

NBER WORKING PAPER SERIES

SUFFICIENT STATISTICS FOR MEASURING FORWARD-LOOKING WELFARE

David Baqaee
Ariel Burstein
Yasutaka Koike-Mori

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

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

We thank Andy Atkeson, Adrien Auclert, Maya Eden, Valentin Haddad, Oleg Itskhoki, Chad Jones, Patrick Kehoe, Narayana Kocherlakota, John Fernald, Kieran Larkin, Hannes Malmberg, Ricardo Reis, Pierre-Olivier Weill, and Ivan Werning for comments. We acknowledge financial support from NSF grant No. 1947611 and the Alfred P. Sloan Foundation. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

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

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

ABSTRACT

We provide a method to measure 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. We find that static measures overstate cost-of-living increases for most households, particularly younger and poorer households. Our estimates can be used to study the welfare consequences of dynamic stochastic shocks that affect households along different margins and time horizons. For example, we find that involuntary job loss is associated with a 20% reduction in money-metric utility for households younger than 60 years old.

David Baqaee
Department of Economics
University of California, Los Angeles
Bunche Hall
Los Angeles, CA 90095
and CEPR
and also NBER
baqaee@econ.ucla.edu

Yasutaka Koike-Mori
University of California, Los Angeles
University Village
3360 Sawtelle Blvd # 305
Los Angeles, CA 90066
yasutaka.koike.mori@gmail.com

Ariel Burstein
Department of Economics
Bunche Hall 8365
Box 951477
UCLA
Los Angeles, CA 90095-1477
and NBER
arielb@econ.ucla.edu

1 Introduction

Measuring the welfare of households 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.¹

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.² For example, welfare may rise even though real income or real consumption falls.³

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

¹The money-metric was introduced by McKenzie (1957), but popularized by Samuelson (1974). For a textbook presentation, see Deaton and Muellbauer (1980).

²Real 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.

³Expenditures 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.

of welfare but applied empiricists use income-like measures, has long been a source of 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. 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, 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

⁴To 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.

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

(e.g. rational expectations), and financial frictions (e.g. complete markets), or first-order approximations.

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. This is critical because consumption-savings decisions are highly non-homothetic and respond strongly to permanent income in the data (see, e.g., Straub, 2019). 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).

Changes in consumption-wealth ratios are not sufficient to back out welfare for non-rentiers. 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 conditional on relative prices, then two households facing the same relative prices are on the same intertemporal indifference curve if, and only if, their budget shares 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,

⁶Our 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.

⁷This 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).

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 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 conventional CPI measures overestimate the true dynamic cost-of-living price index in our sample. Mechanically, this is because the EIS is less than one and there is an increase over time in compensated consumption-wealth ratios.⁸ Furthermore, we find more heterogeneity in the dynamic cost-of-living index than in the static cost-of-living index across both the wealth distribution and by age group. Static measures of inflation typically find that inflation is overstated for rich households and understated for poor households.⁹ In contrast, we find that dynamic inflation rates tend to be lower for poorer households. Furthermore, dynamic non-homotheticities are much more powerful than static ones in our data. However, standard errors for the dynamic price index are larger due to small sample size.

Our methodology is useful as an input for reduced-form empirical work studying the welfare effects of dynamic treatments with uncertain outcomes. Many policies and shocks have complex effects that affect households in ex-ante uncertain ways along many dimensions and at different time horizons. We provide a way to study the welfare effects of such treatments without requiring that researchers enumerate, estimate, and aggregate all the possible ways the treatment affects the household and how the household's beliefs and contingent plans change in response to that treatment.

To illustrate this approach, we study the welfare consequences of job loss in our sample. To do so, we regress our estimates of money-metric utility on job loss. We find that involuntary job loss is associated with a roughly 20% reduction in money-metric utility 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 households discount the future using market rates, or that households

⁸This does not imply that consumption-wealth ratios rise at the individual level since individual-level consumption wealth ratios are not compensated. Furthermore, at the individual level, consumption wealth ratios may also respond to lifecycle considerations associated with aging.

⁹For example, see Jaravel and Lashkari (2024).

have perfect foresight about their future path of earnings.

Related Literature. This paper is closely related to the literature on consumer price indices. Although the vast majority of papers in this literature abstract from time and risk, we use insights from this literature, originally developed for static problems, to study dynamic ones. Feenstra (1994) inverts CES demand curves to infer the value of new goods from changes in expenditures on existing goods.¹⁰ 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.¹¹ 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 using static budget shares. This means that we do not have to estimate or correct for substitution effects, as in Atkin et al. (2024).

Our paper is also related to a literature that non-parametrically measures non-homothetic price indices 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

¹⁰This approach is frequently used to infer the value of new goods or quality change in static settings, see, for example Broda and Weinstein (2006), Broda and Weinstein (2010), Blaum et al. (2018), Aghion et al. (2019), Argente and Lee (2021), and Argente et al. (2023). This approach can also be used to measure the gains from trade, i.e. the value of imported goods, as in Arkolakis et al. (2012).

¹¹This 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).

be summarized by the net-present value of technology shocks plus the initial capital stock. Jones and Klenow (2016) construct country-level welfare measures that account for mortality, leisure, and risk-aversion by fully specifying preferences and calibrating stochastic processes for the determinants of utility.

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 impose perfect foresight, complete markets, or rely on first-order approximations and certainty equivalence. Second, 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. Last, our method accounts for non-homotheticities, which only matter beyond first-order approximations. However, our approach does not unpack which changes in future returns or prices or probabilities are responsible for the changes in measured welfare.

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 constant relative risk aversion (CRRA) 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 provides some extensions of the basic framework such as allowing for secular changes in mortality risk and labor-leisure choice. Section 6 constructs a measure of dynamic welfare for households in the PSID. Section 7 uses these measures to study the dynamic welfare losses from job loss. We conclude in Section 8.

2 Complete Markets and CRRA Preferences

In this section, we consider a special case that demonstrates one of the key ideas of this paper. We make some strong assumptions: there is only one consumption good in each period, preferences are homothetic with constant relative risk aversion, and financial markets are complete. We relax these assumptions in Section 3.

Suppose there are cohorts with planning horizon T at each calendar date τ . Intertemporal preferences are represented by

$$\mathcal{U}(c, \pi) = \sum_{t=0}^T \sum_{s_t} \beta(s^t) \pi(s^t | \tau) \frac{c(s^t)^{1-1/\sigma}}{1-1/\sigma}, \quad \text{where} \quad \sigma \neq 1, \quad (1)$$

where T is the agent's planning horizon, t denotes the number of periods after the start date, s_t is the state in period t , the history of realizations up to period t are denoted by s^t , state-contingent intertemporal discounting (which could depend on life-cycle) are captured by $\beta(s^t)$, the probability of history s^t being realized, conditional on starting at date τ , is $\pi(s^t | \tau)$, consumption in period t given history s^t is $c(s^t)$, and $\sigma \neq 1$ is the EIS.

Assume that financial markets are complete. Suppose at each point in time τ , there is a cohort with preferences (1) facing the following sequence of budget constraints over their life:

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

$$p(s^t | \tau) c(s^t) + \sum_{k \in S_{t+1}} a_k(s^t) = R_{s_t}(s^{t-1} | \tau) a_{s_t}(s^{t-1}), \quad (3)$$

$$p(s^T|\tau)c(s^T) \leq R_{s_T}(s^{T-1}|\tau)a_{s_T}(s^{T-1}). \quad (4)$$

The first constraint ensures that consumption in the period zero plus purchases of Arrow securities sum to initial wealth, where S_1 is the set of possible realizations of the state in period 1. The second constraint ensures that, at each subsequent history s^t , consumption plus purchases of Arrow securities that pay off in $t + 1$ are equal to the returns from previously purchased Arrow securities. The final constraint is a no-Ponzi condition that ensures consumption in every terminal history must be less than the returns from previously purchased Arrow securities. Without loss of generality, we assume that initial wealth w includes present and discounted future income earned by households.

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

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

Our objective is to compare the choice set facing different cohorts of households, always keeping the preference parameters, like the planning horizon T and the taste shifters $\beta(s^t)$ fixed. For example, we compare the value function of 50 year olds in 2005 to the value function of 50 year olds in 2019. This comparison has to take into account not only wealth and present prices but also future prices, probabilities, 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 life-cycle.¹²

For each τ and w , the τ_0 -money-metric utility function, u , is defined implicitly via

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

The equation above defines a scalar-valued function, which we denote by $u(\tau, w|\tau_0)$. In words, $u(\tau, w|\tau_0)$ is the amount of initial wealth in τ_0 such that the household is indifferent between (τ, w) and (τ_0, u) . To compare two different choice sets, say (τ, w) and (τ', w') , we compare $u(\tau, w|\tau_0)$ and $u(\tau', w'|\tau_0)$. The choice set defined by (τ', w') is preferred to (τ, w) if, and only if, $u(\tau', w'|\tau_0)$ is higher than $u(\tau, w|\tau_0)$. Since both of these numbers are in terms of τ_0 dollars, we can calculate the rate of growth of money-metric wealth between (τ, w) and (τ', w') as $u(\tau', w'|\tau_0)/u(\tau, w|\tau_0)$. We can also calculate the change in the cost of

¹²As explained by Fisher and Shell (1968), or more recently Baqaee and Burstein (2023), 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. In this example, the preference relation ranks lotteries over consumption streams and hence welfare comparisons of lotteries are meaningful. However, the planning horizon T , discount parameters $\beta(s^t)$, and the elasticity σ are parameters of the preference relation. The preference relation does not rank its parameters, and hence, they must be held constant in any welfare comparison.

living between τ_0 and τ'_0 as $u(\tau, w|\tau'_0)/u(\tau, w|\tau_0)$. This is the change in the initial wealth needed to reach a (τ, w) -indifference curve in τ'_0 relative to τ_0 . To streamline notation, we frequently suppress τ_0 and write $u(\tau, w)$ in place of $u(\tau, w|\tau_0)$, with the understanding that τ_0 is implicitly held constant.

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

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

where $c(\cdot|\tau_0, 1)$ is the consumption plan of an agent in τ_0 with unit wealth. In words, the $u(\tau, w)$ is the proportional increase in the consumption profile of an agent with unit wealth in τ_0 required to make the agent indifferent to the optimal allocation given wealth w and starting date τ .

Combine the Euler equation,

$$\beta(s^t)\pi(s^t|\tau) \left(\frac{c(s^t|\tau, w)}{c(s^0|\tau, w)} \right)^{-\frac{1}{\sigma}} = \prod_{l=0}^t R_{s_{l+1}}^{-1}(s^l|\tau) \frac{p(s^t|\tau)}{p(s^0|\tau)},$$

with the intertemporal budget constraint,

$$p(s^0|\tau) c(s^0|\tau, w) + \sum_{t=1}^T \sum_{s^t} p(s^t|\tau) \prod_{l=0}^t R_{s_{l+1}}^{-1}(s^l|\tau) c(s^t|\tau, w) = w,$$

to solve for consumption at each state in terms of primitives. Substituting the resulting solutions into (6) and solving for the money-metric yields¹³

$$u(\tau, w) = w \left/ \left[\frac{\sum_{t=0}^T \sum_{s^t} (\beta(s^t)\pi(s^t|\tau))^\sigma p(s^t|\tau)^{1-\sigma} \prod_{l=0}^t R_{s_{l+1}}(s^l|\tau)^{\sigma-1}}{\sum_{t=0}^T \sum_{s^t} (\beta(s^t)\pi(s^t|\tau_0))^\sigma p(s^t|\tau_0)^{1-\sigma} \prod_{l=0}^t R_{s_{l+1}}(s^l|\tau_0)^{\sigma-1}} \right]^{\frac{1}{1-\sigma}} \right. \quad (7)$$

This equation shows that the money-metric can be obtained by appropriately deflating wealth. The difficulties of calculating dynamic money-metrics mentioned by Samuelson

¹³Using $c(s^t|\tau, w) = c(s^t|\tau, 1)w$, (6) can be expressed as $u(\tau, w) = w \left(\frac{\sum_{t=0}^T \sum_{s^t} \beta(s^t)\pi(s^t|\tau)c(s^t|\tau, 1)^{1-1/\sigma}}{\sum_{t=0}^T \sum_{s^t} \beta(s^t)\pi(s^t|\tau_0)c(s^t|\tau_0, 1)^{1-1/\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$. From the Euler equation we obtain $c(s^t|\tau, 1) = \beta(s^t)^\sigma \pi(s^t|\tau)^\sigma \left(\prod_{l=0}^t R_{s_{l+1}}^{-1}(s^l|\tau) \frac{p(s^t|\tau)}{p(s^0|\tau)} \right)^{-\sigma} c(s^0|\tau, 1)$, and substituting this into the intertemporal budget constraint, we can solve for $c(s^t|\tau, 1)$ for all s^t in terms of primitives. Substituting this into the expression for $u(\tau, w)$ above yields (7).

(1961) are apparent. Even in this highly stylized example, computing (7) requires knowledge of the entire path of discounting, probabilities, prices, and returns households expect over their life at both τ and τ_0 , none of which is easy to observe. Note that, if agents do not have perfect foresight about the future, then data from the future (i.e. at $\tau + 1$, $\tau + 2$, etc.) is of no use since we only observe a single sample-path ex-post (and agent's ex-ante beliefs need not coincide with ex-post realizations).

Note the dramatic simplification that comes from assuming that the problem is static. In this case, (7) simplifies to

$$u(\tau, w) = w \Big/ \frac{p(s^0|\tau)}{p(s^0|\tau_0)},$$

which is just nominal wealth in τ deflated by consumer price inflation between τ and τ_0 .

To circumvent the intractable measurement problem in (7), we take advantage of the fact that information about the future can be gleaned from households' consumption-savings choices. To that end, denote the consumption-wealth ratio by

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

where the first equality, implying independence from w , follows from homotheticity of preferences. We use the symbol B^P because B stands for budget share out of total wealth and P indicates the present as opposed to the future. Using the Euler equation and the intertemporal budget constraint, we can show that the consumption-wealth ratio satisfies

$$B^P(\tau) = \frac{p(s^0)^{1-\sigma}}{p(s^0)^{1-\sigma} + \sum_{t=1}^T \sum_{s^t} (\beta(s^t) \pi(s^t|\tau))^\sigma p(s^t|\tau)^{1-\sigma} \prod_{l=0}^t R_{s_{l+1}} (s^l|\tau)^{\sigma-1}}. \quad (8)$$

Intuitively, an increase in the price of consumption in period 0, $p(s^0)$, raises the consumption wealth ratio if, and only if, $\sigma < 1$. Similarly, a decrease in future prices, or an increase in future rates of return, or an increase in the probability of a state with cheap low prices being realized in the future will raise the consumption wealth ratio if, and only if, $\sigma < 1$.¹⁴

Crucially, the denominator in (8) carries information about the household's expectations about the future. Substituting (8) into (7) yields the following proposition.

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

¹⁴This logic builds on the approach in Feenstra (1994), which infers changes in relative prices from changes in relative expenditures given knowledge of the elasticity of substitution.

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

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

To understand Proposition 1, suppose that $\sigma < 1$ so that consumption goods in different periods are complements. In this case, if cohort τ 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 relative present 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 between the two cohorts in future prices relative to current prices.

To arrive at (9), we made some very strong assumptions. We assumed that there is only one consumption good in each period, that financial markets are complete, that household preferences are homothetic, and that the utility function exhibits constant relative risk aversion. In what follows, we relax these assumptions.

3 General Environment and Measure of Welfare

In this section, we set up the general environment. We first define preferences. We then introduce the decision problem of households, allowing for incomplete financial markets. Given preferences and constraints, we then generalize the notion of money-metric welfare. Our main theoretical results are in the next section, after we set up the environment in this section.

3.1 Time-Separable Preferences

Agents have preferences $\succeq$ over stochastic consumption streams represented by a utility function $\mathcal{U}(c, \pi)$, which need not be CRRA. The set of goods available each period is N and we denote the consumption of good n in history s^t by $c_n(s^t)$. Letting S^T be the set of all possible terminal histories implies that c has dimension $|S^T| \times |N|$.

Define the following transformation of the utility function:

$$e(q, \pi, U) = \min_c \{q \cdot c : \mathcal{U}(c, \pi) = U\}, \quad (10)$$

where q has the same dimensionality as c . We refer to e as the *shadow* intertemporal ex-

penditure function and to 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 (10) 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 (10) can be written as

$$e(q, \pi, U) = \tilde{e}\left(P(q(s^0), U), F(\{q(s^t)\}_{t>0}, \pi, U), U\right), \quad (11)$$

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

We say that preferences are time separable whenever the shadow expenditure function e is separable between present shadow prices and future probabilities and 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 (11) 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.)

In Appendix B we provide an alternative equivalent definition of time separability in terms of the direct utility function (and quantities) instead of the expenditure function (and prices).¹⁵ Condition (11) is not equivalent to separability of the direct utility function unless preferences are homothetic. Formally, when preferences are homothetic, separability of the expenditure function is equivalent to requiring that the utility function be expressible as

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

where U is increasing and homogenous of degree one in its arguments, and $\tilde{P}$ and $\tilde{F}$ are increasing and homogeneous of degree one in c .¹⁶ For example, since the preferences in the illustrative example of the previous section, (1), are homothetic and directly separable, they also satisfy our notion of time separability.

¹⁵The notion of separability we use is sometimes referred to as “quasi-separability” (see, e.g., Atkin et al., 2024 who use this terminology in a static context).

¹⁶See Theorem 4.3 in Blackorby et al. (1998).

The following is a parametric example of a utility function that satisfies Definition 1 but is non-homothetic.

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

$$U^{\frac{\sigma-1}{\sigma}} = \sum_{t=0}^T \beta^t U^{\varepsilon_t} \sum_{s_t} \pi(s^t) \left(\sum_n \omega_{nt} U^{\varepsilon_n} c_n(s^t)^{\frac{\gamma-1}{\gamma}} \right)^{\frac{\gamma}{\gamma-1} \frac{\sigma-1}{\sigma}}, \quad (12)$$

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 .¹⁷ These preferences are time separable.¹⁸

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 (10) in practice.

3.2 Decision Problem of Households

We assume that decision makers face a finite planning horizon, i.e. $T < \infty$, and index each cohort by the start date τ .¹⁹ Consumers choose their consumption decisions and portfolio of assets to maximize utility subject to a sequence of state-contingent budget constraints. The first period budget constraint requires that the sum of consumption and asset purchases equals initial wealth,

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

¹⁷Shifts in ω_{nt} over t are life-cycle effects and not taste shocks, since intertemporal preferences are stable.

¹⁸We can also consider a non-homothetic extension of Epstein and Zin (1989):

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

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

The elasticity of substitution across time, σ , can differ from that across states, ρ . If $\sigma = \rho$, we obtain (12).

¹⁹In our baseline, we assume that mortality risk does not depend on τ . See Section 5 for a discussion of how our results can be extended to account for secular trends in mortality risk.

where $p_n(s^0|\tau)$ and $c_n(s^0)$ is the price and consumption of good n in the initial period given starting point τ . There are K different asset types, which may not span the state space, and the quantity purchased is denoted by $a_k(s^0)$. The price of every asset is normalized to be one. 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

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

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. the stock of human capital that pays a wage every period). If $y(s^t|\tau) = 0$ for every s^t , we say that the household is a *rentier*.²⁰

Households face borrowing constraints

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

for some exogenous state-contingent borrowing constraint $X(s^t|\tau) \geq 0$. We impose a no-Ponzi condition that $X(s^T) = 0$ for every terminal history s^T to ensure the agent cannot end the problem in debt.

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

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

The value function ranks decision problems according to underlying preference relation.

²⁰Under complete markets, every household is a rentier since all assets can be bought and sold. In Section 5, 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 ad-hoc borrowing constraints.

3.3 Measuring Welfare and the Cost-of-Living

We measure welfare using a generalization of the money-metric that allows 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: $\{p(\cdot|\tau_0), R(\cdot|\tau_0), \pi(\cdot|\tau_0)\}$, with $p(\cdot|\tau_0) > 0$ and $R(\cdot|\tau_0) > 0$. The τ_0 -money-metric, u , associated with a decision problem (τ, w, y) is defined implicitly via

$$V(\tau, w, y) = V(\tau_0, u, 0).$$

In words, the money-metric, $u(\tau, w, y|\tau_0)$, maps the decision problem (τ, w, y) into the equivalent one-off lump-sum payment the household would need, under the baseline τ_0 , to ensure indifference. In a static deterministic world, the generalized money-metric we define coincides with the textbook money-metric (e.g., Deaton and Muellbauer, 1980).²¹

Proposition 2 (Money-metric cardinalizes utility). *The money-metric utility function represents the same preference ordering as the original value function in (16).*

We require that base prices and returns be non-zero, $p(\cdot|\tau_0) > 0$ and $R(\cdot|\tau_0) > 0$, 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). The change in the cost-of-living between two dates, τ_0 and τ'_0 , for some reference problem (τ, w, y) is

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

That is, to compare the cost-of-living between τ_0 and τ'_0 , we use the money-metric to convert $\{p(\cdot|\tau), R(\cdot|\tau), y, \pi(\cdot|\tau), X(\cdot|\tau), w\}$ into equivalent lump-sum payments in τ_0 and τ'_0 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 prices of goods and services over time.

²¹With incomplete markets, there are many different, non-equivalent, ways one could measure welfare depending. For example, we could convert each problem into a 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.

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

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

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

where $c_n^*(s^t|\mathbf{q}^*, \boldsymbol{\pi}, V(\tau, w, \mathbf{y}))$ are quantities that minimize the shadow expenditure function (10). Moreover, we can set shadow prices for goods in the first period equal to their observed prices:

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

for every $n \in N$.

In other words, if the household faced shadow prices $\mathbf{q}^*(\tau, w, \mathbf{y})$ and minimized shadow expenditures subject to a utility constraint, then the consumption plan the household would choose coincides with the optimal allocation from (16).

In a static deterministic environment, $c_n(s^t|\tau, w, \mathbf{y})$ collapses to uncompensated (Marshallian) demand for good n . On the other hand, the shadow quantity $c_n^*(s^t|\mathbf{q}^*, \boldsymbol{\pi}, V(\tau, w, \mathbf{y}))$ collapses to compensated (Hicksian) demand for good n . 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 can be household-specific.

Proposition 3 makes it possible to define a notion of *compensated* elasticity of intertemporal substitution for an agent facing the problem $(\tau, w, \mathbf{y})$. Denote spending on the present relative to shadow wealth by

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

Accordingly, the share of savings, i.e. spending on the future (superscript F), is

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

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

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

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

That is, the compensated EIS for a household facing $(\tau, w, \mathbf{y})$ 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 Obtaining the Money Metric for Rentiers

The next proposition limits attention to the subset of rentiers. Rentiers are households for whom $\mathbf{y}$ 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 $\mathbf{y}$, in period τ , denote current expenditures by

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

and the share of budget spent on good n by

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

Denote the consumption to wealth ratio by

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

The following proposition characterizes the money-metric utility of rentiers as the solution to a fixed point problem. 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, then the money-metric satisfies the following fixed point problem*

$$\log u(\tau, w, \mathbf{0}|\tau_0) = \log w - \int_{\tau_0}^{\tau} \left(\sum_{n \in N} B_n(x, w_x^*, \mathbf{0}) \frac{d \log p_n}{dx} - \frac{d \log B^P(x, w_x^*, \mathbf{0})/dx}{1 - \sigma(x, w_x^*, \mathbf{0})} \right) dx, \quad (17)$$

where w_x^* solves the equation

$$u(x, w_x^*, \mathbf{0}|\tau_0) = u(\tau, w, \mathbf{0}|\tau_0). \quad (18)$$

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

To streamline notation, we suppress the dependence of the money-metric on the base period, τ_0 , where possible. Before discussing the general case, we build intuition by considering the special case of Proposition 4 where preferences are homothetic.

Corollary 1 (Homothetic Preferences). *If preferences are homothetic, Proposition 4 implies that the money-metric is given by*

$$\log u(\tau, w, \mathbf{0}) = \log w - \underbrace{\int_{\tau_0}^{\tau} \sum_n B_n(x) \frac{d \log p_i}{dx} dx}_{\text{chained inflation}} + \underbrace{\int_{\tau_0}^{\tau} \frac{1}{1 - \sigma(x, \mathbf{0})} \frac{d \log B^P(x, \mathbf{0})}{dx} dx}_{\text{adjustment for future}}. \quad (19)$$

This is quite similar to our starting example in Proposition 1. The only differences relative to that example are that (1) the EIS need not to be constant; (2) there may be multiple goods per period and the price index for present consumption is a chain-weighted average of individual goods prices.

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

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

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 .

We now discuss the general case with non-homothetic preferences. In this case, the relative budget shares on goods in each period, the consumption-wealth ratio, and the EIS all depend on wealth. Proposition 4 shows that it is the compensated version of these objects that must be used when preferences are non-homothetic. This is ensured by equation (18), which requires that the household is kept indifferent along the path between $(\tau, w, \mathbf{0})$ and $(\tau_0, u(\tau, w, \mathbf{0}), \mathbf{0})$. That is, w_x^* is the level of wealth required at period $x \in [\tau_0, \tau]$ such that the consumer stays on the same intertemporal indifference curve. The crucial fact about Proposition 4 is that this defines a fixed point problem. 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 $u(\tau, w, \mathbf{0})$.

Proposition 4, which generalizes static results in Baqaee et al. (2024), is a fixed-point problem that depends on observables and the EIS. The observables are wealth w , budget shares B_n on goods as a function of time and wealth, changes in goods prices from period to period $d \log p_n / dt$, and 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 (17) for the generalized money-metric.

Solution Method. To apply Proposition 4, we can follow the same procedures explained in detail in Baqaee et al. (2024). We begin by guessing a candidate solution $u^0(\tau, w, \mathbf{0})$, for example, by deflating nominal wealth at each date using the chained inflation index and ignoring the role of the future. We then use this initial guess on the right-hand side of (17) to get a new guess. We then iterate on this until convergence. This procedure will always converge since the fixed point to (17) is, locally, a contraction mapping.

Boundaries. Proposition 4 can only be applied reliably inside a suitable boundary. This is because the budget shares $B(\tau, w, \mathbf{0})$ and consumption-wealth ratios $B^P(\tau, w, \mathbf{0})$ are observed only for some subset of wealth levels, say $w \in [\underline{w}_x, \bar{w}_x]$ for $x \in [\tau_0, \tau]$. This limits the range of values of w for which we can calculate the money-metric without out-of-sample extrapolation. Intuitively, if for cohort τ and wealth w , the money-metric value $u(\tau, w, \mathbf{0})$ is not in $[u(x, \underline{w}_x, \mathbf{0}), u(x, \bar{w}_x, \mathbf{0})]$ 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 (17) 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).

Justifying static price indices. There are some 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, $\sigma = \infty$. In this case, (17) simplifies to

$$\log u(\tau, w, \mathbf{0}) = \log w - \int_{\tau_0}^{\tau} \sum_{n \in N} B_n(x, w_x^*, \mathbf{0}) \frac{d \log p_n}{dx} dx. \quad (20)$$

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 . This justifies using real wealth to measure welfare.

The second case is when the EIS is zero and preferences are homothetic, in which case (17) 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 inflation between τ and τ' . 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.)

The final case we consider is when preferences are non-homothetic and the EIS is zero. In this case, (17) 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. Hence, the empirical approach taken by Baqaee et al.

(2024), which ignores time, is justifiable if the EIS is zero and all households are rentiers.

To better understand Proposition 4, we sketch the proof. (The formal proof is in Appendix A). Readers that are not interested in the proof can safely skip the rest of this section.

Proof sketch for Proposition 4. 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)

$$q^*(s^t|\tau, w, \mathbf{0}) = q^*(s^t|\tau, V(\tau, w, \mathbf{0})). \quad (21)$$

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 (21) need not hold for non-rentiers. For 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

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

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

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

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.²²

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

²²In a static, deterministic environment, this price deflator collapses to an ideal price index, also known as a Konüs (1939) price index.

of calculus, as

$$\begin{aligned} \log u(\tau, w, \mathbf{0}) = \log w + \int_{\tau_0}^{\tau} \sum_{t=0}^T \sum_{s^t} & \left(\frac{\partial \log e(q^*(s^t|x, u(\tau, w, \mathbf{0})), \pi(s^t|x), u(\tau, w, \mathbf{0}))}{\partial \log q^*(s^t|x, u(\tau, w, \mathbf{0}))} \cdot \frac{d \log q^*(s^t|x, u(\tau, w, \mathbf{0}))}{dx} \right. \\ & \left. + \frac{\partial \log e(q^*(s^t|x, u(\tau, w, \mathbf{0})), \pi(s^t|x), u(\tau, w, \mathbf{0}))}{\partial \pi(s^t|x)} \cdot \frac{d \pi(s^t|x)}{dx} \right) dx. \end{aligned} \quad (22)$$

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, summing over all possible dates and states. Equation (22) 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

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

When preferences are time separable, the complicated integrand in (22) can be rewritten as

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

leaving the details of this derivation in the appendix. Intuitively, changes in shadow prices and probabilities change the shadow expenditure function to the extent that they move the compensated future price aggregator $F(\{q(s^t)\}_{t>0}, \pi, U)$ relative to the present

aggregator $P(q(s^0), U)$. 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 (22) and manipulating leads to Proposition 4. With the addition of (18), we can treat (22) as a fixed point problem.

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. $\mathbf{y} \neq \mathbf{0}$). Fortunately, we can exploit non-homotheticity of preferences to extend $u(\tau, w, \mathbf{0})$ to non-rentiers. 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:*

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

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

Importantly, Lemma 1 implies that the budget share of each good in the present $B_n(\tau, w, \mathbf{y})$ does not directly depend on wealth w and the stream of payoffs $\mathbf{y}$.

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

Proposition 5 (Money Metric is a Function of Budget Shares and Time). *Suppose that the vector-valued function $\mathbf{b}(\mathbf{p}, V)$ is a one-to-one function of V . Then, there exists a function m mapping budget shares and time into money-metric utility values in that period:*

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

for every τ , w , and $\mathbf{y}$.

The compensated budget shares $\mathbf{b}(\mathbf{p}, V)$ are a one-to-one function of V if 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 $m(\mathbf{B}, \tau)$ mapping vectors of budget shares $\mathbf{B}$ at date τ into the equivalent lump sum wealth at date τ (i.e. $u(\tau, w, \mathbf{y}|\tau)$).

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

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

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

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 $u(\tau, w, \mathbf{y}|\tau)$ for non-rentiers given their static budget shares. Intuitively, we infer the money-metric wealth for non-rentiers using their static budget shares. Proposition 6 formalizes this idea.

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

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

where $[\underline{w}_\tau, \bar{w}_\tau]$ 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

$$m(\cdot, \tau)|_{\mathcal{B}(\tau)} \in \arg \min_{\hat{m} \in M} \int_{\underline{w}_\tau}^{\bar{w}_\tau} (w - \hat{m}(B(\tau, w, \mathbf{0}), \tau))^2 dw.$$

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, $b_i(\mathbf{p}(s^0|\tau), V(\tau, w, \mathbf{y}))$, is strictly monotone in V . Then*

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

In this simple case, if the compensated budget share of i is monotone in utility, then we can deduce that two households $(\tau, w, \mathbf{y})$ 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 $u(\tau, w, \mathbf{y})$ for every $u(\tau, w, \mathbf{y})$ inside the boundary where we can solve (17) (without extrapolation).

5 Extensions

In this section, we discuss some extensions to our basic framework. These extensions are: relaxing common prices and probabilities, allowing rentiers to earn risk-free labor income, allowing for changes in mortality, and allowing leisure.

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

However, non-rentiers' future state variables (i.e. beliefs, prices, returns, borrowing constraints, cash flow) need not be the same as those of the rentiers, nor do they need to be the same for all non-rentiers. For example, rentiers may have access to different assets or hold different beliefs about the returns on those assets than non-rentiers. To understand why, recall that by Lemma 1, 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. 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,

²³To 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. If preference shocks are independent and identically distributed across households, then averaging across households with the same observables eliminates them by the law of large numbers. For this reason, in our empirical applications, we average over households with similar characteristics. However, if the preference shocks are systematic, then this can cause biases. Note that preference stability is a typical maintained assumption 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).

$y(s^t|\tau) = y(t|\tau)$, 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 $\{1, \dots, T\}$.

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

$$\sum_{n \in N} p_n(s^0|\tau) c_n(s^0|\tau) + \sum_{k \in K} a_k(s^0|\tau) + \sum_{m=1}^T B_m(s^0|\tau) = w + y(s^0|\tau), \quad (23)$$

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

$$\sum_{n \in N} p_n(s^t|\tau) c_n(s^t|\tau) + \sum_{k \in K} a_k(s^t|\tau) = \sum_{k \in K} R_k(s^t|\tau) a_k(s^{t-1}|\tau) + B_m(s^0|\tau) R(t|\tau) + y(t|\tau). \quad (24)$$

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.

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

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

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

where $\tilde{P}$ and $\tilde{F}$ are present and future aggregators that are homogeneous of degree one in quantities (for a given U), parameters λ_P and λ_F are the probabilities of surviving in the

present and the future, β is a discount factor, and σ is the constant EIS. Upon death, the household receives a consumption-equivalent payoff of $\bar{c}$ — the lower is $\bar{c}$, the higher is the value of statistical life.²⁴ The parameter ε determines wealth effects for the value of life. All of these parameters may vary as a function of observable characteristics (principally, age).

The shadow intertemporal expenditure function associated with these preferences is

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

where P and F are aggregators that are homogeneous of degree one in prices. For this problem, the lotteries the household faces are parameterized by the stochastic process over prices and returns, $\boldsymbol{\pi}$, as well as mortality risk, $\boldsymbol{\lambda}$. While the expenditure function is time separable in shadow prices, $\mathbf{q}$, and probabilities over market outcomes, $\boldsymbol{\pi}$, it is not time separable in mortality risk $\boldsymbol{\lambda}$.

As long $\boldsymbol{\lambda}$ is constant as a function of calendar time, our results can be used. However, if $\boldsymbol{\lambda}$ changes over time, then time-separability is violated and our results cannot be used. Survival probabilities are not isomorphic to the other probabilities in the model because upon death, households do not make choices subject to a budget constraint. This causes the shadow intertemporal expenditure function to become non-separable in survival probabilities.

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

$$\Phi_P(\tau, w) = \frac{d \log e(\mathbf{p}, \boldsymbol{\pi}, \boldsymbol{\lambda}, V(\tau, w))}{d \log \lambda_0}, \quad \text{and} \quad \Phi_F(\tau, w) = \frac{d \log e(\mathbf{p}, \boldsymbol{\pi}, \boldsymbol{\lambda}, V(\tau, w))}{d \log \lambda_1},$$

where the derivatives of the shadow intertemporal expenditure function are taken at the shadow prices that rationalize consumer choices (i.e. as in Proposition 3). The marginal willingness to pay to reduce the probability of death is an important input into cost-

²⁴Some papers effectively assume that $\bar{c}^{(\sigma-1)/\sigma} = 0$, but this is arbitrary. For example, it implies that the value of life switches from positive to negative as σ goes from above one to below one. Because $\bar{c}^{(\sigma-1)/\sigma}$ determines the household's willingness to pay to reduce the probability of death, its value should instead be disciplined by the value of statistical life (e.g. as in Jones and Klenow, 2016). Our formula in equation (25) directly uses information about the value of statistical life, instead of calibrating $\bar{c}$.

benefit analysis.²⁵ If preferences are non-homothetic, then these elasticities depend on initial wealth.²⁶

The extension of Proposition 4 to this environment is that the money-metric $u(\tau, w, \mathbf{0})$ solves the following fixed point problem:

$$\begin{aligned} \log u(\tau, w, \mathbf{0}) = & \log w - \int_{\tau_0}^{\tau} \left(\sum_{n \in N} B_n(x, w_x^*) \frac{d \log p_n}{dx} - \frac{1}{1-\sigma} \frac{d \log B^P(x, w_x^*)}{dx} \right) dx \\ & - \int_{\tau_0}^{\tau} \left(\Phi_P(x, w_x^*) \frac{d \log \lambda_P(x)}{dx} + \Phi_F(x, w_x^*) \frac{d \log \lambda_F(x)}{dx} \right) dx \\ & + \int_{\tau_0}^{\tau} \frac{\sigma}{1-\sigma} (1 - B^P(x, w_x^*)) \frac{d \log \lambda_F(x)}{dx} dx, \end{aligned} \quad (25)$$

where $u(x, w_x^*, \mathbf{0}) = u(\tau, w, \mathbf{0})$ for each $x \in [\tau_0, \tau]$. The first line of (25) is identical to the one in Proposition 4. The remaining two lines adjust for changes in mortality risk. As expected, if mortality risk does not vary as a function of calendar date, $d \log \lambda_F/dx = d \log \lambda_P/dx = 0$, then (25) simplifies to the expression in Proposition 4.

The intuition for the new terms is the following. The second line of (25) adds the value of increased survival to the money-metric. This is very similar to how price changes must be accounted for: we integrate the demand curve for the value of life with respect to changes in mortality risk. The subtlety is that, as with prices, we must integrate compensated demand curves rather than Marshallian demand curves. This implies that there is a fixed point term in the second line unless preferences are homothetic. The final line of (25) accounts for the fact that, holding all else fixed, changes in the future survival probability changes the consumption wealth ratio. Since the second line is fully accounting for the welfare changes caused by changes in the future survival probability, the third line purges from $\frac{1}{1-\sigma} \frac{d \log B^P(x, w_x^*)}{dx}$ those changes caused by changes in future survival probability.

A simplification of (25) results if the utility payoff of death is zero: $\bar{c}^{\frac{\sigma-1}{\sigma}} = 0$. This is because, in this case, the marginal value of increasing survival probability in the future exactly offsets the adjustment in the consumption wealth ratio:

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

²⁵See Ashenfelter (2006) for a discussion of how such statistics can be estimated.

²⁶If the aggregators P and F are independent of U and $\varepsilon = 0$, then the value of life relative to wealth does not vary as a function of wealth.

This means that (25) simplifies to

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

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

To recap, our results can be extended to account for changing mortality provided with additional information on the value of statistical life as a function of wealth and age. In our empirical application, we abstract from these issues and treat mortality risk as constant over time. We think this is reasonable because the time series dimension of our sample is relatively short (15 years). We emphasize that this does not imply that we take mortality risk to be constant as a function of age. Variation in mortality rates due to changes in age are fully accounted for by our baseline result in Proposition 4. Rather, the assumption we make in our empirical results is that survival probabilities are constant as a function of time conditional on age.

Leisure. Our baseline framework treats hours worked as exogenous. In other words, we require that, conditional on observables like age, leisure choices do not change as a function of calendar time. For rentiers, this can be justified if their leisure choices are at a corner and equal to the time endowment or if their labor productivity is assumed to be zero.²⁷ If rentiers' leisure choices are not changing as function of calendar time, then Proposition 4 can be applied without modification.

We can allow non-rentiers to make labor-leisure choices and still use Proposition 5 to infer their welfare, if relative static budget shares only depend on utility and static prices of goods and services. This assumption rules out non-separabilities between consumption choices and leisure (e.g. we rule out complementarities between leisure and recreational travel).

For example, consider the utility function

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

²⁷The key assumption is that the number of hours rentiers work does not vary as a function of calendar time. So, if rentiers do earn labor income, as in the extension with risk-free cashflows, then to ensure that hours do not vary as a function of calendar time, we require that the number of hours they work is fixed conditional on age (e.g. a nine to five job prior to retirement).

where $\tilde{P}$ and $\tilde{F}$ are aggregators over consumption of goods and services for the present and the future, and $\tilde{H}$ is an aggregator over leisure choices. These preferences have the useful feature that Lemma 1 is still valid and static relative budget shares depend only on static relative prices and U . This means that Proposition 5 can be used without modification.

6 Illustrative Application to US Data

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

6.1 Data

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

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

The PSID includes household expenditure surveys, broken down into seven categories. A major omission is the user cost of owner-occupied housing. To remedy this, we impute equivalent owner-occupied housing costs by matching home owners in each period to renters with similar observable characteristics and spending behavior. That is, we predict rental expenditures using household characteristics and spending behavior using a regression estimated on renters in the same period. This procedure is theoretically justified by Proposition 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

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

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 C for details.²⁹

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

6.2 Constructing Wealth and 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 predicted 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 C.

Figure 1 plots total consumption expenditures relative to proxy wealth in our data. On average, households spend around 3% of their total proxy wealth in the present. The aggregate consumption to wealth ratio rose during the great recession, as wealth shrank relative to consumption, but fell during the stock market boom in the mid 2010s. Over the whole sample, the aggregate consumption-wealth ratio fell.

If the EIS is less than one, then a decrease in the compensated consumption-wealth ratio implies that the dynamic cost-of-living index is rising more quickly than traditional inflation. If we assume that all households are identical (i.e. no lifecycle considerations), markets are complete, and preferences are homothetic, then the compensated consumption-wealth ratio is equal to the aggregate consumption-wealth ratio depicted in Figure 1. This would suggest that inflation according to the dynamic cost-of-living index should be higher than inflation according to static indices. However, in the next section, we show that once we account for non-homotheticity, incomplete markets, and life-cycle dynamics, dynamic cost-of-living inflation is lower than static inflation for the majority

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

Figure 1: Aggregate consumption wealth ratio over time



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

of households.


Since Proposition 4 is only valid for rentiers, we must separate households into rentiers and non-rentiers. We say that a household is a rentier if net financial assets constitute more than 90% of proxy total wealth. We also exclude from the rentier set those households whose net assets are in the top and bottom 2.5%.³⁰ Since there are few rentiers younger than 40 years old in our sample, we do not report estimates for households below 40.

6.3 Money Metric Wealth for Rentiers

Figure 2 shows a scatterplot of log consumption wealth ratios against log wealth for rentiers in the first and last year of the sample. Both panels of Figure 2 show a strong decreasing relationship between consumption-wealth ratios and wealth. The relationship is approximately loglinear. This is consistent with the finding in Straub (2019) 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 2, is relatively understudied. However, as our estimates of money-metric wealth show, in our dataset, dynamic non-homotheticity is an order of magnitude more powerful than static non-homotheticity.

³⁰Our results are similar if we further restrict the set of potential rentiers to those whose financial assets constitute more than 95% of proxy total wealth.

Figure 2: Log consumption-wealth against log wealth for rentiers



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 non-parametrically regress the consumption-wealth ratios on log wealth, age group, and year using kernel regression (npregress in STATA). Similarly, for each good $i \in N$, we regress the budget share on a quadratic function of age, log wealth, and year using LASSO. In both cases, the smoothing parameters are chosen by cross-validation.³¹

We apply Proposition 4 to our estimated cross-sectional curves to recover money-metric utility as a function of date, wealth, and age group.³² 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 EIS, σ , equal to 0.1, the benchmark estimate from Best et al. (2020). Best et al. (2020) estimate that the EIS is relatively homogeneous in the cross-section of households, with point estimates that are uniformly between 0.05 and 0.15 across different quartiles of age and income.

Unfortunately, Best et al. (2020) do not estimate Hicksian elasticities but Marshallian

³¹We use npregress since it offers a disciplined way to trade off non-parametric fit as a function of age, wealth, and time against overfitting. For the static budget shares, we use LASSO instead of npregress, since LASSO is computationally cheaper than npregress, and makes bootstrapping much less time consuming. Although we do not report it, our point estimates of the money-metric are very similar for all age groups and wealth levels if we use npregress rather than LASSO to fit the static budget shares.

³²To 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).

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. Specifically, our results are similar if we set $\sigma = 0.2$ or $\sigma = 0.05$. Of course, the results are highly sensitive to the assumption that σ is not close to one.³³

The left column of Figure 3 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. Since there are fewer young renters, sampling uncertainty is higher for younger age groups. In all cases, the dynamic money-metric wealth is higher than wealth deflated by CPI.

The right column of Figure 3 plots the dynamic annualized inflation rate (log difference between nominal wealth in 2019 dollars and money-metric wealth in 2005 dollars). CPI inflation over this period was 2%, but the dynamic inflation rate is always below 2%. Furthermore, the dynamic inflation rates vary as a function of wealth and age highlighting the importance of non-homotheticity and lifecycle considerations.

To understand the pattern in Figure 3, we decompose the dynamic inflation rate, the deflator in equation (17), into a static and a forward-looking part. Specifically, for a household with wealth w in 2019, the change in the ideal cost-of-living index between 2005 and 2019 is:


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

where w_x^* ensures that we are using compensated consumption-wealth ratios and budget shares. The first summand is a static measure of inflation. 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.


As an example, this decomposition is shown in Figure 4 for the 60 – 69 year old group. The static inflation term is not exactly the same as aggregate CPI because it weighs

³³Even 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 3: Conversion of 2019 dollars to 2005 dollars




(a) money-metric 40 – 49 year olds




(b) annualized inflation 40 – 49 year olds




(c) money-metric 50 – 59 year olds




(d) annualized inflation 50 – 59 year olds




(e) money-metric 60 – 69 year olds




(f) annualized inflation 60 – 69 year olds




(g) money-metric 70 – 79 year olds



(h) annualized inflation 70 – 79 year olds



(i) money-metric 80 – 89 year olds



(j) annualized inflation 80 – 89 year olds

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

Figure 4: Decomposing inflation for 60-69 year olds





Figure 5: Changes in log consumption wealth ratios for 60 – 69 year olds



changes in static prices using compensated budget shares rather than aggregate budget shares. Nevertheless, the static component is very close to aggregate CPI inflation at around 2% per year for all wealth levels. The very slight downward slope reflects the non-homotheticity of static preferences, and static inflation is slightly higher for poorer households, consistent with other studies of the US, like Jaravel and Lashkari (2024), that show that the static cost-of-living index has tended to rise more quickly for poorer households. Nevertheless, the slope of the static inflation line is mild compared to the slope of the dynamic inflation measure.³⁴

The component corresponding to future inflation in Figure 4 is not zero, which means that future inflation is not equal to static inflation. For all wealth levels, future prices are expected to rise by less than static prices, explaining why the dynamic all-encompassing cost-of-living index is below the static inflation line in Figure 4. Moreover, the future inflation term exhibits more dependence on wealth than the static term.

The future component of the 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 60 – 69 year olds. Since compensated consumption-wealth ratios rose for all wealth levels, the dynamic inflation rate is lower than the static one.

The uncompensated change in the consumption wealth ratio is 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 same nominal 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.³⁵

³⁴There 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.

³⁵Note that changes in the compensated and uncompensated consumption-wealth ratios are different to changes in the consumption-wealth ratio at the individual level. The reason is that the compensated ratio holds utility constant and the uncompensated ratio holds nominal wealth constant, while at the individual

Our methodology does not identify which future prices or beliefs are responsible for the patterns in Figure 3. 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 and rich rentiers or young and old rentiers are differentially reliant on, say, returns to real estate, equities, or bonds to finance their consumption, then changes in returns will differentially affect dynamic inflation rates for these households.

6.4 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 we fit log money-metric values to static budget shares and age group indicators for the rentiers:

$$\log w_{h,\tau} = \alpha'_\tau \mathbf{X}_{h,\tau} + \text{error}_{h,\tau}, \quad (26)$$

where $\mathbf{X}$ are budget shares and age groups. We then use the predicted values from this regression to impute money-metric wealth for the non-rentiers conditional on their budget shares, age, and date.³⁶


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 all age groups. Figure 7 plots nominal log money-metric wealth against nominal capitalized (proxy)

level both are changing over time. In particular, this implies that the aggregate consumption to wealth ratio displayed in Figure 1 is not informative about changes in either the compensated or uncompensated ratio.

³⁶We also consider an additional robustness check where we exclude from the set of rentiers outliers in regression (26). That is, if a potential rentier's predicted and measured total wealth differ significantly, as measured by a Cook's distance value greater than one, then we exclude these households from the set of rentiers. The results are very similar.


Figure 6: Distribution of money-metric wealth in 2019



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

wealth for rentiers and non-rentiers. If financial markets are complete and absent measurement error, these two figures should be 45 degree lines.³⁷

Figure 7: Capitalized (proxy) wealth against money-metric wealth (in logs) for all years



7 Treatment Effect on Welfare

Many policies and shocks affect households along many margins simultaneously. For example, job training programs, changes in tax policy, changes in monetary policy, or job

³⁷See Appendix C for more details.

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

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 reports the results. The outcome in the first two columns is log nominal money-metric wealth, $\log u(\tau, w, y|\tau)$. Column (1) shows that job loss is associated with an approximately 20 log point reduction in nominal money-metric nominal wealth. This effect is statistically significant and economically large. Column (2) shows that the effects are much weaker for older workers (above 60 years old). The identifying variation in these regressions comes from differences in static budget shares between households who lost their job and those who did not, given other controls. The spending patterns of households that lose their job differ from those that did not lose their job in a way that, had they been rentiers, their total wealth would be 20% lower. For older household heads, differences in wealth associated to job loss are smaller compared to younger households. The impact of job loss on wealth is much smaller than the effect on measures of contemporaneous household income (which is around 85% lower for households that lost their job).

Columns (5) and (6) repeat these regressions but use log money-metric wealth in 2019 dollars, $\log u(\tau, w, y|2019)$, as the dependent variable instead of log nominal money-metric, $\log u(\tau, w, y|\tau)$. Since we do not estimate the 2019 money-metric for the full sample, the last two columns are for the restricted subset of households where we can infer the money-

Table 1: Log money-metric wealth and job loss

	log nominal money metric				log money metric 2019 dollars	
	(1)	(2)	(3)	(4)	(5)	(6)
Job Loss	-0.197 (0.031)	-0.218 (0.034)	-0.099 (0.031)	-0.108 (0.038)	-0.098 (0.031)	-0.107 (0.038)
Job loss $\times$ 1(age $\geq$ 60)		0.180 (0.083)		0.050 (0.080)		0.045 (0.074)
Lagged LHS	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	48,357	48,357	33,987	33,987	33,987	33,987
	Full Sample		Overlapping Support			

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

metric in 2019 dollars without extrapolation. We do not estimate $\log u(\tau, w, y|2019)$ for the full sample because, as discussed above, Proposition 4 can only be applied inside a suitable boundary (without extrapolation). The point estimates for this restricted subset are about half as large. The reason is that the sample of households for which we can calculate the money-metric wealth in 2019 dollars is, on the whole, much older. Since older households are less affected by job loss, when we restrict the sample, we find smaller point estimates. To show this, columns (3) and (4) repeat the same regression as in columns (1) and (2) but restrict the sample to be the same as the one in columns (5) and (6). We find that the point estimates are almost exactly the same. The fact that the point estimates in (3) and (4) are similar to those in (5) and (6) suggests that there is no correlation between job loss and welfare-relevant inflation rates. Hence, the point estimates for nominal and real losses are similar. If job loss is uncorrelated with welfare-relevant inflation rates for the whole sample, then we can extrapolate the point estimates in (1) and (2) to apply to money-metric wealth in 2019 dollars for the whole sample as well.

An alternative approach in the literature for estimating the dynamic consequences of job loss is due to Davis and Von Wachter (2011). They estimate the present-value of earnings losses after mass-layoff events to be around 12% of counterfactual earnings (using a 5% discount rate). 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 even with the same data it is only under very strong assumptions that our estimates should coincide with their present-value calculation.

For example, 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. Furthermore, Davis and Von Wachter (2011) estimate ex-post earnings losses whereas we estimate ex-ante welfare losses, and households' ex-ante beliefs about the consequences of job loss may be more pessimistic than the outcomes Davis and Von Wachter (2011) estimate. Additionally, we do not assume exponential discounting — if households are present-biased, then welfare losses from job loss are amplified since households care more about the near-term, when earnings are low. Finally, they focus on mass lay-off events whereas we consider any job loss. It is plausible that the welfare losses associated with unconditional job loss are different to those caused by mass lay-offs. Despite all these differences, the estimates from Davis and Von Wachter (2011) provide some quantitative context about how the magnitude of our estimates compare to the standard empirical approach.

8 Conclusion

We provide a methodology for measuring welfare and cost-of-living for households accounting 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 face negligible idiosyncratic undiversifiable income risk (e.g. have negligible risky labor income). Our approach provides a way to non-parametrically measure dynamic welfare without fully specifying a structural model. This makes it useful for reduced-form analysis of the welfare effect of shocks that have dynamic stochastic effects on many variables.

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., C.-T. Hsieh, and M. Lee (2023). Measuring the cost of living in mexico and the united states. *American Economic Journal: Macroeconomics* 15(3), 43–63.
- Argente, D. and M. Lee (2021). Cost of living inequality during the great recession. *Journal of the European Economic Association* 19(2), 913–952.
- Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012). New trade models, same old gains? *American Economic Review* 102(1), 94–130.
- Ashenfelter, O. (2006). Measuring the value of a statistical life: problems and prospects. *The Economic Journal* 116(510), C10–C23.
- Atkin, D., B. Faber, T. Fally, and M. Gonzalez-Navarro (2024). Measuring welfare and

- inequality with incomplete price information. *The Quarterly Journal of Economics* 139(1), 419–475.
- Baqae, D. R. and A. Burstein (2023). Welfare and output with income effects and taste shocks. *The Quarterly Journal of Economics* 138(2), 769–834.
- Baqae, D. R., A. T. Burstein, and Y. Koike-Mori (2024). Measuring welfare by matching households across time. *The Quarterly Journal of Economics* 139(1), 533–573.
- Basu, S., L. Pascali, F. Schiantarelli, and L. Serven (2022, 01). Productivity and the Welfare of Nations. *Journal of the European Economic Association* 20(4), 1647–1682.
- Best, M. C., J. S. Cloyne, E. Ilzetzki, and H. J. Kleven (2020). Estimating the elasticity of intertemporal substitution using mortgage notches. *The Review of Economic Studies* 87(2), 656–690.
- Blackorby, C., D. Primont, and R. R. Russell (1998). Separability: A survey. *Handbook of utility theory* 1, 51–92.
- Blaum, J., C. Lelarge, and M. Peters (2018). The gains from input trade with heterogeneous importers. *American Economic Journal: Macroeconomics* 10(4), 77–127.
- Blundell, R. W., M. Browning, and I. A. Crawford (2003). Nonparametric engel curves and revealed preference. *Econometrica* 71(1), 205–240.
- Broda, C. and D. E. Weinstein (2006). Globalization and the gains from variety. *The Quarterly journal of economics* 121(2), 541–585.
- Broda, C. and D. E. Weinstein (2010). Product creation and destruction: Evidence and price implications. *The American economic review* 100(3), 691–723.
- Catherine, S., P. Sodini, and Y. Zhang (2022). Countercyclical income risk and portfolio choices: Evidence from sweden. *Swedish House of Finance Research Paper* (20-20).
- Comin, D., D. Lashkari, and M. Mestieri (2021). Structural change with long-run income and price effects. *Econometrica* 89(1), 311–374.
- Cooper, D., K. E. Dynan, and H. Rhodenhiser (2019). Measuring household wealth in the panel study of income dynamics: The role of retirement assets.
- Cornes, R. (1992). *Duality and modern economics*. CUP Archive.
- 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.
- Jones, C. I. and P. J. Klenow (2016). Beyond gdp? welfare across countries and time. *American Economic Review* 106(9), 2426–2457.
- Jorgenson, D. W. (2018). Production and welfare: progress in economic measurement. *Journal of Economic Literature* 56(3), 867–919.
- Kimball, M. S. (1995). The quantitative analytics of the basic neomonetarist model. *Journal of Money, Credit, and Banking* 27(4).
- Konüs, A. A. (1939). The problem of the true index of the cost of living. *Econometrica: Journal of the Econometric Society*, 10–29.
- Lucas, R. E. (1987). *Models of business cycles*, Volume 26. Basil Blackwell Oxford.
- McKenzie, L. (1957). Demand theory without a utility index. *The Review of Economic Studies* 24(3), 185–189.
- Nakamura, E., J. Steinsson, and M. Liu (2016). Are chinese growth and inflation too smooth? evidence from engel curves. *American Economic Journal: Macroeconomics* 8(3), 113–44.
- Pollack, R. A. (1975). The intertemporal cost of living index. In *Annals of Economic and Social Measurement*, Volume 4, number 1, pp. 179–198. NBER.

Reis, R. (2005). A dynamic measure of inflation.

Samuelson, P. A. (1961). The evaluation of ‘social income’: Capital formation and wealth. In *The Theory of Capital: Proceedings of a Conference held by the International Economic Association*, pp. 32–57. Springer.

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

Schelling, T. C. (1968). The life you save may be your own. *Problems in public expenditure*, 127–162.

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 Proofs

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

Proof of Proposition 3. The existence of q^* follows from the separating hyperplane theorem, since the constraint set and indifference curves are both convex (the constraint set is an intersection of convex sets). Furthermore, since the solution is a convex optimization problem, the Karush-Kuhn-Tucker conditions must be satisfied. The Lagrangian for households is:

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

$$\begin{aligned}
& - \sum_{s^t=s^0}^{s^T} \lambda(s^t|\tau) \sum_{k \in K} a_k(s^t|\tau) + \sum_{s^t} \lambda(s^t|\tau) \sum_{k \in K} R_k(s^t|\tau) a_k(s^{t-1}|\tau) \\
& - \sum_{s^t} \mu(s^t|\tau) \sum_k a_k(s^t|\tau) + \sum_{s^t} \mu(s^t|\tau) X(s^t|\tau)
\end{aligned}$$

The first order conditions for asset holdings are

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

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

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

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

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

From standard duality, we know that we can also write

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

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

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

Set

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

and

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

Hence

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

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

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

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

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

■

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

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

We proceed in steps, using a series of lemmas.

Lemma 2. *If preferences are time separable, then the following holds*

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

and

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

Proof. By the envelope theorem,

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

and for $t > 0$

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

Then, we know that

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

and

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

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

Next, we show that

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

To do this, use the following equality,

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

Rearranging yields

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

■

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

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

is given by

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

Proof. We start with

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

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

Summing over all $n \in N$ yields

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

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

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

By definition,

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

■

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

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

Proof. From Lemma 2, we know that

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

Next, from homogeneity of degree one, we know that

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

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

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

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

$$\begin{aligned} b^P d \log b^P(q, \pi, U) &= \frac{\partial^2 \log e(q, \pi, U)}{(\partial \log P)^2} \sum_{n \in N} \frac{\partial \log P}{\partial \log q_n(s^0)} d \log q_n(s^0) \\ &+ \frac{\partial^2 \log e}{\partial \log F \partial \log P} \sum_{s^t | t > 0} \left(\sum_{n \in N} \frac{\partial \log F}{\partial \log q_n(s^t)} d \log q_n(s^t) + \frac{\partial \log F}{\partial \pi(s^t)} d \pi(s^t) \right) \\ &= \frac{\partial^2 \log e(q, \pi, U)}{(\partial \log P)^2} \left[\frac{\sum_{n \in N} b_n(q, U) d \log q_n(s^0) -}{\sum_{s^t | t > 0} \left(\sum_{n \in N} \frac{\partial \log F}{\partial \log q_n(s^t)} d \log q_n(s^t) + \frac{\partial \log F}{\partial \pi(s^t)} d \pi(s^t) \right)} \right]. \end{aligned} \tag{27}$$

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

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

Rearranging this gives

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

Plug this back into (27) to get the desired result. ■

Lemma 5. *The shadow prices $q^*(\tau, w, \mathbf{0})$ can be written as a function of τ and $V(\tau, w, \mathbf{0})$. That is, we can write*

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

Furthermore,

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

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

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

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

Hence

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

Lemma 6. *The following holds*

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

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

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

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

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

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

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

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

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

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

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

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

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

Using the shadow expenditure function, we can write

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

where the last line uses Lemma 6. Logging both sides gives

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

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

Using Lemma 4, we can rewrite the last line as

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

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

Proof of Lemma 1. Need to show that

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

By Proposition 3, we know that

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

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

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

Hence, we have that

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

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

Proof of Proposition 5. From Lemma 1, we know that

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

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

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

Since b is an injective function, we can write

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

Since V is monotone in wealth, we can write

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

■

Proof of Corollary 2. Lemma 1 shows that

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

Hence, if b_i is monotone in V , then

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

if, and only if,

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

■

B Alternative Definition of Time Separability

In this appendix, we provide an alternative equivalent definition of time separability that does not explicitly use the shadow intertemporal expenditure function. To do so, consider some preference relation $\succeq$ with a utility function representation $\mathcal{U}(\mathbf{c}, \boldsymbol{\pi})$. Define the distance function to be

$$D(\mathbf{c}, \boldsymbol{\pi}, U) = \max_{\lambda \in \mathbb{R}} \{\lambda : \mathcal{U}(\mathbf{c}/\lambda, \boldsymbol{\pi}) = U\}.$$

The distance function is standard in duality theory and, by construction, is homogeneous of degree one in quantities and decreasing in U (as long as $\mathbf{c}$ are goods). See Blackorby et al. (1998) or Cornes (1992) for textbook treatments.

The distance function carries all the same information as the utility function. In fact, the utility function $U(\mathbf{c}, \boldsymbol{\pi})$ can be recovered from the distance function by solving for U in the following identity

$$D(\mathbf{c}, \boldsymbol{\pi}, U) = 1.$$

The solution gives us $\mathcal{U}(\mathbf{c}, \boldsymbol{\pi})$. This “implicit” representation of the utility function is common in the macroeconomics literature though it is rarely referred to as a distance function. For example, Kimball (1995) preferences and non-homothetic CES (e.g. Comin et al., 2021) preferences are usually defined via a distance function.

Theorem 4.1 in Blackorby et al. (1998) implies that separability in the expenditure function, Definition 1, is equivalent to separability in the distance function. Specifically, Definition 1 is equivalent to requiring that there exists a utility function representation of

$\geq$ satisfying

$$D(\mathbf{c}, \boldsymbol{\pi}, U) = \tilde{D}(\tilde{P}(\mathbf{c}(s^0), U), \tilde{F}(\{\mathbf{c}(s^t)\}_{t>0}, \boldsymbol{\pi}, U), U) = 1,$$

where $\tilde{P}$ and $\tilde{F}$ are homogeneous of degree 1 in $\mathbf{c}$ and $\tilde{D}$ is homogeneous of degree 1 in its first two arguments (all holding U constant) and D is decreasing in U . An example of this is provided in (12) where $\tilde{D}$, $\tilde{P}$, and $\tilde{F}$ are nested non-homothetic CES aggregators. Specifically,

$$\tilde{D} = \frac{1}{U} \left(\tilde{P}^{\frac{\sigma-1}{\sigma}} + \tilde{F}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

where

$$\tilde{P} = U^{\epsilon_0} \left(\sum_n \omega_{n0} U^{\epsilon_n} c_n(s^0)^{\frac{\gamma-1}{\gamma}} \right)^{\frac{\gamma}{\gamma-1}},$$

and

$$\tilde{F} = \left(\sum_{t=1}^T \beta^t U^{\epsilon_t} \sum_{s_t} \pi(s_t) \left(\sum_n \omega_{nt} U^{\epsilon_n} c_n(s_t)^{\frac{\gamma-1}{\gamma}} \right)^{\frac{\gamma}{\gamma-1} \frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}.$$

C Data Construction

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

Net Assets:

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

separately in the PSID) to incorporate as much of the household's assets as possible.³⁸

Capitalized wealth proxy:

We construct a proxy for total wealth by adding the capitalized value of labor income and transfers to net assets. Define household income as labor income plus the variables recorded as social security income and other welfare income. First, we estimate the age-specific income profile for each period τ using cross-sectional data. To do this, we regress a quadratic of the age of the head of household on log income controlling for household characteristics (marital status, state of residence, race of household head, gender, and occupation) and year fixed effects. We then use this regression to predict each household's income profile as their age increases. We inflate these predictions of the household's income in the future by an estimate of expected nominal per capita GDP growth. The expected growth in nominal GDP comes from the Congressional Budget Office's real-time (contemporaneous) forecast of nominal GDP growth and the population growth rate uses realized population growth rates for the United States, assuming a constant growth after 2019. We discount these nominal income flows back to the present using a nominal rate of 6%, consisting of a 4% real rate, following Catherine et al. (2022), and a 2% expected inflation rate. We assume that income flows are zero beyond age 90.

Owner-occupied housing:

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

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

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

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

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

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

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

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

Budget shares:

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

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