Sufficient Statistics for Measuring Forward-Looking Welfare

Lautaro Adamovsky — UCLA
David R. Baqaee — UCLA
Ariel Burstein — UCLA
Yasu Koike* — UNC Chapel Hill

June 18, 2026

Abstract

We provide a sufficient-statistic method for measuring forward-looking welfare in money-metric terms. The method adapts the money metric to dynamic environments with uncertainty, incomplete markets, and borrowing constraints. Under time separability and using households with negligible idiosyncratic income risk as a reference group, we infer the welfare-relevant value of future consumption from consumption-saving behavior, given an estimate of the intertemporal elasticity of substitution (EIS). In simulations disciplined by PSID income dynamics, the method tracks true welfare more accurately than net-present-value calculations. Applying the method to PSID households from 2005 to 2019, we estimate dynamic cost-of-living indices below the static CPI benchmark. This is because, controlling for age and real wealth, unconstrained households in 2019 spend a greater share of wealth on consumption than comparable households in 2005. With an EIS below one, this implies a lower perceived cost of future consumption relative to present consumption. We also illustrate how the resulting welfare measure can be used as an outcome in reduced-form work by estimating welfare losses associated with job loss.

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

1 Introduction

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

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

Therefore, when dynamic considerations are important, researchers tend to rely on one of two approaches. The first is to work with fully-specified dynamic models, where we specify every parameter pertaining to preferences and beliefs. We can then measure welfare using the model structure. Such a dynamic model can, of course, be used to compute welfare directly if it is specified correctly. The problem is that correctly specifying and calibrating such a nonlinear model is a difficult task fraught with specification and estimation error.

The second approach is less structural. It instead computes the net-present value of real wealth, using asset markets and envelope arguments to approximate complete contingent-claims markets. While this avoids the strong functional form assumptions of the other approach, it relies on local approximation around complete markets. It also still requires forecasting future expectations about cashflows and prices, and assumptions about how these should be discounted. (For examples, see the discussion of the related literature below).

We offer an alternative to both approaches. We provide assumptions under which a forward-looking money-metric welfare measure can be computed without fully specifying functional forms or the stochastic processes for future income, goods prices, and asset

[1] A money metric converts expenditures under some vector of prices into equivalent expenditures given some baseline vector of prices. For a textbook presentation, see Deaton and Muellbauer (1980).

returns. The method can be applied even when financial markets are far from complete.

The first step is to define the welfare object targeted by this measurement exercise. We measure welfare using a dynamic money metric. For any household in any year, the dynamic money metric is the amount of wealth that would make the household indifferent between two situations: facing the original dynamic problem, with its income risk and future opportunities, and facing some base-year environment with that amount of wealth and no idiosyncratic income risk. This measure converts both price changes and idiosyncratic risk into base-year dollars. The comparison is between decision problems evaluated by a fixed preference relation, not between different people. This dynamic money metric is a measurement device, not a literal policy experiment.

We then provide a sufficient-statistics method to estimate this measure of welfare for a sample of consumers with common preferences. Besides commonality of preferences, which is a typical maintained assumption in the literature, our analysis relies on two other key assumptions. The first assumption is a type of separability in preferences between the present and the future (formally, separability of the distance function). The second assumption is that there is a reference subset of households that face negligible idiosyncratic undiversifiable risk. We call these reference households the "rentiers". (Rentiers can and do face aggregate risk). Given these two assumptions, we can obtain forward-looking measures of welfare without further assumptions about utility functions (e.g. CES across goods), risk preferences (e.g. expected utility), time preferences (e.g. exponential discounting), beliefs (e.g. the stochastic processes of prices, returns, and incomes), financial frictions (e.g. complete markets), and without resorting to first-order approximations.

It is useful to compare these assumptions with those of the alternatives. Relative to static money-metric measures, which underlie measures of real consumption growth and inflation, our approach is more general because it allows for forward-looking behavior. Relative to fully specified dynamic models, it is less parametric because it does not require fully specifying preferences or the stochastic environment. Relative to net-present-value calculations, it allows for large departures from complete markets and does not require the researcher to estimate future beliefs and behavior. We provide simulations showing that the method outperforms net-present-value calculations even when the rentier reference group is imperfect and time separability is violated.

We sketch the basic idea of our approach. We infer movements in the intertemporal price of future versus present consumption from changes in consumption-saving choices of rentiers, given an estimate of the elasticity of intertemporal substitution (EIS). When $EIS < 1$ (the empirically relevant case), an increase in the consumption-to-wealth ratio for rentiers reveals that the price of consuming in the future relative to the present must have fallen. Hence, a forward-looking measure of welfare will be higher than a static measure that only compares contemporaneous price changes.[2]

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

For non-rentiers, changes in consumption-to-wealth ratios cannot be used to infer expectations about future prices, because consumption-to-wealth ratios also respond to variation in the bindingness of financial constraints and to temporary fluctuations in risky income. Therefore, to recover money-metric utility for non-rentiers, we use a different strategy. Under our assumption of time separability, budget shares in the present can only be a function of relative prices in the present and overall utility. Hence, if budget shares in the present are a one-to-one function of utility conditional on relative prices—for example, if some Engel curve is monotone—then two households facing the same relative prices with the same budget shares must be on the same intertemporal indifference curve. This allows us to construct money-metric values for non-rentiers by matching them with rentiers that have similar static budget shares in the same period.[3]

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 EIS (which could be a constant or vary as a function of wealth and prices).

Before taking this method to the data, we evaluate it in simulated incomplete market economies disciplined by PSID income dynamics. Since the model is fully specified in the simulation, we can compare our estimated money metric to the true welfare object. The method tracks true welfare closely and outperforms net-present-value calculations, especially for financially constrained households exposed to uninsurable income risk. The

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

[3]This does not hold exactly if preferences are not time-separable. In Sections 4.3 and 5, we discuss the implications of violating time-separability and the performance of our method in that case.

simulations also show that the method performs well even when the key assumptions are violated: namely, when time separability of preferences is imperfect and there are no pure rentiers. Specifically, we find that in the simulated ergodic distribution there are households with relatively high assets and low income that can proxy for rentiers even if no household in the sample is truly free from uninsurable idiosyncratic risk.

We then apply the method to PSID households from 2005 to 2019. We find that conventional static consumer price indices overestimate the true dynamic price index in our sample. This means that, over time, households became more optimistic about the future relative to the present. This would be consistent, for example, with an increase in expected returns over time, which would lower the price of consuming in the future relative to the present. Mechanically, our finding follows from the increase in the compensated consumption-to-wealth ratio that we estimate, in combination with an EIS calibrated to be less than one.[4]

Furthermore, we find more heterogeneity in the dynamic cost-of-living index than in the static cost-of-living index across both the wealth distribution and by age group. Static studies document meaningful inflation heterogeneity across households, often with lower-income households facing higher inflation in particular periods or product markets (Hobijn and Lagakos, 2005; Kaplan and Schulhofer-Wohl, 2017; Argente and Lee, 2021; Jaravel, 2019). Our dynamic measure points to a different pattern: in our application, dynamic inflation rates tend to be lower for lower-wealth households. This difference arises because the forward-looking component, inferred from compensated changes in consumption-to-wealth ratios, is much more heterogeneous than the static component.

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

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

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

The first branch uses exact fully-specified models. A prototypical example is described in the AEA presidential address by Lucas (2003). Since this approach is ubiquitous in macroeconomics, we do not provide a systematic review, and instead focus on methodological papers that extend growth and inflation measures to account for dynamics. For example, to construct dynamic measures of inflation and welfare, Reis (2005) and Aoki and Kitahara (2010) calibrate models of household preferences and beliefs, and compute aggregate cost-of-living indices by feeding in the path of observed prices. In a different vein, Jones and Klenow (2016) construct country-level welfare measures that account for mortality, leisure, and risk by fully specifying preferences and calibrating stochastic processes for the determinants of utility.

An issue with this approach is that the researcher has to make assumptions about everything. An advantage of the second approach, which uses net-present value calculations, is that it does not require a fully specified model of preferences. In this approach, researchers forecast the future and discount variables back to the present using asset prices. For example, at the macroeconomic level, Basu et al. (2022) show that, to a first-order, the welfare of a country's infinitely-lived representative consumer can be summarized by the net-present value of technology shocks plus the initial capital stock. At the microeconomic level, Fagereng et al. (2022) and Del Canto et al. (2023) use Taylor approximations to calculate how individual consumer welfare responds to shocks to asset prices and monetary policy, respectively. Fagereng et al. (2022) estimate how various asset prices changed over time and weigh these by discounted net asset sales. Del Canto et al. (2023) use local projections to forecast how monetary shocks change goods and asset prices and weigh these price changes by discounted expected budget shares. Pallotti et al. (2024) use a similar methodology to study the welfare effects of the recent surge in inflation in the Euro area across the household distribution.[5] There are also papers that combine the structural modelling approach with the net-present value one. For example, Huggett and

[5] 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.

Kaplan (2016) use the stochastic discount factor from a structural model to discount the value of future labor income.

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

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

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, 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.[6] We extend this idea to infer the value of missing future prices relative to present prices using changes in consumption-to-wealth ratios.

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

Third, to infer welfare of financially constrained agents, we build on Hamilton (2001), who infers changes in welfare by tracking changes in the budget share of food. The basic idea is that consumers spend a smaller share of their budget on food as they get

[6]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).

richer, so changes in the budget share of food are informative about how rich consumers are.[7] Recently, Atkin et al. (2024) provide a micro-foundation for this approach and use it to calculate welfare across the income distribution. We use a similar idea to Hamilton (2001), because we match unconstrained rentiers to constrained non-rentiers within the same period using static budget shares. However, since rentiers and non-rentiers in the same period face the same (static) relative prices, we do not need to estimate or correct for substitution effects, as in Atkin et al. (2024).

Roadmap. Section 2 sets up the general environment and defines dynamic money-metric welfare. Section 3 contains the main results of the paper: we first show how to recover dynamic money-metric welfare for households that do not bear idiosyncratic risk, and then extend the argument to households that do face such risks. Section 4 provides extensions of the basic framework, including departures from common prices and beliefs, risk-free labor income streams for rentiers, and violations of time separability; additional extensions to time-inconsistency, changing mortality risk, and labor-leisure choice are collected in the appendix. Section 5 evaluates the performance of the method in simulated data, including its robustness to violations in our assumptions. Section 6 applies the method to construct a measure of dynamic welfare for households in the PSID. Section 7 uses this measure to study the dynamic welfare losses from job loss. Section 8 concludes.

2 General Environment and Measure of Welfare

In this section, we set up the general environment. We define preferences, decision problems, and the dynamic money metric. The main theoretical results are in the next section.

2.1 Decision Problem of Households

Consider an agent with continuous, non-satiated, ordinal preferences $\succeq$ over state-contingent consumption streams, $\boldsymbol{c}$, with probabilities, $\boldsymbol{\pi}$. We use the utility function $\mathcal{U}(\boldsymbol{c}, \boldsymbol{\pi})$ to represent these preferences. 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 $\boldsymbol{c}$ and $\boldsymbol{\pi}$ have dimension $|S^J| \times |N|$.

[7] 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).

We assume a finite and fixed planning horizon $J < \infty$.[8] We index decision problems by the start date τ. Consumers choose consumption and asset portfolios to maximize utility subject to a sequence of state-contingent budget constraints. The first-period budget constraint requires that spending on consumption and asset purchases equals initial wealth:

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

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

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

$$\sum_{n \in N} p_n(s^j|\tau) c_n(s^j) + \sum_{k \in K} a_k(s^j) = \sum_{k \in K} R_k(s^j|\tau) a_k(s^{j-1}) + y(s^j|\tau), \tag{2}$$

where $R_k(s^j|\tau)$ is the return on asset k in history s^j, and $y(s^j|\tau)$ is exogenous labor income.[9]

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

$$\sum_{k} a_k(s^j) \geq -X(s^j|\tau). \tag{3}$$

We impose a no-Ponzi condition by setting $X(s^J|\tau) = 0$ for every terminal history s^J, so that the agent cannot end the problem in debt.

The decision problem faced by a household depends on prices, returns, income, 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 $\boldsymbol{X}$ are functions of the start date τ, while the income process $\boldsymbol{y}$ and initial wealth w are household-specific. Hence, each household's problem can be indexed by the tuple $(\tau, w, \boldsymbol{y})$. Define the value function

[8] Our results can be extended to infinite horizon. In our baseline, we also assume that mortality risk does not depend on calendar time τ. See Section 4 for a discussion of how our results can be extended to account for secular trends in mortality risk.

[9] In Section 4, we discuss how to extend the model to allow for endogenous labor-leisure choice.

associated with $(\tau, w, \boldsymbol{y})$ by

$$V(\tau, w, \boldsymbol{y}) = \max_{\boldsymbol{c},\boldsymbol{a}} \left\{ \boldsymbol{\mathcal{U}}(\boldsymbol{c}, \boldsymbol{\pi}) \text{ subject to (1), (2), and (3)} \right\}. \tag{4}$$

The magnitude of the value function is not interpretable, since utility is only determined up to a monotone increasing transformation. Hence, to measure changes in the value function, we convert utils in (4) into dollars using a money metric.

2.2 Measuring Welfare and the Cost of Living

To express welfare in dollars, we compare each decision problem to a reference problem. The reference problem is one in which the household starts with some amount of wealth in a base year and bears no idiosyncratic income risk.[10] We call such a problem a *rentier* problem. Formally, a consumer is a rentier if $y(s^j|\tau) = 0$ in every state and date. A rentier in date τ with wealth w faces the problem $(\tau, w, \boldsymbol{0})$.

Definition 1 (Dynamic Money Metric)**.** Consider a base period τ_0, with prices, returns, probabilities, and borrowing constraints

$$\{\boldsymbol{p}(\cdot|\tau_0), \boldsymbol{R}(\cdot|\tau_0), \boldsymbol{\pi}(\cdot|\tau_0), \boldsymbol{X}(\cdot|\tau_0)\},$$

where $\boldsymbol{p}(\cdot|\tau_0) > 0$ and $\boldsymbol{R}(\cdot|\tau_0) > 0$. The τ_0-*money metric* for the problem $(\tau, w, \boldsymbol{y})$ is defined implicitly by

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

Thus, $m(\tau, w, \boldsymbol{y}|\tau_0)$ is the base-year rentier-equivalent wealth of the original problem. It converts a dynamic stochastic choice set, including future opportunities and idiosyncratic income risk, into base-year dollars. The rentier-equivalent problem itself has no idiosyncratic income risk: the effect of such risk in the original problem is reflected in the amount of base-year wealth that makes the household indifferent. When the base period is clear from context, we suppress τ_0 and write $m(\tau, w, \boldsymbol{y})$ instead of $m(\tau, w, \boldsymbol{y}|\tau_0)$.

In a static deterministic world, the dynamic money metric coincides with the textbook money metric (e.g., Deaton and Muellbauer, 1980). We extend this textbook definition to allow for market incompleteness and changes in probabilities, which do not have counterparts in static consumer theory. The following shows that the dynamic money metric $m(\tau, w, \boldsymbol{y}|\tau_0)$ is a cardinalization of the value function.

[10] It would be straightforward to allow idiosyncratic but insurable risk in the rentier problem. For example, idiosyncratic health or property risk could be included as long as it is fully insurable.

Proposition 1 (Money metric cardinalizes utility)**.** *For any base year* τ_0*,*

$$m(\tau, w, \boldsymbol{y}|\tau_0) \geq m(\tau', w', \boldsymbol{y}'|\tau_0)$$

if and only if

$$V(\tau, w, \boldsymbol{y}) \geq V(\tau', w', \boldsymbol{y}').$$

Given the money metric, we can also define changes in the cost of living in the usual way.

Definition 2 (Dynamic Cost of Living)**.** The change in the cost of living between two base 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 calculate the ratio of lump sums required in the two base environments to reach the indifference curve indexed by $V(\tau, w, \boldsymbol{y})$. In a static deterministic environment, Definition 2 is the traditional ideal (Konüs, 1939) price index from consumer theory.

The goal of the rest of the paper is to infer the money metric function $m(\tau, w, \boldsymbol{y}|\tau_0)$, which depends not only on current prices, wealth, and income, but also on expectations, preferences, and constraints in the future.

2.3 Time-Separable Preferences

We impose the following assumption on preferences. Without loss of generality, note that the preference relation $\succeq$ can also be represented by an implicitly defined utility function:

$$U = D(\boldsymbol{c}, \boldsymbol{\pi}, U), \tag{5}$$

where D is homogeneous of degree one in $\boldsymbol{c}$.[11]

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

[11] To see this, let $\mathcal{U}(\boldsymbol{c}, \boldsymbol{\pi})$ be some utility function that represents $\succeq$. Define the distance function $\tilde{D}(\boldsymbol{c}, \boldsymbol{\pi}, U) = \max_{\alpha>0}\{\alpha : \mathcal{U}(\boldsymbol{c}/\alpha, \boldsymbol{\pi}) \geq U\}$, which is homogeneous of degree one in $\boldsymbol{c}$ by construction. The following is an identity: $\tilde{D}(\boldsymbol{c}, \boldsymbol{\pi}, \mathcal{U}(\boldsymbol{c}, \boldsymbol{\pi})) = 1$. Now define $D(\boldsymbol{c}, \boldsymbol{\pi}, U) \equiv U\tilde{D}(\boldsymbol{c}, \boldsymbol{\pi}, U)$. Then D is homogeneous of degree one in $\boldsymbol{c}$ and satisfies $\mathcal{U}(\boldsymbol{c}, \boldsymbol{\pi}) = D(\boldsymbol{c}, \boldsymbol{\pi}, \mathcal{U}(\boldsymbol{c}, \boldsymbol{\pi}))$. Thus, we obtain equation (5).

Definition 3 (Time Separability)**.** The preference relation $\succeq$ is *time separable* if it has a utility function representation that can be written as

$$U = D(\underbrace{P(\mathbf{c}(s^0), U)}_{\substack{\text{present}\\ \text{bundle}}}, \underbrace{F(\{\mathbf{c}(s^j)\}_{j>0}, \{\pi(s^j)\}_{j>0}, U)}_{\substack{\text{future}\\ \text{bundle}}}, U), \tag{6}$$

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

Definition 3 is equivalent to imposing a certain type of separability on the compensated demand curves generated by $\succeq$.[12] Specifically, Definition 3 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. If preferences are homothetic, then we can drop U from the right-hand side of (6). Time separability in Definition 3 is a restriction on the ordinal preference relation, not on its cardinal representations.

Standard constant-relative-risk-aversion (CRRA) expected-utility preferences satisfy time separability, but the condition is considerably more general. We provide some examples below.

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

- Epstein and Zin (1989) preferences:

$$U_j(s^j) = \left[(1-\beta)c(s^j)^{\frac{\sigma-1}{\sigma}} + \beta\left(\sum_{s^{j+1}} \pi(s^{j+1} \mid s^j) U_{j+1}(s^{j+1})^{\frac{\gamma-1}{\gamma}}\right)^{\frac{\gamma}{\gamma-1}\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{7}$$

  where σ controls the EIS and γ controls risk aversion.

- Variable intertemporal elasticity of substitution:

$$U = \frac{c(s^0)^{1-1/\sigma(U)}}{1-1/\sigma(U)} + \sum_{j=1}^{J} \beta^j \sum_{s^j} \pi(s^j) \frac{c(s^j)^{1-1/\sigma(U)}}{1-1/\sigma(U)}, \tag{8}$$

[12] Atkin et al. (2024) call this property quasi-separability. This notion of separability is also sometimes called separability of the distance function or separability of the expenditure function. See Blackorby et al. (1998) for theoretical background.

where the elasticity of substitution across time can vary as a function of utility.

- Variable income effects across goods within a period:

$$U = \frac{1}{1-1/\sigma}\mathbb{E}_0 \sum_{j=0}^{J} \beta^j C_j^{1-\frac{1}{\sigma}}, \quad \text{where} \quad C_j = \left[\sum_n \omega_n^{\frac{1}{\gamma}} \left[\frac{c_{nj}}{U^{\epsilon_n}}\right]^{\frac{1-\gamma}{\gamma}}\right]^{\frac{\gamma}{\gamma-1}}, \tag{9}$$

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

Appendix B provides some examples that are non-separable.

Although we impose time separability for most of the paper, we discuss how our results are affected if preferences are not perfectly time-separable in Section 4.

3 Main Theoretical Results

In this section, we present the main theoretical results. We begin by describing the observables and maintained assumptions that allow us to apply the theory to data. We then proceed in two steps. First, we show how to recover the money metric for rentier households. Second, we extend the procedure to non-rentier households.

3.1 Observables and Maintained Assumptions

We now describe the observables used by our approach and state the maintained assumptions on the data-generating process.

Fix a vector of observable household characteristics x, such as age. We allow preferences to vary with x: for example, households of different ages may have different planning horizons, discounting, or bequest motives. The maintained stability assumption is that, conditional on x, preferences do not vary directly as a function of calendar time.[13] This is a standard assumption in the literature that attempts to measure dynamic or static welfare using revealed preference arguments.

We observe repeated cross-sections of households with characteristics x. The method does not require that we follow the same households over time. We call the households

[13] In the empirical application, we condition on age groups and treat each age group as a separate problem. With richer data, one could also condition on other observables, such as household size, location, education, or family structure.

with characteristics x observed at calendar date τ the *date-τ cohort*. Although these cohorts share the same ordinal intertemporal preference relation (they rank bundles in the same way), their decision problems may change over time: prices, returns, beliefs, borrowing constraints, and income risks may all vary with calendar time.

For each household in each cohort, we observe financial wealth and current spending on each good, and hence current budget shares and consumption-to-wealth ratios. We also observe good-level inflation rates, which are common across households in the same cohort. Finally, we observe, or can classify, whether each household belongs to the rentier reference group, and we take as given an estimate of the elasticity of intertemporal substitution.

The rest of this section shows how to recover the money metric from these observables, under Definition 3 and the maintained assumptions on the data-generating process. We begin with rentiers and then extend the construction to non-rentiers.

3.2 Obtaining the Money Metric for Rentiers

We begin with rentier households, the reference group that bears negligible idiosyncratic income risk. For concreteness, we treat rentiers as households with no idiosyncratic labor income, so that $\boldsymbol{y} = \mathbf{0}$ in every state. Section 4 shows the argument extends to households with nonzero but risk-free labor income.

The sufficient statistics for rentiers are current budget shares, the share of wealth spent on current consumption, and the EIS, where "current" refers to the initial period s^0 of the problem. To define these objects, let current expenditures for a household facing problem $(\tau, w, \boldsymbol{y})$ be

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

and let the budget share of good n be

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

Let the current expenditure-to-wealth ratio, or current consumption-to-wealth ratio, be

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

Denote the compensated elasticity of intertemporal substitution (EIS) for a household facing problem $(\tau, w, \boldsymbol{y})$ by $\sigma(\tau, w, \boldsymbol{y})$.[14]

To build intuition, we first consider a sequence of special cases leading up to the general result. We begin with the homothetic one-good case.

Proposition 2. *Suppose preferences are homothetic, the EIS is constant, $\sigma(\tau, w, \boldsymbol{y}) = \sigma$, and there is only one consumption good per period. Then, for rentiers,*

$$\log m(\tau, w, \mathbf{0}) = \underbrace{\log w - \log \frac{p(s^0|\tau)}{p(s^0|\tau_0)}}_{\text{static "real wealth"}} + \underbrace{\frac{\log\left(B^P(\tau)/B^P(\tau_0)\right)}{1-\sigma}}_{\substack{\text{future-price} \\ \text{adjustment}}}, \tag{10}$$

where homotheticity allows us to suppress the dependence of B^P on wealth.

This proposition shows that dynamic welfare equals static real wealth plus an adjustment for the relative price of the future, inferred from the consumption-to-wealth ratio. The first term is the familiar static money metric: nominal wealth deflated by the change in the current price of consumption. The second term adjusts for changes in the relative price of future consumption, which are inferred from changes in the consumption-to-wealth ratio. Even in this special case, the result does not require complete markets: financial markets may be incomplete, and preferences over future consumption can be more general than expected utility with a single discount factor and risk aversion governed by the inverse of the EIS.

For example, suppose preferences are Epstein–Zin as in (7). These preferences are homothetic, satisfy Definition 3, and have a constant EIS equal to σ. Hence, the assumptions of Proposition 2 are satisfied, and (10) applies directly. Notice what is, and is not, part of the expression. The discount factor β and the risk-aversion parameter γ in (7) affect households' consumption-saving choices, but they do not enter (10) directly. Their relevant effect on money-metric welfare is summarized by the consumption-to-wealth ratio.

Example 2 (Two-period CRRA)**.** This example shows why the consumption-to-wealth ratio contains information about unobserved future prices. Consider a two-period complete-markets economy with one good per period and cohorts of households with CRRA expected utility,

$$U = \frac{c(s^0)^{1-1/\sigma}}{1-1/\sigma} + \beta \sum_{s^1} \pi(s^1|\tau) \frac{c(s^1)^{1-1/\sigma}}{1-1/\sigma}. \tag{11}$$

[14]See Lemma 4 in the appendix for a formal definition of the EIS. Intuitively, the compensated EIS controls the change in spending on consumption versus savings if the price of every consumption good in the present rises by the same amount, holding utility constant.

Complete markets allow us to collapse the sequence of budget constraints into a single intertemporal budget constraint,

$$p(s^0|\tau)c(s^0) + \sum_{s^1} q(s^1|\tau)c(s^1) \leq w,$$

where $q(s^1|\tau)$ is the date-τ state price of one unit of the good delivered in future history s^1. Because of complete markets, this problem is isomorphic to a static CES consumption problem in which commodities are indexed by date and state and all purchases are made in the first period. Let $P(\tau)$ be the CES ideal price index associated with this problem at date τ:

$$P(\tau) = \left[p(s^0|\tau)^{1-\sigma} + \beta^\sigma \sum_{s^1} \pi(s^1|\tau)^\sigma q(s^1|\tau)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}.$$

This price index depends on unobserved future state prices and probabilities.

In this special case, money-metric welfare is nominal wealth deflated by the change in the intertemporal price index:

$$m(\tau, w, \mathbf{0}) = \frac{w}{P(\tau)/P(\tau_0)}.$$

Moreover, CES demand implies that the share of wealth spent on present consumption satisfies

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

Thus, the consumption-to-wealth ratio reveals the intertemporal price index up to the observed current price:

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

Substituting this expression into $m(\tau, w, \mathbf{0}) = w/[P(\tau)/P(\tau_0)]$ yields (10). This derivation is useful for intuition, but it relies on complete markets and homothetic CRRA preferences. The general result below obtains the same sufficient-statistic logic without those restrictions; its proof is in Appendix A.

The homothetic case in Proposition 2 hides an important empirical issue: with non-homothetic preferences, changes in the consumption-to-wealth ratio combine substitution effects with wealth effects. The next special case keeps one good and a constant EIS, but allows preferences to be non-homothetic.

Proposition 3. *Suppose the EIS is constant, $\sigma(\tau, w, \boldsymbol{y}) = \sigma$, and there is only one consumption* *good per period. Then, for rentiers, the money metric satisfies the following fixed-point problem:*

$$\log m(\tau, w, \mathbf{0}) = \underbrace{\log w - \log \frac{p(s^0|\tau)}{p(s^0|\tau_0)}}_{\text{static "real wealth"}} + \underbrace{\frac{\log\left[B^P(\tau, w, \mathbf{0})/B^P(\tau_0, m(\tau, w, \mathbf{0}), \mathbf{0})\right]}{1-\sigma}}_{\substack{\text{future-price} \\ \text{adjustment}}}. \tag{12}$$

When preferences are non-homothetic, the consumption-to-wealth ratio $B^P(\tau, w, \mathbf{0})$ can change with τ for two reasons: substitution effects and wealth effects. Substitution effects are changes in B^P caused by changes in the price of current consumption relative to future consumption, holding utility fixed. Wealth effects are changes in B^P caused by movements across indifference curves: a household with wealth w in period τ need not be on the same indifference curve as a household with wealth w in period τ_0, because the real purchasing power of w may have changed.

Since we use changes in B^P to infer changes in the relative price of future consumption, we must purge the component of those changes due to wealth effects. To do so, Equation (12) compares the observed consumption-to-wealth ratio $B^P(\tau, w, \mathbf{0})$ with the ratio that would be chosen in the base period by a rentier on the same indifference curve, $B^P(\tau_0, m(\tau, w, \mathbf{0}), \mathbf{0})$. Because these two rentier problems lie on the same indifference curve, their difference in consumption-to-wealth ratios reflects substitution effects rather than wealth effects. The wealth of the base-period rentier is precisely the unknown money metric, so the formula is a fixed-point problem that implicitly pins down $m(\tau, w, \mathbf{0})$.

The final rentier result adds back the two features suppressed so far: multiple goods in each period and state, and an EIS that can vary with date and wealth.[15]

Proposition 4. *The money metric satisfies the following fixed-point problem:*

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

where w_x^* solves

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

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

Relative to Proposition 3, we now allow multiple goods in each period and state, and a varying EIS. The first term inside the integral is static inflation, computed using

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

compensated current budget shares. The second term is the forward-looking adjustment, inferred from compensated changes in the consumption-to-wealth ratio and scaled by the EIS.

The matching condition (14) keeps the household on the same indifference curve along the integration path. Thus, w_x^* is the wealth a rentier would need at date x to have the same money-metric value as the original rentier at (τ, w). This compensation purges wealth effects from budget shares and consumption-to-wealth ratios. Because the compensation itself depends on the unknown money metric, Proposition 4 is a fixed-point problem. Appendix C.1 describes the solution method and the support restrictions under which the fixed point can be implemented without out-of-sample extrapolation.

Brief proof sketch. To prove Proposition 4, we rely on the intertemporal expenditure function. First, for each decision problem, there exist shadow intertemporal prices that rationalize the household's choices. For rentiers, these shadow prices depend only on calendar time and the household's utility level. Given these shadow prices, the expenditure function can be used to compute the rentier money metric. The envelope theorem and the fundamental theorem of calculus then express changes in the money metric as the integral of intertemporal compensated demands with respect to changes in shadow prices. Time separability allows us to replace the welfare-relevant movements in unobserved future shadow prices with movements in the observed consumption-to-wealth ratio, adjusted by the EIS. Finally, the matching condition in (14) recovers the relevant compensated budget shares and compensated consumption-to-wealth ratios from cross-sectional variation, as in Baqaee et al. (2024).

3.3 Obtaining the Money Metric for Non-Rentiers

A challenge for applying Proposition 4 is that it applies only to rentiers. It does not apply directly to households with $\boldsymbol{y} \neq \mathbf{0}$, since these households may face binding borrowing constraints and their total wealth, inclusive of risky labor income, is not priced by financial markets. We therefore cannot infer their welfare from consumption-saving behavior in the same way.

Instead, we use their current budget shares. Under time separability, current budget shares depend on future opportunities only through overall utility. Thus, conditional on current prices, two households with the same current budget shares are on the same intertemporal indifference curve. This allows us to assign a non-rentier the rentier-equivalent wealth of a rentier with the same budget shares.

Lemma 1 (Compensated Budget Shares)**.** *If preferences are time separable, then the current budget share of each good can be expressed as a function only of current 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 function for good n.

Importantly, Lemma 1 implies that the current budget share of each good $B_n(\tau, w, \boldsymbol{y})$ depends on financial wealth w and the income process $\boldsymbol{y}$ only through overall utility $V(\tau, w, \boldsymbol{y})$.

The next proposition makes it possible to extend $m(\tau, w, \mathbf{0})$ to non-rentier households.

Proposition 5 (Budget Shares Identify Current-Date Welfare)**.** *Fix a date τ and current prices $\boldsymbol{p}(s^0|\tau)$. Suppose the compensated budget share function is one-to-one in utility. Then any two households observed at date τ with the same current budget shares have the same current-date money metric:*

$$\boldsymbol{B}(\tau, w, \boldsymbol{y}) = \boldsymbol{B}(\tau, w', \boldsymbol{y}') \quad \Longrightarrow \quad m(\tau, w, \boldsymbol{y}|\tau) = m(\tau, w', \boldsymbol{y}'|\tau).$$

Equivalently, there exists a date-specific function $\tilde{m}_\tau$ such that

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

In words, at a fixed date, current budget shares can be used as a welfare index. If two households face the same current prices and choose the same budget shares, time separability implies that they are on the same intertemporal indifference curve. This condition requires some non-homotheticity: under homothetic preferences, budget shares do not vary with utility at a fixed date.

For rentiers, the identity $m(\tau, w, \mathbf{0}|\tau) = w$ identifies the shape of $\tilde{m}_\tau$ on the rentier support. Intuitively, rentiers reveal the mapping between current budget shares and current-date money-metric values because their current-date money metric equals their observed financial wealth. Hence, we can learn the function $\tilde{m}_\tau$ by fitting rentier wealth to rentier budget shares at date τ. We then use the fitted function to estimate current-date money metrics for non-rentiers whose budget shares lie in the rentier support.

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*

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

In this simple case, if the compensated budget share of good i is monotone in utility, then 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. Proposition 5 generalizes this idea to the full vector of current budget shares.

Propositions 4 and 5 can be combined to extend the money metric from rentiers to non-rentiers whenever the relevant budget-share support conditions are satisfied. First, recover current-date money metrics for non-rentiers by budget-share matching. Then use Proposition 4 to express those values in any base-period dollars.

4 Extensions

Before presenting applications, we discuss three extensions and robustness exercises that clarify the scope of the method. First, we discuss how to relax the common-price and common-belief assumptions. Second, we show that the rentier reference group can include households with positive labor income, provided those income streams are non-risky and can be capitalized. Third, we study what happens when non-homotheticities are not time separable, using the preferences of Comin et al. (2021). Formal statements and proofs are in Appendix D. The appendix also discusses other extensions: time-inconsistent preferences, labor-leisure choice, and calendar-time variation in mortality risk.

4.1 Relaxing Common Prices and Beliefs

We assume that current 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, indexed by τ, but within a date, 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 see why, recall that by Lemma 1, relative spending shares on goods in the present only depend on static prices and utility. Therefore, as long as this function is one-to-one in utility, two households facing the same static prices at a point in time choose the same current budget shares across goods if and only if they are on the same indifference curve.

These restrictions still allow substantial heterogeneity. Current prices may vary with observables such as date, location, or age, and rentier beliefs may change over calendar time. Non-rentiers' future state variables, including beliefs, prices, returns, borrowing constraints, and income streams, need not coincide with those of rentiers or with each other. Moreover, we do not require that households' beliefs about the future be objective. All that matters is that $\boldsymbol{\pi}(\cdot|\tau)$ is the lottery that rentiers in cohort τ believe they face, whether or not those beliefs are generated by a rational expectations equilibrium.

4.2 Rentiers with Risk-Free Labor Income

For simplicity, we define rentier households as those with zero exogenous income: $y(s^j|\tau) = 0$ for every history s^j. What matters, however, is that the reference group bears negligible idiosyncratic risk, not that its exogenous income is literally zero. If exogenous income streams are risk-free, pledgeable, and spanned by traded assets, then they can be capitalized and added to financial wealth.

To see this, consider households with nonzero, time-varying, but risk-free exogenous labor income, $y(s^j|\tau) = y(j|\tau)$. For example, pensioners with defined benefits, tenured professors, and public-sector workers with stable contracts may bear little idiosyncratic labor income risk. To include these households in the set of rentiers, assume that these income streams are risk-free, pledgeable, spanned by available bonds, and not subject to binding ad-hoc borrowing constraints.

For concreteness, suppose the first-period budget constraint, previously (1), is now

$$\sum_{n\in N} p_n(s^0|\tau)c_n(s^0|\tau) + \sum_{k\in K} a_k(s^0|\tau) + \sum_{l=1}^{J} A_l(s^0|\tau) = w + y(0|\tau), \tag{15}$$

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

$$\sum_{n\in N} p_n(s^j|\tau)c_n(s^j|\tau) + \sum_{k\in K} a_k(s^j|\tau) = \sum_{k\in K} R_k(s^j|\tau)a_k(s^{j-1}|\tau) + A_j(s^0|\tau)R(j|\tau) + y(j|\tau). \tag{16}$$

For simplicity of notation, we assume these bonds are available only in the initial period. Allowing households to trade these bonds after the initial period does not alter the result.

Under these assumptions, the household's problem is isomorphic to that of a rentier with zero exogenous income and augmented initial wealth

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

Thus, households with risk-free exogenous labor income and no binding ad-hoc borrowing constraints can be treated as rentiers whose wealth includes the present discounted value of that income.

4.3 Non-Time-Separable Non-Homotheticities

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

$$U = \frac{1}{1 - 1/\theta} \mathbb{E}_0 \sum_{j=0} \beta^j C_j^{1-\frac{1}{\theta}}, \quad \text{where} \quad C_j = \left[\sum_n \omega_n^{\frac{1}{\gamma}} \left[\frac{c_{nj}}{C_j^{\epsilon_n}} \right]^{\frac{1-\gamma}{\gamma}} \right]^{\frac{\gamma}{\gamma-1}}. \tag{17}$$

Under these preferences, within-period non-homotheticities are governed by the period-specific aggregator C_j rather than by lifetime utility U. Hence, these preferences are not time separable in the sense of Definition 3. The question is whether this violation is quantitatively important for our sufficient-statistic procedure.

The main lesson is that it need not be. In Appendix D.2, we show that, for rentiers, the sufficient-statistic formula in Proposition 4 continues to hold up to an additional error term. This error term is scaled by the first-period Frisch elasticity of intertemporal substitution.[16] Therefore, when intertemporal substitution is small, the approximation error from imposing time separability is also small. Intuitively, if households are reluctant to substitute consumption across periods, then misspecifying how within-period Engel curves interact across time has only a limited effect on the intertemporal welfare calculation.

The error term has a transparent economic interpretation and consists of two components. The first component compares a covariance in the present with a net-present-value average of analogous future covariances. These covariances are between the slope of Engel curves, disciplined by $\epsilon_n/\mathbb{E}_B[\epsilon_n]$, and changes in relative goods prices. Thus, a nonzero covariance between Engel slopes and price changes is not enough, by itself, to generate a large error. What matters is whether this covariance is stronger in the present

[16]The Frisch EIS measures intertemporal substitution holding the marginal utility of wealth fixed.

than along the future path, after applying the appropriate net-present-value weights. If the same goods whose prices rise today are also the goods whose prices rise in the future, in roughly the same Engel-slope-weighted sense, this component is small because the present and future covariances largely offset each other.

The second component depends on the covariance between variation in the Frisch EIS across future states and changes in the real interest rate. This component is small when the Frisch EIS does not vary sharply across dates and states. As noted above, both components are scaled by the first-period Frisch EIS.

A similar logic applies to non-rentiers. In the limiting case in which the first-period Frisch EIS is zero, preferences become Leontief across periods, so $C_j = C$ for all j, and the common value C is a monotone function of lifetime utility. In that limit, the budget-share matching argument in Proposition 5 applies exactly. Away from the limit, the argument is approximate.

Section 5 evaluates the quantitative importance of this approximation as part of a broader simulation exercise. In our calibration, the sufficient-statistic procedure continues to track true welfare closely across several outcomes: welfare levels, responses to idiosyncratic income shocks, and responses to price transitions. The formal propositions and the exact error expression are collected in Appendix D.2.

5 Performance in Simulated Data

Before applying the method to the PSID, we evaluate its performance in simulated data. We simulate agents whose income processes are disciplined by PSID household labor-income data, estimate their dynamic money-metric utility using our sufficient-statistic procedure, and compare the estimates with the true money-metric values computed from the solved model. We also compare our estimates to a net-present-value (NPV) benchmark, as in Basu et al., 2022; Fagereng et al., 2022; Del Canto et al., 2023; Pallotti et al., 2024, which values current financial assets plus expected future labor income discounted at the risk-free rate.

We organize the exercises in three steps. First, we study the benchmark case in which preferences satisfy Definition 3 and the researcher observes true rentiers, who earn no labor income. Second, we remove true rentiers from the estimation sample and instead use households with high financial-wealth-to-total-wealth ratios, mirroring the empirical implementation in Section 6.2. Third, we violate time separability by simulating agents with the preferences in Comin et al. (2021). Across the exercises, we study welfare levels, welfare changes after idiosyncratic income shocks, and welfare changes after a large price transition.

Because this is a Monte Carlo exercise, we use the fully solved structural model to compute the true welfare object against which the estimator is evaluated. The estimator itself does not use the structural value function. It uses the same objects as our empirical procedure: expenditure shares, prices, observed wealth for rentiers or rentier-like households, and the EIS.

5.1 Benchmark: Time-Separable Preferences and True Rentiers

We start with the benchmark case in which preferences are time separable and the researcher observes a population of true rentiers. Suppose preferences take the infinite-horizon analogue of Example 1:

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

where σ is the EIS, γ controls the elasticity of substitution between goods within each period and state, ϵ_n shapes the wealth elasticity of good n, and ω_n is a share parameter. Consumers face sequential budget constraints

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

where R is the risk-free return and y_j is an individual-specific risky income process. Consumers also face an exogenous ad-hoc borrowing constraint $a_j \geq -X$.

Calibration. For the baseline welfare-level and idiosyncratic-income-shock exercises, we abstract from aggregate risk and keep prices and returns constant over time. The 12,000 simulated non-rentiers are assigned to one of six latent income groups estimated from PSID household labor-income data. Each group has a quarterly log labor-income AR(1) process, with group-specific mean, persistence, and shock volatility; the details are described in Appendix E.1. Households can borrow up to five times their type's average quarterly income. We define total wealth in the simulation as current financial wealth plus the net-present value of expected future labor income; for true rentiers, total wealth coincides with financial wealth. The annual risk-free rate is 3%, so $R = 1.0075$. The EIS is set to 0.1, following the estimates of Best et al. (2020); this is also the value we use in the empirical application. Appendix E.3 repeats the exercises with an EIS close to one.[17] We set the intratemporal elasticity of substitution parameter to $\gamma = 0.25$. Our calibration has three goods within each period, with $\epsilon = (-0.6, -0.3, 0)$, and the quarterly discount factor is $\beta = 0.99$.

We stand at time $\tau = 0$, with agents in the ergodic distribution. The simulated economy contains true rentiers, who earn no labor income, and non-rentiers, who face risky labor income and borrowing constraints. By Definition 1, the true dynamic money-metric value of a simulated household with state $(w, \boldsymbol{y})$ is $m(\tau, w, \boldsymbol{y}|\tau)$. For true rentiers, this is simply current financial wealth. For non-rentiers, the true money-metric value is the amount of financial wealth that delivers the same value as the non-rentier's current value function when labor income is set to zero.

We begin with the level of money-metric welfare. To estimate welfare for non-rentiers, we apply Proposition 5. In practice, this means using true rentiers to learn the relationship between budget shares and current-date money-metric values at date τ. Since rentiers have no labor income, their current-date money metric equals their financial wealth. We regress rentier asset holdings on a flexible polynomial of log budget shares, where budget shares are the fractions of total expenditure spent on each good. We then use the fitted relationship to impute the money-metric value of each non-rentier from their budget shares. The exercise asks whether this imputed value is close to the true amount of rentier wealth that would make the household indifferent to its risky-income problem.

We then study welfare changes after idiosyncratic income shocks. We estimate each non-rentier's money-metric welfare at $\tau = 0$ and again at $\tau = 1$, after new individual income shocks are realized, and compare the estimated change with the true change implied by the value function. This exercise is demanding because welfare can move substantially from one period to the next for households that receive large idiosyncratic income shocks, especially when they are close to the borrowing constraint.

Finally, we study an unexpected price transition. At $\tau = 1$, households learn that goods 1 and 2 have been hit by a price shock and that prices will continue adjusting over future periods until they reach a new steady-state price vector. We use only the expenditure shares observed before the shock and immediately after households learn about the shock to estimate the welfare change. This is possible because forward-looking households adjust their consumption choices immediately, so their post-shock expenditure shares embed information about the future price path they now face.

[17] We provide this robustness check to show that the method continues to work well even when the EIS is close to one. The only requirement for the method to perform correctly is that the EIS is known and not equal to one.

Results. Since preferences satisfy Definition 3 and the estimation sample contains true rentiers, Proposition 5 implies that the error of our method should be zero. Figure 1a shows that the absolute value of the log error in level money-metric values using our method are indeed very close to zero.

For comparison, Figure 1b shows the much larger errors associated with the net-present-value (NPV) benchmark. The NPV benchmark is defined as:

$$m^{NPV}(\tau, w, y \mid \tau) \equiv Ra_0 + \sum_{j=0}^{\infty} R^{-j} E_0[y_j],$$

where a_0 is the household's beginning-of-period asset position, R is the gross risk-free return, Ra_0 is financial wealth after returns are realized, and $E_0[y_j]$ is expected labor income j periods ahead. This benchmark values current financial assets plus expected future labor income discounted at the risk-free rate. In our implementation, we compute this NPV measure using the true conditional income process from the simulation. If financial markets were complete, then this object would coincide with the dynamic money metric. With incomplete markets, however, the NPV-based measure does not account for the welfare cost of uninsurable idiosyncratic risk or the value of liquidity when borrowing constraints bind. It therefore tends to overstate welfare for non-rentiers, especially for households with little financial wealth relative to risky human wealth.

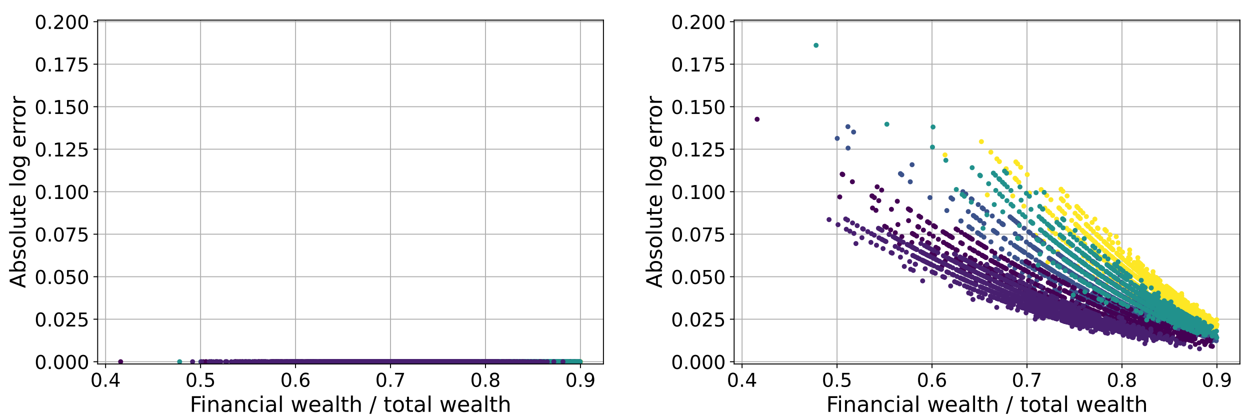

(a) Our method.

(b) NPV approach.

Figure 1: Absolute log error in level money-metric utility for non-rentiers.

The next exercise asks whether the method also captures welfare changes after idiosyncratic income shocks. Figure 2a and Figure 2b display the error in the log change in money-metric utility due to idiosyncratic income shocks for our method and the NPV approach. In practice, this means estimating welfare at $\tau = 0$ and again at $\tau = 1$ after the new idiosyncratic shocks, and comparing each agent's estimated welfare change with the true change. These true welfare changes can be large from one period to the next, especially for financially constrained households that experience large income shocks. Our method nevertheless estimates these changes very precisely. Errors are smaller for the NPV approach in changes than in levels, but they can still be large for households that experience large income shocks and are financially constrained.

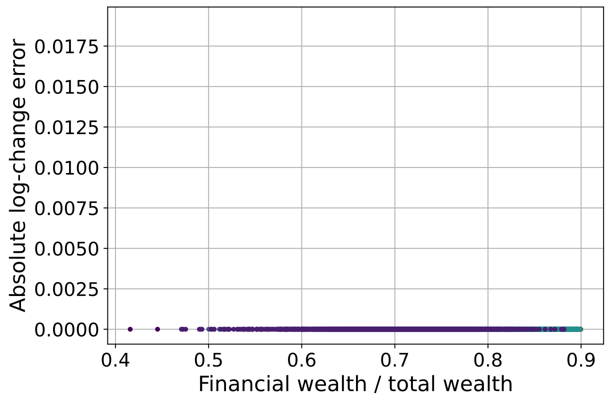

(a) Our method.

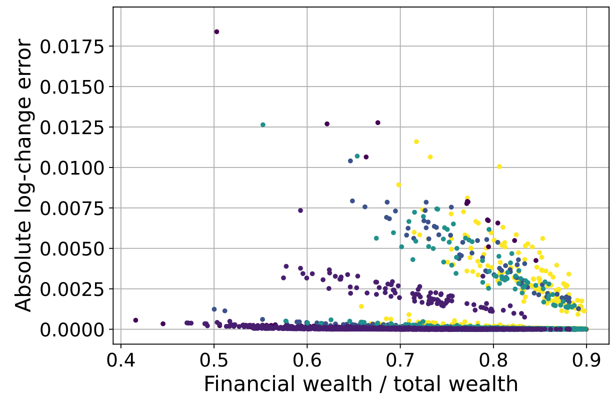

(b) NPV approach.

Figure 2: Absolute log error in changes in money-metric utility for non-rentiers.

We now turn to the price-transition exercise. This exercise asks how well our method tracks welfare changes when agents face a large unexpected price shock at $\tau = 1$ that continues to affect prices for several periods. The true welfare effect is obtained by comparing each household's dynamic money-metric utility immediately after the shock with the value for the same household absent the shock, holding fixed its realized assets and income. To apply the method, we only need households' expenditures at $\tau = 0$ and $\tau = 1$. Even though the shock has future consequences, the welfare-relevant information is reflected in spending patterns on impact. Figure 3 shows that the estimated welfare changes track the true ones closely. The small errors are due to the discrete-time approximation to continuous time.

5.2 Using Rentier-Like Households Instead of True Rentiers

The benchmark exercise assumes that true rentiers are observed. In practice, however, there are very few households with literally zero risky labor income, and in many datasets there may be no pure rentiers at all. In our empirical implementation, we therefore use households whose financial wealth is large relative to total wealth as rentier-like households. These households are not necessarily wealthier than non-rentiers; rather,

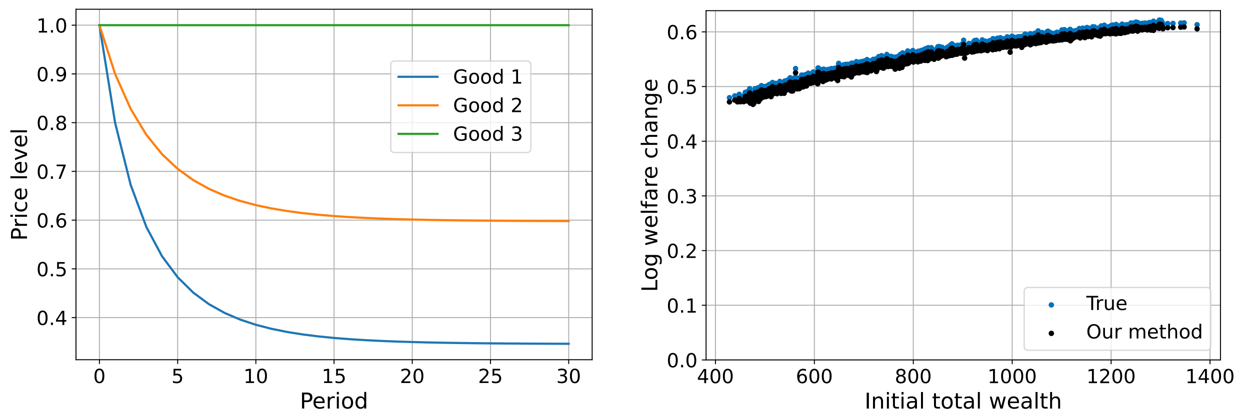

(a) Price paths.

(b) Welfare changes for non-rentiers.

Figure 3: Price paths and welfare effects under the unexpected price transition. Panel (a) displays the price transition. Panel (b) compares true welfare changes with estimates from our method.

they differ in the composition of their wealth, with a larger share held in financial assets.

We reproduce that problem in the simulation by excluding true rentiers from the estimation step and using only 'contaminated rentiers": non-rentiers whose financial wealth is at least 90% of total wealth. This mirrors the classification rule used in the empirical application in Section 6.2.[18]

The calibration is the same as in the benchmark time-separable exercise. The only change is the estimation sample used to learn the relationship between budget shares and wealth. Instead of fitting this relationship on true rentiers, we fit it on contaminated rentiers and then use the resulting coefficients to estimate money-metric welfare for the rest of the non-rentier population.

Results. Figure 4a shows the absolute value of the log error in money-metric levels, while Figure 4b shows the log error in estimated money-metric changes after idiosyncratic shocks. The errors increase relative to the true-rentier benchmark, as expected, but remain an order of magnitude smaller than the errors from the NPV approach.

We also repeat the price-transition exercise using contaminated rentiers. As before, the full model is used only to compute the true welfare change. Our sufficient statistics method uses the sample of contaminated rentiers rather than true rentiers as the reference group. Figure 5 shows that the contaminated-rentier implementation remains accurate, despite the fact that the reference group is no longer made up of true rentiers.

[18] Appendix E.3 repeats the contaminated-rentier level, idiosyncratic-shock, and price-transition exercises with an EIS close to one.

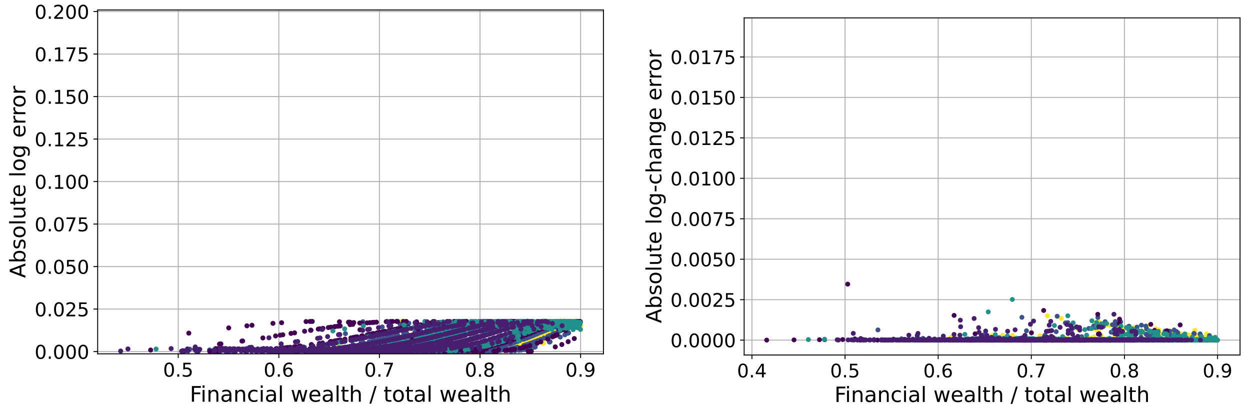

(a) Level errors.

(b) Change errors.

Figure 4: Absolute log errors for our method using contaminated rentiers.

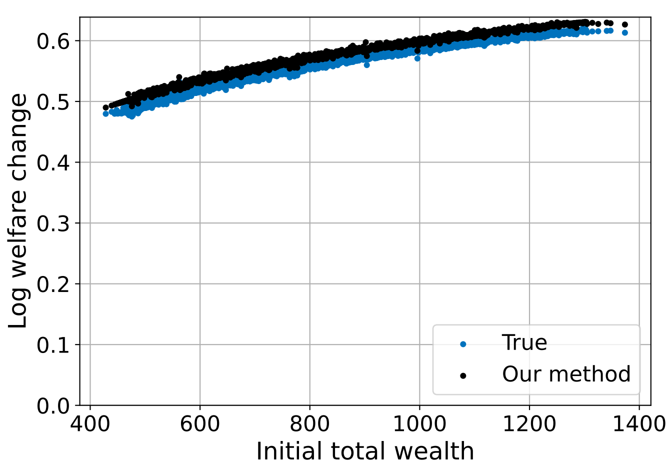

Figure 5: True welfare changes and estimated welfare changes under the unexpected price transition, using contaminated rentiers.

5.3 Violating Time Separability

We finally assess the performance of our method when preferences are not time separable. We use the preferences in (17), where within-period non-homotheticities are governed by the current period consumption aggregator C_j rather than lifetime utility U. These preferences violate Definition 3. The exercise asks whether our sufficient-statistic procedure continues to perform well when this key assumption fails, but the Frisch EIS is small, as suggested by the theoretical discussion in Section 4.3.

We keep the parametrization as close as possible to the time-separable benchmark. We set the intratemporal elasticity of substitution parameter to $\gamma = 0.25$ and set the vector of ϵ's to the values in Comin et al. (2021): $\epsilon = (1, 0.8, 0.2)$. The Frisch EIS is not constant under these preferences, but it depends strongly on the parameter θ. We set $\theta = 0.1$, which implies that the Frisch elasticity of intertemporal substitution, σ_0, varies between 0.06 and 0.09 in the cross-section of simulated households, using Equation (24).[19] This is within the range of values implied by the estimates from Best et al. (2020).

Results. Figures 6 and 7 repeat the level and idiosyncratic-shock exercises under Comin et al. preferences. These figures use contaminated rentiers, meaning agents whose financial wealth represents more than 90% of total wealth, to learn the relationship between budget shares and welfare. Our method again outperforms the NPV approach, especially for lower-income and more constrained households.

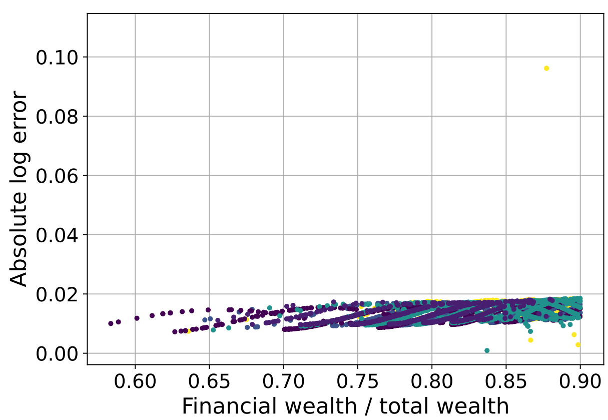

(a) Our method.

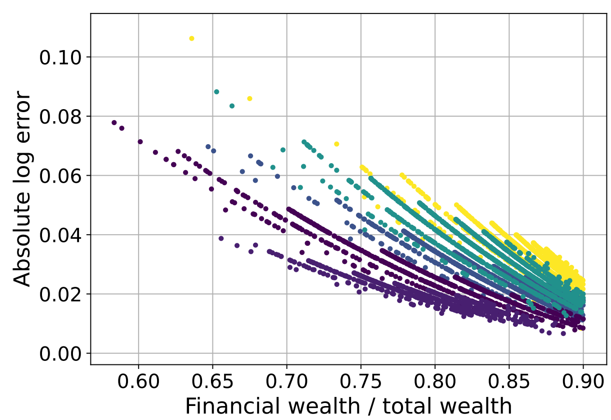

(b) NPV approach.

Figure 6: Absolute log error in level money-metric utility for non-rentiers under Comin et al. preferences.

[19]Figure 15 in the appendix plots the cross-sectional distribution of σ_0 for simulated non-rentiers under this calibration.

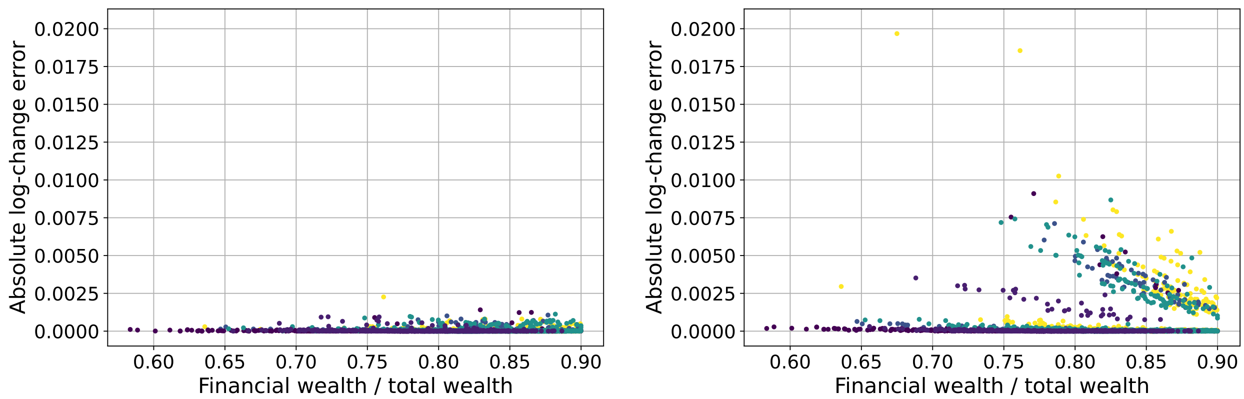

(a) Our method.

(b) NPV approach.

Figure 7: Absolute log error in changes in money-metric utility for non-rentiers under Comin et al. preferences.

Figure 8 repeats the same price transition under Comin et al. (2021) preferences, using the true-rentier implementation.[20] This isolates the error due to non-time-separability, rather than the additional error from using contaminated rentiers (the magnitude the errors are similar if we use contaminated rentiers). The estimated welfare changes track the true changes closely.

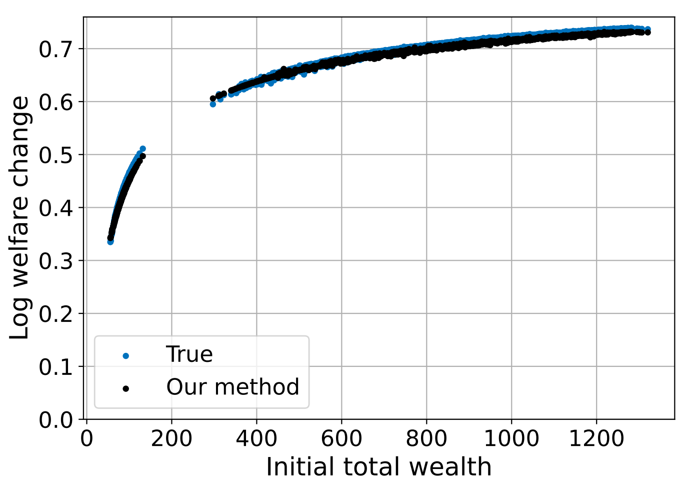

Figure 8: True welfare changes and estimated welfare changes under the unexpected price transition with Comin et al. preferences.

[20]The visual gap in the figure comes from using a finite number of latent income groups. With six groups, the ergodic distribution has a range of financial-asset values with no simulated agents. Simulating more income groups would fill in this region, but is not necessary for the exercise and would not affect the results.

6 Illustrative Application to US Data

We now apply our method to estimate the dynamic money metric using US household and price data. In the next section, we use these estimates to study the welfare effects of shocks.

6.1 Data

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

Since our estimates are at the household level, the household data must be accurate, but the sample need not be representative of the underlying population in terms of sampling frequency. Therefore, we use the raw PSID data without sampling weights. The PSID contains repeated cross-sectional data on household expenditures by category, household-level balance sheets, household income, and demographic information. To account for the effect of age on preferences, including planning horizons, we condition our results on decade of life.[21]

The PSID includes household expenditure surveys broken down into seven categories. A major omission is the user cost of owner-occupied housing.[22] To remedy this, we impute owner-occupied housing costs by matching homeowners in each period to renters with similar observable characteristics and spending behavior. Specifically, we predict rental expenditures using a regression, estimated on renters in the same period, that includes household characteristics and spending behavior. This procedure is theoretically justified by Proposition 5. In the final year of our sample, 2019, the PSID asked homeowners to report the rental value of their property if they were to rent it out. We use the answers to this question to validate our imputation procedure. Regressing surveyed housing costs on our imputed measure of housing costs, with both variables expressed relative to current expenditures, yields a coefficient of 1.03 and an R^2 of 0.59. See Appendix F.1 for details.[23]

We combine the BLS price data with the PSID expenditure survey using a correspondence between PSID spending categories and Consumer Price Index (CPI) categories.

[21]Ideally, with enough observations, we could treat each age separately. Similarly, with more data, we could split the sample along other observed characteristics that may influence preferences, such as gender, household size, or location.

[22]See Chodorow-Reich et al. (2025) for a discussion of how debt-financed durable goods prices, like housing, should be incorporated into price indices.

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

Appendix F.1 provides details on the construction of the variables.

6.2 Constructing Wealth and Classifying Rentiers

To apply Proposition 4, we need to identify a sample of rentier households. We do so by first constructing a proxy for total wealth for all households in the sample. Our proxy equals financial wealth—net asset value including home equity and defined-contribution pension savings—plus the present discounted value of predicted labor and transfer income.

To calculate the present discounted value of labor and transfer income, we predict each household's expected lifetime income profile based on observed characteristics and discount the resulting flows using a real discount rate of 4%, following Catherine et al. (2022). Appendix F.1 details the construction of net assets and capitalized labor and transfer income.

We classify a household as a rentier if net financial assets constitute more than 90% of proxy total wealth. Results are similar if we use a 95% threshold. If the household head is unemployed and looking for a job, we exclude the household from the rentier sample.

Finally, to reduce the influence of outliers, we exclude from the rentier sample households whose net assets are in the top or bottom 5% of the rentier asset distribution. Since there are relatively few rentiers younger than 40 in our sample, we do not report estimates for households below age 40.

6.3 Dynamic Money Metric Wealth over Time

Figure 9 shows scatterplots of log consumption-to-wealth ratios against log wealth for rentiers in the first and last years of the sample. Both panels show a strong decreasing relationship between consumption-to-wealth ratios and wealth. The relationship is approximately loglinear, consistent with Straub (2019), who finds that the consumption-to-wealth ratio declines strongly with permanent income. Whereas static non-homotheticity is a cornerstone of the consumption literature, dynamic non-homotheticity, like the pattern depicted in Figure 9, is relatively understudied. As our estimates below show, in our data dynamic non-homotheticity is an order of magnitude more powerful than static non-homotheticity.

To compute money-metric values for rentiers using Proposition 4, we need to evaluate consumption-to-wealth ratios and static budget shares as functions of date and wealth for each age group. To do this, we non-parametrically regress log consumption-to-wealth ratios on log wealth, age group, and year using kernel regression, implemented with

Figure 9: Log consumption-to-wealth against log wealth for rentiers

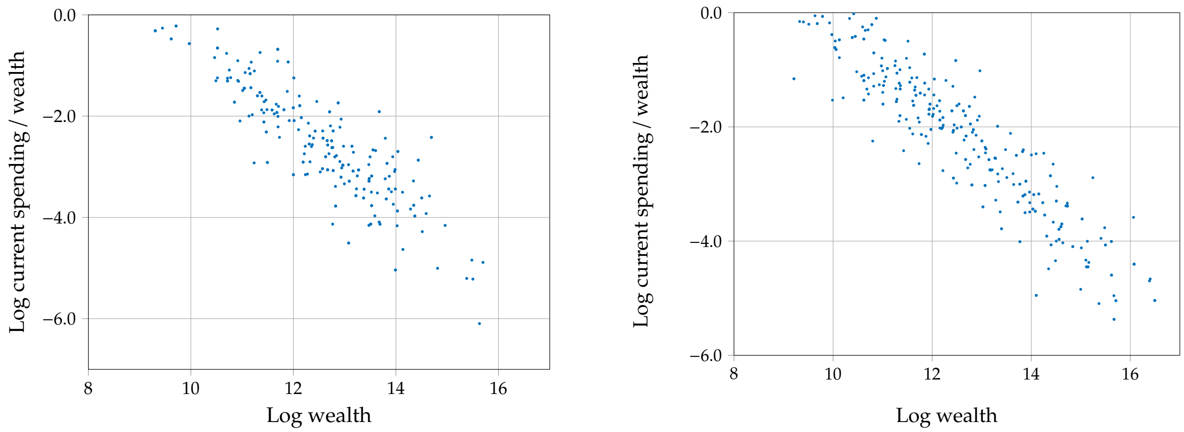

(a) 2005

(b) 2019

npregress in Stata. Similarly, for each good $i \in N$, we regress the log budget share on a quadratic function of age, log wealth, and year using LASSO. In both cases, smoothing parameters are chosen by cross-validation.[24]

We apply Proposition 4 to the estimated cross-sectional curves to recover money-metric utility as a function of date, wealth, and age group.[25] For illustration, we use the initial year, $\tau_0 = 2005$, as the base year, so that money-metric values map nominal wealth in each year τ 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 uniformly between 0.05 and 0.15 across different quartiles of age and income.

Best et al. (2020) do not estimate Hicksian, or compensated, elasticities. However, Slutsky's equation implies that, if consumption is a normal good, then the compensated intertemporal elasticity should be smaller in magnitude than the uncompensated one. We experimented with nearby values of σ, namely $\sigma = 0.2$ and $\sigma = 0.05$, and found that our results are not sensitive to this choice. Of course, the results are quantitatively sensitive to the assumption that σ is not close to one.[26]

[24]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 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.

[25]To recover the dynamic money metric, we need to solve the integral equation in Proposition 4. To do so, we use the "recursive" methodology described in Baqaee et al. (2024).

[26]Even though in this paper we do not estimate the compensated EIS, we note that in principle it can be

The left column of Figure 10 plots the money metric for each age group as a function of 2019 wealth, expressed in 2005 base prices. In other words, for each level of nominal wealth w in 2019, we use Proposition 4 to compute the amount of wealth in the 2005 environment that would make a rentier in the same age group indifferent to having wealth w in 2019. The confidence bands are calculated by bootstrap. Since there are fewer young rentiers, sampling uncertainty is higher for younger age groups. For comparison, we also plot a naive calculation that deflates nominal wealth in 2019 by official CPI inflation between 2005 and 2019. In all cases, dynamic money-metric wealth is higher than CPI-deflated wealth.

The right column of Figure 10 plots the dynamic annualized inflation rate, defined as the log difference between nominal wealth in 2019 dollars and money-metric wealth in 2005 dollars. CPI inflation over this period was 2% per year, but the dynamic inflation rate is always below 2%. Furthermore, dynamic inflation rates vary with wealth and age, highlighting the importance of non-homotheticity and lifecycle considerations.

To understand the pattern in Figure 10, we decompose the dynamic inflation rate, the deflator in (13), into a static component and a forward-looking component. Specifically, for a household with wealth w in 2019, the change in the ideal cost-of-living index between 2005 and 2019 is

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

where w_x^* ensures that we use compensated consumption-to-wealth ratios and budget shares. The first summand is a static measure of inflation. The second summand is expected future inflation relative to present inflation. If the second term is negative, then the price of the future bundle rises less than the price of the present bundle, so overall inflation is lower than static inflation.

As an example, Figure 11 shows this decomposition for the 60–69 year-old group. The static inflation term is not exactly the same as aggregate CPI because it weights 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 slight downward slope reflects the non-homotheticity of static preferences: static inflation is slightly higher for poorer households, consistent with

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 10: Conversion of 2019 dollars to 2005 dollars

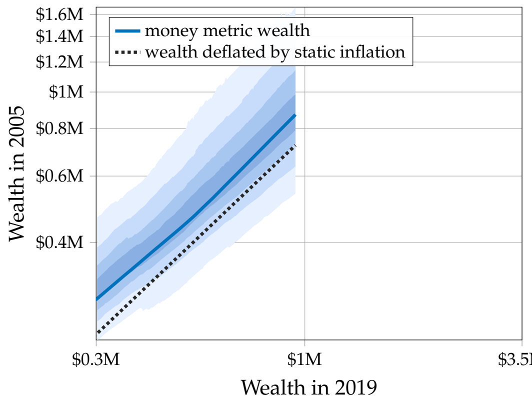

(a) money-metric 40 – 49 year olds

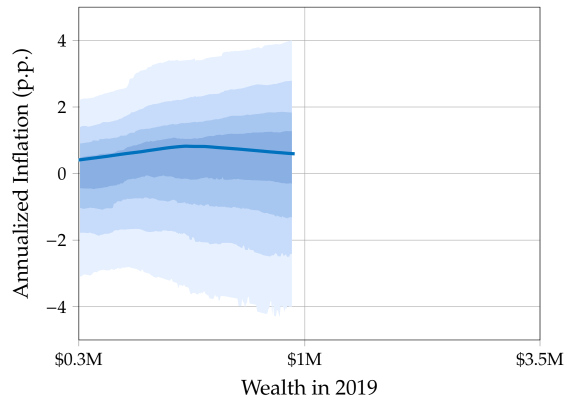

(b) annualized inflation 40 – 49 year olds

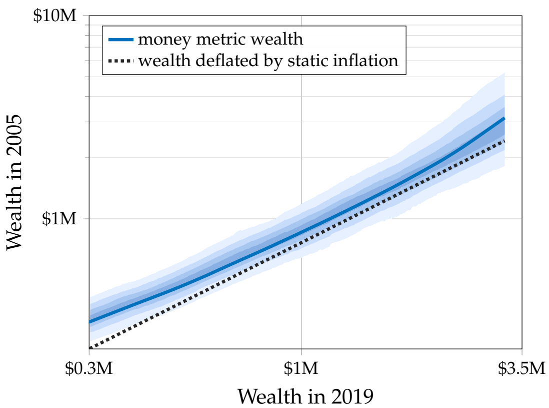

(c) money-metric 50 – 59 year olds

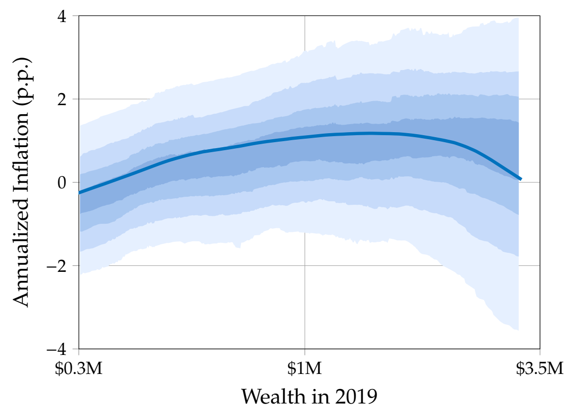

(d) annualized inflation 50 – 59 year olds

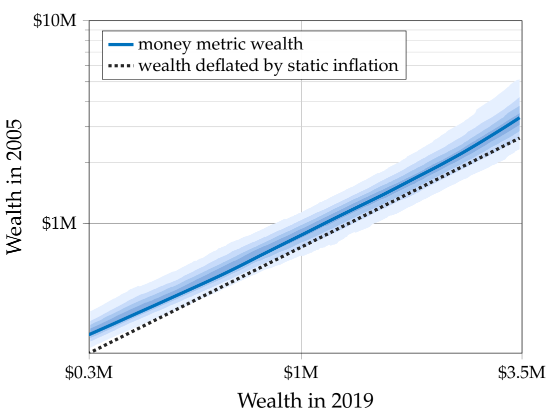

(e) money-metric 60 – 69 year olds

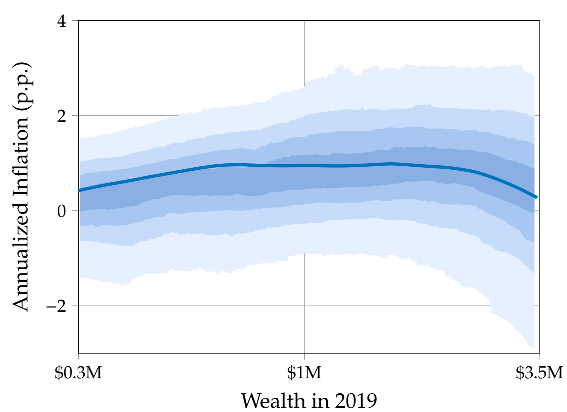

(f) annualized inflation 60 – 69 year olds

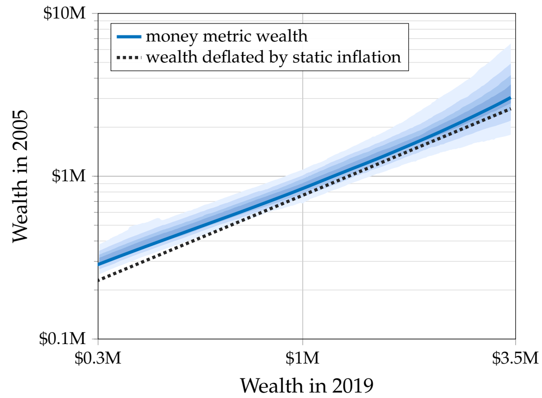

(g) money-metric 70 – 79 year olds

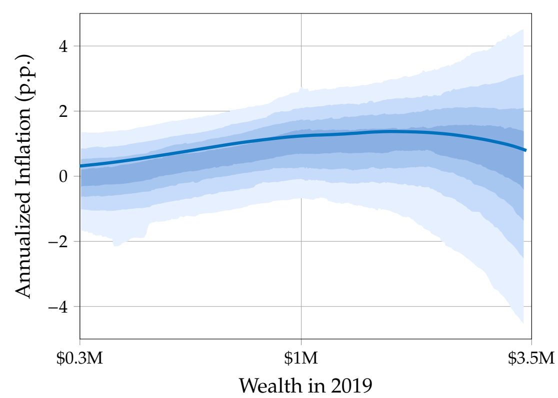

(h) annualized inflation 70 – 79 year olds

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

other studies of the U.S., such asJaravel and Lashkari (2024), which 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.[27]

The component corresponding to future inflation in Figure 11 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 lies below the static inflation line in Figure 11. 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-to-wealth ratio. Figure 12 plots both the compensated and uncompensated log change in the consumption-to-wealth ratio between 2005 and 2019 as a function of nominal wealth for 60–69 year olds. Since compensated consumption-to-wealth ratios rose for all wealth levels, dynamic inflation is less than

[27]There may be several reasons why the contribution of static non-homotheticity is so mild in our exercise. First, 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. Second, 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 with more disaggregated data, static inflation heterogeneity may be stronger, especially because within-category price changes can differ systematically across the quality ladder (Kaplan and Schulhofer-Wohl, 2017; Sangani, 2026).

static inflation.[28]

The uncompensated change in the consumption-to-wealth ratio is more positive than the compensated one. This is because there is a strong wealth effect: the consumption-to-wealth ratio declines as households become richer. Households with nominal wealth w in 2005 are on a higher indifference curve than households with the same nominal wealth in 2019 because of positive inflation. Therefore, the wealth effect implies that such households have higher consumption-to-wealth ratios in 2019 than in 2005, even if relative prices do not change. Changes in compensated consumption-to-wealth ratios, which are purged of wealth effects, are smaller and reflect only substitution effects. This figure underscores the importance of accounting for wealth effects when using consumption-to-wealth ratios to infer changes in relative prices.[29]

Figure 11: Decomposing inflation for 60-69 year olds

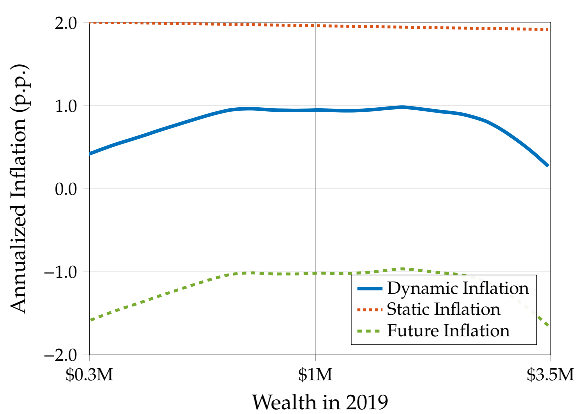

Figure 12: Changes in log consumption-to-wealth ratios for 60 – 69 year olds

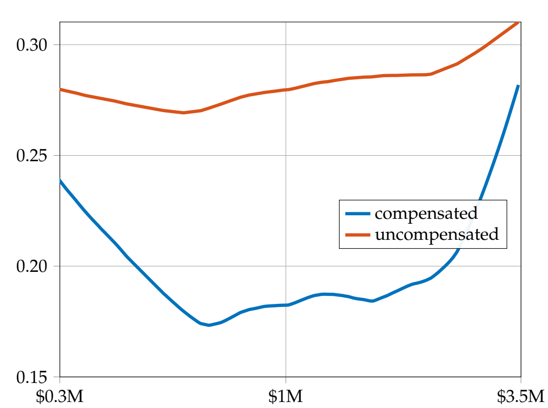

Our method does not identify which future prices or beliefs are responsible for the patterns in Figure 10. However, the differences in the dynamic measure of inflation need not

[28]In Appendix F.2, we also plot the aggregate consumption-to-wealth ratio for our sample. The aggregate patterns over time are very different: whereas the compensated and uncompensated consumption-to-wealth ratios rose between 2005 and 2019 holding fixed utility and nominal wealth, respectively, the aggregate consumption-to-wealth ratio fell between these two dates. This is not surprising. For example, if there is positive economic growth, so households reach higher indifference curves over time, then they move down the consumption-to-wealth Engel curve and the aggregate consumption-to-wealth ratio declines, even if the compensated and uncompensated consumption-to-wealth ratios rise.

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

be caused by differential exposures to future goods prices alone. For example, dynamic inflation is lower than static inflation if expected future returns increase. Furthermore, even if all households are symmetrically exposed to future goods prices, the future component of inflation can differ across households because of differences in expected asset returns; see Fagereng et al. (2022). For example, if poor and rich rentiers, or young and old rentiers, rely differentially on returns to real estate, equities, or bonds to finance consumption, then changes in returns will differentially affect dynamic inflation rates for these households.

6.4 Non-Rentiers

We now turn to the remaining households: the non-rentiers. Proposition 4 does not apply directly to these households, so we recover their money-metric values using Proposition 5. For each date τ, we fit log money-metric values to static budget shares and age-group indicators for rentiers:

$$\log m(\tau, w_{h,\tau}, \mathbf{0}|\tau) = \log w_{h,\tau} = \boldsymbol{\alpha}'_{\tau} \boldsymbol{X}_{h,\tau} + \text{error}_{h,\tau}, \tag{18}$$

where h indexes the household, τ indexes the time period, and $\boldsymbol{X}_{h,\tau}$ includes budget shares and age-group indicators. We then use the predicted values from this regression to impute money-metric wealth for non-rentiers conditional on their budget shares, age, and date.[30] We then convert the imputed wealth into base-year dollars using the Proposition 4 deflator, exactly as for rentiers. Appendix F.2 provides figures of the money-metric values for non-rentiers.

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

[30] As a robustness check, we exclude outliers from the rentier set in regression (18). Specifically, if a potential rentier's predicted and measured total wealth differ substantially, as measured by a Cook's distance greater than one, then we exclude that household from the rentier set. The results are very similar.

7 Using the Measure in Reduced-Form Work

Having constructed money-metric welfare for both rentiers and non-rentiers, we now illustrate how the measure can be used as an outcome in reduced-form empirical work. Many policies and shocks affect households along several margins at once. For example, job loss affects current earnings, future earnings risk, borrowing constraints, consumption, and expectations; a monetary policy shock affects goods prices, labor-market outcomes, equity prices, house prices, bond prices, and other variables. A standard approach is to estimate the dynamic effect of the shock on each relevant margin and then aggregate the discounted present value of these effects using predicted pre-shock behavior. This is the approach taken by, for example, Davis and Von Wachter (2011) and Del Canto et al. (2023). This approach is transparent, but it requires the researcher to enumerate, measure, and forecast all the channels through which the shock affects welfare. Moreover, the estimate is a first-order approximation around perfect-foresight allocations.

Our measure provides a complementary approach. We first construct the dynamic money metric and then use it as the outcome variable. In this section, we use job loss as an example. The goal is not to provide a new identification strategy for the causal effects of job loss. Rather, it is to show how a dynamic welfare measure can summarize the many channels through which a shock affects households.

We implement this idea in the PSID. We regress log money-metric wealth on an indicator equal to one when the household head loses her job and reports searching for a new job in that period. To control in a simple way for confounds and selection, we include year fixed effects, demographic controls, and lagged log money-metric wealth.

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

The regression compares static budget shares across goods between households that lost their job and households that did not, conditional on the controls. In other words, the budget shares of households that lose their job differ from those that do not in a way that, had they been rentiers, would correspond to 20 log points lower total wealth. This exercise is not using the drop in total consumption following job loss as a sufficient statistic for welfare, as in Gruber (1997). Instead, it uses changes in the composition of

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 $\times \mathbf{1}(\text{age} \geq 60)$		0.180^{**}
		(0.083)
Lagged LHS	Yes	Yes
Controls	Yes	Yes
Observations	48,357	48,357

Notes: Bootstrapped standard errors in parentheses. Controls are year fixed effects, age group, marital status of head of household, industry, and education level.

spending across goods. The smaller coefficient for older household heads implies that these budget-share differences are less pronounced for older households.

8 Conclusion

This paper develops a sufficient-statistic approach for measuring forward-looking welfare in money-metric terms. The method extends the logic of static cost-of-living measurement to dynamic environments with uncertainty, incomplete markets, borrowing constraints, and non-homothetic preferences. Rather than specifying a full stochastic model of future prices, income, returns, beliefs, and constraints, the method uses current expenditure behavior and consumption-saving choices to infer the welfare-relevant value of the future. The key requirements are time separability, stable preferences conditional on observables, an estimate of the intertemporal elasticity of substitution, and a reference group of households with negligible idiosyncratic income risk.

We show in simulations disciplined by PSID income dynamics that the method tracks true welfare more accurately than net-present-value calculations, including when the reference group is only approximately free of idiosyncratic risk and when time separability is imperfect. Applying the method to PSID households from 2005 to 2019, we find dynamic cost-of-living indices below the static CPI benchmark and substantially more heterogeneity across wealth and age than implied by static inflation alone. We also illustrate how the resulting welfare measure can be used as an outcome in reduced-form work by studying the welfare losses associated with job loss. The goal is not to replace a correctly specified structural model when one is available, but to provide a practical forward-looking welfare measure for settings where the realistic alternatives are static deflators or net-present-value calculations.

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., 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.

Braxton, J. C., K. Herkenhoff, J. Rothbaum, and L. Schmidt (2025). Changing income risk across the US skill distribution: Evidence from a generalized kalman filter. *American Economic Review 115*(12), 4438–4475.

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).

Chodorow-Reich, G., E. Glaeser, X. Jaravel, and A. Lipton (2025). Homeownership and the cost of living. *Manuscript, Harvard University*.

Comin, D., D. Lashkari, and M. Mestieri (2021). Structural change with long-run income and price effects. *Econometrica 89*(1), 311–374.

Cooper, D., K. E. Dynan, and H. Rhodenhiser (2019). Measuring household wealth in the panel study of income dynamics: The role of retirement assets.

Costa, D. L. (2001). Estimating real income in the united states from 1888 to 1994: Correcting cpi bias using engel curves. *Journal of political economy 109*(6), 1288–1310.

Davis, S. J. and T. M. Von Wachter (2011). Recessions and the cost of job loss. Technical report, National Bureau of Economic Research.

Deaton, A. and J. Muellbauer (1980). *Economics and consumer behavior*. Cambridge university press.

Del Canto, F. N., J. R. Grigsby, E. Qian, and C. Walsh (2023). Are inflationary shocks regressive? a feasible set approach. Technical report, National Bureau of Economic Research.

Epstein, L. G. and S. E. Zin (1989). Substitution, risk aversion, and the temporal behavior of consumption. *Econometrica 57*(4), 937–969.

Fagereng, A., M. Gomez, E. Gouin-Bonenfant, M. Holm, B. Moll, and G. Natvik (2022). Asset-price redistribution. Technical report, Working Paper.

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

Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review 87*(1), 192–205.

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

Hobijn, B. and D. Lagakos (2005). Inflation inequality in the united states. *Review of Income and Wealth 51*(4), 581–606.

Huggett, M. and G. Kaplan (2016). How large is the stock component of human capital? *Review of Economic Dynamics 22*, 21–51.

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

Jaravel, X. (2019). The unequal gains from product innovations: Evidence from the us retail sector. *The Quarterly Journal of Economics 134*(2), 715–783.

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.

Kaplan, G. and S. Schulhofer-Wohl (2017). Inflation at the household level. *Journal of Monetary Economics 91*, 19–38.

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. (2003). Macroeconomic priorities. *American economic review 93*(1), 1–14.

Nakamura, E., J. Steinsson, and M. Liu (2016). Are chinese growth and inflation too smooth? evidence from engel curves. *American Economic Journal: Macroeconomics 8*(3), 113–44.

Pallotti, F., G. Paz-Pardo, J. Slacalek, O. Tristani, and G. L. Violante (2024). Who bears the costs of inflation? euro area households and the 2021–2023 shock. *Journal of Monetary Economics 148*, 103671.

Pollack, R. A. (1975). The intertemporal cost of living index. In *Annals of Economic and Social Measurement, Volume 4, number 1*, pp. 179–198. NBER.

Reis, R. (2005). A dynamic measure of inflation. Working Paper 11746, National Bureau of Economic Research.

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

Sangani, K. (2026). Cheapflation cycles. Working Paper 35235, National Bureau of Economic Research.

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

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