Sufficient Statistics for Measuring Forward-Looking Welfare

David R. Baqaee
UCLA

Ariel Burstein
UCLA

Yasutaka Koike-Mori
UNC Chapel Hill*

July 13, 2025

Abstract

We provide a method to measure welfare in money-metric terms, accounting for future expectations. Our approach relies on two key assumptions: (1) separability of the expenditure function between present and future, and (2) existence of households without idiosyncratic, undiversifiable risk. We infer expectations about the future from observed consumption-savings choices of this subset of households. Our sufficient statistics methodology accommodates incomplete markets, lifecycle motives, non-rational expectations, non-exponential time discounting, and arbitrary functional forms. Applying the method requires price changes over time, repeated cross-sectional data on household income, balance sheets, and expenditures, and the elasticity of intertemporal substitution. Using PSID data from the United States, we find that static measures overstate cost-of-living increases, especially for 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.

*We thank Andy Atkeson, Adrien Auclert, Maya Eden, Valentin Haddad, Oleg Itskhoki, Chad Jones, Patrick Kehoe, Narayana Kocherlakota, John Fernald, Kieran Larkin, Jay Lu, Hannes Malmberg, Ricardo Reis, Pierre-Olivier Weill, and Ivan Werning for comments. We acknowledge financial support from NSF grant No. 2343691 and the Alfred P. Sloan Foundation.

1 Introduction

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

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

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

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

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

equivalent lump-sum. If financial markets are complete, then this definition coincides with the usual money-metric definition.

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

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, the empirically relevant case, then an increase in the consumption-to-wealth ratio suggests that the price of consuming in the future relative to the present must be falling. This means that forward-looking measures of welfare, that take future expected prices into account, will be higher than static measures that only compare how static prices change.[2] The basic idea, which infers the value of future prices from changes in expenditures in the present, is a dynamic version of the Feenstra (1994) approach to inferring the value of new goods in a static setting.

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. It is only changes in consumption-to-wealth ratios due to substitution effects that are informative about how relative prices are changing. This is quantitatively 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-to-wealth ratio nonparametrically, without first specifying and estimating a dynamic model, we match rentiers in the cross-section over time, extending ideas from Baqaee et al. (2024) to a dynamic setting.

Changes in consumption-wealth ratios are not sufficient to back out welfare for non-rentiers. For non-rentiers, changes in consumption-to-wealth ratios can be driven by variation in the bindingness of financial constraints or temporary fluctuations in risky

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

income. To recover money-metric utility for non-rentiers, we rely 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.[3] This is similar to how expenditures on food can be used to infer changes in welfare in static settings, as in Hamilton (2001).

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

To illustrate our method, we apply our results to the United States using the Panel Study of Income Dynamics (PSID) and price data from the Bureau of Labor Statistics. We begin by selecting a subsample of rentiers. We do this by computing a proxy for the present value of expected future labor and transfer income for each household. We say that a household is a rentier if the present value of their future labor and transfer income 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.[4] Furthermore, we find more heterogeneity in the dynamic cost-of-living index than in the static cost-of-living index across both the wealth distribution and by age group. Static measures of inflation typically find that inflation is overstated for rich households and understated for poor households.[5] In contrast, we find that dynamic inflation rates tend to be lower for poorer households. Furthermore, dynamic non-homotheticities are much more powerful than

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

[4]This 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, e.g., aging.

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

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 have perfect foresight about their future path of earnings.

Related Literature. Our paper is related to the literature on dynamic price indices and measures of welfare. Compared to the literature on static price indices, the literature on dynamic price indices is comparatively small. Nevertheless, one can trace its origins to the inception of index number theory. For example, Irving Fisher and Arthur Pigou both recognized that the ideal cost-of-living index should incorporate information about the future, though they did not offer a specific remedy. As mentioned before, this literature can be divided into two branches: papers that rely on some variation of the net-present value approach, and those that rely on fully-specified dynamic models. We discuss these two branches in turn.

The advantage of using asset prices to measure welfare is that asset prices can obviate the need to directly specify a utility function. An important early paper in the first branch is by Alchian and Klein (1973), which argues that a proper definition of the consumer price index must be based on intertemporal consumption. The paper proposes including asset prices in the CPI, since asset prices measure the present discounted value of future service flows. A related point is made by Hulten (1979) in the context of national income accounting. This paper points out that productivity shocks today drive capital accumulation in the future, and so the Solow (1957) residual understates the importance of technological change. The paper proposes using interest rates to calculate a dynamic technology residual. More recently and along similar lines, Basu et al. (2022) show that,

to a first-order, the welfare of a country's infinitely-lived representative consumer can be summarized by the net-present value of technology shocks plus the initial capital stock.[6]

At the microeconomic level, Fagereng et al. (2022) and Del Canto et al. (2023) use Taylor approximations to calculate how individual consumer welfare responds to shocks to asset prices and monetary policy, respectively. Both papers rely on forecasting the future and discounting the future to the present. Fagereng et al. (2022) estimate how various asset prices changed over time and weigh these by discounted net asset sales. Del Canto et al. (2023) use local projections to forecast how monetary shocks change goods and asset prices and weigh these price changes by discounted expected budget shares.

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

The approaches above break down how specific future change in prices or incomes affects welfare, since the money metric is calculated by summing up discounted changes in future cashflows and prices. In contrast, our approach does not automatically unpack changes in measured welfare into different channels. Instead, to understand the effect of a specific shock on welfare, we first compute money-metric welfare, and then use a regressions of changes in money-metric welfare on the shock of interest to estimate the dynamic treatment effect (subject to the usual caveats about interpreting regression coefficients).

The second branch uses an exact fully-specified model, instead of relying on approximations. For example, 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 consid-

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

ers only financial wealth, whereas Aoki and Kitahara (2010) use Epstein and Zin (1989) preferences and allow for both financial and non-financial wealth. Both papers use homothetic preferences and assume that all assets can be traded — that is, they abstract from idiosyncratic uninsurable risk and borrowing constraints. In contrast, our method accommodates uninsurable risk, borrowing constraints, and non-homothetic preferences. In a different vein, Jones and Klenow (2016) construct country-level welfare measures that account for mortality, leisure, and risk-aversion by fully specifying preferences and calibrating stochastic processes for the determinants of utility.

Although our focus is on dynamics, in terms of methods, our paper is related to the literature on static consumer price indices. We build on tools from this literature, tools originally developed for static problems, to study dynamic problems. Specifically, our method builds on three different ideas from the static literature. The first is Feenstra (1994), who inverts CES demand curves to infer the value of new goods from changes in expenditures on existing goods.[7] We extend this idea to infer the value of missing future prices relative to present prices using changes in consumption-to-wealth ratios.

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

Third, to infer welfare of financially constrained agents, we build on Hamilton (2001), who infers welfare over time by tracking budget shares on food. The basic idea is that consumers spend a smaller share of their budget on food as they get richer, so changes in the budget share of food are informative about how rich consumers are.[8] Recently, Atkin et al. (2024) provide a micro-foundation for this approach and use it to calculate welfare across the income distribution. We use a similar idea to Hamilton (2001), because we match unconstrained rentiers to constrained non-rentiers within the same period using static budget shares. However, since rentiers and non-rentiers in the same period face the same (static) relative prices, we do not need to estimate or correct for substitution effects, as in Atkin et al. (2024).

[7]This 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).

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

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

2 Stripped-Down Example

In this section, we consider simple example that demonstrates one of the key ideas of this paper. To keep things simple, 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.

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

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

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

Assume that financial markets are complete, and households at date τ face the following sequence of budget constraints:

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

$$p(s^j|\tau)c(s^j|\tau) + \sum_{k \in S_{t+1}} a_k(s^j|\tau) = R_k(s^{j-1}|\tau)a_k(s^{j-1}|\tau), \tag{3}$$

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

The price of consumption in history s^j, conditional on starting at τ, is denoted by $p(s^j|\tau)$. The holdings of and gross return on asset k in that history are denoted by $a_k(s^j|\tau)$ and

$R_k(s^j|\tau)$. The first constraint ensures that consumption in the period zero plus purchases of Arrow securities sum to initial wealth, where S_1 is the set of possible realizations of the state in period 1. The second constraint ensures that, at each subsequent history s^j, consumption plus purchases of Arrow securities that pay off in $j+1$ are equal to the returns from previously purchased Arrow securities. The final constraint is a no-Ponzi condition that ensures consumption in every terminal history must be less than the returns from previously purchased Arrow securities. 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

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

Our objective is to compare the choice set facing different cohorts of households, always keeping the preference parameters, like the planning horizon J and the taste shifters β fixed. For example, we compare the value function of 50 year olds in 2005 to the value function of 50 year olds in 2019. This comparison has to take into account not only wealth and present prices but also future prices, probabilities, and returns.[9]

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

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

The equation above defines a scalar-valued function, which we denote by $m(\tau, w|\tau_0)$. In words, $m(\tau, w|\tau_0)$ is the amount of wealth in τ_0 that makes the household indifferent to having wealth w at date τ. To compare two different choice sets, say (τ, w) and (τ', w'), we compare $m(\tau, w|\tau_0)$ and $m(\tau', w'|\tau_0)$. The choice set defined by (τ', w') is preferred to (τ, w) if, and only if, $m(\tau', w'|\tau_0)$ is higher than $m(\tau, w|\tau_0)$. That is, the money-metric function is a utility function that represents the household's preferences.

Since $m(\tau, w|\tau_0)$ has interpretable units, in terms of τ_0 dollars, we can calculate the rate of growth of money-metric wealth between (τ, w) and (τ', w') as $m(\tau', w'|\tau_0)/m(\tau, w|\tau_0)$. We can also calculate the change in the cost of living between τ_0 and τ_0' as $m(\tau, w|\tau_0')/m(\tau, w|\tau_0)$.

[9]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. As 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 J, discount parameter β, and the elasticity σ are parameters of the preference relation. The preference relation does not rank its parameters, and hence, these parameters must be held constant in any welfare comparison.

This is the change in the initial wealth needed to reach a (τ, w)-indifference curve in τ_0' relative to τ_0. To streamline notation, we frequently suppress dependence on the base period, τ_0, and write $m(\tau, w)$ in place of $m(\tau, w|\tau_0)$, with the understanding that τ_0 is implicitly held constant.[10]

Since financial markets are complete, by repeated substitution, we can combine all budget constraints into a single intertemporal budget constraint:

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

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

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

$$P(\tau) = \left[\frac{\sum_{j=0}^{J}\sum_{s_j}\left(\beta^j\pi(s^j|\tau)\right)^\sigma p\left(s^j|\tau\right)^{1-\sigma}\prod_{l=0}^{j} R_{s_{l+1}}\left(s^l|\tau\right)^{\sigma-1}}{\sum_{j=0}^{J}\sum_{s_j}\left(\beta^j\pi(s^j|\tau_0)\right)^\sigma p\left(s^j|\tau_0\right)^{1-\sigma}\prod_{l=0}^{j} R_{s_{l+1}}\left(s^l|\tau_0\right)^{\sigma-1}}\right]^{\frac{1}{1-\sigma}},$$

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

From the analogy with the isomorphic static problem, we know that money metric utility, $m(\tau, w)$, can be computed by deflating nominal wealth in τ by the inflation in $P(\tau)$:

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

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

Note the dramatic simplification that comes from assuming that the problem is static.

[10] 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

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

where $\boldsymbol{c}(\cdot|\tau_0, 1)$ is the consumption plan of an agent in τ_0 with unit wealth. In words, money-metric utility $u(\tau, w)$ 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.

In this case, (7) simplifies to

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

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

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

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

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

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

To understand the intuition, consider the realistic case where the intertemporal elasticity of substitution is less than one, $\sigma < 1$. Then, the household consumes more out of their wealth in τ relative to τ_0 if, and only if, the price of consumption in the present relative to the overall price index rises. This is because consumption in the present and the future are gross complements, so an increase in the price of consumption in the present relative to the future raises spending in the present.

We can rewrite this expression to solve for the change in the intertemporal price index:

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

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

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

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

To understand Proposition 1, suppose that $\sigma < 1$. If consumers at date τ consume a larger fraction of their wealth than consumers at date τ_0, then this indicates that the price of consuming in the future is lower than consuming in the present for cohort τ than for cohort τ_0. Hence, the overall price index for households in τ is relatively cheaper than τ_0 compared to a comparison of static inflation rates $p(s^0|\tau)/p(s^0|\tau_0)$.[11]

The advantage of Proposition 1 is that it can be implemented without direct observation of households' expectations about the future. Of course, to arrive at (10), we made 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, discounting is exponential, that the utility function exhibits constant relative risk aversion, and so on. 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 and introduce the decision problem of households. Given preferences and constraints, we generalize the definition 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 continuous, non-satiated, preferences $\succeq$ over consumption streams, $\boldsymbol{c}$, with beliefs, $\boldsymbol{\pi}$. The set of goods available each period is N and we denote the consumption of good n in history s^j by $c_n(s^j)$. Letting S^J be the set of all possible terminal histories implies that $\boldsymbol{c}$ has dimension $|S^J| \times |N|$.

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

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

where D is homogeneous of degree one in $\boldsymbol{c}$.[12]

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

[11] This logic builds on the approach in Feenstra (1994), which infers changes in the ideal price index due to new goods from changes in expenditures on continuing goods given knowledge of the elasticity of substitution.

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

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

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

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

Definition 1 is equivalent to imposing a certain type of separability on the compensated demand curves generated by $\succeq$. Specifically, Definition 1 is equivalent to assuming that spending on i relative to j in a given block (the present or the future bundle) is only a function of relative prices in that block and utility. Hence, the only way future prices affect relative budget shares in the present is through wealth effects (i.e. by changing U). Conversely, relative spending shares in the future (across dates, states, and goods) depend on present prices only through wealth effects. [13] If preferences are homothetic, then we can drop U from the right-hand side of (12).

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

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

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

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

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

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

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

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

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

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

where budget shares on different commodities within a period can vary as a function of utility. The parameter γ captures the compensated elasticity of substitution between goods within a state and period.

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

3.2 Decision Problem of Households

Consumers face a finite planning horizon $J < \infty$ and are indexed by the start date τ.[14] 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,

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

where $p_n(s^0|\tau)$ and $c_n(s^0)$ is the price and consumption of good n in the initial period given starting point τ. There are K different asset types, which may not span the state space, and the quantity 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^j, the agent faces the budget constraint

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

where $R_k(s^j|\tau)$ is the return of asset k in history s^j and $y(s^j|\tau)$ is an exogenous payoff. We think of $y(s^j|\tau)$ as exogenous labor income. We say that the household is a *rentier* if labor

[14] In 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.

income, $y(s^j|\tau)$, is zero in every state.[15]

Consumers face borrowing constraints

$$\sum_k a_k(s^j) \geq -X(s^j|\tau) \tag{15}$$

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

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

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

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

3.3 Measuring Welfare and the Cost-of-Living

We measure welfare using a generalization of the money-metric.

Definition 2 (Dynamic Money Metric)**.** Consider a reference period τ_0, with reference prices, returns, and probabilities about the future: $\{\boldsymbol{p}(\cdot|\tau_0), \boldsymbol{R}(\cdot|\tau_0), \boldsymbol{\pi}(\cdot|\tau_0)\}$, with $\boldsymbol{p}(\cdot|\tau_0) > 0$ and $\boldsymbol{R}(\cdot|\tau_0) > 0$. The τ_0-*money-metric*, m, associated with a decision problem $(\tau, w, \boldsymbol{y})$ is defined implicitly via

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

In words, the money-metric, $m(\tau, w, \boldsymbol{y}|\tau_0)$, maps the decision problem $(\tau, w, \boldsymbol{y})$ into the equivalent one-off lump-sum payment the consumer would need, under the baseline τ_0, to ensure indifference. That is, what is the rentier-equivalent wealth in τ_0 that would make the household indifferent. In a static deterministic world, the generalized money-metric we define coincides with the textbook money-metric (e.g., Deaton and Muellbauer, 1980).

[15] In Section 5, we discuss how to extend the model to allow for endogeneous labor-leisure choice. In that section, we also provide conditions under which consumers earning risk-free labor income can be treated as-if they are rentiers.

We extend this textbook definition to allow for market incompleteness and changes in probabilities, which do not have counterparts in static consumer theory.

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

Proposition 2 (Money-metric cardinalizes utility)**.** *The money-metric utility function represents the same preference ordering as the original value function in* (16)*, as long as prices* $\boldsymbol{p}(\cdot|\tau_0)$ *and returns* $\boldsymbol{R}(\cdot|\tau_0)$ *are positive.*

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

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

Definition 3 (Dynamic Cost-of-Living)**.** The change in the cost-of-living between two dates, τ_0 and τ_0', for some reference problem $(\tau, w, \boldsymbol{y})$ is

$$\frac{m\left(\tau, w, \boldsymbol{y}|\tau_0'\right)}{m\left(\tau, w, \boldsymbol{y}|\tau_0\right)}.$$

That is, to compare the cost-of-living between τ_0 and τ_0', we use the money-metric to convert $\{\boldsymbol{p}(\cdot|\tau), \boldsymbol{R}(\cdot|\tau), \boldsymbol{y}, \boldsymbol{\pi}(\cdot|\tau), \boldsymbol{X}(\cdot|\tau), w\}$ 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.

4 Main Results

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

4.1 Obtaining the Money Metric for Rentiers

The next proposition limits attention to the subset of rentiers. Rentiers are consumers for whom risky labor income, $\boldsymbol{y}$, is identically equal to $\mathbf{0}$. These consumers are denoted by $(\tau, w, \mathbf{0})$.

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

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

and the share of budget spent on good n by

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

Denote the consumption to wealth ratio by

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

Denote the compensated elasticity of intertemporal substitution (EIS) for a household facing problem $(\tau, w, \boldsymbol{y})$ by $\sigma(\tau, w, \boldsymbol{y})$.[16]

To help build intuition, we considering a series of special cases that build up to the general result. To streamline notation, we suppress the dependence of the money-metric on the base period, writing $m(\tau, w, \boldsymbol{y})$ where possible.

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

Proposition 3. *Suppose preferences are homothetic, the EIS is constant, and there is only one consumption good per period. Then, for rentiers:*

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

This is quite similar to our starting example in Proposition 1. The only differences relative to that example are that (1) financial markets may not be complete (2) preferences over the future can be more general (e.g. the discount factor need not be the same in every period and the risk aversion need not be controlled by the inverse of the EIS).

Next, we relax the assumption of homotheticity in Proposition 3

Proposition 4. *Suppose the EIS is constant and there is only one consumption good per period.*

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

Then, for rentiers:

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

When preferences are non-homothetic, the consumption-to-wealth ratio can change for two reasons: (1) because the price of consumption in the present relative to consumption in the future changed, or (2) because the household in τ is differently wealthy to the one in τ_0 (i.e. is on a different indifference curve). To construct the money-metric, we must use changes in the consumption-to-wealth ratio only due to substitution effects between the present and the future, not wealth effects. To accomplish this, Equation (17) uses the consumption wealth ratio, $B^P(\tau_0, m(\tau, w, \mathbf{0}))$, of a consumer in τ_0 who is on the same indifference curve as $(\tau, w, \mathbf{0})$. For every τ and w, Equation (17) implicitly pins down $m(\tau, w, \mathbf{0})$.

The following proposition is the general case, which nests Proposition 3 and Proposition 4, characterizes the money-metric utility of rentiers as the solution to a fixed point problem.[17]

Proposition 5 (Money Metric for Rentiers)**.** *If preferences are time separable and $\sigma(x, w, \mathbf{0}) \neq 1$ almost everywhere for $x \in [\tau_0, \tau]$, then the money-metric satisfies the following fixed point problem*

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

where w_x^* solves the equation

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

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

Relative to Proposition 4, we now allow multiple goods per period, and we allow the EIS to vary. Hence, in this case, the relative budget shares on goods in each period, B_n, the consumption-wealth ratio, B^P, and the EIS, σ, all depend on wealth. Proposition 5 shows that it is the compensated version of these objects that must be used when preferences are non-homothetic. This is ensured by equation (19), which requires that the household is kept indifferent along the integration path between $(\tau, w, \mathbf{0})$ and $(\tau_0, m(\tau, w, \mathbf{0}), \mathbf{0})$. That is, w_x^* is the level of wealth required at period $x \in [\tau_0, \tau]$ such that the consumer stays on

[17] We impose the following technical assumption. For every state s^j, assume prices $p_n(s^j|\tau)$, asset returns $R_k(s^j|\tau)$, and probabilities $\pi(s^j|\tau)$ are absolutely continuous functions of calendar time τ.

the same intertemporal indifference curve. The crucial fact about Proposition 5 is that this defines a fixed point problem. To arrive at $m(\tau, w, \mathbf{0})$, we need to integrate compensated budget shares and compensated changes in savings rate. On the other hand, to perform the necessary compensation, we need to know $m(\tau, w, \mathbf{0})$.

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

Solution Method. To apply Proposition 5, we can follow the same procedures explained in detail in Baqaee et al. (2024). We begin by guessing a candidate solution $m^0(\tau, w, \mathbf{0})$, for example, by deflating nominal wealth at each date using the chained inflation index and ignoring the role of the future. We then use this initial guess on the right-hand side of (18) to get a new guess. We then iterate on this until convergence. This procedure will always converge since the fixed point to (18) is, locally, a contraction mapping.

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

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

4.2 Non-Rentiers

A challenge for the applicability of Proposition 5 is that, in practice, many households in the sample are non-rentiers (i.e. $\boldsymbol{y} \neq 0$). Fortunately, we can exploit non-homotheticity of preferences to extend $m(\tau, w, \mathbf{0})$ to non-rentiers. To do so, observe the following.

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

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

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

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

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

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

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

for every τ, w, and $\boldsymbol{y}$.

The compensated budget shares $\boldsymbol{b}(\boldsymbol{p}, V)$ are a one-to-one function of V if no two distinct values of V result in the same vector of budget shares. Notably, this rules out homothetic preferences, since once we fix time τ, then the budget shares are constant for every value of V. In words, Proposition 6 implies that, if budget shares are one-to-one with V, then holding time τ fixed, there exists a function $\tilde{m}(\boldsymbol{B}, \tau)$ mapping vectors of budget shares $\boldsymbol{B}$ at date τ into the equivalent lump sum wealth at date τ (i.e. $m(\tau, w, \boldsymbol{y}|\tau)$).

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

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

problem

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

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

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

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

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

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

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

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

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

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

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

4.3 Monte Carlo Test

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

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

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

Consumers face sequential budget constraints of the form:

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

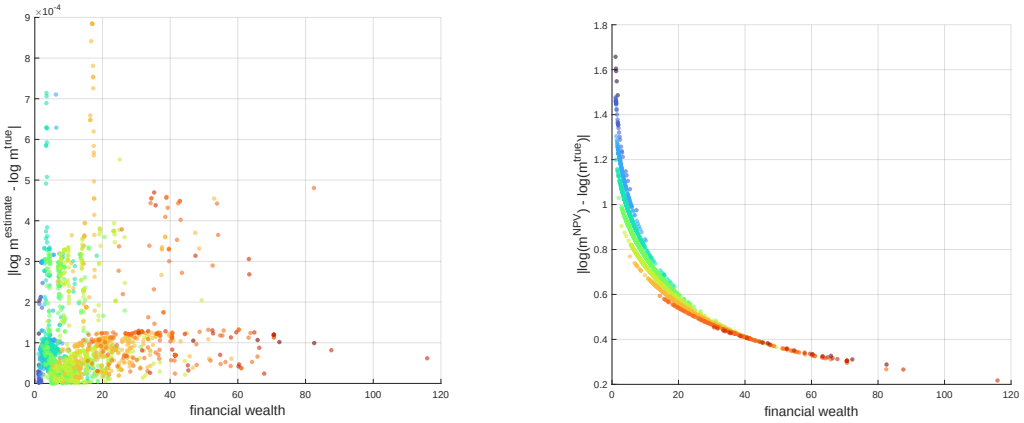
where R is the risk-free rate and y_j is an individual-specific income process. Consumers also face an exogenous borrowing constraint of $-X$.

We calibrate the parameters, simulate a finite population, and then apply our methodology. For simplicity, we assume away aggregate risk and consider a scenario where prices and returns are constant through time. Our simulated sample has two sub-populations, a group of rentiers who earn no labor income, and a group of non-rentiers whose labor income process follows an AR(1) process in logs, with persistence parameter 0.975 and standard deviation 0.15, to match estimates of the quarterly persistence and cross-sectional standard deviation of the persistent component of log income. We set the exogenous borrowing limit to -5, so households can borrow at most 5 times the average quarterly income. The annual risk free rate is 3% (so $R = 1.075$). The EIS is set to 0.1 following the estimates of Best et al. (2020), the within-period elasticity γ is equal to 0.25, following Comin et al. (2021). Our calibration has three goods within each period, and we calibrate ϵ_n to match the slope of Engel curves in the Panel Study of Income Dynamics (PSID) across three aggregated consumption goods. Finally we set the discount factor $\beta = 0.96$.

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

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

If financial markets were complete, $m^{NPV}(\tau, w, \mathbf{y}) = m(\tau, w, \mathbf{y})$. However, since that is not the case, the NPV approach significantly overstates money-metric utility for non-rentiers, especially for poorer and more financially constrained rentiers exposed to more idiosyncratic risk.



(a) Our method (b) NPV

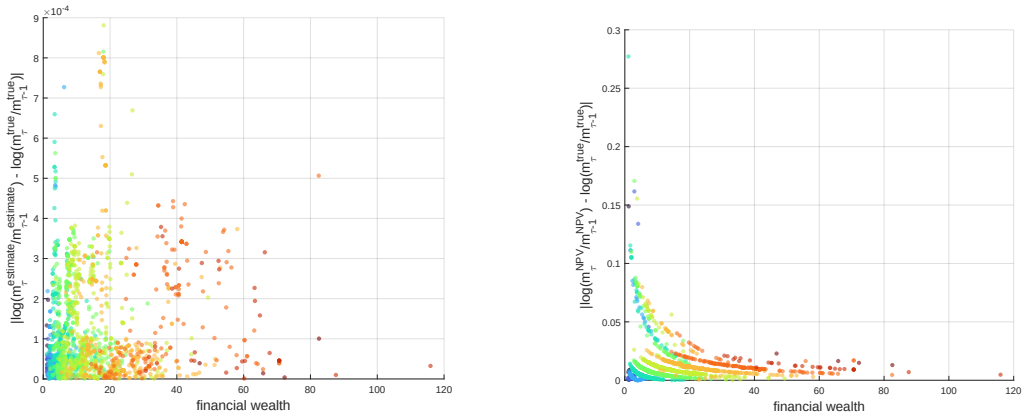
Figure 1: Comparison of errors in estimating money metric utility

Note: Colors indicate idiosyncratic income draws.

Figure 2 displays the error in the log change in money metric utility due to idiosyncratic income shocks both for our method and the NPV approach. Our method continues to perform very well. The errors, in changes, are smaller for the NPV approach but can still be large for certain households, especially those that experience large income shocks and are financially constrained.

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, allowing leisure, allowing time-inconsistencies, and considering a popular functional form of preferences, due to Comin et al. (2021), that do not satisfy our notion of time-separability.



(a) Our method (b) NPV

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

Note: Colors indicate idiosyncratic income draws.

5.1 Relaxing Common Prices and Probabilities

We assume that, conditional on observables (like age, gender, location) households have the same preference relation within and across cohorts.[18] 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

[18]To see what can go wrong, suppose that the preference relation of cohorts is changing over time. Then, changes in consumption-wealth ratios or budget shares may be due to changes in preference parameters (e.g. discount factor changes), and not due to changes in welfare-relevant variables like prices, returns, or probabilities. 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).

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.

5.2 Rentiers that Earn Risk-free Cash Flows

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

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

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

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

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

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

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

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

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

5.3 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 exogenous changes in mortality risk, consider the utility function

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

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

For this problem, the lotteries the household faces are parameterized by the stochastic process over prices and returns, $\boldsymbol{\pi}$, as well as mortality risk, $\boldsymbol{\lambda}$. While the expenditure function is time separable in shadow prices, $\boldsymbol{q}$, 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 households cannot alter the pay-off upon death by saving.

To extend our results to this setting, we again index decision problems by calendar date, τ, which now also indexes survival probabilities, λ_P and λ_F. Following Schelling (1968), denote the value of statistical life for cohort τ to be the marginal willingness to pay for increasing survival probabilities in the present and the future relative to initial wealth w. Formally, note that the money-metric utility function now depends on survival probabilities: $m(\tau, w, \boldsymbol{y}, \boldsymbol{\lambda})$. Define the marginal willingness to pay for increasing survival

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

probabilities by

$$\Phi_P(\tau, w) = \frac{d \log m(\tau, w, \boldsymbol{y}, \boldsymbol{\lambda})}{d \log \lambda_0}, \qquad \text{and} \qquad \Phi_F(\tau, w) = \frac{d \log m(\tau, w, \boldsymbol{y}, \boldsymbol{\lambda})}{d \log \lambda_1},$$

The marginal willingness to pay to reduce the probability of death is an important input into cost-benefit analysis of policies that affect longevity.[20] If preferences are non-homothetic, then these elasticities depend on initial wealth.[21]

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

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

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

The intuition for the new terms is the following. The second line of (22) adds the value of increased survival to the money-metric. This is very similar to how price changes must accounted for: we integrate the demand curve for the value of life with respect to changes in mortality risk. The subtlety is that, as with prices, we must integrate compensated demand curves rather than Marshallian demand curves. This implies that there is a fixed point term in the second line unless preferences are homothetic. The final line of (22) accounts for the fact that, holding all else fixed, changes in the future survival probability 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 (22) results if the utility payoff of death is zero: $\bar{c}^{\frac{\sigma-1}{\sigma}} = 0$. This is because, in this case, the marginal value of increasing survival probability in the future

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

[21] If 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.

exactly offsets the adjustment in the consumption wealth ratio:

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

This means that (22) 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 5. Rather, the assumption we make in our empirical results is that survival probabilities are constant as a function of time conditional on age.

5.4 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.[22] If rentiers' leisure choices are not changing as function of calendar time, then Proposition 5 can be applied without modification.

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

[22] The 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).

choices and leisure (e.g. we rule out complementarities between leisure and recreational travel).

For example, consider the utility function

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

where $\tilde{P}$ and $\tilde{F}$ are aggregators over consumption of goods and services for the present and the future, and $\tilde{H}$ is an aggregator over leisure choices. These preferences have the useful feature that Lemma 1 is still valid and static relative budget shares depend only on static relative prices and U. This means that Proposition 6 can be used without modification.

5.5 Time-Inconsistent Preferences

In forming the consumers' value function, we assume that the consumer makes decisions as-if the consumption plan chosen in the first period is followed in subsequent periods. However, if preferences are not time-consistent, and the household cannot commit, then the consumer's future selves would not stick to time-0's optimal consumption plan. In the appendix, we prove that, as long as time-separability holds, then our results can be applied without change, even if household behavior is time inconsistent. Specifically, we require that spending on any commodity k relative to any other commodity l in the present only depend on relative prices in the present and overall utility. Conversely, we require that spending on any commodity k relative to any other commodity l in the future, chosen by future selves, also only depends on relative prices in the future and overall utility. Under this assumption, our methodology applies without change.

5.6 Non-Time-Separable Non-Homotheticities

The last extension we consider is a case where time-separability is violated. Specifically, we consider the following class of preferences, due to Comin et al. (2021):

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

These preferences are superficially quite similar to the ones we used in the Monte Carlo exercise in Section 4.3, but there is an important difference. Under these preferences, non-homotheticities within a period are determined by C_j rather than U. This makes these preferences not time-separable according to Definition 1. Nevertheless, we show in

this section, and in the accompanying appendix, that our method continues to perform well even though the assumptions under which we derived our results are violated. Specifically, we show that when the EIS is close to zero, which is empirically realistic, then the errors in our procedure are also close to zero. We then verify this quantitatively using a Monte Carlo exercise.

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

Theoretical results. We start with a preliminary result characterizing the EIS for this class of preferences. Recall that the (compensated) EIS is defined via the change in consumption relative to savings induced by a uniform change in all prices in the present, holding the household on the same intertemporal indifference curve.

Lemma 2. *The EIS for the preferences above is*

$$\sigma(\tau, w, \boldsymbol{y}) = \left[(1-\gamma) \frac{Var_{B(s^0)}(\epsilon_n)}{\mathbb{E}_{B(s^0)}[\epsilon_n]^2} + 1 - \frac{\left[1 - \frac{1}{\theta}\right]}{\mathbb{E}_{B(s^0)}[\epsilon_n]} \right]^{-1}.$$

where the variance and expectation use period 0 *budget shares, denoted* $B(s^0)$*, as weights. Since* $B(s^0)$ *vary as a function of* $(\tau, w, \boldsymbol{y})$*, the EIS also varies.*

The EIS is not, in general, equal to the usual parameter θ because, under these preferences, the curvature induced by the ϵ_n's interacts with intertemporal choices. Note that, in the special case where $\epsilon_n = 1$ for every n, we recover the usual homothetic preferences and the EIS, $\sigma(\tau, w, \boldsymbol{y})$, is identically equal to θ.

Now, we show how our results are altered for rentiers. For simplicity, we assume complete markets, though the result can be extended to allow for incomplete markets.

Proposition 8. *If preferences are time separable and* $\sigma(x, w, \mathbf{0}) \neq 1$ *almost everywhere for* $x \in [\tau_0, \tau]$*, then the money-metric satisfies the following fixed point problem*

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

where w_x^* *solves the equation* $m(x, w_x^*, \mathbf{0}) = m(\tau, w, \mathbf{0})$ *for each* $x \in [\tau_0, \tau]$*. The boundary condition is that* $m(\tau_0, w, \mathbf{0}) = w$*. The error term is*

$$error = \int_{\tau_0}^{\tau} \frac{\sigma(x, w_x^*, \mathbf{0})}{\sigma(x, w_x^*, \mathbf{0}) - 1} \left[Cov_{B(s^0)}\left(\frac{\epsilon_i}{\mathbb{E}_{B(s^0)}[\epsilon_i]}, \frac{d \log p_n(s^0)}{dx}\right) - \sum_{j=0}^{J} \frac{E(s^j)/R(s^j)}{\sum_{j'} E(s^{j'})/R(s^j)} Cov_{B(s^j)}\left(\frac{\epsilon_n}{\mathbb{E}_{B(s^j)}[\epsilon_n]}, \frac{d \log p_n(s^j)}{dx}\right) \right] dx.$$

Proposition 8 is the same as Proposition 5 except that there is an error term. The error term depends on the difference between a covariance in the present and the net-present value of future covariances. The covariance is between the slope of Engel curves, disciplined by $\epsilon_n/\mathbb{E}_B(\epsilon_n)$, and the change in prices, given by $d \log p_n(s^j)/dx$. The difference between these covariances is multiplied by $\sigma/(\sigma - 1)$, where σ is the EIS.

Three comments are in order. First, as the EIS tends to zero, this error term disappears and Proposition 5 applies without change. In our Monte Carlo exercises, where $\sigma \approx 0.1$, we find that these error terms are very small even for large shocks.[23] Second, for shocks that move intra-temporal prices together, like interest rate shocks, these covariances are automatically zero, and hence the error is zero regardless of the value of σ. Lastly, even if price shocks are correlated with the slope of Engel curves, as long as the contemporaneous and future correlation are similar, the error term is likely to be small.

Just as Proposition 5 applies to rentiers with Comin et al. (2021) preferences as the EIS tends to zero, the same is true with Proposition 7, which applies to non-rentiers with Comin et al. (2021) preferences as the EIS tends to zero.

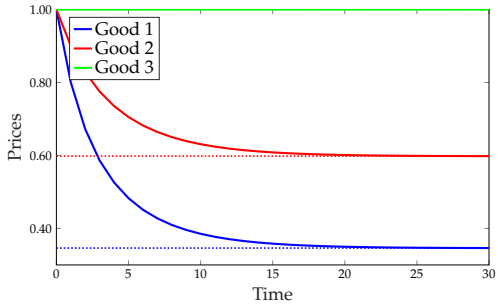
Proposition 9. *Proposition 7 applies to non-rentiers with preferences defined in Equation* (23) *when* $\sigma = 0$ *(which happens if, for example,* $\theta \to 0$*).*

We leave the proofs and formal discussions in the appendix.

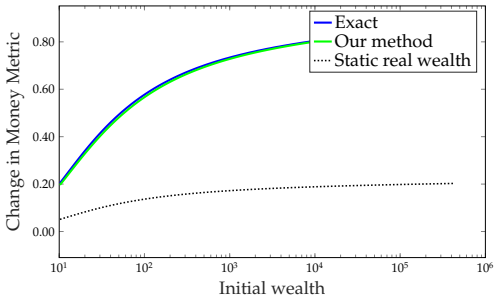
Monte Carlo Simulations. We calibrate the preference parameters in (8), simulate a finite population, and then apply our methodology. As in Section 4.3, our simulated sample has two sub-samples: a group of rentiers who face no idiosyncratic risk, and a group of non-rentiers that earn risky labor income. We leave the details of the calibration in the appendix and simply report our results for both rentiers and non-rentiers. The most important parameter is the EIS, which we calibrate to be around 0.1 (note that, by Lemma 2, the EIS is not a constant and varies slightly in the cross-section).

Figure 3 shows the performance of our method against the exact answer and changes in static real wealth in response to a large price shock. The shock is plotted in the left-panel

[23] Note that $\sigma(\tau, w, \boldsymbol{y}) \to 0$ whenever $\theta \to 0$. That is, even though the EIS is not equal to θ everywhere, the two coincide in the special case when θ is driven to zero.



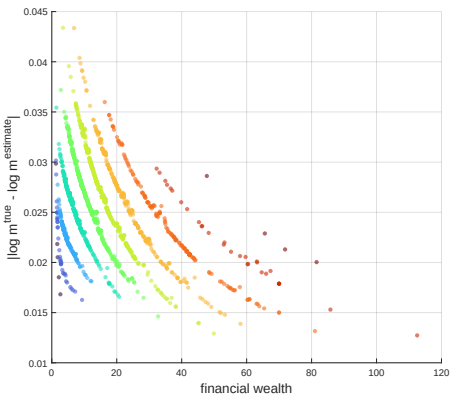
(a) Price shock



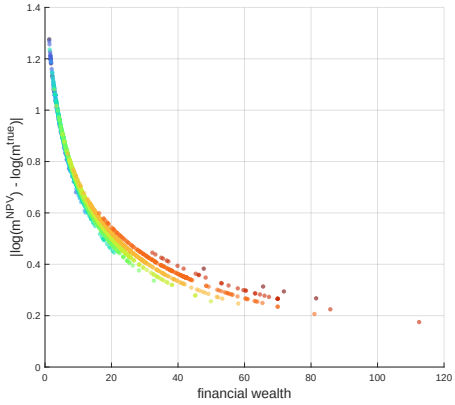
(b) Terminal money metric change

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

and the change in the money metric in logs is plotted in the right panel. As the figure shows, even though the price shocks are correlated with the slope of the Engel curves, the errors are extremely small and almost imperceptible even for very large shocks.



(a) Our method

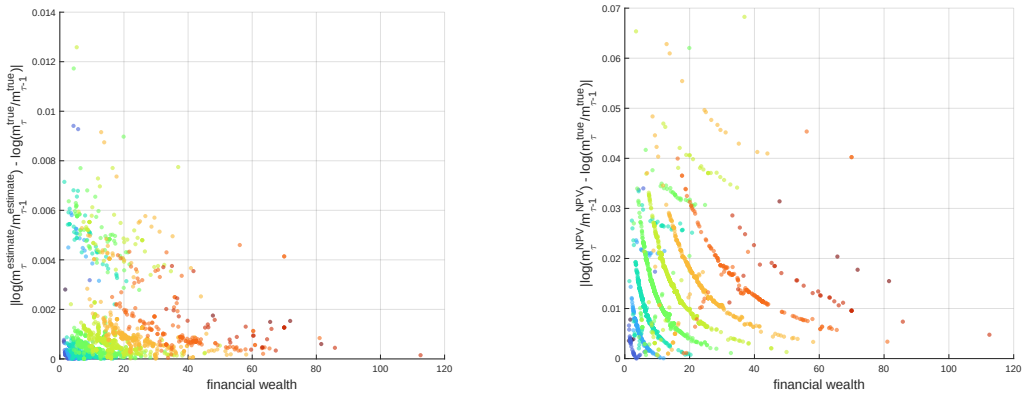


(b) NPV

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

Note: Colors indicate idiosyncratic income draws.

Figures 4 and 5 plot the error both in the level of the money metric and its change in response to idiosyncratic labor income shocks according to our method. While the errors are larger than they are when time separability holds, as in Figure 1 and 2, they are nevertheless fairly small, especially in changes. It is worth noting that the shocks experienced by the non-rentiers are quite large, and a sizeable portion of households receive income shocks that move the exact money-metric utility function by more than 20 log points. For comparison, we include the change in the net present value of total wealth, discounted using the risk-free rate, in the right panel. In both levels and changes,



(a) Our method (b) NPV

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

Note: Colors indicate idiosyncratic income draws.

the errors are much larger.

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.[24]

[24]Ideally, with enough observations, we could treat each age separately. Similarly, with more data, we

The PSID includes household expenditure surveys, broken down into seven categories. A major omission is the user cost of owner-occupied housing. To remedy this, we impute equivalent owner-occupied housing costs by matching home owners in each period to renters with similar observable characteristics and spending behavior. That is, we predict rental expenditures using household characteristics and spending behavior using a regression estimated on renters in the same period. This procedure is theoretically justified by Proposition 6. In the final year of our sample, 2019, the PSID asked home owners to report the rental value of their property if they were to rent it out. We use the answers to this question, in 2019, to validate our imputation procedure. When we regress surveyed housing costs on our imputed measure of housing costs (both relative to current expenditures), we find a coefficient of 1.03 and an R^2 value is 0.59. See Appendix C for details.[25]

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 5, we need to observe a sample of rentier households. To classify rentiers, we first estimate a proxy measure of total wealth for all households in the sample. Our proxy for total wealth is the sum of financial wealth (net asset value including home equity and defined contribution pension savings) and the present discounted value of predicted labor and transfer income. 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 6 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

could split the sample along other observed characteristics that may influence preferences, like gender, household size, location, etc.

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

Figure 6: 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.

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 6. 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 inflation is lower than static inflation for the majority of households.

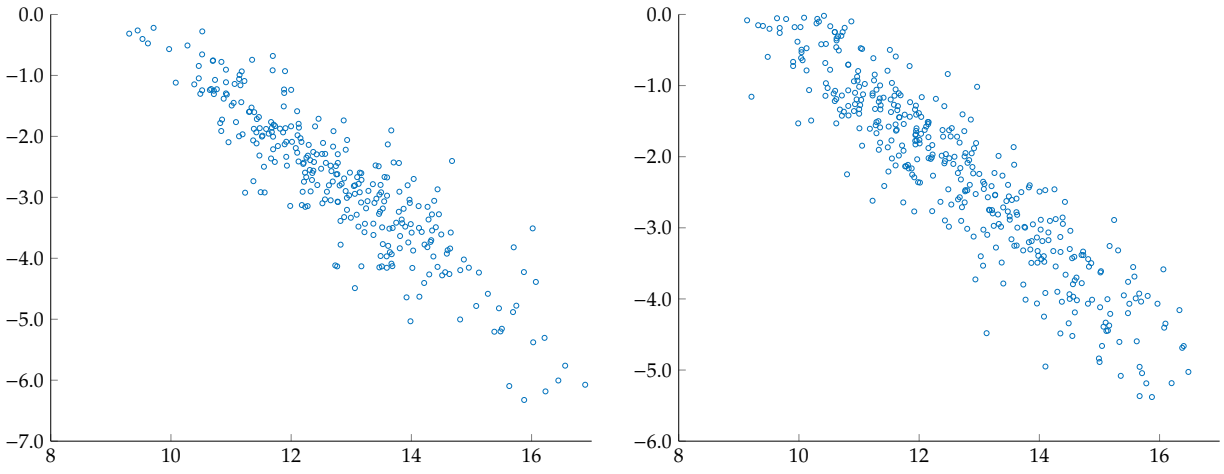
Since Proposition 5 is only valid for rentiers, we must separate households into rentiers and non-rentiers. We say that a household is a rentier if net financial assets constitute more than 90% of proxy total wealth. We also exclude from the rentier set those households whose net assets are in the top and bottom 2.5%.[26] Since there are few rentiers younger than 40 years old in our sample, we do not report estimates for households below 40.

[26]Our 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.

6.3 Money Metric Wealth for Rentiers

Figure 7 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 7 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 7, is relatively understudied. However, as our estimates of money-metric wealth show, in our dataset, dynamic non-homotheticity is an order of magnitude more powerful than static non-homotheticity.

Figure 7: Log consumption-wealth against log wealth for rentiers



(a) 2005 (b) 2019

To compute money-metric values for rentiers using Proposition 5, 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.[27]

[27]We 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.

We apply Proposition 5 to our estimated cross-sectional curves to recover money-metric utility as a function of date, wealth, and age group.[28] For illustration, we use the initial year, $t_0 = 2005$, as the base year, so that money-metric wealth values map nominal wealth in each year t into equivalent wealth in 2005. For our benchmark results, we set the 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.

Best et al. (2020) do not estimate Hicksian (compensated) elasticities. Luckily, Slutsky's equation implies that, if consumption is a normal good, then the compensated intertemporal elasticity should be smaller in magnitude than the uncompensated one. 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.[29]

The left column of Figure 8 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 rentiers, sampling uncertainty is higher for younger age groups. In all cases, the dynamic money-metric wealth is higher than wealth deflated by CPI.

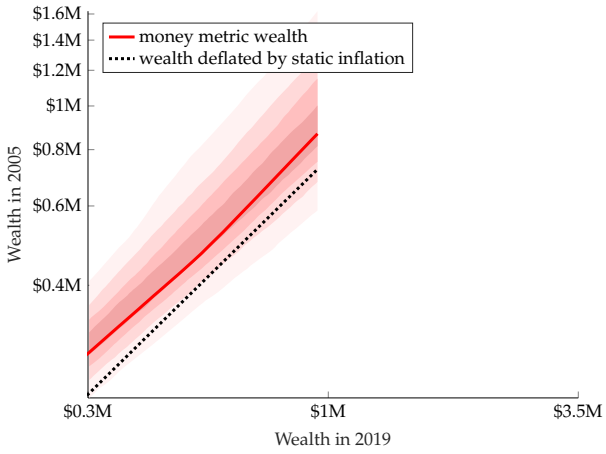
The right column of Figure 8 plots the dynamic annualized inflation rate (log difference between nominal wealth in 2019 dollars and money-metric wealth in 2005 dollars). CPI inflation over this period was 2%, but the dynamic inflation rate is always below 2%. Furthermore, the dynamic inflation rates vary as a function of wealth and age highlighting the importance of non-homotheticity and lifecycle considerations.

To understand the pattern in Figure 8, we decompose the dynamic inflation rate, the deflator in equation (18), into a static and a forward-looking part. Specifically, for a household with wealth w in 2019, the change in the ideal cost-of-living index between

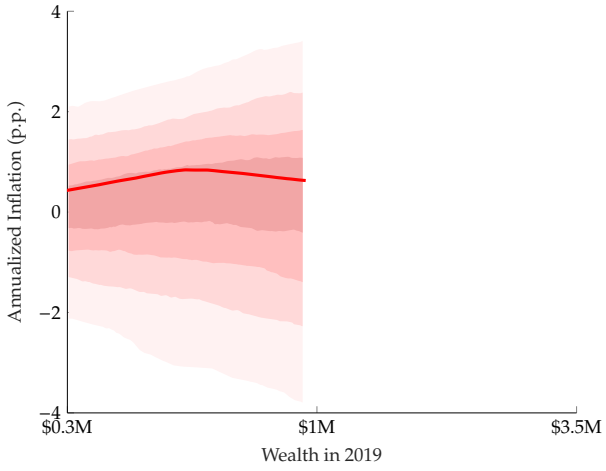
[28] To recover the money-metric, we need to solve the integral equation in Proposition 5. To do so, we use the "recursive" methodology described in Baqaee et al. (2024).

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

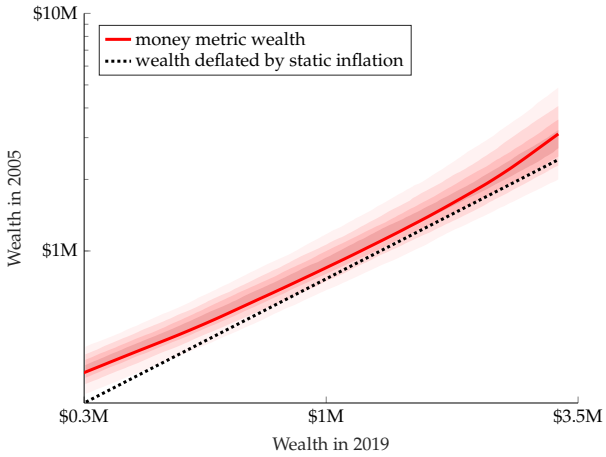
Figure 8: Conversion of 2019 dollars to 2005 dollars



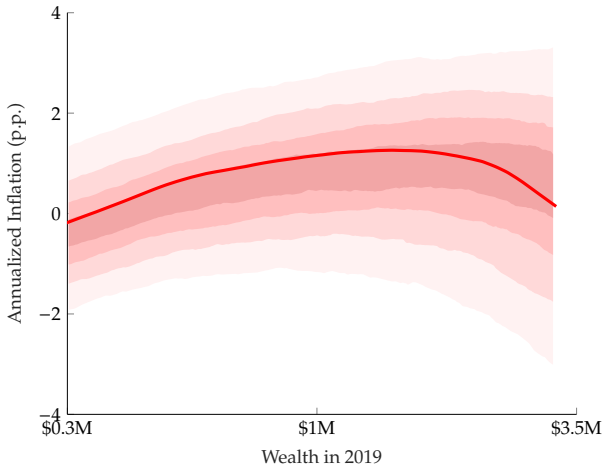
(a) money-metric 40 – 49 year olds



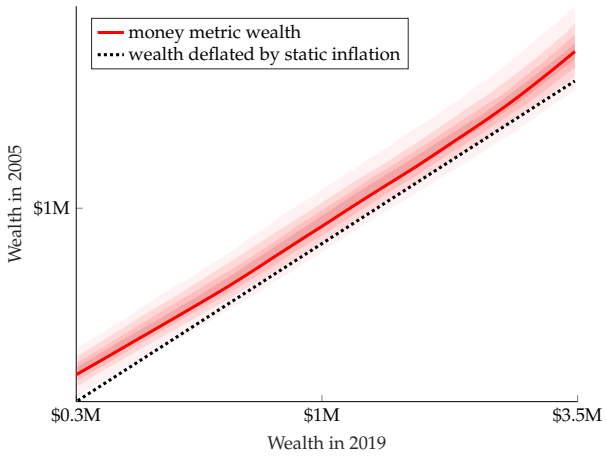
(b) annualized inflation 40 – 49 year olds



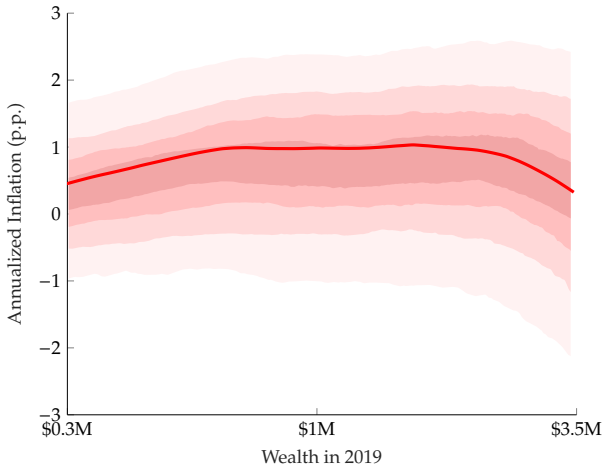
(c) money-metric 50 – 59 year olds



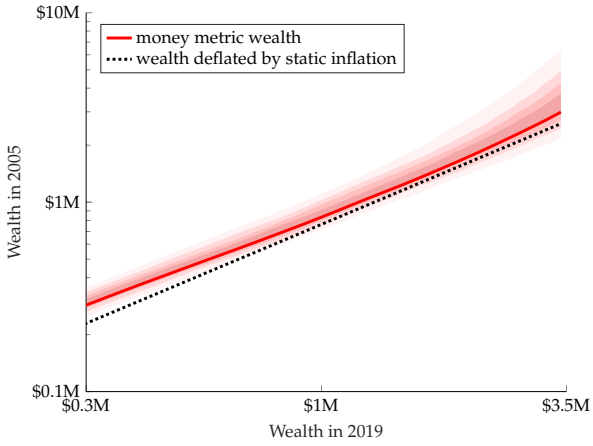
(d) annualized inflation 50 – 59 year olds



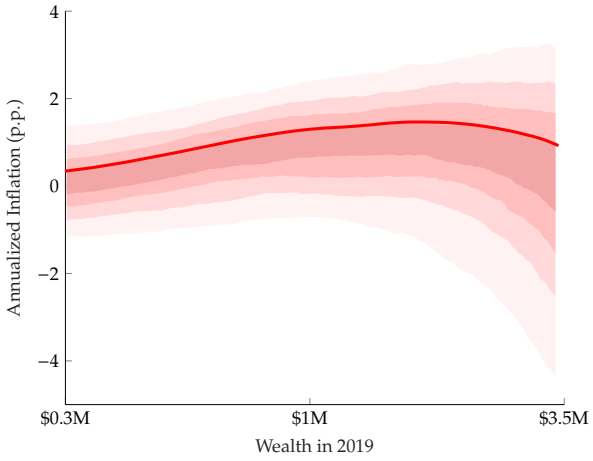
(e) money-metric 60 – 69 year olds



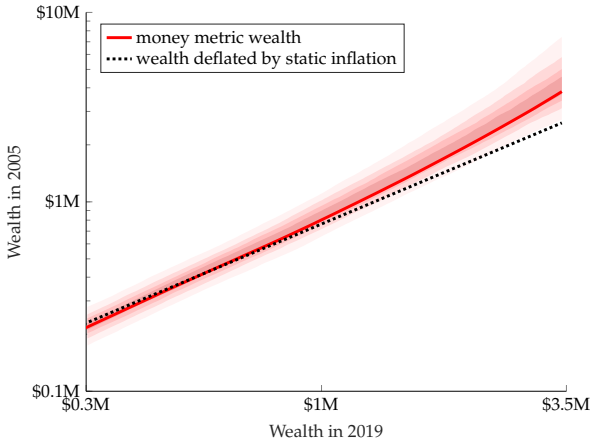
(f) annualized inflation 60 – 69 year olds



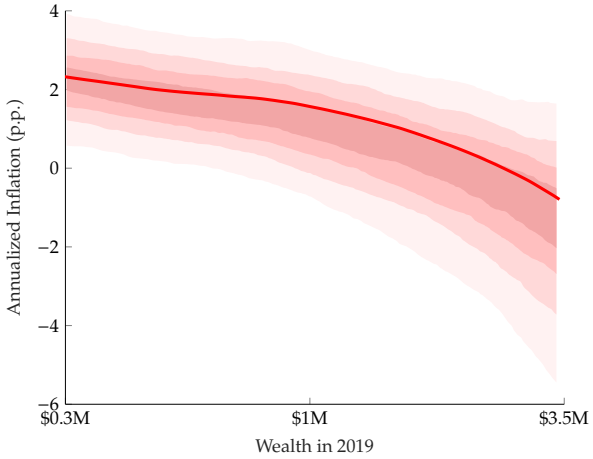
(g) money-metric 70 – 79 year olds



(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 9: Decomposing inflation for 60-69 year olds

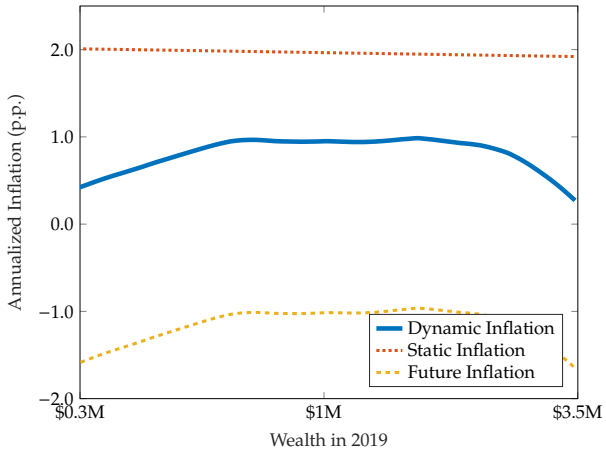
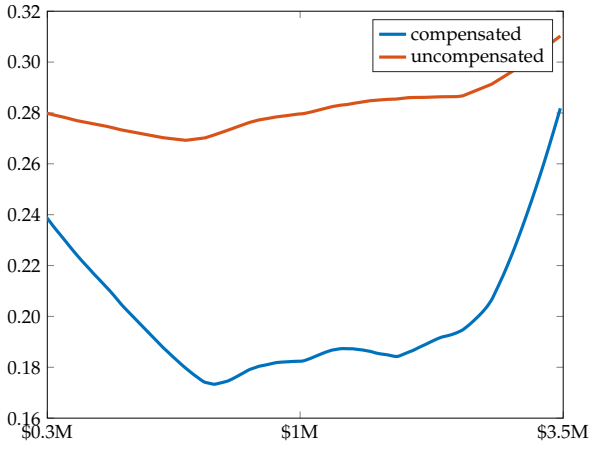


Figure 10: Changes in log consumption wealth ratios for 60 – 69 year olds



2005 and 2019 is:

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

where w_x^* ensures that we are using compensated consumption-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 9 for the $60 - 69$ year old group. The static inflation term is not exactly the same as aggregate CPI because it weighs changes in static prices using compensated budget shares rather than aggregate budget shares. Nevertheless, the static component is very close to aggregate CPI inflation at around 2% per year for all wealth levels. The very slight downward slope reflects the non-homotheticity of static preferences, and static inflation is slightly higher for poorer households, consistent with other studies of the US, like Jaravel and Lashkari (2024), that show that the static cost-of-living index has tended to rise more quickly for poorer households. Nevertheless, the slope of the static inflation line is mild compared to the slope of the dynamic inflation measure.[30]

The component corresponding to future inflation in Figure 9 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 9. 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 10 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

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

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.[31]

Our methodology does not identify which future prices or beliefs are responsible for the patterns in Figure 8. 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 5 does not apply to these households. To recover the money-metric for these households, we rely on Propositions 6 and 7 instead. For each date τ we fit log money-metric values to static budget shares and age group indicators for the rentiers:

$$\log w_{h,\tau} = \boldsymbol{\alpha}_\tau' \boldsymbol{X}_{h,\tau} + \text{error}_{h,\tau}, \tag{24}$$

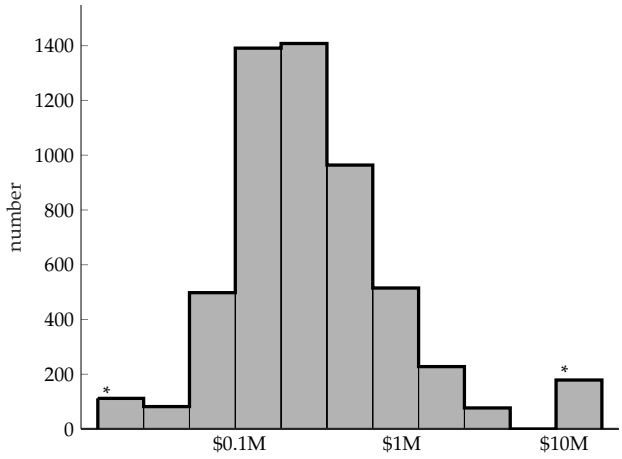
where $\boldsymbol{X}$ are budget shares and age groups. We then use the predicted values from this regression to impute money-metric wealth for the non-rentiers conditional on their budget

[31] Note 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 level both are changing over time. In particular, this implies that the aggregate consumption to wealth ratio displayed in Figure 6 is not informative about changes in either the compensated or uncompensated ratio.

shares, age, and date.[32]

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 11: Distribution of money-metric wealth in 2019



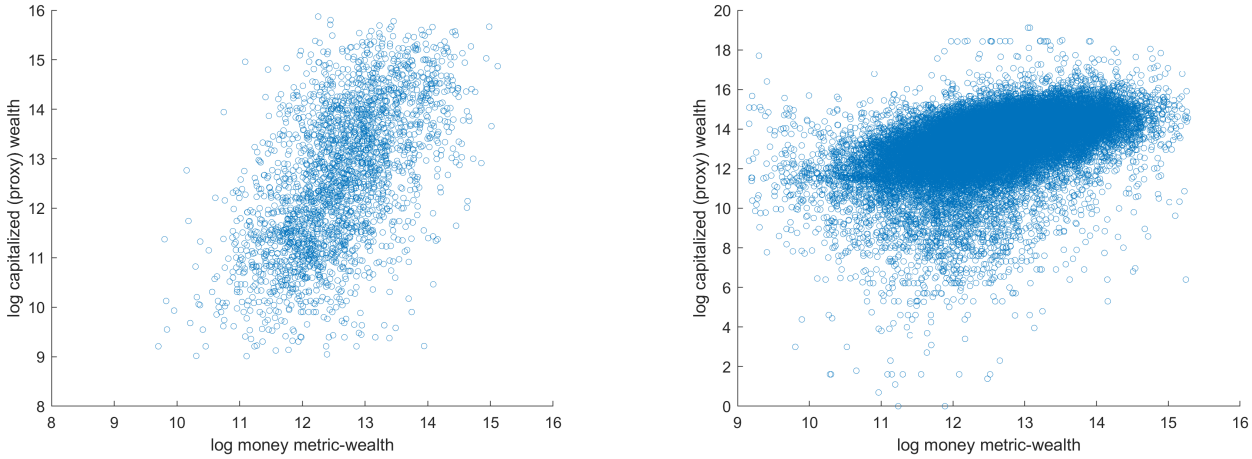
Notes: Bars with asterisk are households whose estimated money-metric wealth in $2019 dollars is outside of the wealth range of rentiers in that year.

Figure 11 displays the distribution of money-metric wealth, in 2019, for all age groups. Figure 12 plots nominal log money-metric wealth against nominal capitalized (proxy) wealth for rentiers and non-rentiers. If financial markets are complete and absent measurement error, these two figures should be 45 degree lines.[33]

[32]We also consider an additional robustness check where we exclude from the set of rentiers outliers in regression (24). 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.

[33]See Appendix C for more details.

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



(a) Rentiers (b) Non-rentiers

7 Treatment Effect on Welfare

Many policies and shocks affect households along many margins simultaneously. For example, job training programs, changes in tax policy, changes in monetary policy, or job loss all plausibly have dynamic effects on many different relevant variables for households. 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: Log money-metric wealth and job loss

	log nominal money metric	
	(1)	(2)
Job Loss	−0.197***	−0.218***
	(0.031)	(0.034)
Job Loss ×**1**(age ≥ 60)		0.180**
		(0.083)
Lagged LHS	Yes	Yes
Controls	Yes	Yes
Observations	48,357	48,357

Standard errors in parentheses.

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

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

Table 1 reports the results. The outcome in the first two columns is log nominal money-metric wealth, $\log m(\tau, w, \boldsymbol{y}|\tau)$. 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).

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.

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 method 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 welfare without fully specifying a structural model. This makes it useful for ex-post analysis of the welfare effect of shocks that have dynamic stochastic effects on many welfare-relevant margins.

References

Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li (2019). Missing growth from creative destruction. *American Economic Review 109*(8), 2795–2822.

Alchian, A. A. and B. Klein (1973). On a correct measure of inflation. *Journal of Money, Credit and Banking 5*(1), 173–191.

Almås, I. (2012). International income inequality: Measuring ppp bias by estimating engel curves for food. *American Economic Review 102*(2), 1093–1117.

Almås, I., T. K. Beatty, and T. F. Crossley (2018). Lost in translation: What do engel curves tell us about the cost of living?

Alvarez, F. and U. J. Jermann (2004). Using asset prices to measure the cost of business cycles. *Journal of Political economy 112*(6), 1223–1256.

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

Argente, D., C.-T. Hsieh, and M. Lee (2023). Measuring the cost of living in mexico and the united states. *American Economic Journal: Macroeconomics 15*(3), 43–63.

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

Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012). New trade models, same old gains? *American Economic Review 102*(1), 94–130.

Ashenfelter, O. (2006). Measuring the value of a statistical life: problems and prospects. *The Economic Journal 116*(510), C10–C23.

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

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

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

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

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

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

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

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

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

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.

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.

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.

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.

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.

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*.