

NBER WORKING PAPER SERIES

SUFFICIENT STATISTICS FOR MEASURING FORWARD-LOOKING WELFARE

David Baqaee
Ariel Burstein
Yasutaka Koike-Mori

Working Paper 32567
<http://www.nber.org/papers/w32567>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2024, Revised August 2025

We thank Andy Atkeson, Adrien Auclert, Maya Eden, Valentin Haddad, Oleg Itskhoki, Chad Jones, Greg Kaplan, Patrick Kehoe, Narayana Kocherlakota, John Fernald, Kieran Larkin, Jay Lu, Hannes Malmberg, Ricardo Reis, Pierre-Olivier Weill, and Ivan Werning for comments. We acknowledge financial support from NSF grant No. 2343691 and the Alfred P. Sloan Foundation. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by David Baqaee, Ariel Burstein, and Yasutaka Koike-Mori. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Sufficient Statistics for Measuring Forward-Looking Welfare
David Baqaee, Ariel Burstein, and Yasutaka Koike-Mori
NBER Working Paper No. 32567
June 2024, Revised August 2025
JEL No. C14, D0, D10, D11, E0, E01, E20, I3, J6

ABSTRACT

We propose a sufficient-statistics method to measure a money-metric index of dynamic welfare. The method relies on two assumptions: (i) individuals' preferences are separable between present and future consumption, and (ii) a subset of households bear no undiversifiable idiosyncratic risk. For this group, we infer the value of the future relative to the present from consumption-saving behavior. We then pin down the welfare of all other households by matching their budget shares. Our method accommodates incomplete markets, life-cycle motives, non-rational expectations, and non-exponential discounting—without specifying functional forms. Implementation needs only time series of prices, repeated cross-sectional income-balance-sheet- spending data, and an estimate of the intertemporal-substitution elasticity. In simulations, our method tracks true welfare better than net-present-value calculations, especially under credit constraints. Applied to the PSID, our measure differs sharply from a static index. Our estimates can be used to study the welfare effect of shocks without enumerating every margin and forecasting every horizon through which those shocks operate.

David Baqaee
University of California, Los Angeles
Department of Economics
and CEPR
and also NBER
baqaee@econ.ucla.edu

Yasutaka Koike-Mori
University of North Carolina at Chapel Hill
Department of Economics
yasutaka.koike.mori@gmail.com

Ariel Burstein
University of California, Los Angeles
Department of Economics
and NBER
arielb@econ.ucla.edu

1 Introduction

Measuring changes in consumers' welfare is an essential task of economics. Consumers are exposed to thousands of prices that affect them in different ways, and it is the economist's job to compress all this information down into a single number that summarizes each consumer's well-being. In static settings, the money-metric utility function accomplishes this task, and there are well-known sufficient statistics for computing it. To measure changes in welfare using the money-metric, we need to deflate changes in income by average price changes, weighing each individual price change by its budget share.¹

In dynamic settings, welfare depends not just on income and prices today, but also on expected state-contingent income and prices in the future. If a full set of contingent claims markets existed, then measuring welfare in dynamic settings would be as straightforward as measuring welfare in static settings. However, to quote Samuelson (1961): *"the futures prices needed for making the requisite wealth-like comparisons are simply unavailable. So it would be difficult to make operational the theorists' desired measures. . . . the national income statistician is very far from having even an approximation to the data needed for these comparisons."*

Therefore, when dynamic considerations are important, researchers tend to rely on one of two approaches. The first is to work with fully-specified dynamic models, where we specify preferences, beliefs, and all other parameters, and then measure welfare using the model structure. The second is to compute the net-present value of real wealth, using asset markets to approximate complete contingent-claims markets. This entails forecasting future cashflows and prices, and discounting them back to the present using some discount factor. Both approaches require taking a stance on future state-contingent prices, cashflows, returns, probabilities, and plans.

This paper develops an alternative approach to measuring dynamic welfare, sidestepping the need to take a stance on expectations about the future. To do this, we adapt the standard definition of money-metrics to settings with uncertainty and market incompleteness. To quantify the welfare of an agent facing a dynamic stochastic problem in year τ , we ask: what is the one-time lump sum payment the agent must receive in some base year (with no other income sources thereafter) such that the agent is indifferent between their initial problem and this counterfactual problem. This amount of wealth is the money-metric utility associated with the problem the agent faces in year τ , in base year, τ_0 , dollars. This measure is both a "price"-equivalent, because it converts τ prices into τ_0

¹The money-metric converts expenditures under some vector of prices into equivalent expenditures given some baseline vector of prices. Changes in welfare can then be quantified by comparing changes in equivalent expenditures under some fixed (base) price vector. For a textbook presentation, see Deaton and Muellbauer (1980).

prices, and a “risk”-equivalent, because it converts risky labor income into an equivalent lump-sum payment. If financial markets are complete, then this definition coincides with the usual money-metric definition.

We then propose a sufficient-statistics method to estimate this forward-looking measure of welfare for a sample of consumers with common preferences. Our analysis relies on two key assumptions. The first assumption is a type of separability in preferences between the present and the future. The second assumption is that there is a subset of households that do not face idiosyncratic undiversifiable risk, which we call “rentier” households. (Rentiers can and do face aggregate risk). Given these two assumptions, we can obtain forward-looking measures of welfare without further assumptions about utility functions (e.g. CES across goods), risk preferences (e.g. expected utility), time preferences (e.g. exponential discounting), beliefs (e.g. rational expectations), and financial frictions (e.g. complete markets), and without resorting to first-order approximations.

We sketch the basic idea of our approach, starting with the rentier subsample. From changes in consumption–saving choices—given an estimate of the elasticity of intertemporal substitution (EIS)—we infer movements in the intertemporal price of future versus present consumption. When $EIS < 1$ (the empirically relevant case), a rise in the consumption-to-wealth ratio reveals that the price of consuming in the future relative to the present must be falling. Hence, a forward-looking measure of welfare, that takes future expected prices into account, will be higher than a static measure that only compares contemporaneous price changes.²

A caveat is that the relevant object is the *compensated* consumption-to-wealth ratio, which strips out wealth effects and isolates substitution effects. Only substitution-driven changes are informative about movements in relative intertemporal prices. This distinction is quantitatively important because consumption–saving behavior is highly non-homothetic and sensitive to permanent income (see, e.g., Straub, 2019). We recover the compensated ratio nonparametrically—without specifying or estimating a dynamic model—by matching rentiers across time in the cross-section, extending Baqaee et al. (2024) to a dynamic setting.

For non-rentiers, changes in consumption-to-wealth ratios cannot be used to infer expectations about future prices, because consumption-to-wealth ratios also respond to variation in the bindingness of financial constraints and to temporary fluctuations in risky income. Therefore, to recover money-metric utility for non-rentiers, we use a different

²Our sufficient-statistic formulas take observed changes in goods-and-services prices and in consumption-to-wealth ratios as given. Counterfactual exercises require counterfactual prices and counterfactual changes in consumption-to-wealth ratios, which in turn call for a fully specified structural model.

strategy. Under our assumption of time separability, budget shares in the present can only be a function of relative prices in the present and overall utility. Hence, if budget shares in the present are a one-to-one function of utility conditional on relative prices, then two households facing the same relative prices with the same budget shares must be on the same intertemporal indifference curve. This allows us to construct money-metric values for non-rentiers by matching them with rentiers that have similar static budget shares in the same period.³ We use Monte Carlo simulations to show that our method tracks welfare much better than present value calculations, especially for financially constrained consumers.

Our method requires three pieces of information. First, a repeated cross-sectional survey of static household expenditures that includes some rentier households whose wealth is observed. Second, a time series of static price changes. Finally, knowledge of the compensated elasticity of intertemporal substitution (which could be a constant or vary as a function of wealth and prices). We also require that households' preferences be stable functions of observable characteristics. That is, we rule out unobservable taste shocks, where households with similar observable characteristics have different preferences.

To illustrate our method, we apply our results to the United States using the Panel Study of Income Dynamics (PSID) and price data from the Bureau of Labor Statistics. We find that conventional static consumer price indices overestimate the true dynamic price index in our sample. This means that, over time, households became more optimistic about the future relative to the present. This would be consistent, for example, with an increase in expected returns over time, which would lower the price of consuming in the future relative to the present. Mechanically, our finding follows from the increase in the compensated consumption-to-wealth ratio that we estimate, in combination with an EIS calibrated to be less than one.⁴ Furthermore, we find more heterogeneity in the dynamic cost-of-living index than in the static cost-of-living index across both the wealth distribution and by age group. Static measures of inflation typically find that inflation is overstated for rich households and understated for poor households.⁵ In contrast, we find that dynamic inflation rates tend to be lower for poorer households. Furthermore, dynamic non-homotheticities are much more powerful than static ones in our data.

Our method is useful as an input for reduced-form empirical work studying the welfare

³This does not hold exactly if preferences are not time-separable in the way we define. In Section 5.6, we discuss the performance of our method when time-separability is violated.

⁴This does not imply that consumption-to-wealth ratios rose, over time, at the individual level since observed individual-level consumption-to-wealth ratios are not compensated. Furthermore, at the individual level, consumption-wealth ratios may also respond to lifecycle considerations associated with, e.g., aging.

⁵For example, see Jaravel and Lashkari (2024).

effects of dynamic treatments with uncertain outcomes. Many policies and shocks have complex effects that affect households in ex-ante uncertain ways along many dimensions and at different time horizons. We provide a way to study the welfare effects of such treatments without requiring that researchers enumerate, estimate, and aggregate all the possible ways the treatment affects the household.

To illustrate this approach, we study the welfare consequences of job loss in our sample. To do so, we regress our estimates of money-metric utility on job loss. We find that involuntary job loss is associated with a roughly 20% reduction in money-metric utility. The effects vary by age, and are much milder for households above 60 years old.

Related Literature. Our paper is related to the literature on dynamic measures of welfare. Some early theoretical papers in this literature include Samuelson (1961), Alchian and Klein (1973), Pollack (1975), and Hulten (1979). As mentioned before, the subsequent applied literature can be divided into two branches: papers that use fully-specified dynamic models, and papers that rely on some variation of the net-present value approach. We discuss these two branches in turn.

The first branch uses exact fully-specified models. A prototypical example is described in the AEA presidential address by Lucas (2003). Since this approach is ubiquitous in macroeconomics, we do not provide a systematic review, and instead focus on methodological papers that extend growth and inflation measures to account for dynamics. For example, to construct dynamic measures of inflation and welfare, Reis (2005) and Aoki and Kitahara (2010) calibrate parametric models of household preferences and beliefs, and compute aggregate cost-of-living indices by feeding in the path of observed prices. Reis (2005) uses additively time-separable preferences and considers only financial wealth, whereas Aoki and Kitahara (2010) use Epstein and Zin (1989) preferences and allow for both financial and non-financial wealth. Both papers use homothetic preferences and assume that all assets can be traded — that is, they abstract from idiosyncratic uninsurable risk and borrowing constraints. In contrast, our method accommodates uninsurable risk, borrowing constraints, and non-homothetic preferences. In a different vein, Jones and Klenow (2016) construct country-level welfare measures that account for mortality, leisure, and risk by fully specifying preferences and calibrating stochastic processes for the determinants of utility.⁶

The issue with this approach is that the researcher has to make assumptions about everything. The advantage of the second approach, which uses net-present value calcula-

⁶There are papers that combine the structural modelling approach with the net-present value one. For example, Huggett and Kaplan (2016) use the stochastic discount factor from a structural model to discount the value of future labor income.

tions, is that it does not require a fully specified economic model. By forecasting the future and discounting using asset prices, we can avoid the need to model how the future is determined or what the utility function is. For example, at the macroeconomic level, Basu et al. (2022) show that, to a first-order, the welfare of a country's infinitely-lived representative consumer can be summarized by the net-present value of technology shocks plus the initial capital stock. At the microeconomic level, Fagereng et al. (2022) and Del Canto et al. (2023) use Taylor approximations to calculate how individual consumer welfare responds to shocks to asset prices and monetary policy, respectively. Both papers rely on forecasting the future and discounting the future to the present. Fagereng et al. (2022) estimate how various asset prices changed over time and weigh these by discounted net asset sales. Del Canto et al. (2023) use local projections to forecast how monetary shocks change goods and asset prices and weigh these price changes by discounted expected budget shares. Pallotti et al. (2024) use a similar methodology to study the welfare effects of the recent surge in inflation in the Euro area across the household distribution.⁷

Our paper differs from these papers in some important ways. Some benefits of our approach are that: (1) we do not impose perfect foresight, complete markets, or rely on approximate certainty equivalence; (2) we do not directly estimate or model future asset or goods prices, beliefs, discount factors, and future holdings of assets or purchases of goods. Instead, we back out the welfare impact of future shocks from changes in consumption-savings decisions. This means our measure can accommodate, for example, non-rational expectations; (3) our method accounts for non-homotheticities, which are important beyond the first-order approximations typically used in this literature. We are able to do this because we assume preferences are time-separable, so that all welfare-relevant information about the future can be inferred from choices made in the present.

The net-present value approach automatically unpacks how specific future changes in prices or incomes affect welfare. In contrast, our approach does not automatically break changes in measured welfare down into different channels. Instead, to understand the effect of a specific shock on welfare, we first compute money-metric welfare, and then use a regression of changes in money-metric welfare on the shock of interest to estimate the dynamic treatment effect (subject to the usual caveats about interpreting regression coefficients).

Although our focus is on dynamics, in terms of methods, our paper is related to the literature on static consumer price indices. We build on tools from this literature, tools

⁷There is a broader literature that uses asset prices to answer specific welfare-related questions, but we do not provide a systematic review of this literature. For example, Alvarez and Jermann (2004) measure the welfare cost of consumption fluctuations using asset prices without needing to specify a utility function.

originally developed for static problems, to study dynamic problems. Specifically, our method builds on three different ideas from the static literature. The first is Feenstra (1994), who inverts CES demand curves to infer the value of new goods from changes in expenditures on existing goods.⁸ We extend this idea to infer the value of missing future prices relative to present prices using changes in consumption-to-wealth ratios.

Second, since only changes in consumption-to-wealth ratios due to substitution effects are relevant for inferring changes in relative prices, we build on ideas from Baqaee et al. (2024) to non-parametrically correct for non-homotheticities using cross-sectional data. Other papers that non-parametrically deal with non-homotheticities using cross-sectional data in a static setting are Blundell et al. (2003) and Jaravel and Lashkari (2024).

Third, to infer welfare of financially constrained agents, we build on Hamilton (2001), who infers changes in welfare by tracking changes in the budget share of food. The basic idea is that consumers spend a smaller share of their budget on food as they get richer, so changes in the budget share of food are informative about how rich consumers are.⁹ Recently, Atkin et al. (2024) provide a micro-foundation for this approach and use it to calculate welfare across the income distribution. We use a similar idea to Hamilton (2001), because we match unconstrained rentiers to constrained non-rentiers within the same period using static budget shares. However, since rentiers and non-rentiers in the same period face the same (static) relative prices, we do not need to estimate or correct for substitution effects, as in Atkin et al. (2024).

Roadmap. Section 2 illustrates a key idea in our method using a simple stripped-down complete-markets example. Section 3 generalizes the environment to allow for incomplete markets, uninsurable risk, borrowing constraints, and non-homothetic preferences. Section 4 contains the main results of the paper showing how to extend the basic idea in Section 2 to this more complex environment, first for rentiers and then non-rentiers. We also provide Monte Carlo simulations in this section. Section 5 provides some extensions of the basic framework such as allowing for labor-leisure choice, time inconsistency, secular changes in mortality risk, and violations of time-separability. Section 6 constructs a measure of dynamic welfare for households in the PSID. Section 7 uses this measure to study the dynamic welfare losses from job loss. Section 8 concludes.

⁸This approach is frequently used to infer the value of new goods or quality change in static settings, see, for example Broda and Weinstein (2006), Broda and Weinstein (2010), Blaum et al. (2018), Aghion et al. (2019), Argente and Lee (2021), and Argente et al. (2023). This approach can also be used to measure the gains from trade, i.e. the value of imported goods, as in Arkolakis et al. (2012).

⁹This approach, especially paired with an AIDS functional form, is frequently used to measure inflation in historical settings and settings where data quality is low. See, for example, Costa (2001), Almås (2012), Almås et al. (2018), and Nakamura et al. (2016).

2 Stripped-Down Example

In this section, we demonstrate a key idea using a simple example. We assume there is one consumption good per period, preferences are homothetic with constant EIS, and financial markets are complete. We relax these strong assumptions in the next sections.

Consider consumers with planning horizon J , living at calendar date τ , with preferences represented by

$$\mathcal{U}(c, \pi) = \sum_{j=0}^J \sum_{s_j} \beta^j \pi(s^j|\tau) \frac{c(s^j|\tau)^{1-1/\sigma}}{1-1/\sigma}, \quad (1)$$

where j denotes the number of periods after the start date, the history of realizations up to period j are denoted by s^j , $\pi(s^j|\tau)$ is the probability of history s^j being realized given the start date τ , consumption in period j given history s^j is $c(s^j|\tau)$, β is the discount factor, and σ is the EIS.¹⁰

Households at date τ face the following sequence of budget constraints:

$$p(s^0|\tau)c(s^0|\tau) + \sum_{k' \in S_1} a_{k'}(s^0|\tau) = w, \quad (2)$$

$$p(s^j|\tau)c(s^j|\tau) + \sum_{k' \in S_{j+1}} a_{k'}(s^j|\tau) = R_k(s^{j-1}|\tau)a_k(s^{j-1}|\tau), \quad (0 < j < J) \quad (3)$$

$$p(s^J|\tau)c(s^J|\tau) \leq R_k(s^{J-1}|\tau)a_k(s^{J-1}|\tau). \quad (4)$$

The price of consumption in history s^j , conditional on starting at τ , is denoted by $p(s^j|\tau)$. The holdings of Arrow security k purchased in history s^j are $a_k(s^j|\tau)$ with return $R_k(s^j|\tau)$ in the pay-off state k . The price of every asset is normalized to be one. The first constraint ensures that consumption in the period zero plus purchases of Arrow securities sum to initial wealth, where S_1 is the set of possible realizations of the state in period 1. The second constraint ensures that, at each subsequent history s^j , consumption plus purchases of Arrow securities that pay off in $j+1$ are equal to the returns from previously purchased Arrow securities. The final constraint is a no-Ponzi condition that ensures consumption in every terminal history must be less than the returns from previously purchased Arrow securities. The consumer's financial wealth at the beginning is w .

The value function for a household with initial wealth w in cohort τ is

$$V(\tau, w) = \max_{c, a} \{\mathcal{U}(c, \pi) \text{ subject to (2), (3), and (4)}\}. \quad (5)$$

¹⁰The history s^j is (s^{j-1}, s_j) , where s_j is the realization of the state in period j .

Our objective is to compare the choice set facing different cohorts of households, always keeping the preference parameters, like the planning horizon J and the taste shifters β fixed. For example, we compare the value function of 50 year olds in 2005 to the value function of 50 year olds in 2019. This comparison has to take into account not only wealth and present prices but also expectations about future prices, probabilities, and returns.¹¹

The τ_0 -*money-metric* utility function, $m(\tau, w|\tau_0)$, is defined implicitly via

$$V(\tau, w) = V(\tau_0, m(\tau, w|\tau_0)).$$

In words, $m(\tau, w|\tau_0)$ is the amount of wealth in τ_0 that makes the consumer indifferent to having wealth w at date τ . To compare two different choice sets, say (τ, w) and (τ', w') , we compare $m(\tau, w|\tau_0)$ and $m(\tau', w'|\tau_0)$. The value of (τ', w') is higher than (τ, w) if, and only if, $m(\tau', w'|\tau_0)$ is higher than $m(\tau, w|\tau_0)$. That is, the money-metric is itself a value function that represents the household's preferences.

Since $m(\tau, w|\tau_0)$ has interpretable units, in terms of τ_0 dollars, we can calculate the rate of growth of money-metric wealth between (τ, w) and (τ', w') as $m(\tau', w'|\tau_0)/m(\tau, w|\tau_0)$. We can also calculate the change in the cost of living between τ_0 and τ'_0 as $m(\tau, w|\tau'_0)/m(\tau, w|\tau_0)$. This is the change in the initial wealth needed to reach a (τ, w) -indifference curve in τ'_0 relative to τ_0 . To streamline notation, we frequently suppress dependence on the base period, τ_0 , and write $m(\tau, w)$ in place of $m(\tau, w|\tau_0)$, with the understanding that τ_0 is implicitly held constant.¹²

Since financial markets are complete, all constraints can be combined into a single intertemporal budget constraint:

$$p(s^0|\tau)c(s^0|\tau, w) + \sum_{j=1}^J \sum_{s_j} p(s^j|\tau) \prod_{l=0}^{j-1} R_{s_{l+1}}^{-1}(s^l|\tau) c(s^j|\tau, w) \leq w.$$

¹¹We do not compare the welfare of a single household at different points in their life because preferences and planning horizons change along the life-cycle. As explained by Fisher and Shell (1968), or more recently Baqaee and Burstein (2023), welfare comparisons always involve comparisons of choice sets for some fixed preference relation — they do not involve intertemporal comparisons of interpersonal utility values. In this example, the preference relation ranks lotteries over consumption streams and hence welfare comparisons of lotteries are meaningful. However, the planning horizon J , discount parameter β , and the elasticity σ are parameters of the preference relation. The preference relation does not rank its parameters, and hence, these parameters must be held constant in any welfare comparison.

¹²When preferences are homothetic, the money-metric is also related to the consumption-equivalent variation defined by Lucas (1987). To see this, define $c(s^t|\tau, w)$ to be the optimal consumption in history s^t of cohort τ with initial wealth w — the maximizers in (5). It is easy to verify that $m(\tau, w)$ is the scalar that solves

$$\mathcal{U}(m(\tau, w) \times c(\cdot|\tau_0, 1)) = \mathcal{U}(c(\cdot|\tau, w)), \quad (6)$$

where $c(\cdot|\tau_0, 1)$ is the consumption plan of an agent in τ_0 with unit wealth.

The consumer's problem is therefore isomorphic to a static problem where preferences are CES with elasticity of substitution σ across dates and states.

Let $P(\tau)$ denote the intertemporal CES ideal price index for consumers at date τ :

$$P(\tau) = \left[\sum_{j=0}^J \sum_{s_j} (\beta^j \pi(s^j|\tau))^\sigma p(s^j|\tau)^{1-\sigma} \prod_{l=0}^{j-1} R_{s_{l+1}} (s^l|\tau)^{\sigma-1} \right]^{\frac{1}{1-\sigma}},$$

which depends on all state-contingent prices, returns, and probabilities.

By analogy with the isomorphic static problem, the money metric utility, $m(\tau, w)$, can be computed by deflating nominal wealth by inflation in the intertemporal price index:

$$m(\tau, w) = \frac{w}{P(\tau)/P(\tau_0)}. \quad (7)$$

Even in this highly stylized example, computing the money metric is challenging because it requires knowledge of the entire path of state-contingent probabilities, prices, and returns at both τ and τ_0 , none of which is easy to elicit. Note that, if agents do not have perfect foresight about the future, then realizations from the future (i.e. at $\tau + 1$, $\tau + 2$, etc.) may not be enough since we only observe a single sample-path ex-post. Moreover, the agent's ex-ante beliefs need not coincide with ex-post realizations.¹³

To make progress on the missing information in (7), we take advantage of the fact that information about the future can be gleaned from households' consumption-savings choices. To that end, denote the consumption-to-wealth ratio by

$$B^P(\tau, w) = B^P(\tau) = \frac{p(s^0|\tau) c(s^0|\tau, w)}{w},$$

where the first equality, implying independence from w , follows from homotheticity of preferences. We use the symbol B^P because the consumption-to-wealth ratio is the budget share of the present out of total wealth. From standard CES consumer demand, we know that the consumption-to-wealth ratio depends on the ratio of prices in the present relative to the overall life-time price index:

$$\frac{B^P(\tau)}{B^P(\tau_0)} = \left[\frac{p(s^0|\tau)/P(\tau)}{p(s^0|\tau_0)/P(\tau_0)} \right]^{1-\sigma}. \quad (8)$$

To understand the intuition, consider the realistic case where the EIS is less than one, $\sigma < 1$.

¹³Note the dramatic simplification from assuming that the problem is static, i.e. $J = 0$. In this case, (7) simplifies to $m(\tau, w) = \frac{w}{p(s^0|\tau)/p(s^0|\tau_0)}$, which is nominal wealth deflated by consumer price inflation.

Then, the household consumes more out of their wealth in τ relative to τ_0 if, and only if, the price of consumption in the present relative to the overall price index rises. This is because consumption in the present and the future are complements, so an increase in the price of consumption in the present relative to the future raises spending in the present.

We can rewrite (8) to solve for the change in the intertemporal price index:

$$\log \frac{P(\tau)}{P(\tau_0)} = \log \frac{p(s^0|\tau)}{p(s^0|\tau_0)} - \frac{1}{1-\sigma} \log \frac{B^p(\tau)}{B^p(\tau_0)}. \quad (9)$$

That is, we can infer the change in the intertemporal price index given information about the change in prices in the present, the EIS, and the change in the consumption-to-wealth ratio. Substituting (9) into (7) yields the following proposition.

Proposition 1 (Money Metric for Special Case). *Money-metric welfare for a household in cohort τ with wealth w , in terms of τ_0 dollars, is given by:*

$$\log m(\tau, w) = \log w - \log \frac{p(s^0|\tau)}{p(s^0|\tau_0)} + \frac{1}{1-\sigma} \log \frac{B^p(\tau)}{B^p(\tau_0)}. \quad (10)$$

To understand Proposition 1, suppose that $\sigma < 1$. If consumers at date τ consume a larger fraction of their wealth than consumers at date τ_0 , then this indicates that the price of consuming in the future is lower than consuming in the present for cohort τ than for cohort τ_0 . Hence, the overall price index for households in τ is relatively cheaper than τ_0 compared to a comparison of static inflation rates $p(s^0|\tau)/p(s^0|\tau_0)$. This happens if, for example, expected returns in the future rise or if expected prices fall.¹⁴

The advantage of Proposition 1 is that it can be implemented without direct observation of households' expectations about the future. Of course, to arrive at (10), we made strong assumptions. In what follows, we show how the dynamic money metric can be calculated when we relax some of these assumptions.

3 General Environment and Measure of Welfare

In this section, we set up the general environment. We define preferences and constraints. We then generalize the definition of money-metric welfare. The main theoretical results are in the next section.

¹⁴This logic is similar to Feenstra (1994), who infers changes in the ideal price index due to new goods from changes in expenditures on continuing goods given knowledge of the elasticity of substitution.

3.1 Time-Separable Preferences

Agents have continuous, non-satiated, preferences \succeq over consumption streams, \mathbf{c} , with beliefs, $\boldsymbol{\pi}$. The set of goods available each period is N and we denote the consumption of good n in history s^j by $c_n(s^j)$. Letting S^J be the set of all possible terminal histories implies that \mathbf{c} has dimension $|S^J| \times |N|$.

Without loss of generality, the preference relation \succeq can be represented by an implicitly defined utility function:

$$U = D(\mathbf{c}, \boldsymbol{\pi}, U), \quad (11)$$

where D is homogeneous of degree one in \mathbf{c} .¹⁵

Throughout the paper, we impose the following time-separability condition on preferences.

Definition 1 (Time Separability). The preference relation \succeq is *time separable* if it has a utility function representation that can be written as

$$U = D(\underbrace{P(\mathbf{c}(s^0), U)}_{\text{present bundle}}, \underbrace{F(\{\mathbf{c}(s^j)\}_{j>0}, \{\boldsymbol{\pi}(s^j)\}_{j>0}, U)}_{\text{future bundle}}, U), \quad (12)$$

where D , P , and F are scalar valued function that are increasing and homogeneous of degree one in \mathbf{c} .

Definition 1 is equivalent to imposing a certain type of separability on the compensated demand curves generated by \succeq . Specifically, Definition 1 is equivalent to assuming that spending on i relative to j in a given block (the present or the future bundle) is only a function of relative prices in that block and utility. Hence, the only way future prices affect relative budget shares in the present is through wealth effects (i.e. by changing U). Conversely, relative spending shares in the future (across dates, states, and goods) depend on present prices only through wealth effects.¹⁶ If preferences are homothetic, then we can drop U from the right-hand side of (12).

It is easy to see that the preferences in the stripped down example, defined by equation (1), satisfy time-separability. We provide some additional examples below.

¹⁵To see this, let $\mathcal{U}(\mathbf{c}, \boldsymbol{\pi})$ be some utility function that represents \succeq . Define the distance function to be $\tilde{D}(\mathbf{c}, \boldsymbol{\pi}, U) = \max_{\alpha} \{\alpha : \mathcal{U}(\mathbf{c}/\alpha, \boldsymbol{\pi}) \geq U\}$, which is homogeneous of degree one in \mathbf{c} by construction. The following is an identity $\tilde{D}(\mathbf{c}, \boldsymbol{\pi}, \mathcal{U}(\mathbf{c}, \boldsymbol{\pi})) = 1$. Multiplying both sides by $U(\mathbf{c}, \boldsymbol{\pi})$ yields the desired result.

¹⁶This notion of separability is also sometimes called separability of the distance function or separability of the expenditure function. See Blackorby et al. (1998) for theoretical background. In a static setting, Atkin et al. (2024) refer to this type of separability as “quasi-separability.”

Example 1 (Examples of Time-Separable Preferences). The following preferences are all time separable. Age-dependant discounting:

$$U = \frac{c(s^0)^{1-1/\sigma}}{1-1/\sigma} + \sum_{j=1}^J \beta_j \frac{c(s^j)^{1-1/\sigma}}{1-1/\sigma},$$

where the discount factor can vary as a function of the age of the consumer. Variable risk-aversion:

$$U = \frac{c(s^0)^{1-1/\sigma}}{1-1/\sigma} + \sum_{j=1}^J \beta_j \frac{\left[\sum_{s^j} \pi(s^j) c(s^j) \right]^{\frac{1-1/\sigma}{\gamma(U)}}}{1-1/\sigma},$$

where the risk-aversion parameter, here $\gamma(U)$, can vary as a function of the consumer's utility. Variable elasticity intertemporal of substitution:

$$U = \frac{c(s^0)^{1-1/\sigma(U)}}{1-1/\sigma(U)} + \sum_{j=1}^J \beta_j \frac{c(s^j)^{1-1/\sigma(U)}}{1-1/\sigma(U)},$$

where the elasticity of substitution across time can vary as a function utility. Variable income effects across goods within a period:

$$U = \frac{1}{1-1/\sigma} \mathbb{E}_0 \sum_{j=0}^J \beta^j C_j^{1-\frac{1}{\sigma}}, \quad \text{where} \quad C_j = \left[\sum_n \omega_n^{\frac{1}{\gamma}} \left[\frac{c_{nj}}{U \epsilon_n} \right]^{\frac{1-\gamma}{\gamma}} \right]^{\frac{\gamma}{\gamma-1}},$$

where budget shares on different commodities within a period can vary as a function of utility. The parameter γ captures the compensated elasticity of substitution between goods within a state and period. Appendix A provides some examples that are non-separable.

We impose time separability of preferences throughout the rest of the paper. We discuss how our results are affected if preferences are not time-separable in Section 5.

3.2 Decision Problem of Households

Consumers face a finite planning horizon $J < \infty$ and are indexed by the start date τ .¹⁷ Consumers choose their consumption decisions and portfolio of assets to maximize utility subject to a sequence of state-contingent budget constraints. The first period budget

¹⁷Our results can be extended to infinite horizon. In our baseline, we also assume that mortality risk does not depend on calendar time τ . See Section 5 for a discussion of how our results can be extended to account for secular trends in mortality risk.

constraint requires that the sum of consumption and asset purchases equals initial wealth,

$$\sum_{n \in N} p_n(s^0|\tau) c_n(s^0) + \sum_{k \in K} a_k(s^0) = w, \quad (13)$$

where $p_n(s^0|\tau)$ and $c_n(s^0)$ is the price and consumption of good n in the initial period given starting point τ . There are K different asset types, which may not span the state space, and the quantity of asset k purchased is $a_k(s^0)$. If there are durable goods, say housing, then the stock of durables must be included as an asset, a_k , and the user-cost of service flows must be included as a price, p_k . This is how we treat housing in our empirical application.

At each subsequent history s^j , the agent faces the budget constraint

$$\sum_{n \in N} p_n(s^j|\tau) c_n(s^j) + \sum_{k \in K} a_k(s^j) = \sum_{k \in K} R_k(s^j|\tau) a_k(s^{j-1}) + y(s^j|\tau). \quad (14)$$

where $R_k(s^j|\tau)$ is the return of asset k in history s^j and $y(s^j|\tau)$ is exogenous labor income.¹⁸

There are exogenous and potentially state-contingent ad-hoc borrowing constraints, $X(s^j|\tau) \geq 0$, requiring that

$$\sum_k a_k(s^j) \geq -X(s^j|\tau). \quad (15)$$

We impose a no-Ponzi condition that $X(s^j) = 0$ for every terminal history s^j to ensure the agent cannot end the problem in debt.

The decision problem faced by households depend on the tuple of prices, returns, incomes, probabilities, borrowing constraints, and wealth: $\{\mathbf{p}, \mathbf{R}, \mathbf{y}, \boldsymbol{\pi}, \mathbf{X}, w\}$. Prices, \mathbf{p} , returns, \mathbf{R} , probabilities, $\boldsymbol{\pi}$, and borrowing constraints, \mathbf{X} , are functions of the start date τ , but the stream of cashflows, \mathbf{y} , and initial wealth, w , are consumer-specific. Hence, each household's problem can be indexed by (τ, w, \mathbf{y}) . Define the value function associated with (τ, w, \mathbf{y}) by

$$V(\tau, w, \mathbf{y}) = \max_{c, a} \{\mathcal{U}(c, \boldsymbol{\pi}) \text{ subject to (13), (14), and (15)}\}. \quad (16)$$

3.3 Measuring Welfare and the Cost-of-Living

We measure welfare using a generalization of the money-metric. To do so, we first define rentiers.

Definition 2 (Rentiers). We say that the consumer is a *rentier* if labor income, $y(s^j|\tau)$, is zero

¹⁸In Section 5, we discuss how to extend the model to allow for endogenous labor-leisure choice.

in every state.¹⁹ The problem of a rentier in date τ with wealth w is denoted by $(\tau, w, \mathbf{0})$.

We use rentiers to define money-metric wealth.

Definition 3 (Dynamic Money Metric). Consider a reference period τ_0 , with prices, returns, probabilities and constraints: $\{\mathbf{p}(\cdot|\tau_0), \mathbf{R}(\cdot|\tau_0), \boldsymbol{\pi}(\cdot|\tau_0), \mathbf{X}(\cdot|\tau_0)\}$, with $\mathbf{p}(\cdot|\tau_0) > 0$ and $\mathbf{R}(\cdot|\tau_0) > 0$. The τ_0 -money-metric associated with a decision problem (τ, w, \mathbf{y}) is defined implicitly via

$$V(\tau, w, \mathbf{y}) = V(\tau_0, m(\tau, w, \mathbf{y}|\tau_0), \mathbf{0}).$$

In words, the money-metric, $m(\tau, w, \mathbf{y}|\tau_0)$, maps the decision problem (τ, w, \mathbf{y}) into the equivalent one-off lump-sum payment the consumer would need in τ_0 to ensure indifference. That is, $m(\tau, w, \mathbf{y}|\tau_0)$ is the rentier-equivalent wealth in τ_0 that would make the household indifferent. In a static deterministic world, the generalized money-metric we define coincides with the textbook money-metric (e.g., Deaton and Muellbauer, 1980). We extend this textbook definition to allow for market incompleteness and changes in probabilities, which do not have counterparts in static consumer theory.²⁰

As with classical money metrics, $m(\tau, w, \mathbf{y})$ ranks decision problems in the same order as the underlying preferences.

Proposition 2 (Money-metric cardinalizes utility). *The money metric ranks $m(\tau, w, \mathbf{y}|\tau_0) \geq m(\tau', w', \mathbf{y}'|\tau_0)$ if, and only if, $V(\tau, w, \mathbf{y}) \geq V(\tau', w', \mathbf{y}')$.*

Given the money-metric, we can also define changes in the cost-of-living for different cohorts in the following way.

Definition 4 (Dynamic Cost-of-Living). The change in the cost-of-living between two dates, τ_0 and τ'_0 , for some reference problem (τ, w, \mathbf{y}) is

$$\frac{m(\tau, w, \mathbf{y}|\tau'_0)}{m(\tau, w, \mathbf{y}|\tau_0)}.$$

That is, to compare the cost-of-living between τ_0 and τ'_0 , we use the money-metric to calculate the ratio of lump-sums required in τ_0 and τ'_0 to reach the same indifference curve

¹⁹In Section 5, we provide conditions under which consumers earning risk-free labor income can be treated as-if they are rentiers.

²⁰With incomplete markets, there is not a single budget constraint, so there are many different, non-equivalent, ways one could measure welfare. For example, we could convert each problem into a certainty equivalent annuity value. We instead focus on the equivalent one-off lump-sum payment because this is the welfare measure that we can recover from the data given our assumptions. If preferences are homothetic, then growth rates computed using our definition of the money-metric coincide with growth-rates calculated using the Lucas consumption-equivalent definition, even when markets are incomplete.

$V(\tau, w, \mathbf{y})$. In a static deterministic environment, Definition 4 is the same as the traditional ideal (Konüs) price index from consumer theory.

The goal of the rest of the paper is to infer the money-metric function $m(\tau, w, \mathbf{y}|\tau_0)$.

4 Main Results

We present our main result in steps. In Section 4.1 we present a method for recovering the money-metric for rentiers. In Section 4.2 we show how to recover the money-metric for non-rentiers. In Section 4.3, we provide simulation results comparing the performance of our method to the net-present value approach.

4.1 Obtaining the Money Metric for Rentiers

The next proposition limits attention to the subset of rentiers: consumers for whom risky labor income, \mathbf{y} , is identically equal to $\mathbf{0}$. Inferring the money-metric is simpler for these households because their entire wealth is financial and the ad-hoc borrowing constraints never bind.

Before stating the proposition, we introduce some notation. For each household with wealth w , cashflows \mathbf{y} , in period τ , denote current expenditures by

$$E(\tau, w, \mathbf{y}) = \sum_{n \in N} p_n(s^0|\tau) c_n(s^0|\tau, w, \mathbf{y})$$

and the share of budget spent on good n in the present by

$$B_n(\tau, w, \mathbf{y}) = \frac{p_n(s^0|\tau) c_n(s^0|\tau, w, \mathbf{y})}{E(\tau, w, \mathbf{y})}.$$

Denote the consumption-to-wealth ratio by

$$B^P(\tau, w, \mathbf{y}) = \frac{E(\tau, w, \mathbf{y})}{w}.$$

Denote the compensated elasticity of intertemporal substitution (EIS) for a household facing problem (τ, w, \mathbf{y}) by $\sigma(\tau, w, \mathbf{y})$.²¹

To help build intuition, we consider a series of special cases that build up to the general

²¹See the appendix for a formal definition. Loosely, the compensated EIS controls the change in spending on consumption versus savings if the price of every consumption good in the present rises by the same amount, holding utility constant.

result. To streamline notation, we suppress the dependence of the money-metric on the base period, writing $m(\tau, w, \mathbf{y})$ instead of $m(\tau, w, \mathbf{y}|\tau_0)$ where possible.

We begin by a special case that is quite close to Proposition 1.

Proposition 3. *Suppose preferences are homothetic, the EIS is constant, $\sigma(\tau, w, \mathbf{y}) = \sigma$, and there is only one consumption good per period. Then, for rentiers:*

$$\log m(\tau, w, \mathbf{0}) = \underbrace{\log w - \log \frac{p(s^0|\tau)}{p(s^0|\tau_0)}}_{\text{static "real wealth"}} + \underbrace{\frac{\log (B^P(\tau)/B^P(\tau_0))}{1 - \sigma}}_{\text{adjustment for future}}.$$

This is quite similar to our starting example in Proposition 1. The only differences relative to that example are that (1) financial markets may be incomplete (because assets may not span the state space); (2) preferences over the future can be more general (e.g. the discount factor need not be the same in every period and risk-aversion need not be controlled by the inverse of the EIS). The intuition is similar to Proposition 1, although the proof is slightly more involved.

The assumption of homotheticity in Proposition 3 is strong and counterfactual (as we show in our empirical application). Hence, we relax it.

Proposition 4. *Suppose the EIS is constant, $\sigma(\tau, w, \mathbf{y}) = \sigma$, and there is only one consumption good per period. Then, for rentiers, the money-metric satisfies the following fixed-point problem:*

$$\log m(\tau, w, \mathbf{0}) = \underbrace{\log w - \log \frac{p(s^0|\tau)}{p(s^0|\tau_0)}}_{\text{static "real wealth"}} + \underbrace{\frac{\log [B^P(\tau, w, \mathbf{0})/B^P(\tau_0, m(\tau, w, \mathbf{0}), \mathbf{0})]}{1 - \sigma}}_{\text{adjustment for future}}. \quad (17)$$

When preferences are non-homothetic, the consumption-to-wealth ratio, $B^P(\tau, w, \mathbf{0})$, can change as a function of τ for two reasons: (1) substitution effects, or (2) wealth effects. Substitution effects are changes in B^P due to changes in the price of consumption in the present relative to the future, holding fixed utility. Wealth effects are changes in B^P due to the fact that a household with wealth w in τ may be on a different indifference curve than a household with wealth w in τ_0 , since the real purchasing power of w may have changed.

Since we are using changes in B^P to back-out changes in the relative price of future consumption, we must purge changes in B^P that are caused by wealth effects. To accomplish this, Equation (17) uses the consumption-to-wealth ratio, $B^P(\tau_0, m(\tau, w, \mathbf{0}), \mathbf{0})$, of a rentier in τ_0 on the same indifference curve as $(\tau, w, \mathbf{0})$. For every τ and w , Equation (17) implicitly pins down $m(\tau, w, \mathbf{0})$ as the solution to a fixed-point problem. One way to find a fixed

point is to start with some guess of $m(\tau, w, \mathbf{0})$ (e.g. deflating nominal wealth w by the static inflation rate $p(s^0|\tau)/p(s^0|\tau_0)$), using that guess to compute the right-hand side of (17), and then using the result to update the initial guess, and continuing until convergence.

The following proposition is the fully general case, nesting both Proposition 3 and Proposition 4, and dropping the requirements that the EIS be constant and there be a single consumption good per period.²²

Proposition 5. *The money-metric satisfies the following fixed point problem*

$$\log m(\tau, w, \mathbf{0}) = \log w - \int_{\tau_0}^{\tau} \left(\sum_{n \in N} B_n(x, w_x^*, \mathbf{0}) \frac{d \log p_n}{dx} - \frac{d \log B^P(x, w_x^*, \mathbf{0})/dx}{1 - \sigma(x, w_x^*, \mathbf{0})} \right) dx, \quad (18)$$

where w_x^* solves the equation

$$m(x, w_x^*, \mathbf{0}) = m(\tau, w, \mathbf{0}). \quad (19)$$

for each $x \in [\tau_0, \tau]$. The boundary condition is that $m(\tau_0, w, \mathbf{0}) = w$.

Relative to Proposition 4, we now allow multiple goods per period, and we allow the EIS to vary. The budget shares on goods in present expenditures, B_n , the consumption-to-wealth ratio, B^P , and the EIS, σ , all depend on wealth. Proposition 5 shows that it is the compensated version of these objects that must be used when preferences are non-homothetic. This is ensured by equation (19), which requires that the household is kept indifferent along the integration path between τ_0 and τ . That is, w_x^* is the level of wealth required at period $x \in [\tau_0, \tau]$ such that the consumer stays on the same intertemporal indifference curve. The crucial fact about Proposition 5 is that this defines a fixed point problem. To arrive at $m(\tau, w, \mathbf{0})$, we need to integrate compensated budget shares and compensated changes in savings rate. On the other hand, to perform the necessary compensation, we need to know $m(\tau, w, \mathbf{0})$.

Proposition 5, which generalizes static results in Baqaee et al. (2024), is a fixed-point problem that depends on observables and the EIS. The observables are wealth w , budget shares B_n on goods as a function of time and wealth, changes in goods prices from period to period $d \log p_n/dt$, and consumption-to-wealth ratios, B^P , as a function of time and wealth. Given these observables, and estimates of the EIS, we can solve (18) for the money-metric.

Solution Method. To apply Proposition 5, we can follow the same procedures explained in detail in Baqaee et al. (2024). We begin by guessing a candidate solution $m^0(\tau, w, \mathbf{0})$, for

²²We impose the following technical assumption: $p_n(\cdot|\tau)$, asset returns $R_k(\cdot|\tau)$, and probabilities $\pi(\cdot|\tau)$ are absolutely continuous functions of calendar time τ . We also assume that $\sigma(x, w, \mathbf{0}) \neq 1$ almost everywhere for $x \in [\tau_0, \tau]$.

example, by deflating nominal wealth at each date using the chained inflation index and ignoring the role of the future. We then use this initial guess on the right-hand side of (18) to get a new guess. We then iterate on this until convergence.

Boundaries. Proposition 5 can only be applied inside a suitable boundary. This is because the budget shares $B(\tau, w, \mathbf{0})$ and consumption-to-wealth ratios $B^P(\tau, w, \mathbf{0})$ are observed only for some subset of wealth levels, say $w \in [\underline{w}_x, \bar{w}_x]$ for $x \in [\tau_0, \tau]$. This limits the range of values of w for which we can calculate the money-metric without out-of-sample extrapolation. Intuitively, if for cohort τ and wealth w , the money-metric value $m(\tau, w, \mathbf{0})$ is not in $[m(x, \underline{w}_x, \mathbf{0}), m(x, \bar{w}_x, \mathbf{0})]$ for some $x \in [\tau_0, \tau]$, then we cannot recover $m(\tau, w, \mathbf{0})$ without extrapolation. This is because there are no households in cohort $x \in [\tau_0, \tau]$ that are on the same indifference curve as the rentier with wealth w at time τ .

Fortunately, Proposition 5 automatically provides the boundary over which the money-metric can be calculated without extrapolation. The initial boundary at τ_0 is just the range in the data: $[\underline{w}_{\tau_0}, \bar{w}_{\tau_0}]$. As we solve (18) forward, for each $\tau > \tau_0$, we can update the boundary because the information required to compute it only depends on previous values of the money-metric (see Baqaee et al., 2024 for a discussion of this issue in a static context).

Brief proof sketch. To prove Proposition 5, we rely on the intertemporal expenditure function. We first show that, for each consumer, there exist some shadow intertemporal prices that can rationalize that consumer's choices. For rentiers, these shadow prices depend only on calendar time τ and overall utility. We then show that the expenditure function, given these shadow prices, can be used to compute the money-metric for rentiers. We then apply the envelope theorem and the fundamental theorem of calculus to re-express the money metric as the area under intertemporal Hicksian demand curves with respect to changes in shadow prices. Finally, we use time-separability of preferences to substitute for expected future shadow prices using the consumption-to-wealth ratio. The final step requires recognizing that the area under the Hicksian demand curves can be inferred from cross-sectional data as in Baqaee et al. (2024).

4.2 Obtaining the Money Metric for Non-Rentiers

A challenge for the applicability of Proposition 5 is that, in practice, many households in the sample may be non-rentiers (i.e. $y \neq 0$). This proposition does not apply to non-rentiers since they face potentially binding borrowing constraints and their total wealth,

inclusive of risky labor income, is not priced by financial markets. Fortunately, we can exploit non-homotheticity of preferences to extend $m(\tau, w, \mathbf{0})$ to non-rentiers. To do so, observe the following.

Lemma 1 (Compensated Budget Shares). *If preferences are time separable, then the budget share of each good in the initial period, τ , can be expressed as a function of only present prices and overall utility:*

$$B_n(\tau, w, \mathbf{y}) = b_n(\mathbf{p}(s^0|\tau), V(\tau, w, \mathbf{y})).$$

We refer to b_n as the compensated budget share of n .

Importantly, Lemma 1 implies that the budget share of each good in the present $B_n(\tau, w, \mathbf{y})$ depends on financial wealth w and the stream of payoffs \mathbf{y} only through overall utility $V(\tau, w, \mathbf{y})$.

The next proposition makes it possible to extend $m(\tau, w, \mathbf{0})$ to cover non-rentier households. Recall that $m(\tau, w, \mathbf{y}|\tau_0)$ denotes the money-metric value for the problem faced by an agent at calendar time τ , with wealth w , and payoff streams \mathbf{y} in τ_0 dollars.

Proposition 6 (Money Metric is a Function of Budget Shares and Time). *Suppose that the compensated budget share function, $\mathbf{b}(\mathbf{p}, V)$, is a one-to-one function of V . Then, there exists a function \tilde{m} mapping budget shares and time into money-metric utility values in that period:*

$$m(\tau, w, \mathbf{y}|\tau) = \tilde{m}(\mathbf{B}(\tau, w, \mathbf{y}), \tau),$$

for every τ , w , and \mathbf{y} .

The compensated budget shares $\mathbf{b}(\mathbf{p}, V)$ are a one-to-one function of V if each value of V maps to a different vector of static budget shares. Notably, this rules out homothetic preferences, since once we fix time τ , then the budget shares are constant for every value of V . In words, Proposition 6 implies that, if budget shares are one-to-one with V , then holding time τ fixed, there exists a function $\tilde{m}(\mathbf{B}, \tau)$ mapping vectors of budget shares \mathbf{B} at date τ into the equivalent lump sum wealth at date τ (i.e. $m(\tau, w, \mathbf{y}|\tau)$).

Hence, if we know the function \tilde{m} , and we observe budget shares $\mathbf{B}(\tau, w, \mathbf{y})$ at time τ , then we can deduce the money-metric utility $m(\tau, w, \mathbf{y}|\tau)$ for a household facing the problem (τ, w, \mathbf{y}) . Given $m(\tau, w, \mathbf{y}|\tau)$ we can then use Proposition 5 to convert this to money-metric utility for some other base date $m(\tau, w, \mathbf{y}|\tau_0)$.

How do we learn the shape of the function \tilde{m} ? We use the identity that $m(\tau, w, \mathbf{0}|\tau) = w$ for rentiers. Given this, we can learn the shape of \tilde{m} by solving the following least-squares problem

$$\arg \min_{\hat{m} \in \mathcal{M}} \|w - \hat{m}(\mathbf{B}(\tau, w, \mathbf{0}), \tau)\|,$$

where M is a candidate set of functions that contains \tilde{m} . In words, for every date τ , we fit a flexible function that relates budget shares to wealth for rentiers. We then use this fitted relationship to impute $m(\tau, w, \mathbf{y}|\tau)$ for non-rentiers given their static budget shares. Intuitively, we infer the money-metric wealth for non-rentiers using their static budget shares. Proposition 7 formalizes this idea.

Proposition 7 (Money Metric for Non-Rentiers). *Denote the set of observed rentier budget shares by*

$$\mathcal{B}(\tau) = \{B(\tau, w, \mathbf{0}) : w \in [\underline{w}_\tau, \overline{w}_\tau]\},$$

where $[\underline{w}_\tau, \overline{w}_\tau]$ is the support of the wealth distribution of the rentiers at date τ . Let $\tilde{m}|_{\mathcal{B}(\tau)}$ be the function \tilde{m} restricted to the domain $\mathcal{B}(\tau)$. This restricted function can be recovered by nonlinear least squares:

$$\tilde{m}(\cdot, \tau)|_{\mathcal{B}(\tau)} \in \operatorname{argmin}_{\hat{m} \in M} \int_{\underline{w}_\tau}^{\overline{w}_\tau} (w - \hat{m}(B(\tau, w, \mathbf{0}), \tau))^2 dw.$$

A special case is when the budget share of a specific good, say food, is known to be strictly monotone in utility. (The empirical regularity that the budget share of food declines for richer consumers is called Engel's law.)

Corollary 1 (Engel's Law). *Suppose that there exists a good $i \in N$ whose budget share, $b_i(p(s^0|\tau), V(\tau, w, \mathbf{y}))$, is strictly monotone in V . Then*

$$m(\tau, w, \mathbf{y}) = m(\tau, w^*, \mathbf{0}), \quad \text{if, and only if,} \quad B_i(\tau, w, \mathbf{y}) = B_i(\tau, w^*, \mathbf{0}).$$

In this simple case, if the compensated budget share of i is monotone in utility, then we can deduce that two households (τ, w, \mathbf{y}) and $(\tau, w^*, \mathbf{0})$ have the same utility if, and only if, their budget shares on good i coincide. Proposition 7 is just a generalization of this basic idea.

Propositions 5, 6, and 7 can be combined to recover $m(\tau, w, \mathbf{y})$ for every (τ, w, \mathbf{y}) inside the boundary where we can solve (18) (without extrapolation).

4.3 Monte Carlo Test

We show the finite sample performance of our method using a simple Monte Carlo exercise. Suppose preferences take the form in Example 1:

$$U = \frac{1}{1 - 1/\sigma} \mathbb{E}_0 \left[\sum_{j=0}^{\infty} \beta^j C_j^{1 - \frac{1}{\sigma}} \right], \quad \text{where} \quad C_j = \left[\sum_n \omega_n^{\frac{1}{\gamma}} \left[\frac{c_{nj}}{U \epsilon_n} \right]^{\frac{1-\gamma}{\gamma}} \right]^{\frac{\gamma}{\gamma-1}},$$

where σ is the EIS and γ controls the elasticity of substitution between goods within each period and state. The parameter ϵ_n shapes the wealth elasticity of good n and ω_n is a share parameter.

Consumers face sequential budget constraints of the form:

$$\sum_n p_{nj} c_{nj} + a_{j+1} = Ra_j + y_j,$$

where R is the risk-free rate and y_j is an individual-specific income process. Consumers also face an exogenous borrowing constraint of $-X$.

We calibrate the parameters, simulate a finite population, and then apply our method. For this example, we assume away aggregate risk and consider a scenario where prices and returns are constant through time. Our simulated sample has two sub-populations: a group of rentiers who earn no labor income, and a group of non-rentiers whose labor income process follows an AR(1) process in logs, with persistence parameter 0.975 and standard deviation 0.15, to match estimates of the quarterly persistence and cross-sectional standard deviation of the persistent component of log income. We set the exogenous borrowing limit to -5 , so households can borrow at most 5 times the average quarterly income. The annual risk free rate is 3% (so $R = 1.0075$). The EIS is set to 0.1 following the estimates of Best et al. (2020), the within-period elasticity γ is equal to 0.25. Our calibration has three goods within each period, and we set the vector of ϵ 's to be $[0.0, -0.3, -0.6]$. Finally we set the discount factor $\beta = 0.96$.

Figure 1 shows the absolute value of the log error in money-metric values using our methodology for non-rentiers, which is quite small, as expected. Figure 1 also shows the error associated with using a naive net-present value (NPV) calculation to estimate money-metric utility for non-rentiers. That is,

$$m^{NPV}(\tau, w, \mathbf{y}|\tau) = Ra_0 + \sum_{j=0}^{\infty} R^{-j} \mathbb{E}_0[y_j].$$

If financial markets were complete then $m^{NPV}(\tau, w, \mathbf{y}|\tau) = m(\tau, w, \mathbf{y}|\tau)$. Without complete markets, the NPV approach drastically overstates money-metric utility for non-rentiers, especially for poorer and more financially constrained rentiers exposed to more idiosyncratic risk.

Figure 2 displays the error in the log change in money metric utility due to idiosyncratic income shocks both for our method and the NPV approach. Our method continues to perform very well. In changes, the errors are smaller for the NPV approach but can still

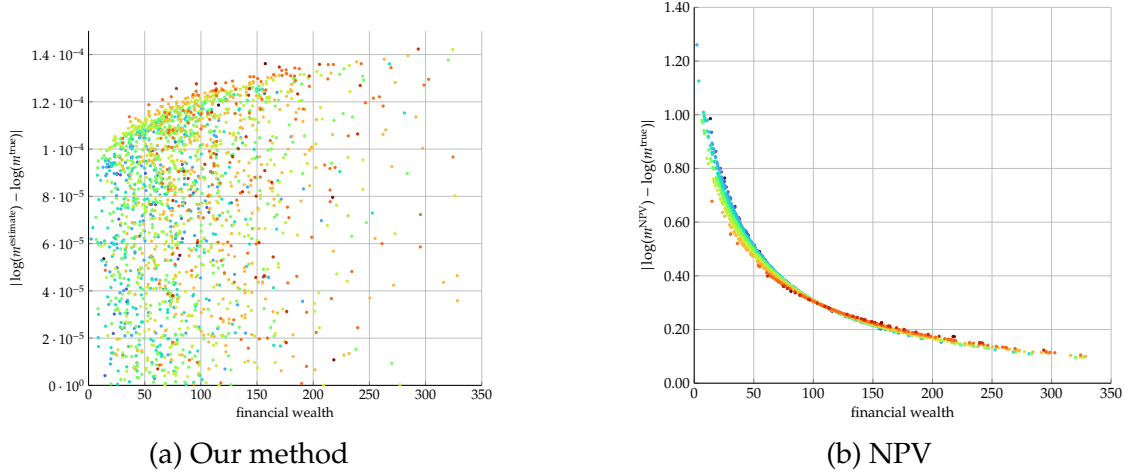


Figure 1: Comparison of errors in estimating money metric utility

Note: Colors indicate idiosyncratic income draws.

be large for certain households, especially those that experience large income shocks and are financially constrained.

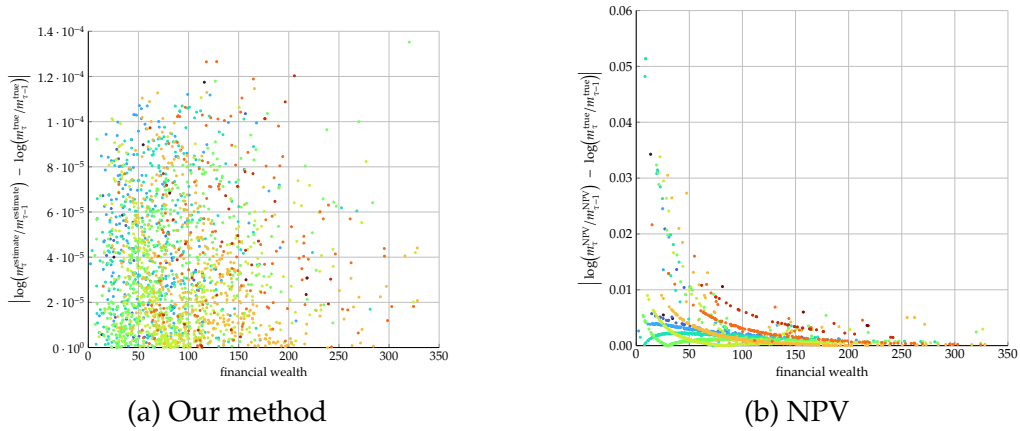


Figure 2: Comparison of errors in estimating changes in money metric utility

Note: Colors indicate idiosyncratic income draws.

5 Extensions and Robustness

Before presenting the empirical application, we discuss some extensions of the framework. These are: (1) we relax the assumption of common prices and probabilities; (2) we discuss when our method can accommodate time-inconsistency; (3) we permit rentiers to earn

risk-free labor income; (4) we introduce an explicit labor–leisure choice; and, (5) we allow mortality risk to vary as a function of calendar time. We end this section by (6) examining how our method performs when non-homotheticity is non-time-separable, using the functional form proposed by Comin et al. (2021) as a test case. Formal statements and proofs for the extensions are in Appendix C.

5.1 Relaxing Common Prices and Probabilities

We assume that, conditional on observables (like age, gender, location) households have the same preference relation within and across cohorts.²³ We also assume that static prices only vary as a function of observables (e.g. time or location). Similarly, cohorts of rentiers at each point in time must hold common beliefs about future prices and rates of return. Beliefs can change over time but, within a period, they can only vary for rentiers as a function of observable characteristics (e.g. age).

However, non-rentiers’ future state variables (i.e. beliefs, prices, returns, borrowing constraints, cash flow) need not be the same as those of the rentiers, nor do they need to be the same for all non-rentiers. For example, rentiers may have access to different assets or hold different beliefs about the returns on those assets than non-rentiers. To understand why, recall that by Lemma 1, relative spending shares on goods in the present only depend on static prices and utility. Therefore, as long as this function is one-to-one, two households facing the same static prices at a point in time choose the same budget shares across goods in the present if, and only if, they are on the same indifference curve.

Finally, we do not require that households’ beliefs about the future be objective. All that matters is that $\pi(\cdot|\tau)$ is the lottery that households in cohort τ believe they face — this may or may not be the result of a rational expectations equilibrium.

5.2 Time-Inconsistent Preferences

In defining the consumers’ value function, we assume that the consumer makes decisions as-if the consumption plan chosen in the first period is followed in subsequent periods.

²³To see what can go wrong, suppose that the preference relation of cohorts is changing over time. Then, changes in consumption-to-wealth ratios or budget shares may be due to changes in preference parameters (e.g. discount factor changes), and not due to changes in welfare-relevant variables like prices, returns, or probabilities. If preference shocks are independent and identically distributed across households, then averaging across households with the same observables eliminates them by the law of large numbers. For this reason, in our empirical applications, we average over households with similar characteristics. However, if the preference shocks are systematic, then this can cause biases. Note that preference stability is a typical maintained assumption even in the literature on static welfare measures. See Baqaee and Burstein (2023) for more details.

However, if preferences are not time-consistent, and the household cannot commit, then the consumer's future selves may not follow the time-0 optimal consumption plan. In Appendix C, we show that as long as time-separability holds, our results can be applied without change, even if household behavior is time inconsistent.

5.3 Rentiers that Earn Risk-free Cash Flows

Rentier households are defined to be those with zero exogenous cash flows: $y(s^j|\tau) = 0$ for every s^j . However, we can expand the set of rentiers to include households with non-zero, time-varying, but risk-free exogenous cash flows, $y(s^j|\tau) = y(j|\tau)$, as well. Examples could include public sector employees, members of teachers' unions, pensioners on defined benefits, and tenured professors. To treat these households as rentiers, we assume that they do not face ad-hoc borrowing constraints and that they can access bonds with maturities $\{1, \dots, J\}$.

Specifically, the first period budget constraint, previously (13), is now

$$\sum_{n \in N} p_n(s^0|\tau) c_n(s^0|\tau) + \sum_{k \in K} a_k(s^0|\tau) + \sum_{l=1}^J B_l(s^0|\tau) = w + y(0|\tau), \quad (20)$$

where $B_l(s^0|\tau)$ is the quantity of bond of maturity $l \in \{1, \dots, J\}$ purchased in period τ , with return $R(l|\tau)$ at date $\tau + l$. At each subsequent history s^j , the budget set is then

$$\sum_{n \in N} p_n(s^j|\tau) c_n(s^j|\tau) + \sum_{k \in K} a_k(s^j|\tau) = \sum_{k \in K} R_k(s^j|\tau) a_k(s^{j-1}|\tau) + B_j(s^0|\tau) R(j|\tau) + y(j|\tau). \quad (21)$$

For simplicity of notation, we assume these bonds are only available to purchase at time τ (allowing access to these bonds after τ does not alter the results).

Under these assumptions, it is straightforward to show that these households' problem is isomorphic to that of a rentier (with exogenous cash flows equal to zero) but where wealth is instead defined to be

$$w(s^0|\tau) + \sum_{j=0}^J \frac{y(j|\tau)}{R(j|\tau)}.$$

That is, households with risk-free exogenous cash flows and no ad-hoc borrowing constraints are isomorphic to rentiers whose wealth is augmented by the present discounted value of exogenous cash flows. We do not implement this extension in our empirical illustration.

5.4 Leisure

Our baseline framework abstracts from labor-leisure choice. For rentiers, this assumption can be replaced by the assumption that rentiers' choose to not work given their labor productivity. In this case, Proposition 5 applies without modification.²⁴

We can allow non-rentiers to make labor-leisure choices and still use Proposition 6 to infer their welfare, if relative static budget shares only depend on utility and static prices of goods and services. This assumption rules out non-separabilities between consumption choices and leisure. For example, consider the utility function

$$U^{\frac{\sigma-1}{\sigma}} = \tilde{P}(c(s^0), U)^{\frac{\sigma-1}{\sigma}} + \tilde{F}(\{c(s^t)\}_{t>0}, \pi, U)^{\frac{\sigma-1}{\sigma}} + \tilde{H}(\{l(s^t)\}_{t \geq 0}, \pi, U),$$

where \tilde{P} and \tilde{F} are aggregators over consumption of goods and services for the present and the future, and \tilde{H} is an aggregator over leisure choices. These preferences have the useful feature that Lemma 1 is still valid and static relative budget shares depend only on static relative prices and U . This means that Proposition 6 can be used without modification. We do not pursue this extension in our empirical application.

5.5 Changes in Mortality

While the baseline model accommodates changes in mortality risk as a function of age, it assumes that mortality risk is fixed as a function of calendar time. Allowing for secular changes in mortality risk can alter the results because changes in welfare caused by changes in mortality risk are not necessarily reflected in changes of the consumption-to-wealth ratio.

To extend the results to allow for exogenous changes in mortality risk, consider the utility function

$$U^{\frac{\sigma-1}{\sigma}} = \lambda_P \tilde{P}(c(s^0), U)^{\frac{\sigma-1}{\sigma}} + \lambda_P \lambda_F \beta \tilde{F}(\{c(s^t)\}_{t>0}, \pi, U)^{\frac{\sigma-1}{\sigma}} + [1 - \lambda_P + \lambda_P(1 - \lambda_F)\beta] \bar{c}^{\frac{\sigma-1}{\sigma}} U^{\frac{\sigma-1}{\sigma} + \varepsilon},$$

where \tilde{P} and \tilde{F} are present and future aggregators that are homogeneous of degree one in quantities (for a given U), parameters λ_P and λ_F are the probabilities of surviving in the present and the future, β is a discount factor, and σ is the constant EIS. Upon death, the household receives a consumption-equivalent payoff of \bar{c} — the lower is \bar{c} , the higher is the

²⁴If rentiers do earn riskfree labor income, as in the extension with risk-free cashflows, then we instead need to assume that hours do not vary as a function of calendar time (they do not have to be zero). This requires assuming that the number of hours worked by rentiers is fixed conditional on age (e.g. a nine-to-five job prior to retirement).

value of statistical life.²⁵ The parameter ε determines wealth effects for the value of life. All of these parameters may vary as a function of observable characteristics (principally, age).

For this problem, the lotteries the household faces are parameterized by the stochastic process over prices and returns, π , as well as mortality risk, λ . Whereas preferences are time-separable in consumption choices, c , and probabilities over market outcomes, π , they are not time-separable in mortality risk λ .

As long as λ is constant as a function of calendar time, our results hold without modification. However, if λ changes over time, then time-separability is violated and our results cannot be used. Survival probabilities are not isomorphic to the other probabilities in the model because households cannot alter the pay-off upon death by saving.

To extend our results to this setting, we again index decision problems by calendar date, τ , which now also indexes survival probabilities, λ_P and λ_F . Following Schelling (1968), denote the value of statistical life for cohort τ to be the marginal willingness to pay for increasing survival probabilities in the present and the future relative to initial wealth w :

$$\Phi_P(\tau, w) = \frac{\partial \log m(\tau, w, \mathbf{y})}{\partial \log \lambda_P}, \quad \text{and} \quad \Phi_F(\tau, w) = \frac{\partial \log m(\tau, w, \mathbf{y})}{\partial \log \lambda_F}.$$

The marginal willingness to pay to reduce the probability of death is an important input into cost-benefit analysis of policies that affect longevity.²⁶ If preferences are non-homothetic, then these elasticities depend on initial wealth.²⁷

The extension of Proposition 5 to this environment is that the money-metric for rentiers, $m(\tau, w, \mathbf{0})$, solves the following fixed point problem:

$$\begin{aligned} \log m(\tau, w, \mathbf{0}) = & \log w - \int_{\tau_0}^{\tau} \left(\sum_{n \in \mathbb{N}} B_n(x, w_x^*) \frac{d \log p_n}{dx} - \frac{1}{1 - \sigma} \frac{d \log B^P(x, w_x^*)}{dx} \right) dx \quad (22) \\ & - \int_{\tau_0}^{\tau} \left(\Phi_P(x, w_x^*) \frac{d \log \lambda_P(x)}{dx} + \Phi_F(x, w_x^*) \frac{d \log \lambda_F(x)}{dx} \right) dx \\ & + \int_{\tau_0}^{\tau} \frac{\sigma}{1 - \sigma} (1 - B^P(x, w_x^*)) \frac{d \log \lambda_F(x)}{dx} dx, \end{aligned}$$

²⁵Some papers effectively assume that $\bar{c}^{(\sigma-1)/\sigma} = 0$, but this is arbitrary. For example, it implies that the value of life switches from positive to negative as σ goes from above one to below one. Because $\bar{c}^{(\sigma-1)/\sigma}$ determines the household's willingness to pay to reduce the probability of death, its value should instead be disciplined by the value of statistical life (e.g. as in Jones and Klenow, 2016). Our formula in equation (22) directly uses information about the value of statistical life, instead of calibrating \bar{c} .

²⁶See Ashenfelter (2006) for a discussion of how such statistics can be estimated.

²⁷If the aggregators P and F are independent of U and $\varepsilon = 0$, then the value of life relative to wealth does not vary as a function of wealth.

where $m(x, w_x^*, \mathbf{0}) = m(\tau, w, \mathbf{0})$ for each $x \in [\tau_0, \tau]$. The first line of (22) is identical to the one in Proposition 5. The remaining two lines adjust for changes in mortality risk. As expected, if mortality risk does not vary as a function of calendar date, $d \log \lambda_F / dx = d \log \lambda_P / dx = 0$, then (22) simplifies to the expression in Proposition 5.

The intuition for the new terms is the following. The second line of (22) adds the value of increased survival to the money-metric. This is very similar to how price changes must be accounted for: we integrate the demand curve for the value of life with respect to changes in mortality risk. The subtlety is that, as with prices, we must integrate compensated demand curves rather than Marshallian demand curves. This implies that there is a fixed point term in the second line unless preferences are homothetic. The final line of (22) accounts for the fact that, holding all else fixed, changes in the future survival probability change the consumption-wealth ratio. Since the second line is fully accounting for the welfare changes caused by changes in the future survival probability, the third line purges from $\frac{1}{1-\sigma} \frac{d \log B^P(x, w_x^*)}{dx}$ those changes caused by changes in future survival probability.

A simplification of (22) results if the utility payoff of death is zero: $\bar{c}^{\frac{\sigma-1}{\sigma}} = 0$. This is because, in this case, the marginal value of increasing survival probability in the future exactly offsets the adjustment in the consumption-wealth ratio:

$$\Phi_F(\tau, w) = \frac{\sigma}{1-\sigma} [1 - B^P(\tau, w)].$$

This means that (22) simplifies to

$$\log m(\tau, w) = \log w - \int_{\tau_0}^{\tau} \left(\sum_{n \in N} B_n(x, w_x^*) \frac{d \log p_n}{dx} - \frac{1}{1-\sigma} \frac{d \log B^P(x, w_x^*)}{dx} + \Phi_P(x, w_x^*) \frac{d \log \lambda_P(x)}{dx} \right) dx,$$

so that changes in the consumption-wealth ratio appropriately account for changes in welfare caused by changes in future survival probabilities. Although assuming $\bar{c}^{\frac{\sigma-1}{\sigma}} = 0$ is common in the literature and simplifies our results, there is no empirically compelling reason to treat this as a benchmark.

To recap, our results can be extended to account for changing mortality provided with additional information on the value of statistical life as a function of wealth and age. In our empirical application, we abstract from these issues and treat mortality risk as constant over time. We emphasize that this does not imply that we take mortality risk to be constant as a function of age. Variation in mortality rates due to changes in age are fully accounted for by our baseline result in Proposition 5. Rather, the assumption we make in our empirical results is that survival probabilities are constant as a function of time conditional on age, and we do not pursue this extension in our empirical illustration.

5.6 Non-Time-Separable Non-Homotheticities

We end this section by considering what happens when non-homotheticities are not time-separable. To do so, we consider the following class of preferences, due to Comin et al. (2021):

$$U = \frac{1}{1-1/\theta} \mathbb{E}_0 \sum_{j=0} \beta^j C_j^{1-\frac{1}{\theta}}, \quad \text{where} \quad C_j = \left[\sum_n \omega_n^{\frac{1}{\gamma}} \left[\frac{c_{nj}}{C_j \epsilon_n} \right]^{\frac{1-\gamma}{\gamma}} \right]^{\frac{\gamma}{\gamma-1}}. \quad (23)$$

These preferences are superficially quite similar to the ones we used in the Monte Carlo exercise in Section 4.3, but there is an important difference. Under these preferences, non-homotheticities within a period are determined by C_j rather than U . These preferences are not time-separable in the sense of Definition 1. Nevertheless, we show that our method continues to perform well even though the assumptions under which we derived our results are violated. Specifically, we prove that when the EIS is close to zero, which is empirically realistic, then the errors in our procedure are close to zero. We then verify this using quantitative Monte Carlo exercises.

We separate our discussion into two parts. First, we discuss the problem of rentiers and non-rentiers theoretically, showing that our method should perform well for realistic calibrations. Second, we provide some Monte Carlo simulations to validate our theoretical results.

Theoretical results. Since these preferences are additively separable, we use the Frisch elasticity of intertemporal substitution, instead of the compensated (Hicksian) elasticity of substitution to define our sufficient statistics. The *Frisch* elasticity of intertemporal substitution (EIS) controls how spending responds to a uniform change in prices, holding the marginal utility of wealth constant. We start with a preliminary result characterizing the Frisch EIS for this class of preferences.

The Frisch EIS, in period t , for the preferences above is

$$\sigma_t(\tau, w, \mathbf{y}) = \left[(1-\gamma) \frac{\text{Var}_{B(s^t)}(\epsilon_n)}{\mathbb{E}_{B(s^t)}[\epsilon_n]^2} + 1 - \frac{\left[1 - \frac{1}{\theta}\right]}{\mathbb{E}_{B(s^t)}[\epsilon_n]} \right]^{-1}. \quad (24)$$

where the variance and expectation use period t budget shares, denoted $B(s^t)$, as weights. Since budget shares vary as a function of (τ, w, \mathbf{y}) , the Frisch EIS also varies as a function of these variables. The Frisch EIS for the first period, period 0, is denoted by $\sigma_0(\tau, w, \mathbf{y})$.

The Frisch EIS is not, in general, equal to the usual parameter θ in (23) because the curvature induced by the ϵ_n 's interacts with intertemporal choices. Note that, in the

special case of homothetic preferences, where $\epsilon_n = 1$ for every n , we recover the usual fact that the Frisch EIS, $\sigma_0(\tau, w, \mathbf{y})$, is identically equal to θ .

Now, we show how our results are altered for rentiers. For simplicity, we assume complete markets, though the result can be extended to allow for incomplete markets along similar lines as in Proposition 5.

Proposition 8. *The money-metric for rentiers with preferences (23) satisfies the following*

$$\log m(\tau, w, \mathbf{0}) = \log w - \int_{\tau_0}^{\tau} \left(\sum_{n \in N} B_n(x, w_x^*, \mathbf{0}) \frac{d \log p_n}{dx} - \frac{d \log B^P(x, w_x^*, \mathbf{0})/dx}{1 - \sigma_0(x, w_x^*, \mathbf{0})} \right) + \text{error } dx.$$

where w_x^* solves the equation $m(x, w_x^*, \mathbf{0}) = m(\tau, w, \mathbf{0})$ for each $x \in [\tau_0, \tau]$. The boundary condition is that $m(\tau_0, w, \mathbf{0}) = w$. The error term shrinks to zero as $\sigma_0 \rightarrow 0$, which happens as $\theta \rightarrow 0$. The error term is

$$\begin{aligned} \text{error} = -\sigma_0 \left[(1 - \gamma) \left(\text{Cov}_{B_0} \left(\frac{\epsilon_i}{\mathbb{E}_{B_0}[\epsilon_i]}, d \log p_{0i} \right) - \mathbb{E}_{\sigma_t}^{NPV} \left[\text{Cov}_{B_t} \left(\frac{\epsilon_i}{\mathbb{E}_B[\epsilon_i]}, d \log p_{ti} \right) \right] \right) \right. \\ \left. + \text{Cov}^{NPV} \left(\frac{\sigma_t}{\bar{\sigma}}, \mathbb{E}_{B_t} \left[d \log \frac{p_t}{1 + r_t} \right] \right) \right]. \end{aligned}$$

See the appendix for more discussion of the error term.

Proposition 8 is the same as Proposition 5 except that there is an error term (which varies by τ and w). The most important observation is that, since the error term is multiplied by σ_0 , as the EIS tends to zero, the error also tends to zero. Moreover, even disregarding the fact that the error term scales in σ_0 , the remaining terms are likely to be quantitatively small in practice.

The first component of the error term depends on the difference between a covariance in the present and the net-present value of future covariances weighted by the Frisch EIS in those periods ($\mathbb{E}_{\sigma_t}^{NPV}$ denotes a net-present value average weighted by σ_t). These are the covariances between the slope of Engel curves, disciplined by $\epsilon_n/\mathbb{E}_B(\epsilon_n)$, and the change in prices, given by $d \log p_n(s^j)/dx$. Hence, if the covariance of price shocks with the slope of the Engel curves is roughly constant, this term is small.

The second component of the error term depends on the covariance of variations in the Frisch EIS (measured by $\sigma_t/\bar{\sigma}$, where $\bar{\sigma}$ is net-present value average of all σ 's) with price and return shocks. This term too is likely to be small in practice, since empirically, the EIS does not vary much (Best et al., 2020). Both of these error terms are multiplied by the Frisch EIS in the present, σ_0 , which is also likely to be small.

In our Monte Carlo exercises, where $\sigma_0 \approx 0.1$, we find that these error terms are extremely small even for very large shocks.

The next proposition shows that Proposition 7 also applies to non-rentiers with Comin et al. (2021) preferences as the Frisch EIS tends to zero.

Proposition 9. *Proposition 7 applies to non-rentiers with preferences defined in Equation (23) when $\sigma_0 = 0$ (which happens if, for example, $\theta \rightarrow 0$).*

Intuitively, when the first period Frisch EIS satisfies $\sigma_0 = 0$, preferences become Leontief across periods and so $C_j = C$ for all j and the common value C is a monotone function of lifetime utility. We leave the proofs and additional discussion in Section C.

Monte Carlo Simulations. To validate these theoretical results, we calibrate the preferences in (23), simulate a finite population, and then apply our method. As in Section 4.3, our simulated sample has two sub-samples: a group of rentiers who face no idiosyncratic risk, and a group of non-rentiers that earn risky labor income. We keep many of the same parameter values as in Section 4.3: the risky income process for non-rentiers has persistence 0.975 and standard deviation of 0.15 in logs, discount factor $\beta = 0.96$, the borrowing constraint is -5 (5 quarters of average labor income), the risk-free rate $R = 1.0075$ (or about 3% at annual frequency), and the within-period elasticity parameter $\gamma = 0.25$. We set the vector of ϵ 's equal to Comin et al. (2021): $\epsilon = (1, 0.8, 0.2)$. The most important parameter is θ , because it has a strong effect on the Frisch EIS. We set it equal to $\theta = 0.1$, which ensures σ_0 varies between 0.06 and 0.1 in the cross-section of households in our simulated sample (using equation 24).

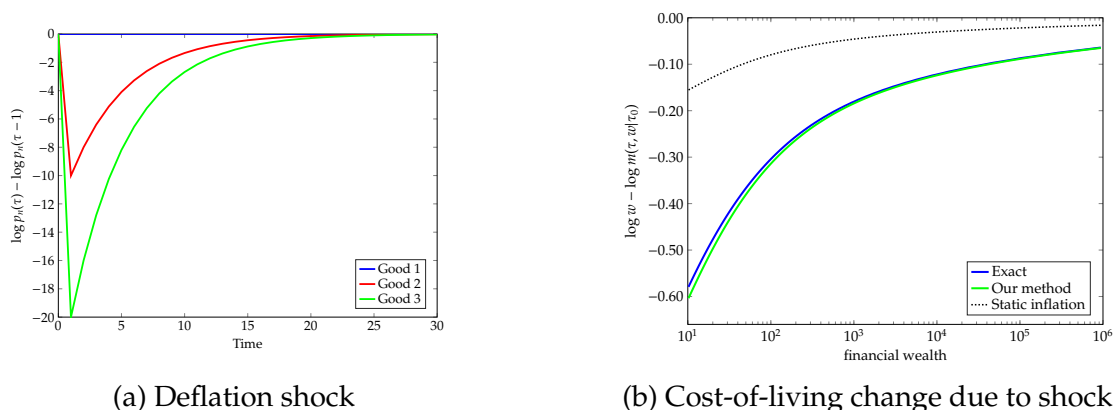


Figure 3: Performance for rentiers with Comin et al. (2021) preferences due to price shock

Figure 3 considers a very large persistent deflationary shock, realized in period $\tau = 1$, to the path of goods prices. The left panel shows the path of inflation, which households

learn about in period $\tau = 1$. The right panel shows for the rentier subsample the change in the cost-of-living in period $\tau = 1$ relative to before the shock is realized in $\tau_0 = 0$ — that is, $\log m(\tau, w|\tau) - \log m(\tau, w|\tau_0) = \log w - \log m(\tau, w|\tau_0)$. Even though the price shocks are correlated with the slope of the Engel curves, the errors are quite small. For comparison, the figure also reports static inflation index between $\tau = 1$ and $\tau_0 = 0$ for different wealth levels. Static deflation is much smaller than the dynamic index, since the static measure does not take into account that the prices will continue to fall for some time in the future.

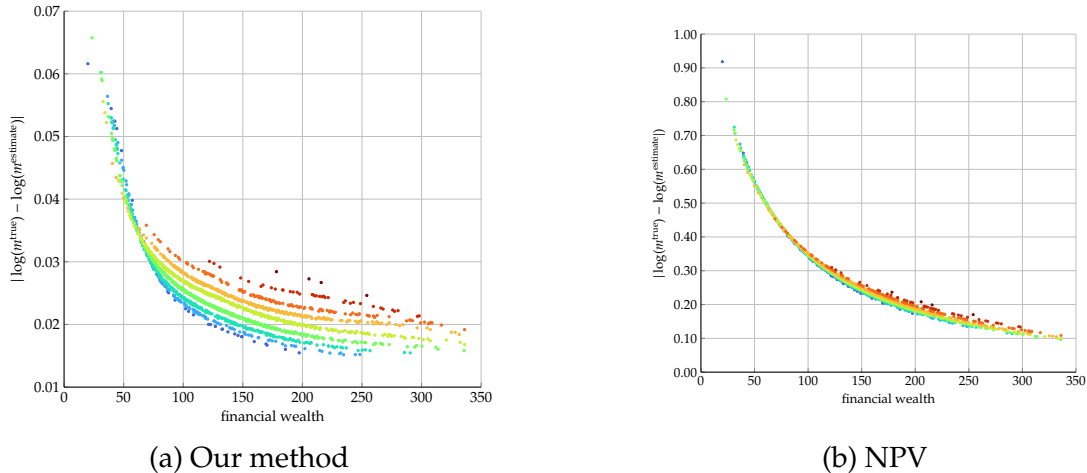


Figure 4: Performance for non-rentiers with Comin et al. (2021) preferences in levels

Note: Colors indicate idiosyncratic income draws.

Figure 4 plots the error in the money metric for non-rentiers according to our method. For comparison, we include the net present value of total wealth, discounted using the risk-free rate, in the right panel. While our method now has errors, the errors are significantly smaller than for the NPV approach. Figure 5 plots the error in the change in the money metric for non-rentiers in response to idiosyncratic income shocks. It is worth noting that the shocks experienced by the non-rentiers are quite large, and a sizeable portion of households receive income shocks that move the exact money-metric utility function by more than 20 log points. Nevertheless, in changes, the errors are small.

6 Illustrative Application to US Data

In this section, we apply our method to estimate the money-metric using data from the US. In the next section, we use these estimates to study the response of welfare to shocks.

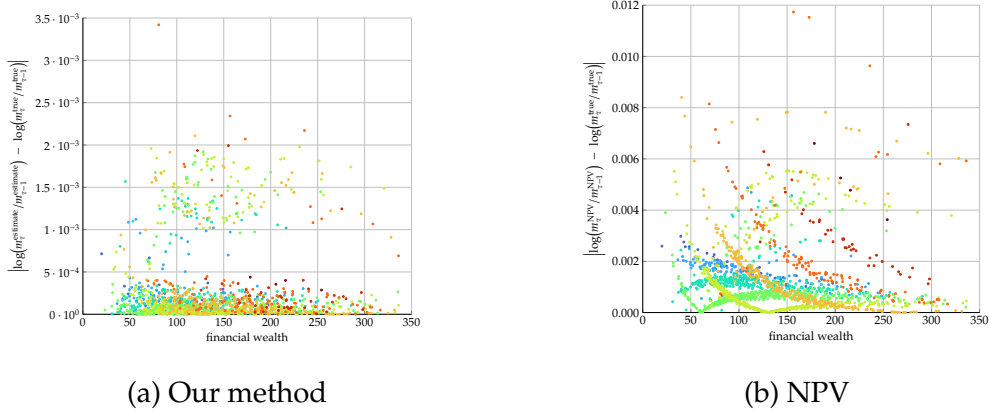


Figure 5: Performance for non-rentiers with Comin et al. (2021) preferences in changes

Note: Colors indicate idiosyncratic income draws.

6.1 Data

We require data on households and on prices. For data on households, we use the Panel Study of Income Dynamics (PSID) spanning the years 2005 to 2019. For the price data, we use consumption price data from the Bureau of Labor Statistics (BLS). We describe each dataset in turn.

The household data must be accurate, but it does not need to be representative of the underlying population in terms of sampling frequency. Therefore, we can use the raw data from the PSID without sampling weights. The PSID contains repeated cross-sectional data on household expenditures by category, household-level balance sheets (assets and liabilities), household incomes, and demographic information. To account for how age affects preferences (including planning horizons), we condition our results on decade of life.²⁸

The PSID includes household expenditure surveys, broken down into seven categories. A major omission is the user cost of owner-occupied housing. To remedy this, we impute equivalent owner-occupied housing costs by matching home owners in each period to renters with similar observable characteristics and spending behavior. That is, we predict rental expenditures using household characteristics and spending behavior using a regression estimated on renters in the same period. This procedure is theoretically justified by Proposition 6. In the final year of our sample, 2019, the PSID asked home owners to report the rental value of their property if they were to rent it out. We use the

²⁸Ideally, with enough observations, we could treat each age separately. Similarly, with more data, we could split the sample along other observed characteristics that may influence preferences, like gender, household size, location, etc.

answers to this question, in 2019, to validate our imputation procedure. When we regress surveyed housing costs on our imputed measure of housing costs (both relative to current expenditures), we find a coefficient of 1.03 and an R^2 value is 0.59. See Appendix D for details.²⁹

We combine the price data from the BLS with the expenditure survey from the PSID via a correspondence between PSID spending categories and categories of goods in the Consumer Price Index (CPI). For more details about how specific variables are constructed, see Appendix D.

6.2 Constructing Wealth and Classifying Rentiers

To apply Proposition 5, we need to observe a sample of rentier households. To classify rentiers, we first estimate a proxy measure of total wealth for all households in the sample. Our proxy for total wealth is the sum of financial wealth (net asset value including home equity and defined contribution pension savings) and the present discounted value of predicted labor and transfer income.

To calculate the present discounted value of labor and transfer income, we predict each household's expected lifetime income profile based on observed characteristics, and discount the resulting flows using a real discount rate of 4% following Catherine et al. (2022). The construction of net assets and capitalized labor and transfer income is detailed in Appendix D.

We say that a household is a rentier if net financial assets constitute more than 90% of proxy total wealth. (Results are similar if we use a 95% threshold to define rentiers). If the head of the household is unemployed and looking for a job, then we exclude this household from the sample of rentiers. Last, to deal with outliers, we exclude from the rentier set those households whose net assets are in the top and bottom 2.5%. Since there are few rentiers younger than 40 years old in our sample, we do not report estimates for households below 40.

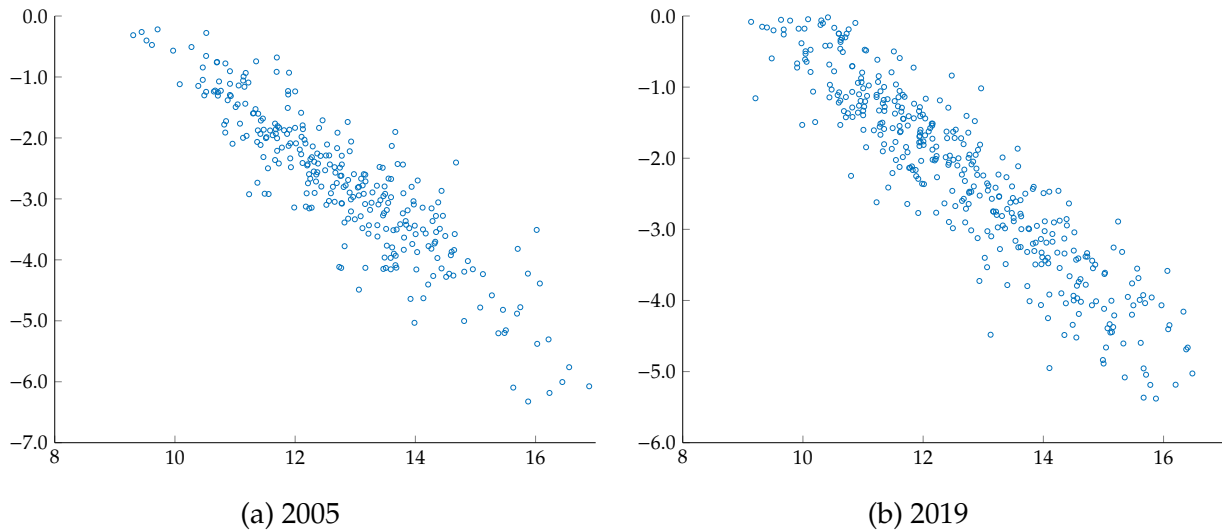
6.3 Money Metric Wealth for Rentiers

Figure 6 shows a scatterplot of log consumption-to-wealth ratios against log wealth for rentiers in the first and last year of the sample. Both panels of Figure 6 show a strong decreasing relationship between consumption-to-wealth ratios and wealth. The relationship is approximately loglinear. This is consistent with the finding in Straub (2019)

²⁹We abstract from other durable consumer goods, for which the user cost would have to be estimated in a similar way.

that the consumption-to-wealth ratio strongly declines in permanent income. Whereas static non-homotheticity is a cornerstone of the literature on consumption, dynamic non-homotheticity, like the one depicted in Figure 6, is relatively understudied. However, as our estimates of money-metric wealth show, in our dataset, dynamic non-homotheticity is an order of magnitude more powerful than static non-homotheticity.

Figure 6: Log consumption-to-wealth against log wealth for rentiers



To compute money-metric values for rentiers using Proposition 5, we need to evaluate consumption-to-wealth ratios and static budget shares as a function of date and wealth for each age group. To do this, we non-parametrically regress the log of consumption-to-wealth ratios on log wealth, age group, and year using kernel regression (`npregress` in STATA). Similarly, for each good $i \in N$, we regress the log of the budget share on a quadratic function of age, log wealth, and year using LASSO. In both cases, the smoothing parameters are chosen by cross-validation.³⁰

We apply Proposition 5 to our estimated cross-sectional curves to recover money-metric utility as a function of date, wealth, and age group.³¹ For illustration, we use the initial year, $t_0 = 2005$, as the base year, so that money-metric values map nominal wealth in each year t into equivalent wealth in 2005. For our benchmark results, we set the EIS, σ ,

³⁰We use `npregress` since it offers a disciplined way to trade off non-parametric fit as a function of age, wealth, and time against overfitting. For the static budget shares, we use LASSO instead of `npregress`, since LASSO is computationally cheaper than `npregress`, and makes bootstrapping much less time consuming. Although we do not report it, our point estimates of the money-metric are very similar for all age groups and wealth levels if we use `npregress` rather than LASSO to fit the static budget shares.

³¹To recover the money-metric, we need to solve the integral equation in Proposition 5. To do so, we use the “recursive” methodology described in Baqaee et al. (2024).

equal to 0.1, the benchmark estimate from Best et al. (2020). Best et al. (2020) estimate that the EIS is relatively homogeneous in the cross-section of households, with point estimates that are uniformly between 0.05 and 0.15 across different quartiles of age and income.

Best et al. (2020) do not estimate Hicksian (compensated) elasticities. Luckily, Slutsky's equation implies that, if consumption is a normal good, then the compensated intertemporal elasticity should be smaller in magnitude than the uncompensated one. We experimented with lower values of σ and found that our results are not sensitive to using values of σ lower than 0.1. Specifically, our results are similar if we set $\sigma = 0.2$ or $\sigma = 0.05$. Of course, the results are highly sensitive to the assumption that σ is not close to one.³²

The left column of Figure 7 plots the money-metric, for each age group, as a function of 2019 wealth for 2005 base prices. The confidence bands are calculated by bootstrap. Since there are fewer young rentiers, sampling uncertainty is higher for younger age groups. For comparison, we also plot a naive calculation that deflates nominal wealth in 2019 by the official CPI inflation between 2005 and 2019. In all cases, the dynamic money-metric wealth is higher than wealth deflated by CPI.

The right column of Figure 7 plots the dynamic annualized inflation rate (log difference between nominal wealth in 2019 dollars and money-metric wealth in 2005 dollars). CPI inflation over this period was 2%, but the dynamic inflation rate is always below 2%. Furthermore, the dynamic inflation rates vary as a function of wealth and age highlighting the importance of non-homotheticity and lifecycle considerations.

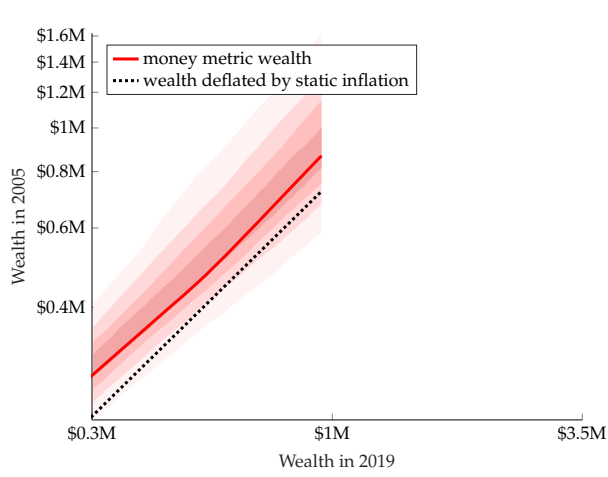
To understand the pattern in Figure 7, we decompose the dynamic inflation rate, the deflator in equation (18), into a static and a forward-looking part. Specifically, for a household with wealth w in 2019, the change in the ideal cost-of-living index between 2005 and 2019 is:

$$\log \frac{w}{m(2019, w, \mathbf{0} | 2005)} = \underbrace{\int_{2005}^{2019} \sum_{n \in N} B_n(x, w_x^*, \mathbf{0}) \frac{d \log p_n}{dx} dx}_{\text{static chained inflation}} + \underbrace{\frac{1}{\sigma - 1} \log \left(\frac{B^P(2019, w, \mathbf{0})}{B^P(2005, w_{2005}^*, \mathbf{0})} \right)}_{\text{future relative to static inflation}}$$

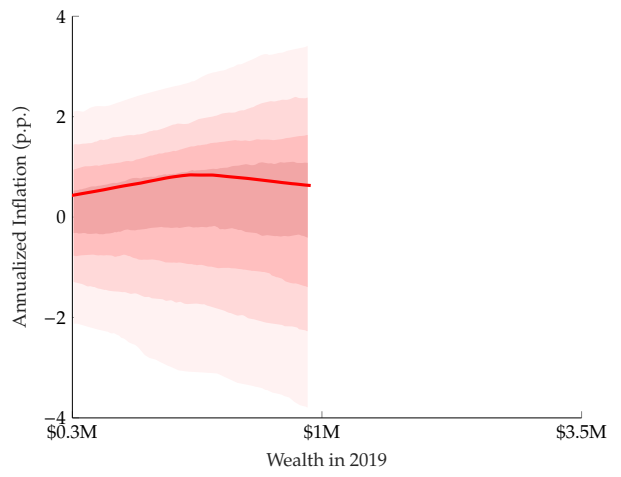
where w_x^* ensures that we are using compensated consumption-to-wealth ratios and budget shares. The first summand is a static measure of inflation. The second summand is expected future inflation relative to present inflation. If the second term is negative, then the rate at which the price of the future bundle changes is lower than the rate at which the

³²Even though in this paper we do not estimate the compensated EIS, we note that in principle it can be estimated using changes in present prices, without knowledge of unobserved future prices and beliefs (for a related discussion, see Proposition 6 in Baqaee et al. (2024).


Figure 7: Conversion of 2019 dollars to 2005 dollars




(a) money-metric 40 – 49 year olds




(b) annualized inflation 40 – 49 year olds



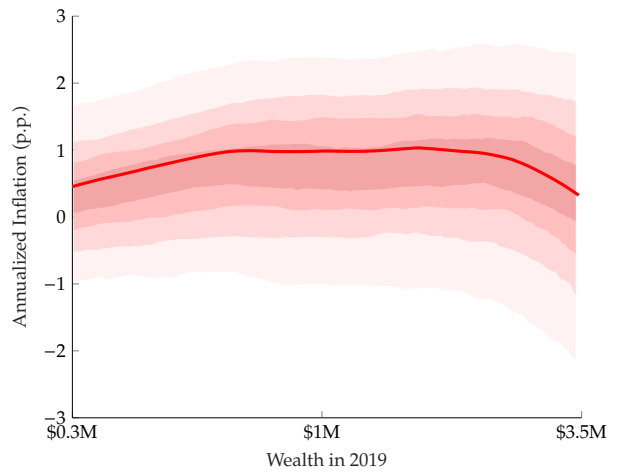
(c) money-metric 50 – 59 year olds



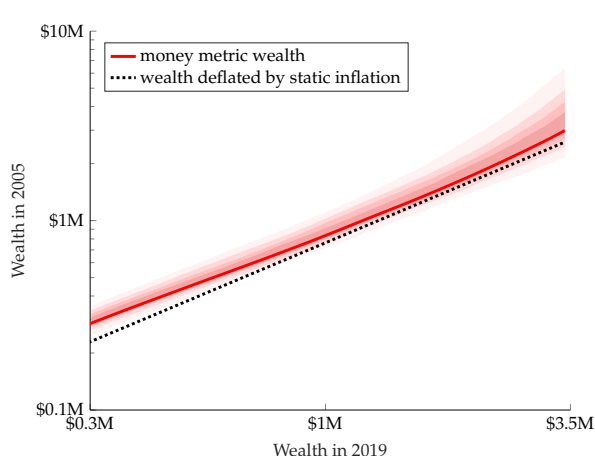
(d) annualized inflation 50 – 59 year olds



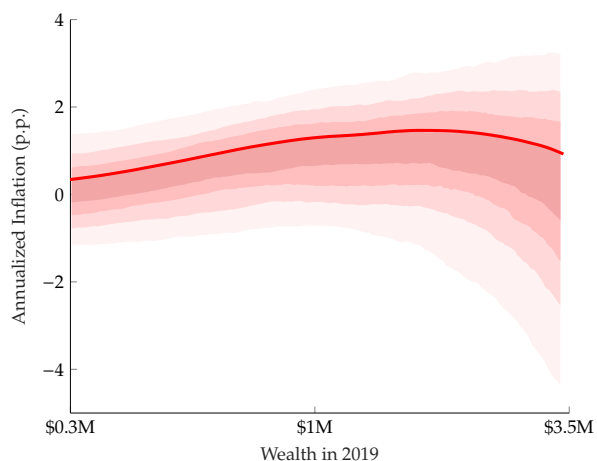
(e) money-metric 60 – 69 year olds



(f) annualized inflation 60 – 69 year olds



(g) money-metric 70 – 79 year olds



(h) annualized inflation 70 – 79 year olds

Notes: Shaded area depicts 90% confidence intervals from bootstrap.

price of the present bundle changes. Therefore, the overall rate of inflation is lower than the static one.

As an example, this decomposition is shown in Figure 8 for the 60 – 69 year old group. The static inflation term is not exactly the same as aggregate CPI because it weighs changes in static prices using compensated budget shares rather than aggregate budget shares. Nevertheless, the static component is very close to aggregate CPI inflation at around 2% per year for all wealth levels. The very slight downward slope reflects the non-homotheticity of static preferences, and static inflation is slightly higher for poorer households, consistent with other studies of the US, like Jaravel and Lashkari (2024), that show that the static cost-of-living index has tended to rise more quickly for poorer households. Nevertheless, the slope of the static inflation line is mild compared to the slope of the dynamic inflation measure.³³

The component corresponding to future inflation in Figure 8 is not zero, which means that future inflation is not equal to static inflation. For all wealth levels, future prices are expected to rise by less than static prices, explaining why the dynamic all-encompassing cost-of-living index is below the static inflation line in Figure 8. Moreover, the future inflation term exhibits more dependence on wealth than the static term.

³³There may be several reasons why the contribution of static non-homotheticity is so mild in our exercise. First, this sample is limited to rentiers — this means that we are looking at a relatively rich set of households compared to studies focusing on static inflation, which typically include very poor households in the sample. Second, we construct a price index as a function of wealth, rather than as a function of current expenditures, as is done in static studies of the cost-of-living. Third, our sample period of fourteen years is reasonably short compared to previous studies, which compute changes over 50 years or longer. Finally, we have only seven spending categories, and static non-homotheticities may be stronger at more disaggregated levels.

The future component of the dynamic inflation measure is proportional to the compensated change in the log consumption-to-wealth ratio. Figure 9 plots both the compensated and uncompensated log change in the consumption-to-wealth ratio between 2005 and 2019 as a function of nominal wealth for 60 – 69 year olds. Since compensated consumption-to-wealth ratios rose for all wealth levels, dynamic inflation is less than static inflation.³⁴

The uncompensated change in the consumption-to-wealth ratio is more positive than the compensated one. This is because there is a strong wealth effect whereby the consumption-to-wealth ratio declines as households become richer. Households with nominal wealth w in 2005 are on a higher indifference curve than households with that same nominal level of wealth in 2019 because of positive inflation. Therefore, the wealth effect means that such households have higher consumption-to-wealth ratios in 2019 than in 2005, even if relative prices do not change. Changes in compensated consumption-to-wealth ratios, which are purged of wealth effects, are lower and reflect only substitution effects. This figure underscores the importance of accounting for wealth effects when using consumption-to-wealth ratios to infer changes in relative prices.³⁵

Our method does not identify which future prices or beliefs are responsible for the patterns in Figure 7. However, the differences in the dynamic measure of inflation need not be caused by differential exposures to future goods prices alone. For example, dynamic inflation is lower than static inflation if there is, for example, an increase in expected future returns. Furthermore, even if all households are symmetrically exposed to future goods prices, the future component of inflation can differ across households because of differences in expected returns of assets (see Fagereng et al., 2022). For example, if poor and rich rentiers or young and old rentiers are differentially reliant on, say, returns to real estate, equities, or bonds to finance their consumption, then changes in returns will differentially affect dynamic inflation rates for these households.

³⁴In Appendix E we also plot the aggregate consumption-to-wealth ratio for our sample. The aggregate patterns over time are very different — whereas the compensated and uncompensated consumption-to-wealth ratios rose between 2005 and 2019 (holding fixed utility and nominal wealth respectively), the aggregate consumption-to-wealth ratio fell between these two dates. This is not surprising. For example, if there is positive economic growth, so households reach higher indifference curves over time, then they move down the consumption-to-wealth Engel curve and aggregate consumption-to-wealth ratio declines, even if the compensated and uncompensated consumption-to-wealth ratio rises.

³⁵Note that changes in the compensated and uncompensated consumption-to-wealth ratios are different to changes in the consumption-to-wealth ratio at the individual level. The reason is that the compensated ratio holds utility constant and the uncompensated ratio holds nominal wealth constant, while at the individual level both are changing over time. That is, consumption-to-wealth ratios could be falling for every household in the sample in the panel, and yet the change in the compensated and uncompensated consumption-to-wealth ratio can be positive for every date, age, and wealth level.

Figure 8: Decomposing inflation for 60-69 year olds

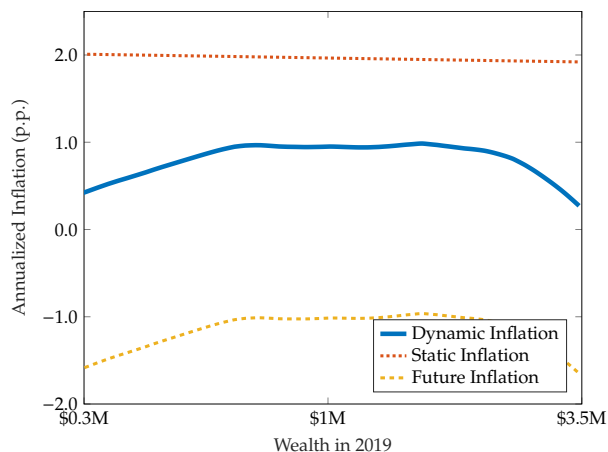
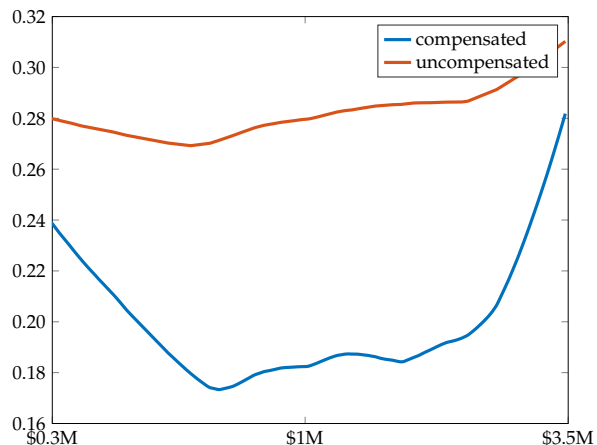


Figure 9: Changes in log consumption-to-wealth ratios for 60 – 69 year olds



6.4 Non-Rentiers

We now turn our attention to the remaining households — the non-rentiers. Proposition 5 does not apply to these households. To recover the money-metric for these households, we rely on Propositions 6 and 7 instead. For each date τ we fit log money-metric values to static budget shares and age group indicators for the rentiers:

$$\log m(\tau, w_{h,\tau}, \mathbf{0}|\tau) = \log w_{h,\tau} = \boldsymbol{\alpha}'_{\tau} \mathbf{X}_{h,\tau} + \text{error}_{h,\tau,\tau} \quad (25)$$

where h indexes the consumer, τ the time period, and $\mathbf{X}_{h,\tau}$ are budget shares and age group. We then use the predicted values from this regression to impute money-metric wealth for the non-rentiers conditional on their budget shares, age, and date.³⁶ We can then convert the imputed wealth into base year dollars using the Proposition 5 deflator, exactly as for rentiers. Appendix E provides figures of the money-metric values for non-rentiers.

This procedure is analogous to Hamilton (2001), and more recently Atkin et al. (2024), who use relative budget shares within a subset of goods, in their case food, to infer changes in welfare in a static context. Unlike Atkin et al. (2024), who compare relative budget shares across time (adjusted for substitution effects) to infer changes in money-metric income over time, we compare relative budget shares within each period across rentier and non-rentier households. Since rentiers and non-rentiers face the same relative prices

³⁶We consider an additional robustness check where we exclude from the set of rentiers the outliers in regression (25). That is, if a potential rentier's predicted and measured total wealth differ significantly, as measured by a Cook's distance value greater than one, then we exclude these households from the set of rentiers. The results are very similar.

at each point in time, we do not have to correct relative budget shares for substitution effects and can deduce money-metric wealth for non-rentiers from the rentiers.

7 Treatment Effect on Welfare

Many policies and shocks affect households along many margins simultaneously. For example, job training programs, changes in tax policy, changes in monetary policy, or job loss all plausibly have dynamic effects on many different relevant variables for households. Del Canto et al. (2023), for example, document that monetary policy shocks affect households through goods price inflation, labor market outcomes, changes in equity prices, house prices, bond prices, and so on.

To understand the welfare effect of such shocks, researchers can estimate the dynamic effects of the shock on each of the different relevant variables and then use changes in those variables, weighted by predicted pre-shock household behavior, to calculate the welfare effect. This is the approach taken by Del Canto et al. (2023). Other than requiring the researcher to enumerate, measure, and estimate all the relevant variables through which the shock affects households, the resulting welfare estimates are first-order approximations around perfect-foresight allocations.

We provide a complementary approach that sidesteps enumerating and estimating all potentially relevant margins. We illustrate this by studying the welfare effect associated with job loss in the PSID. To do so, we regress log money-metric wealth for households on a dummy variable for job loss for the head of the household. Our measure of job loss is equal to one if the head of household loses her job and reports that she is searching for a new job in that period. To control in a simple way for confounds and selection, we include year fixed effects, demographic controls, and lagged log money-metric wealth.

Table 1 reports the results. The outcome is log nominal money-metric wealth, $\log m(\tau, w, y|\tau)$. Column (1) shows that job loss is associated with an approximately 20 log point reduction in nominal money-metric nominal wealth. This effect is statistically significant and economically large. Column (2) shows that the effects are much weaker for older workers (above 60 years old). The impact of job loss on wealth is much smaller than the effect on measures of contemporaneous household income (which is around 85% lower for households that lost their job).³⁷

³⁷An alternative approach in the literature for estimating the dynamic consequences of job loss is due to Davis and Von Wachter (2011). They estimate the present-value of earnings losses after mass-layoff events to be around 12% of counterfactual earnings (using a 5% discount rate). The fact that our point estimates are somewhat different to those in Davis and Von Wachter (2011) is not surprising, because even with the same data it is only under very strong assumptions that our estimates should coincide. For example,

Table 1: Log money-metric wealth and job loss

	log nominal money metric	
	(1)	(2)
Job Loss	-0.197*** (0.031)	-0.218*** (0.034)
Job Loss $\times 1(\text{age} \geq 60)$		0.180** (0.083)
Lagged LHS	Yes	Yes
Controls	Yes	Yes
Observations	48,357	48,357

Standard errors in parentheses.

$p < 0.10$, * $p < 0.05$, *** $p < 0.01$.

Notes: Standard errors are bootstrapped. Controls are year fixed effects, age group, marital status of head of household, industry, and education level.

The identifying variation in these regressions comes from differences in static budget shares between households who lost their job and those who did not, given other controls. That is, the spending patterns of households that lose their job differ from those that did not lose their job in a way that, had they been rentiers, their total wealth would be 20% lower. For older household heads, differences in wealth associated to job loss are smaller compared to younger households.

8 Conclusion

We provide a method for measuring welfare and cost-of-living for households accounting for dynamics, uncertainty, market incompleteness, borrowing constraints, and non-

our calculation does not assume complete markets and if the household's marginal utility is high in states where earnings are low, as is probably realistic, then calculating net present values using a constant discount factor understates the welfare losses of job loss. Furthermore, Davis and Von Wachter (2011) estimate ex-post earnings losses whereas we estimate ex-ante welfare losses, and households' ex-ante beliefs about the consequences of job loss may be more pessimistic than the outcomes Davis and Von Wachter (2011) estimate. Additionally, we do not assume exponential discounting — if households are present-biased, then welfare losses from job loss are amplified since households care more about the near-term, when earnings are low. Finally, they focus on mass lay-off events whereas we consider any job loss. It is plausible that the welfare losses associated with unconditional job loss are different to those caused by mass lay-offs. Despite all these differences, the estimates from Davis and Von Wachter (2011) provide some quantitative context about how the magnitude of our estimates compare to the standard empirical approach.

homotheticities. Our methodology requires repeated household consumption, income, and wealth surveys, as well as prices. The key assumptions we make are that preferences are intertemporally separable and, given observable characteristics, all households have common preferences and face the same prices and beliefs in each period. To calculate money-metrics and cost-of-living, we also require that some subset of households face negligible idiosyncratic undiversifiable income risk (e.g. have negligible risky labor income). Our approach provides a way to non-parametrically measure welfare without fully specifying a structural model. This makes it useful for ex-post analysis of the welfare effect of shocks that have dynamic stochastic effects on many welfare-relevant margins.

References

- Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li (2019). Missing growth from creative destruction. *American Economic Review* 109(8), 2795–2822.
- Alchian, A. A. and B. Klein (1973). On a correct measure of inflation. *Journal of Money, Credit and Banking* 5(1), 173–191.
- Almås, I. (2012). International income inequality: Measuring ppp bias by estimating engel curves for food. *American Economic Review* 102(2), 1093–1117.
- Almås, I., T. K. Beatty, and T. F. Crossley (2018). Lost in translation: What do engel curves tell us about the cost of living?
- Alvarez, F. and U. J. Jermann (2004). Using asset prices to measure the cost of business cycles. *Journal of Political economy* 112(6), 1223–1256.
- Aoki, S. and M. Kitahara (2010). Measuring a dynamic price index using consumption data. *Journal of Money, Credit and Banking* 42(5), 959–964.
- Argente, D., C.-T. Hsieh, and M. Lee (2023). Measuring the cost of living in mexico and the united states. *American Economic Journal: Macroeconomics* 15(3), 43–63.
- Argente, D. and M. Lee (2021). Cost of living inequality during the great recession. *Journal of the European Economic Association* 19(2), 913–952.
- Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012). New trade models, same old gains? *American Economic Review* 102(1), 94–130.
- Ashenfelter, O. (2006). Measuring the value of a statistical life: problems and prospects. *The Economic Journal* 116(510), C10–C23.
- Atkin, D., B. Faber, T. Fally, and M. Gonzalez-Navarro (2024). Measuring welfare and inequality with incomplete price information. *The Quarterly Journal of Economics* 139(1), 419–475.
- Baqee, D. R. and A. Burstein (2023). Welfare and output with income effects and taste shocks. *The Quarterly Journal of Economics* 138(2), 769–834.
- Baqee, D. R., A. T. Burstein, and Y. Koike-Mori (2024). Measuring welfare by matching households across time. *The Quarterly Journal of Economics* 139(1), 533–573.
- Basu, S., L. Pascali, F. Schiantarelli, and L. Serven (2022, 01). Productivity and the Welfare of Nations. *Journal of the European Economic Association* 20(4), 1647–1682.

- Best, M. C., J. S. Cloyne, E. Ilzetzki, and H. J. Kleven (2020). Estimating the elasticity of intertemporal substitution using mortgage notches. *The Review of Economic Studies* 87(2), 656–690.
- Blackorby, C., D. Primont, and R. R. Russell (1998). Separability: A survey. *Handbook of utility theory* 1, 51–92.
- Blaum, J., C. Lelarge, and M. Peters (2018). The gains from input trade with heterogeneous importers. *American Economic Journal: Macroeconomics* 10(4), 77–127.
- Blundell, R. W., M. Browning, and I. A. Crawford (2003). Nonparametric engel curves and revealed preference. *Econometrica* 71(1), 205–240.
- Broda, C. and D. E. Weinstein (2006). Globalization and the gains from variety. *The Quarterly journal of economics* 121(2), 541–585.
- Broda, C. and D. E. Weinstein (2010). Product creation and destruction: Evidence and price implications. *The American economic review* 100(3), 691–723.
- Catherine, S., P. Sodini, and Y. Zhang (2022). Countercyclical income risk and portfolio choices: Evidence from sweden. *Swedish House of Finance Research Paper* (20-20).
- Comin, D., D. Lashkari, and M. Mestieri (2021). Structural change with long-run income and price effects. *Econometrica* 89(1), 311–374.
- Cooper, D., K. E. Dynan, and H. Rhodenhiser (2019). Measuring household wealth in the panel study of income dynamics: The role of retirement assets.
- Costa, D. L. (2001). Estimating real income in the united states from 1888 to 1994: Correcting cpi bias using engel curves. *Journal of political economy* 109(6), 1288–1310.
- Davis, S. J. and T. M. Von Wachter (2011). Recessions and the cost of job loss. Technical report, National Bureau of Economic Research.
- Deaton, A. and J. Muellbauer (1980). *Economics and consumer behavior*. Cambridge university press.
- Del Canto, F. N., J. R. Grigsby, E. Qian, and C. Walsh (2023). Are inflationary shocks regressive? a feasible set approach. Technical report, National Bureau of Economic Research.
- Epstein, L. G. and S. E. Zin (1989). Substitution, risk aversion, and the temporal behavior of consumption. *Econometrica* 57(4), 937–969.
- Fagereng, A., M. Gomez, E. Gouin-Bonenfant, M. Holm, B. Moll, and G. Natvik (2022). Asset-price redistribution. Technical report, Working Paper.
- Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *The American Economic Review*, 157–177.
- Fisher, F. M. and K. Shell (1968). *Taste and quality change in the pure theory of the true cost of living index*. Value, Capital and Growth. Edinburgh University Press.
- Hamilton, B. W. (2001, June). Using engel’s law to estimate cpi bias. *American Economic Review* 91(3), 619–630.
- Huggett, M. and G. Kaplan (2016). How large is the stock component of human capital? *Review of Economic Dynamics* 22, 21–51.
- Hulten, C. R. (1979). On the “importance” of productivity change. *The American economic review* 69(1), 126–136.
- Jaravel, X. and D. Lashkari (2024). Measuring growth in consumer welfare with income-dependent preferences: Nonparametric methods and estimates for the united states. *The Quarterly Journal of Economics* 139(1), 477–532.
- Jones, C. I. and P. J. Klenow (2016). Beyond gdp? welfare across countries and time. *American Economic Review* 106(9), 2426–2457.
- Konüs, A. A. (1939). The problem of the true index of the cost of living. *Econometrica: Journal of the Econometric*

- Society*, 10–29.
- Lucas, R. E. (1987). *Models of business cycles*, Volume 26. Basil Blackwell Oxford.
- Lucas, R. E. (2003). Macroeconomic priorities. *American economic review* 93(1), 1–14.
- Nakamura, E., J. Steinsson, and M. Liu (2016). Are chinese growth and inflation too smooth? evidence from engel curves. *American Economic Journal: Macroeconomics* 8(3), 113–44.
- Pallotti, F., G. Paz-Pardo, J. Slacalek, O. Tristani, and G. L. Violante (2024). Who bears the costs of inflation? euro area households and the 2021–2023 shock. *Journal of Monetary Economics* 148, 103671.
- Pollack, R. A. (1975). The intertemporal cost of living index. In *Annals of Economic and Social Measurement*, Volume 4, number 1, pp. 179–198. NBER.
- Reis, R. (2005). A dynamic measure of inflation. Working Paper 11746, National Bureau of Economic Research.
- Samuelson, P. A. (1961). The evaluation of ‘social income’: Capital formation and wealth. In *The Theory of Capital: Proceedings of a Conference held by the International Economic Association*, pp. 32–57. Springer.
- Schelling, T. C. (1968). The life you save may be your own. *Problems in public expenditure*, 127–162.
- Straub, L. (2019). Consumption, savings, and the distribution of permanent income. *Unpublished manuscript*, Harvard University.

A Non-Separable Preferences

Example 1 provides examples of preferences that satisfy the time separability condition in Definition 1. In this appendix, we provide some examples that do not satisfy this condition.

Example 2 (Examples Violating Time-Separability). Dynamic Stone-Geary:

$$U = \frac{(c(s^0) + a_0)^{1-\sigma}}{1-\sigma} + \sum_{j=1}^J \beta^j \frac{(c(s^j) + a_j)^{1-\sigma}}{1-\sigma},$$

and dynamic Addilog:

$$U = \frac{c(s^0)^{1-\sigma_0}}{1-\sigma_0} + \sum_{j=1}^J \beta^j \frac{c(s^j)^{1-\sigma_j}}{1-\sigma_j},$$

are not time-separable unless $a_j = 0$ and $\sigma_j = \sigma$ for every j .

The direct utility function of these preferences is separable across time, in the sense that utility can be written as a function of a present block and a future block. That is, one can write $U = \tilde{D}(\tilde{P}(c(s^0)), \tilde{F}(\{c(s^j)\}_{j>0}))$. However, these preferences are not time-separable in the sense of Definition 1 since the functions \tilde{D} , \tilde{P} , and \tilde{F} will not be homogeneous of degree one.

In the language of utility theory from Blackorby et al. (1998), the examples above are *directly separable* but they are not *separable in the expenditure/distance function*. The different notions of separability—separability in the direct utility function, separability in the indirect utility function, and separability in the expenditure/distance function—do not nest each other and have different implications for separability of the Frisch, Marshallian, and Hicksian demand systems, respectively. They do, however, coincide if, and only if, preferences are homothetic.

Although the examples above are not time separable in the sense defined in Definition 1, they can be modified to become time-separable. For example, the following “pseudo” Stone-Geary functional form is time separable:

$$U = \frac{(c(s^0) + a_0)^{1-\sigma}}{1-\sigma} + \frac{\left[\left(\sum_{j=1}^J \beta_j c(s^j)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} + a_1 \right]^{1-\sigma}}{1-\sigma}.$$

Similarly, the following “pseudo” Addilog functional form is also time-separable:

$$U = \frac{c(s^0)^{1-\sigma_0}}{1-\sigma_0} + \frac{F\left(\left\{c(s^j)\right\}_{j>0}\right)^{1-\sigma_1}}{1-\sigma_1},$$

as long as F is homogeneous of degree one.

Another example that does not satisfy time separability in Definition 1 is the intertemporal non-homothetic CES preferences studied in Comin et al. (2021). We discuss the performance of our method under these preferences in Section 5.6.

B Proofs

Proof of Proposition 2. Since $\partial V/\partial w > 0$ as long as $p(\tau_0) \neq \mathbf{0}$ and $\mathbf{R}(\tau_0) > 0$, u is monotone increasing in V . ■

Proofs of Proposition 3 & Proposition 4. These are corollaries of Proposition 5. ■

Proof of Proposition 5. This proof is lengthy, so we break it down into several steps.

1. We define the expenditure function, and discuss the implications of time separability.
2. We prove a series of lemmas about the expenditure function. For rentiers, the expenditure function can be used to calculate the money metric utility function.
3. Finally, we combine the lemmas to prove Proposition 5.

A key challenge, which lengthens the proof, is that because financial markets may be incomplete, we cannot rely on the existence of a single intertemporal budget constraint to solve for money-metric utility.

Throughout this proof, we use a familiar transformation of the utility function. Define the *shadow* intertemporal expenditure function to be

$$e(\mathbf{q}, \boldsymbol{\pi}, U) = \min_c \{\mathbf{q} \cdot \mathbf{c} : \mathcal{U}(\mathbf{c}, \boldsymbol{\pi}) = U\}, \quad (26)$$

where \mathbf{q} has the same dimensionality as the state-contingent commodity space. We refer to e as the *shadow* intertemporal expenditure function and to \mathbf{q} as *shadow* prices. We use the qualifier “shadow” to emphasize that e is a purely theoretical construct and agents need not be solving the expenditure minimization problem defined in (26) in practice.

Note that Definition 1 is equivalent to the following.

Proposition 10 (Time-Separability of Expenditure Function). *Suppose that the preference relation \succeq is time-separable in the sense of Definition 1. Then the shadow intertemporal expenditure function must also be time separable in the sense that it can be written as*

$$e(\mathbf{q}, \boldsymbol{\pi}, U) = \tilde{e}\left(\tilde{P}\left(\mathbf{q}(s^0), U\right), \tilde{F}\left(\left\{\mathbf{q}(s^t)\right\}_{t>0}, \boldsymbol{\pi}, U\right), U\right), \quad (27)$$

where \tilde{e} , \tilde{P} , and \tilde{F} are scalar-valued functions. The function \tilde{e} is increasing in all three arguments and homogeneous of degree one in the first two arguments. The functions \tilde{P} and \tilde{F} are increasing and homogeneous of degree one in \mathbf{q} , and non-decreasing in U .

This proposition is a consequence of Theorem 4.3 in Blackorby et al. (1998).

To prove Proposition 5, we use the shadow intertemporal expenditure function, and proceed by first proving a series of lemmas.

Since each decision problem is indexed by (τ, w, \mathbf{y}) , denote consumption of good n in history s^t by $c_n(s^t|\tau, w, \mathbf{y})$. The next lemma shows that for every decision problem (τ, w, \mathbf{y}) , there exists a set of shadow prices $\mathbf{q}^*(\tau, w, \mathbf{y})$ that rationalize the allocations $\mathbf{c}(\cdot|\tau, w, \mathbf{y})$.

Lemma 2 (Dual Problem). *There exist $\mathbf{q}^*(\tau, w, \mathbf{y})$ such that, for every s^t and n ,*

$$c_n^*(s^t|\mathbf{q}^*, \boldsymbol{\pi}, V(\tau, w, \mathbf{y})) = c_n(s^t|\tau, w, \mathbf{y}),$$

where $c_n^*(s^t|\mathbf{q}^*, \boldsymbol{\pi}, V(\tau, w, \mathbf{y}))$ are quantities that minimize the shadow expenditure function (26). Moreover, we can set shadow prices for goods in the first period equal to their observed prices:

$$q_n^*(s^0|\tau, w, \mathbf{y}) = p_n(s^0|\tau),$$

for every $n \in N$.

Proof of Lemma 2. The existence of \mathbf{q}^* follows from the separating hyperplane theorem, since the constraint set and indifference curves are both convex (the constraint set is an intersection of convex sets). We can provide a more constructive proof by using the Karush-Kuhn-Tucker (KKT) conditions. Since the expenditure minimization problem is convex, the KKT conditions must be satisfied. The Lagrangian for households is:

$$\mathcal{L}(\mathbf{p}, \mathbf{R}, \mathbf{y}, \boldsymbol{\pi}, w) = \mathcal{U}(\mathbf{c}, \boldsymbol{\pi}) - \lambda(s^0|\tau) \left[\sum_{n \in N} p_n(s^0|\tau) c_n(s^0|\tau) + \sum_{k \in K} a_k(s^0|\tau) - w \right]$$

$$\begin{aligned}
& + \sum_{s^t} \lambda(s^t|\tau) \left[\sum_{n \in N} p_n(s^t|\tau) c_n(s^t|\tau) + \sum_{k \in K} a_k(s^t|\tau) - \sum_{k \in K} R_k(s^t|\tau) a_k(s^{t-1}|\tau) + y(s^t|\tau) \right] \\
& - \sum_{s^t} \mu(s^t|\tau) \left[\sum_k a_k(s^t|\tau) - X(s^t|\tau) \right] \\
& = \mathcal{U}(c, \pi) + \lambda(s^0|\tau)w + \sum_{s^t} \lambda(s^t|\tau) y(s^t|\tau) \\
& - \sum_{s^t=s^0}^{s^T} \lambda(s^t|\tau) \sum_{n \in N} p_n(s^t|\tau) c_n(s^0|\tau) \\
& - \sum_{s^t=s^0}^{s^T} \lambda(s^t|\tau) \sum_{k \in K} a_k(s^t|\tau) + \sum_{s^t} \lambda(s^t|\tau) \sum_{k \in K} R_k(s^t|\tau) a_k(s^{t-1}|\tau) \\
& - \sum_{s^t} \mu(s^t|\tau) \sum_k a_k(s^t|\tau) + \sum_{s^t} \mu(s^t|\tau) X(s^t|\tau)
\end{aligned}$$

The first order conditions for asset holdings are

$$- \left[\lambda(s^t|\tau) + \mu(s^t|\tau) \right] = \sum_{s^{t+1}} \lambda(s^{t+1}|\tau) R_k(s^{t+1}|\tau)$$

Substituting this back in, we get that the Lagrangian is equal to

$$\mathcal{L}(p, R, y, \pi, w) = \mathcal{U}(c, \pi) + \lambda(s^0|\tau)w + \sum_{s^t} \lambda(s^t|\tau) y(s^t|\tau) - \sum_{s^t=s^0}^{s^T} \lambda(s^t|\tau) \sum_{n \in N} p_n(s^t|\tau) c_n(s^0|\tau) + \sum_{s^t} \mu(s^t|\tau) X(s^t|\tau).$$

Define the indirect utility function to be v that satisfies this equation:

$$e(q, \pi, v) = W.$$

From standard duality, we know that we can also write

$$v(q, \pi, W) = \max_c \{ \mathcal{U}(c, \pi) : q \cdot c = W \}.$$

Call the maximizers above $c^{**}(q, \pi, W)$. The Lagrangian for intertemporal indirect utility function is

$$\mathcal{L}^{**}(q, \pi, W) = \mathcal{U}(\{c, \pi\}) - \mu [q \cdot c - W].$$

Set

$$q_n(s^t) = \frac{\lambda(s^t|\tau)}{\lambda(s^0|\tau)} p_n(s^t|\tau)$$

and

$$W = w + \sum_{s^t} \frac{\lambda(s^t|\tau)}{\lambda(s^0|\tau)} y(s^t|\tau) + \sum_{s^t} \frac{\mu(s^t|\tau)}{\lambda(s^0|\tau)} X(s^t|\tau).$$

Hence

$$\mathcal{L}^{**}(\mathbf{q}, \boldsymbol{\pi}, W) = \mathcal{U}(\{\mathbf{c}, \boldsymbol{\pi}\}) + \mu \left[w + \sum_{s^t} \frac{\lambda(s^t|\tau)}{\lambda(s^0|\tau)} y(s^t|\tau) + \sum_{s^t} \frac{\mu(s^t|\tau)}{\lambda(s^0|\tau)} X(s^t|\tau) - \sum_{s^t} \sum_{n \in N} \frac{\lambda(s^t|\tau)}{\lambda(s^0|\tau)} p_n(s^t|\tau) c_n(s^t|\tau) \right].$$

These problems have the same solution because the Lagrangian is the same. Hence

$$\mathbf{c}^{**}(\mathbf{q}, \boldsymbol{\pi}, W) = \mathbf{c}(s^t|\tau, w, \mathbf{y}),$$

where $q_n(s^t) = \lambda(s^t|\tau) p_n(s^t|\tau)$ and $W = \lambda(s^0|\tau)w + \sum_{s^t} \lambda(s^t|\tau) y(s^t|\tau) + \sum_{s^t} \mu(s^t|\tau) X(s^t|\tau)$. By standard duality arguments, we also know that

$$\mathbf{c}^{**}(\mathbf{q}, \boldsymbol{\pi}, W) = \mathbf{c}^*(\mathbf{q}, \boldsymbol{\pi}, v(\mathbf{q}, \boldsymbol{\pi}, W)) = \mathbf{c}^*(\mathbf{q}, \boldsymbol{\pi}, V(\mathbf{q}, \boldsymbol{\pi}, W)).$$

■

Next, define the following function, called the compensated budget share of n :

$$b_n(\mathbf{q}(s^0), U) = \frac{c_n(s^0) q_n(s^0)}{e(\mathbf{q}, \boldsymbol{\pi}, U) b^P(\mathbf{q}, \boldsymbol{\pi}, U)},$$

where we suppress the dependence of $c_n(s^0)$ on $(\mathbf{q}(s^0), U)$.

Lemma 3. *If preferences are time separable, then the following holds*

$$b^P(\mathbf{q}, \boldsymbol{\pi}, U) \equiv \sum_{n \in N} \frac{c_n(s^0) q_n(s^0)}{e(\mathbf{q}, \boldsymbol{\pi}, U)} = \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \tilde{P}},$$

$$b^F(\mathbf{q}, \boldsymbol{\pi}, U) \equiv 1 - b^P(\mathbf{q}, \boldsymbol{\pi}, U) = \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \tilde{F}},$$

and

$$\frac{\partial \log \tilde{P}}{\partial \log q_n(s^0)} = \frac{c_n(s^0) q_n(s^0)}{e(\mathbf{q}, \boldsymbol{\pi}, U) b^P(\mathbf{q}, \boldsymbol{\pi}, U)}.$$

Proof. By the envelope theorem,

$$\frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)} = \frac{c_n(s^0) q_n(s^0)}{e(\mathbf{q}, \boldsymbol{\pi}, U)},$$

and for $t > 0$

$$\frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^t)} = \frac{c_n(s^t)q_n(s^t)}{e(\mathbf{q}, \boldsymbol{\pi}, U)}.$$

Then, we know that

$$b^P(\mathbf{q}, \boldsymbol{\pi}, U) = \sum_{n \in N} \frac{c_n(s^0)q_n(s^0)}{e(\mathbf{q}, \boldsymbol{\pi}, U)} = \sum_{n \in N} \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)} = \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \tilde{P}} \sum_{n \in N} \frac{\partial \log \tilde{P}}{\partial \log q_n(s^0)} = \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \tilde{P}}$$

and

$$\begin{aligned} b^F(\mathbf{q}, \boldsymbol{\pi}, U) &= \sum_{s^t | t > 0} \sum_{n \in N} \frac{c_n(s^t)q_n(s^t)}{e(\mathbf{q}, \boldsymbol{\pi}, U)} = \sum_{s^t | t > 0} \sum_{n \in N} \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^t)} \\ &= \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \tilde{F}} \sum_{s^t | t > 0} \sum_{n \in N} \frac{\partial \log \tilde{F}}{\partial \log q_n(s^0)} = \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \tilde{F}} \end{aligned}$$

where the last steps use homogeneity of degree 1 in \mathbf{q} of \tilde{P} and \tilde{F} .

Next, we show that

$$\frac{\partial \log \tilde{P}}{\partial \log q_n(s^0)} = b_n(\mathbf{q}(s^0), U).$$

To do this, use the following equality,

$$\begin{aligned} \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)} &= \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \tilde{P}} \frac{\partial \log \tilde{P}}{\partial \log q_n(s^0)} \\ &= \frac{c_n(s^0)q_n(s^0)}{e(\mathbf{q}, \boldsymbol{\pi}, U)} \\ &= b^P(\mathbf{q}, \boldsymbol{\pi}, U) \frac{\partial \log \tilde{P}}{\partial \log q_n(s^0)}. \end{aligned}$$

Rearranging yields

$$\frac{\partial \log \tilde{P}}{\partial \log q_n(s^0)} = \frac{c_n(s^0)q_n(s^0)}{e(\mathbf{q}, \boldsymbol{\pi}, U)b^P(\mathbf{q}, \boldsymbol{\pi}, U)}.$$

■

Lemma 4. *When preferences are time separable, the elasticity of intertemporal substitution*

$$\sigma^*(\mathbf{q}, \boldsymbol{\pi}, U) \equiv 1 - \sum_{n \in N} \frac{\partial \log b^P(\mathbf{q}, \boldsymbol{\pi}, U)/b^F(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)}.$$

is given by

$$\sigma^*(\mathbf{q}, \boldsymbol{\pi}, U) = 1 - \frac{\partial^2 \log e / (\partial \log \tilde{P})^2}{b^F(\mathbf{q}, \boldsymbol{\pi}, U) b^P(\mathbf{q}, \boldsymbol{\pi}, U)}.$$

Proof. We start with

$$\begin{aligned} \frac{\partial \log b^P(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)} &= \frac{1}{b^P(\mathbf{q}, \boldsymbol{\pi}, U)} \frac{\partial}{\partial \log q_n(s^0)} \left[\sum_{k \in N} \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \tilde{P}} \frac{\partial \log \tilde{P}}{\partial \log q_k(s^0)} \right], \\ &= \frac{1}{b^P(\mathbf{q}, \boldsymbol{\pi}, U)} \frac{\partial}{\partial \log q_n(s^0)} \left[\sum_{k \in N} \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \tilde{P}} b_k(\mathbf{q}, U) \right], \\ &= \frac{1}{b^P(\mathbf{q}, \boldsymbol{\pi}, U)} \left[\sum_{k \in N} \frac{\partial}{\partial \log q_n(s^0)} \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \tilde{P}} b_k(\mathbf{q}, U) + \sum_{k \in N} \frac{\partial \log e}{\partial \log \tilde{P}} \frac{\partial b_k(\mathbf{q}, U)}{\partial \log q_n(s^0)} \right], \\ &= \frac{1}{b^P(\mathbf{q}, \boldsymbol{\pi}, U)} \left[\sum_{k \in N} \frac{\partial^2 \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{(\partial \log \tilde{P})^2} b_n(\mathbf{q}, U) b_k(\mathbf{q}, U) + \frac{\partial \log e}{\partial \log \tilde{P}} \frac{\sum_{k \in N} \partial b_k(\mathbf{q}, U)}{\partial \log q_n(s^0)} \right], \\ &= \frac{1}{b^P(\mathbf{q}, \boldsymbol{\pi}, U)} \left[\frac{\partial^2 \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{(\partial \log \tilde{P})^2} b_n(\mathbf{q}, U) \sum_{k \in N} b_k(\mathbf{q}, U) \right], \\ &= \frac{1}{b^P(\mathbf{q}, \boldsymbol{\pi}, U)} \frac{\partial^2 \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{(\partial \log \tilde{P})^2} b_n(\mathbf{q}, U). \end{aligned}$$

Summing over all $n \in N$ yields

$$\begin{aligned} \sum_n \frac{\partial \log b^P(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)} &= \frac{1}{b^P(\mathbf{q}, \boldsymbol{\pi}, U)} \frac{\partial^2 \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{(\partial \log \tilde{P})^2} \sum_n b_n(\mathbf{q}, U), \\ &= \frac{1}{b^P(\mathbf{q}, \boldsymbol{\pi}, U)} \frac{\partial^2 \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{(\partial \log \tilde{P})^2}. \end{aligned}$$

Since $b^P + b^F = 1$, we have that

$$\frac{\partial \log b^F(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)} = - \frac{b^P(\mathbf{q}, \boldsymbol{\pi}, U)}{b^F(\mathbf{q}, \boldsymbol{\pi}, U)} \frac{\partial \log b^P(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)}.$$

By definition,

$$\begin{aligned} 1 - \sigma(\mathbf{q}, \boldsymbol{\pi}, U) &= \sum_n \frac{\partial \log b^P(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)} - \sum_n \frac{\partial \log b^F(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)}, \\ &= \frac{1}{b^P(\mathbf{q}, \boldsymbol{\pi}, U)} \frac{\partial^2 \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{(\partial \log \tilde{P})^2} + \frac{b^P(\mathbf{q}, \boldsymbol{\pi}, U)}{b^F(\mathbf{q}, \boldsymbol{\pi}, U)} \frac{\partial \log b^P(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)}, \end{aligned}$$

$$= \frac{1}{b^F(\mathbf{q}, \boldsymbol{\pi}, U)b^P(\mathbf{q}, \boldsymbol{\pi}, U)} \frac{\partial^2 \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{(\partial \log \tilde{P})^2}.$$

■

Lemma 5. *When preferences are time separable, the following equation holds:*

$$\frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \mathbf{q}} d \log \mathbf{q} + \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \boldsymbol{\pi}} d \boldsymbol{\pi} = -\frac{d \log b^P(\mathbf{q}, \boldsymbol{\pi}, U)}{1 - \sigma^*(\mathbf{q}, \boldsymbol{\pi}, U)} + \sum_{n \in \mathbb{N}} b_n(\mathbf{q}(s^0), U) d \log q_n(s^0)$$

Proof. From Lemma 3, we know that

$$\begin{aligned} \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \mathbf{q}} d \log \mathbf{q} + \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \boldsymbol{\pi}} d \boldsymbol{\pi} &= b^P(\mathbf{q}, \boldsymbol{\pi}, U) \sum_{n \in \mathbb{N}} b_n(\mathbf{q}, U) d \log q(s^0) \\ &+ b^F(\mathbf{q}, \boldsymbol{\pi}, U) \sum_{s^t | t > 0} \left(\sum_{n \in \mathbb{N}} \frac{\partial \log F}{\partial \log q_n(s^t)} d \log q_n(s^t) + \frac{\partial \log F}{\partial \boldsymbol{\pi}(s^t)} d \boldsymbol{\pi}(s^t) \right). \end{aligned}$$

Next, from homogeneity of degree one, we know that

$$\frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \tilde{P}} + \frac{\partial \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \tilde{F}} = 1.$$

Differentiating this identity with respect to P and F yields the following equation

$$\frac{\partial^2 \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{(\partial \log \tilde{P})^2} = -\frac{\partial^2 \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{\partial \log \tilde{P} \partial \log \tilde{F}} = \frac{\partial^2 \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{(\partial \log \tilde{F})^2}.$$

Hence, fixing utility, the total derivative of $b^P(\mathbf{q}, \boldsymbol{\pi}, U)$ with respect to \mathbf{q} and $\boldsymbol{\pi}$ is

$$\begin{aligned} b^P d \log b^P(\mathbf{q}, \boldsymbol{\pi}, U) &= \frac{\partial^2 \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{(\partial \log \tilde{P})^2} \sum_{n \in \mathbb{N}} \frac{\partial \log \tilde{P}}{\partial \log q_n(s^0)} d \log q_n(s^0) \\ &+ \frac{\partial^2 \log e}{\partial \log \tilde{F} \partial \log \tilde{P}} \sum_{s^t | t > 0} \left(\sum_{n \in \mathbb{N}} \frac{\partial \log \tilde{F}}{\partial \log q_n(s^t)} d \log q_n(s^t) + \frac{\partial \log \tilde{F}}{\partial \boldsymbol{\pi}(s^t)} d \boldsymbol{\pi}(s^t) \right) \\ &= \frac{\partial^2 \log e(\mathbf{q}, \boldsymbol{\pi}, U)}{(\partial \log \tilde{P})^2} \left[\begin{array}{c} \sum_{n \in \mathbb{N}} b_n(\mathbf{q}, U) d \log q_n(s^0) - \\ \sum_{s^t | t > 0} \left(\sum_{n \in \mathbb{N}} \frac{\partial \log \tilde{F}}{\partial \log q_n(s^t)} d \log q_n(s^t) + \frac{\partial \log \tilde{F}}{\partial \boldsymbol{\pi}(s^t)} d \boldsymbol{\pi}(s^t) \right) \end{array} \right]. \end{aligned} \tag{28}$$

From Lemma 3 and Lemma 4, we can rewrite this as

$$\frac{d \log b^P(\mathbf{q}, \boldsymbol{\pi}, U)}{(1 - \sigma^*(\mathbf{q}, \boldsymbol{\pi}, U))} = (1 - b^P(\mathbf{q}, \boldsymbol{\pi}, U)) \left[\frac{\sum_{n \in N} b_n(\mathbf{p}, U) d \log q_n(s^0) -}{\sum_{s^t | t > 0} \left(\sum_{n \in N} \frac{\partial \log \tilde{F}}{\partial \log q_n(s^t)} d \log q_n(s^t) + \frac{\partial \log \tilde{F}}{\partial \pi(s^t)} d \pi(s^t) \right)} \right],$$

Rearranging this gives

$$b^P(\mathbf{q}, \boldsymbol{\pi}, U) \sum_{n \in N} b_n(\mathbf{p}, U) d \log q_n(s^0) - b^F(\mathbf{q}, \boldsymbol{\pi}, U) \times \sum_{s^t | t > 0} \left(\sum_{n \in N} \frac{\partial \log \tilde{F}}{\partial \log q_n(s^t)} d \log q_n(s^t) + \frac{\partial \log \tilde{F}}{\partial \pi(s^t)} d \pi(s^t) \right) = - \frac{d \log b^P(\mathbf{q}, \boldsymbol{\pi}, U)}{1 - \sigma^*(\mathbf{q}, \boldsymbol{\pi}, U)} + \sum_{n \in N} b_n(\mathbf{q}(s^0), U) d \log q_n(s^0).$$

Plug this back into (28) to get the desired result. ■

Lemma 6. *The shadow prices $\mathbf{q}^*(\tau, w, \mathbf{0})$ can be written as a function of τ and $V(\tau, w, \mathbf{0})$. That is, we can write*

$$\mathbf{q}^*(\tau, w, \mathbf{0}) = \mathbf{q}^*(\tau, V(\tau, w, \mathbf{0})).$$

Furthermore,

$$V(\tau, w, \mathbf{0}) = v(\mathbf{q}^*(\tau, V(\tau, w, \mathbf{0})), \boldsymbol{\pi}(\cdot | \tau), w)$$

for every τ and w .

Proof. The first part follows from the fact that the value function $V(\tau, w, \mathbf{0})$ is monotone in w . Hence, we can substitute the inverse of $V(\tau, w, \mathbf{0})$ with respect to w into $\mathbf{q}^*(\tau, w, \mathbf{0})$ to get $\mathbf{q}^*(\tau, V(\tau, w, \mathbf{0})) = \mathbf{q}^*(\tau, V^{-1}(V(\tau, w, \mathbf{0})), \mathbf{0})$.

For the second part, we know from Proposition 2, that

$$c(\tau, w, \mathbf{0}) = c^*(\mathbf{q}^*(\tau, V(\tau, w, \mathbf{0})), \boldsymbol{\pi}(\cdot | \tau), V(\tau, w, \mathbf{0})).$$

Hence

$$\begin{aligned} V(\tau, w, \mathbf{0}) &= \mathcal{U}(c(\tau, w, \mathbf{0}), \boldsymbol{\pi}(\cdot | \tau)) \\ &= \mathcal{U}(c^*(\mathbf{q}^*(\tau, V(\tau, w, \mathbf{0})), \boldsymbol{\pi}(\cdot | \tau), V(\tau, w, \mathbf{0})), \boldsymbol{\pi}(\cdot | \tau)) \\ &= v(\mathbf{q}^*(\tau, V(\tau, w, \mathbf{0})), \boldsymbol{\pi}(\cdot | \tau), w). \end{aligned}$$

■

We now narrow our focus on rentiers. The following lemma shows that for the rentiers, the expenditure function, evaluated at shadow prices and the intertemporal indifference curve $m(\tau, w, \mathbf{0})$ coincides with that level of financial wealth.

Lemma 7. *The following holds*

$$e(\mathbf{q}^*(\tau, m(\tau, w, \mathbf{0})), \boldsymbol{\pi}(\cdot|\tau), m(\tau, w, \mathbf{0})) = w.$$

Proof. From the proof of Proposition 2, we know that

$$e(\mathbf{q}^*(\tau, m(\tau, w, \mathbf{y})), \boldsymbol{\pi}(\cdot|\tau), m(\tau, w, \mathbf{y})) = w + \sum_{s^t} \lambda(s^t|\tau) y(s^t|\tau) + \sum_{s^t} \mu(s^t|\tau) X(s^t|\tau),$$

where $\lambda(s^t|\tau)$ are lagrange multipliers on state-contingent budget constraints and $\mu(s^t|\tau)$ are lagrange multipliers on borrowing constraints. Since $y(s^t|\tau) = 0$, we know that

$$e(\mathbf{q}^*(\tau, m(\tau, w, \mathbf{y})), \boldsymbol{\pi}(\cdot|\tau), m(\tau, w, \mathbf{y})) = w + \sum_{s^t} \mu(s^t|\tau) X(s^t|\tau).$$

We prove the desired result by showing that $\mu(s^t) \equiv 0$. To do this, we use backward induction. Suppose that for some t , we know that, for every $t' > t$, we have $\sum_k a_k(s^{t'}|\tau) \geq 0$. That is, the borrowing constraint is slack for every $s^{t'}$ following s^t . For the sake of deriving a contradiction, suppose that $\mu(s^t|\tau) \neq 0$. Then

$$\sum_{n \in \mathbb{N}} p_n(s^{t+1}|\tau) c_n(s^{t+1}|\tau) + \sum_k a_k(s^{t+1}|\tau) = \sum_{k \in K} R_k(s^t|\tau) a_k(s^{t-1}|\tau) < - \left[\min_k R_k(s^t|\tau) \right] X(s^{t-1}|\tau) < 0.$$

This implies that

$$\sum_k a_k(s^{t+1}|\tau) < 0,$$

which is a contradiction. Hence, we know that

$$\sum_k a_k(s^{t+1}|\tau) \geq 0.$$

This implies that $\mu(s^t|\tau) = 0$. We finish by observing that we know that for every s^T , the no-Ponzi scheme condition implies that

$$\sum_k a_k(s^T|\tau) \geq 0.$$

This is the first step of the backward induction. ■

With these preliminaries out of the way, we are ready to prove Proposition 5. We start

with the definition of the money-metric. That is, $m(\tau, w, \mathbf{0})$ solves the following equation:

$$V(\tau, w, \mathbf{0}) = V(\tau_0, m(\tau, w, \mathbf{0}), \mathbf{0}).$$

From Lemma 6, we know

$$v(q^*(\tau, V(\tau, w, \mathbf{0})), \pi(\cdot|\tau), w) = V(\tau, w, \mathbf{0}) = V(\tau_0, m(\tau, w, \mathbf{0}), \mathbf{0}) = v(q^*(\tau_0, V(\tau, w, \mathbf{0}), \mathbf{0}), \pi(\cdot|\tau_0), m(\tau, w, \mathbf{0})).$$

Hence, $m(\tau, w, \mathbf{0})$ solves

$$v(q^*(\tau, V(\tau, w, \mathbf{0})), \pi(\cdot|\tau), w) = v(q^*(\tau_0, V(\tau, w, \mathbf{0}), \mathbf{0}), \pi(\cdot|\tau_0), m(\tau, w, \mathbf{0})).$$

Without loss of generality, by Proposition 2, cardinalize the value function using the money-metric (since the value function is only defined up to monotone transformations).

Therefore

$$v(q^*(\tau, m(\tau, w, \mathbf{0})), \pi(\cdot|\tau), w) = v(q^*(\tau_0, m(\tau, w, \mathbf{0}), \mathbf{0}), \pi(\cdot|\tau_0), m(\tau, w, \mathbf{0})).$$

Using the shadow expenditure function, we can write

$$\begin{aligned} m(\tau, w, \mathbf{0}) &= e(q^*(\tau_0, m(\tau, w, \mathbf{0})), \pi(\cdot|\tau_0), m(\tau, w, \mathbf{0})), \\ &= e(q^*(\tau_0, m(\tau, w, \mathbf{0})), \pi(\cdot|\tau_0), m(\tau, w, \mathbf{0})) \frac{e(q^*(\tau, m(\tau, w, \mathbf{0})), \pi(\cdot|\tau), m(\tau, w, \mathbf{0}))}{e(q^*(\tau, m(\tau, w, \mathbf{0})), \pi(\cdot|\tau), m(\tau, w, \mathbf{0}))}, \\ &= e(q^*(\tau, m(\tau, w, \mathbf{0})), \pi(\cdot|\tau), m(\tau, w, \mathbf{0})) \frac{e(q^*(\tau_0, m(\tau, w, \mathbf{0})), \pi(\cdot|\tau_0), m(\tau, w, \mathbf{0}))}{e(q^*(\tau, m(\tau, w, \mathbf{0})), \pi(\cdot|\tau), m(\tau, w, \mathbf{0}))}, \\ &= w \frac{e(q^*(\tau_0, m(\tau, w, \mathbf{0})), \pi(\cdot|\tau_0), m(\tau, w, \mathbf{0}))}{e(q^*(\tau, m(\tau, w, \mathbf{0})), \pi(\cdot|\tau), m(\tau, w, \mathbf{0}))}, \end{aligned}$$

where the last line uses Lemma 7. Logging both sides gives

$$\begin{aligned} \log m(\tau, w, \mathbf{0}) &= \log w + \log \frac{e(q^*(\tau_0, m(\tau, w, \mathbf{0})), \pi(\cdot|\tau_0), m(\tau, w, \mathbf{0}))}{e(q^*(\tau, m(\tau, w, \mathbf{0})), \pi(\cdot|\tau), m(\tau, w, \mathbf{0}))}, \\ &= \log w + \int_{\tau}^{\tau_0} \left(\frac{\partial \log e(q^*(x, m(\tau, w, \mathbf{0})), \pi(\cdot|x), m(\tau, w, \mathbf{0}))}{\partial \log q^*} \frac{d \log q^*}{dx} \right. \\ &\quad \left. + \frac{\partial \log e(q^*(x, m(\tau, w, \mathbf{0})), \pi(\cdot|x), m(\tau, w, \mathbf{0}))}{\partial \log \pi(\cdot|x)} \frac{d \log \pi(\cdot|x)}{dx} \right) dx, \end{aligned}$$

where the second equality uses the fundamental theorem of calculus for line integrals.

Using Lemma 5, we can rewrite the last line as

$$\begin{aligned}\log m(\tau, w, \mathbf{0}) &= \log w - \int_{\tau_0}^{\tau} \left(\sum_{n \in N} b_n(\mathbf{p}(\cdot|x), m(\tau, w, \mathbf{0})) \frac{d \log p_n}{dx} \right. \\ &\quad \left. + \frac{d \log b^P(\mathbf{q}^*(x, m(\tau, w, \mathbf{0})), \boldsymbol{\pi}(\cdot|x), m(\tau, w, \mathbf{0}))}{\sigma^*(\mathbf{q}^*(\tau, m(\tau, w, \mathbf{0})), \boldsymbol{\pi}(\cdot|x), m(\tau, w, \mathbf{0})) - 1} \frac{1}{dx} \right) dx, \\ &= \log w - \int_{\tau_0}^{\tau} \left(\sum_{n \in N} B_n(x, w_x^*, \mathbf{0}) \frac{d \log p_n}{dx} + \frac{1}{\sigma(x, w_x^*, \mathbf{0}) - 1} \frac{d \log B^P(x, w_x^*, \mathbf{0})}{dx} \right) dx,\end{aligned}$$

where for the last step, we replaced compensated budget share with uncompensated budget share for wealth w_x^* , defined in the statement of the proposition. This completes the proof of Proposition 5. ■

Proof of Lemma 1. We need to show that

$$B_n(\tau, w, \mathbf{y}) = b_n(\mathbf{p}(s^0|\tau), V(\tau, w, \mathbf{y})).$$

By Lemma 2, we know that

$$B_n(\tau, w, \mathbf{y}) = \frac{p_n(s^0|\tau) c_n(s^0|\tau, w, \mathbf{y})}{\sum_{m \in N} p_m(s^0|\tau) c_m(s^0|\tau, w, \mathbf{y})} = \frac{q_n^*(s^0) c_n^*(s^0|\mathbf{q}^*, \boldsymbol{\pi}, V(\tau, w, \mathbf{y}))}{\sum_{m \in N} q_m^*(s^0) c_m^*(s^0|\mathbf{q}^*, \boldsymbol{\pi}, V(\tau, w, \mathbf{y}))} \equiv b_n(\mathbf{q}^*, \boldsymbol{\pi}, V(\tau, w, \mathbf{y})).$$

Next, we know, from Shephard's lemma that for each $n \in N$

$$\begin{aligned}\frac{q_n^*(s^0) c_n^*(\mathbf{q}^*, \boldsymbol{\pi}, V(\tau, w, \mathbf{y}))}{e(\mathbf{q}^*, \boldsymbol{\pi}, V(\tau, w, \mathbf{y}))} &= \frac{\partial \log e(\mathbf{q}^*, \boldsymbol{\pi}, V(\tau, w, \mathbf{y}))}{\partial \log q_n^*(s^0)} \\ &= \frac{\partial \log e(P(\mathbf{q}^*(s^0), V(\tau, w, \mathbf{y})), F(\{\mathbf{q}^*(s^t)\}_{t>0}, \boldsymbol{\pi}, V(\tau, w, \mathbf{y})), V(\tau, w, \mathbf{y}))}{\partial \log P} \\ &= \frac{\partial \log P(\mathbf{q}^*(s^0), V(\tau, w, \mathbf{y}))}{\partial \log q_n^*(s^0)}.\end{aligned}$$

Hence, we have that

$$\frac{q_n^*(s^0) c_n^*(s^0|\mathbf{q}^*, \boldsymbol{\pi}, V(\tau, w, \mathbf{y}))}{\sum_{m \in N} q_m^*(s^0) c_m^*(s^0|\mathbf{q}^*, \boldsymbol{\pi}, V(\tau, w, \mathbf{y}))} = \frac{\frac{\partial \log P(\mathbf{q}^*(s^0), V(\tau, w, \mathbf{y}))}{\partial \log q_n^*(s^0)}}{\sum_{m \in N} \frac{\partial \log P(\mathbf{q}^*(s^0), V(\tau, w, \mathbf{y}))}{\partial \log q_m^*(s^0)}},$$

which is only a function of $\mathbf{q}^*(s^0) = \mathbf{p}(s^0|\tau)$ and $V(\tau, w, \mathbf{y})$ as needed. ■

Proof of Proposition 6. From Lemma 1, we know that

$$\mathbf{B}(\tau, w, \mathbf{y}) = b_n(\mathbf{p}(s^0|\tau), V(\tau, w, \mathbf{y})).$$

By definition of $m(\tau, w, \mathbf{y}|\tau)$, it follows that

$$\mathbf{B}(\tau, w, \mathbf{y}) = b_n(\mathbf{p}(s^0|\tau), V(\tau, m(\tau, w, \mathbf{y}|\tau), \mathbf{0})).$$

Since b is an injective function, we can write

$$V(\tau, m(\tau, w, \mathbf{y}|\tau), \mathbf{0}) = b_n^{-1}(\mathbf{p}(s^0|\tau), \mathbf{B}(\tau, w, \mathbf{y})).$$

Since V is monotone in wealth, we can write

$$m(\tau, w, \mathbf{y}|\tau) = V^{-1}(\tau, b_n^{-1}(\mathbf{p}(s^0|\tau), \mathbf{B}(\tau, w, \mathbf{y})), \mathbf{0}) = m(\mathbf{B}(\tau, w, \mathbf{y}), \tau).$$

■

Proof of Corollary 1. Lemma 1 shows that

$$B_i(\tau, w, \mathbf{y}) = b_i(\mathbf{p}(s^0|\tau), V(\tau, w, \mathbf{y})).$$

Hence, if b_i is monotone in V , then

$$B_i(\tau, w, \mathbf{y}) = b_i(\mathbf{p}(s^0|\tau), V(\tau, w, \mathbf{y})) = b_i(\mathbf{p}(s^0|\tau), V(\tau, w^*, \mathbf{0})) = B_i(\tau, w^*, \mathbf{0})$$

if, and only if,

$$V(\tau, w, \mathbf{y}) = V(\tau, w^*, \mathbf{0}).$$

■

C More Details about Extensions

In this appendix, we present some additional more formal details about the extensions discussed Section 5.

C.1 Changes in Mortality

Consider preferences represented by the utility function with constant EIS,

$$U^{\frac{\sigma-1}{\sigma}} = \lambda_P \tilde{P}(c(s^0), U)^{\frac{\sigma-1}{\sigma}} + \lambda_P \lambda_F \beta \tilde{F}(\{c(s^t)\}_{t>0}, \pi, U)^{\frac{\sigma-1}{\sigma}} + [1 - \lambda_P + \lambda_P(1 - \lambda_F)\beta] \bar{c}^{\frac{\sigma-1}{\sigma}} U^{\frac{\sigma-1}{\sigma} + \varepsilon}$$

where \tilde{P} and \tilde{F} are homogeneous of degree one in quantities (for a given U). The probability of surviving in the present is λ_P , and the probability of surviving in the future is $\lambda_P \lambda_F$. An increase in either λ_P or λ_F raise longevity, and λ_F tilts the increase in lifespan to the future. If the household does not survive, consumption is equal to \bar{c} . The parameter ε determines income effects in the value of life.

The expenditure function associated to these preferences is

$$e(\mathbf{q}, \pi, \lambda, U) = \lambda_P^{\frac{\sigma}{1-\sigma}} \left(P(\mathbf{p}(s^0), U)^{1-\sigma} + \lambda_F^\sigma \beta^\sigma F(\{\mathbf{q}(s^t)\}_{t>0}, \pi, U)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \left[1 - [1 - \lambda_P + \lambda_P(1 - \lambda_F)\beta] \bar{c}^{\frac{\sigma-1}{\sigma}} U^\varepsilon \right]^{\frac{\sigma}{\sigma-1}} U,$$

where P and F are homogeneous of degree one in prices.

Define the Hicksian money-value of increases in λ_P and λ_F are

$$\phi_P(\tau, U) = \frac{d \log e(\mathbf{q}, \pi, \lambda, U)}{d \log \lambda_P} = \frac{\sigma}{1 - \sigma} \left[1 + \frac{\lambda_P (1 - (1 - \lambda_F)\beta) \bar{c}^{\frac{\sigma-1}{\sigma}} U^\varepsilon}{1 - [1 - \lambda_P + \lambda_P(1 - \lambda_F)\beta] \bar{c}^{\frac{\sigma-1}{\sigma}} U^\varepsilon} \right]$$

and

$$\phi_F(\tau, U) = \frac{d \log e(\mathbf{q}, \pi, \lambda, U)}{d \log \lambda_F} = \frac{\sigma}{1 - \sigma} \left[1 - b^p(\tau, U) + \frac{\lambda_P \lambda_F \beta \bar{c}^{\frac{\sigma-1}{\sigma}} U^\varepsilon}{1 - [1 - \lambda_P + \lambda_P(1 - \lambda_F)\beta] \bar{c}^{\frac{\sigma-1}{\sigma}} U^\varepsilon} \right].$$

If $\phi_P(\tau, U) < 0$ or $\phi_F(\tau, U) < 0$, then households are willing to give up initial wealth to increase longevity via a marginal increase in λ_P or λ_F .

Define the Marshallian value of life for households in cohort τ with initial wealth w to be

$$\Phi_0(\tau, w) = \frac{d \log e(\mathbf{p}, \pi, \lambda, V(\tau, w))}{d \log \lambda_P}$$

and

$$\Phi_1(\tau, w) = \frac{d \log e(\mathbf{p}, \pi, \lambda, V(\tau, w))}{d \log \lambda_P}.$$

Note that

$$m(\tau, w) = \tilde{H}(V(\tau, w))$$

where \tilde{H} is strictly increasing. Hence,

$$\Phi_P(\tau, w) = \phi_P(\tau, V(\tau, w)) = \phi_P(\tau, \tilde{H}^{-1}(m(\tau, w))) \equiv \tilde{\phi}_P(\tau, m(\tau, w))$$

and

$$\Phi_F(\tau, w) = \phi_F(\tau, V(\tau, w)) = \phi_F(\tau, \tilde{H}^{-1}(m(\tau, w))) \equiv \tilde{\phi}_F(\tau, m(\tau, w))$$

Following similar arguments to the proof of Proposition 5, we can show that

$$\begin{aligned} \log m(\tau, w) &= \log w - \log \frac{P(\tau, m(\tau, w))}{P(\tau_0, m(\tau, w))} + \frac{1}{1-\sigma} \log \frac{b^p(\tau, m(\tau, w))}{b^p(\tau_0, m(\tau, w))} \\ &- \int_{\tau_0}^{\tau} \Phi_P(x, m(\tau, w)) d \log \lambda_P(x) dx - \int_{\tau_0}^{\tau} \left[\phi_F(x, V(\tau, w)) - \frac{\sigma}{1-\sigma} (1 - b^p(x, m(\tau, w))) \right] d \log \lambda_F(x) dx, \end{aligned}$$

or

$$\begin{aligned} \log m(\tau, w) &= \log w - \log \frac{P(\tau, m(\tau, w))}{P(\tau_0, m(\tau, w))} + \frac{1}{1-\sigma} \log \frac{b^p(\tau, m(\tau, w))}{b^p(\tau_0, m(\tau, w))} \\ &- \int_{\tau_0}^{\tau} \tilde{\phi}_0(x, m(\tau, w)) d \log \lambda_P(x) dx - \int_{\tau_0}^{\tau} \left[\tilde{\phi}_1(x, m(\tau, w)) - \frac{\sigma}{1-\sigma} (1 - b^p(x, m(\tau, w))) \right] d \log \lambda_F(x) dx, \end{aligned}$$

or

$$\begin{aligned} \log m(\tau, w) &= \log w - \int_{\tau_0}^{\tau} \left(\sum_{n \in \mathbb{N}} B_n(x, w_x^*) \frac{d \log p_n}{dx} - \frac{1}{1-\sigma} \frac{d \log B^P(x, w_x^*)}{dx} \right) dx - \int_{\tau_0}^{\tau} \Phi_P(x, w_x^*) d \log \lambda_P(x) dx \\ &- \int_{\tau_0}^{\tau} \left[\Phi_1(x, w_x^*) - \frac{\sigma}{1-\sigma} (1 - B^P(x, w_x^*)) \right] d \log \lambda_F(x) dx. \end{aligned}$$

C.2 Time-Inconsistent Preferences

Here we argue that our results continue to hold when consumers are time-inconsistent, as long as time-separability holds. We impose again that the consumer's spending on good i relative to good j in the present is only a function of prices in the present and overall U . In this case, it is obvious that Proposition 6 continues to hold. Moreover, if the current self expects spending on future good i relative to good j , decided by future selves, to also only be a function of prices in the future and overall U , then Proposition 5 also continues to apply.

To see this, and for concreteness, consider the following simple case (the argument

generalizes). Suppose preferences are of the form

$$U^{\frac{\sigma(U)-1}{\sigma(U)}} = \left(\sum_{n'} \omega_{n'0} U^{\varepsilon_{n'0}} c_{n'0}^{\frac{\gamma-1}{\gamma}} \right)^{\frac{\gamma}{\gamma-1} \frac{\sigma(U)-1}{\sigma(U)}} + \beta(U) F(c_{j>0}, U)^{\frac{\sigma(U)-1}{\sigma(U)}}$$

where F is homogeneous of degree 1. Assume further that there is no risk and there is a single risk free interest rate between each period (for simplicity).

Suppose that a consumer in the initial period plans as if choices in future periods $j > 0$ are made by future selves according to

$$c_{nj>0} = f_{nj}(\mathbf{p}_{j>0}, U) \frac{w_1}{p_{nj}}$$

where w_1 is wealth at the first period in the future ($j = 1$), and f_{nj} is the budget share on good n in each future period $j > 0$, which satisfies

$$\sum_{n, j>0} f_{nj}(\mathbf{p}_{j>0}, U) = 1.$$

Choices made by future selves need not be consistent with the choices that the household would have made under commitment. The fact that future choices made by future selves depend only on future relative prices and U is the essence of the time-separability assumption.

The household in the present has initial wealth w and picks savings a_1 with one period return r_1 , satisfying the budget constraint:

$$\sum_{n'} p_{n'0} c_{n'0} + a_1 = w,$$

and the initial wealth in $j = 1$ is

$$w_1 = r_1 a_1.$$

Replacing future choices into the utility function (given optimal allocation of consumption in the present) yields

$$U^{\frac{\sigma(U)-1}{\sigma(U)}} = \left(\frac{w - a_1}{\left(\sum_{n'} \omega_{n'0}^\gamma U^{\varepsilon_{n'0} \gamma} p_{n'0}^{1-\gamma} \right)^{\frac{1}{1-\gamma}}} \right)^{\frac{\sigma(U)-1}{\sigma(U)}} + \beta(U) (r_1 a_1 G(\mathbf{p}_{j>0}, U))^{\frac{\sigma(U)-1}{\sigma(U)}}, \quad (29)$$

where

$$r_1 a_1 G(\mathbf{p}_{j>0}, U) = F\left(f_{nj}(\mathbf{p}_{j>0}, U) \frac{w_1}{p_{ns}}, U\right) = r_1 a_1 F\left(f_{nj}(\mathbf{p}_{j>0}, U) \frac{1}{p_{ns}}, U\right).$$

The household in the present chooses a_1 to maximize U , which implies the first order condition:

$$\left(\frac{w - a_1}{\left(\sum_{n'} \omega_{n'0}^\gamma U^{\epsilon_{n'} \gamma} p_{n'0}^{1-\gamma}\right)^{\frac{1}{1-\gamma}}}\right)^{\frac{\sigma(U)-1}{\sigma(U)}} (w - a_1)^{-\frac{1}{\sigma}} = \beta(U) r_1^{1-\frac{1}{\sigma}} a_1^{-\frac{1}{\sigma}} \left(G(\mathbf{p}_{j>0}, U)\right)^{\frac{\sigma(U)-1}{\sigma(U)}}$$

Substituting the optimal choice a_1 into (29) and simplifying yields

$$U = \left(\frac{w}{\left(\sum_{n'} \omega_{n'0}^\gamma U^{\epsilon_{n'} \gamma} p_{n'0}^{1-\gamma}\right)^{\frac{1}{1-\gamma}}}\right) (B^p)^{\frac{1}{1-\sigma(U)}},$$

where $B^p = \frac{\sum_{n'} \epsilon_{n'0} p_{n'0}}{w}$ is the consumption-wealth ratio in the present, which encodes future prices and returns.

Consider now some base prices (present and future) and returns at τ_0 , \mathbf{p}_{τ_0} and $r_{1\tau_0}$. To obtain the money metric, we solve for the wealth level w_{τ_0} that keeps utility unchanged at U under base prices:

$$\left(\frac{w}{\left(\sum_{n'} \omega_{n'0}^\gamma U^{\epsilon_{n'} \gamma} p_{n'0}^{1-\gamma}\right)^{\frac{1}{1-\gamma}}}\right) (B^p(U))^{\frac{1}{1-\sigma(U)}} = \left(\frac{w_{\tau_0}}{\left(\sum_{n'} \omega_{n'0}^\gamma U^{\epsilon_{n'} \gamma} p_{n'0\tau_0}^{1-\gamma}\right)^{\frac{1}{1-\gamma}}}\right) (B_{\tau_0}^p(U))^{\frac{1}{1-\sigma(U)}},$$

where $B_{\tau_0}^p(U)$ is the consumption-to-wealth ratio of a consumer in the same indifference curve U at τ_0 . Hence,

$$w_{\tau_0} = w \left(\frac{\sum_{n'} \omega_{n'0}^\gamma U^{\epsilon_{n'} \gamma} p_{n'0\tau_0}^{1-\gamma}}{\sum_{n'} \omega_{n'0}^\gamma U^{\epsilon_{n'} \gamma} p_{n'0}^{1-\gamma}}\right)^{\frac{1}{1-\gamma}} \left(\frac{B^p(U)}{B_{\tau_0}^p(U)}\right)^{\frac{1}{1-\sigma(U)}}.$$

The first term in the right hand side is static inflation for a consumer in indifference curve U , and the second term in brackets is the change in consumption-wealth ratio for this consumer. As in the model with commitment, the change in the consumption-wealth ratio is a sufficient statistic for how changes in future prices and rates of returns affects the household in the present. From this equation, we can then derive the expression in Proposition 5 for this case. This argument can be generalized to allow for more general preferences, risk, etc. The key is that, given the assumption on separability (future selves

consumption choices only depend on future prices, returns, and U) the expression for dynamic welfare is the same with or without commitment.

C.3 Non-Time-Separable Non-Homotheticity

We discuss the problem of rentiers first, before considering the non-rentiers.

Rentiers' problem. For simplicity, we assume that rentiers solve a complete markets problem. (The results could be extended along similar lines as the main paper to incomplete markets.) The intertemporal expenditure function for rentiers is defined by the following constrained optimization problem:

$$\min_{E_t} \sum_{t=0}^{\infty} \frac{E_t}{(1+r_t)}$$

such that

$$U = \left[\sum_t \beta^t C_t^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}},$$

and expenditures in period t satisfy

$$E_t = \left[\sum_i \omega_{ti} [C_t^{\epsilon_i} p_{ti}]^{1-\gamma} \right]^{\frac{1}{1-\gamma}}.$$

The Lagrangian is

$$\sum_{t=0}^{\infty} \frac{E_t}{(1+r_t)} - \lambda \left[U - \left[\sum_t \beta^t C_t^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}} \right] + \sum_t \mu_t \left[E_t - \left[\sum_i \omega_{ti} [C_t^{\epsilon_i} p_{ti}]^{1-\gamma} \right]^{\frac{1}{1-\gamma}} \right].$$

The first order conditions are

$$\frac{1}{(1+r_t)} = \mu_t,$$

and

$$\lambda \frac{\theta}{\theta-1} \left[\sum_t \beta^t C_t^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}-1} \beta^t \frac{\theta-1}{\theta} C_t^{\frac{\theta-1}{\theta}-1} = \mu_t \frac{1}{1-\gamma} \left[\sum_i \omega_{ti} [C_t^{\epsilon_i} p_{ti}]^{1-\gamma} \right]^{\frac{1}{1-\gamma}-1} \sum_i \omega_{ti} [C_t^{\epsilon_i} p_{ti}]^{1-\gamma} (1-\gamma) \epsilon_i \frac{1}{C_t}.$$

Combine these first order conditions to get the equation,

$$\frac{1}{(1+r_t)} E_t^\gamma \sum_i \omega_{ti} [C_t^{\epsilon_i} p_{ti}]^{1-\gamma} \epsilon_i \frac{1}{C_t} = \lambda U^{\frac{1}{\theta}} \beta^t C_t^{\frac{-1}{\theta}}$$

or

$$\frac{1}{(1+r_t)} E_t^\gamma \left[\sum_i \omega_{ti} [C_t^{\epsilon_i} p_{ti}]^{1-\gamma} \right] \sum_i \frac{\omega_{ti} [C_t^{\epsilon_i} p_{ti}]^{1-\gamma}}{\sum_i \omega_{ti} [C_t^{\epsilon_i} p_{ti}]^{1-\gamma}} \epsilon_i \frac{1}{C_t} = \lambda U^{\frac{1}{\theta}} \beta^t C_t^{\frac{-1}{\theta}}.$$

Using the fact that budget shares are:

$$b_{ti} = \frac{\partial \log E_t}{\partial \log p_{ti}} = \frac{\omega_{ti} [C_t^{\epsilon_i} p_{ti}]^{1-\gamma}}{\sum_i \omega_{ti} [C_t^{\epsilon_i} p_{ti}]^{1-\gamma}},$$

we can rewrite the first-order condition as:

$$\frac{1}{(1+r_t)} \frac{E_t}{C_t} \sum_i b_{ti} \epsilon_i = \lambda U^{\frac{1}{\theta}} \beta^t C_t^{\frac{-1}{\theta}}.$$

The log-linearized demand system is

$$-d \log(1+r_t) + d \log E_t - d \log C_t + \frac{\text{Cov}_b(\epsilon_i, d \log b_{ti})}{\mathbb{E}_b[\epsilon_i]} = d \log \lambda + \frac{1}{\theta} d \log U - \frac{1}{\theta} d \log C_t,$$

or

$$\left[1 - \frac{1}{\theta} \right] d \log C_t - d \log E_t = -d \log \lambda - \frac{1}{\theta} d \log U - d \log(1+r_t) + \frac{\text{Cov}_b(\epsilon_i, d \log b_{ti})}{\mathbb{E}_b[\epsilon_i]},$$

where

$$d \log E_t = \mathbb{E}_{b_t} [d \log p_{ti}] + \mathbb{E}_{b_t} [\epsilon_i] d \log C_t,$$

and

$$d \log b_{ti} = (1-\gamma) [d \log p_{ti} - d \log E_t] + (1-\gamma) \epsilon_i d \log C_t$$

To keep the household on the same indifference curve, we also require that

$$\sum_t \frac{E_t / (1+r)^t}{\sum_{t'} E_{t'} / (1+r)^{t'}} (d \log E_t - \mathbb{E}_{b_t} [d \log p_t]) = 0,$$

which is a consequence of Shephard's lemma. Hence, more generally, for any compensated variation, the following equations must hold:

$$\left[1 - \frac{1}{\theta}\right] d \log C_t - d \log E_t = -d \log \lambda - d \log(1 + r_t) + \frac{\text{Cov}_b(\epsilon_i, d \log b_{ti})}{\mathbb{E}_b[\epsilon_i]},$$

$$\frac{1}{\mathbb{E}_{b_t}[\epsilon_i]} [d \log E_t - \mathbb{E}_{b_t}[d \log p_{ti}]] = d \log C_t,$$

$$d \log b_{ti} = (1 - \gamma) [d \log p_{ti} - d \log E_t] + (1 - \gamma)\epsilon_i d \log C_t,$$

$$\sum_t \frac{E_t/(1 + r_t)}{\sum_{t'} E_{t'}/(1 + r_{t'})} (d \log E_t - \mathbb{E}_{b_t}[d \log p_t]) = 0.$$

Combine these equations to get

$$\frac{\left[1 - \frac{1}{\theta}\right]}{\mathbb{E}_{b_t}[\epsilon_i]} [d \log E_t - \mathbb{E}_{b_t}[d \log p_{ti}]] - d \log E_t = -d \log \lambda - d \log(1 + r_t) + \frac{\text{Cov}_b(\epsilon_i, d \log b_{ti})}{\mathbb{E}_b[\epsilon_i]},$$

$$d \log b_{ti} = -(1 - \gamma) [d \log E_t - d \log p_{ti}] + (1 - \gamma) \frac{\epsilon_i}{\mathbb{E}_{b_t}[\epsilon_i]} [d \log E_t - \mathbb{E}_{b_t}[d \log p_{ti}]],$$

and

$$\sum_t \frac{E_t/(1 + r_t)}{\sum_{t'} E_{t'}/(1 + r_{t'})} (d \log E_t - \mathbb{E}_{b_t}[d \log p_t]) = 0.$$

Or

$$\frac{\left[1 - \frac{1}{\theta}\right]}{\mathbb{E}_{b_t}[\epsilon_i]} [d \log E_t - \mathbb{E}_{b_t}[d \log p_{ti}]] - d \log E_t = -d \log \lambda - d \log(1 + r_t) + \frac{(1 - \gamma)}{\mathbb{E}_b[\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) + (1 - \gamma) \frac{\text{Var}_b(\epsilon_i)}{\mathbb{E}_b[\epsilon_i]^2} [d \log E_t - \mathbb{E}_{b_t}[d \log p_{ti}]],$$

and

$$\sum_t \frac{E_t/(1 + r_t)}{\sum_{t'} E_{t'}/(1 + r_{t'})} (d \log E_t - \mathbb{E}_{b_t}[d \log p_t]) = 0.$$

Define the inverse Frisch elasticity $\kappa_t = \left[\frac{\left[1 - \frac{1}{\theta}\right]}{\mathbb{E}_{b_t}[\epsilon_i]} - (1 - \gamma) \frac{\text{Var}_b(\epsilon_i)}{\mathbb{E}_b[\epsilon_i]^2} \right]$. Then we can write

$$\kappa_t [d \log E_t - \mathbb{E}_{b_t}[d \log p_{ti}]] = d \log E_t - d \log \lambda - d \log(1 + r_t) + \frac{(1 - \gamma)}{\mathbb{E}_b[\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}),$$

and

$$\sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} (d \log E_t - \mathbb{E}_{b_t} [d \log p_t]) = 0.$$

Substituting one equation into the second implies that

$$\sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} \kappa_t^{-1} \left(d \log E_t - d \log \lambda - d \log(1+r_t) + \frac{(1-\gamma)}{\mathbb{E}_b[\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) \right) = 0.$$

Denote by $\bar{\kappa} = \left[\sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} \kappa_t^{-1} \right]^{-1}$ the harmonic net present value weighted average of κ_t . Then we can rewrite the previous equation as

$$\begin{aligned} \bar{\kappa} \sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} \kappa_t^{-1} \left(d \log E_t - d \log(1+r_t) + \frac{(1-\gamma)}{\mathbb{E}_b[\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) \right) &= d \log \lambda, \\ \mathbb{E}_{\bar{\kappa}}^{NPV} \left[\mathbb{E}_{b_t} [d \log p_t] - d \log(1+r_t) + \frac{(1-\gamma)}{\mathbb{E}_b[\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) \right] &= d \log \lambda, \end{aligned}$$

where $\mathbb{E}_{\bar{\kappa}}^{NPV}$ is an average that uses weights $\frac{\frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} \kappa_t^{-1}}{\left[\sum_{t'} \frac{E_{t'}/(1+r_{t'})}{\sum_{t''} E_{t''}/(1+r_{t''})} \kappa_{t'}^{-1} \right]}$, which is a net-present value weighted κ_t . Hence,

$$\begin{aligned} [\kappa_t - 1] d \log E_t - [\kappa_t - 1] \mathbb{E}_{b_t} [d \log p_{ti}] &= \mathbb{E}_{b_t} [d \log p_{ti}] - d \log(1+r_t) + \frac{(1-\gamma)}{\mathbb{E}_b[\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) \\ &\quad - \mathbb{E}_{\bar{\kappa}}^{NPV} \left[\mathbb{E}_{b_t} [d \log p_t] - d \log(1+r_t) + \frac{(1-\gamma)}{\mathbb{E}_b[\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) \right] \end{aligned}$$

Evaluate this for $t = 0$ to get

$$\begin{aligned} [\kappa_0 - 1] [d \log E_0 - \mathbb{E}_{b_0} [d \log p_{0i}]] &= \mathbb{E}_{b_0} [d \log p_{0i}] + \frac{(1-\gamma)}{\mathbb{E}_b[\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) \\ &\quad - \mathbb{E}_{\bar{\kappa}}^{NPV} \left[\mathbb{E}_{b_t} [d \log p_t] - d \log(1+r_t) + \frac{(1-\gamma)}{\mathbb{E}_b[\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) \right], \end{aligned}$$

or

$$\begin{aligned} d \log E_0 &= [\kappa_0 - 1]^{-1} \kappa_0 \mathbb{E}_{b_0} [d \log p_{0i}] + [\kappa_0 - 1]^{-1} \frac{(1-\gamma)}{\mathbb{E}_b[\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) \\ &\quad - [\kappa_0 - 1]^{-1} \mathbb{E}_{\bar{\kappa}}^{NPV} \left[\mathbb{E}_{b_t} [d \log p_t] - d \log(1+r_t) + \frac{(1-\gamma)}{\mathbb{E}_b[\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) \right]. \end{aligned}$$

Let W denote wealth, and note that the consumption-to-wealth ratio satisfies

$$\begin{aligned} d \log E_0 - d \log W &= [\kappa_0 - 1]^{-1} \kappa_0 \mathbb{E}_{b_0} [d \log p_{0i}] + [\kappa_0 - 1]^{-1} \frac{(1 - \gamma)}{\mathbb{E}_b [\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) \\ &\quad - [\kappa_0 - 1]^{-1} \mathbb{E}_{\kappa^{-1}}^{NPV} \left[\mathbb{E}_{b_t} [d \log p_t] - d \log(1 + r_t) + \frac{(1 - \gamma)}{\mathbb{E}_b [\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) \right] \\ &\quad - \sum_t \frac{E_t/(1 + r_t)}{\sum_{t'} E_{t'}/(1 + r_{t'})} (\mathbb{E}_{b_t} [d \log p_t] - d \log(1 + r_t)), \end{aligned}$$

where the last line uses the fact that $d \log W = \sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} (\mathbb{E}_{b_t} [d \log p_t] - d \log(1 + r_t))$. This fact is a consequence of two equations: (1) the accounting identity $d \log W = \sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} (d \log E_t - d \log(1 + r_t))$, and (2) the fact that any compensated variation must satisfy $\sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} (d \log E_t - \mathbb{E}_{b_t} [d \log p_t]) = 0$. Combining these two equations yields the desired result.

$$\begin{aligned} d \log E_0 - d \log W &= [\kappa_0 - 1]^{-1} \kappa_0 \mathbb{E}_{b_0} [d \log p_{0i}] + [\kappa_0 - 1]^{-1} \frac{(1 - \gamma)}{\mathbb{E}_b [\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) \\ &\quad - [\kappa_0 - 1]^{-1} \mathbb{E}_{\kappa^{-1}}^{NPV} \left[\mathbb{E}_{b_t} [d \log p_t] - d \log(1 + r_t) + \frac{(1 - \gamma)}{\mathbb{E}_b [\epsilon_i]} \text{Cov}_b(\epsilon_i, d \log p_{ti}) \right] \\ &\quad - \sum_t \frac{E_t/(1 + r_t)}{\sum_{t'} E_{t'}/(1 + r_{t'})} (\mathbb{E}_{b_t} [d \log p_t] - d \log(1 + r_t)), \end{aligned}$$

or

$$\begin{aligned} d \log E_0 - d \log W &= [\kappa_0 - 1]^{-1} \kappa_0 \mathbb{E}_{b_0} [d \log p_{0i}] \\ &\quad + \frac{(1 - \gamma)}{\kappa_0 - 1} \left[\text{Cov}_b\left(\frac{\epsilon_i}{\mathbb{E}_b [\epsilon_i]}, d \log p_{ti}\right) - \mathbb{E}_{\kappa^{-1}}^{NPV} \left[\text{Cov}_b\left(\frac{\epsilon_i}{\mathbb{E}_b [\epsilon_i]}, d \log p_{ti}\right) \right] \right] \\ &\quad - \frac{\bar{\kappa}}{\kappa_0 - 1} \sum_t \frac{E_t/(1 + r_t)}{\sum_{t'} E_{t'}/(1 + r_{t'})} \left[\frac{1}{\kappa_t} + \frac{[\kappa_0 - 1]}{\bar{\kappa}} \right] (\mathbb{E}_{b_t} [d \log p_t] - d \log(1 + r_t)) \end{aligned}$$

where $\bar{\kappa} = \left[\sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} \kappa_t^{-1} \right]^{-1}$ is the harmonic net present value weighted average of κ_t . Then we can write

$$\begin{aligned} d \log E_0 - d \log W &= \frac{\kappa_0}{\kappa_0 - 1} \mathbb{E}_{b_0} [d \log p_{0i}] \\ &\quad - \frac{\kappa_0}{\kappa_0 - 1} \sum_t \frac{E_t/(1 + r_t)}{\sum_{t'} E_{t'}/(1 + r_{t'})} \left[\frac{1}{\kappa_0} \frac{\bar{\kappa}}{\kappa_t} + \left[1 - \frac{1}{\kappa_0} \right] \right] (\mathbb{E}_{b_t} [d \log p_t] - d \log(1 + r_t)) \\ &\quad + \frac{(1 - \gamma)}{\kappa_0 - 1} \left[\text{Cov}_b\left(\frac{\epsilon_i}{\mathbb{E}_b [\epsilon_i]}, d \log p_{ti}\right) - \mathbb{E}_{\kappa^{-1}}^{NPV} \left[\text{Cov}_b\left(\frac{\epsilon_i}{\mathbb{E}_b [\epsilon_i]}, d \log p_{ti}\right) \right] \right], \end{aligned}$$

or

$$\begin{aligned}
d \log E_0 - d \log W &= \frac{\kappa_0}{\kappa_0 - 1} \mathbb{E}_{b_0} [d \log p_{0i}] - \frac{\kappa_0}{\kappa_0 - 1} \sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} (\mathbb{E}_{b_t} [d \log p_t] - d \log(1+r_t)) \\
&\quad - \frac{1}{\kappa_0 - 1} \sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} \left[\frac{\bar{\kappa} - \kappa_t}{\kappa_t} \right] (\mathbb{E}_{b_t} [d \log p_t] - d \log(1+r_t)) \\
&\quad + \frac{(1-\gamma)}{\kappa_0 - 1} \left[\text{Cov}_b \left(\frac{\epsilon_i}{\mathbb{E}_b [\epsilon_i]}, d \log p_{ti} \right) - \mathbb{E}_{\kappa^{-1}}^{NPV} \left[\text{Cov}_b \left(\frac{\epsilon_i}{\mathbb{E}_b [\epsilon_i]}, d \log p_{ti} \right) \right] \right],
\end{aligned}$$

Therefore,

$$\begin{aligned}
\frac{\kappa_0 - 1}{\kappa_0} [d \log E_0 - d \log W] &= \mathbb{E}_{b_0} [d \log p_{0i}] - \sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} (\mathbb{E}_{b_t} [d \log p_t] - d \log(1+r_t)) \\
&\quad - \frac{1}{\kappa_0} \sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} \left[\frac{\bar{\kappa} - \kappa_t}{\kappa_t} \right] (\mathbb{E}_{b_t} [d \log p_t] - d \log(1+r_t)) \\
&\quad + \frac{(1-\gamma)}{\kappa_0} \left[\text{Cov}_b \left(\frac{\epsilon_i}{\mathbb{E}_b [\epsilon_i]}, d \log p_{ti} \right) - \mathbb{E}_{\kappa^{-1}}^{NPV} \left[\text{Cov}_b \left(\frac{\epsilon_i}{\mathbb{E}_b [\epsilon_i]}, d \log p_{ti} \right) \right] \right],
\end{aligned}$$

or

$$\begin{aligned}
\sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} (\mathbb{E}_{b_t} [d \log p_t] - d \log(1+r_t)) &= \mathbb{E}_{b_0} [d \log p_{0i}] - \frac{\kappa_0 - 1}{\kappa_0} [d \log E_0 - d \log W] \\
&\quad + \frac{(1-\gamma)}{\kappa_0} \left[\text{Cov}_b \left(\frac{\epsilon_i}{\mathbb{E}_b [\epsilon_i]}, d \log p_{ti} \right) - \mathbb{E}_{\kappa^{-1}}^{NPV} \left[\text{Cov}_b \left(\frac{\epsilon_i}{\mathbb{E}_b [\epsilon_i]}, d \log p_{ti} \right) \right] \right] \\
&\quad - \frac{1}{\kappa_0} \sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} \left[\frac{\bar{\kappa} - \kappa_t}{\kappa_t} \right] (\mathbb{E}_{b_t} [d \log p_t] - d \log(1+r_t)).
\end{aligned}$$

Hence, from Shephard's lemma, the change in money metric is in the change in wealth deflated by average prices:

$$\begin{aligned}
d \log m &= d \log W - \sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} (\mathbb{E}_{b_t} [d \log p_t] - d \log(1+r_t)) \\
&= d \log W - \mathbb{E}_{b_0} [d \log p_{0i}] + \left[1 - \frac{1}{\kappa_0} \right] \left[d \log \frac{E_0}{W} \right] \\
&\quad - \frac{(1-\gamma)}{\kappa_0} \left[\text{Cov}_{b_0} \left(\frac{\epsilon_i}{\mathbb{E}_{b_0} [\epsilon_i]}, d \log p_{0i} \right) - \mathbb{E}_{\kappa^{-1}}^{NPV} \left[\text{Cov}_b \left(\frac{\epsilon_i}{\mathbb{E}_b [\epsilon_i]}, d \log p_{ti} \right) \right] \right] \\
&\quad + \frac{1}{\kappa_0} \sum_t \frac{E_t/(1+r_t)}{\sum_{t'} E_{t'}/(1+r_{t'})} \left[\frac{\bar{\kappa} - \kappa_t}{\kappa_t} \right] (\mathbb{E}_{b_t} [d \log p_t] - d \log(1+r_t)).
\end{aligned}$$

Integrating this expression, using the fundamental theorem of calculus, yields the desired result. Note that as $\theta \rightarrow 0$, we have

$$\lim_{\theta \rightarrow 0} \kappa_0 = \lim_{\theta \rightarrow 0} \left[\frac{\left[1 - \frac{1}{\theta}\right]}{\mathbb{E}_{b_t} [\epsilon_i]} - (1 - \gamma) \frac{\text{Var}_b(\epsilon_i)}{\mathbb{E}_b [\epsilon_i]^2} \right] = -\infty.$$

Hence, in this limit, the error term vanishes:

$$\lim_{\theta \rightarrow 0} \log m = \log W - \int \left(\mathbb{E}_{b_0} [d \log p_{0i}] - \left[d \log \frac{E_0}{W} \right] \right).$$

Non-Rentiers. Now consider the the non-rentiers, who face a similar problem except receive risky labor income and are subject to borrowing constraints (as usual, borrowing constraints cannot bind for rentiers, so we can ignore them in the previous calculation). Non-rentiers maximize utility

$$\begin{aligned} \max U &= \mathbb{E} \left[\sum_t \beta^t C_t^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}} \\ E_t &= \left[\sum_i \omega_{ti} [C_t^{\epsilon_i} p_{ti}]^{1-\gamma} \right]^{\frac{1}{1-\gamma}} \\ E_t(s) + \sum_k a_{kt}(s) &= \sum_k R_{k,t}(s) a_{kt-1}(s) + y_t(s) \\ \sum_k a_{kt}(s) &\leq -B_t(s). \end{aligned}$$

Note that as $\theta \rightarrow 0$, the utility function converges to $U \rightarrow \max [\min_{t \geq 0} C_t]$. Furthermore, it is easy to see that for this problem,

$$\max \left[\min_{t \geq 0} C_t \right] = C_0.$$

This can be proved by contradiction. Suppose that a feasible consumption plan obtains its minimum at some history $s^{f^*} > 0$. The agent can reduce consumption in 0, increase savings, and increase $C(s^{f^*})$, which raises utility. Hence, the postulated consumption plan cannot be optimal. This argument holds as long as rates of return are positive, so that savings at time 0 can raise consumption in history s^{f^*} (the presence of a riskfree asset guarantees this time).

D Data Construction

We use two different datasets. One is a household-level survey (PSID) and the other is data on prices of different categories of goods (CPI). The PSID is a longitudinal survey, interviewing households annually until 1997 and biennially thereafter. Each sample includes about 7,000-9,000 households. We use seven spending categories and merge them with CPI categories. We describe how we construct the variables needed for our methodology below.

Net Assets:

The wealth module of the PSID tracks the value of components of household balance sheets (business equity, stocks, mutual funds, bonds, automobiles, pensions, cash, etc.). Home equity data are recorded as the value of a household's home minus its mortgage obligations. The PSID aggregates these variables, imputes missing values, and reports the comprehensive variables WEALTH1 and WEALTH2. WEALTH1 represents wealth excluding home equity, while WEALTH2 is the sum of WEALTH1 and home equity. As Cooper et al. (2019) note, these measures exclude the value of defined-contribution (DC) account. We define net assets as WEALTH2 plus the value of DC account (recorded separately in the PSID) to incorporate as much of the household's assets as possible.³⁸

Capitalized wealth proxy:

We construct a proxy for total wealth by adding the capitalized value of labor income and transfers to net assets. Define household income as labor income plus the variables recorded as social security income and other welfare income. First, we estimate the age-specific income profile for each period τ using cross-sectional data. To do this, we regress a quadratic of the age of the head of household on log income controlling for household characteristics (marital status, state of residence, race of household head, gender, and occupation) and year fixed effects. We then use this regression to predict each household's income profile as their age increases. We inflate these predictions of the household's income in the future by an estimate of expected nominal per capita GDP growth. The expected growth in nominal GDP comes from the Congressional Budget Office's real-time (contemporaneous) forecast of nominal GDP growth and the population growth rate uses realized population growth rates for the United States, assuming a constant growth after

³⁸Cooper et al. (2019) report that adding DC account information to WEALTH2 generally matches the total assets reported in the Survey of Consumer Finances (SCF). If no value was provided and the value was given in bins, the median household value between the bins was used for imputation.

2019. We discount these nominal income flows back to the present using a nominal rate of 6%, consisting of a 4% real rate, following Catherine et al. (2022), and a 2% expected inflation rate. We assume that income flows are zero beyond age 90.

Owner-occupied housing:

For renters, we use the housing expenditures variable in the PSID (which includes utilities). For owner-occupied housing, we impute housing costs by matching homeowners to renters using static budget shares in each period. This procedure should yield accurate estimates as long as preferences are time separable.

Specifically, for each year, we run the following regression for renters:

$$housing_{h,\tau} = \sum_{i \neq \text{housing}} \alpha_{i,\tau} spending_{i,h,\tau} + \beta_{1,\tau} age_{h,\tau} + \beta_{2,\tau} age_{h,\tau}^2 + stateFE_{h,\tau} + \epsilon_{h,\tau},$$

where the left-hand side variable is expenditures on housing (including utilities), and covariates are households' spending on non-housing categories, age, and state fixed effects. We then use this regression to impute (predict) rental expenditures for homeowners based on their age, spending on non-housing categories, and state of residence.

In 2019, a new question was added to the PSID survey which asks the following:

If someone were to rent this (apartment/mobile home/home) today, how much do you think it would rent for per month, unfurnished and without utilities?

We use the responses to this question to validate our procedure. A regression of the survey values (including utilities) on our imputed values, both relative to current expenditures, has a coefficient of 1.03 with an R^2 value of 0.59. This suggests that our imputation performs well.

Budget shares:

We align the seven categories of the PSID (food, housing, transportation, education, health, clothing, and recreation) with the CPI.³⁹ As mentioned above, for homeowners, we impute housing costs. The relative budget share is defined as the spending on each category divided by total spending. We compute the consumption-wealth ratio of households by dividing total spending in each year by wealth.

³⁹The corresponding codes for CPI are CPIFABSL, CPIHOSSL, CPITRNSL, CPIEDSL, CPIMEDSL, CPI-APPSL, and CPIRECSL, respectively. Education includes child care. Recreation includes Trips & vacations and Recreation & entertainment in PSID.

E Additional Figures

Figure 10: Aggregate consumption-to-wealth ratio over time



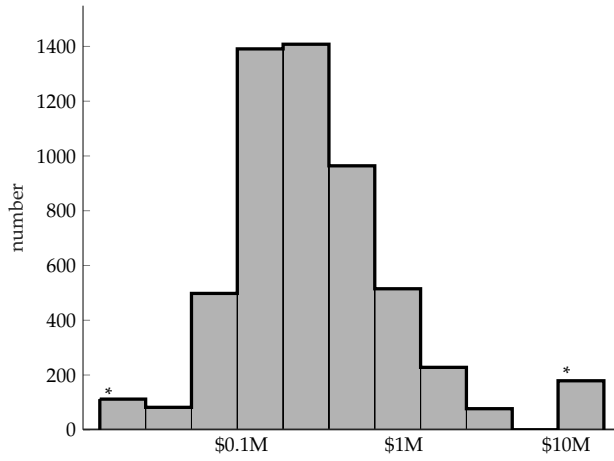
Notes: Aggregate consumption is the sum of consumption expenditures for all households in our sample. Aggregate wealth is the sum of proxy wealth for all households in our sample.

Figure 10 plots total consumption expenditures relative to proxy wealth in our data. On average, households spend around 3% of their total proxy wealth in the present. The aggregate consumption to wealth ratio rose during the great recession, as wealth shrank relative to consumption, but fell during the stock market boom in the mid 2010s. Over the whole sample, the aggregate consumption-wealth ratio fell.

Figure 11 displays the distribution of money-metric wealth, in 2019, for all age groups, including rentiers and non-rentiers. Figure 12 plots nominal log money-metric wealth against nominal capitalized (proxy) wealth for rentiers and non-rentiers. If financial markets are complete and absent measurement error, these two figures should be 45 degree lines.⁴⁰

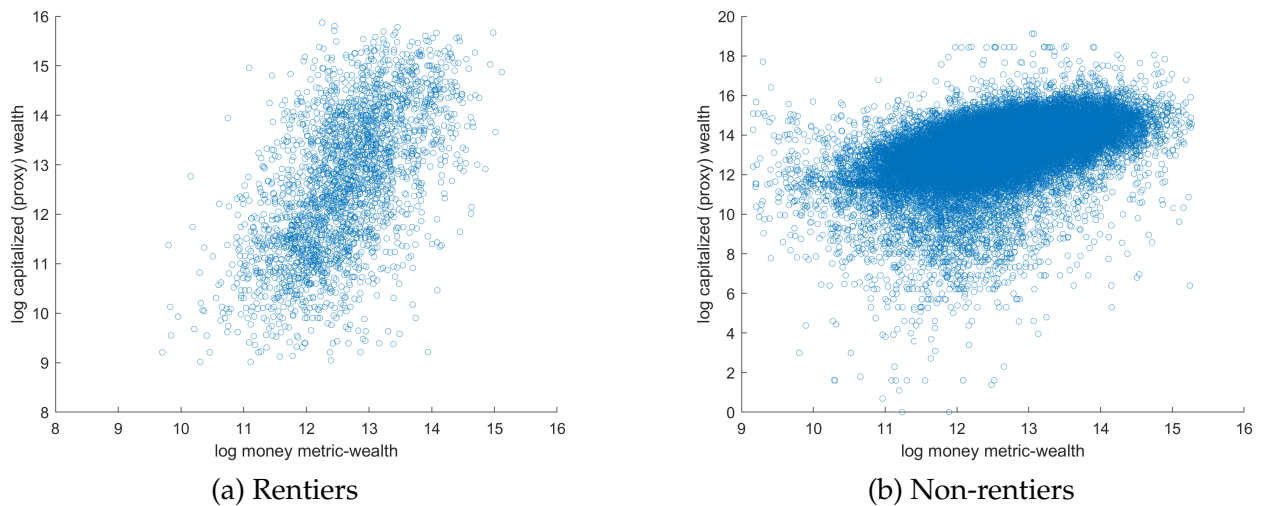
⁴⁰See Appendix D for more details.

Figure 11: Distribution of money-metric wealth in 2019



Notes: Bars with asterisk are households whose estimated money-metric wealth in \$2019 dollars is outside of the wealth range of rentiers in that year.

Figure 12: Capitalized (proxy) wealth against money-metric wealth (in logs) for all years



Notes: All variables in this figure are in nominal terms (not base year prices).