A Sufficient Statistics Approach to Measuring Forward-Looking Welfare

David R. Baqaee UCLA

Ariel Burstein UCLA

Yasutaka Koike-Mori UCLA*

[Preliminary and Incomplete]

Abstract

We provide a method to measure forward-looking consumer welfare, in money metric terms, without imposing parametric assumptions about beliefs or the utility function beyond time separability. Forward-looking consumer welfare takes into account beliefs about the future. Our methodology allows for non-rational expectations, non-homotheticicities (within and across time), uninsurable risk, incomplete markets, and borrowing constraints and does not require us to estimate beliefs, contingent plans, or any parameters of the utility function besides the intertemporal elasticity of substitution. Our method requires goods' prices over time and repeated cross-sectional household expenditures, financial wealth, and income data. We apply our method to the Panel Study of Income Dynamics from the United States, and use our results to study growth and inflation between 2005 and 2019. We find that forward-looking measures of welfare growth tend to be lower than standard static measures like real consumption growth. Our methodology is well-suited to studying the welfare consequences of complex shocks that affect households along different margins and time horizons. As an example, we study the welfare effect of job loss and find that involuntary job loss is associated with a 7% reduction in money metric utility.

*We thank Andy Atkeson, Valentin Haddad, Oleg Itskhoki, Pablo Fajgelbaum, Kieran Larkin, and Pierre-Olivier Weill for comments. We acknowledge financial support from NSF grant No. 1947611 and the Alfred P. Sloan Foundation.

1 Introduction

Measuring consumer welfare is a central task of economics. In macroeconomics, it is widely recognized that a comprehensive measure of welfare must be forward-looking — it must account not just for present consumption, but also expected future consumption. The typical structural macroeconomic model explicitly models intertemporal decision-making, and welfare calculations, as in Lucas (1987), take into account both time and uncertainty.

In this sense, theory is far ahead of measurement practice. Both at the individual and the aggregate level, standard statistical measures of wellbeing, like real income or real consumption, are founded on a static and deterministic formulation of price theory. For example, the *Consumer Price Index* manual (IMF, 2004), a 500-page tome detailing the theory and practice of constructing consumer price indices and measures of consumer welfare, has almost nothing to say about the theory or the practice of how expectations about the future should influence the measurement of welfare and inflation.

Static measures are constructed by deflating nominal static income or consumption by a static price deflator. Such measures do not account for how expectations about the future affect consumption decisions. This means that they can easily produce misleading estimates of changes in welfare. If forward-looking measures of welfare and inflation are so theoretically appealing, why then are they almost totally unheard of outside of fully specified structural models? The main difficulty is that much of what one needs to observe in order to compute a measure of welfare is not observable. Static measures require knowing expenditures on and prices of goods and services in different periods. In contrast, dynamic measures require knowing future state-contingent consumption plans, future state-contingent goods and asset prices, and probabilities of different states being realized. This intractable measurement problem is perhaps the central reason why forward-looking measures of real wealth have remained firmly in the domain of economic theory rather than measurement practice.

This state of affairs, where structural work uses forward-looking wealth-like measures of welfare but applied empiricists use income-like measures, has long been a source of tension in economics. For example, Samuelson (1961) concludes with this:

> "When we work with simple and exact models, in which no extraneous statistical difficulties of measurement could arise, the only valid measure of welfare comes from compute *wealth-like* magnitudes not income magnitudes. In the absence of perfect certainty, 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. But operational practicality aside, if the theorist specifies in detail the dynamic technology of his model, he will meet none of the pitfalls that come from an attempt to summarize his model by various crude approximations. ... I must stress and restress that although wealth-like magnitudes offer no difficulties in theoretical principle as compared with the statical case, the national income statistician is very far from having even an approximation to the data needed for these comparisons. A vital difficulty is the hard and unchangeable fact of uncertainty. Futures markets might enable us to salvage something even in the presence of uncertainty; but futures markets are themselves of little quantitative importance in present-day economies."

This paper develops a new approach to overcome some of these conceptual and measurement challenges. Since variables involving the future are missing, we back them out from the household's consumption-savings choices, conditional on estimates of the elasticity of intertemporal substitution (EIS). We do not put parametric restrictions on preferences, but our results do rely critically on the assumption that preferences are indirectly separable over time and that elasticity of intertemporal substitution is different from one.

Since our methodology does not require researchers to fully spell out a structural model of preferences, beliefs, and prices, it is well-suited for use in reduced-form empirical work for understanding the welfare effects of complex shocks. Many policies and shocks have complex effects that affect households along many dimensions and at different time horizons. Our paper provides a way to study the welfare effects of treatments on households without requiring that researchers enumerate and estimate all the possible ways the treatment affects the household and how the household's beliefs and contingent plans change in response to the treatment.

We sketch the basic idea of our approach. Consider a perfect foresight consumption-savings problem with complete markets and homothetic time-separable preferences. In this case, changes in the continuation value of the future can be backed out from changes in the consumption-wealth ratio given knowledge of the EIS. If the EIS is less than one, then an increase in the consumption-wealth ratio suggests that inflation in the future bundle is falling relative to inflation in the present bundle. This means that forward-looking measures of inflation that account for how all prices, not just contemporaneous prices, change will be lower than static measures of inflation. Accordingly, real wealth will be higher than what is implied by a static measure.

This basic idea continues to be useful in more complex environments. To show this, we first generalize the definition of money metric utility from its static deterministic formulation (see e.g., Samuelson, 1974) to an environment with incomplete markets and uncertainty. We then consider households that make dynamic decisions facing incomplete financial markets with borrowing constraints.

We define a household to be a *rentier* if it possess negligible non-financial assets. The consumption-wealth ratio of rentiers can change due to substitution effects or wealth effects. The compensated consumption-wealth ratio, which neutralizes wealth effects, responds only due to substitution effects. If we can identify the compensated consumption-wealth ratio, then we can use it to back out changes in the marginal value of consumption today relative to the future. Given this information, we can then construct money metric utility for rentier households. But given the money metric utility function, we can identify the compensated consumption-wealth ratio. As in Baqaee et al. (2023), this is a fixed point problem with a unique solution.

If preferences are non-homothetic within the period, then we can extend our estimates of money metric utility to non-rentiers as well. To do this, we rely on a generalization of Engel's law. Specifically, if the vector of budget shares is a one-to-one function of utility, then two households are on the same intertemporal indifference curve if, and only if, their budget shares in the same period are the same. This allows us to construct money-metric values for financially constrained households by matching their static budget shares to the wealth of rentiers.

To apply our results, we require 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 over time. Finally, we require knowledge of the compensated elasticity of intertemporal substitution (which could be a constant or vary as a function of wealth and time). Our theoretical results also assume that households' preference relations are stable functions of observable characteristics. That is, we rule out unobservable taste shocks, where households with similar characteristics have different preferences.

As a proof of concept, we apply our theoretical results to the United States using the Panel Study of Income Dynamics (PSID) and price data from the Bureau of Labor Statistics. We begin by selecting a subsample of rentiers. We do this by computing a proxy for the present value of expected future labor and transfer income for each household. We say that a household is a rentier if the present value of their future labor and transfer income is less than 10% of their total wealth. These households tend to be older and more wealthy. We apply our methodology to recover an ideal cost-of-living index for rentiers. We then extend the money metric to cover non-rentiers by matching households in the same period together via static budget shares as described above.

We find that static CPI measures are a poor approximation to the true dynamic cost-of-living price index. Forward-looking inflation is higher, and growth is lower, than static measures. Furthermore, we find significant heterogeneity in the cost-of-living index across both the wealth and age distribution.

We then apply our method to study the welfare consequences of job loss in our sample. We find that involuntary job loss is associated with a roughly 7% reduction in money metric utility. That is, households that lose their jobs experience a reduction in their lifetime equivalent wealth of roughly 7%. Unlike alternative numbers from the literature, this number does not assume that all households discount the future using market rates, or even that households have rational expectations about the future.

Related Literature. The literature on consumer price indices is vast but the majority of it abstracts from time. The literature on dynamic price indices is comparatively small, but can trace its origins to the inception of consumer price indices. For example, Fisher and Pigou both recognized that the ideal cost-of-living index should incorporate information about the future, though they did not offer a specific remedy. Alchian and Klein (1973) argue that a proper definition of the price index must be based on intertemporal consumption, and they propose including asset prices in the CPI. Pollack (1975) studies conditions under which the intertemporal cost-of-living index can be broken up into sub-indices. In the context of national income accounting, Hulten (1979) points out that productivity shocks today drive capital accumulation in the future, and so the Solow (1957) residual understates the importance of technological change. He proposes to use interest rates to calculate a dynamic technology residual. Relatedly, 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.

To measure dynamic measures of inflation and welfare, Reis (2005) and Aoki and Kitahara (2010) calibrate parametric models of household preferences, beliefs, and compute aggregate cost-of-living indices by feeding in the path of observed prices. Reis (2005) uses additively time-separable homothetic 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.

More recently, Fagereng et al. (2022) and Del Canto et al. (2023) use Taylor approximations, around the perfect foresight steady-state, to calculate how consumer welfare responds to changes in asset prices and monetary policy, respectively. Fagereng et al. (2022) estimate how various asset prices changed over time in Norway and weigh these changes in asset prices by discounted net holdings. Del Canto et al. (2023) estimate local projections of how monetary shocks change goods and asset prices, and then weigh these price changes by discounted budget shares of households. Both of these papers construct a dynamic price index in a way that is analogous to how static price indices are calculated — weighing price changes by budget shares — except they estimate how future (discounted) budget shares and future (discounted) prices will change.

Our paper differs from all of these papers in some important ways. First, we do not directly attempt to estimate or model future asset or goods prices, beliefs, discount factors, and future holdings of assets or purchases of goods. Instead, we use changes in consumption-savings decisions to back out implied changes instead. Second, our approach does allow some households to have non-infinitesimal idiosyncratic and undiversifable labor income risk, which is ruled out in all previous work. Finally, we allow for preferences to be non-homothetic both across and within each period. We document strong intertemporal non-homotheticities whereby wealthier households consume a much smaller share of their wealth per period. Indeed, we take advantage of these non-homotheticities to deal with uninsurable risk.

Finally, our paper is also related to generalizations of price theory to incomplete markets, like Farhi et al. (2022), who decompose changes in demand in an incomplete market world into income and substitution effects. As pointed out by Baqaee and Burstein (2023), money metric measures of welfare require knowledge of compensated demand. That is, money metric measures of welfare, and by extension ideal price indices, treat income and substitution effects differently. We extend the complete market analysis in that paper to a world with incomplete markets and show that a similar intuition persists.

Our paper builds on tools from the literature on static price indices. Feenstra (1994) inverts CES demand curves to infer the value of new goods and other missing prices.[1] Baqaee et al. (2023) generalize this approach to allow for separable non-homothetic and non-CES preferences. In this paper, we build on this idea to infer the value of the future from consumption-saving choices. As in Hamilton (2001), we use budget shares to match households who are on the same indifference curve.[2] In contrast to Hamilton (2001), we

[1]This approach is frequently used to infer the value of new goods or quality change in static settings, see, for example Broda and Weinstein (2010) and Aghion et al. (2019).

[2]This 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).

match households in the cross-section within time periods, as opposed to across time periods. Finally, we apply results from Baqaee et al. (2023) who show that money metric utility can be represented as the solution to a fixed point problem. Our paper is also related to recent contributions by Atkin et al. (2020) and Jaravel and Lashkari (2022) who analyze static price indices with non-homotheticities and cross-sectional data.

Roadmap. Section 2 defines the compensated elasticity of intertemporal substitution, time separability of preferences, and introduces useful notation. Section 3 uses a simple complete-markets example, with additively time-separable utility, to illustrate how changes in consumption-wealth ratios can be used to infer welfare changes. Section 4 introduces the actual decision problem we are interested in, which features incomplete markets, uninsurable risk, borrowing constraints, and non-homothetic preferences. Section 5 contains the main results of the paper showing how to extend the basic idea in Section 3 to this more complex environment. Section 6 constructs a measure of dynamic welfare for households in the PSID. Section 7 uses these measures to study the dynamic welfare losses from job loss. We conclude in Section 8.

2 Time Separable Preferences

In this section, we define preferences and introduce a notion of time separability that we make use of throughout the paper. This section only describes preferences, and we delay specifying the decision problem to later sections. Agents have preferences over consumption goods represented by a utility function

$$\mathcal{U}\left(\{\boldsymbol{c}, \boldsymbol{\pi}\}\right), \tag{1}$$

where $\boldsymbol{c}$ is a state-contingent stream of consumption bundles and $\boldsymbol{\pi}$ is the probability distribution over states. Denote the sequence of shock realizations up to period t by $s^t = (s_0, \ldots, s_t)$, the set of goods available each period by N, and the consumption of good n in history s^t by $c_n(s^t)$. Let $\pi(s^t)$ denote the probability of drawing history s^t. The preferences in (1) nest the common exponential discounting expected utility function as a special case.

Given some preferences (1), define the following transformation of the utility function:

$$e(\boldsymbol{q}, \boldsymbol{\pi}, U) = \min_{\boldsymbol{c}} \left\{\boldsymbol{q} \cdot \boldsymbol{c} : \mathcal{U}\left(\{\boldsymbol{c}, \boldsymbol{\pi}\}\right) = U\right\}, \tag{2}$$

where $\boldsymbol{q}$ has the same dimensionality as $\boldsymbol{c}$. This transformation takes as an argument a utility function $\boldsymbol{\mathcal{U}}$ and probabilities $\boldsymbol{\pi}$ and returns a new function e. We refer to e as the *shadow* intertemporal expenditure function and to $\boldsymbol{q}$ 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 (2) in practice.

Write $q_n(s^t)$ for the element of $\boldsymbol{q}$ corresponding to $c_n(s^t)$. Denote the maximizers in (2) by $c_n^*(s^t|\boldsymbol{q}, \boldsymbol{\pi}, U)$, and the share of spending in the first period, the present, to be

$$b^P(\boldsymbol{q}, \boldsymbol{\pi}, U) = \frac{\sum_{n\in N} q_n(s^0) c_n^*(s^0|\boldsymbol{q}, \boldsymbol{\pi}, U)}{e(\boldsymbol{q}, \boldsymbol{\pi}, U)}.$$

Define the share of savings for the future to be

$$b^F(\boldsymbol{q}, \boldsymbol{\pi}, U) = 1 - b^P(\boldsymbol{q}, \boldsymbol{\pi}, U).$$

Definition 1 (Compensated Elasticity of Intertemporal Substitution)**.** The compensated elasticity of intertemporal substitution (EIS), $\sigma^*(\boldsymbol{q}, \boldsymbol{\pi}, U)$, is

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

That is, the compensated EIS captures how spending on consumption versus savings changes if the shadow price of every consumption good in the present rises by the same amount, holding overall utility constant.

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

Definition 2 (Time Separability)**.** Preferences are (indirectly) *time separable* if the function defined in (2) can be written as

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

where P and F are increasing and homogeneous of degree one in $\boldsymbol{q}$ and non-decreasing in U.

That is, preferences are time separable if, holding utility constant, the shadow expenditure function e is separable between present (superscript P) and future (superscript F) shadow prices. When preferences are homothetic, as is the case in the vast majority of the macroeconomics literature, indirect time separability is equivalent the common assumption that the utility function is separable between present and future consumption quantities.

Definition 3 (Homothetic Separability)**.** If preferences are homothetic, then indirectly time separability is equivalent to direct time separability. That is, the utility function can be written as

$$\mathcal{U}(\boldsymbol{c},\boldsymbol{\pi}) = U\left(\mathcal{P}\left(\boldsymbol{c}(s^0)\right), \mathcal{F}\left(\left\{\boldsymbol{c}(s^t)\right\}_{t>0}, \boldsymbol{\pi}\right)\right),$$

where U is increasing and homogenous of degree one in its arguments, and $\mathcal{P}$ and $\mathcal{F}$ are increasing and homogeneous of degree one in $\boldsymbol{c}$.

Proof. This follows from Theorem 4.3 in Blackorby et al. (1998). ■

Definition 3 implies that the typical preferences commonly used in macroeconomics, like exponentially discounted expected utility or Epstein and Zin (1989), satisfy indirect time separability since they are homothetic and directly time separable.

The following is a parametric example that is indirectly time separable but non-homothetic.

Example 1 (Exponential Discounting Expected Utility)**.** Preferences that take the form

$$1 = \sum_{t=0} \beta^t \sum_{s_t} \pi(s^t) \left(\sum_{n} \omega_{nt} U^{\varepsilon_n} \left(c_n(s^t) \right)^{1-\gamma_n(s^t)} \right)^{\rho_t} \tag{3}$$

are indirectly time separable.

We impose time separability of preferences throughout the rest of the paper. Note that time separability is a primitive condition on preferences, not on the decision problems agent face. That is, preferences can be indirectly time separable without agents facing the decision problem defined by (2) in practice.

3 An Illustrative Example

In this section, we consider a highly restrictive special case that demonstrates one of the key ideas of this paper. We make some very strong assumptions: we assume there is only one consumption good in each period, there is no uncertainty, preferences are homothetic with constant relative risk aversion, and financial markets are complete. The fact that financial markets are complete implies that households do not have idiosyncratic uninsurable risk like non-pledgeable labor income. These are assumptions we relax in Section 4.

Suppose that preferences, in (1), take the form

$$\mathcal{U}(\boldsymbol{c}) = \sum_{t=\tau}^{T} \beta^t \left(c(t)\right)^{\frac{\sigma-1}{\sigma}}, \tag{4}$$

where T is the agent's horizon (potentially infinite), β is the discount factor, and σ is the compensated elasticity of intertemporal substitution. Calendar time is indexed by τ. Agents face the budget constraint

$$p(t|\tau)c(t|\tau) + a(t+1|\tau) = R(t|\tau)a(t|\tau)$$

each period, where $c(t|\tau)$ and $a(t+1|\tau)$ are consumption and savings, $R(t|\tau)$ is the gross rate of return on savings, and $p(t|\tau)$ is the price of the consumption good in period $\tau + t$. All variables depend on τ, which indexes the first period and corresponds to calendar time. To keep the problem well-defined, savings in the final period cannot be negative: $a(T|\tau) \geq 0$. Denote initial wealth by $w = a(0|\tau)$. Households choose consumption and savings to maximize utility, taking prices and returns as given.

Consider an overlapping generations structure — at each calendar date $\tau \in \mathbb{R}$, there is a cohort of agents with planning horizon T facing some vector of intertemporal prices and returns $\{p(t|\tau), R(t|\tau)\}_{t=0}^{T}$. To keep the notation simple, we assume that each cohort makes discrete time decisions, $t \in \{0, \ldots, T\}$, but that cohorots and prices evolve in continuous time $\tau \in \mathbb{R}$.[3] The value function depends on τ, which determines the path of prices and rates of return agents face, as well as initial wealth w. Denote this value function by $V(\tau, w)$.

Our objective is to compare changes in welfare for different cohorts of households. To do this, we use the money metric, which is a popular way to measure welfare in microeconomics (see Deaton and Muellbauer, 1980).[4] The money metric converts wealth for a cohort at date τ into equivalent wealth under a common base date τ_0. That is, for each τ and w, we define a *money metric* function $u(\tau, w)$ that solves

$$V(\tau, w) = V(\tau_0, u(\tau, w)).$$

To compare the wealth of two households in different cohorts, (τ, w) and (τ', w'), we can compare $u(\tau, w)$ and $u(\tau', w')$. The constraint set defined by (τ', w') is preferred to (τ, w) if, and only if, $u(\tau', w')$ is higher than $u(\tau, w)$. Since both of these numbers are in terms of

[3] All our results extend in an obvious way if households make decisions in continuous time instead.

[4] The money metric was introduced by McKenzie (1957), but popularized by Samuelson (1974).

τ_0 dollars, we can also meaningfully compare their magnitude and calculate the rate of growth between (τ, w) and (τ', w') as $u(\tau', a_0')/u(\tau, a_0)$.

Since financial markets are complete, we know that the shadow intertemporal prices for consumption at t for a household in cohort τ is just

$$q(t|\tau) = \left(\prod_{x=0}^{t} R(x|\tau)^{-1}\right) p(t|\tau).$$

The wealth required in τ_0 to reach the same indifference curve as $V(\tau, w)$ is just:

$$u(\tau, w) = e(\boldsymbol{q}(\cdot|\tau_0), V(\tau, w)) = e(\boldsymbol{q}(\cdot|\tau_0))V(\tau, w) = w\frac{e(\boldsymbol{q}(\cdot|\tau_0, 1))}{e(\boldsymbol{q}(\cdot|\tau, 1))}, \tag{5}$$

where the first equality follows from the definition, the second follows from the fact that preferences are homothetic, and the third equality follows from the fact that $w = e(q(\cdot|\tau))V(\tau, w)$. In words, (5) shows that the money metric $u(\tau, w)$ deflates nominal wealth w at date τ using a price index, $e(\boldsymbol{q}(\cdot|\tau_0, 1))/e(\boldsymbol{q}(\cdot|\tau, 1))$, between τ and τ_0. This price index depends on how prices and rates of return change between τ and τ_0.

Denote the discounted expenditures of households in cohort τ in each period t by

$$b(t|\tau) = \left(\prod_{x=0}^{t} R(t|\tau)^{-1}\right) \frac{p(t|\tau)c(t|\tau, w)}{w}.$$

These discounted budget shares are not a function of wealth since preferences are homothetic. Using the fundamental theorem of calculus and Shephard's lemma, rewrite (5) as

$$\log u(\tau, w) = \log w - \int_{\tau_0}^{\tau} \sum_{t=0}^{T} b(t|\xi) \left[\frac{d \log p(t|\xi)}{d\xi} - \sum_{x=0}^{t} \frac{d \log R(x|\xi)}{d\xi}\right] d\xi. \tag{6}$$

In words, order cohorts between τ and τ_0. For each cohort ξ between τ and τ_0, calculate the rate of change in prices $d \log p(t|\xi)$ and returns $d \log R(t|\xi)$ cohort ξ faces and average these changes using discounted budget shares for each period of life t. Cumulate these changes over all cohorts ξ between τ_0 and τ to arrive at the cumulative change in the cost-of-living between τ_0 and τ and use this to deflate nominal wealth at τ.

In a static world, where the horizon is just $T = 1$, (6) simplifies to a trivial calculation:

$$\log u(\tau, w) = \log w - \int_{\tau_0}^{\tau} \frac{d \log p(0|\xi)}{d\xi} d\xi = \log w - \log \frac{p(0|\tau)}{p(0|\tau_0)}.$$

That is, in a static world, we simply need to deflate nominal wealth by the change in the price of the consumption good between τ and τ_0. However, even in the simplest dynamic setting ($T > 1$), with no uncertainty, complete markets, homothetic preferences, and a single consumption good in each period, measuring (6) directly is difficult. This is because it requires knowing discounted future budget shares, future prices, and future rates of return for each cohort.

One of the key ideas in this paper is to infer changes in future rates of returns and future prices by relying on changes in consumption-savings choices. Changes in the consumption-wealth ratio over time depend on the change in the price of the consumption good in period 0 relative to a budget-share weighted average of the price of consumption across all future dates:

$$\frac{d\log b(0|\tau)}{d\tau} = (1-b(0|\tau))(1-\sigma)\left[\frac{d\log p(0|\tau)}{d\tau} - \sum_{t>0}^{T}\frac{b(t|\tau)}{1-b(0|\tau)}\left[\frac{d\log p(t|\tau)}{d\tau} - \sum_{x=0}^{t}\frac{d\log R(x|\tau)}{d\tau}\right]\right].$$

This expression can be rearranged to solve for the budget-share weighted changes in future prices:

$$\sum_{t>0}^{T} b(t|\tau)\left[\frac{d\log p(t|\tau)}{d\tau} - \sum_{x=0}^{t}\frac{d\log R(x|\tau)}{d\tau}\right] = (1-b(0|\tau))\frac{d\log p(0|\tau)}{d\tau} - \frac{1}{1-\sigma}\frac{d\log b(0|\tau)}{d\tau}.$$

Substitute this into (6) to arrive at the following simple result.

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:*

$$\log u(\tau, w) = \log w - \log\frac{p(0|\tau)}{p(0|\tau_0)} - \frac{1}{\sigma-1}\log\frac{b_0(\tau)}{b_0(\tau_0)}. \tag{7}$$

To understand Proposition 1, suppose that $\sigma < 1$ so that consumption goods in different periods are complements. In this case, if cohort τ is saves a smaller fraction of wealth than cohort 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. This allows us to back out the change in future prices, comparing cohort τ and τ_0, using the change in savings rate and the elasticity of intertemporal substitution. The bigger the difference in savings rates between the two cohorts, the bigger is the difference in the future prices relative to current prices.

To arrive at (7), we made some very strong assumptions. We assumed that there is only one consumption good in each period, we assumed away uncertainty, we assumed that financial markets are complete, we assumed that household preferences are homothetic, and we assumed that the utility function exhibits constant relative risk aversion (CRRA). The rest of the paper is devoted to relaxing all of these assumptions and showing that when these assumptions are relaxed, the basic intuition behind (7) still applies and can be used to measure money metric growth and inflation.

Proposition 1 can also be derived as a consequence of Feenstra (1994). Our proof is lengthier than the one in Feenstra (1994), however, our proof generalizes when we relax the assumptions of this example. See Appendix D for more discussion.

4 Environment and Measure of Welfare

In this section, we set up the economic environment. We relax the assumption that preferences are homothetic and CRRA, that financial markets are complete, that all assets are pledgeable, and that agents have perfect foresight. We generalize the notion of money metric welfare to this environment and use it to define our measure of welfare and the cost-of-living.

4.1 Decision Problem of Households

Decision makers face a planning horizon of length $T < \infty$. Preferences, the money metric, and choices are all indexed by the length of the planning horizon T, which reflects the household's age. Welfare comparisons are carried out holding T constant. To streamline notation, we omit the dependence on T.

Let the first date be τ. We index the consumer's decision problem using the start date. That is, define $s^t(\tau)$ to be history of shock realizations t periods after the start date τ. Let $\pi(s^t|\tau)$ be the probability of history s^t being realized conditional on starting at τ. Let $\boldsymbol{c}(s^t|\tau) \in \mathbb{R}^N$ be the vector of consumption goods in history s^t conditional on starting at τ.

Consumers choose their consumption decisions and portfolio of assets to maximize utility (1) subject to a sequence of state-contingent budget constraints. Denote the price of good $n \in N$ in period t given history s^t with initial condition τ by $p_n(s^t|\tau)$. The first period budget constraint is

$$\sum_{n\in N} p_n(s^0|\tau)c_n(s^0|\tau) + \sum_{k\in K} a_k(s^0|\tau) = w, \tag{8}$$

where, with some abuse of notation, $p_n(s^0|\tau)$ and $c_n(s^0|\tau)$ is the price and consumption of good n in the initial period, which we assume are known with certainty. There are K different asset types and the quantity the household chooses to purchase is denoted by $a_k(s^0|\tau)$. The price of every asset is normalized to be one, and the scalar w is the initial wealth.

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

$$\sum_{n\in N} p_n(s^t|\tau)c_n(s^t|\tau) + \sum_{k\in K} a_k(s^t|\tau) = \sum_{k\in K} R_k(s^t|\tau)a_k(s^{t-1}|\tau) + y(s^t|\tau), \tag{9}$$

where $R_k(s^t|\tau)$ is the return of asset k in history s^t and $y(s^t|\tau)$ is an exogenous payoff. We think of $y(s^t|\tau)$ as a the flow of payments from assets that cannot be traded. If $y(s^t|\tau) = 0$ for every s^t, we say that the household is a rentier.

We also impose borrowing constraints requiring that

$$\sum_{k} a_k(s^t|\tau) \geq -X(s^t|\tau) \tag{10}$$

for some exogenous state-contingent borrowing constraint $X(s^t|\tau) \geq 0$. We require that $X(s^T) = 0$ for every s^T to ensure the agent cannot end the problem in debt.

The decision problem faced by households are indexed by the tuple of prices, returns, probabilities, borrowing constraints, and wealth: $\{\boldsymbol{p}, \boldsymbol{R}, \boldsymbol{\pi}, \boldsymbol{X}, w\}$. Define the value function associated with each problem to be

$$V\left(\{\boldsymbol{p}, \boldsymbol{R}, \boldsymbol{y}, \boldsymbol{\pi}, \boldsymbol{X}, w\}\right) = \max_{\boldsymbol{c},\boldsymbol{a}} \left\{\boldsymbol{\mathcal{U}}\left(\{\boldsymbol{c}, \boldsymbol{\pi}\}\right) : \text{constraints (8), (9), (10) are satisfied}\right\}. \tag{11}$$

The value function ranks decision problems according to underlying preference relation. In a static, deterministic environment, the value function in (11) collapses to the indirect utility function in consumer theory, which ranks static budgets sets, defined by static prices and wealth, into utils.

4.2 Measuring Welfare and the Cost-of-Living

We measure welfare using a notion of the money metric generalized to allow for forward-looking decisions, uncertainty, and incomplete markets.

Definition 4 (Dynamic Money Metric)**.** Consider a reference period τ_0, with reference prices, returns, and probabilities about the future: $\{\boldsymbol{p}(\cdot|\tau_0), \boldsymbol{R}(\cdot|\tau_0), \boldsymbol{\pi}(\cdot|\tau_0)\}$, with $\boldsymbol{p}(\cdot|\tau_0) > 0$ and $\boldsymbol{R}(\cdot|\tau_0) > 0$. The *money metric*, in τ_0 dollars, associated with a decision problem $\{\boldsymbol{p}, \boldsymbol{R}, \boldsymbol{y}, \boldsymbol{\pi}, \boldsymbol{X}, w\}$ is a scalar-valued function u that satisfies the following equation

$$V\left(\{\boldsymbol{p}, \boldsymbol{R}, \boldsymbol{y}, \boldsymbol{\pi}, \boldsymbol{X}, w\}\right) = V\left(\{\boldsymbol{p}(\cdot|\tau_0), \boldsymbol{R}(\cdot|\tau_0), \boldsymbol{\pi}(\cdot|\tau_0), \boldsymbol{0}, \boldsymbol{0}, u\}\right).$$

In words, the money metric, u, maps the decision problem $\{\boldsymbol{p}, \boldsymbol{R}, \boldsymbol{y}, \boldsymbol{\pi}, \boldsymbol{X}, w\}$ into the equivalent one-off lump-sum payment the household would need, under the baseline τ_0, to ensure indifference. Denote this money metric by $u\left(\{\boldsymbol{p}, \boldsymbol{R}, \boldsymbol{y}, \boldsymbol{\pi}, \boldsymbol{X}, w\}|\tau_0\right)$. In a static deterministic world, the generalized money metric coincides with the traditional money metric.[5]

In classical consumer theory, the money metric cardinalizes utility. The same holds for the generalized notion of the money metric defined above.

Proposition 2 (Money metric cardinalizes utility)**.** *The money metric is a cardinalization of the value function.*

We require that $\boldsymbol{p}(\cdot|\tau_0) > 0$ and $\boldsymbol{R}(\cdot|\tau_0) > 0$ to ensure the value function is well-defined under baseline prices and returns.

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

Definition 5 (Dynamic Cost-of-Living)**.** Consider two cohorts τ and τ', each with reference prices, returns, and probabilities about the future. Define the change in the cost-of-living between τ and τ' for a household facing problem $\{\boldsymbol{p}(\cdot|\tau), \boldsymbol{R}(\cdot|\tau), \boldsymbol{y}, \boldsymbol{\pi}(\cdot|\tau), \boldsymbol{X}(\cdot|\tau), w\}_{s^t}$ in cohort τ to be

$$\frac{u\left(\{\boldsymbol{p}(\cdot|\tau), \boldsymbol{R}(\cdot|\tau), \boldsymbol{y}, \boldsymbol{\pi}(\cdot|\tau), \boldsymbol{X}(\cdot|\tau), w\}|\tau'\right)}{u\left(\{\boldsymbol{p}(\cdot|\tau), \boldsymbol{R}(\cdot|\tau), \boldsymbol{y}, \boldsymbol{\pi}(\cdot|\tau), \boldsymbol{X}(\cdot|\tau), w\}|\tau\right)}.$$

That is, we use the money metric to convert $\{\boldsymbol{p}(\cdot|\tau), \boldsymbol{R}(\cdot|\tau), \boldsymbol{y}, \boldsymbol{\pi}(\cdot|\tau), \boldsymbol{X}(\cdot|\tau), w\}_{s^t}$ into equivalent lump-sum payments in τ and τ' and compare the ratio of these numbers. In a static deterministic environment, the change in the cost-of-living collapses to the traditional ideal (Konüs) price index of consumer theory.

The objective in this paper is to infer the money metric u by combining cross-sectional survey data on household consumption and finances along with goods and services prices over time.

5 Main Results

We present our main result in steps. First, we start with some preliminaries in Section 5.1. In Section 5.2 we present a method for recovering the generalized money metric for households whose non-market cashflows are negligible. In Section 5.3 we show how to recover the money metric for households whose non-market cashflows are not negligible.

[5]In principle, there are many different ways one could measure welfare. For example, we could convert each problem into an certainty equivalent annuity value. We focus on equivalent one-off lump-sum payments because this is the welfare measure that we can recover from the data given our assumptions.

5.1 Preliminary Result

In this section, we establish that for every decision problem there is a corresponding dual shadow expenditure minimization problem where the shadow prices and shadow wealth rationalize the household's consumption choices. This duality is useful because it allows us to define the notion of a "compensated" elasticity of intertemporal substitution and "compensated" budget shares. These objects are important in allowing us to recover the generalized money metric.

Since prices, returns, and budget constraints are indexed by the initial condition, τ, with some abuse of notation, we write the value function as $V(\tau, w, \boldsymbol{y})$, where τ indexes the goods and asset prices and budget constraints, given initial wealth w and state-contingent cashflows from non-marketable assets $\boldsymbol{y}$. The next proposition shows that for every decision problem $(\tau, w, \boldsymbol{y})$, there exists a set of shadow prices $\boldsymbol{q}^*(\tau, w, \boldsymbol{y})$ that rationalize the allocations in (11).

Proposition 3 (Dual Problem)**.** *There exist $\boldsymbol{q}^*(\tau, w, \boldsymbol{y})$ such that, for every s^t and n, consumption choices are the same*

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

Moreover, we can set shadow prices for goods in the first period equal to their observed prices:

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

for every $n \in N$.

In other words, if the household faced shadow prices $\boldsymbol{q}^*(\tau, w, \boldsymbol{y})$ and minimized shadow expenditures subject to a utility constraint, then the consumption plan the household would choose coincides with the ones they choose when facing (11).

In a static deterministic environment, $c_n(s^t|\tau, w, \boldsymbol{y})$ collapses to uncompensated (Marshallian) demand for good n. On the other hand, the shadow quantity $c_n^*(s^t|\boldsymbol{q}^*, \boldsymbol{\pi}, V(\tau, w, \boldsymbol{y}))$ collapses to compensated (Hicksian) demand for good n when the environment is static and deterministic. One of the important differences between c_n^* and standard Hicksian demand is that it depends on shadow prices, rather than actual prices, and these shadow prices, in principle, depend on the household's indifference curve.

Proposition 3 makes it possible to define a notion of *compensated* elasticity of intertemporal substitution for an agent facing the problem $(\tau, w, \boldsymbol{y})$.

Definition 6 (Elasticity of Intertemporal Substitution)**.** The compensated EIS for a household facing problem $(\tau, w, \boldsymbol{y})$ is defined to be

$$\sigma(\tau, w, \boldsymbol{y}) = \sigma^*(\boldsymbol{q}^*(\cdot|\tau, w, \boldsymbol{y}), \boldsymbol{\pi}(\cdot|\tau), V(\tau, w, \boldsymbol{y})),$$

where q^* are shadow prices given in Proposition 3.

That is, the compensated EIS for a household facing $(\tau, w, \boldsymbol{y})$ is defined to be how spending on consumption versus savings changes, for this household, if the shadow price of every consumption good in the present rises by the same amount, holding utility constant. The compensated EIS is a crucial statistic that we will need for our main results.

5.2 Recovering Money Metric for Rentiers

The next proposition limits attention to a subset of households whose entire wealth can be bought or sold on financial markets. We call these households *rentiers*. Before stating the proposition, we introduce some notation.

For each household with wealth w, cashflows $\boldsymbol{y}$, in period τ, denote current expenditures by

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

and the share of budget spent on good n by

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

Denote the consumption to wealth ratio by

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

Consider the subset of households for whom $\boldsymbol{y}$ is identically equal to zero $\mathbf{0}$ in every state of nature and denote them by $(\tau, w, \mathbf{0})$.

Proposition 4 (Money Metric for Rentiers). *If preferences are time separable and $\sigma(t, w, \mathbf{0}) \neq 1$ almost everywhere, then the money metric satisfies the following integral equation*

$$\log u(\tau, w, \mathbf{0}) = \log w - \int_{\tau_0}^{\tau} \left(\sum_{n \in N} B_n(t, w_t^*, \mathbf{0}) \frac{d \log p_n}{dt} + \frac{1}{\sigma(t, w_t^*, \mathbf{0}) - 1} \frac{d \log B^P(t, w_t^*, \mathbf{0})}{dt} \right) dt, \quad (12)$$

where w_t^* solves the equation

$$u(t, w_t^*, \mathbf{0}) = u(\tau, w, \mathbf{0}). \quad (13)$$

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

Proposition 4 generalizes the results in Baqaee et al. (2023) to a dynamic stochastic environment. It shows that the money metric can be recovered by deflating nominal wealth at τ using cumulative inflation between τ and τ_0 with an adjustment for changes in expenditures out of wealth. The cumulative inflation rate between τ and τ_0 uses compensated household budget shares. Similarly changes in expenditures out of wealth must be compensated. That is, budget shares and changes in savings rates must reflect only substitution effects, due to changes in prices and probabilities, and not income effects.

Proposition 4 is a fixed-point problem in terms of 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 changes in expenditures relative to wealth B^P as a function of time and wealth. Given these observables, and estimates of the EIS, we can solve (12) for the generalized money metric.

Solution Method. To apply Proposition 4, we begin by guessing a solution $u^0(\tau, w, \mathbf{0})$, for example, using a static price index. We then use this initial guess on the right-hand side of (12) to get a new guess. We then iterate on this until convergence. This procedure will always converge since the fixed point to (12) is, locally, a contraction mapping. Details are provided in the appendix.

Boundaries. Proposition 4 can only be applied reliably inside a suitable boundary. This is because the budget shares $\boldsymbol{B}(\tau, w, \mathbf{0})$ and consumption-wealth ratios $B^P(\tau, w, \mathbf{0})$ are observed only for some subset of time periods, say $\tau \in [\underline{\tau}, \overline{\tau}]$ and some subset of wealth levels, say $w \in [\underline{w}_t, \overline{w}_t]$. This limits the range of values of w and τ 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 $u(\tau, w, \mathbf{0})$ is not in $[u(s, \underline{w}_s, \mathbf{0}), u(s, \overline{w}_s, \mathbf{0}]$ for some $s \in [\tau_0, \tau]$, then we cannot reocover $u(\tau, w, \mathbf{0})$ without extrapolation. This is because there are no households in cohort $s \in [\tau_0, \tau]$ that are on the same indifference curve as the rentier with wealth w at time τ. This means we cannot calculate the compensated budget shares and consumption-wealth ratio for these households, and hence cannot calculate their money metric utility value.

Proof Sketch. To better understand Proposition 4, we sketch the proof. (The formal proof is in Appendix C). First, we establish that the dual shadow prices, defined by Proposition 3, associated with $(\tau, w, \mathbf{0})$ can alternately be written to depend on τ and $V(\tau, w, \mathbf{0})$ instead of

w directly. That is, for each history s^t, we can write

$$\boldsymbol{q}^*(s^t|\tau, w, \mathbf{0}) = \boldsymbol{q}^*\left(s^t|\tau, V(\tau, w, \mathbf{0})\right). \tag{14}$$

This is intuitive since w and $V(\tau, w, \mathbf{0})$ are monotone. Therefore, we can think of q^* as a "Hicksian" or compensated shadow price because it depends on utility rather than wealth. One of the reasons we focus on households with only marketable wealth is that (14) need not hold for households with non-marketable wealth. For households with non-marketable assets, the shadow prices do not just depend on wealth and calendar time, they also depend on expected cashflows and borrowing constraints.

Next, using the Hicksian shadow prices, we show that the money metric, $u(\tau, w, \mathbf{0})$ can be expressed using the shadow expenditure function as

$$u(\tau, w, \mathbf{0}) = e(\boldsymbol{q}^*\left(\cdot|\tau_0, u(\tau, w, \mathbf{0})\right), \boldsymbol{\pi}(\cdot|\tau_0), u(\tau, w, \mathbf{0})).$$

That is, for households without non-marketable assets, the money metric value coincides with the shadow expenditures that a household would need to be given to reach the utility level $u(\tau, w, \mathbf{0})$ when facing Hicksian shadow prices. We can manipulate this expression to get

$$\log u(\tau, w, \mathbf{0}) = \log w - \log \frac{e(\boldsymbol{q}^*\left(\cdot|\tau, u(\tau, w, \mathbf{0})\right), \boldsymbol{\pi}(\cdot|\tau), u(\tau, w, \mathbf{0}))}{e(\boldsymbol{q}^*\left(\cdot|\tau_0, u(\tau, w, \mathbf{0})\right), \boldsymbol{\pi}(\cdot|\tau_0), u(\tau, w, \mathbf{0}))}.$$

That is, the money metric is nominal wealth, at date τ, deflated using an appropriate price index that holds utility fixed and tracks changes in Hicksian shadow prices and probabilities.[6]

Next, we re-express the last term in the equation above, using the fundamental theorem of calculus, as

$$\begin{aligned}\log u(\tau, w, \mathbf{0}) = \log w + \int_{\tau}^{\tau_0} \sum_{t=x}^{T} \sum_{s^t} \Bigg(& \frac{\partial \log e(\boldsymbol{q}^*\left(s^t|x, u(\tau, w, \mathbf{0})\right), \boldsymbol{\pi}(s^t|x), u(\tau, w, \mathbf{0}))}{\partial \log \boldsymbol{q}^*(s^t|x, u(\tau, w, \mathbf{0}))} \cdot \frac{d \log \boldsymbol{q}^*\left(s^t|x, u(\tau, w, \mathbf{0})\right)}{dx} \\ & + \frac{\partial \log e(\boldsymbol{q}^*\left(s^t|x, u(\tau, w, \mathbf{0})\right), \boldsymbol{\pi}(s^t|x), u(\tau, w, \mathbf{0}))}{\partial \log \boldsymbol{\pi}(s^t|x)} \cdot \frac{d \log \boldsymbol{\pi}(s^t|x)}{dx} \Bigg) dx. \end{aligned} \tag{15}$$

The integral, which is equal to the change in the ideal price index, consists of two sets of terms. The first set of integrands, on the top line, track how the expenditure function responds to changes in shadow prices in all possible times and states as calendar time, indexed by x, moves from the base year τ_0 to τ. In a static deterministic environment,

[6] In a static, deterministic environment, this price deflator collapses to an ideal price index, also known as a Konüs (1939) price index.

this collapses to how the expenditure function responds to changes in static prices. The second set of integrands, on the second line, track how the expenditure function responds to changes in probabilities in all possible future dates and states as calendar time, indexed by x, moves from the base year τ_0 to τ. These terms have no counterparts in the standard static deterministic framework.

These summands in the integral are very high-dimensional, potentially infinite-dimensional, sums over all possible dates and states. Equation (15) elucidates the enormous complexity of forward-looking measures of inflation as compared to the traditional static objects. The forward-looking measure depends on how all possible future shadow prices and probabilities change as time moves forward. This complexity is compounded by the fact that we must weigh changes in all of these unobservable shadow prices and probabilities by the elasticities of the shadow expenditure function with respect to shadow prices and probabilities respectively.

Fortunately, we can cut through much of this complexity as long as preferences are time separable. When preferences are time separable, the complicated integrand in (15) can be rewritten as

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

Intuitively, changes in shadow prices and probabilities change the shadow expenditure function to the extent that they move the compensated future bundle $F\left(\{\boldsymbol{q}(s^t)\}_{t>0}, \boldsymbol{\pi}, U\right)$ relative to the present bundle $P\left(\boldsymbol{q}(s^0), U\right)$. This compensated relative price, in turn, changes consumption relative to savings rates if the EIS is not equal to one. Hence, we can infer changes in the prices and probabilities relevant for the future by observing changes in saving behavior, as long as we know the EIS.

Plugging this equation back into (15) and manipulating leads to Proposition 4. The last step uses the insight from Baqaee et al. (2023) that, with the addition of (13), we can treat (15) as a fixed point problem. Proposition 4 implies that if we can identify households for whom non-marketable wealth is negligible, then we can recover $u(\tau, w, \boldsymbol{0})$ as a function of time τ and wealth w.

A crucial fact about Proposition 4 is that it defines a fixed point problem. This is because, to arrive at $u(\tau, w, \boldsymbol{0})$, we need to integrate compensated budget shares and compensated changes in savings rate. However, to perform the necessary compensation, we need to know $u(\tau, w, \boldsymbol{0})$. This fixed point problem disappears when we specialize preferences to be homothetic.

Special Cases. To build intuition, we consider a series of homothetic special cases of Proposition 4 below.

Corollary 1 (Homothetic Preferences)**.** *If preferences are homothetic, then the money metric satisfies the following equation*

$$\log u(\tau, w, \mathbf{0}) = \log w - \int_0^\tau \left(\sum_i B_n(t) \frac{d \log p_i}{dt} + \frac{1}{\sigma(t, \mathbf{0}) - 1} \frac{d \log B^P(t, \mathbf{0})}{dt} \right) dt.$$

When preferences are homothetic, the integrand in (12) simplifies. First, the share of spending on each good $B_n(t, w_t^*, \boldsymbol{y})$ is only a function of the time period, which we write as $B_n(t)$. This is a consequence of homotheticity and time separability. Time separability means that budget shares on present consumption do not respond to changes in future prices. Homotheticity implies that budget shares on present consumption do not depend on wealth. Since all households face the same within-period relative prices, this means that $B_n(t, w, \boldsymbol{y})$ is the same for all households at time t.

Second, the compensated change in the consumption share of wealth $d \log B^P(t, w_t^*, \mathbf{0})/dt$ simplifies to the uncompensated change in the consumption share of wealth $d \log B^P(t, \mathbf{0})/dt$. This is because the consumption share of wealth $B^P(t, w, \boldsymbol{y})$ is the same for all households for whom $\boldsymbol{y} = \mathbf{0}$. This is a consequence of homotheticity, whereby the rate at which households with only marketable wealth substitute between spending and saving is the same regardless of their level of marketable wealth. Notably, we still require that non-marketable wealth, $\boldsymbol{y}$, be zero.

Next, we consider how Corollary 1 simplifies further if we assume that the EIS is constant.

Corollary 2 (Homothetic CES Preferences)**.** *If preferences are homothetic and the EIS is constant, then the money metric satisfies the following equation*

$$\log u(\tau, w, \mathbf{0}) = \log w - \int_0^\tau \sum_i B_n(t) \frac{d \log p_i}{dt} dt + \frac{\log \left(B^P(\tau, \mathbf{0})/B^P(\tau, \mathbf{0})\right)}{\sigma - 1}.$$

Compared to Corollary 1, the final term is now a simple log difference in the consumption share of wealth, comparing households in the initial period, τ_0, to households in period τ.

Finally, we consider how Corollary 2 simplifies further if we assume that the present and the future are perfect substitutes.

Corollary 3 (Homothetic CES Preferences with infinite EIS)**.** *If, in addition, the intertemporal elasticity of substitution is infinite, then*

$$\log u(\tau, w, \mathbf{0}) = \log w - \int_0^\tau \sum_i B_n(t) \frac{d \log p_i}{dt} dt. \tag{16}$$

Equation (16) is familiar. It is just nominal wealth deflated using a chain-weighted (or Divisia) consumption price index. In other words, Corollary 3 provides assumptions under which the naive application of a static consumption price deflator is valid for converting wealth in τ into equivalent value in τ_0. The assumptions required are very strong, requiring that preferences be homothetic both within and across time periods (so that uncompensated budget shares can be used), and that the intertemporal elasticity of substitution be infinite (so that the household weakly always prefers to spend all of its wealth in the present). Under these incredibly assumptions, the naive practice of deflating wealth by a Divisia consumption price deflator is valid.

5.3 Non-Rentiers

A challenge for the applicability Proposition 4 is that, in practice, many households in the sample may have non-negligible non-marketable wealth (i.e. $\boldsymbol{y} \neq 0$). Fortunately, we can use non-homotheticity of preferences to extend $u(\tau, w, \mathbf{0})$ to households with non-marketable assets. To do so, we first make the following observation.

Proposition 5 (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:*

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

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

Importantly, Proposition 5 implies that the budget share of each good in the present $B_n(\tau, w, \boldsymbol{y})$ does not directly depend on marketable wealth w and the stream of payoffs from non-marketable asset $\boldsymbol{y}$.

The next proposition makes it possible to extend $u(\tau, w, \mathbf{0})$ to cover households with non-marketable assets. To do so, with some abuse of notation, index money-metrics by their base period. That is, let $u_{\tau_0}(\tau, w, \mathbf{0})$ denote the money metric value for the problem faced by an agent at calendar time τ, with marketable wealth w, non-marketable cashflows $\boldsymbol{y}$ in terms of the base period τ_0.

Proposition 6 (Money Metric is a Function of Budget Shares and Time)**.** *Suppose that the vector-valued function $\boldsymbol{b}(\boldsymbol{p}, V)$ is an injective function of V. Then, there exists a function m satisfying*

$$u_\tau(\tau, w, \boldsymbol{y}) = m(\boldsymbol{B}(\tau, w, \boldsymbol{y}), \tau),$$

for every τ, w, and $\boldsymbol{y}$.

The compensated budget shares $\boldsymbol{b}(\boldsymbol{p}, V)$ are an injective function of V if no two distinct values of V result in the same vector of 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 $m(\boldsymbol{B}, \tau)$ mapping vectors of budget shares $\boldsymbol{B}$ at date τ into the equivalent lump sum wealth at date τ (i.e. $u_\tau(\tau, w, \boldsymbol{y})$).

Hence, if we know the function m, and we observe budget shares $\boldsymbol{B}(\tau, w, \boldsymbol{y})$ at time τ, then we can deduce the money-metric utility $u_\tau(\tau, w, \boldsymbol{y})$ for a household facing the problem $(\tau, w, \boldsymbol{y})$. Given $u_\tau(\tau, w, \boldsymbol{y})$ we can then use Proposition 4 to convert this to money-metric utility for some other base date $u_{\tau_0}(\tau, w, \boldsymbol{y})$.

How do we learn the shape of the function m? We use the identity that for households with only marketable wealth $u_\tau(\tau, w, \boldsymbol{0}) = w$. Given this, we can learn the shape of m by solving the following optimization problem

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

where M is a set of functions that contains m. In words, we fit a flexible function that relates budget shares to wealth for households with negligible non-marketable wealth. We then use this fitted relationship to impute the $u_\tau(\tau, w, \boldsymbol{y})$ for households with non-marketable assets. Proposition 7 formalizes this idea.

Proposition 7 (Money Metric for Non-Rentiers)**.** *Let*

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

Let $m|_{\mathcal{B}(\tau)}$ be the function m restricted to the domain $\mathcal{B}(\tau)$. We have that

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

A special case of Proposition 6 and Proposition 7 is the case where the budget share of a specific good, usually food, is known to be strictly monotone in utility.

Corollary 4 (Engel's Law)**.** *Suppose that there exists a good $i \in N$ whose budget share, $b_i(\boldsymbol{p}(s^0|\tau), V(\tau, w, \boldsymbol{y}))$, is strictly monotone in V. Then*

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

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

Propositions 4, 6, and 7 can be combined to recover $u(\tau, w, \boldsymbol{y})$ for every $u(\tau, w, \boldsymbol{y})$ inside a suitable boundary.

Discussion. We use observed changes in consumption-wealth ratio to infer changes in households' beliefs about future prices and rates of return. This requires assuming that the preference relation of cohorts is not changing over time. Preferences can vary in the cross-section by observable characteristic (e.g. age). However, holding observable characteristics constant, consumption-savings choices change *only* due to changes in wealth, prices, rates of returns, and beliefs about the future — they do not change due to changes in preference parameters.

Throughout, our approach does not require that households' beliefs about the future be "objective" in any sense. 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. We do, however, require that all households in cohort τ face the same prices and beliefs.

6 Measuring Money Metric Utility

In this section, we apply our method to data from the US. We require data on households and on prices. Importantly, our methodology can recover the money metric for a subsample of households even if that subsample does not sample households at the same frequency as the population.

For our household data, we use the Panel Study of Income Dynamics (PSID) spanning the years 2005 to 2019. These are the years with detailed expenditure surveys. One advantage of this dataset, compared to other US household surveys like the Consumer Expenditure Survey (CEX) or the Survey of Consumer Finances (SCF), is that it has a comprehensive record of the data needed for our method. This includes expenditure categories, income and household financial assets, and demographic information, all within a single dataset.

For owner-occupied housing costs, we impute equivalent rental expenditures, in a theory-consistent way, by matching home owners in each period to renters with similar spending behavior.[7] For the price data, we create 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 A.

We begin by computing the money metric wealth for rentier households. We then focus on the rest of the sample.

6.1 Rentiers

To apply Proposition 4, we need to observe a sample of rentier households. To create such a subsample, we first impute 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 labor and transfer income. If the head of the household is unemployed and looking for a job, then we exclude this household from the sample of rentiers.

To capitalize labor and transfer income for each household, we predict 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 income is detailed in Appendix A. We say that a household is a rentier if net assets—–or marketable wealth—–constitute more than 90% of total wealth.

Figure 1 shows the empirical distribution of consumption-wealth ratios in the rentier subsample (pooled across all years). To limit mismeasurement, we further exclude households if their current consumption-wealth ratio is more than 20% on the basis that, for these households, we may be underestimating the value of non-marketable wealth. To further limit the role of outliers, we also exclude households whose net assets are in the top and bottom 2.5%.[8] The rest of Section 6.1 limits attention to this rentier subsample. We treat the remaining, non-rentier, households separately in the next section.

In an average year, approximately 5 percent of PSID households above 45 years old are rentiers by our definition. Figure 2 shows the age and wealth distributions for these

[7] In 2019, the PSID asked home owners to report the rental value of their property. We use the answers to this question, in 2019, to validate our imputation procedure. When we regress surveyed housing costs on our imputed measure of housing costs without an intercept term, we find a coefficient of 1.03 and an R^2 value is 0.85. See Figure 9 and Appendix A for details.

[8] In Appendix B, we show that our results are robust to varying these cut-off values.

Figure 1: Consumption-wealth ratio for prospective rentiers, before dropping outliers.

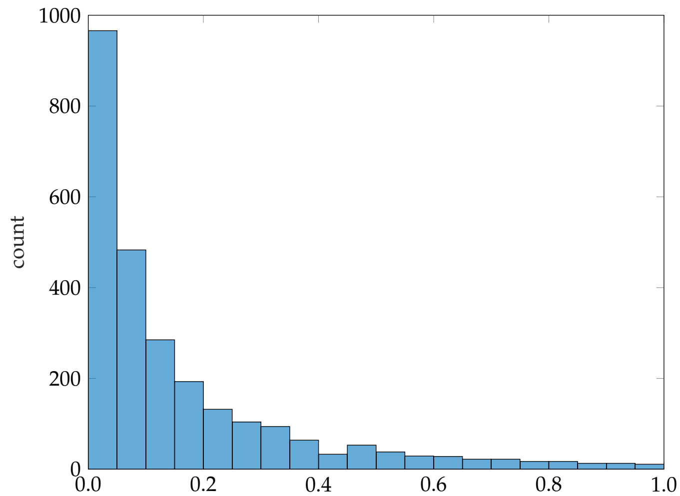

Notes: This figure pools all years.

households. The sample is concentrated around households in late middle age. The median wealth, in 2019, is around $700,000 dollars. Our methodology does not require the rentier distribution to be representative of the non-rentier distribution in terms of wealth or demographics. However, there are very few rentiers below 45 years old, making our estimates for ages below 45 quite noisy. For this reason, we only report results for households that are 45 years and older.

Figure 2: Empirical distributions for households with negligible non-marketable wealth

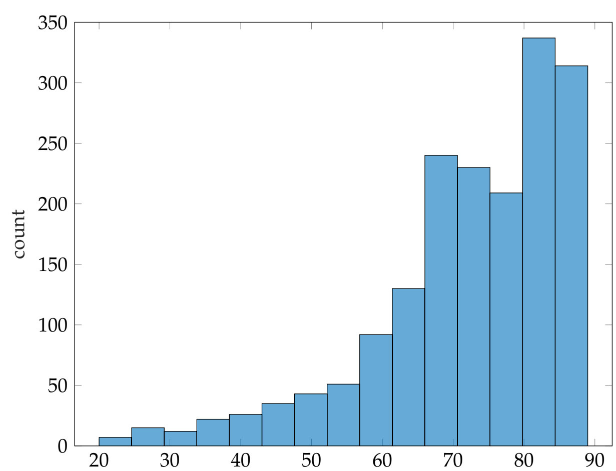

(a) Age

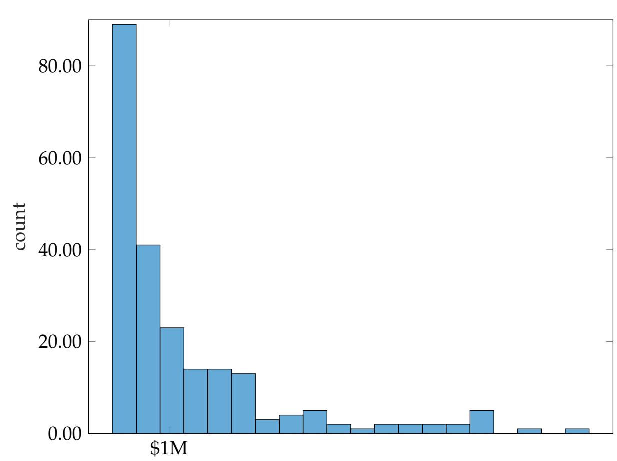

(b) Wealth

Notes: Age is pooled across all years. Wealth is only shown for the final year in the sample, 2019.

To compute money metric values for rentiers, we need to be able to evaluate consumption-wealth ratios and relative budget shares as a function of calendar date, age, and wealth. To do this, we fit smooth curves that relate consumption-wealth ratios and budget shares to age and wealth level in each period. That is, let h index individual households and τ index calendar dates. For each date τ, we use locally weighted scatterplot smoothing (LOWESS) to fit cross-sectional mapping from wealth and age to consumption-wealth ratios:

$$\log B^P_{h,\tau} = \hat{B}_\tau(\log \text{wealth}_{h,\tau}, \text{age}_{h,\tau}) + \epsilon_{h\tau}, \tag{17}$$

where we assume that $age_{h,\tau}$ is one-to-one with planning horizon of household h at date τ. Similarly, for each good $i \in N$, we use LOWESS to fit the cross-sectional mapping from wealth and age to budget shares on each good at each date. We choose the smoothing parameter of LOWESS by cross-validation.

Figure 3: Consumption-wealth ratios

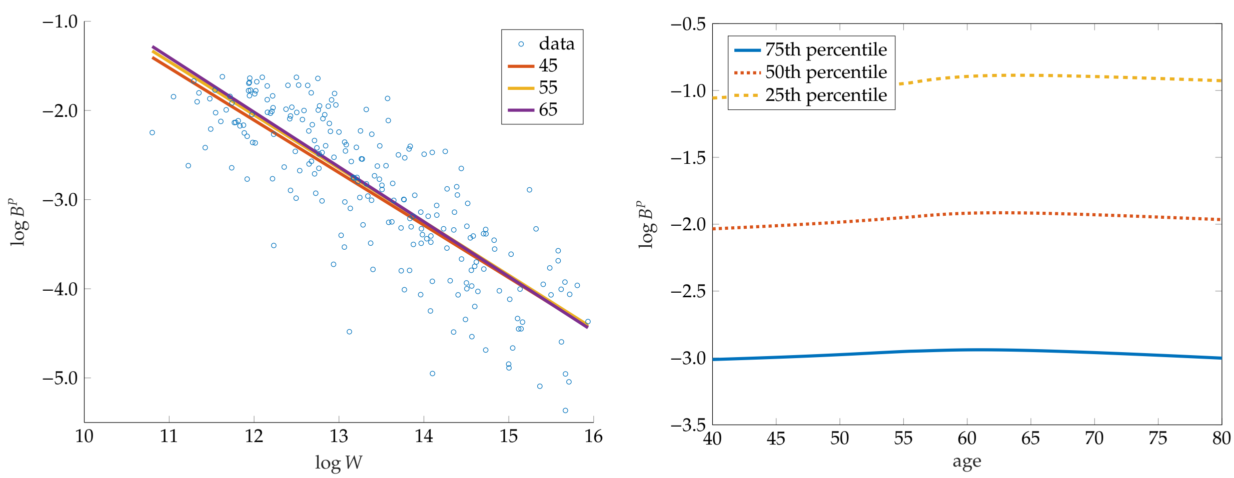

(a) Log consumption-wealth against log wealth for different ages.

(b) Log consumption-wealth against age for different wealth percentiles.

Notes: This figure uses data from the final year in the sample, 2019.

Figures 3 shows different cross-sections of the estimated relationship in (17). Figure 3a plots the fitted curves relating the consumption-wealth ratio to wealth, in 2019, for a selection of different ages along with a scatterplot of the raw data. We see a strong downward sloping loglinear relationship where wealthier households save a greater portion of their wealth.

Figure 3b shows the age profile of the log consumption-wealth ratio as a function of age for different percentiles of the wealth distribution. The consumption-wealth ratio is much less sensitive to age than to wealth. For the poorest households, the consumption-wealth ratio rises slightly as households age, but for richer households, it is mostly flat.

We use our estimated cross-sectional curves in Proposition 4 to recover money metric utility as a function of calendar date, age, and wealth.[9] When we apply Proposition 4, we only compare households of a certain age to households of that same age at different dates. We never compare households with different ages to one another since we allow households to have different preferences at different ages. In particular, if lifespan is finite, then households face different decision problems at different points in their life. For this reason, the age profile of the consumption-wealth ratio, depicted in Figure 2b, is not directly used for our results.

We use $t_0 = 2005$ as the base year, so that, for each age, the money metric value maps wealth in each year t into equivalent wealth in 2005. To avoid out-of-sample extrapolation, we limit the money metric values to be within the convex hull of the wealth and age distributions for each year. For our benchmark results, we set the EIS, σ, equal to 0.1 following the estimates of Best et al. (2020).

For illustration, Figure 4a plots the money metric as a function of 2019 wealth for 55 year old households. For comparison, we also plot a naive calculation that deflates nominal wealth in 2019 by CPI inflation. We find that the naive measure overstates the money metric for all households, especially for richer households. That is, all households experienced a higher effective inflation rate over this period than what is implied by CPI inflation, particularly richer households.

To understand why, we decompose the welfare-relevant inflation rate, in equation (12), into a present and a future part. Specifically, for a household with wealth w in 2019, the change in the ideal cost-of-living index between 2005 and 2019 is:

$$\log \frac{w}{u(2019, w)} = \underbrace{\int_{2005}^{2019} \sum_{n \in N} B_n(t, w_t^*, \mathbf{0}) \frac{d \log p_n}{dt} dt}_{\text{static inflation}} + \underbrace{\frac{1}{\sigma - 1} \log\left(\frac{B^P(t, w_\tau, \mathbf{0})}{B^P(t, w_{2005}^*, \mathbf{0})} \right)}_{\text{future relative to static inflation}},$$

where w_t^* ensures that we are using compensated consumption-wealth ratios and budget shares. The first summand is a static measure of inflation (as in, e.g. Baqaee et al., 2023). The first summand tracks how the compensated price of the present bundle is changing over time. The second summand is related to expected future inflation relative to present inflation. If the second term is positive, then the future inflation (the rate at which the price of the future bundle changes) is higher than static inflation (the rate at which the price of the present bundle changes).

[9] To recover the money metric, we need to solve the integral equation in Proposition 4. To do so, we use the "recursive" methodology described in Baqaee et al. (2023).

Figure 4: Money metric and cost-of-living inflation for 55 year olds

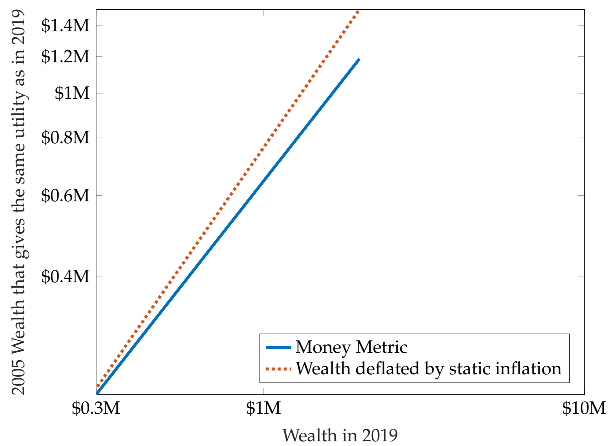

(a) The money metric and wealth deflated by official CPI between 2005 to 2019 (dollar, log scale). This figure converts wealth in 2005 into equivalent wealth in 2019 and vice versa.

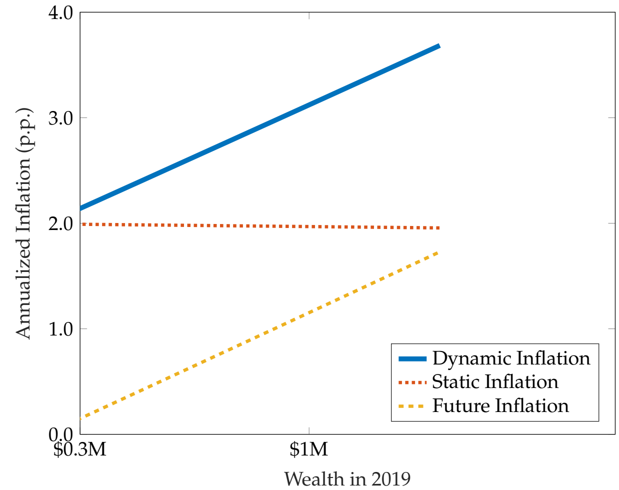

(b) Decomposition of dynamic inflation from 2005 to 2019: the first term (static inflation) and the second term (future inflation) in the Proposition 4 integral, and their sum (dynamic inflation).

The decomposition is shown in Figure 4b. We discuss the static and future parts in turn. The static inflation term differs from the aggregate CPI because it weighs prices changes in static prices using compensated budget shares rather than aggregate budget shares. Nevertheless, it is 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.

The downward slope in the static inflation term is consistent with previous studies, like Blundell et al. (2003), Jaravel and Lashkari (2022), and Baqaee et al. (2023), 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 extremely mild, and to the extent that there is a slope, it goes in the opposite direction when compared to the overall forward-looking cost-of-living index.[10]

The fact that the dynamic inflation rate in Figure 4b is positive and has a positive slope with wealth is driven by changes in the "future" inflation component. All households anticipate higher increases in the price of consumption in the future. This means that the

[10]There 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 expenditures as is done in static studies of the cost-of-living. Finally, our sample period of fourteen years is reasonably short compared to previous studies, which compute changes over 50 years or longer.

overall cost-of-living index is higher than the static inflation rate. This gap is larger for richer households. This finding comes about because compensated consumption-wealth ratios in 2019 are lower than they were in 2005. This indicates that households are more pesimistic about the future relative to the present in 2019 than they were in 2005.

Figure 5 illustrates this by showing the compensated and uncompensated consumption-wealth ratio over time for the median 55 year old household in 2019. The uncompensated consumption-wealth ratio has been fairly constant over time. On the other hand, the compensated consumption-wealth ratio, which eliminates income effects, ratio has fallen over time. The compensated consumption-wealth ratio also has more extreme fluctuations than the uncompensated one. This is because the EIS is less than one. To see why, consider a decrease in the consumption-wealth ratio, which all else equal, indicates an increase in future inflation. This increase in future inflation, holding nominal wealth constant, implies a decrease in utility. To compensate for this, and keep the household on the same indifference curve, we must increase nominal wealth. Increasing wealth causes the consumption-wealth ratio to fall due to the fact that the consumption-wealth Engel curve is decreasing in wealth. Therefore, the compensated consumption-wealth ratio falls by more than the uncompensated one. The same logic works, but in reverse, for increases in the consumption-wealth ratio.

Figure 5: Compensated and uncompensated consumption-wealth ratio for median rentier in 2019.

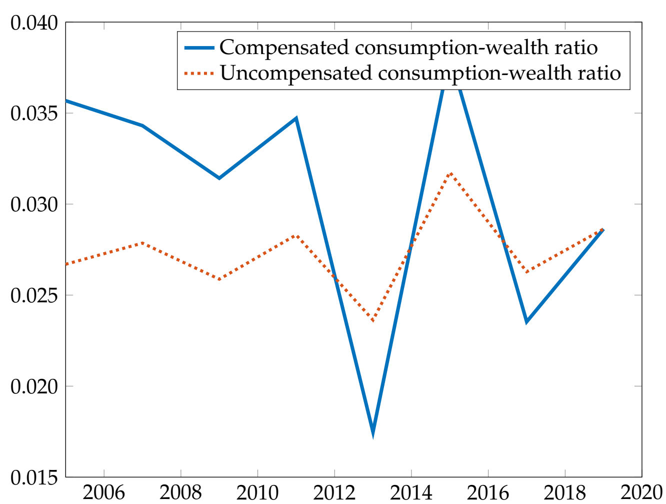

Notes: For a household with median wealth in 2019 and age 55.

Our methodology does not identify which future prices or beliefs are responsible for the patterns in Figure 4. However, the differences in the dynamic measure of inflation for rich and poor households need not be caused by differential exposures to future goods prices alone. These differences can also be driven by differential expected returns of assets used by different households to save. For example, if richer households save primarily through equities and poorer households save primarily through bonds, then differential trends in the rates of return in the two asset classes, for instance a declining equity risk premium, can generate Figure 4.

Figure 6 shows the dynamic cost-of-living index as a function of wealth for different ages. We see that the patterns are not the same for all ages. For all ages, the cost-of-living inflation is rising in wealth. Once again, our methodology does not identify which future prices are responsible for these trends. However, as before, these differences could be driven by the fact that rich and poor young households are differentially exposed to expected asset returns compared to rich and poor old households.

Figure 6: Cost-of-living inflation by age

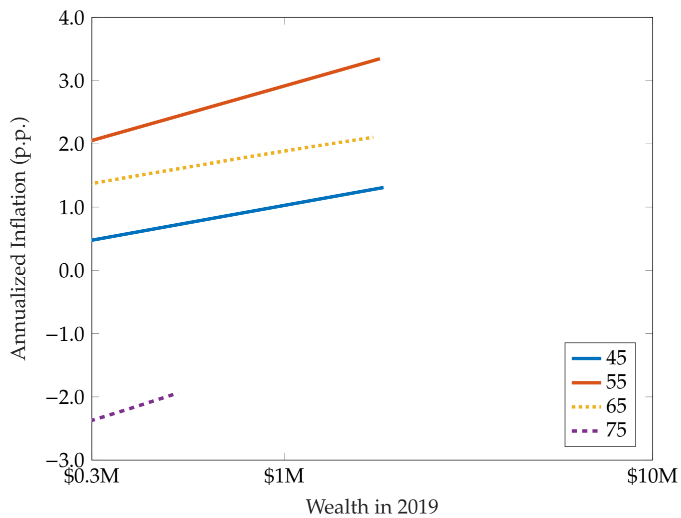

6.2 Non-Rentiers

We now turn our attention to the remaining households — the non-rentiers. Proposition 4 does not apply to these households. To recover the money metric for these households, we rely on Propositions 6 and 7 instead. To start with, for rentiers, we regress measured wealth at date τ on a polynomial of budget shares and age at date τ:

$$w_{h,\tau} = \alpha_{0,\tau} + \sum_{k=1}^{2} \sum_{i=1}^{N} \alpha_{i,\tau,k} B_{i,h,\tau}^{k} + \alpha_{0,\tau}^{ages} age_{h,\tau} + \alpha_{1,\tau}^{ages} age_{h,\tau}^{2} + \text{error}_{h,\tau}, \tag{18}$$

where α's are regression coefficients. We use this estimated relationship to impute equivalent wealth for the non-rentiers conditional on their budget shares and age.

To check the performance of our estimates of wealth, we compare median wealth for rentiers against the predicted median wealth following Proposition 7 (for illustration, we consider 55 year olds). If preferences are homothetic within period, then budget shares cannot be used to predict wealth, and hence we cannot recover measures of wealth from (18). However, we see a fairly good fit in Figure 7a. The median wealth imputed from (18) is slightly less than observed wealth for most year, but the difference is small and the trend in the two figures is very similar.

Instead of using (18) to estimate nominal money metric wealth for non-rentiers, we could instead rely on a discounted present value calculation. If financial markets are complete, we can also nominal money metric wealth by forecasting future expected labor income and transfer payments, discounting those payments back to the present using market interest rates, and adding this to net financial assets (e.g. as in Catherine et al., 2022).[11] We call this *capitalized* wealth and compare it to our imputed wealth from (18) for non-rentier households. If markets are incomplete, then even absent any mismeasurement or misspecification error, these two measures do not coincide.

Figure 7b compares median nominal money metric wealth against median capitalized wealth for the non-rentier subsample. Unlike in Figure 7a, we find large differences. The median equivalent wealth we estimate is both more volatile than capitalized wealth and lower. This makes sense since our imputed wealth is a certainty-equivalent and reflects the uninsurable risk borne by households. That is, if households are exposed to uninsurable idiosyncratic risk, then they should be willing to take a lump sum payment that is worth less than the discounted present value of their labor and transfer income.

Figure 8 plots the growth rate for the median wealth as a function of age between 2005 – 2019. The growth rate for all households has been positive, but by different amounts depending on age. The annual real growth rate, in money metric terms, has been highest for 75 year olds at almost 6% per year. The growth rate for 55 year olds, on the other hand, has only been around 0.5% per year.

For completeness, Figure 8 also shows the median growth in wealth for the whole sample and the rentier subsample deflated by CPI. Since the deflator is the same, this shows that nominal wealth for the rentier sample has been growing at a similar rate than nominal money metric wealth for the non-rentier sample.

The last bar graph shows growth in nominal median income deflated by CPI for our sample of households. This is a static flow measure, which simply deflates income flows

[11]See Appendix A for more details.

Figure 7: Median equivalent and capitalized wealth

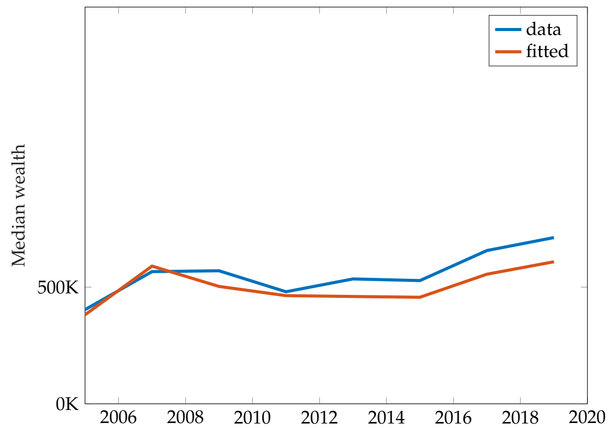

(a) Rentier subsample

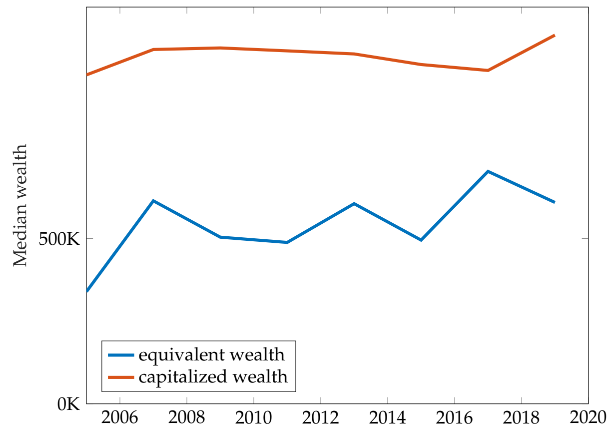

(b) Non-rentier subsample

Figure 8: Money metric growth for the median household by age

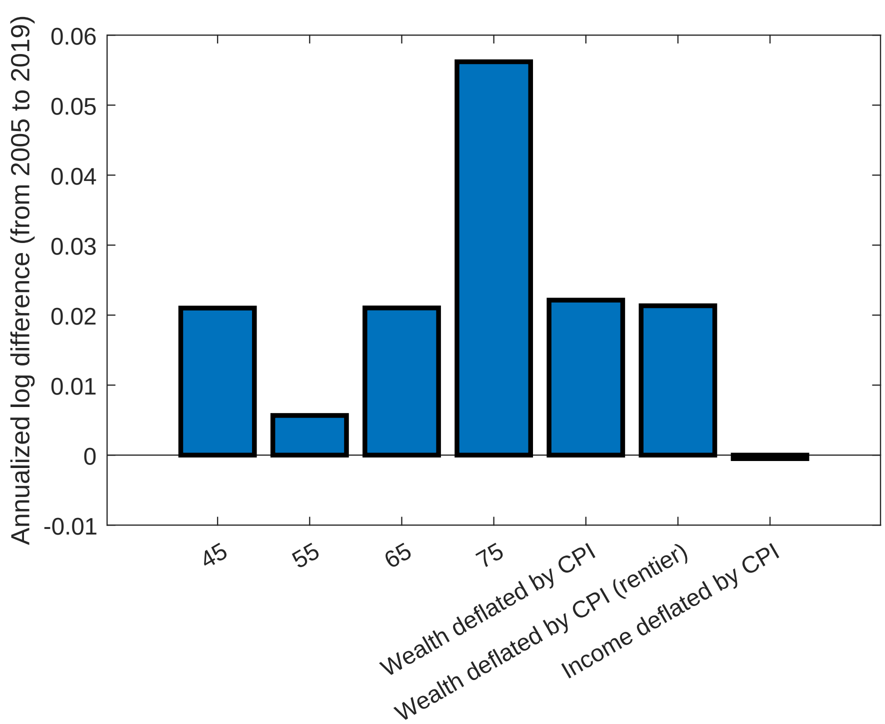

by static prices. This measure has stagnated between 2005 and 2019. For comparison, the median growth in real income deflated by CPI reported by the US Census Bureau over this period has been 0.8% per year. Of course, the number from Census Bureau is not directly comparable to ours since we work with a more restricted sample (excluding households younger than 40 for example).

7 Welfare Treatment Effects

In this section, we show how our methodology can be used to study the welfare effect of different shocks. Many, if not most shocks, affect households along many different dimensions. For example, job training programs, educational investments, tax policy, monetary policy all plausibly have dynamic effects on many different relevant variables for households. For example, Del Canto et al. (2023) show that monetary policy shocks affect households through many different channels: goods price inflation, labor market outcomes, changes in equity prices, house prices, bond prices, and so on.

To understand the welfare effect of a complex shocks like the ones described above, one option is to estimate the dynamic effects of the shock on 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 complete-markets (i.e. ruling out non-negligible uninsurable risk).

Our methodology provides an alternative approach whereby we measure welfare and equivalent-wealth without directly estimating the dynamic impact of the shock on all relevant variables or spending plans in the future. In this section, we illustrate this idea by considering the welfare effects of a job loss in the PSID.

For this illustration, we regress money metric utility 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 in that period and reports that she is searching for a new job. To control for confounds and selection, we include year fixed effects, demographic controls, and we control for the lagged value of money metric utility.

The results are reported in Table 1, which shows that job loss is associated with a roughly 7% reduction in money metric utility . This effect is large, but much smaller than the effect on income (which mechanically falls by 100% in that period). We can compare this to the dynamic consequences of job loss in Davis and Von Wachter (2011). They

Table 1: Percent change in equivalent wealth due to job loss

	log u					
	(1)	(2)	(3)	(4)	(5)	(6)
Job Loss	-0.073***	-0.045**	-0.079***	-0.070***	-0.072***	-0.089***
	(0.023)	(0.020)	(0.020)	(0.020)	(0.020)	(0.021)
Laged log u	No	Yes	Yes	Yes	Yes	No
Married	No	Yes	Yes	Yes	Yes	Yes
Age	No	No	Yes	Yes	Yes	Yes
Education	No	No	No	Yes	Yes	Yes
Industry	No	No	No	No	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	23,240	23,240	23,240	23,240	23,240	23,240

Standard errors in parentheses

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

estimate the present-value of earnings losses after mass-layoff events to be around 12% of counterfactual earnings (using a 5% discount rate). Of course, our numbers are not directly comparable to them since our sample years, sample selection criteria, and definition of layoffs are different. Furthermore, we consider the reduction in real money metric value losses, which are not the same as counterfactual earnings unless financial assets are closer to zero. Nevertheless, it is interesting that our methodology results in numbers that are, in terms of order of magnitude, not completely dissimilar to a present-value calculation.

8 Conclusion

We provide a methodology for measuring welfare and the cost-of-living for households that accounts for dynamics, uncertainty, market incompleteness, borrowing constraints, and non-homotheticities. Our methodology requires repeated household consumptin, income, and wealth surveys, as well as prices. The key assumptions we make are that preferences are common across all households, preferences are intertemporally separable, and all households in a given period face the same prices. To calculate money-metrics and cost-of-living, we require that some subset of households be rentiers, with negligible labor income, and knowledge of intertemporal elasticity of substitution. Our approach provides a way to measure dynamic welfare with relatively few assumptions, making it useful for studying the welfare effect of shocks and policies that have dynamic stochastic effects on many variables that affect households.

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?

Aoki, S. and M. Kitahara (2010). Measuring a dynamic price index using consumption data. *Journal of Money, Credit and Banking 42*(5), 959–964.

Atkin, D., B. Faber, T. Fally, and M. Gonzalez-Navarro (2020, March). Measuring Welfare and Inequality with Incomplete Price Information. NBER Working Papers 26890, National Bureau of Economic Research, Inc.

Baqaee, D. R. and A. Burstein (2023). Welfare and output with income effects and taste shocks. *The Quarterly Journal of Economics 138*(2), 769–834.

Baqaee, D. R., A. Burstein, and Y. Koike-Mori (2023). Measuring welfare by matching households across time.

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.

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 (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).

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.

Farhi, E., A. Olivi, and I. Werning (2022). Price theory for incomplete markets. Technical report, National Bureau of Economic Research.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *The American Economic Review*, 157–177.

Hamilton, B. W. (2001, June). Using engel's law to estimate cpi bias. *American Economic Review 91*(3), 619–630.

Hulten, C. R. (1979). On the" importance" of productivity change. *The American economic review 69*(1), 126–136.

IMF (2004). *Consumer Price Index Manual: Theory and Practice*. USA: International Monetary Fund.

Jaravel, X. and D. Lashkari (2022). Nonparametric measurement of long-run growth in consumer welfare. Discussion Paper 1859, Center for Economic Performance.

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.

McKenzie, L. (1957). Demand theory without a utility index. *The Review of Economic Studies 24*(3), 185–189.

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.

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.

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.

Samuelson, P. A. (1974). Complementarity: An essay on the 40th anniversary of the hicks-allen revolution in demand theory. *Journal of Economic literature 12*(4), 1255–1289.

Samuelson, P. A. and S. Swamy (1974). Invariant economic index numbers and canonical duality: survey and synthesis. *The American Economic Review 64*(4), 566–593.

Solow, R. M. (1957). Technical change and the aggregate production function. *The review of Economics and Statistics*, 312–320.

A 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.[12]

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

[12]Cooper 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.

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, at each date, 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). 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 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:

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

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. Figure 9 shows a binned scatter plot of our imputed values for housing against housing costs (including utilities) from the survey. A regression of the survey values on our imputed values (without an intercept) has a coefficient of 1.03 with an R^2 value of 0.85. This suggests that our imputation performs well.

Figure 9: Imputed housing costs against surveyed housing costs

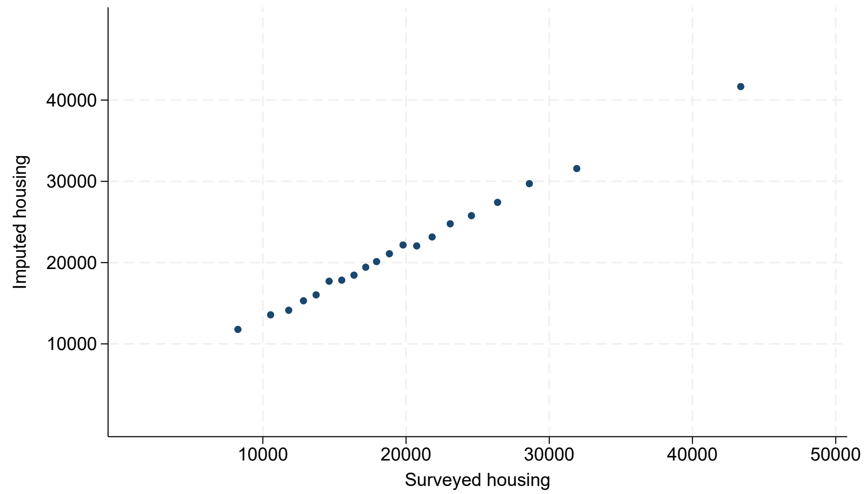

Notes: Surveyed housing is the sum of the answer to the question in the text (above) plus utilities. We annualize spending on housing. This graph uses only data for homeowners from 2019. The top and bottom 1% values for the surveyed housing costs are dropped.

Budget shares:

We align the seven categories of the PSID (food, housing, transportation, education, health, clothing, and recreation) with the CPI.[13] 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.

B Additional figures

To be added.

[13]The corresponding codes for CPI are CPIFABSL, CPIHOSSL, CPITRNSL, CPIEDSL, CPIMEDSL, CPIAPPSL, and CPIRECSL, respectively. Education includes child care. Recreation includes Trips & vacations and Recreation & entertainment in PSID.

C Proofs

Proof of Example 1. For any s^t ($t > 0$) and any $i \neq j$, we know that

$$\frac{\partial\left(\frac{\partial \mathcal{U}(\{c,\pi\})}{\partial c_i(s^0)} / \frac{\partial \mathcal{U}(\{c,\pi\})}{\partial c_j(s^0)}\right)}{\partial c_l(s^t)} = \frac{\partial\left(\frac{\tilde{u}(c_1(s^0),\ldots,c_N(s^0))}{\partial c_i(s^t)} / \frac{\partial\tilde{u}(c_1(s^0),\ldots,c_N(s^0))}{\partial c_j(s^t)}\right)}{\partial c_l(s^t)}, \\ = 0$$

and for any s^t and s^k ($t, k > 0$)

$$\frac{\partial\left(\frac{\partial \mathcal{U}(\{c,\pi\})}{\partial c_i(s^t)} / \frac{\partial \mathcal{U}(\{c,\pi\})}{\partial c_j(s^k)}\right)}{\partial c_l(s^0)} = \frac{\partial\left\{\left(\beta^t \pi(s^t)\frac{\tilde{u}(c_1(s^t),\ldots,c_N(s^t))}{\partial c_i(s^t)}\right) / \left(\beta^k \pi(s^k)\frac{\partial\tilde{u}(c_1(s^k),\ldots,c_N(s^k))}{\partial c_j(s^k)}\right)\right\}}{\partial c_l(s^0)}, \\ = 0$$

Thus the Leontief-Sono condition for separability (Blackorby et al. (1978), p.53) is satisfied. ■

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

Proof of Proposition 3. The existence of 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). Furthermore, since the solution is a convex optimization problem, the Karush-Kuhn-Tucker conditions must be satisfied. The Lagrangian for households is:

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

$$
-\sum_{s^t=s^0}^{s^T} \lambda(s^t|\tau) \sum_{k\in K} a_k(s^t|\tau) + \sum_{s^t} \lambda(s^t|\tau) \sum_{k\in K} R_k(s^t|\tau) a_k(s^{t-1}|\tau)
$$

$$
-\sum_{s^t} \mu(s^t|\tau) \sum_{k} a_k(s^t|\tau) + \sum_{s^t} \mu(s^t|\tau) X(s^t|\tau)
$$

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

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

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

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

From standard duality, we know that we can also write

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

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

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

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

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

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

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

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

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

■

Proof of Proposition 4. For the proof, we define the following function:

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

We proceed in steps, using a series of lemmas.

Lemma 1. *If preferences are indirectly time separable, then the following holds*

$$b^P(\boldsymbol{q}, \boldsymbol{\pi}, U) \equiv \sum_{n \in N} \frac{c_n(s^0)q_n(s^0)}{e(\boldsymbol{q}, \boldsymbol{\pi}, U)} = \frac{\partial \log e(\boldsymbol{q}, \boldsymbol{\pi}, U)}{\partial \log P},$$
$$b^F(\boldsymbol{q}, \boldsymbol{\pi}, U) \equiv 1 - b^P(\boldsymbol{q}, \boldsymbol{\pi}, U) = \frac{\partial \log e(\boldsymbol{q}, \boldsymbol{\pi}, U)}{\partial \log F},$$

and

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

Proof. By the envelope theorem,

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

and for $t > 0$

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

Then, we know that

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

and

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

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

Next, we show that

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

To do this, use the following equality,

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

Rearranging yields

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

■

Lemma 2. *When preferences are time separable, the intertemporal elasticity of substitution is given by*

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

Proof. We start with

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

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

$$= \frac{1}{b^P(\boldsymbol{q}, \boldsymbol{\pi}, U)} \left[\frac{\partial^2 \log e(\boldsymbol{q}, \boldsymbol{\pi}, U)}{(\partial \log P)^2} b_n(\boldsymbol{q}, U) \sum_{k \in N} b_k(\boldsymbol{q}, U) \right],$$

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

Summing over all $n \in N$ yields

$$\sum_n \frac{\partial \log b^P(\boldsymbol{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)} = \frac{1}{b^P(\boldsymbol{q}, \boldsymbol{\pi}, U)} \frac{\partial^2 \log e(\boldsymbol{q}, \boldsymbol{\pi}, U)}{(\partial \log P)^2} \sum_n b_n(\boldsymbol{q}, U),$$

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

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

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

By Definition 1,

$$1 - \sigma(\boldsymbol{q}, \boldsymbol{\pi}, U) = \sum_n \frac{\partial \log b^P(\boldsymbol{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)} - \sum_n \frac{\partial \log b^F(\boldsymbol{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)},$$

$$= \frac{1}{b^P(\boldsymbol{q}, \boldsymbol{\pi}, U)} \frac{\partial^2 \log e(\boldsymbol{q}, \boldsymbol{\pi}, U)}{(\partial \log P)^2} + \frac{b^P(\boldsymbol{q}, \boldsymbol{\pi}, U)}{b^F(\boldsymbol{q}, \boldsymbol{\pi}, U)} \frac{\partial \log b^P(\boldsymbol{q}, \boldsymbol{\pi}, U)}{\partial \log q_n(s^0)},$$

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

■

Lemma 3. *When preferences are indirectly time separable, the following equation holds:*

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

Proof. From Lemma 1, we know that

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

$$+b^F(\boldsymbol{q},\boldsymbol{\pi},U)\sum_{s^t|t>0}\left(\sum_{n\in N}\frac{\partial \log F}{\partial \log q_n(s^t)}d\log q_n(s^t)+\frac{\partial \log F}{\partial \log \boldsymbol{\pi}(s^t)}d\log \boldsymbol{\pi}(s^t)\right).$$

Next, from homogeneity of degree one, we know that

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

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

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

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

$$\begin{aligned}
b^P d\log b^P(\boldsymbol{q},\boldsymbol{\pi},U) &= \frac{\partial^2 \log e(\boldsymbol{q},\boldsymbol{\pi},U)}{(\partial \log P)^2}\sum_{n\in N}\frac{\partial \log P}{\partial \log q_n(s^0)}d\log q_n(s^0)\\
&+\frac{\partial^2 \log e}{\partial \log F\partial \log P}\sum_{s^t|t>0}\left(\sum_{n\in N}\frac{\partial \log F}{\partial \log q_n(s^t)}d\log q_n(s^t)+\frac{\partial \log F}{\partial \log \pi(s^t)}d\log \pi(s^t)\right)\\
&=\frac{\partial^2 \log e(\boldsymbol{q},\boldsymbol{\pi},U)}{(\partial \log P)^2}\left[\begin{array}{c}\sum_{n\in N} b_n(\boldsymbol{q},U)d\log q_n(s^0)-\\ \sum_{s^t|t>0}\left(\sum_{n\in N}\frac{\partial \log F}{\partial \log q_n(s^t)}d\log q_n(s^t)+\frac{\partial \log F}{\partial \log \pi(s^t)}d\log \pi(s^t)\right)\end{array}\right].
\end{aligned} \tag{19}$$

From Lemma 1 and Lemma 2, we can rewrite this as

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

Rearranging this gives

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

Plug this back into (19) to get the desired result. ∎

Lemma 4. *The shadow prices $\boldsymbol{q}^*(\tau,w,\mathbf{0})$ can be written as a function of τ and $V(\tau,w,\mathbf{0})$. That is, we can write*

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

Furthermore,

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

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 $\boldsymbol{q}^*(\tau, w, \mathbf{0})$ to get $\boldsymbol{q}^* (\tau, V(\tau, w, \mathbf{0})) = q^*(\tau, V^{-1}(V(\tau, w, \mathbf{0})), \mathbf{0})$.

For the second part, we know from Proposition 3, that

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

Hence

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

∎

Lemma 5. *The following holds*

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

Proof. From the proof of Proposition 3, we know that

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

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

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

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

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

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

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

From Lemma 4, we know

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

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

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

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

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

Using the shadow expenditure function, we can write

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

where the last line uses Lemma 5. Logging both sides gives

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

where the second equality uses the fundamental theorem of calculus for line integrals. Using Lemma 3, we can rewrite the last line as

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

where for the last step, we replaced compensated budget share with uncompensated budget share.

■

Proof of Proposition 5. Need to show that

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

By Proposition 3, we know that

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

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

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

Hence, we have that

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

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

Proof of Proposition 6. From Proposition 5, we know that

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

By definition of $u_\tau(\tau, w, \boldsymbol{y})$, it follows that

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

Since b is an injective function, we can write

$$V(\tau, u_\tau(\tau, w, \boldsymbol{y}), \boldsymbol{0})) = b_n^{-1}\left(\boldsymbol{p}(s^0|\tau), \boldsymbol{B}(\tau, w, \boldsymbol{y})\right).$$

Since V is monotone in wealth, we can write

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

■

Proof of Corollary 4. Proposition 5 shows that

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

Hence, if b_i is monotone in V, then

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

if, and only if,

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

■

D Relating Proposition 1 to Feenstra (1994)

Proposition 1 is also a consequence of the Feenstra (1994) approach to imputing the value of missing prices. Feenstra (1994) introduced this approach to adjust CES price indices for the value of new goods. The Feenstra (1994) approach applies because with complete markets, the consumers' problem is equivalent to a static problem where consumers make all their consumption choices at date 0. In this case, the preferences in (4) are a CES aggregator over dates. Hence, demand for consumption in the first period, relative to total wealth, follows CES demand:

$$\log \frac{b(0|\tau)}{b(0|\tau_0)} = (1-\sigma)\left[\log \frac{p(0|\tau)}{p(0|\tau_0)} - \frac{e(\boldsymbol{q}(\cdot|\tau_0, 1))}{e(\boldsymbol{q}(\cdot|\tau, 1))}\right]. \tag{20}$$

Rearrange this equation and combine it with (5) to arrive at (7). The proof we offer for Proposition 1 is different and much longer. Instead of inverting the demand curve to solve for the ideal price index, as in (20), we use demand for time 0 consumption to solve for relative prices between time 0 and the future. We then substitute this expression into (6) and integrate. The reason we do so is because this alternative, lengthier, proof generalizes when we relax the assumptions in this section. Outside of the CES special case, the demand curve for present consumption does not directly depend on the ideal price deflator as in (20). Hence, we cannot simply invert the demand curve to solve for the ideal price index.