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 consumer welfare, in money metric terms, taking into account expectations about the future. Our two key assumptions are that (1) the expenditure function is separable between the present and the future, and (2) there are some households that do not face idiosyncratic undiversifiable risk. Our sufficient statistics methodology allows for incomplete markets, lifecycle motives, non-rational expectations, non-exponential time discounting, and arbitrary functional forms. To apply our formulas, we require estimates of the elasticity of intertemporal substitution, goods and services' prices over time, and repeated cross-sectional information on households' income, balance sheets, and expenditures. We illustrate our method using the PSID from the United States and find large deviations from CPI-based measures. Our estimates can be used to study the welfare consequences of complex shocks that affect households along different margins and time horizons. For example, involuntary job loss is associated with a 20% reduction in money metric utility for households younger than 60 years old.

*We thank Andy Atkeson, Valentin Haddad, Oleg Itskhoki, Pablo Fajgelbaum, John Fernald, 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 macroeconomic model explicitly incorporates intertemporal decision-making, fully specifies preferences, beliefs, and parameters, and performs welfare counterfactuals taking into account both time and uncertainty (e.g. Lucas, 1987).

In contrast, conventional sufficient-statistic measures of welfare used in applied and empirical work, like real income and real consumption, eschew fully specifying every aspect of the problem. These measures can be calculated without specifying the demand system or providing a model of income and prices. They are constructed by deflating nominal expenditure or income by a price deflator, and can be micro-founded using the money-metric utility function.

Despite their generality, these sufficient-statistic measures apply only to static and deterministic decision problems and can easily produce misleading estimates of changes in intertemporal welfare.[1] For example, welfare may rise even though real income or real consumption falls.[2]

Therefore, when dynamic considerations are important, researchers tend to rely on fully-specified models to measure welfare. The main difficulty for the sufficient-statistic approach is that much of what one needs to observe in order to compute a measure of dynamic welfare is not easily observable. Static sufficient-statistic measures require knowing prices and spending on goods and services in different periods. By contrast, dynamic measures require knowing state-contingent spending plans, 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

[1] Real income and real consumption use consumer price deflators and money-metric utility uses a Konüs (1939) price deflator. These deflators are constructed for static decision problems. For example, the *Consumer Price Index* manual (IMF, 2004) or the survey paper by Jorgenson (2018) on "Production and Welfare Measurement", have little to say about how expectations about the future should influence the measurement of welfare and cost-of-living indices.

[2] Expenditures can fall even though welfare rises if, for example, there is a beneficial shock in the future and the elasticity of intertemporal substitution is greater than one or if there is a negative shock today and a positive shock in the future and the elasticity of intertemporal substitution is less than one.

tension in economics. For example, Samuelson (1961) concludes his paper on dynamic welfare measurement 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 computing *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. [...] the national income statistician is very far from having even an approximation to the data needed for these comparisons."

This paper develops a new approach to overcome some of these challenges. To do so, we first generalize the definition of money-metric utility so that it can be applied in settings with incomplete markets. In static settings with complete markets, welfare is measured in money metric terms by the expansion in the budget constraint (income) required to reach a given indifference curve.[3] When markets are incomplete, there is no single intertemporal budget constraint and the expansion required to reach a given indifference curve depends on which budget constraint is shifted. We define the money metric utility of each allocation to be the equivalent one-time lump sum wealth required to reach the indifference curve associated with that allocation.[4] If markets are complete and there is no risk, then this definition coincides with the usual one.

We then propose a sufficient-statistics methodology to estimate this forward-looking measure of welfare for a population with common preferences. In order to do this, we make two key assumptions. The first assumption is that preferences are separable between the present and the future and that the elasticity of intertemporal substitution is not equal to one.[5] The second assumption is that there is a subset of households that do not face idiosyncratic undiversifiable risk, which we call "rentier" households. Given these two assumptions, we can obtain forward-looking measures of welfare without further assumptions about utility functions (e.g. CES across goods), risk preferences (e.g. expected utility), time preferences (e.g. exponential discounting), beliefs (e.g. rational expectations), and financial frictions (e.g. complete markets), or first-order approximations.

[3]The money metric was introduced by McKenzie (1957), but popularized by Samuelson (1974). For a textbook presentation, see also Deaton and Muellbauer (1980).

[4]To be more precise, for an agent facing a dynamic stochastic problem in year τ, we ask: what is the one-time lump sum payment the agent must receive in some base year (with no other income sources thereafter) such that the agent is indifferent between their initial problem and this counterfactual problem. This amount of wealth is the money metric utility associated with the problem the agent faced in year τ, in base year dollars.

[5]Specifically, our notion of time separability is that the expenditure function implied by intertemporal preferences be separable between current and future prices.

We sketch the basic idea of our approach. First, for the rentier subsample, we back out the change in continuation value of the future relative to the present using changes in consumption-savings choices, conditional on estimates of the elasticity of intertemporal substitution (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, money metric utility will be higher than what is implied by deflating wealth by static inflation.[6]

Importantly, to calculate dynamic welfare we must use changes in the *compensated* consumption-wealth ratio of rentiers, which neutralizes wealth effects and responds only to substitution effects. To back-out the compensated consumption-wealth ratio without first specifying and estimating a dynamic model, we match rentiers in the cross-section over time, building on ideas from Baqaee et al. (2024).

Next, we recover money metric utility for non-rentiers by relying 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 non-rentiers by matching them with rentiers with similar static budget shares.[7]

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

To illustrate our method, we apply our results to the United States using the Panel Study of Income Dynamics (PSID) and price data from the Bureau of Labor Statistics. We 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

[6]Our sufficient statistic formulas take changes in prices of goods and services and in consumption-wealth ratios as given. For answers to counterfactual questions, one would have to provide counterfactual prices and counterfactual changes in consumption-wealth ratios, which requires a more fully-specified structural model.

[7]This requires that rentiers and non-rentiers are drawn from the same population (i.e. same preferences, beliefs, and prices), given observed characteristics (e.g. gender, age, and location).

is less than 10% of their total wealth. Taking estimates of the EIS from Best et al. (2020), we recover an ideal cost-of-living index using changes in consumption-wealth ratios for these households. 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. Furthermore, we find significant heterogeneity in the cost-of-living index across both the wealth distribution and by age group, although standard errors are large for some groups due to small sample size.

Our methodology is useful as an input for reduced-form empirical work studying the welfare effects of dynamic treatments. Many policies and shocks have complex effects that affect households in ex-ante uncertain ways along many dimensions and at different time horizons. We provide a way to study the welfare effects of such treatments without requiring that researchers enumerate, estimate, and aggregate all the possible ways the treatment affects the household and how the household's beliefs and contingent plans change in response to that treatment. To illustrate, we study the welfare consequences of job loss in our sample. To do so, we regress our estimates of money-metric utility on job loss. We find that involuntary job loss is associated with a roughly 20% reduction in money metric utility for households younger than 60. For households above 60, the losses are smaller and less statistically significant. Unlike alternative numbers from the literature, our results do not assume that all households discount the future using market rates, or that households have perfect foresight about the future.

Related Literature. Our paper is closely related to the vast literature on consumer price indices. The majority of papers in this literature abstract from time and risk. Nevertheless, we use several tools from this literature. Feenstra (1994) inverts CES demand curves to infer the value of new goods and other missing prices.[8] We extend this idea to infer the value of missing future prices relative to present prices using changes in consumption-wealth ratios.

Hamilton (2001) matches households that are on the same indifference curve over time via budget shares on food.[9] Atkin et al. (2024) provide a micro-foundation for this approach and use it to calculate welfare across the income distribution. We use a similar idea to Hamilton (2001), but we match rentiers to non-rentiers within the same period

[8]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), Aghion et al. (2019), Blaum et al. (2018), Argente and Lee (2021), and Redding and Weinstein (2020).

[9]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).

using static budget shares. This means that we do not have to correct for substitution effects, as in Atkin et al. (2024).

Our paper is also related to a literature that non-parametrically measures static money metric utility using repeated cross-sectional data, like Blundell et al. (2003), Jaravel and Lashkari (2024), and Baqaee et al. (2024). Specifically, we generalize Baqaee et al. (2024) to environments with intertemporal preferences and incomplete markets.

The literature on dynamic price indices is comparatively small, but can trace its origins to the inception of index number theory. 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 and 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 to calculate how consumer welfare responds to shocks to 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 asset sales. 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.

Our paper differs from these papers in some important ways. First, we do not directly estimate or model future asset or goods prices, beliefs, discount factors, and future holdings of assets or purchases of goods. Instead, we back out the welfare impact of changes in future prices and probabilities from changes in consumption-savings decisions. Second, our approach allows for aggregate and idiosyncratic risk without assuming a specific functional form for the stochastic discount factor. Third, we model non-homotheticities and do not rely on first-order approximations (for which non-homotheticites do not matter).[10]

Finally, our paper is also related to generalizations of price theory to incomplete markets. A notable example is Farhi et al. (2022), who decompose changes in demand in an incomplete market world into income and substitution effects. Our paper is related to this task since, to construct a money metric measure of welfare, we define and use a notion of compensated demand that treats income and substitution effects differently.

Roadmap. Section 2 illustrates a key idea in our method using a complete-markets example with homothetic and time-separable preferences: namely, that changes in consumption-wealth ratios can be used to infer dynamic welfare changes. Section 3 generalizes the environment to allow for incomplete markets, uninsurable risk, borrowing constraints, and non-homothetic preferences. Section 4 contains the main results of the paper showing how to extend the basic idea in Section 2 to this more complex environment, first for rentiers and then non-rentiers. Section 5 constructs a measure of dynamic welfare for households in the PSID. Section 6 uses these measures to study the dynamic welfare losses from job loss. We conclude in Section 7.

2 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 3.

[10] We document strong intertemporal non-homotheticities whereby wealthier households consume a smaller share of their wealth per period. This is consistent with Straub (2019).

Suppose that intertemporal preferences, in (5), take the form

$$\mathcal{U}(\boldsymbol{c}) = \sum_{t=0}^{T} \beta^t \left(c(t)\right)^{\frac{\sigma-1}{\sigma}}, \tag{1}$$

where T is the agent's horizon (potentially infinite), β is the discount factor, and $\sigma \neq 1$ is the elasticity of intertemporal substitution.

Consider an overlapping generations structure — at each calendar date $\tau \in \mathbb{R}$, there is a cohort of agents with planning horizon T. We index each cohort's decision problem using the start date τ and define intertemporal prices and returns t periods after τ by $\{p(t|\tau), R(t|\tau)\}_{t=0}^{T}$. Although each cohort makes discrete time decisions, $t \in \{0, \ldots, T\}$, new cohorts arrive every instant $\tau \in \mathbb{R}$. Prices and returns are continuous functions of calendar time τ.[11]

Each cohort τ faces a sequence of budget constraints, one for each period $t \in \{0, \ldots, T\}$:

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

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 calendar time $\tau + t$. 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.

Our objective is to compare the choice set facing different cohorts of households, always keeping the planning horizon fixed. For example, we compare the value of the choice set available to 60 year olds in 2005 to value of the choice set available to 60 year olds in 2019. This comparison has to take into account not only wealth and present prices but also future prices and returns. We do not compare the welfare of a single household at different points in their life because preferences and planning horizons change along the lifecycle.[12] Since utility is only defined up to monotone transformations, we quantify the value of choice sets using the money metric

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, define the τ_0 *money metric* function

[11] Specifically, for every $t \geq 0$, prices $p(t|\tau)$ and returns $R(t|\tau)$ in $t + \tau$ are absolutely continuous functions of calendar time τ.

[12] As explained by Fisher and Shell (1968), welfare comparisons, based on revealed preference theory, always involve comparisons of choice sets for the same preference relation — they do not involve intertemporal comparisons of interpersonal utility values.

$u(\tau, w)$ as a solution to

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

To compare two different choice sets, say (τ, w) and (τ', w'), we compare $u(\tau, w)$ and $u(\tau', w')$. The choice 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 τ_0 dollars, we can also meaningfully compare their magnitude and calculate the rate of growth between (τ, w) and (τ', w') as $u(\tau', w')/u(\tau, w)$.

Since financial markets are complete, we can rewrite the household's problem as-if each cohort τ chose consumption at each age t subject to an intertemporal budget constraint with intertemporal prices for age t consumption given by

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

Let $e(\boldsymbol{q}, U) = \min_{c} \{\boldsymbol{q} \cdot \boldsymbol{c} : \boldsymbol{\mathcal{U}}(\boldsymbol{c}) = U\}$ be the associated intertemporal expenditure function. The wealth required in τ_0 to reach the same indifference curve as $V(\tau, w)$ is just:

$$u(\tau, w) = e(\boldsymbol{q}(\cdot|\tau_0), V(\tau, w)) = e(\boldsymbol{q}(\cdot|\tau_0), 1)V(\tau, w) = w\frac{e(\boldsymbol{q}(\cdot|\tau_0, 1))}{e(\boldsymbol{q}(\cdot|\tau, 1))}, \tag{2}$$

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 $w = e(q(\cdot|\tau), 1)V(\tau, w)$. In words, (2) 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 relative to wealth by

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

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 (2) as

$$\log u(\tau, w) = \log w - \int_{\tau_0}^{\tau} \sum_{t=0}^{T} b(t|x) \left[\frac{d \log p(t|x)}{dx} - \sum_{z=0}^{t} \frac{d \log R(z|x)}{dx}\right] dx. \tag{3}$$

In words, order cohorts between τ and τ_0 by date. For each cohort x between τ and τ_0, calculate the rate of change in prices $d\log p(t|x)$ and returns $d\log R(t|x)$ cohort x faces and average these changes using discounted budget shares for each period of life t. Cumulate these changes over all cohorts x 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.[13]

In a static world, where the horizon is just $T = 1$, (3) simplifies to a trivial calculation:

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

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 (3) 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:

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

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

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

Substitute this into (3) to arrive at the following simple result.

Proposition 1 (Money Metric for Special Case)**.** *Money metric welfare for a household in cohort*

[13] To express $u(\tau, w)$ in (3) as a chained-weighted average of present and future price changes, we have not imposed CRRA preferences (1). Alternatively, imposing CRRA and applying exact hat-algebra, we can express $u(\tau, w)$ as $\log u(\tau, w) = \log w - \frac{1}{1-\sigma} \log \sum_{t=0}^{T} b(t|\tau_0)\left(\frac{q(t|\tau)}{q(t|\tau_0)}\right)^{1-\sigma}$, which requires knowledge of σ. We provide a lengthier derivation of Proposition 1 because it generalizes beyond this example. See also Appendix C.

τ with wealth w, in terms of τ_0 dollars, is given by:

$$\log u(\tau, w) = \log w - \log \frac{p(0|\tau)}{p(0|\tau_0)} + \frac{1}{1-\sigma} \log \frac{b(0|\tau)}{b(0|\tau_0)}. \tag{4}$$

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 (4), 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 (4) 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). However, this lengthier proof generalizes when we relax the assumptions of this example. See Appendix C for more discussion.

3 General Environment and Measure of Welfare

In this section, we set up the more general economic environment. We we first define preferences, relaxing CRRA, homotheticity, and perfect foresight, and introduce a notion of time separability. We then introduce the decision problem of households, allowing for incomplete financial markets. We then generalize the notion of money metric welfare to this environment and use it to define our measure of welfare and the cost-of-living.

3.1 Time Separable Preferences

Agents have preferences $\succeq$ over stochastic consumption streams represented by a utility function

$$\mathcal{U}\left(\{\boldsymbol{c}, \boldsymbol{\pi}\}\right), \tag{5}$$

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 (5) nest the common exponential discounting expected utility function as a special case.

Let Given some preferences (5), define the following transformation of the utility function:

$$e(\boldsymbol{q}, \boldsymbol{\pi}, U) = \min_{\boldsymbol{c}} \{\boldsymbol{q} \cdot \boldsymbol{c} : \boldsymbol{\mathcal{U}}(\{\boldsymbol{c}, \boldsymbol{\pi}\}) = U\}, \tag{6}$$

where $\boldsymbol{q}$ has the same dimensionality as $\boldsymbol{c}$. 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 (6) in practice.

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

Definition 1 (Time Separability)**.** The expenditure function is *time separable* if the function defined in (6) can be written as

$$e(\boldsymbol{q}, \boldsymbol{\pi}, U) = e\left(P\left(\boldsymbol{q}(s^0), U\right), F\left(\left\{\boldsymbol{q}(s^t)\right\}_{t>0}, \boldsymbol{\pi}, U\right), U\right), \tag{7}$$

where P and F are scalar-valued functions that are increasing in $\boldsymbol{q}$, homogeneous of degree one in $\boldsymbol{q}$, and non-decreasing in U.[14]

We say that preferences are time separable whenever the shadow expenditure function e is separable between present and future shadow prices. This implies that spending shares across goods in the present depend on future prices *only* through changes in U (i.e. wealth effects). Conversely, spending shares in the future (across dates, states, and goods) depend on present prices only through wealth effects. Condition (7) is violated if changes in intertemporal prices, say the real interest rate, cause spending shares in the present to change even holding intertemporal utility constant. (See Lemma 1 in Section 4 for more details.)

Equation (7) implies that future prices and probabilities are aggregated together into a single scalar F, which acts like a certainty-equivalent that accounts for how changes in future prices and risk affect expenditures on the future.

[14] Separability of the expenditure function is sometimes referred to as "quasi-separability" (see, e.g., Atkin et al., 2024 who use this terminology in a static context).

Condition (7) is not equivalent to direct separability of the utility function unless preferences are homothetic. When preferences are homothetic, as is the case in the vast majority of the macroeconomics literature, then separability of the expenditure function is equivalent to requiring that the utility function be expressible 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 $\boldsymbol{c}$. That is, separability of the expenditure function is equivalent to separability of the direct utility function whenever preferences are homothetic.[15]

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

Example 1 (Exponential Discounting Expected Utility)**.** Suppose that the utility function solves the following equation

$$U^{\frac{\sigma-1}{\sigma}} = \sum_{t=0}^{T} \beta^t U^{\varepsilon_t} \sum_{s_t} \pi(s^t) \left(\sum_n \omega_{nt} U^{\varepsilon_n} c_n(s^t)^{\frac{\gamma-1}{\gamma}} \right)^{\frac{\gamma}{\gamma-1} \frac{\sigma-1}{\sigma}}, \tag{8}$$

where t denotes age, β is the discount factor, ω_{nt} are taste parameters for good n at age t, γ and ρ control substitution elasticities across goods, states, and time, and ε_t and ε_n govern wealth effects for consumption at age t and consumption of good n.[16] These preferences are time separable.[17]

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 time separable without agents facing the decision problem defined by (6) in practice.

[15]See Theorem 4.3 in Blackorby et al. (1998).

[16]Note that life-cycle shifts in ω_{nt} are not taste shocks since intertemporal preferences are stable.

[17]We can also consider a non-homothetic extension of Epstein and Zin (1989):

$$U\left(s^t\right)^{\frac{\sigma-1}{\sigma}} = C\left(s^t\right)^{\frac{\sigma-1}{\sigma}} + \beta_{t+1} U^{\varepsilon_{t+1}} \left[\mathbb{E}U\left(s^{t+1}\right)\right]^{\frac{\sigma-1}{\sigma}}, \quad \text{where} \quad U = U(s^0),$$

$$\mathbb{E}U\left(s^{t+1}\right) = \left[\sum_{s_{t+1}} \pi(s^{t+1}) U\left(s^{t+1}\right)^{\frac{\rho-1}{\rho}}\right]^{\frac{\rho}{\rho-1}}, \quad \text{and} \quad C\left(s^t\right) = \left(\sum_n \omega_{nt} U^{\varepsilon_n} c_n(s^t)^{\frac{\gamma-1}{\gamma}}\right)^{\frac{\gamma}{\gamma-1}}.$$

The elasticity of substitution across time (σ) can differ from that across states (ρ). If $\sigma = \rho$, then we obtain (8).

3.2 Decision Problem of Households

Decision makers face a planning horizon of length $T < \infty$.[18] 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 (5) 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

$$\sum_{n \in N} p_n(s^0|\tau) c_n(s^0|\tau) + \sum_{k \in K} a_k(s^0|\tau) = w, \tag{9}$$

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. If there are durable goods, say housing, then the stock of durables must be included as an asset, a_k, and the user-cost of service flows must be included as a price, p_k. This is how we treat housing in our empirical application.

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

$$\sum_{n \in N} p_n(s^t|\tau) c_n(s^t|\tau) + \sum_{k \in K} a_k(s^t|\tau) = \sum_{k \in K} R_k(s^t|\tau) a_k(s^{t-1}|\tau) + y(s^t|\tau), \tag{10}$$

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 the payoff from assets that cannot be bought or sold (e.g. human capital that pays a wage every period but cannot be sold). If $y(s^t|\tau) = 0$ for every s^t, we say that

[18] Our theoretical results can be extended to infinite horizon (e.g. bequests) but we abstract from these considerations in this paper.

the household is a *rentier*.[19] We also impose borrowing constraints requiring that

$$\sum_k a_k(s^t|\tau) \geq -X(s^t|\tau) \tag{11}$$

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, incomes, probabilities, borrowing constraints, and wealth: $\{\boldsymbol{p}, \boldsymbol{R}, \boldsymbol{y}, \boldsymbol{\pi}, \boldsymbol{X}, w\}$. Define the value function associated with each problem to be

$$V(\{\boldsymbol{p}, \boldsymbol{R}, \boldsymbol{y}, \boldsymbol{\pi}, \boldsymbol{X}, w\}) = \max_{\boldsymbol{c},\boldsymbol{a}} \{\boldsymbol{\mathcal{U}}(\{\boldsymbol{c}, \boldsymbol{\pi}\}) : \text{constraints (9), (10), (11) are satisfied}\}. \tag{12}$$

The value function ranks decision problems according to underlying preference relation. In a static, deterministic environment, the value function in (12) collapses to the indirect utility function in consumer theory, which ranks static budgets sets, defined by static prices and wealth, into utils.

3.3 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 2 (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

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

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 $u(\{\boldsymbol{p}, \boldsymbol{R}, \boldsymbol{y}, \boldsymbol{\pi}, \boldsymbol{X}, w\}|\tau_0)$. In a static deterministic world, the generalized money metric coincides with the traditional money metric.[20]

[19]Under complete markets, every household is a rentier. In Section 4.4, we extend the set of rentiers to include households with a risk free cash flow $y(s^t|\tau) = y(t|\tau)$ that do not face borrowing constraints (outside of the no-Ponzi condition).

[20]In principle, there are many different ways one could measure welfare. For example, we could

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 3 (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.

4 Main Results

We present our main result in steps. First, we start with some preliminaries in Section 4.1. In Section 4.2 we present a method for recovering the generalized money metric for rentiers. In Section 4.3 we show how to recover the money metric for non-rentiers.

4.1 Shadow Prices and Compensated Demand

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

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.

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 (12). Let $c_n^*(s^t|\boldsymbol{q}^*, \boldsymbol{\pi}, V(\tau, w, \boldsymbol{y}))$ be the shadow quantities that minimize the shadow expenditure function (6).

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

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

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 (12).

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 $(\tau, w, \boldsymbol{y})$. Denote spending on the present relative to shadow wealth by

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

The share of spending on the future (superscript F) is

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

Definition 4 (Elasticity of Intertemporal Substitution)**.** The compensated EIS for a household facing problem $(\tau, w, \boldsymbol{y})$ is defined to be

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

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.

4.2 Recovering Money Metric for Rentiers

The next proposition limits attention to the subset of rentiers. Rentiers are households for whom $\boldsymbol{y}$ is identically equal to $\mathbf{0}$ in every state of nature, and are denoted by $(\tau, w, \mathbf{0})$.

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

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

The following proposition characterizes the money metric utility of rentiers as the solution to an integral equation. To do this, assume that, for every state s^t, prices $p_n(s^t|\tau)$, asset returns $R_k(s^t|\tau)$, and probabilities $\pi(s^t|\tau)$ are absolutely continuous functions of calendar time τ.

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

$$\log u(\tau, w, \mathbf{0}) = \log w - \int_{\tau_0}^{\tau} \left(\sum_{n \in N} B_n(x, w_x^*, \mathbf{0}) \frac{d \log p_n}{dx} + \frac{1}{\sigma(x, w_x^*, \mathbf{0}) - 1} \frac{d \log B^P(x, w_x^*, \mathbf{0})}{dx} \right) dx, \tag{13}$$

where w_x^* solves the equation

$$u(x, w_x^*, \mathbf{0}) = u(\tau, w, \mathbf{0}). \tag{14}$$

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

Proposition 4 generalizes the results in Baqaee et al. (2024) 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. Both inflation and changes in expenditures out of wealth at $x \in [\tau_0, \tau]$ must be calculated using compensated demand, at $u(\tau, w, \mathbf{0})$, because the household is kept indifferent between $(\tau, w, \mathbf{0})$ and $(\tau_0, u(\tau, w, \mathbf{0}), \mathbf{0})$.

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 (13) 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 (13) to get a new guess. We then iterate on this until convergence. This procedure will always converge since the fixed point to (13) 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 wealth levels, say $w \in [\underline{w}_x, \overline{w}_x]$ for $x \in [\tau_0, \tau]$. This limits the range of values of w for which we can calculate the money metric without out-of-sample extrapolation. Intuitively, if for cohort τ and wealth w, the money metric value $u(\tau, w, \mathbf{0})$ is not in $[u(x, \underline{w}_x, \mathbf{0}), u(x, \overline{w}_x, \mathbf{0})]$ for some $x \in [\tau_0, \tau]$, then we cannot recover $u(\tau, w, \mathbf{0})$ without extrapolation. This is because there are no households in cohort $x \in [\tau_0, \tau]$ that are on the same indifference curve as the rentier with wealth w at time τ.

Fortunately, Proposition 4 automatically provides the boundary over which the money metric can be calculated without extrapolation. The initial boundary at $t = \tau_0$ is just the range in the data: $[\underline{w}_{\tau_0}, \overline{w}_{\tau_0}]$. As we solve (13) forward, for each $\tau > \tau_0$, we can update the boundary because the information required to compute only depends on previous values of the money-metric (see Baqaee et al., 2024 for a discussion of this issue in a static context).

Sketch of Proof. To better understand Proposition 4, we sketch the proof. (The formal proof is in Appendix B). 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, for rentiers we can write (in an abuse of notation)

$$\boldsymbol{q}^*(s^t|\tau, w, \mathbf{0}) = \boldsymbol{q}^*\left(s^t|\tau, V(\tau, w, \mathbf{0})\right). \tag{15}$$

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 rentiers is that (15) need not hold for non-rentiers. For households with idiosyncratic undiversifiable income (i.e. non-rentiers), 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

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

That is, for rentiers, 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

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

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.[21]

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

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

of calculus, as

$$\log u(\tau, w, \mathbf{0}) = \log w + \int_{\tau}^{\tau_0} \sum_{t=0}^{T} \sum_{s^t} \left(\frac{\partial \log e(\boldsymbol{q}^* (s^t|x, u(\tau, w, \mathbf{0})), \boldsymbol{\pi}(s^t|x), u(\tau, w, \mathbf{0}))}{\partial \log \boldsymbol{q}^*(s^t|x, u(\tau, w, \mathbf{0}))} \cdot \frac{d \log \boldsymbol{q}^* (s^t|x, u(\tau, w, \mathbf{0}))}{dx} \right.$$
$$\left. + \frac{\partial \log e(\boldsymbol{q}^* (s^t|x, u(\tau, w, \mathbf{0})), \boldsymbol{\pi}(s^t|x), u(\tau, w, \mathbf{0}))}{\partial \boldsymbol{\pi}(s^t|x)} \cdot \frac{d\boldsymbol{\pi}(s^t|x)}{dx} \right) dx. \tag{16}$$

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, 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 (16) 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. Denote the compensated budget share on good n to be

$$b_n(\boldsymbol{q}^*, U) = \frac{q_n^*(s^0)c_n^*(\boldsymbol{q}^*, U)}{\sum_{n' \in N} q_{n'}^*(s^0)c_{n'}^*(\boldsymbol{q}^*, U)}.$$

When preferences are time separable, the complicated integrand in (16) can be rewritten as

$$-\frac{d \log b^P(\boldsymbol{q}^*, \boldsymbol{\pi}, V(\tau, w, \mathbf{0}))}{1 - \sigma(\tau, w, \mathbf{0})} + \sum_{n \in N} b_n(\boldsymbol{q}^*(s^0), \boldsymbol{\pi}, V(\tau, w, \mathbf{0})) d \log q_n^*(s^0),$$

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 (16) and manipulating leads to Proposition 4. The last step uses the insight from Baqaee et al. (2024) that, with the addition of (14), we can treat (16) as a fixed point problem. Proposition 4 implies that for rentiers, we can recover $u(\tau, w, \mathbf{0})$ 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 $u(\tau, w, \mathbf{0})$, we need to integrate compensated budget shares and compensated changes in savings rate. On the other hand, to perform the necessary compensation, we need to know $u(\tau, w, \mathbf{0})$. This fixed point problem disappears when we specialize preferences to be homothetic.

To build intuition, we consider the homothetic special case of Proposition 4.

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

$$\log u(\tau, w, \mathbf{0}) = \log w - \int_{\tau_0}^{\tau} \left(\sum_n B_n(x) \frac{d \log p_i}{dx} + \frac{1}{\sigma(x, \mathbf{0}) - 1} \frac{d \log B^P(x, \mathbf{0})}{dx} \right) dx. \tag{17}$$

When preferences are homothetic, the integrand in (13) simplifies. First, the share of spending on each good $B_n(x, w_x^*, \boldsymbol{y})$ is only a function of the time period x, so we write it as $B_n(x)$. 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 $B_n(x, w, \boldsymbol{y})$ is the same for all households at time x.

Second, the compensated change in the consumption share of wealth $d \log B^P(x, w_x^*, \mathbf{0})/dx$ simplifies to the uncompensated change in the consumption share of wealth $d \log B^P(x, \mathbf{0})/dx$. This is because the consumption share of wealth $B^P(x, w, \boldsymbol{y})$ 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.

Corollary 1 simplifies further if we assume that the EIS is constant (as in the illustrative example in Section 2):

$$\log u(\tau, w, \mathbf{0}) = \log w - \int_{\tau_0}^{\tau} \sum_i B_n(x) \frac{d \log p_i}{dx} dx + \frac{\log \left(B^P(\tau, \mathbf{0}) / B^P(\tau_0, \mathbf{0}) \right)}{1 - \sigma}.$$

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.

Justifying use of static deflators. There are two cases of Proposition 4 where the use of a static deflator for dynamic welfare is justifiable (outside of the obvious case where households are fully myopic). The first one is if the EIS is infinite, in which case, (13) simplifies to

$$\log u(\tau, w, \mathbf{0}) = \log w - \int_{\tau_0}^{\tau} \sum_{n \in N} B_n(x, w_x^*, \mathbf{0}) \frac{d \log p_n}{dx} dx. \tag{18}$$

That is, dynamic welfare is given by nominal wealth deflated using a deflator that relies only on static price changes between τ_0 and τ.

The second case is when the EIS is zero and preferences are homothetic, in which case (13) can be rewritten as

$$\log \frac{u(\tau', w', \mathbf{0})}{u(\tau, w, \mathbf{0})} = \log \frac{E(\tau', w', \mathbf{0})}{E(\tau, w, \mathbf{0})} - \int_{\tau}^{\tau'} \sum_{n \in N} B_n(x, \mathbf{0}) \frac{d \log p_n}{dx} dx.$$

That is, the growth in money metric wealth from (τ, w) to (τ', w') is just the growth in nominal expenditures deflated by a chained static price deflator between τ and τ' (i.e. chained real consumption growth). This is consistent with Reis (2005), who argues that when the EIS is zero and agents are rentiers, then static real consumption growth coincides with dynamic welfare growth. This result breaks down however, even when the EIS is zero, if households are non-rentiers or if preferences are non-homothetic.

When preferences are non-homothetic and the EIS is zero, then (13) can be rewritten as

$$\log E(\tau_0, u(\tau, w, \mathbf{0}), \mathbf{0}) = \log E(\tau, w, \mathbf{0}) - \int_{\tau_0}^{\tau} \sum_{n \in N} B_n(x, w_x^*, \mathbf{0}) \frac{d \log p_n}{dx} dx.$$

This implies that nominal expenditures deflated by static inflation (calculated using compensated budget shares) is equal to the expenditures of the compensated household in τ_0. If expenditures are increasing in wealth, then this is also a money metric since it is a monotone transformation of utility. This would justify the empirical approach taken by Baqaee et al. (2024), which ignores time, as long as we also assume that all households are rentiers.

4.3 Non-Rentiers

A challenge for the applicability of Proposition 4 is that, in practice, many households in the sample are non-rentiers (i.e. $\boldsymbol{y} \neq 0$). Fortunately, we can exploit 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.

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

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

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

Importantly, Lemma 1 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 non-rentier households. To do so, with some abuse of notation, index money-metrics by their base period. That is, let $u_{\tau_0}(\tau, w, \boldsymbol{y})$ 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 5 (Money Metric is a Function of Budget Shares and Time)**.** *Suppose that the vector-valued function $\boldsymbol{b}(\boldsymbol{p}, V)$ is a one-to-one 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 a one-to-one 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 5 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 $u_\tau(\tau, w, \mathbf{0}) = w$ for rentiers. Given this, we can learn the shape of m by solving the following least-squares problem

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

where M is a set of functions that contains m. In words, we fit a flexible function that relates budget shares to wealth for rentiers. We then use this fitted relationship to impute $u_\tau(\tau, w, \boldsymbol{y})$ for non-rentiers given their static budget shares. Intuitively, if there exists a subset of rentiers and non-rentiers with the same preferences, beliefs about the future, and static prices, given observable characteristics, then we can infer their money metric wealth using their static budget shares. Proposition 6 formalizes this idea.

Proposition 6 (Money Metric for Non-Rentiers)**.** *Let*

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

where $[\underline{w}_\tau, \overline{w}_\tau]$ is the support of the wealth distribution of the rentiers at date τ. Let $m|_{\mathcal{B}(\tau)}$ be the function m restricted to the domain $\mathcal{B}(\tau)$. We have that

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

A special case of Proposition 5 and Proposition 6 is the case where the budget share of a specific good, usually food, is known to be strictly monotone in utility.

Corollary 2 (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, 5, and 6 can be combined to recover $u(\tau, w, \boldsymbol{y})$ for every $u(\tau, w, \boldsymbol{y})$ inside the boundary where we can solve (13) (without extrapolation).

4.4 Extensions

Relaxing common prices and probabilities. We assume that, conditional on observables (like age, gender, location, and rentier-status) households have the same preference relation within and across cohorts.[22] We also assume that static prices only vary as a function of observables (e.g. time or location). Similarly, cohorts of rentiers at each point in time must hold common beliefs about future prices and rates of return. Beliefs can change over time but, within a period, they can only vary for rentiers as a function of observable characteristics (e.g. age).

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

Finally, we do not require that households' beliefs about the future be "objective" 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.

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

Specifically, the first period budget constraint (9) is

$$\sum_{n \in N} p_n(s^0|\tau) c_n(s^0|\tau) + \sum_{k \in K} a_k(s^0|\tau) + \sum_{m=1}^{T} B_m(s^0|\tau) = w + y(s^0|\tau), \tag{19}$$

[22] To see what can go wrong, suppose that the preference relation of cohorts is changing over time. Then, changes in consumption-wealth ratios or budget shares may be due to changes in preference parameters (e.g. discount factor changes), and not due to changes in welfare-relevant variables like prices, returns, or probabilities. Note that preference stability is a typical maintained assumption in the literature on static welfare measures. To deal with preferences instability, e.g. taste shocks over time, we would need to specify a model of demand to purge out changes in choices that are driven by taste shocks as opposed to income and substitution effects, as discussed in Baqaee and Burstein (2023).

where $B_m(s^0|\tau)$ is the quantity of bond of maturity $m \in \{1, \ldots, T\}$ purchased in period τ, with return $R(m|\tau)$ at date $\tau + m$. At each subsequent history s^t, the budget set is then

$$\sum_{n \in N} p_n(s^t|\tau)c_n(s^t|\tau) + \sum_{k \in K} a_k(s^t|\tau) = \sum_{k \in K} R_k(s^t|\tau)a_k(s^{t-1}|\tau) + B_m(s^0|\tau)R(t|\tau) + y(t|\tau). \tag{20}$$

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

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

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

That is, households with risk-free exogenous cash flows and no ad-hoc borrowing constraints are isomorphic to rentiers whose wealth is augmented by the present discounted value of exogenous cash flows.

5 Illustrative Application to US Data

In this section, we apply our method to data from the US. In the next section, we use our estimates of money metric wealth to study how welfare responds to job loss.

5.1 Data

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

The household data must be accurate, but it does not need to be representative of the underlying population in terms of sampling frequency. Therefore, we can use the raw data from the PSID without sampling weights. The PSID contains repeated cross-sectional data on household expenditures by category, household-level balance sheets (assets and liabilities), household incomes, and demographic information. Since lifecycle considerations are important, we split the household sample by age. Ideally, with enough observations, we could treat each age separately. Unfortunately, our dataset is not very large, so we split the sample into only two groups: above or below sixty years old, and assume that households in each age group at each point have the same preferences and face the same prices and probabilities. We provide sensitivity with respect to this cutoff in the appendix.[23]

The PSID includes household expenditure surveys, broken down into seven categories. A major omission is the user cost of owner-occupied housing. To remedy this, we impute equivalent owner-occupied housing costs by matching home owners in each period to renters with similar observable characteristics and spending behavior. That is, we predict rental expenditures using household characteristics and spending behavior using a regression estimated on renters in the same period. This procedure is theoretically justified by Proposition 5. In the final year of our sample, 2019, the PSID asked home owners to report the rental value of their property if they were to rent it out. We use the answers to this question, in 2019, to validate our imputation procedure. When we regress surveyed housing costs on our imputed measure of housing costs (both relative to current expenditures), we find a coefficient of 1.03 and an R^2 value is 0.59. See Appendix A for details.[24]

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

5.2 Classifying Rentiers

To apply Proposition 4, we need to observe a sample of rentier households. To classify rentiers, we first estimate a proxy measure of total wealth for all households in the sample. Our proxy for total wealth is the sum of financial wealth (net asset value including home equity and defined contribution pension savings) and the present discounted value of 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 calculate the present discounted value of labor and transfer income, we predict each household's expected lifetime income profile based on observed characteristics, and discount the resulting flows using a real discount rate of 4% following Catherine et al. (2022). The construction of net assets and capitalized income is detailed in Appendix A. We say that a household is a potential rentier if net financial assets constitute more than 90% of total wealth.

[23]Similarly, with more data, we could split the sample along other observed characteristics that may influence preferences, like gender, household size, location, etc.

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

In practice, we do not know with certainty who the rentiers are. Our results can suffer from bias if we include non-rentiers in the sample of rentiers. On the other hand, our estimates will have higher variance if we accidentally exclude true rentiers from the sample of rentiers. Since sampling uncertainty can be quantified, we choose to be more conservative in terms of minimizing bias by further limiting the set of rentiers. To do so, we exclude households from the rentier set if their current consumption-wealth ratio is more than 20% on the basis that, for these households, we may be underestimating the value of their non-financial wealth. To further limit the role of outliers, we also exclude households whose net assets are in the top and bottom 2.5%.[25]

For the sample of potential rentiers, following Propositions 5 and 6, we regress measured total wealth at date τ on a polynomial of budget shares at date τ for each age group separately:

$$w_{h,\tau} = \alpha_{0,\tau} + \sum_{k=1}^{2} \sum_{i=1}^{N} \alpha_{i,\tau,k} B_{i,h,\tau}^{k} + \text{error}_{h,\tau}, \tag{21}$$

where α's are regression coefficients. If a potential rentier is an outlier in this regression, so predicted and measured total wealth differ significantly, then we exclude these households from the set of rentiers. A potential rentier is an outlier if the Cook's distance value is greater than one (we check sensitivity to this assumption by varying this cutoff value in the appendix).

5.3 Money Metric Wealth for Rentiers

Figure 1 shows a scatterplot of log consumption wealth ratios against log wealth for rentiers in the two age groups. As might be expected, there are many more rentiers in the older age group, since this group contains retirees as compared to the younger group. As we show later, this results in much more sampling uncertainty of money metric wealth for the younger group than for the older group.[26]

Both panels of Figure 1 show a strong decreasing relationship between consumption-wealth ratios and wealth. The relationship is approximately loglinear and the slope is roughly the same for both age groups. These findings are consistent with those of Straub (2019) who documents that the consumption-wealth ratio strongly declines in permanent income. Whereas static non-homotheticity is a cornerstone of the literature on consumption, dynamic non-homotheticity, like the one depicted in Figure 1, is relatively understudied. However, as our estimates of money metric wealth show, in our dataset,

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

[26] Our sample includes 300 rentiers below sixty years old (38 per year on average) and 1610 above sixty years old (201 per year).

Figure 1: Log consumption-wealth against log wealth

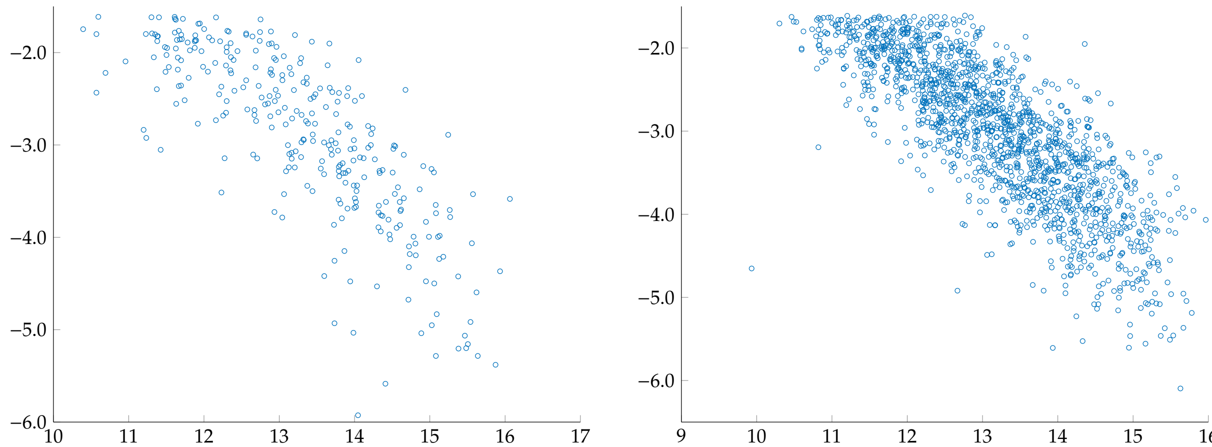

(a) < 60 years old

(b) ≥ 60 years old

dynamic non-homotheticity is an order of magnitude more powerful than static non-homotheticity.[27]

To compute money metric values for rentiers using Proposition 4, we need to evaluate consumption-wealth ratios and static budget shares as a function of date and wealth for each age group. To do this, we fit smooth curves that relate consumption-wealth ratios and budget shares to date and wealth level in each period. That is, let h index rentiers and τ index dates. For each age group, we use locally weighted scatterplot smoothing (LOWESS) to fit cross-sectional mapping from wealth and age to consumption-wealth ratios:

$$\log B^{P}_{h,\tau} = \hat{B}(\log \text{wealth}_{h,\tau}, \tau) + \text{error}_{h,\tau},. \tag{22}$$

Similarly, for each good $i \in N$, we use LOWESS to fit the cross-sectional mapping from wealth and date to budget shares on each good for each age group. We choose the smoothing parameter of LOWESS by cross-validation (i.e. maximize out-of-sample predictive success).

We use our estimated cross-sectional curves in Proposition 4 to recover money metric utility as a function of date and wealth for each age group.[28] For illustration, we use the initial year, $t_0 = 2005$, as the base year, so that money metric wealth values map nominal wealth in each year t into equivalent wealth in 2005. For our benchmark results, we set the

[27]As we discuss below, a possible reason why static non-homotheticity is mild in our dataset is that our expenditure data is heavily aggregated (seven expenditure categories). It is possible that with more disaggregated data, static non-homotheticity might play a more important role.

[28]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. (2024).

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

Unfortunately, Best et al. (2020) do not estimate Hicksian elasticities but Marshallian ones. Luckily, theoretically, Slutsky's equation implies that, if consumption is a normal good, then the Hicksian intertemporal elasticity should be smaller in magnitude than the Marshallian one. Since the elasticity is bounded below by zero, we experiment with lower values of σ in the appendix and find that our results are not sensitive to using values of σ lower than 0.1. In the appendix we provide sensitivity of our results to setting $\sigma = 0.2$ and $\sigma = 0.05$.[29]

Figure 2 plots the money metric, for each age group, as a function of 2019 wealth for 2005 base prices. For comparison, we also plot a naive calculation that deflates nominal wealth in 2019 by the official CPI inflation between 2005 and 2019. The confidence bands are calculated by bootstrap. As expected, sampling uncertainty is higher for the younger age group where we have significantly fewer rentiers than for the older age group. In both cases, non-homotheticity is important since both red lines cross the black line. When the red line is above the black line, static official inflation overstates the true inflation rate. Conversely, when the red line is below the black line, the official inflation rate understates the true inflation rate.

For the younger group, true inflation is higher for most wealth levels than CPI inflation. On the other hand, for the older group, true inflation is lower than CPI inflation for households with wealth above 2 million dollars and higher for poorer households. Of course, all of these conclusions are subject to a high degree of sampling uncertainty (this sampling uncertainty would diminish if our dataset were larger).

The true inflation rate between τ_0 and τ for a household with wealth w at time τ, in log terms, is defined as $100 \times \log\left(w/u_{\tau_0}(\tau, w)\right)$. We plot this in Figure 3 for $\tau_0 = 2005$ and $\tau = 2019$. For younger households, the inflation rate is relatively stable as a function of wealth, ranging from 2% for the poorest households to 3% for the richest households. On the other hand, for the older group, there is a stronger negative relationship with wealth, and inflation ranges from 3% for the poorest to 1% for the richest. As in Figure 2, sampling uncertainty is significantly higher for the younger group.

To understand the patterns in Figure 3, we decompose the welfare-relevant inflation

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

Figure 2: Money metric wealth in 2005 base prices

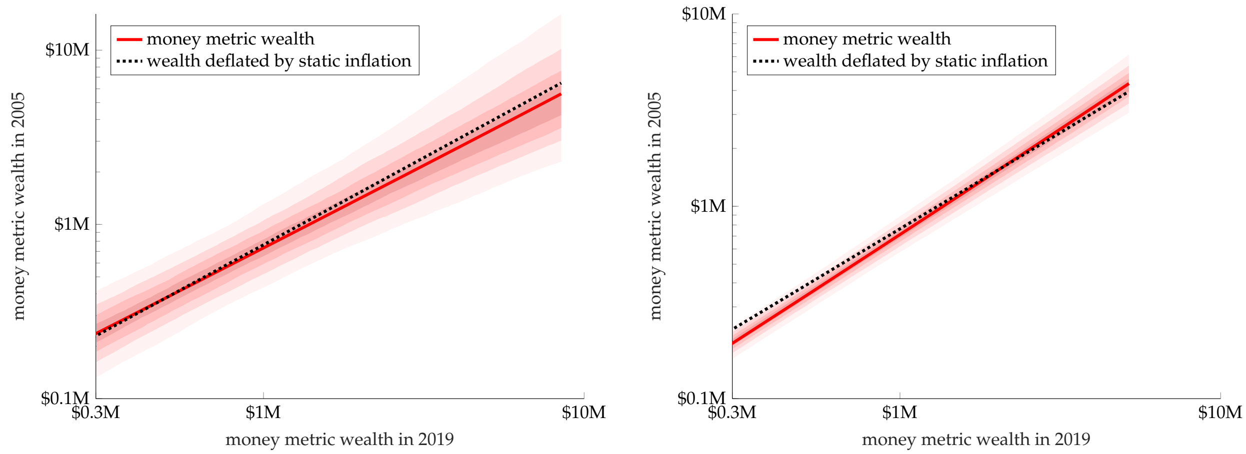

(a) < 60 years old (b) ≥ 60 years old

Figure 3: Annualized dynamic inflation between 2005 and 2019 as a function of wealth in 2019

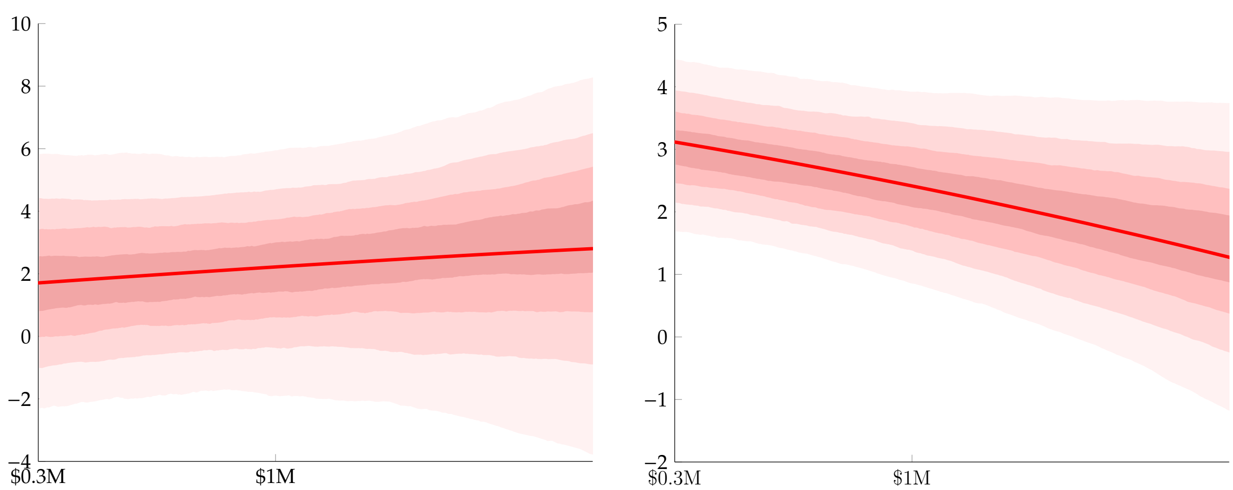

(a) < 60 years old (b) ≥ 60 years old

rate, in equation (13), 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:

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

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

Figure 4: Decomposition of annualized inflation between 2005 and 2019 for each wealth in 2019

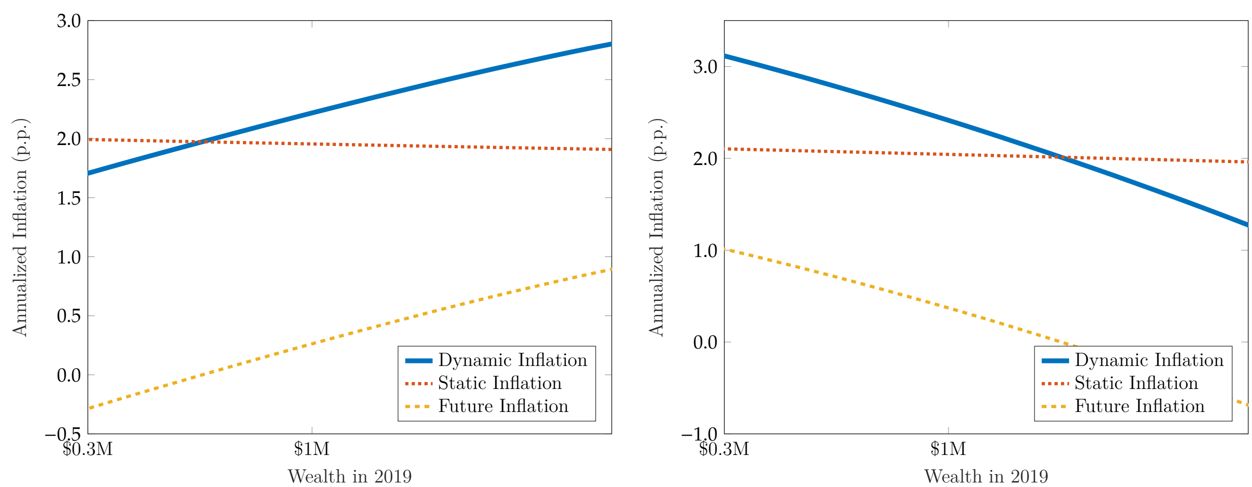

(a) < 60 years old (b) ≥ 60 years old

This decomposition is shown in Figure 4. The static inflation term is not exactly the same as aggregate CPI because it weighs changes in static prices using compensated budget shares rather than aggregate budget shares. Nevertheless, the static component is very close to aggregate CPI inflation at around 2% per year for all wealth levels. The very slight downward slope reflects the non-homotheticity of static preferences, and static inflation is slightly higher for poorer households, consistent with other studies of the US, like Jaravel and Lashkari (2024), that show that the static cost-of-living index has tended to rise more quickly for poorer households. Nevertheless, the slope of the static inflation line is very mild compared to the slope of the dynamic inflation measure.[30]

The component corresponding to future inflation is different for the two age groups and exhibits a strong dependence on wealth. For most of the households in the younger group, future prices are expected to rise more quickly than present prices, which explains why the dynamic all-encompassing cost-of-living index is mostly above the static inflation line in panel 6a. Furthermore, the richer households expect higher inflation than the poorer households, which is why the line has an upward slope. On the other hand, for the older group, these patterns are reversed. Future inflation is expected to be lower than static inflation for very wealthy households and higher than static inflation for poorer households.

The future component of our dynamic inflation measure is proportional to the compensated change in the log consumption wealth ratio. Figure 5 plots both the compensated and uncompensated log change in the consumption wealth ratio between 2005 and 2019 as a function of nominal wealth for both age groups. In both cases, the uncompensated change in the consumption wealth ratio is much more positive than the compensated one. This is because there is a strong wealth effect whereby the consumption wealth ratio declines as households become richer. For a given nominal level of wealth, households in 2005 are on a higher indifference curve than households with that level of wealth in 2019 because of positive inflation. Therefore, the wealth effect means that such households would have higher consumption wealth ratios in 2019 than in 2005, even if relative prices do not change. The changes in compensated consumption wealth ratios, which are purged of wealth effects, are lower and reflect only substitution effects. This figure underscores the importance of accounting for wealth effects when using consumption wealth ratios to infer changes in relative prices.

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 need not be caused by differential exposures to future goods prices alone. Even if all households are symmetrically exposed to future goods prices, the future component of inflation can differ across households because of differences in expected returns of assets (see Fagereng et al., 2022). For example, if poor rentiers in the younger group and rich rentiers in the older group are more reliant on, say, real estate to finance their consumption,

[30]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 as a function of current expenditures, as is done in static studies of the cost-of-living. Third, our sample period of fourteen years is reasonably short compared to previous studies, which compute changes over 50 years or longer. Finally, we have only seven spending categories, and static non-homotheticities may be stronger at more disaggregated levels.

Figure 5: Change in log consumption wealth ratios between 2005 and 2019 for each wealth in 2019

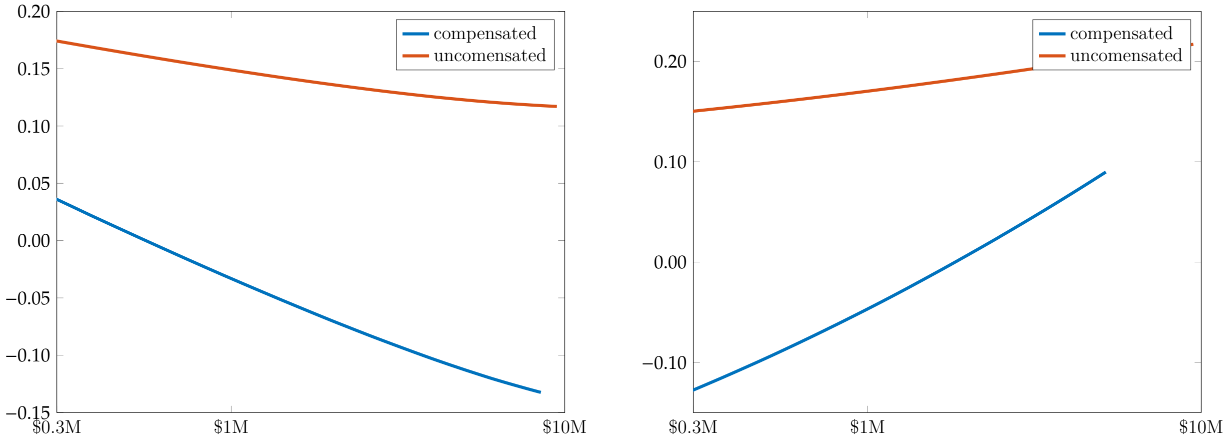

(a) < 60 years old

(b) ≥ 60 years old

then an increase in return to real estate will lower the future component of inflation for such households.

5.4 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 5 and 6 instead. For each date τ and each age group, we use (21) to predict money metric wealth for the non-rentiers conditional on their budget shares.

This is analogous to Hamilton (2001), and more recently Atkin et al. (2024), who use relative budget shares within a subset of goods, in their case food, to infer changes in welfare in a static context. Unlike Atkin et al. (2024), who compare relative budget shares across time (adjusted for substitution effects) to infer changes in money metric income over time, we compare relative budget shares within each period across rentier and non-rentier households. Since rentiers and non-rentiers face the same relative prices at each point in time, we do not have to correct relative budget shares for substitution effects and can deduce money metric wealth for non-rentiers from the rentiers.

Figure 6 displays the distribution of money metric wealth, in 2019, for both age groups. The median household in the younger age group is richer than the median household in the older age group.[31] As explained above, we are unable to estimate money metric wealth

[31]Even for the set of rentiers, the median household in the younger age group is wealthier than the median household in the older age group.

Figure 6: Distribution of money metric wealth in 2019 by age group

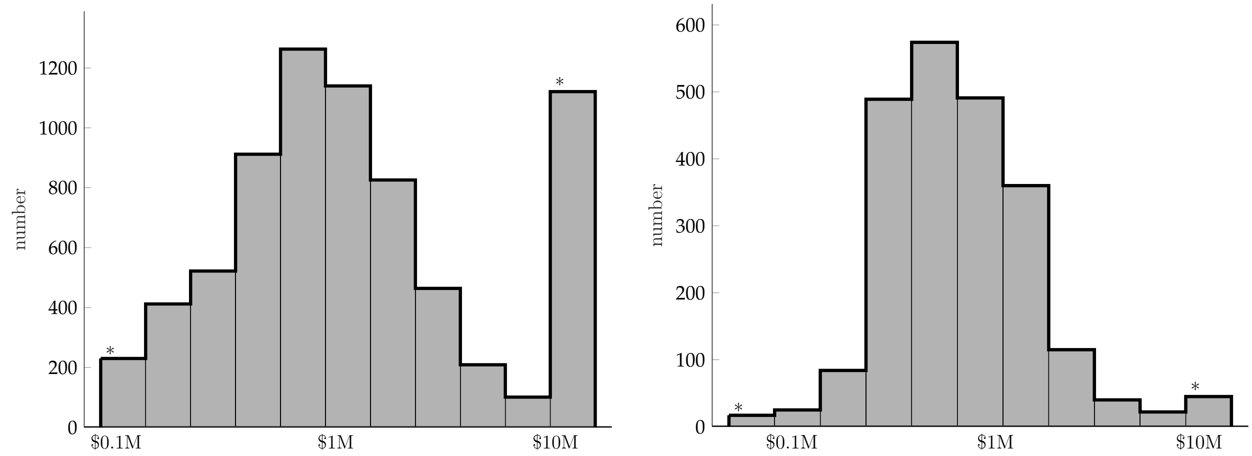

(a) < 60 years old

(b) ≥ 60 years old

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

for households whose estimated wealth values are outside of the support of the rentier wealth distribution (unless we extrapolate). Since there are many fewer rentiers in the younger age group, the set of households for whom we cannot estimate money metric wealth values is much larger.

If financial markets are complete, instead of using Equation (21), we can calculate 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).[32] This is the proxy wealth measure we initially use to separate rentiers from non-rentiers. Figure 7 plots the money metric wealth, from (21), against this alternative measure of wealth.

For the rentier subsample, the line of best fit is the 45-degree line for both age groups. This is mechanical because (21) is estimated using the rentiers. The fact that the R^2 in these regressions is positive shows that budget shares do contain information about wealth, as in Proposition 5. For non-rentiers, the fit is much worse. Notably, there are some households that seem to have little wealth as judged by their assets and future predicted income flows, but have high wealth as judged by their static spending shares.

Once in possession of money metric wealth for non-rentiers at each date, we convert these estimates into a single base year using the money metric functions in Figure 2. Figure 8 plots the growth rate for the median money metric wealth for both age groups. The

[32] See Appendix A for more details.

Figure 7: Capitalized wealth against money metric wealth (in logs)

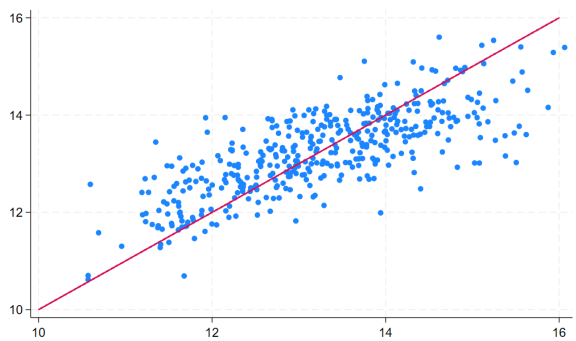

(a) Rentiers, < 60 years old

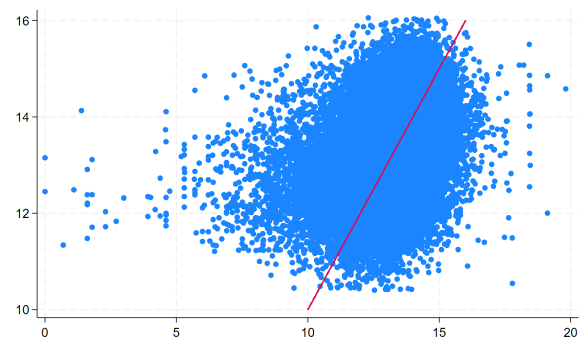

(b) Non-rentiers, < 60 years old

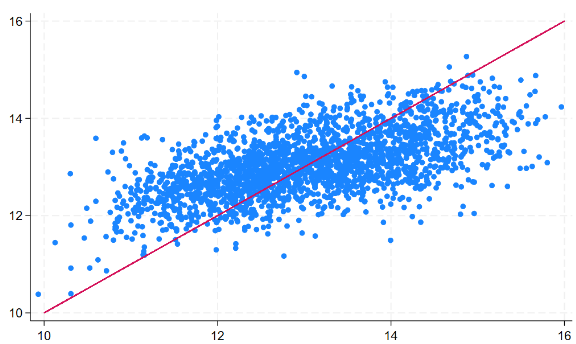

(c) Rentiers, ≥ 60 years old

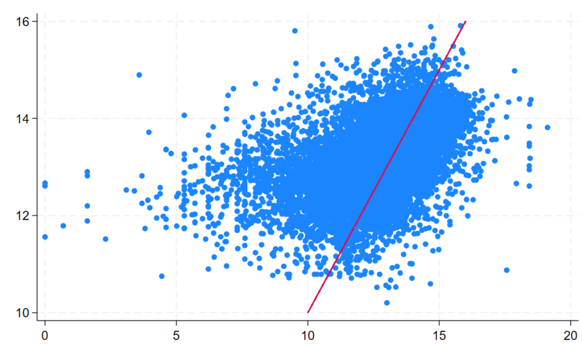

(d) Non-rentiers, ≥ 60 years old

median grew by around 1% per year for both age groups. Figure 8 also shows the growth rate if we deflate median wealth using official CPI inflation instead. The CPI implies very different growth rates for the two age groups, with the median for the younger group growing more than 2% per year and the median for the older group being close to zero. However, according to our estimates, the CPI understates inflation for the median in the younger group and overstates inflation for the median in the older group. The last two bar graphs show growth in nominal median income deflated by CPI for our sample of households. This is a static flow measure, which simply deflates income flows by static inflation. This measure has stagnated between 2005 and 2019 for the median household if we pool age groups, but it has been negative for the younger group and positive for the older group.

Figure 8: Growth for the median household

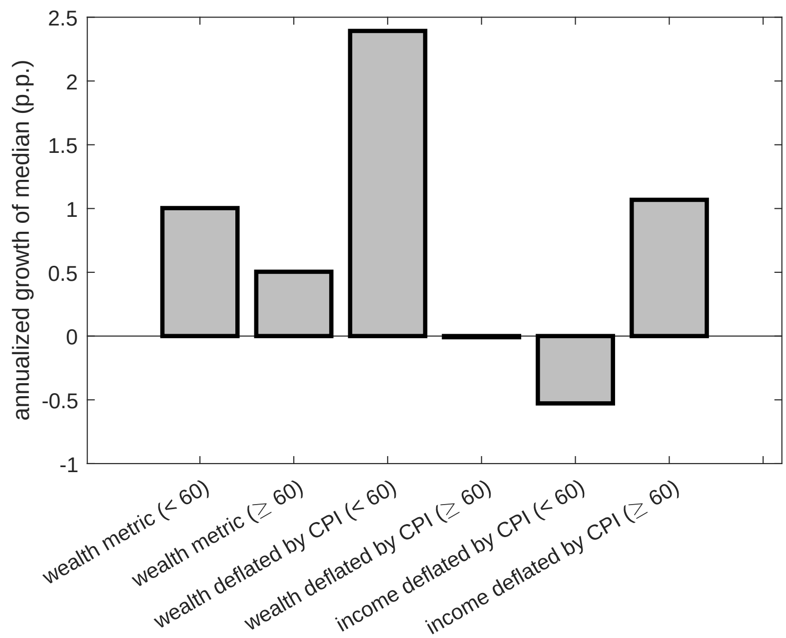

6 Treatment Effect on Welfare

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

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

Our methodology provides a complementary approach. Instead of enumerating and estimating all potentially relevant margins, we back out the component of welfare that depends on expectations about the future from observed changes in consumption-savings behavior. We illustrate this by studying the welfare effects of job loss using the PSID. We regress log money metric wealth for households on a dummy variable for job loss for the head of the household. Our measure of job loss is equal to one if the head of household loses her job and reports that she is searching for a new job in that period. To control for confounds and selection, we include year fixed effects, demographic controls, and lagged log money metric wealth.

Table 1: Percent change in money metric wealth due to job loss

	log nominal money metric wealth			log money metric wealth in 2019 dollars		
	(1)	(2)	(3)	(4)	(5)	(6)
Job loss	-0.203***	-0.209***	-0.212***	-0.194***	-0.200***	-0.214***
	(0.014)	(0.015)	(0.015)	(0.020)	(0.021)	(0.021)
Job loss × $\mathbf{1}(\text{age} \geq 60)$		0.127***	0.114**		0.112**	0.106**
		(0.046)	(0.048)		(0.046)	(0.049)
Lagged LHS	Yes	Yes	No	Yes	Yes	No
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	84,398	84,398	84,398	61,321	61,321	61,321

Standard errors of the estimates coefficients across 500 bootstraps in parentheses

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

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

The first three columns of Table 1 report the results in nominal terms and the last three columns report the results in terms of 2019 dollars. The results are very similar. In both cases, job loss is associated with a reduction in money metric wealth of around 20 log points for the younger age group. This reduction is only around 8 log points for the older age group (and less statistically significant). The identifying variation in these regressions comes from changes in household's static budget shares. When household's lose their job, their spending patterns change in a way that, had they been rentiers, would suggest a 20% reduction in their total wealth. For older household heads, the change in spending patterns following job loss suggest a less drastic reduction in wealth than for younger households. However, for both groups, the welfare losses are very large. Nevertheless, they are still much smaller than the effect on household income (which falls by around 85% in that period).

We can compare our estimate to the dynamic consequences of job loss in Davis and Von Wachter (2011). They estimate the present-value of earnings losses after mass-layoff events to be around 12% of counterfactual earnings (using a 5% discount rate). This is an alternative approach to our methodology for condensing the effect of a dynamic shock, like job loss, into a single welfare-relevant number. Our point estimates are larger than the present-value earnings losses estimated by Davis and Von Wachter (2011).

The fact that our point estimates are somewhat different to theirs is not surprising, because it is only under very strong assumptions that our estimates would coincide with their present-value calculation. First, our calculation does not assume complete markets and if the household's marginal utility is high in states where earnings are low, as is probably realistic, then a constant discount factor understates the welfare losses of job loss. Second, Davis and Von Wachter (2011) estimate ex-post earnings losses whereas we estimate ex-ante welfare losses. If households do not have perfect foresight, then ex-ante welfare losses can be larger than average ex-post losses due to risk-aversion. Relatedly, we do not impose rational expectations, so households' ex-ante beliefs about the consequences of job loss may be more pessimistic than what Davis and Von Wachter (2011) estimate. Third, we do not assume exponential discounting — if households are present-biased, then welfare losses from job loss are amplified since households care more about the near-term, when earnings are low. Finally, they focus on mass lay-off events whereas we consider any job loss. It is plausible that the welfare losses associated with unconditional job loss are different to those caused by mass lay-offs.

7 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 consumption, income, and wealth surveys, as well as prices. The key assumptions we make are that preferences are intertemporally separable and, given observable characteristics, all households have common preferences and face the same prices and beliefs in each period. To calculate money-metrics and cost-of-living, we also require that some subset of households be rentiers, with negligible idiosyncratic undiversifiable income risk (e.g. risky labor income). Our approach provides a way to non-parametrically measure dynamic welfare, making it useful for studying the welfare effect of shocks that have dynamic stochastic effects on many variables that affect households and researchers do not wish to fully specify the economic environment.

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.

Argente, D. and M. Lee (2021). Cost of living inequality during the great recession. *Journal of the European Economic Association 19*(2), 913–952.

Atkin, D., B. Faber, T. Fally, and M. Gonzalez-Navarro (2024). Measuring welfare and inequality with incomplete price information. *The Quarterly Journal of Economics 139*(1), 419–475.

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. T. Burstein, and Y. Koike-Mori (2024). Measuring welfare by matching households across time. *The Quarterly Journal of Economics 139*(1), 533–573.

Basu, S., L. Pascali, F. Schiantarelli, and L. Serven (2022, 01). Productivity and the Welfare of Nations. *Journal of the European Economic Association 20*(4), 1647–1682.

Best, M. C., J. S. Cloyne, E. Ilzetzki, and H. J. Kleven (2020). Estimating the elasticity of intertemporal substitution using mortgage notches. *The Review of Economic Studies 87*(2), 656–690.

Blackorby, C., D. Primont, and R. R. Russell (1998). Separability: A survey. *Handbook of utility theory 1*, 51–92.

Blaum, J., C. Lelarge, and M. Peters (2018). The gains from input trade with heterogeneous importers. *American Economic Journal: Macroeconomics 10*(4), 77–127.

Blundell, R. W., M. Browning, and I. A. Crawford (2003). Nonparametric engel curves and revealed preference. *Econometrica 71*(1), 205–240.

Broda, C. and D. E. Weinstein (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.

Fisher, F. M. and K. Shell (1968). *Taste and quality change in the pure theory of the true cost of living index*. Value, Capital and Growth. Edinburgh University Press.

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

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 (2024). Measuring growth in consumer welfare with income-dependent preferences: Nonparametric methods and estimates for the united states. *The Quarterly Journal of Economics 139*(1), 477–532.

Jorgenson, D. W. (2018). Production and welfare: progress in economic measurement. *Journal of Economic Literature 56*(3), 867–919.

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.

Redding, S. J. and D. E. Weinstein (2020). Measuring aggregate price indices with taste shocks: Theory and evidence for ces preferences. *The Quarterly Journal of Economics 135*(1), 503–560.

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.

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

Straub, L. (2019). Consumption, savings, and the distribution of permanent income. *Unpublished manuscript, Harvard University*.

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.[33]

Capitalized wealth proxy:

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

[33]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.

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

Budget shares:

We align the seven categories of the PSID (food, housing, transportation, education, health, clothing, and recreation) with the CPI.[34] 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 Proofs

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. ∎

[34] 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.

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:

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

The first order conditions for asset holdings are

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

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

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

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 2. *If preferences are 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 3. *When preferences are time separable, the elasticity of intertemporal substitution*

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

is given by

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

Proof. We start with

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

Summing over all $n \in N$ yields

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

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,

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

■

Lemma 4. *When preferences are time separable, the following equation holds:*

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

Proof. From Lemma 2, we know that

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

Next, from homogeneity of degree one, we know that

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

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

$$
\begin{aligned}
b^P d\log b^P(\boldsymbol{q},\boldsymbol{\pi},U) &= \frac{\partial^2 \log e(\boldsymbol{q},\boldsymbol{\pi},U)}{(\partial \log P)^2}\sum_{n\in N}\frac{\partial \log P}{\partial \log q_n(s^0)}d\log q_n(s^0) \\
&\quad + \frac{\partial^2 \log e}{\partial \log F \partial \log P}\sum_{s^t|t>0}\left(\sum_{n\in N}\frac{\partial \log F}{\partial \log q_n(s^t)}d\log q_n(s^t) + \frac{\partial \log F}{\partial \pi(s^t)}d\pi(s^t)\right) \\
&= \frac{\partial^2 \log e(\boldsymbol{q},\boldsymbol{\pi},U)}{(\partial \log P)^2}\left[\begin{array}{c} \sum_{n\in N} b_n(\boldsymbol{q},U)d\log q_n(s^0)- \\ \sum_{s^t|t>0}\left(\sum_{n\in N}\frac{\partial \log F}{\partial \log q_n(s^t)}d\log q_n(s^t) + \frac{\partial \log F}{\partial \pi(s^t)}d\pi(s^t)\right)\end{array}\right].
\end{aligned} \tag{23}
$$

From Lemma 2 and Lemma 3, we can rewrite this as

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

Rearranging this gives

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

Plug this back into (23) to get the desired result. ■

Lemma 5. *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,

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

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

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 6. *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 5, 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 6. Logging both sides gives

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

$$= \log w + \int_{\tau}^{\tau_0} \left(\frac{\partial \log e(\boldsymbol{q}^*(x, u(\tau, w, \mathbf{0})), \boldsymbol{\pi}(\cdot|x), u(\tau, w, \mathbf{0}))}{\partial \log \boldsymbol{q}^*} \frac{d \log \boldsymbol{q}^*}{dx} \right.$$
$$\left. + \frac{\partial \log e(\boldsymbol{q}^*(x, u(\tau, w, \mathbf{0})), \boldsymbol{\pi}(\cdot|x), u(\tau, w, \mathbf{0}))}{\partial \log \boldsymbol{\pi}(\cdot|x)} \frac{d \log \boldsymbol{\pi}(\cdot|x)}{dx} \right) dx,$$

where the second equality uses the fundamental theorem of calculus for line integrals. Using Lemma 4, we can rewrite the last line as

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

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

■

Proof of Lemma 1. 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 $n \in N$

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

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 5. From Lemma 1, 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 2. Lemma 1 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

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

if, and only if,

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

■

C 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 (1) are a CES aggregator over dates. Hence, demand for consumption in the first period, relative to total wealth, follows CES demand:

$$\log \frac{b(0|\tau)}{b(0|\tau_0)} = (1-\sigma)\left[\log \frac{p(0|\tau)}{p(0|\tau_0)} - \log \frac{e(\boldsymbol{q}(\cdot|\tau,1))}{e(\boldsymbol{q}(\cdot|\tau_0,1))}\right]. \tag{24}$$

Rearrange this equation and combine it with (2) to arrive at (4). 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 (24), we use demand for time 0 consumption to solve for relative prices between time 0 and the future. We then substitute this expression into (3) 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 (24). Hence, we cannot simply invert the demand curve to solve for the ideal price index.

D Sensitivity Analysis for Results in Section 5

TBA