (2004), Separate Neural Systems Value Immediate and Delayed Monetary Rewards, *Science*, 306 p503-507.
- [41] Samuel McClure, Ericson, David Laibson, George Lowenstein and Jonathan Cohen (2004), Time Discounting for Primary Rewards, *The Journal of Neuroscience*, 27.21 p5796-5804.
- [42] Stephen Marglin (1963), The Social Rate of Discount and The Optimal Rate of Investment, *The Quarterly Journal of Economics*, 77.1, p95-111.
- [43] Albert Marcet and Ramon Marimon (1998), Recursive Contracts, Manuscript, Universitat Pompeu Fabra.
- [44] Charles Manski (2007), Identification for Prediction and Decision, Cambridge, MA: Harvard University Press.
- [45] Janet Metcalfe and Walter Mischel (1999), A Hot/Cool System Analysis of Delay of Gratification: Dynamics of Willpower, *Psychological Review*, 106, p3-19.

- [46] Phillip Nelson (1970), Information and Consumer Behavior, *Journal of Political Economy*, 78.2 p311-329.
- [47] Ted O'Donoghue and Matthew Rabin (1999), Doing It Now or Later, *American Economic Review*, 89.1 p103-124.
- [48] Ted O'Donoghue and Matthew Rabin (2001), Choice and Procrastiation, *Quarterly Journal of Economics*, 116.1 p121-160.
- [49] Karl Popper (1959), *The Logic of Scientific Discovery*, London: Hutchinson and Company.
- [50] Paul Samuelson (1937), A Note on the Measurement of Utility, *Review of Economic Studies*, 4, p155-161.
- [51] Simone G. Schaner (2012), Do Opposites Detract? Intrahousehold Preference Heterogeneity and Inefficient Strategic Savings, Mimeo.
- [52] Reinhard Selten (1991), Properties of a Measure of Predictive Success, *Mathematical Social Sciences*, 21, p153-167.
- [53] Avner Shaked and John Sutton (1984), Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model, *Econometrica*, 52.6 p1351-1364.
- [54] Richard Thaler and Shlomo Benartzi (2004), Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving, *Journal of Political Economy*, 112. 1 p164-187.
- [55] Richard Thaler and Hersh Shefrin (1981), An Economic Theory of Self Control, *Journal of Political Economy*, 89.2, p392-410.
- [56] Hal Varian (1982), The Nonparametric Approach to Demand Analysis, *Econometrica*, 50, 945-973.
- [57] Hal Varian (1985), Nonparametric Analysis of Optimizing Behavior with Measurement Error, *Journal of Econometrics*, 30, 445-458.
- [58] Yoram Weiss and Robert Willis (1997), Match Quality, New Information and Martial Dissolution, *Journal of Labor Economics*, 15.1 Part 2: Essays in Honor of Yoram Ben-Porath p293-329.
- [59] Stephane Zuber (2010), The Aggregation of Preferences: Can We Ignore the Past?, Mimeo.

Appendix A: Proofs

This appendix gives the proofs of our Theorems 1-3.

Proof of Theorem 1

(Necessity) Suppose the set of observations S can be rationalised by the time consistency model. Let η denote the Lagrange multiplier associated with the household budget constraint. We get the following first order constraints for the household optimisation problem (with $\frac{\partial u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t)}{\partial \mathbf{q}_t^m}$ ($m \in \{A, B\}$) and $\frac{\partial u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t)}{\partial \mathbf{Q}_t}$ the subgradients for the function u^H at $(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t)$):

$$\begin{aligned} \beta^{t-1} \frac{\partial u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t)}{\partial \mathbf{q}_t^m} &\leq \eta \mathbf{p}_t \quad (m \in \{A, B\}) \text{ and} \\ \beta^{t-1} \frac{\partial u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t)}{\partial \mathbf{Q}_t} &\leq \eta \mathbf{P}_t. \end{aligned}$$

Under concavity of the function u^H , we have

$$\begin{aligned} u^H(\mathbf{q}_s^A, \mathbf{q}_s^B, \mathbf{Q}_s) - u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t) &\leq \left(\frac{\partial u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t)}{\partial \mathbf{q}_t^A} \right)' (\mathbf{q}_s^A - \mathbf{q}_t^A) + \\ &\quad \left(\frac{\partial u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t)}{\partial \mathbf{q}_t^B} \right)' (\mathbf{q}_s^B - \mathbf{q}_t^B) + \\ &\quad \left(\frac{\partial u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t)}{\partial \mathbf{Q}_t} \right)' (\mathbf{Q}_s - \mathbf{Q}_t). \end{aligned}$$

Combining the different inequalities leads to

$$u^H(\mathbf{q}_s^A, \mathbf{q}_s^B, \mathbf{Q}_s) - u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t) \leq \frac{\eta}{\beta^{t-1}} [\mathbf{p}'_t(\mathbf{q}_s - \mathbf{q}_t) + \mathbf{P}'_t(\mathbf{Q}_s - \mathbf{Q}_t)].$$

By using $\frac{u^H(\mathbf{q}_s^A, \mathbf{q}_s^B, \mathbf{Q}_s)}{\eta} = u_s^H$, we obtain the inequalities in Theorem 1.

(Sufficiency) Suppose the inequalities in Theorem 1 hold. Then, define the following felicity function:

$$u^H(\mathbf{x}^A, \mathbf{x}^B, \mathbf{X}) = \min_s \left(u_s^H + \frac{1}{\beta^{s-1}} [\mathbf{p}'_s(\mathbf{x}^A - \mathbf{q}_s^A) + \mathbf{p}'_s(\mathbf{x}^B - \mathbf{q}_s^B) + \mathbf{P}'_s(\mathbf{X} - \mathbf{Q}_s)] \right).$$

Using a straightforwardly similar argument as Varian (1982), we can derive $u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t) = u_t^H$.

Consider any consumption plan $\{\mathbf{x}_t^A, \mathbf{x}_t^B, \mathbf{X}_t\}_{t \in T}$ such that

$$\sum_{t \in T} [\mathbf{p}'_t(\mathbf{x}_t^A - \mathbf{q}_t^A) + \mathbf{p}'_t(\mathbf{x}_t^B - \mathbf{q}_t^B) + \mathbf{P}'_t(\mathbf{X}_t - \mathbf{Q}_t)] \leq 0,$$

i.e. the consumption plan $\{\mathbf{x}_t^A, \mathbf{x}_t^B, \mathbf{X}_t\}_{t \in T}$ is affordable given the outlay associated with $\{\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t\}_{t \in T}$. Then, we need to show that

$$\sum_{t \in T} \beta^{t-1} u^H(\mathbf{x}_t^A, \mathbf{x}_t^B, \mathbf{X}_t) \leq \sum_{t \in T} \beta^{t-1} u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t).$$

Using the utility function defined above, we obtain

$$\begin{aligned} &\sum_{t \in T} \beta^{t-1} u^H(\mathbf{x}_t^A, \mathbf{x}_t^B, \mathbf{X}_t) \\ &\leq \sum_{t \in T} \beta^{t-1} \left(u_t^H + \frac{1}{\beta^{t-1}} [\mathbf{p}'_t(\mathbf{x}_t^A - \mathbf{q}_t^A) + \mathbf{p}'_t(\mathbf{x}_t^B - \mathbf{q}_t^B) + \mathbf{P}'_t(\mathbf{X}_t - \mathbf{Q}_t)] \right) \\ &= \sum_{t \in T} \beta^{t-1} u_t^H + \sum_{t \in T} [\mathbf{p}'_t(\mathbf{x}_t^A - \mathbf{q}_t^A) + \mathbf{p}'_t(\mathbf{x}_t^B - \mathbf{q}_t^B) + \mathbf{P}'_t(\mathbf{X}_t - \mathbf{Q}_t)] \\ &\leq \sum_{t \in T} \beta^{t-1} u_t^H, \end{aligned}$$

so that $u^H(\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t) = u_t^H$ gives the wanted conclusion.

Proof of Theorem 2

(Necessity) Suppose the set of observations S can be rationalised by the full efficiency model. Let η denote the Lagrange multiplier associated with the household budget constraint. We get the following first order constraints (with $\frac{\partial u^m(\mathbf{q}_t^m, \mathbf{Q}_t)}{\partial \mathbf{q}_t^m}$ and $\frac{\partial u^m(\mathbf{q}_t^m, \mathbf{Q}_t)}{\partial \mathbf{Q}_t}$ ($m \in \{A, B\}$) the subgradients for the function u^m at bundle

$(\mathbf{q}_t^m, \mathbf{Q}_t)$:

$$\begin{aligned}\beta_A^{t-1} \frac{\partial u^A(\mathbf{q}_t^A, \mathbf{Q}_t)}{\partial \mathbf{q}_t^A} &\leq \eta \mathbf{P}_t, \\ \mu \beta_B^{t-1} \frac{\partial u^B(\mathbf{q}_t^B, \mathbf{Q}_t)}{\partial \mathbf{q}_t^B} &\leq \eta \mathbf{P}_t, \\ \beta_A^{t-1} \frac{\partial u^A(\mathbf{q}_t^A, \mathbf{Q}_t)}{\partial \mathbf{Q}_t} + \mu \beta_B^{t-1} \frac{\partial u^B(\mathbf{q}_t^B, \mathbf{Q}_t)}{\partial \mathbf{Q}_t} &\leq \eta \mathbf{P}_t.\end{aligned}$$

Under concavity of each felicity functions u^m , we have

$$u^m(\mathbf{q}_s^m, \mathbf{Q}_s) - u^m(\mathbf{q}_t^m, \mathbf{Q}_t) \leq \left(\frac{\partial u^m(\mathbf{q}_t^m, \mathbf{Q}_t)}{\partial \mathbf{q}_t^m} \right)' (\mathbf{q}_s^m - \mathbf{q}_t^m) + \left(\frac{\partial u^m(\mathbf{q}_t^m, \mathbf{Q}_t)}{\partial \mathbf{Q}_t} \right)' (\mathbf{Q}_s - \mathbf{Q}_t).$$

Now define, for each $t \in T$,

$$\begin{aligned}\lambda^A &= \eta, \lambda^B = \frac{\eta}{\mu}, \\ \mathbf{P}_t^A &= \frac{\beta_A^{t-1}}{\lambda^A} \frac{\partial u^A(\mathbf{q}_t^A, \mathbf{Q}_t)}{\partial \mathbf{Q}_t}, \\ \mathbf{P}_t^B &= \mathbf{P}_t - \mathbf{P}_t^A.\end{aligned}$$

Combination yields

$$u^m(\mathbf{q}_s^m, \mathbf{Q}_s) - u^m(\mathbf{q}_t^m, \mathbf{Q}_t) \leq \frac{\lambda^m}{\beta_m^{t-1}} \mathbf{P}_t'(\mathbf{q}_s^m - \widehat{\mathbf{q}}_t^m) + \frac{\lambda^m}{\beta_m^{t-1}} (\mathbf{P}_t^m)'(\mathbf{Q}_s - \mathbf{Q}_t).$$

By using $u^m(\mathbf{q}_s^m, \mathbf{Q}_s)/\lambda^m = u_s^m$, we obtain the inequalities in Theorem 2.

(Sufficiency) Suppose the inequalities in Theorem 2 hold. then, define the following felicity function for each member $m \in \{A, B\}$:

$$u^m(\mathbf{x}^m, \mathbf{X}) = \min_s \left(u_s^m + \frac{\lambda^m}{\beta_m^{s-1}} [\mathbf{p}'_s(\mathbf{x}^m - \mathbf{q}_s^m) + \mathbf{P}'_s(\mathbf{X} - \mathbf{Q}_s)] \right).$$

Using a straightforwardly similar argument as Varian (1982), we can derive $u^m(\mathbf{q}_t^m, \mathbf{Q}_t) = u_t^m$.

Consider any consumption plan $\{\mathbf{x}_t^A, \mathbf{x}_t^B, \mathbf{X}_t\}_{t \in T}$ such that

$$\sum_{t \in T} [\mathbf{p}'_t(\mathbf{x}_t^A - \mathbf{q}_t^A) + \mathbf{p}'_t(\mathbf{x}_t^B - \mathbf{q}_t^B) + \mathbf{P}'_t(\mathbf{X}_t - \mathbf{Q}_t)] \leq 0,$$

i.e. the consumption plan $\{\mathbf{x}_t^A, \mathbf{x}_t^B, \mathbf{X}_t\}_{t \in T}$ is affordable given the outlay associated with $\{\mathbf{q}_t^A, \mathbf{q}_t^B, \mathbf{Q}_t\}_{t \in T}$. Then, for $\omega^A = 1$ and $\omega^B = \omega$ we need to show that

$$\sum_{m \in \{A, B\}} \sum_{t \in T} \omega^m \beta_m^{t-1} u^m(\mathbf{x}_t^m, \mathbf{X}_t) \leq \sum_{m \in \{A, B\}} \sum_{t \in T} \omega^m \beta_m^{t-1} u^m(\mathbf{q}_t^m, \mathbf{Q}_t).$$

Without losing generality, we can assume $\eta = 1$ or, equivalently, $\omega^m = \frac{1}{\lambda^m}$. As such, we obtain

$$\begin{aligned}
& \sum_{m \in \{A, B\}} \sum_{t \in T} \omega^m \beta_m^{t-1} u^m(\mathbf{x}_t^m, \mathbf{X}_t) \\
& \leq \sum_{m \in \{A, B\}} \sum_{t \in T} \omega^m \beta_m^{t-1} \left(u_t^m + \frac{\lambda^m}{\beta_m^{t-1}} [\mathbf{p}'_t(\mathbf{x}_t^m - \mathbf{q}_t^m) + \mathbf{P}'_t(\mathbf{X}_t - \mathbf{Q}_t)] \right) \\
& = \sum_{m \in \{A, B\}} \sum_{t \in T} \omega^m \beta_m^{t-1} u_t^m + \sum_{m=A, B} \sum_{t \in T} [\mathbf{p}'_t(\mathbf{x}_t^m - \mathbf{q}_t^m) + \mathbf{P}'_t(\mathbf{X}_t - \mathbf{Q}_t)] \\
& \leq \sum_{m \in \{A, B\}} \sum_{t \in T} \omega^m \beta_m^{t-1} u_t^m,
\end{aligned}$$

so that $u^m(\mathbf{q}_t^m, \mathbf{Q}_t) = u_t^m$ gives the wanted conclusion.

Proof of Theorem 3

The result uses that, for a given partition \mathbf{T} , consistency with the no-commitment model requires that, for each subset T_τ ($\tau \in \{1, \dots, \Upsilon\}$), that the corresponding subset of observations $\{\mathbf{q}_t, \mathbf{Q}_t; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T_\tau}$ can be rationalized by the full commitment model. Then, the result follows directly from Theorem 2.

Appendix B: Measurement error and imperfect capital markets

In our assessment of the time consistent model we also considered (in Section 3.2) robustness checks with respect to measurement error and imperfect capital markets. In this appendix we provide a formal presentation of the methodology we used for this purpose.

Measurement error

As explained in the main text, our measurement error procedure assumes the following multiplicative error structure:

$$q_{t,n}^* = (1 + \epsilon_{t,n}) q_{t,n} \text{ and } Q_{t,k}^* = (1 + \varepsilon_{t,k}) Q_{t,n},$$

for $q_{t,n}^*$ and $Q_{t,n}^*$ the true (but unobserved) values of the private and public quantities, respectively. Let us assume that the *true* data set $S^* = \{\mathbf{q}_t^*, \mathbf{Q}_t^*; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T}$ is rationalisable by the time consistent model. Unfortunately, we only observe the *actual* set data $S = \{\mathbf{q}_t, \mathbf{Q}_t; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T}$, which may not be theory-rationalisable.

Define the following perturbation to observed quantities:

$$\tilde{q}_{t,n} = (1 + \tilde{\epsilon}_{t,n}) q_{t,n} \text{ and } \tilde{Q}_{t,k} = (1 + \tilde{\varepsilon}_{t,k}) Q_{t,n},$$

While we cannot observe the actual errors $\epsilon_{t,n}$ and $\varepsilon_{t,k}$, we can calculate the smallest perturbations $\tilde{\epsilon}_{t,n}$ and $\tilde{\varepsilon}_{t,k}$ necessary such that the perturbed data set $\tilde{S} = \{\tilde{\mathbf{q}}_t, \tilde{\mathbf{Q}}_t; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T}$ is consistent with the time consistent model. Specifically, we minimise the sum of squared error terms:

$$\min \tilde{\mathbf{V}} = \sum_{t \in T} \left(\sum_{n=1}^N (\tilde{\epsilon}_{t,n})^2 + \sum_{k=1}^K (\tilde{\varepsilon}_{t,k})^2 \right),$$

subject to the constraint that \tilde{S} satisfies the model.

$\tilde{\mathbf{V}}$ can be used to assess the null hypothesis that the true data S can be rationalised by the model as soon as we account for measurement error in observed quantities. If we assume that the true errors $\epsilon_{t,n}$ and

$\varepsilon_{t,k}$ are independently normally distributed with zero mean and constant variance σ^2 , then we know that

$$\frac{\mathbf{V}}{\sigma^2} \sim \chi_{120}^2,$$

where \mathbf{V} is the sum of squared errors associated with the true quantities, $\mathbf{V} = \sum_{t \in T} (\sum_{n=1}^N (\varepsilon_{t,n})^2 + \sum_{k=1}^K (\varepsilon_{t,k})^2)$, and the degrees of freedom equal the number of quantity errors, $|T| * N = 8 * 15$.¹⁶ If \mathbf{V} and σ^2 were observable, the null hypothesis would be rejected if $\mathbf{V}/\sigma^2 \geq C_\alpha$, where C_α represents the critical value from the chi-squared distribution at significance level α .

Given that neither \mathbf{V} nor σ^2 are observable in practice, we use $\tilde{\mathbf{V}}$ to approximate the sum of squared measurement errors. As $\tilde{\mathbf{V}}$ corresponds to the minimal perturbation of the quantities such that the data are rationalisable, we must have $\tilde{\mathbf{V}} \leq \mathbf{V}$. Therefore, our test statistic $\tilde{\mathbf{V}}/\sigma^2$ is conservative. Finally, given that the true extent of the measurement error, σ^2 , is unobserved, we follow the procedure suggested by Varian (1985). We calculate the critical standard deviation σ_α needed such that the null hypothesis of theory-rationalisable behaviour cannot be rejected at some significance level α :

$$\sigma_\alpha = \sqrt{\frac{\tilde{\mathbf{V}}}{C_\alpha}}.$$

In words, we reject rationalisability at a significance level α if and only if σ_α exceeds our prior beliefs about the likely magnitude of σ .

Imperfect capital markets

Under the assumption of perfect capital markets, households are able to borrow and save at the same interest rate, and there are no limits regarding the maximum amount they can borrow. If we denote savings in period t by D_t , then this simply assumes that $D_t \in \mathbb{R}$, where borrowing corresponds to negative saving. When combined with the assumption of perfect foresight, this ideal setting implies perfect *ex ante* consumption smoothing across time periods or, equivalently, that the marginal utility of wealth in each period t (denoted by η_t) is kept constant over time, i.e. $\eta_t = \eta$ for all $t \in T$.

To account for potentially binding borrowing constraints (and thus imperfect capital markets), we adopt a procedure similar to the one of Demuyne and Verriest (2012). Specifically, we assume that the household's ability to achieve a perfectly smooth consumption path is hindered by the fact that D_t cannot fall below some (unobserved) maximum borrowing amount, d_t . If this constraint is binding in some period (i.e. $D_t = -d_t$ for some $t \in T$), then the household is forced to consume less than what would be optimal without borrowing constraints.

In our revealed preference characterisation of the time consistent model, we can include the possibility of borrowing constraints by adding the condition $\eta_t \geq \eta_{t+1}$. Then, we have $\eta_t = \eta_{t+1}$ (only) if $D_t > -d_t$, which means the borrowing constraint is not binding. Conversely, $\eta_t > \eta_{t+1}$ whenever $D_t = -d_t$, which complies with imperfect consumption smoothing.

Summarising, by allowing the marginal utility of wealth to decrease (but not increase) over time, we can effectively account for potentially binding borrowing constraints. Including this argument in the proof of Theorem 1 yields the following revealed preference conditions for the time consistent model under imperfect capital markets.

Corollary 1 *The set of observations $S = \{\mathbf{q}_t, \mathbf{Q}_t; \mathbf{p}_t, \mathbf{P}_t\}_{t \in T}$ can be rationalised by the time consistency model with (unobserved) borrowing constraints if and only if there exist, for all $t \in T$, a utility number $u_t^H \in \mathbb{R}$, a marginal utility of wealth $\eta_t > 0$ and a positive constant $\beta \in (0, 1]$ that, for any $s, t \in T$, satisfy*

$$\begin{aligned} u_s^H - u_t^H &\leq \frac{\eta_t}{\beta^t} [\mathbf{p}'_t(\mathbf{q}_s - \mathbf{q}_t) + \mathbf{P}'_t(\mathbf{Q}_s - \mathbf{Q}_t)], \\ \eta_t + 1 &\leq \eta_t. \end{aligned}$$

¹⁶Note that since prices are strictly positive, the relative errors are bounded from below; $\varepsilon_{t,n}, \varepsilon_{t,k} \in]-1, +\infty[$. However, imposing normality should not be a concern, since these errors are typically not too far removed from zero.

These new revealed preference conditions are generally weaker than the ones in Theorem 1. The two sets of conditions (only) coincide in the limiting case where the (unobserved) borrowing constraints are assumed never to be binding (this corresponds to $\eta_t = \eta$ for all $t \in T$, which was used in the proof of Theorem 1). Interestingly, the conditions remain linear in unknowns for a given value of β .

