Consumption Wedges: Measuring and Diagnosing Distortions *

Sasha Indarte,[†] Raymond Kluender,[‡] Ulrike Malmendier,[§] and Michael Stepner[¶]

March 19, 2025
(for latest draft, click here)

Abstract

Deviations from canonical consumption-savings models have been attributed to a wide range of distortions, including financial constraints and behavioral preferences. We develop a new sufficient statistics approach to measure the impact of such distortions on consumption as a *micro-level wedge* between actual consumption and a counterfactual "frictionless" consumption. Our approach is applicable to a broad class of models and, unlike standard wedge measurement approaches, does not rely on an assumption of full-information rational expectations (FIRE), which allows us to isolate the influence of frictions and behavioral preferences from deviations from FIRE. Since different frictions imply different properties of wedges, the estimates of wedges can be used as a diagnostic to distinguish between models. To implement this approach, we field a new survey of economic beliefs, which we link to bank account transactions data for a population of predominantly middle-income US consumers with low liquid wealth. We find that consumption choices are significantly distorted both upward and downward. The median wedge is 52% of frictionless consumption in absolute value, with 75% having negative wedges (under-consuming) and 25% having positive wedges (over-consuming). Since financial constraints only generate negative wedges, additional or alternative distortions (such as present bias or consumer inertia) are necessary to rationalize the consumption decisions of low-liquidity households.

*We are grateful to Manish Amde, Stephan Floyd, Angela Hung, Ben Larocco, Ray Sin, and Yang Tian of EarnIn for their help compiling the data and executing the survey required for this project. For helpful comments and feedback, we thank David Berger, John Beshears, Gideon Bornstein, Taha Choukhmane, Francesco D'Acunto (discussant), Tim de Silva, Marty Eichenbaum, Ben Enke, Joao Gomes, Yuriy Gorodnichenko, Fatih Guvenen, Ben Keys, Dirk Krueger, Igor Livshits, Emi Nakamura, Matt Notowidigdo, Michael Roberts, Nick Roussanov, and Jón Steinsson. Nikki Azerang, Chazel Hakim, Clint Hamilton, Eleanor Jenke, Emma Lee, and Karan Makkar provided outstanding research assistance. We gratefully acknowledge financial support from the Wharton School's Jacobs Levy Center.

[†]The Wharton School, University of Pennsylvania, aindarte@wharton.upenn.edu

[‡]Entrepreneurial Management Unit, Harvard Business School, rkluender@hbs.edu

[§]Department of Economics and Haas School of Business, University of California Berkeley, NBER, and CEPR, ulrike@econ.berkeley.edu

[¶]Department of Economics, University of Toronto and Opportunity Insights, michael.stepner@utoronto.ca

1 Introduction

Financial constraints play a central role in theories of consumer behavior in macroeconomics and household finance (e.g., Kaplan and Violante, 2014). A key motivation for the focus on financial constraints is empirical evidence that households' marginal propensities to consume (MPCs) are much larger than predicted by canonical frictionless models, especially among low-liquidity households (e.g., Johnson, Parker and Souleles, 2006). However, alternative distortions— such as present bias, consumption adjustment costs, and bounded rationality—can also generate large MPCs (Lee and Maxted, 2023; Maxted, Laibson and Moll, 2024; Beraja and Zorzi, 2024; Ilut and Valchev, 2023). Importantly, these alternative models imply different distributional and aggregate consequences of fiscal policy, monetary policy, and business cycle fluctuations compared to models featuring only financial constraints.[1] To better guide theory and ultimately policy, more evidence is needed on which distortions influence consumption and by how much.

To this end, we develop a new sufficient statistics approach to measure the impact of distortions as *consumption wedges* between actual consumption and a counterfactual "frictionless" consumption. Consumption wedges quantify the total impact of distortions on consumption, including both frictions (such as borrowing constraints) and behavioral preferences that result in "as if" constrained behavior (e.g., present bias or bounded rationality). Because alternative distortions have different predictions for the properties of consumption wedges—such as their signs, correlates, and responses to shocks—they can be used as a diagnostic to discriminate between models of consumer behavior. Unlike existing wedge-based analyses (e.g., Chari, Kehoe and McGrattan, 2007; Berger, Bocola and Dovis, 2023), we measure wedges using subjective expectations data, which allows us to avoid assuming full-information rational expectations (FIRE) and to isolate the influence of distortions on consumption separately from the influence of deviations from FIRE. We also differ by measuring wedges at the household (micro) level rather than the aggregate (macro) level, enabling us to document heterogeneity in wedges.

We implement our approach using new data on consumer expectations linked to administrative transactions data for a population of predominantly middle-income US consumers with low liquid wealth. We find that distortions are an important determinant of consumption for households in our sample: the median absolute value wedge is 52% of frictionless consumption. Moreover, we find that 75% of consumers have negative wedges (under-consumption) and 25% have positive wedges (over-consumption). The results also highlight the value of studying micro-

[1] For example, Lee and Maxted (2023) shows that adding present bias to a model with financial constraints leads to large consumption responses to both positive and negative wealth shocks. With only financial constraints, the response to negative shocks is much larger compared to positive shocks.

level wedges, as opposed to only macro-level wedges. Small aggregate wedges could potentially mask significant cross-sectional heterogeneity. Indeed, the average wedge is less than half of the median absolute value wedge (-19% versus 52%). Additionally, the mix of positive and negative wedges implies that financial constraints—the primary friction used to explain high MPCs—are not the dominant distortion for at least 25% of our sample. This is because financial constraints can only distort consumption downward and thus only create negative wedges. Present bias and sources of consumer inertia (such as adjustment costs or bounded rationality) are plausible additional or alternative frictions that can rationalize our results. Taken together, our results indicate that distortions, including those beyond financial constraints, play an important role in shaping the consumption of low-liquidity households.

We begin by characterizing frictionless consumption in a stylized model. A household chooses consumption and saving (or borrowing) via a risky asset given their realized income and wealth. Under the benchmark model, households face no frictions (borrowing constraints, consumption adjustment costs, etc.) and have standard preferences.[2] Differently from approaches in the prior literature, our benchmark model allows household beliefs to flexibly deviate from FIRE. As a result, the wedge between a household's actual and "frictionless" consumption quantifies the degree to which households over- or under-consume due to any frictions or non-standard preferences, isolating their impact from the influence of any deviations from FIRE.

In the benchmark model, frictionless consumption is characterized by an Euler equation and a budget constraint. Under a first-order approximation, frictionless consumption is a log-linear function of net worth (divided by income) and beliefs about future nominal income growth, nominal interest rates, and inflation. The coefficients in this formula depend on two preference parameters (the discount factor β and the inverse intertemporal elasticity of substitution γ) and the choice of two approximation points (the ratios of wealth to income and consumption to income). Our characterization relies on the Euler equation and budget constraint being necessary conditions for optimality, but they need not be sufficient. As a result, our formula for frictionless consumption is robust to a variety of model extensions, such as additional household choices (e.g., labor supply or default) and a richer asset environment (including the case of complete markets). We intentionally exclude frictions and behavioral preferences from our benchmark; as a result, the wedge between actual and frictionless consumption quantifies the total net impact on consumption of *all* distortions (frictions and behavioral preferences) directly affecting consumption.

To empirically measure micro-level consumption wedges, we use new data linking subjec-

[2] We use the term "standard preferences" to refer to preferences that are time consistent, time separable, strictly increasing, and strictly concave.

tive expectations to administrative consumer transactions data. The transactions data come from EarnIn, an American financial technology company that offers millions of users early access to their wages prior to their regularly scheduled payday. Compared to the broader US population, EarnIn users are younger, more likely to identify as female, predominantly middle income, and have lower net worth and liquid assets. EarnIn fielded three surveys to its users in September 2022, July 2024, and November 2024 that solicited expectations over inflation, nominal income growth, and nominal interest rates (for both saving and borrowing). We study merged data linking 14,817 respondents' de-identified transactions data to their surveyed expectations.

Eliciting subjective beliefs is important for two reasons. First, subjective expectations data enable us to avoid assuming FIRE. Prior analyses of economic wedges assume FIRE because it allows estimating rational expectations by averaging realized outcomes (e.g., Chari, Kehoe and McGrattan, 2007; Berger, Bocola and Dovis, 2023). However, behaviors that are consistent with constrained optimization can also generally be rationalized by some set of beliefs. Prior wedge analyses may therefore conflate the effects of frictions with deviations from FIRE. Our wedge analysis is the first, to our knowledge, to relax the assumption of FIRE, allowing us to quantify the impact of frictions and non-standard preferences separately from the effects of any deviations from FIRE.

The second advantage of subjective beliefs data is that it allows us to calculate individual-level rather than aggregate wedges. Even if households actually form FIRE, measuring individual-level wedges based on realized (future) individual-level outcomes would conflate the impact of frictions with prediction error. To overcome this, prior work averaged wedges across households with similar observable characteristics, and then focused on aggregate or group-level wedges (Berger et al., 2023). Subjective expectations data allow us to directly calculate individual-level wedges.

We have two main results. First, we find that the typical consumption wedge is large. In absolute value, the median wedge is 52% of frictionless consumption. The magnitude of this estimate implies that distortions are first-order determinants of consumption among low-liquidity households. Hence, incorporating frictions or behavioral preferences into theories of consumer behavior is necessary to generate a realistic cross-section of consumption. The significant cross-sectional heterogeneity also implies that frictions and/or behavioral preferences are important determinants of the distribution of welfare. And because heterogeneity in consumer behavior can matter for aggregate outcomes (e.g., Kaplan et al., 2018), this heterogeneity may also be important for explaining aggregate consumption and income. Additionally, without first taking its absolute value, the average wedge is much smaller (19% of frictionless consumption). This highlights the value of studying micro-level wedges, as aggregate or average wedges may understate the importance of frictions in the cross-section.

Second, the observed distribution of wedges rejects financial constraints as the dominant friction for at least one quarter of our sample. We find that 25% of wedges are positive (over-consumption) and 75% are negative (under-consumption). Since financial constraints can only generate negative wedges, and therefore cannot account for the 25% of over-consumers, additional or alternative distortions are necessary to generate positive wedges. We identify two directions for theory to rationalize our findings. The first is to augment models featuring financial constraints to also include frictions that generate positive wedges, such as present bias (e.g., Lee and Maxted, 2023; Maxted, Laibson and Moll, 2024). The second is to include frictions that generate both positive and negative wedges, such as consumption adjustment costs and bounded rationality (e.g., Beraja and Zorzi, 2024; Ilut and Valchev, 2023). Both are sources of consumer inertia, which can create both positive and negative wedges by limiting the consumption response to shocks. In addition to these qualitative diagnostic implications, our quantitative results can also be used to calibrate structural models.

We verify that our two main findings are robust to alternative calibration choices regarding the preference parameters and the approximation points. A sensitivity analysis generally finds mild sensitivity and, at worst, that our baseline calibration is conservative regarding the typical size of wedges and degree of over-consumption. Our baseline calibration also assumes preference homogeneity and our sensitivity analysis further suggests that preference heterogeneity is unlikely to account for the bulk of the wedges we measure. We also assess the potential impact of measurement error in consumption, wealth, and beliefs on our results in several analyses. We generally find that it would take a substantial degree of measurement error to meaningfully change our results.

Our last analyses provide further evidence on the interpretation and sources of the wedges. We begin by relating individuals' wedges to both their actual and hypothetical consumption behavior. We measure individual-level MPCs in response to stimulus checks received in 2021 and find that MPCs are positively correlated with wedges. This suggests that a force capable of generating positive wedges, such as present bias or consumer inertia, is likely a key reason for high MPCs out of transitory income shocks. We also compare wedges against answers to questions about hypothetical behavior, such as whether one would save more or save less if they expected higher inflation. Here we too find a positive association between wedges and hypothetically preferring to save less. This suggests that forces generating positive wedges are also important for consumers' responses to expected inflation.

We next examine the relationship between individuals' wedges and measures of their financial distress, housing tenure, and consumption commitments. For the first, we elicit measures of both perceived financial distress (such as anxiety about finances) and observable proxies (such as regularly having bank account balances below $500). We find that all of our proxies for financial distress are strongly positively correlated with wedges. To a less extent, larger (in magnitude) negative wedges are also associated with greater financial distress among those with negative wedges. This pattern potentially be rationalized by either of our two proposed alternative mechanisms: households facing both financial constraints and present bias or consumer inertia.

Separately comparing respondents with and without mortgages (as a proxy for homeownership), we find that the bulk of positive wedges are attributable to people without mortgages. People with mortgages are the primary source of the larger negative wedges in our sample. This is consistent with financial constraints being the dominant friction for homeowners, while present bias and/or inertia being most influential among households that do not own a home. Regarding the potential role of consumption commitments, we construct a proxy based on the share of income spent on housing and childcare, two expenses that are typically large and difficult to adjust. We find that consumption wedges are positively correlated with consumption commitments, suggesting that they are another key ingredient to models of consumption choices. Taking stock, our results suggest that it is important to consider distortions beyond financial constraints, such as present bias and consumer inertia, to better understand consumer behavior.

Related Literature. Our paper contributes to several literatures. First, we build on the empirical macroeconomics literature studying the determinants of consumption. A central finding of this literature is large MPCs, especially among consumers with low income and low liquid wealth (Johnson, Parker and Souleles, 2006; Ganong and Noel, 2019; Baker, 2018). These cross-sectional patterns have served as important motivation for the inclusion of wealth heterogeneity and financial frictions in macro models (e.g., Kaplan and Violante, 2014; Koşar, Melcangi, Pilossoph and Wiczer, 2023). However, recent work has also found high MPCs among high-earning and wealthy households, which has motivated behavioral explanations, such as bounded rationality and present bias (e.g., Ilut and Valchev, 2023; Boutros, 2022; Lian, 2023; Maxted, Laibson and Moll, 2024).

Our findings can help refine the design of models by providing new data points that speak to extent and direction of consumption distortions among low-liquidity households. The large size of the distortions we document, relative to a benchmark that does not require FIRE, highlights the importance of frictions in determining the consumption for low-liquidity households. Additionally, the heterogeneity in (positive and negative) wedges and their correlation with MPCs support alternative models of frictions. Quantitatively, our results may be useful calibration targets for models as well.

Methodologically, our analysis and findings also suggest that measuring micro-level wedges and using them to test alternative models of frictions could be a promising direction for future work. Our empirical approach is applicable to other settings, e.g., using only survey data. Notably, it does not require quasi-experimental variation, unlike the empirical literature analyzing MPCs.

Second, we add to the empirical macroeconomics literature on consumer expectations. This literature has documented the importance of beliefs, including departures from FIRE, in explaining consumer behavior. D'Acunto et al. (2023) and Weber et al. (2022) provide recent reviews of this area. We complement recent papers that have linked consumption and beliefs data using, for example, survey measures (Coibion et al., 2023; D'Acunto et al., 2022), grocery shopping data through the Nielsen panel (Weber et al., 2023), German bank data (Hackethal et al., 2023), and credit card transactions (Kanz et al., 2021). Consumer beliefs appear to deviate from FIRE. For example, inflation expectations are excessively influenced by grocery prices (D'Acunto et al., 2021) and D'Acunto et al. (2024) finds evidence of extrapolative income expectations. Such findings motivate our use of subjective beliefs data to isolate the effects of frictions and behavioral preferences from distorted beliefs. Our consumption wedge analysis provides a novel demonstration and application of the value of consumer expectations.

Third, we contribute to the literature on wedge measurement by relaxing assumptions of FIRE and measuring wedges at the individual level. The business cycle accounting methodology of Chari et al. (2007) first popularized studying wedges between between actual and frictionless values of aggregate variables. Subsequent work on wedges has focused on quantifying the importance of misallocation across firms and risk-sharing across households for growth and business cycles (e.g. Hsieh and Klenow, 2009; Baqaee and Farhi, 2020; Berger et al., 2023). Recently, Choukhmane and de Silva (2024) demonstrates an alternative approach to quantifying frictions that exploits quasi-experimental variation in constraints to separate the influence of beliefs and preferences from constraints, which they apply to study the determinants of stock market participation. Our approach to measuring wedges does not require quasi-experimental data and complements this approach by separating the influence of beliefs from frictions and behavioral preferences.

Outline. We begin by introducing our frictionless benchmark and wedge measurement approach in Section 2. Section 3 describes our survey and linked transactions data. Section 4 presents our analysis of consumption wedges and Section 5 concludes.

2 Theory: Measuring Consumption Wedges

This section develops our approach to measuring consumption wedges. We begin by presenting a frictionless benchmark model, which we use to characterize frictionless consumption. We then define consumption wedges as the difference between actual and frictionless consumption. These wedges can be calculated using data on consumption, income, wealth, and beliefs over future inflation, income growth, and interest rates. These variables are sufficient statistics for consumption wedges in a large class of models; we discuss the robustness of our formula for a variety of model extensions.

2.1 Frictionless Benchmark

Consumption-Savings Problem. A consumer lives for T periods. She chooses consumption C_t and savings A_{t+1} to maximize her expected utility subject to a budget constraint, solving:

$$V_t(Y_t, A_t, P_t, R_t) = \max_{\{A_{t+1}, C_t\}} u\left(\frac{C_t}{P_t}\right) + \beta \widetilde{E}_t \left[V_{t+1}(Y_{t+1}, A_{t+1}, P_{t+1}, R_{t+1})\right] \tag{1}$$

$$\text{s.t.} \quad C_t + A_{t+1} = Y_t + A_t R_t. \tag{2}$$

Every period, she receives income Y_t and her start-of-period wealth is $A_t R_t$, where A_t is her previous savings and R_t is the rate of return realized on her wealth. A negative value of A_t corresponds to borrowing. The price level in period t is P_t. Uppercase letters denote nominal variables and lowercase letters their real counterpart (i.e., real consumption is $c_t = \frac{C_t}{P_t}$). We assume the consumer has "standard preferences," which we take to mean time consistent, time separable, strictly increasing, strictly concave, and continuously differentiable.

The expectations operator $\widetilde{E}_t(\cdot)$ denotes the consumer's subjective expectation conditional on her information set at time t. We do not place restrictions on the contents of her information set. Her subjective expectations integrate over a conditional distribution of future possible values of income, wealth, prices, and interest rates. We do not require that her subjective conditional expectation follows Bayes' rule nor that it uses valid probability distributions. That is, she can have non-FIRE expectations.

We refer to the model above as our frictionless benchmark. There are three important features of this benchmark. First, it assumes that there are no frictions. That is, there are no borrowing constraints, transaction or adjustment costs, etc. Second, it assumes standard preferences. As a result, deviations in actual consumption from the benchmark can arise due to either constraints or behavioral preferences that result in "as-if" constrained behavior, such as present bias and habit formation. Third, the frictionless benchmark flexibly allows for deviations from FIRE. The first two features allow the consumption wedge to flexibly capture distortions due to either frictions or behavioral preferences. The third feature ensures that the consumption wedge does not conflate the influence of deviations from FIRE with constraints and behavioral preferences.

Optimal consumption $C_t^\star$ in the frictionless benchmark is characterized by the budget constraint in equation (2) and the Euler equation:

$$u'\left(\frac{C_t^\star}{P_t}\right) = \beta \widetilde{E}_t\left[u'\left(\frac{C_{t+1}}{P_{t+1}}\right)\frac{R_{t+1}}{\pi_{t+1}}\right] \tag{3}$$

where $\pi_{t+1} = \frac{P_{t+1}}{P_t}$ is the inflation rate.

Frictionless Consumption. To obtain an approximate characterization of frictionless consumption, we forward-iterate the budget constraint and Euler equations, log-linearize them, and combine them. This process yields the equation below (with derivation in Appendix D):

$$\ln\left(\frac{C_t^\star}{Y_t}\right) \approx \alpha_0 + \alpha_1\frac{A_tR_t}{Y_t} + \sum_{j=1}^{T}\left\{\left[\alpha_Y\widetilde{E}_t \ln G_{t+j}^Y + \alpha_\pi \widetilde{E}_t \ln \pi_{t+j} + \alpha_R \widetilde{E}_t \ln R_{t+j}\right]\left(\sum_{k=j}^{T}\rho^k\right)\right\} \tag{4}$$

Equation (4) relates the logarithmized APC, $\ln\left(\frac{C_t^\star}{Y_t}\right)$, to three sets of objects: First are two preference parameters, the coefficient of relative risk aversion γ and the discount factor β. Second is start-of-period-wealth (divided by income), $\frac{A_tR_t}{Y_t}$. Third are expectations of gross nominal income growth G_{t+j}^Y, inflation π_{t+j}, and nominal interest rates R_{t+j}.

The parameters $\{\alpha_0, \alpha_1, \alpha_Y, \alpha_\pi, \alpha_R, \rho\}$ are functions of the preference parameters (γ, β) and the approximation points around which we log-linearize. The approximation points are the APC $\left(\frac{C}{Y}\right)$, initial net worth to income $\left(\frac{AR}{Y}\right)$, and the ratio of nominal income growth to the nominal interest rate $\left(\rho \equiv \frac{G^Y}{R}\right)$. Table 1 gives expressions for the parameter values.

In general, frictionless consumption is increasing in initial wealth and expected nominal income growth. It is decreasing in the expected nominal interest rate. Holding constant the expected nominal income and nominal interest rate, higher expected inflation affects consumption through two channels: it lowers real income (reducing consumption through an income effect) and it lowers real interest rates (increasing consumption through a substitution effect). For $\gamma \geq 1$, consumption is decreasing in expected inflation, as this leads the income effect to dominate the

Table 1. Frictionless Consumption Coefficient Formulas

Term	Formula (finite horizon)	Formula (infinite horizon)
α_0	$\left[1-\kappa_1-\frac{C}{Y}\frac{\ln(\beta)}{\gamma}\sum_{j=1}^{T}\left(\sum_{k=j}^{T}\rho^k\right)\right]\left(\frac{C}{Y}\sum_{j=0}^{T}\rho^j\right)^{-1}$	$(1-\kappa_1)\left(\frac{1-\rho}{C/Y}\right)-\frac{\ln(\beta)}{\gamma}\frac{\rho}{(1-\rho)}$
α_1	$\left(\frac{C}{Y}\sum_{j=0}^{T}\rho^j\right)^{-1}$	$\frac{1-\rho}{C/Y}$
α_Y	α_1	unchanged
α_π	$-\alpha_Y\frac{C}{Y}\left(1-\frac{1}{\gamma}\right)$	unchanged
α_R	$-\alpha_Y\left(1-\frac{C}{Y}+\frac{C/Y}{\gamma}\right)$	unchanged
ρ	$\frac{G^Y}{R}$	unchanged
κ_0	$\frac{C}{Y}+\left(\frac{C}{Y}-1\right)\sum_{j=1}^{T}\rho^j$	$\frac{C/Y-\rho}{1-\rho}$
κ_1	$\kappa_0-\frac{C}{Y}\ln\left(\frac{C}{Y}\right)\sum_{j=0}^{T}\rho^j-\left(\frac{C}{Y}-1\right)\ln(\rho)\sum_{j=1}^{T}\left(\sum_{k=j}^{T}\rho^k\right)$	$\kappa_0-\frac{C}{Y}\frac{\ln(C/Y)}{1-\rho}-\left(\frac{C}{Y}-1\right)\ln(\rho)\left[\frac{\rho}{(1-\rho)^2}\right]$

Notes: This table displays expressions relating the coefficients in Equation (4) to underlying parameters.

substitution effect.

2.2 Consumption Wedges

Actual consumption can deviate from frictionless consumption. We calculate the distortions to consumption induced by frictions (including behavioral preferences) for person i as the wedge between actual and frictionless consumption:

$$\eta_{it}=\ln\left(\frac{C_{i,t}}{Y_{i,t}}\right)-\ln\left(\frac{C^{\star}_{i,t}}{Y_{i,t}}\right) \tag{5}$$

The consumption wedge η_{it} for person i in time t is the difference between their actual and frictionless log-APC (where * denotes frictionless). Intuitively, the consumption wedge η_{it} measures how far "off" a consumer is relative to their Euler equation. Since actual income Y_{it} is in the denominator of both the actual and frictionless the APC, we can rewrite η_{it} as $\ln\left(\frac{C_{it}}{C^{\star}_{it}}\right)$. Thus η_{it} describes relative differences in terms of consumption.

Sufficient Statistics for the Consumption Wedge. Frictionless consumption is a known function of two preferences parameters (β,γ), initial wealth (divided by income), and beliefs about income, inflation, and interests rates. With knowledge of these objects, along with actual consumption, it is possible to calculate a household's consumption wedge using Equations (4) and (5). Moreover, we discuss later that this wedge formula is robust to a variety of model extensions. As a result, these variables are sufficient statistics for consumption wedges in a broad class of models.

Interpreting Consumption Wedges. Our benchmark is intentionally excludes real-world frictions. This design ensures that consumption wedges capture the total impact of *any* frictions directly influencing consumption. This includes constraints and adjustment costs, behavioral preferences, and bounded rationality (e.g., household consumption following a "simplified" policy function). Negative wedges corresponds to "under-consumption" (i.e., consuming less than the frictionless benchmark) and positive wedges to "over-consumption."

Example: Financial Constraints. To make the interpretation of wedges more concrete, we discuss several prominent frictions and behavioral preferences and how they relate to consumption wedges. We start with the primary friction considered by macroeconomics and household finance: financial constraints. These are most often modeled as a constant borrowing limit (e.g., Aiyagari, 1994), an endogenous borrowing limit (e.g., Bornstein and Indarte, 2023), or "soft" constraints arising from discrepancies in borrowing and saving rates (e.g., Kaplan et al., 2018). These frictions can introduce a wedge into the Euler equation, relative to the frictionless Euler Equation(3). For example, consider a constant borrowing limit such as $A_{t+1} \geq \overline{A}$. The Euler equation would be:

$$u'\left(\frac{C_t}{P_t}\right) = \beta \widetilde{E}_t \left[u'\left(\frac{C_{t+1}}{P_{t+1}}\right) \frac{R_{t+1}}{\pi_{t+1}} \right] + \mu_t$$

where $\mu_t \geq 0$ is the Lagrange multiplier. The Lagrange multiplier μ_t is positive if and only if the constraint is binding. All else equal, a binding constraint reduces consumption C_t. An important feature of financial constraints is that they only generate negative consumption wedges. Therefore, a testable implication of financial constraints is the sign of the consumption wedges. The presence of positive wedges would indicate that financial constraints are insufficient to rationalize empirical consumption choices.

Example: Present Bias. Present bias is a behavioral preference that features time inconsistency. Consider for example, beta-delta discounting, where agents discount future utility by an additional factor $\overline{\beta} < 1$ relative to the standard exponential discounting model (where δ is the exponential discount factor, corresponding to β in our notation above). Under these preferences, the expectation term in the Euler equation is scaled down by the degree of present bias ($\overline{\beta}$):

$$u'\left(\frac{C_t}{P_t}\right) = \overline{\beta}\delta\widetilde{E}_t\left[u'\left(\frac{C_{t+1}}{P_{t+1}}\right)\frac{R_{t+1}}{\pi_{t+1}}\right]$$

As a result, these preferences cause consumption to be higher relative to a "debiased" ($\overline{\beta} = 1$) consumer (Maxted, 2022), all else equal. Hence, present bias creates positive wedges. Similar to financial constraints, we can use the sign of empirical consumption wedges to test whether present bias is sufficient to rationalize empirical consumption choices.

Example: Inertia. Another class of distortions introduces inertia into consumption choices. One example is consumption commitments, where inertia is generated by consumption adjustment costs (e.g., Chetty and Szeidl, 2007; Beraja and Zorzi, 2024) or Calvo-style adjustment shocks (e.g., Auclert, Rognlie and Straub, 2020; Bornstein, 2021). Another is habit formation, which is a preference-based source of inertia where the utility of current consumption depends on past consumption (e.g., Fuhrer, 2000; Christiano et al., 2005; Smets and Wouters, 2007). This history dependence violates the time separability assumption of our benchmark and will therefore also be captured by our consumption wedge formula. Bounded rationality can similarly create inertia when costly cognition limits or delays consumption adjustments. For example, in Ilut and Valchev (2023), cognition costs limit households' updating of consumption decision rules, leading to inertial behavior. This class of frictions can produce either positive or negative consumption wedges. To see this, consider a consumer facing a convex utility cost of adjust their consumption: $\phi(C_t - C_{t-1})$. In such a case, the Euler equation would have additional terms reflecting this cost:

$$u'\left(\frac{C_t}{P_t}\right) - \phi'(C_t - C_{t-1}) = \beta\widetilde{E}_t\left\{\left[u'\left(\frac{C_{t+1}}{P_{t+1}}\right) - \phi'(C_{t+1} - C_t)\right]\frac{R_{t+1}}{\pi_{t+1}}\right\}.$$

Under these inertial preferences, inertia can limit the downward adjustment of consumption following negative wealth shocks, as adjustments incur a penalty, resulting in positive wedges (over-consumption). Similarly, positive shocks can lead to negative wedges. Empirical findings of both positive and negative wedges could be rationalized by this class of distortions.

Empirical evidence on consumption wedges can help guide the choice and modeling of frictions. Qualitatively, the presence of both positive and negative wedges (i.e., both over- and under-consumption) would indicate that neither financial constraints nor present bias alone are sufficient to explain empirical consumption choices. Quantitatively, estimates of wedges, their distribution, correlations with observables, or reactions to shocks could also be used to calibrate quantitative models and thus also discipline the parameters governing distortions.

2.3 Model Extensions

Our frictionless benchmark intentionally abstracts away from many real-world features for two reasons. First, by positing a benchmark without frictions or behavioral preferences, the benchmark is a ”special case” in a large class of models. That is, the benchmark corresponds to versions of richer models where frictions and behavioral preferences are turned ”off.” For example, in a model with borrowing constraints, our benchmark corresponds to infinite borrowing constraints. In a model with beta-delta discounting, it corresponds to zero present bias (i.e., $\beta = 1$ in the notation of Maxted, 2022). Abstracting away from frictions and behavioral preferences enables the wedges to measure the effect of all such distortions on consumption.

Second, to simplify the exposition we do not explicitly model other choices households make, such as labor supply. This is without loss of generality because our characterization of frictionless consumption only requires that the budget constraint and Euler equations are necessary conditions for optimality. They need not be sufficient. Below, we discuss a variety of model extensions and their implications for interpreting and measuring consumption wedges.

Additional Household Choices. The consumption wedge formula remains unchanged and its interpretation similar when allowing additional household choice variables. These include, for example, labor supply or default/bankruptcy.[3] Adding choice variables results in additional optimality conditions. But as long as an Euler equation and budget constraint continue to be necessary conditions for optimality, the consumption wedge formula is unaltered. However, if these other choices are subject to separate frictions, such as distortionary taxation on labor supply, the impact of those frictions is not captured in the consumption wedge.[4] The wedges reflect only the frictions that alter the Euler equation. In this sense, the distortions measured are specific to the consumption-saving decision.

Additional Assets. We can also enrich the frictionless benchmark to feature a portfolio choice problem where the household chooses a mix of assets and liabilities, including housing. In such a model, there is an Euler equation for each asset. Taking a portfolio-weighted sum of the Euler equations across each asset yields another Euler equation in terms of the portfolio's return. It is this Euler equation, featuring the (possibly leveraged) portfolio return, that one would use to characterize frictionless consumption. In our application we consider two securities: savings and

[3] To see this more concretely, note that the derivation of the log-linearized Euler equation in Lemma D.1.1 does not require that the utility function depend only on consumption. That is, it can depend on other endogenous choices like labor.

[4] Note also that if taxes are present, income should be measured on an after-tax basis.

debt. For each household, we measure their wedge using their expected portfolio return, which depends on their leverage and beliefs about the return to savings ($\widetilde{E}_t R^S_{i,t}$) and cost of debt ($\widetilde{E}_t R^D_{i,t}$) as follows:

$$\widetilde{E}_t R_{i,t} = \frac{S_{i,t}}{S_{i,t} - D_{i,t}} \widetilde{E}_t R^S_{i,t} + \frac{-D_{i,t}}{S_{i,t} - D_{i,t}} \widetilde{E}_t R^D_{i,t} \tag{6}$$

where $S_{i,t}$ is their total assets and $D_{i,t}$ their total liabilities (making the ratios above correspond to portfolio weights). Ideally, one would measure the portfolio return using beliefs data on all individual assets and liabilities facing households. This would be most important when measuring wedges for households with larger and more complex portfolios. However, because the households in our empirical setting have little wealth and likely relatively simpler portfolios, measurement error from focusing on two securities (savings and debt) is likely milder. A majority of respondents in our sample do not have a mortgage (and thus are unlikely to be homeowners); we also study these households separately as measurement error is also likely lower for this group their portfolio returns would not depend on expected house price appreciation.

Durable Goods. The wedge formula is not altered by the presence of durable goods. The logic is similar to the other endogenous choices discussed above. However, durable goods do present measurement challenges for applying the consumption wedge formula. Consumption of durable goods is difficult to measure because they yield a flow of consumption services over time after an initial purchase. To overcome this, Appendix D.2 formally introduces durable goods to our framework. The key assumption we make is that notional (i.e., composite) consumption is a Cobb Douglas aggregate of non-durable and durable consumption flows. This assumption allows us to impute the APC for total consumption from non-durable consumption and an estimate of the non-durable share of expenditures.

Heterogeneity. In a model with a single representative household, the aggregate or average wedge would be the appropriate object of interest. For models with multiple households, a distribution of consumption wedges can be calculated. If one finds empirically that there is significant heterogeneity in wedges, this would indicate that frictions and/or behavioral preferences are important determinants of cross-sectional heterogeneity in consumption. For the class of models where heterogeneity in consumption behavior affects aggregate outcomes (e.g., Kaplan and Violante, 2014), this would imply that correctly modeling frictions and/or behavioral preferences is important for predicting aggregate outcomes (in addition to the distribution of welfare).

Non-Household Agents and General Equilibrium. The frictionless benchmark does not explicitly feature additional agents such as firms or financial intermediaries. Adding these agents does not generally alter the consumption wedge formula. Similarly to adding endogenous household choices like labor supply, if these other agents' decisions are subject to frictions, the impact of those frictions is not reflected in the consumption wedge. We also abstract away from general equilibrium in that we do not explicitly model determinants of prices. Adding such features does not affect the consumption wedge formula so long as one continues to assume that households are price takers, as in our benchmark. While it is rare for models to deviate from the assumption that households are price takers, it is worth noting that such deviations would be conflated with the effects of frictions and behavioral preferences in the consumption wedge.

2.4 Why Wedges?

Consumption wedges have advantages as an object of study compared to widely-used alternatives such as MPCs and proxies for constraints. In contrast to MPCs, estimating consumption wedges does not require quasi-random variation, only observational data. While our analysis uses administrative consumption data to minimize measurement error, in principle one could use a survey to solicit all of the necessary inputs to measure consumption wedges. Additionally, both consumption wedges and MPCs can serve as calibration targets for quantitative models. It is difficult to empirically estimate a distribution of MPCs (and how that distribution varies with the state of the economy). As a result, models are often disciplined by a single estimate of an MPC (or estimates for several groups). Like MPCs, consumption wedges are local to the context in which they are measured. However, consumption wedges can provide additional data points that can help distinguish between models in cases where MPCs cannot.

Household finance research has long relied on proxies like credit card utilization and a lack of liquid wealth to tag people as constrained. However, frictions like present bias and consumption adjustment costs can result in high utilization and low liquid wealth without financial constraints binding. In contrast, the sign of consumption wedges indicates whether a household is under- or over-consuming. Additionally, rather than simply being a proxy for the severity of constraints, consumption wedges directly quantify the impact of frictions and behavioral preferences on consumption.

An important difference in our consumption wedge measurement approach is that we do not assume FIRE. Prior wedge analyses in the style of Chari et al. (2007) and Berger et al. (2023) assume FIRE, which allows these approaches to measure wedges without beliefs data. Under FIRE, one can measure expectations by averaging realized future outcomes for subgroups that provide valid counterfactuals for each other. However, if beliefs do deviate from FIRE, the influence of these deviations on consumption would be conflated with the impact of frictions and behavioral preferences. By calculating wedges with subjective expectations data, the assumption of FIRE can be relaxed and the effects of frictions and behavioral preferences can be separated from the influence of non-FIRE beliefs.

A second advantage of using subjective expectations data to measure consumption wedges is that it enables us to more easily measure micro-level wedges. That is, individual wedges for households. Prior wedge analyses have focused on aggregate/macro-level wedges. One reason for this choice is the estimation of FIRE beliefs by averaging. Because time series data is generally more limited for households than national aggregates, it is more difficult to estimate FIRE beliefs at the household level. Studying micro-level wedges in addition to macro-level wedges is useful because even if aggregate wedges are close to zero, there may be significant heterogeneity in wedges in the cross-section. Heterogeneity in consumer behavior can be an important determinant of macro transmission (e.g., Kaplan and Violante, 2014).

3 Data and Survey Design

Our data come from EarnIn, a US-based financial technology company that provides earned wage access to users with regular pay schedules, a fixed work location or electronic timekeeping system, and a connected bank account.[5] Earned wage access allows users to access their earnings prior to receiving their paycheck. EarnIn maintains an administrative database which includes information about each user and their bank account transactions (categorized by Plaid, a financial services company that links users' bank accounts with EarnIn), bank account balances, earnings, and cashout activity through the smartphone application. For an overview of the structure of the EarnIn data, see Appendix B.

Over two million US residents are contained in the EarnIn transactions database. Data containing both subjective economic expectations and detailed, comprehensive transactions data are rare. Prior studies have linked economic expectations to data on grocery spending (D'Acunto et al., 2021) or credit card spending (Kanz et al., 2021). Our dataset is one of few that links expectations to earnings, spending, and savings data that can paint a near-comprehensive picture of a consumer's economic activity at a high frequency. To our knowledge, we are the first to gather this data for US consumers (in contemporaneous work, Hackethal et al., 2023; D'Acunto et al., 2024, link similar data for users of a German and Chinese bank, respectively).

[5] For more information on EarnIn, see www.earnin.com.

The survey data come from three waves of surveys fielded by EarnIn in September 2022, July 2024, and November 2024. In selecting the subset of EarnIn users eligible to complete the survey, we imposed data quality requirements on the transactions data to ensure the users' linked accounts capture their economic activity. For the September 2022 and November 2024 survey waves, we restricted the sample to users for whom we observe earnings, regular spending, and balances in the 12 months leading up to the survey and for whom EarnIn had not reached their weekly email marketing limit. This resulted in an eligible sample of approximately 500,000 for the September 2022 survey wave and 220,000 for the November 2024 survey wave. The July 2024 survey wave was a follow-up survey for respondents who had completed the September 2022 survey wave, excluding those who did not meet additional survey transactions data quality restrictions or who already hit their email marketing limit (see Appendix B for detail). This resulted in 3,900 users eligible for the wave. See Appendix B for the full data-processing details including sample restrictions, earnings identification, and categorization of outflow transactions.

EarnIn sent qualifying users an invitation to complete the survey in waves spanning September 29 to October 2, 2022, July 12 to July 29, 2024, and November 22 to December 4, 2024, using their standard email marketing channels. Users were invited to complete a five-to-ten-minute (wave 1) or ten-to-fifteen (waves 2 and 3) survey about their current economic well-being and their outlook for the future. They were offered a $5 (waves 1 and 3) or $10 (wave 2) Amazon gift card as an incentive to complete the survey. In total, we received 15,866 responses—10,103 in wave 1, 875 in wave 2, and 4,888 in wave 3.[6]

In each wave, the survey included questions on income and basic demographics, economic expectations, household finances, financial distress, and financial literacy. We ask respondents to forecast percent changes in prices (short-run inflation) and income over the next 12 months, medium-run inflation between 24 and 36 months, and their estimate of inflation over the prior 12 months. The phrasing of these questions was based on similar questions from the University of Michigan Survey of Consumers (MSC) and the Federal Reserve Bank of New York Survey of Consumer Expectations (SCE). We also asked respondents to report the percentage yield they would expect to receive on any extra money saved and the percentage rate they would pay for any additional borrowing.

Our household finances questions include household income, debt, savings, and whether households use alternative financial products such as payday loans or pawn shops. Additionally, we elicit perceived financial distress by asking whether households perceive their debt as

[6] In wave 1, we exhausted our survey incentive budget after sending 250,000 invitations and closed the survey, resulting in a response rate of 4%. In waves 2 and 3, we closed each survey after sending two reminders, resulting in response rates of 22% and 2%, respectively.

manageable, whether they have difficulty borrowing, and whether they have anxiety relating to their finances, among other questions. For financial literacy, we ask the two of the "Big 5" financial literacy questions that ask about interest rates and inflation (Lusardi and Mitchell, 2014). EarnIn linked survey responses to each user, which enables us to study the relationship between individuals' economic beliefs and granular consumption-savings decisions.

The median completion time of our survey was seven minutes. For our analysis, we drop responses that reflect inattention or low-effort by limiting the sample to respondents who spent at least 3.5 minutes on the survey, provided internally consistent responses, and reported expectations within "reasonable" ranges.[7] We also drop users with insufficient transactions data coverage in the 12 months surrounding each survey wave (see Appendix B for details). After imposing all sample restrictions, we have 6,242 survey responses and 5,960 unique survey respondents in our analysis sample.

For all respondents, we observe earnings, spending, and bank account balances from January 2021 to November 2024. Our baseline analysis measures consumption and income over the 12-month period before each survey wave (e.g., for wave 1, this would be spending over October 2021 to September 2022). We classify inflow transactions as earnings (observed post-tax) using a combination of the observed earnings data with the Plaid categorization, memo line, and periodicity of the transaction. We additionally identify unemployment insurance transactions as a secondary source of income. We measure income as the sum of post-tax earnings and any received unemployment insurance payments. We define our measure of spending as outflow transactions that can be categorized as non-durable spending.[8] We focus on non-durable spending because it corresponds more closely to consumption, whereas the relationship between spending and consumption of durables depends on their rate of depreciation and the flow of services rendered. Our primary object of interest is the average propensity to consume (APC) out of income, which we define by dividing annual non-durable spending by annual (post-tax) earnings.

3.1 Summary Statistics, Expectations, and Realized Deviations

Table 2 presents summary statistics on the survey respondents. Respondents are 69 percent female, 49 percent non-white, 24 percent Black, 37 years old on average, and 41 percent college-educated. 42 percent correctly answered two financial literacy questions, 61 percent report having

[7] We limit the survey sample to respondents who reported inflation and income growth expectations between the 3rd and 97th percentiles (-15% and 50% for short-run inflation expectations, -20% and 50% for long-run inflation expectations, and -25% and 80% for income growth expectations) and interest rate expectations between the 1st and 97th percentiles (0% and 25% for savings rate expectations, 0% and 60% for borrowing rate expectations).

[8] See Appendix B for our outflows categorization methodology, which follows Ganong and Noel (2019).

Table 2. Summary Statistics

	Mean (1)	SD (2)	P25 (3)	P50 (4)	P75 (5)	N (6)
Panel A: Demographics						
Female (%)	68.9	.	.	.	.	5,960
White (%)	51.4	.	.	.	.	5,847
Black (%)	23.9	.	.	.	.	5,847
Age	37.2	8.8	31.0	36.0	43.0	5,939
Has children (%)	50.5	.	.	.	.	5,960
Spouse or partner (%)	52.7	.	.	.	.	5,960
College (%)	41.4	.	.	.	.	5,960
Panel B: Household finances						
Liquid assets ($)	1,812	4,199	236	250	1,639	5,956
Total assets ($)	60,483	86,937	12,031	26,672	56,815	5,960
Total debt ($)	40,866	57,640	7,334	15,792	34,141	5,960
Liquid net worth ($)	-39,055	57,452	-33,345	-15,027	-6,373	5,956
Total net worth ($)	19,617	61,911	-7,574	8,574	29,917	5,960
Total spending ($)	33,475	15,357	22,103	30,056	41,292	5,960
Nondurables spending ($)	29,466	13,536	19,454	26,550	36,693	5,960
Income ($)	41,566	18,965	28,388	37,699	50,548	5,960
Liquid net worth to income (%)	-92.1	131.6	-112.7	-41.4	-13.6	5,956
Total net worth to income (%)	44.8	135.1	-22.7	23.7	81.8	5,960
Nondurables C/Y (%)	75.1	27.2	55.6	69.9	87.8	5,960
Has mortgage (%)	17.2	.	.	.	.	5,959
Panel C: Economic expectations						
E(Inflation, 1Y) (%)	9	10	3	6	10	5,960
E(Inflation, 3Y) (%)	6	11	-2	5	10	5,960
E(Income growth) (%)	5	10	2	3	6	5,960
E(Rate on savings) (%)	4	4	1	2	5	5,960
E(Rate on borrowing) (%)	14	10	5	12	22	5,960
E(Net levered return) (%)	4	31	-5	2	15	5,960
Panel D: Perceived constraints						
Difficulty borrowing (%)	46.7	.	.	.	.	5,958
Debt unmanageable (%)	69.8	.	.	.	.	5,957
High financial anxiety (%)	51.2	.	.	.	.	5,958
Used alternative financial services (%)	53.3	.	.	.	.	3,543
High financial literacy (%)	42.3	.	.	.	.	5,960
Panel E: Observed constraints						
Balance often <$500 (%)	79.7	.	.	.	.	5,960
Paid overdraft or late fees (%)	60.6	.	.	.	.	5,960

Notes: The table presents summary statistics of user-level variables derived from the survey and transactions data. Columns (1) through (5) show the distribution of each variable, and column (6) shows the number of users with nonmissing data. In Panel B, spending and income are observed in the transactions data and aggregated over the 12 months prior to the survey. Income reflects the sum of post-tax labor earnings and unemployment insurance benefits. Our liquid assets measure is taken directly from the survey, while total assets and debt are imputed based on survey-reported finances and demographics. "Liquid net worth" is defined as reported liquid assets minus imputed total debt, and "total net worth" is defined as imputed total assets minus imputed total debt. In Panel E, balances are "often" below $500 if the user's observed average weekly balance is below $500 for at least 26 weeks in the 12-month pre-survey period. Includes users across all three survey waves who meet the restrictions outlined in Appendix B.

less than \$500 in savings, 70 percent report having an unmanageable amount of debt, and 51% report having a high anxiety related to their finances. As such, our sample exhibits a high degree of financial distress and skews younger, female, and non-white relative to the US population.

Table 2 summarizes users' finances. The median user has \$26,672 total assets but just \$250 liquid assets, \$15,792 total debt, \$37,699 annual post-tax income, and \$26,550 annual nondurables spending.[9] The median nondurables APC and net worth-to-income ratio are 70% and 24%, respectively. Further, only 17% of users have a mortgage. Appendix C compares the distribution of income and liquid assets within the EarnIn sample to the broader US. Our sample covers a broad range of the US labor income distribution, with a slight skew toward the middle of the distribution. In terms of liquid wealth, the EarnIn sample is skewed towards people with relatively little liquidity. Taking stock, our sample is more likely to be low-liquidity, middle-income, financially distressed, younger, and female than the general population.

Figure 1. Distribution of Inflation Expectations

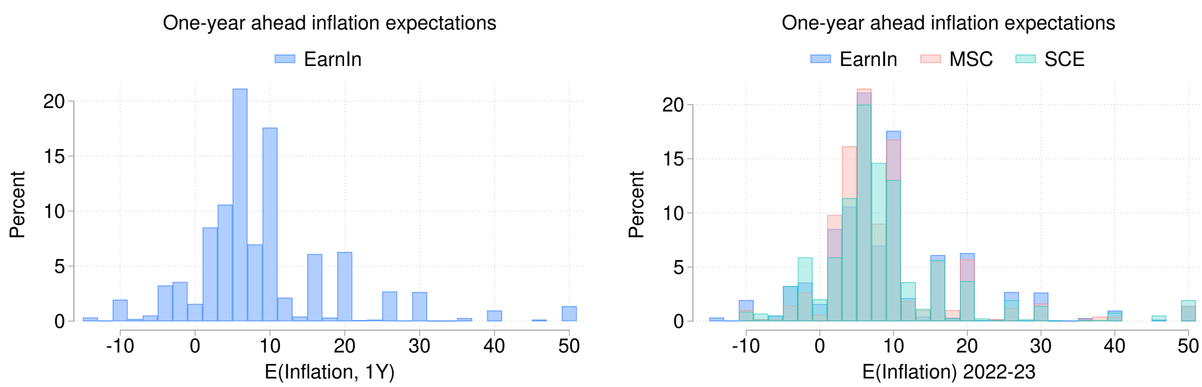

Notes: The figures show the distribution of one-year ahead inflation expectations in the EarnIn sample (left) relative to the distributions in the Michigan Survey of Consumers (MSC) and NY Federal Reserve Survey of Consumer Expectations (SCE) (right). MSC and SCE data are from October 2022. The EarnIn sample includes users across all three survey waves who meet the restrictions outlined in Appendix B.

Figure 1 presents the distribution of inflation expectations collected in the survey. The right panel overlays analogous elicitations of inflation expectations through the MSC and SCE during the same time period. The distribution of inflation expectations is remarkably similar for our survey sample, which suggests the survey instrument is performing similarly to these benchmark surveys and that inflation expectations for our survey sample are not markedly different from these nationally representative samples.

[9]Our survey solicited binned measures of liquid assets and total debt. To obtain a more precise measure of net worth AR/Y, we impute values for total assets and total debt. To do so, we leverage our detailed economic, financial, and demographic data and train a machine learning model (XGBoost) to predict these three balance sheet variables. We

Table 3. Distribution of Economic Expectations

	Ex-Post	Distribution							
		Mean	SD	P10	P25	P50	P75	P90	N
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Expectations									
E(Inflation) 2022-23		9.1	9.6	2.0	5.0	7.0	10.0	20.0	4,354
E(Inflation) 2024-25		5.1	12.1	-7.0	-2.0	5.0	10.0	20.0	4,339
Perceived inflation 2021-22		10.9	12.6	2.0	5.0	8.0	13.0	25.0	4,334
E(Income growth)		5.5	10.0	-3.0	2.0	4.2	10.0	15.0	4,354
E(Real income growth)		-3.6	13.6	-20.0	-8.0	-3.0	2.0	10.0	4,354
E(Interest on savings)		3.5	4.0	0.2	1.0	2.0	5.0	10.0	4,354
E(Interest on borrowing)		14.3	10.0	3.0	5.0	12.0	20.0	28.0	4,354
Panel B: Deviations from ex-post realizations									
Inflation 2022-23	3.7	5.4	9.6	-1.7	1.3	3.3	6.3	16.3	4,354
Inflation 2021-22	8.2	2.7	12.6	-6.2	-3.2	-0.2	4.8	16.8	4,334
Income growth	20.4	-14.8	47.9	-83.4	-24.8	-3.6	10.6	28.4	4,349
Real income growth	16.7	-20.3	47.0	-87.4	-32.3	-9.5	5.6	23.1	4,349
Interest on savings	0.5	3.0	4.0	-0.2	0.6	1.5	4.6	9.6	4,354
Interest on borrowing	21.3	-7.1	10.0	-18.3	-16.3	-9.3	-1.3	6.7	4,354

Notes: The table shows summary statistics for the economic expectations questions (Panel A) and the difference between expectations and realized values (Panel B). Includes only users who met our survey and transactions data quality restrictions (outlined in Appendix B). Column (1) reflects the realized value of each economic variable. For both nominal and real income growth, deviations are measured against each user's realized annual income growth 12 months after the survey, based on the transactions data, and the ex-post value reflects the average earnings growth across users. For the remaining variables, deviations are measured against US economy-wide values. Ex-post inflation reflects annual CPI growth in October 2022 and October 2023 (BLS, 2024). Ex-post interest on savings reflects the September 2023 average national deposit rate on savings (FDIC, 2024). Ex-post interest on borrowing reflects the average commercial bank interest rate on credit card plans, averaged across August and November 2023 (Board of Governors of the Federal Reserve System, 2024). Includes wave 1 users who meet the restrictions outlined in Appendix B.

Table 2 presents summary statistics of surveyed expectations for wave 1 users only, for whom we can compare expectations to realized values. One-year ahead inflation expectations for our sample are around nine percent, but respondents expect inflation to come down, with three-year inflation expectations of about five percent. Respondents perceive inflation over the prior year to be eleven percent which is higher than standard measures of inflation but could be reflective of the consumption baskets of our comparatively middle-income sample. Respondents forecast nominal one-year income growth of around five percent which implies forecasted real income losses of almost four percent. Respondents report quite reasonable interest rates on marginal savings and

estimate the model using data from the Survey of Consumer Finances.

borrowing of three percent and 14 percent, respectively.

Table 3 and Figure 2 present the deviations of the realized inflation, nominal earnings growth, and interest rates from the elicited expectations. To calculate the ex-post realization of inflation, the interest rate on savings, and the interest rate on borrowing, we use annual CPI growth from BLS, the national deposit rate on savings from the FDIC, and the average commercial bank interest rate on credit cards from the Board of Governors of the Federal Reserve System.[10] Measurement of individual-level realizations for these expectations is not feasible in the transactions data, so these deviations reflect the difference between the individual's expectations and aggregate measures of inflation and interest rates. Realized inflation for 2022-2023 is 3.7 percent, 5.4 percent lower than the average respondents' forecast of 9.1 percent. Responses are in line with the relationship between perceived and realized inflation between 2021-22, where respondents reported perceived inflation of 10.9 percent as realized inflation was 8.2 percent and could reflect higher exposure to inflation among our sample. Respondents' expected returns to savings were 3 percentage points higher than the average national deposit rate of 0.45 percent, and they underestimated the cost of borrowing as measured by the average credit card interest rate by 7.1 percentage points.

By leveraging the transactions data for the year following the survey, we can measure individual-level earnings realizations. This allows us to test how accurate individuals' income growth expectations are and how deviations from their expectations corresponds to their spending and savings decisions. Earnings growth expectations are remarkably accurate on average, with mean realized nominal earnings growth (measured as the percent change in earnings in the twelve months before and after the survey) of 5.99 percent relative to expected nominal earnings growth of 5.50 percent. In addition to being approximately mean zero, the distribution of prediction errors is symmetric around zero. The 25th percentile forecasted earnings growth 15 percentage points lower than they obtained, while the 75th percentile forecasted their earnings growth by a similar magnitude of 17 percentage points higher than they obtained. We will leverage this variation in income growth forecast errors in conjunction with spending behavior to test the assumption of full-information rational expectations.

4 Results: Estimated Consumption Wedges

This section measures and analyzes consumption wedges for our survey population. We begin by describing the calibration of the parameters in the wedge formula. We then present the dis-

[10] We use annual CPI growth as of October 2022 for perceived inflation and as of October 2023 for realized inflation. The national deposit rate on savings is as of September 2023. For the interest rate on credit cards, we take the average across August and November 2023 as the data are published on a quarterly basis.

Figure 2. Deviations from Ex-Post Realizations

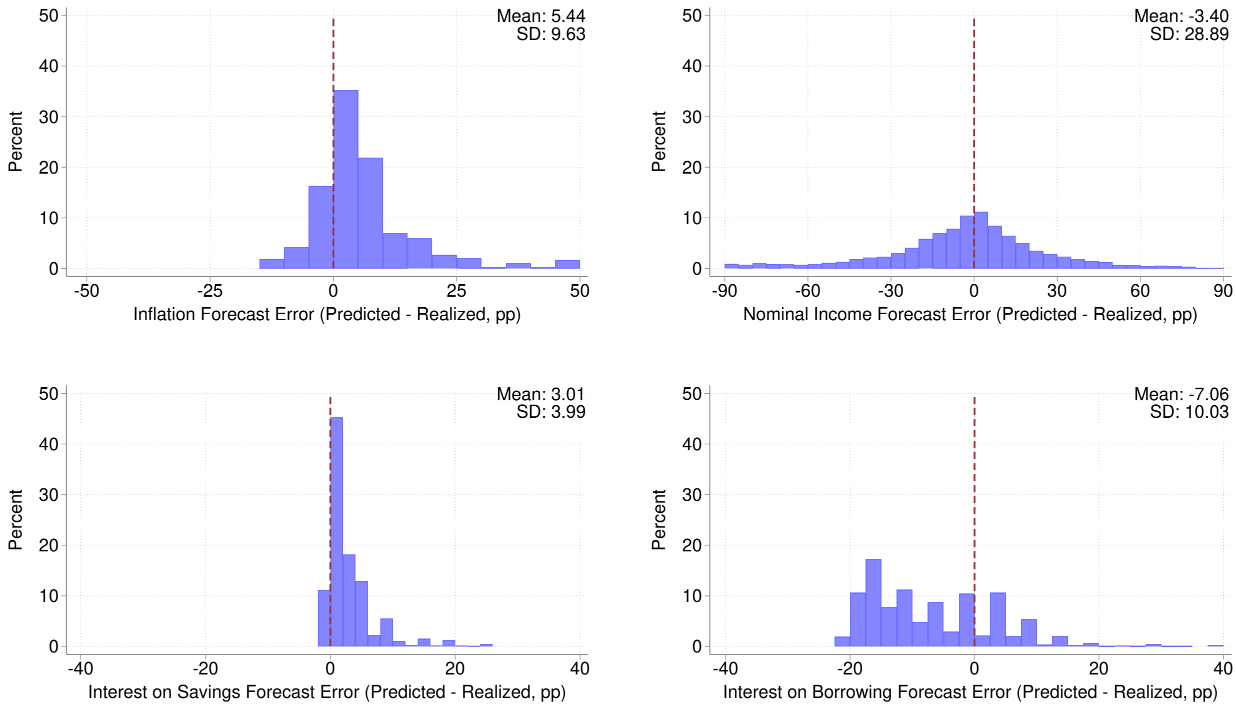

Notes: The figures show the distribution of the difference between expected and realized values of economic variables. Deviations are calculated as described in the notes of Table 3. For income growth forecast errors, we trim users with forecast errors outside of -90 to 90 percentage points. Includes only users who met our survey and transactions data quality restrictions (outlined in Appendix B).

tribution of wedges in our survey population. Notably, the median (absolute value) wedge is 52% of frictionless consumption. We also find that 25% of the sample has positive wedges (over-consumers), which cannot be explained by financial constraints. We then show that our findings of large wedges and a large over-consumer share are robust to model calibration choices and measurement error. Lastly, we correlate wedges with observables to test alternative theories about the sources of the consumption distortions.

4.1 Calibration

To measure consumption wedges for our survey sample, we begin by calculating each respondent's approximate frictionless log-APC, given by Equation (4), which we reproduce below.

$$\ln\left(\frac{C^{\star}_{i,t}}{Y_{ti,}}\right) \approx \alpha_0 + \alpha_1 \frac{A_{i,t}R_{i,t}}{Y_{i,t}} + \sum_{j=1}^{T}\left\{\left[\alpha_Y \widetilde{E}_{i,t} \ln G^{Y}_{i,t+j} + \alpha_\pi \widetilde{E}_{i,t} \ln \pi_{t+j} + \alpha_R \widetilde{E}_{i,t} \ln R_{i,t+j}\right]\left(\sum_{k=j}^{T}\rho^k\right)\right\}$$

Frictionless consumption (specifically, the log-APC) is a linear function of the ratio of net worth to income $\left(\frac{A_{i,t}R_{i,t}}{Y_{i,t}}\right)$ and expectations over nominal income growth ($\widetilde{E}_{i,t} \ln G^{Y}_{i,t+j}$), inflation ($\widetilde{E}_{i,t} \ln \pi_{t+j}$), and returns to wealth ($\widetilde{E}_{i,t} \ln \pi_{t+j}$).

Throughout, we assume that households face an infinite horizon ($T \to \infty$). This allows us to simplify expressions for the α coefficients in the frictionless consumption formula. Table 4 details expressions for the α coefficients and reports their calculated values. Calculating these coefficients requires calibrating four parameters: two preference parameters (discount factor β and the coefficient of relative risk aversion γ) and two approximation points (APC $\frac{C}{Y}$ and net worth to income $\frac{AR}{Y}$). Applying our the wedge formula also requires additional assumptions related to the measurement of consumption, the term structure of beliefs, and the measurement of consumer portfolios and expected returns. We detail our approach below.

Table 4. Calculated Wedge Coefficients

Coefficient	Value	Formula	Multiplicand
α_0	0.225	$(1-\kappa_1)\left(\frac{1-\rho}{C/Y}\right) - \frac{\ln\beta}{\gamma}\frac{\rho}{(1-\rho)}$	intercept
α_1	0.328	$\frac{1-\rho}{C/Y}$	$\frac{A_t R_t}{Y_t}$
α_Y	0.328	$\frac{1-\rho}{C/Y}$	$\widetilde{E}_{i,t} \ln G^{Y}_{i,t+1}$
α_π	-0.145	$-\alpha_Y \frac{C}{Y}\left(1-\frac{1}{\gamma}\right)$	$\widetilde{E}_{i,t} \ln \pi_{t+1}$
α_R	-0.184	$-\alpha_Y\left(1-\frac{C}{Y}+\frac{C/Y}{\gamma}\right)$	$\widetilde{E}_{i,t} \ln R_{i,t+1}$
ρ	0.7110	$\frac{1+\frac{AR}{Y}-\frac{C}{Y}}{\frac{AR}{Y}}$	(geometric common ratio)
κ_0	0.586	$\frac{C}{Y}+\left(\frac{C}{Y}-1\right)\frac{\rho}{1-\rho}$	(intermediate parameter)
κ_1	0.627	$\kappa_0 - \frac{C}{Y}\frac{\ln\frac{C}{Y}}{1-\rho} - \left(\frac{C}{Y}-1\right)\ln\rho\left[\frac{\rho}{(1-\rho)^2}\right]$	(intermediate parameter)

Notes: This table presents formulas for the expressions used in the wedge analysis. It identifies their multiplicand and also notes their calibrated value. The formula for ρ comes from evaluating the budget constraint at the approximation points. Details are available in Appendix D.3.

Calibrating Preference and Approximation Point Parameters. Our baseline calibration uses standard values for the annual discount factor ($\beta = 0.92$) and the inverse intertemporal elasticity of substitution ($\gamma = 2$). For the approximation points, we use typical values for these objects in the EarnIn sample. Specifically, we set $\frac{C}{Y} = 88.04\%$ and $\frac{AR}{Y} = -41.39\%$, which are median values for the ratios of consumption and liquid net worth to income (respectively). We use liquid net worth (liquid assets minus debt) rather than total net worth (i.e., including illiquid assets) because our measure of illiquid assets is prone to over-estimating their value.[11] Table 5 summarizes our calibration. In sensitivity analyses, we later assess the robustness of our main results to alternative parameter values.

Table 5. Calibrated Parameter Values and their Sources

Parameter	Value	Meaning	Source
$\frac{C}{Y}$	88.04%	Steady state ratio of consumption expenditures to income	Median ratio of non-durable spending to income in EarnIn sample (69.87%) divided by non-durable share of expenditures 79.37% (calculated in Beraja and Zorzi, 2024, using the Consumer Expenditure Survey)
$\frac{AR}{Y}$	-41.39%	Steady state ratio of net worth to income	Median ratio of net worth to income in EarnIn sample
γ	2	Inverse intertemporal elasticity of substitution	Standard value
β	0.92	Annual discount factor	Standard value

Notes: This table presents parameters used in the wedge analysis. It details the values used in our preferred specification, the economic meaning of the parameters, and the source of the chosen value.

Expected Returns. Taking our frictionless consumption formula to the data requires additional assumptions related to the measurement of key inputs to the formula. For the expected gross return/interest rate $\widetilde{E}_{i,t} \ln R_{i,t+1}$, one would ideally measure the expected portfolio return. This is because the budget constraint depends on the total net return across all assets and debts. But to measure this requires knowing expected returns for all assets and debts, as well as their portfolio shares. To minimize survey attrition, our survey questions focused on two expected interest rates: the return to saving and the cost of debt. We calculate portfolio weights by dividing both total assets (including imputed illiquid assets) and total debt by total net worth. We then take a weighted average of the two expected returns (as shown in Equation 6).

Term Structure of Beliefs. Frictionless consumption depends not only one-year-ahead beliefs but also the entire term structure of beliefs. Our survey solicits one-year-ahead beliefs over nom-

[11] This results in a lower ρ than we would otherwise obtain using median total net worth to income. We opt for this because the upward bias of illiquid assets can result in a $\rho \geq 1$ and an undefined expression for frictionless consumption. The issue of $\rho \geq 1$ can arise when the measure of net worth to income is biased but consumption to income is not similarly biased.

inal income growth, interest rates, and inflation, as well as a three-year-ahead belief for inflation. To calculate wedges, we impute the remaining term structure. For inflation expectations, we use data on one to thirty-year-ahead expected inflation (measured via inflation swaps). Our approach is motivated by the observation that the ratio of one- to three-year-ahead beliefs is similar to that of the swaps-implied beliefs (measured in the same month of each survey wave). We details our procedure in Appendix **??**. For interest rates, we assume that scale over time in proportion to inflation expectations.[12] Lastly, the term structure of income likely reflects lifecycle considerations. To impute this, we use data from the Survey of Consumer Expectations (SCE), which solicits one-year-ahead income expectations for a broad population of US residents. Using data on all available years, we find that one-year-ahead expected nominal income growth appears to exponentially decay with age. We estimate this rate of decay in the SCE and use it to extrapolate income expectations for our sample.

Income and Consumption Measurement Assumptions. To calculate the wedge, we subtract the frictionless log-APC (implied by the consumer's beliefs and current wealth) from their actual log-APC. We then exponentiate it and subtract one to obtain our preferred measure of consumption wedges. That is:

$$\text{consumption wedge} \equiv \exp(\eta_{i,t}) - 1 = \exp\left[\ln\left(\frac{C_{i,t}}{C^{\star}_{i,t}}\right)\right] - 1 = \frac{C_{i,t} - C^{\star}_{i,t}}{C^{\star}_{i,t}}.$$

To measure the log-APC, we first calculate the annual log-APC for non-durable goods. For the denominator (income), we measure annual income as the sum of all observed (after-tax) income, including government transfers such as UI payments. We calculate the non-durables log-APC using data from the 12 months prior to and including the survey month for each wave (i.e., October 2021 to September 2022 for Wave 1).[13] To obtain a "total" APC (i.e., for all consumption), we divide each respondent's non-durable APC by the expenditure share of non-durable goods (79.37%). Under the assumption that notional ("total") consumption is a Cobb-Douglas aggregate of durable and non-durable good consumption flows, this calculation yields the notional APC (for a proof, see Appendix D.2). We obtain the non-durable expenditure share from Beraja and Zorzi (2024), which calculates it using Consumer Expenditure Survey data.

[12] Specifically: $\widetilde{E}_{i,t} \ln R_{i,t+j+1} = \widetilde{E}_{i,t} \ln R_{i,t+j} \frac{\widetilde{E}_{i,t} \ln \pi_{t+j+1}}{\widetilde{E}_{i,t} \ln \pi_{t+j}}$.

[13] Conceptually, we want to measure time t consumption wedges using time t consumption and time t beliefs about $t+k$ variables. We verify that we obtain similar results when using all 12 months versus fewer (including only September 2022) or when averaging calculated wedges within respondents across all 12 months. For more, see Appendix E.

4.2 Main Results: Empirical Consumption Wedges

Calculating consumption wedges for each respondent, we find large wedges and significant heterogeneity. Figure 3 displays histograms of measured wedges. Our first main finding is that many consumers have significantly distorted consumption. The modal wedge is approximately -55%, indicating that the modal respondent has consumption distorted downwards 55% relative to their counterfactual frictionless consumption. Taking the absolute value of wedges, we find that the median distortion is 52%. Further examining the absolute value wedges, we see that the modal respondent has a distortion around 55%. Not all consumers have large wedges. Approximately 8% of the sample has consumption within 10% of frictionless consumption. These households either face minimal distortions or offsetting distortions—for example, a consumer may be present biased but financial constraints limit their ability to over-consume.

Figure 3. Distribution of Dynamic Consumption Wedges

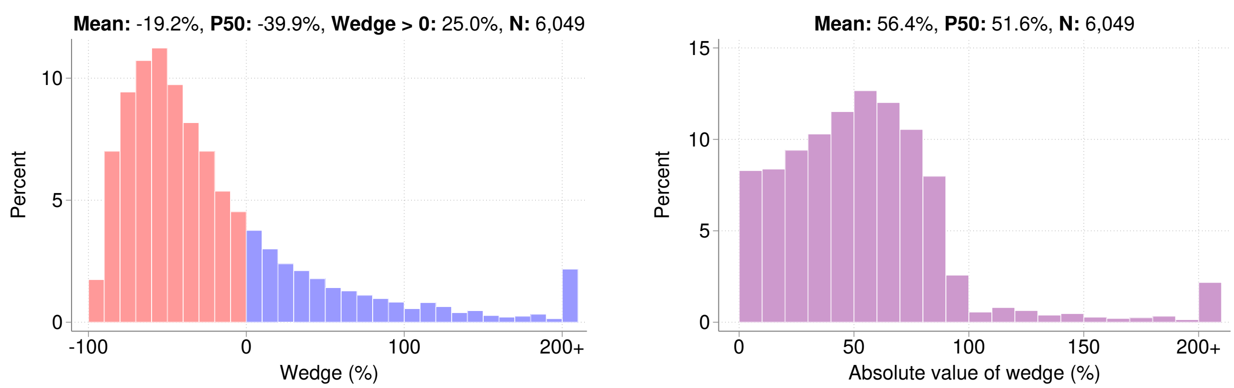

Notes: The figures show the distribution of dynamic consumption wedges (left) and their absolute values (right). Wedges are defined as the percent difference between the observed APC and frictionless APC, and they are trimmed at the 1st and 99th percentiles. Includes users from all three survey waves that meet the restrictions outlined in Appendix B.

Overall, the histograms reveal that many consumers in our sample face large distortions to their consumption. The large size of these distortions indicates that frictions or behavioral preferences are an important determinant of the consumption choices of low-liquidity households. Therefore, frictions and behavioral preferences are important to include in models of consumer behavior featuring such households.

Our findings also highlight the value of studying micro-level wedges, as opposed to macro-level (aggregate or average) wedges. The average wedge is much smaller in magnitude than than the median absolute value wedge (-19% versus 52%). With observing the distribution, one could significantly under-estimate the importance of distortions in the consumption decisions of low-liquidity consumers. In many modern macroeconomic models, heterogeneity in consumption behavior matters in the sense that it influences aggregate consumption (e.g., Kaplan and Violante, 2014; Maxted et al., 2024). For the class of models where such heterogeneity affects the aggregate economy, knowledge of the *distribution* of consumption wedges can be especially helpful for disciplining the choice and quantitative modeling of frictions and behavioral preferences. Indeed, the moments we document here may prove useful for calibrating quantitative models.

A caveat regarding the size of the wedges is that, given the nature of the services provided by EarnIn, our sample may tend to select on people with especially high demand for liquidity. One might therefore expect that wedges are especially large for this population. As a result, the typical wedge sizes we estimate are most likely an upper bounded for the broader population of low-liquidity (but middle-income) US residents. Regarding our second main result on the mix of positive and negative wedges (discussed below), selection is unlikely to be a concern. While EarnIn users likely have higher demand for liquidity that observationally similar peers, it is not obvious that this demand would stem from one particular distortion. Financial constraints, present bias, and inertia can all create demand for liquidity, and earned wage access is not obviously more appealing to people with one liquidity motive versus another.

Our second main finding is that many households have positive wedges (over-consume). Specifically, 25% of households are over-consuming relative to their frictionless consumption. This finding challenges the dominant modeling paradigm in household finance and macroeconomics: that financial constraints are the key friction shaping the consumption choices of low-liquidity households. As discussed in Section 2, financial constraints only generate negative wedges. While financial constraints could explain the 75% of consumers with negative wedges, such constraints cannot be the dominant friction for the other 25%. Financial constraints are therefore unable to account for the behavior of a large fraction of consumers in our sample. By similar logic, the presence of negative wedges also rejects present bias as the sole friction/behavioral preference facing the low-liquidity households in our sample. Neither financial constraints nor present bias alone can rationalize the mix of positive and negative wedges. An alternative model is necessary. One possible solution is that consumers are subject to both financial constraints and present bias. Alternatively, as discussed in Section 2, frictions that generate inertia in consumption can also give rise to a mix of positive and negative wedges in equilibrium. These include consumption adjustment costs, habit formation, and forms of bounded rationality that results in stick behavior (e.g., Ilut and Valchev, 2023).

A valuable endeavor for future research would be to quantitatively investigate the ability of these alternative frameworks to match the wedge distributions and other patterns that we document. We next assess the robustness of the results to calibration choices and measurement error. Following this, we then correlate wedges with other observables to provide additional insights into the nature of the underlying distortions. This yields additional guidance for the choice of distortions. We also provide additional moments that are better-suited targets for calibrating quantitative models because they are less likely to be affected by our analysis sample not being a representative cross-section of the US.

4.3 Robustness

A strength of our wedge measurement approach is that can recover consumption wedges for a broad class of models using the same set of sufficient statistics. But implementing this approach requires several important assumptions that are difficult to directly test/falsify. To gauge the robustness of the main findings of (1) median absolute value wedges of 52% and (2) a 25% share of over-consumers, we conduct several sensitivity and robustness analyses. Overall, we find that our results are strongly robust to both calibration choices and measurement error.

4.3.1 Sensitivity Analysis

Table 6. Sensitivity Analysis: Impact of Alternative Calibration Choices

	Parameter Range			Overconsumer (%)		P50 Abs(Wedge) (%)	
	Calibration	Min	Max	Min	Max	Min	Max
β	0.92	0.80	0.98	18.5	28.3	50.0	56.0
γ	2.0	1.0	5.0	23.1	27.0	45.4	59.9
Approx. point for nondurable $\frac{C}{Y}$	0.6987	0.55	0.75	23.2	28.8	51.5	63.3
Nondurable share of spending	0.7937	0.72	0.90	19.8	33.5	51.7	81.4
Approx. point for $\frac{AR}{Y}$	-0.4139	-0.81	-0.31	24.6	27.0	49.5	64.5

Notes: Table presents the sensitivity of two results to our parameter calibration: (1) the percent of users who overconsume and (2) the median absolute value wedge. Under the "Parameter Range" heading, the "Calibration" column shows our baseline assumed value and the "Min" and "Max" columns show the range of values that we test. Under the "Overconsumer (%)" and "P50 Abs(Wedge) (%)" columns, the "Min" and "Max" columns show the minimum and maximum value of each result that we get across each parameter range. When we vary one parameter, we hold all other parameters at their baseline calibrated values.

Time Preference Parameter. We begin by varying our choice of annual discount factor (β) from 0.80 to 0.98. Figure E.1 displays results. Our median absolute value wedge remains similar, ranging from 50-56%. The over-consumer share is slightly more sensitive and ranges from 19-28%. This share decreases as the discount factor is lowered, as less patience implies a higher level of frictionless consumption for a given set of beliefs. For a given level of actual consumption, higher frictionless consumption in turn implies a lower degree of over-consumption. Because a majority of respondents have negative wedges, lower frictionless consumption implies larger (in magnitude) negative wedges. Hence the median wedge size increases for lower values of β. Thus, for even a very low discount rate, we still find that a significant share of households are over-consuming relative to their hypothetical frictionless consumption, and the impact of consumption generally large for households with either positive or negative wedges.

Risk Preference Parameter. We next examine alternative choices for the coefficient of relative risk aversion/inverse IES (γ). We vary γ from one to five; Figure E.2 reports our results. The share of over-consumers remains close to 25%, ranging from 23-27%. The median absolute value wedge is decreasing in γ, but at a decreasing rate. It begins to plateau around $\gamma = 5$, which implies a median absolute value wedge of around 45%.

Preference Heterogeneity. Our baseline calibration assumes homogeneous time and risk preferences across all respondents. If consumers instead have heterogeneous preferences, the difference between an individual's preference from the calibrated value would appear as a consumption wedge. Hence preference heterogeneity could also be a source of consumption wedges when measured assuming homogeneous preferences. The sensitivity analyses above are informative about the extent to which preference heterogeneity could plausibly be sufficient to explain the wedges that we find. Even though we vary the preference parameters for all users simultaneously, these changes affect our main results in monotonic ways. Therefore, as long as we consider a plausible range of cross-sectional preference parameters, the sensitivity analyses bounds the plausible contribution of preference heterogeneity to the observed wedges. These analyses suggest that such heterogeneity generally would not explain more than 5 pp of neither the typical (absolute value) wedge size nor the share of over-consumers.

Approximation Points. We next vary the approximation points used to calculate the wedges, starting with the non-durable APC. Given the nature of transaction data, measurement error is most likely to result in understating non-durables spending, as some expenditure may occur that is not captured in the consumer's account (e.g., out of cash that is never deposited). However, over-stating is also possible if we misclassify durables expenditure as non-durables. Since measurement error could be positive or negative, we consider a range of approximation point values for the non-durable APC (from 55-75%) that includes our calibrated value of 69.87%. Figure E.3 reports our findings. The over-consumer share is increasing in this approximation point. Since under-stating consumption is the most likely form of measurement error, this suggests that our findings of substantial over-consumption are likely conservative with respect to this limitation. Moreover, this bias is likely small, as the magnitude of the over-consumer share does not vary significantly (ranging from 23-29%). The median absolute value wedges tends to increase for either extremely high or low values for this approximation point. Our calibrated value is among the lowest value we observe, suggesting our findings of large typical wedges are also (at worst) conservative). The typical wedge size remains similar, ranging from 52-63%.

We next vary the approximation point for net worth to income from -81% to -31% (its calibrated value is -41.39%). Figure E.4 displays results showing that the over-consumer share is little-changed, ranging from 25-27%. The median absolute value wedge remains close to 52%, ranging from 50-65%.

Non-Durable Expenditure Share. The last calibration parameter we vary is the non-durable expenditure share. Our calibrated value is 79.37% (obtained from Beraja and Zorzi, 2024). In our sensitivity analysis, we vary it from 72-90%, which more than spans the range of other values reported in the literature.[14] Figure E.5 reports results. The over-consumer share is generally decreasing in the calibrated non-durable expenditure share. This is largely because actual non-durable APCs are divided by this share to infer total APCs, hence smaller shares imply larger total APCs. For given a level of frictionless consumption, this implies more over-consumption/less under-consumption. The relationship to the over-consumer share is relatively insensitive to increases (falling at most to 20%) but more sensitive to increases (rising to at most 34%). The relationship of this parameter to the median absolute value wedge is non-monotonic. This is because decreasing this parameter generally shifts the distribution of wedges to right, but its mode is located to the left of zero. Small shifts to right have little impact on the typical wedge size, as negative wedges becoming smaller is roughly offset by positive wedges becoming larger. For larger shifts left (lower calibrated values of the non-durable share), once the mode is above zero, the typical wedge size grows more (peaking near 81% for extremely low values of non-durable expenditure). The typical wedge size is less sensitive to increases in the calibrated share, peaking near 60%. Our baseline calibration results in almost minimizing the typical wedge size, suggesting that our finding of a large typical wedge is also conservative with respect to this calibration choice.

[14] This choice of range is motivated by other estimates of this expenditure share in the literature. Estimates in Ganong and Noel (2019) using transactions data imply values of 77-85%. Laibson et al. (2022) estimate a value close to 87.5% from aggregate spending data. We obtain a value of 89.7% if we attempt to classify durable expenditure in the EarnIn transactions data.

4.3.2 Measurement Error Robustness

The next analyses examine the sensitivity of our main findings to measurement error.

Subgroups with Milder Measurement Error. We begin by studying subgroups where measurement error is plausibly milder. This helps gauge the plausibility that our results are sensitive to measurement error. Appendix E details the groups that we consider. Examples include dropping users with UI income that, despite the phrasing of our survey questions, may not have included growth in UI in their income growth forecasts.[15] Another analysis omits users who answered one or more financial literacy questions incorrectly.[16] Such users may have a more difficult time identifying their economic expectations. Similarly, we also drop users who reported inflation answers that are multiples of five, as such responses may exhibit rounding. We also consider several subgroups related to consumption measurement error. One example includes users whose cash withdrawals exceed 50% of non-durable spending in at least one month. Such users may have significant spending that we do not capture.

Figure E.6 reports results from excluding these various subgroups. The share of over-consumers remains extremely similar, generally ranging from 24-26%. The median absolute value consumption wedge also remains similar, varying from 51-53% of frictionless consumption.

Adding Noise. To gain a sense of the amount of measurement error necessary to significantly alter our findings, we study the impact of adding noise to the inputs used to measure wedges. We conduct 1,000 simulations where we add random noise to each of the consumption wedge inputs (APC, ratio of net worth to income, and expectations). The noise is drawn from a distribution with mean-zero and a standard deviation of 1.5 percentage points. In each simulation, we re-calculate wedges and measure the share of over-consumers and median absolute value wedge. Figure E.7 reports histograms of our results across the simulations. Overall, the over-consumer share remains near 25% and the median absolute value wedge close to 52%. This suggests it would take substantial measurement error to significantly alter these findings.

4.4 Interpretation: Evidence from Wedge Correlates

We next examine how consumption wedges vary with observable characteristics. We focus in particular on correlations that help validate the interpretation of the wedges (as measures of dis-

[15] 173 users in our sample (4%) have at least one month of UI income in the pre-survey period.

[16] These responses could indicate either low financial literacy or inattention during the survey.

tortions to consumption) and that shed light on the nature of the frictions and/or behavioral preferences distorting consumption.

MPCs. MPCs have received significant attention in both the household finance and macroeconomics literature. Evidence of high MPCs has been a central motivation for incorporating financial constraints into theories of consumer behavior. Motivated by this, we next examine how consumption wedges correlate with individuals' MPCs.

We measure individual-level MPCs based on consumers' non-durable spending responses to the March 2021 stimulus payments. These checks provided $1,400 to each eligible individual, with an additional $1,400 for each dependent.[17] Approximately 68% of the survey analysis sample received a stimulus check. We determine each user's stimulus payment date and amount from the transactions data. For each user, we examine consumption from 28 days before to 27 days after the stimulus check was received. Days -27 through -1 are the "pre" period, and days 0 through 27 are the "post" period. We then use the same date ranges in 2022 and 2023 as comparison periods. We calculate each individual's MPC as follows:

$$MPC_i = \frac{1}{StimulusAmount_i} \times (\Delta Spend_i^{2021} - \frac{\Delta Spend_i^{2022} + \Delta Spend_i^{2023}}{2}) \tag{7}$$

where

$$\Delta Spend_i^t = Spend_i^{Post,t} - Spend_i^{Pre,t} \tag{8}$$

Our MPC measure captures the "excess" consumption associated with receipt of the stimulus check. We note that this measure should be interpreted as at best a proxy for an individual's MPC, as we only have three observations per person. As such, this measure is unlikely an asymptotically valid estimate of the individual's true MPC. The median estimated MPC is 30%. There are a few extreme outliers (e.g., below -500% or above 500%), likely due to large, one-time purchases. Given this feature of the data, we winsorize MPCs at the 5th and 95th percentiles.

Figure 4 displays a binscatter comparing individuals' MPCs against their consumption wedges. As consumption wedges increase, we observe larger MPCs on average. A 25 pp larger wedge is associated with a 1.1 pp% larger MPC. A limitation of our MPC measure is that it is measured for a check received 1.5 to 3.5 years before our surveys (conducted between September 2022 and November 2024). The relationship we measure likely understates the relationship one would find if able to instead use a contemporaneous MPC measure.

These results have two implications. The first is as a validation of the consumption wedges,

[17] The stimulus payment dates range from March 12, 2021 to May 28, 2021.

showing that they are strongly related to an important economic behavioral response. The second is that higher MPCs are associated with over-consumption, rather than the under-consumption. This suggests that the empirical phenomenon of high MPCs among low-liquidity households is largely due to forces that generate positive wedges, such as present bias and consumer inertia, rather than financial constraints.

Figure 4. Relationship Between Dynamic Consumption Wedges and Nondurable MPCs

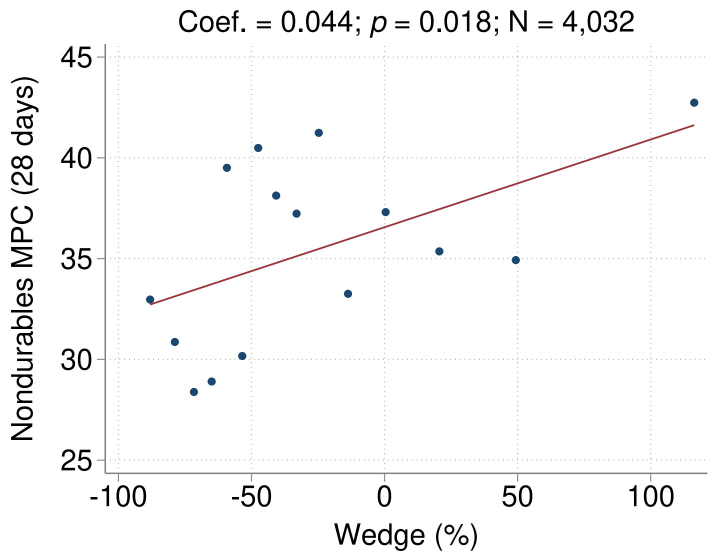

Notes: The figure illustrates the relationship between dynamic consumption wedges and nondurable MPCs. The binned scatterplot plots the average nondurables MPC within quantile-based intervals of consumption wedges, with no control variables. Wedges are defined as the percent difference between the observed APC and frictionless APC, and they are trimmed at the 1st and 99th percentiles. Each user's nondurable MPC is calculated by comparing the change in nondurable spending in the 28 days after their March 2021 stimulus check to the average change in spending during the same 28-day period in 2022 and 2023. MPCs are winsorized at the 5th and 95th percentiles. Includes users from all three survey waves that meet the restrictions outlined in Appendix B and for whom we observe a stimulus check in the transactions data (68.0% of the sample).

Hypothetical Spending/Saving Behavior. We also find wedges are correlated in an intuitive way with consumers' hypothetical behavior. The presents users with a hypothetical scenario that asks how they would adjust their saving behavior in response to higher expected inflation. We find that responding with "save less" is strongly, positively related to consumption wedges. Figure 5 displays a binscatter. This finding demonstrates an internal consistency with users' own mental models of their behavior and the wedges that we measuring using their economic beliefs and actual consumption choices. The survey also solicited rationales for respondents' hypothetical savings behavior. Among respondents selecting "save less,", a large majority (approximately

85%) attributed this dis-saving to an *inability* to reduce spending. This explanation most naturally relates to consumer inertia.

Figure 5. Relationship Between Dynamic Consumption Wedges and Hypothetical Response to Inflation

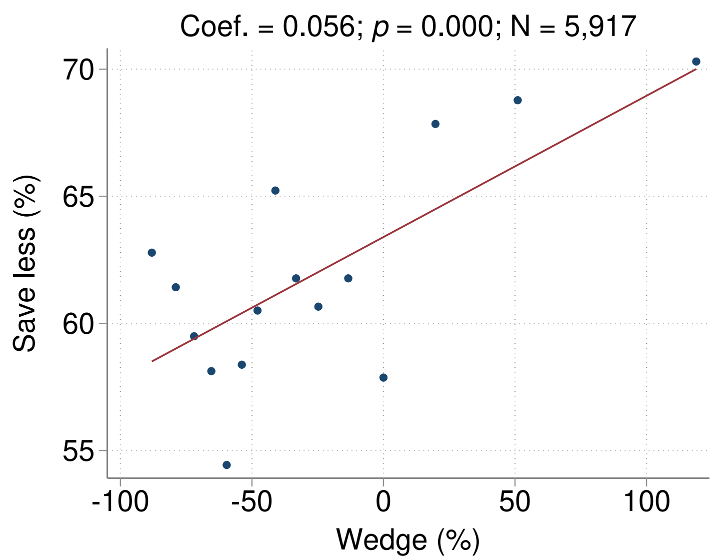

Notes: The figure illustrates the relationship between dynamic consumption wedges and an indicator whether the user would hypothetically "save less" in response to higher inflation. The binned scatterplot plots the average "save less" indicator within quantile-based intervals of consumption wedges, with no control variables. Wedges are defined as the percent difference between the observed APC and frictionless APC, and they are trimmed at the 1^{st} and 99^{th} percentiles. Includes users from all three survey waves that meet the restrictions outlined in Appendix B.

Financial Distress Proxies. We next examine how wedges vary with proxies for financial distress. These include subjective measures, such as ratings of anxiety about finances or the manageability of one's debt. We also use objective measures such as having savings account balances below $500 most of the time. We report correlations in Figure 6 below. For all measures we consider, we find a strong, positive relationship between consumption wedges and financial distress. While the relationship is positive on average, it is sometimes slightly V-shaped. This is most true for the "debt unmanageable" measure. The V-shape suggests that extremely large distortions (both positive and negative) are associated with greater financial distress. If the wedges truly captured distortions to consumption, one would expect them to be related to financial distress. Therefore, one interpretation of this result is as a validation of the measured consumption wedges.

A second interpretation relates to theories of consumer behavior. Overall, financial distress is generally rarer among households with negative consumption wedges. The most financially

Figure 6. Relationship Between Dynamic Consumption Wedges and Financial Distress

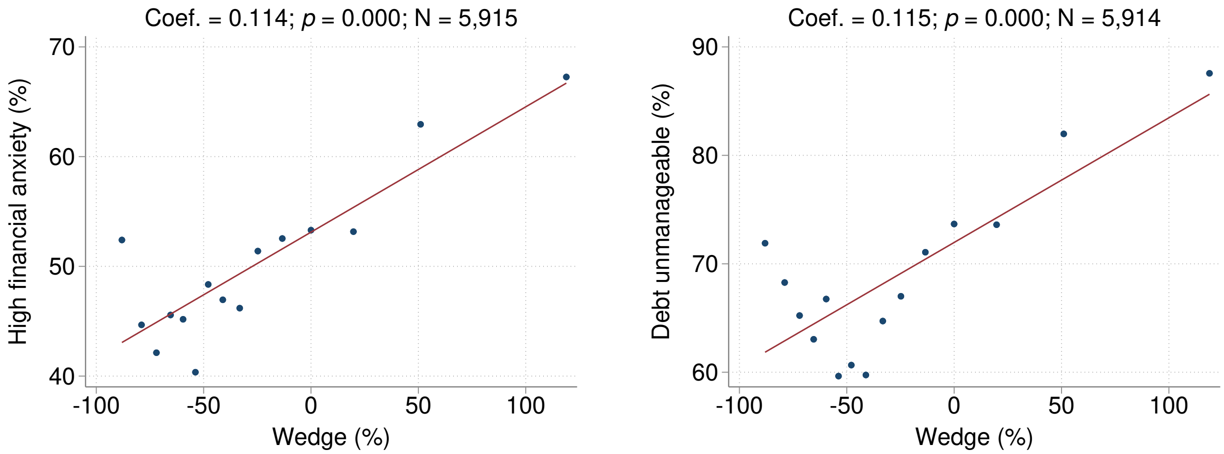

Notes: The figure presents binned scatterplots that illustrate the relationship between dynamic consumption wedges and two measures of financial distress: (1) an indicator for whether the user reports "high" or "very high" financial anxiety and (2) an indicator for whether the user reports having "a bit more" or "far more" debt than is manageable. Each binned scatterplot plots the average value of the financial distress indicator within quantile-based intervals of consumption wedges, with no control variables. Wedges are defined as the percent difference between the observed APC and frictionless APC, and they are trimmed at the 1st and 99th percentiles. Includes users from all three survey waves that meet the restrictions outlined in Appendix B.

distressed consumers are those with positive wedges. One possible rationalization for this phenomenon relates to wealthy hand-to-mouth households. Households able to accumulate substantial illiquid assets, like homes or retirement savings, are likely less financially distressed than those that struggle to acquire such assets. At the same time, constraints on their ability to borrow against or liquidate this wealth may significantly distort their consumption downwards. These households may best resemble the wealthy "hand-to-mouth" of Kaplan and Violante (2014) and indeed have their consumption primarily distorted by financial constraints. The impact of these constraints could be especially strong when the value of illiquid assets is higher. Hence, households with large/moderate negative distortions may on average be less financially distressed than those that are less able to acquire illiquid assets. Households that are both house- and cash-poor may lack the necessary collateral to be as exposed to financial constraints as the wealthy hand-to-mouth. This suggests that financial constraints is a better model for the wealthy hand-to-mouth households compared to the poor hand-to-mouth. Instead, present bias or consumer inertia is more likely the dominant friction affecting households with low liquid wealth.

Lastly, the positive association of over-consumption and financial distress suggests that the underlying frictions or behavioral preferences driving over-consumption are linked with lower wellbeing. To the extent that financial distress is a proxy for higher marginal utility of consumption, those with positive wedges would tend to have the highest marginal utility. This positive correlation between wedges and a potential proxy for marginal utility could be a useful moment for evaluating quantitative macro models.

Figure 7. Distribution of Dynamic Consumption Wedges by Mortgage Possession

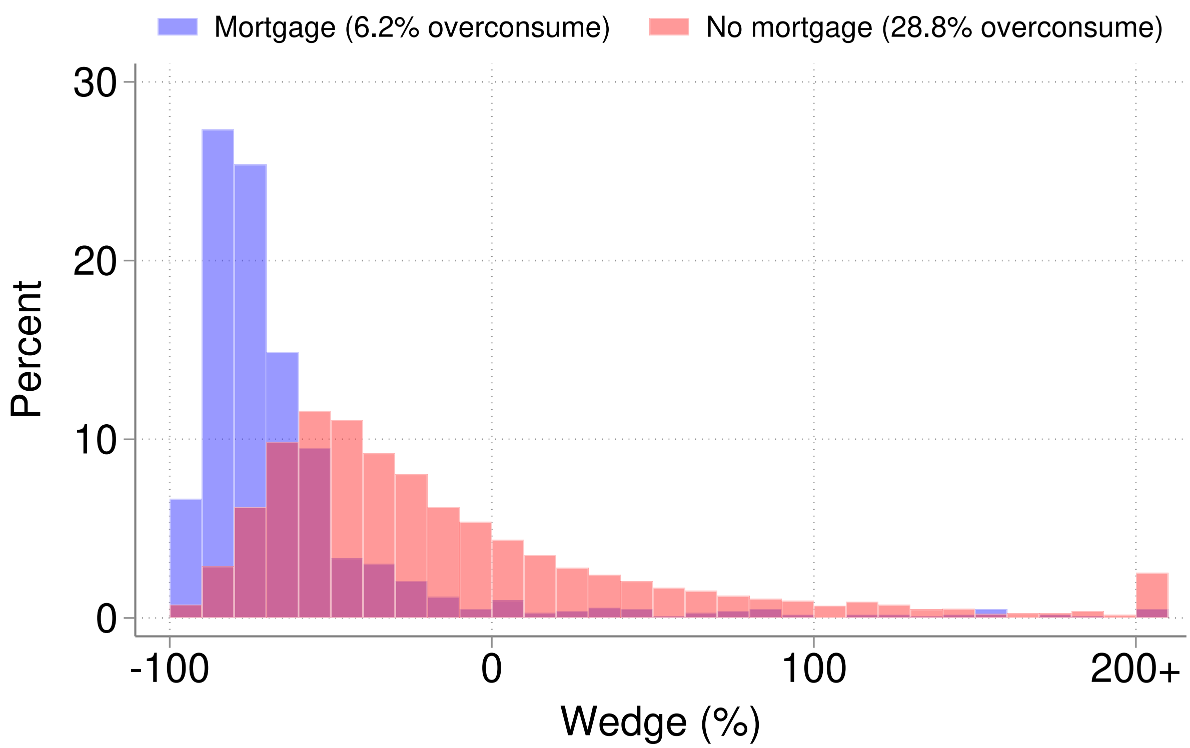

Notes: The figures show the distribution of dynamic consumption wedges, separately for users that have a mortgage (17.2% of the sample) and those that do not have a mortgage (82.8%). Wedges are defined as the percent difference between the observed APC and frictionless APC, and they are trimmed at the 1^{st} and 99^{th} percentiles. Includes users from all three survey waves that meet the restrictions outlined in Appendix B.

Homeownership. Motivated by our results regarding financial distress, we next compare wedges across homeowners and non-homeowners. We proxy for homeownership with with whether the respondent has a mortgage. Consistent with our wealthy hand-to-mouth interpretation above, we find over-consumption is extremely rare among homeowners (6%). Additionally, among the population with negative wedges, homeowners tend to have much larger (in magnitude) wedges. This is consistent with those with illiquid wealth being most distorted by financial constraints. Among those without mortgages, 29% over-consume, and those with negative wedges tend to have smaller wedges. In sum, these results further bolster our interpretation of the financial distress analysis and conclusions related the nature of distortions influencing wealthy versus poor hand-to-mouth households.

Figure 8. Relationship Between Dynamic Consumption Wedges and Consumption Commitments

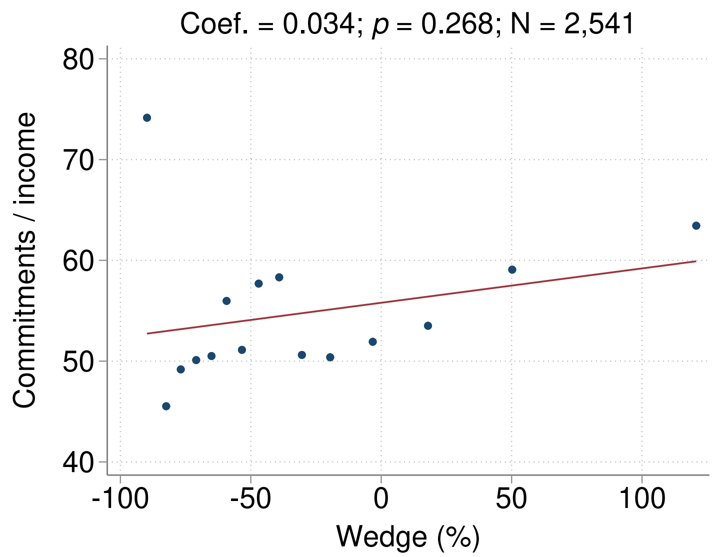

Notes: The figure illustrates the relationship between dynamic consumption wedges and consumption commitments (% income), calculated as the ratio of survey-reported monthly housing plus childcare costs to annual income from the transactions data. The binned scatterplot plots the average consumption commitments (% income) within quantile-based intervals of consumption wedges, with no control variables. Wedges are defined as the percent difference between the observed APC and frictionless APC, and they are trimmed at the 1st and 99th percentiles. Wedges are defined as the percent difference between the observed APC and frictionless APC, and they are trimmed at the 1st and 99th percentiles. Includes users from the second and third survey waves that meet the restrictions outlined in Appendix B.

Consumption Commitments. Lastly, we examine the relationship between consumption wedges and a proxy for consumption commitments in Figure 8. This measure is available only for the waves 2 and 3 survey respondents. These waves added questions soliciting respondents' monthly expenditures on housing and childcare. Our survey focused on these two specific expenses because they are among the largest "consumption commitments" (i.e., difficult to adjust expenditures) and can be relatively more difficult to identify in transactions data. We divide these reported monthly expenditures by median monthly income over the preceding twelve months (i.e., July 2023 to June 2024 for July 2024 respondents), as measured in the transactions data. We interpret a high value of this ratio as a consumer having a high degree of consumption commitments. We find an overall positive relationship between wedges and consumption commitments. A 25 pp larger wedge is associated with a 0.75 pp higher ratio of committed consumption to income.

This pattern suggests consumption commitments are a plausible friction behind the consumption wedges we measure.

5 Conclusion

This paper introduces a novel approach to measure individual-level distortions to consumption. We use a new dataset that links surveyed economic expectations to administrative transactions data for a sample of households that skews low-liquidity and middle-income. We measure the impact of distortions (frictions or behavioral preferences) as a wedge between actual consumption and a counterfactual "frictionless" benchmark. Our benchmark allows households to deviate from full-information rational expectations (FIRE), so that the wedge isolates the influence of frictions and behavioral preferences separately from deviations from FIRE. This is an important innovation to wedge measurement, as there generally exists some set beliefs that can rationalize behavior that could otherwise be explained by frictions or behavioral preferences. Because our benchmark is a special case in a large class of models, our approach makes it possible to measure the total impact of distortions on consumption due to a wide variety of frictions or behavioral preferences.

Our main findings indicate that distortions play and important role in driving household consumption and call into question the dominant role played by financial constraints in explaining the consumption of low-liquidity households. Most households have large distortions in our sample; the median (absolute value) distortion stands at 52% of frictionless consumption. The average distortion is -19%, but this belies significant heterogeneity in the cross-section. In particular, there is a mix of positive and negative wedges. 25% of wedges are positive (over-consumption) while the remaining 75% are negative. Because financial constraints can only generate negative wedges, the 25% of over-consumers cannot be explained by financial constraints. Additional or alternative mechanisms are necessary to explain the consumption choices of low-liquidity households. We identify two promising alternatives. A combination of financial constraints and present bias could potentially generate a similar distribution of wedges, as present bias creates positive wedges. Additionally, consumer inertia (e.g., consumption commitments) can give rise to both positive and negative wedges. Correlating wedges with MPCs, financial distress proxies, homeownership, and proxies for consumption commitments further suggests the two alternative models hold promise to better explain the consumption choices of low-liquidity households.

We outline several directions for future research. Future research could use surveys alone (or in conjunction with administrative transactions data) to measure wedges in other settings. Measuring wedges for consumers at different life cycle stages or in a broader population would be especially valuable. One could also use such measures to document other correlations or possibly estimate the causal effect of various shocks (such as monetary policy or stimulus check receipt) on wedges. Such evidence could help further guide the design of theories of consumer behavior. Our findings of large wedges indicate the importance of incorporating frictions or behavioral preferences into such theories. Another valuable direction for future research would be to study the wedges produced by quantitative structural models and to compare wedges for low-liquidity households with those that we find. Such evidence would help test competing models of frictions and behavioral preferences. Additionally, moments from the wedge distribution we estimate could also be used to calibrate such models, disciplining the quantitative features of the frictions or behavioral preferences present. Lastly, our findings suggest that it is important to devote more attention to distortions other than financial constraints.

References

Aiyagari, S Rao, "Uninsured idiosyncratic risk and aggregate saving," *The Quarterly Journal of Economics*, 1994, *109* (3), 659–684.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub, "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model," Technical Report, National Bureau of Economic Research 2020.

Baker, Scott R, "Debt and the response to household income shocks: Validation and application of linked financial account data," *Journal of Political Economy*, 2018, *126* (4), 1504–1557.

Baqaee, David Rezza and Emmanuel Farhi, "Productivity and misallocation in general equilibrium," *The Quarterly Journal of Economics*, 2020, *135* (1), 105–163.

Beraja, Martin and Nathan Zorzi, "Durables and Size-Dependence in the Marginal Propensity to Spend," Technical Report, National Bureau of Economic Research 2024.

Berger, David, Luigi Bocola, and Alessandro Dovis, "Imperfect risk sharing and the business cycle," *The Quarterly Journal of Economics*, 2023, *138* (3), 1765–1815.

Bledsoe, James, "2024 eCommerce Size and Sales Forecast," 2024. https://www.trade.gov/ecommerce-sales-size-forecast.

BLS, "Consumer Price Index," 2024. https://www.bls.gov/cpi/.

Board of Governors of the Federal Reserve System, "Commercial Bank Interest Rate on Credit Card Plans, Accounts Assessed Interest," 2024. https://fred.stlouisfed.org/series/TERMCBCCINTNS.

Bornstein, Gideon, "Entry and profits in an aging economy: The role of consumer inertia," *Review of Economic Studies, forthcoming*, 2021.

— **and Sasha Indarte**, "The impact of social insurance on household debt," *Available at SSRN 4205719*, 2023.

Boutros, Michael, "Windfall income shocks with finite planning horizons," Technical Report, Bank of Canada 2022.

Chari, Varadarajan V, Patrick J Kehoe, and Ellen R McGrattan, "Business cycle accounting," *Econometrica*, 2007, *75* (3), 781–836.

Chetty, Raj and Adam Szeidl, "Consumption commitments and risk preferences," *The Quarterly Journal of Economics*, 2007, *122* (2), 831–877.

Choukhmane, Taha and Tim de Silva, "What Drives Investors' Portfolio Choices? Separating Risk Preferences from Frictions," Technical Report, National Bureau of Economic Research 2024.

Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans, "Nominal rigidities and the dynamic effects of a shock to monetary policy," *Journal of political Economy*, 2005, *113* (1), 1–45.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, and Maarten Van Rooij, "How does consumption respond to news about inflation? Field evidence from a randomized control trial," *American Economic Journal: Macroeconomics*, July 2023, *15* (3), 109–52.

D'Acunto, Francesco, Michael Weber, and Xiao Yin, "Subjective Income Expectations and Household Debt Cycles," Technical Report 2024.

_ , Ulrike Malmendier, and Michael Weber, "What do the data tell us about inflation expectations?," in "Handbook of economic expectations," Elsevier, 2023, pp. 133–161.

D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber, "IQ, Expectations, and Choice," *The Review of Economic Studies*, 2022.

_ , Michael Weber, and Xiao Yin, "Subjective Income Expectations and Household Debt Cycles," Technical Report, National Bureau of Economic Research 2024.

_ , Ulrike Malmendier, Juan Ospina, and Michael Weber, "Exposure to grocery prices and inflation expectations," *Journal of Political Economy*, 2021, *129* (5), 1615–1639.

FDIC, "National Rates and Rate Caps," 2024. https://www.fdic.gov/national-rates-and-rate-caps.

Fuhrer, Jeffrey C, "Habit formation in consumption and its implications for monetary-policy models," *American economic review*, 2000, *90* (3), 367–390.

Ganong, Peter and Pascal Noel, "Consumer spending during unemployment: Positive and normative implications," *American economic review*, 2019, *109* (7), 2383–2424.

Hackethal, Andreas, Philip Schnorpfeil, and Michael Weber, "Households' Response to the Wealth Effects of Inflation," *Working Paper*, 2023.

Hsieh, Chang-Tai and Peter J Klenow, "Misallocation and manufacturing TFP in China and India," *The Quarterly journal of economics*, 2009, *124* (4), 1403–1448.

Ilut, Cosmin and Rosen Valchev, "Economic agents as imperfect problem solvers," *The Quarterly Journal of Economics*, 2023, *138* (1), 313–362.

Johnson, David S, Jonathan A Parker, and Nicholas S Souleles, "Household expenditure and the income tax rebates of 2001," *American Economic Review*, 2006, *96* (5), 1589–1610.

Kanz, Martin, Ricardo Perez-Truglia, and Mikhail Galashin, "Macroeconomic Expectations and Credit Card Spending," 2021.

Kaplan, Greg and Giovanni L Violante, "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 2014, *82* (4), 1199–1239.

_ , Benjamin Moll, and Giovanni L Violante, "Monetary policy according to HANK," *American Economic Review*, 2018, *108* (3), 697–743.

Koşar, Gizem, Davide Melcangi, Laura Pilossoph, and David G Wiczer, "Stimulus through insurance: The marginal propensity to repay debt," Technical Report 2023.

Laibson, David, Peter Maxted, and Benjamin Moll, "A simple mapping from mpcs to mpxs," Technical Report, National Bureau of Economic Research 2022.

Lee, Sean Chanwook and Peter Maxted, "Credit Card Borrowing in Heterogeneous-Agent Models: Reconciling Theory and Data," *Available at SSRN 4389878*, 2023.

Lian, Chen, "Mistakes in future consumption, high MPCs now," *American Economic Review: Insights*, 2023, *5* (4), 563–581.

Lusardi, Annamaria, "Permanent income, current income, and consumption: Evidence from two panel data sets," *Journal of Business & Economic Statistics*, 1996, *14* (1), 81–90.

_ **and Olivia S Mitchell**, "The economic importance of financial literacy: Theory and evidence," *American Economic Journal: Journal of Economic Literature*, 2014, *52* (1), 5–44.

Maxted, Peter, "Present bias unconstrained: Consumption, welfare, and the present-bias dilemma," *Available here: https://static1. squarespace. com/static/6186b3b155561c2ab5fe4957*, 2022, *62*, 1659150753483.

_ **, David Laibson, and Benjamin Moll**, "Present bias amplifies the household balance-sheet channels of macroeconomic policy," *The Quarterly Journal of Economics*, 2024, p. qjae026.

Smets, Frank and Rafael Wouters, "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American economic review*, 2007, *97* (3), 586–606.

United States Census Bureau, "Current Population Survey Annual Social and Economic Supplements," 2025.

Weber, Michael, Francesco D'Acunto, Yuriy Gorodnichenko, and Olivier Coibion, "The subjective inflation expectations of households and firms: Measurement, determinants, and implications," *Journal of Economic Perspectives*, 2022, *36* (3), 157–184.

_ **, Yuriy Gorodnichenko, and Olivier Coibion**, "The expected, perceived, and realized inflation of us households before and during the covid19 pandemic," *IMF Economic Review*, 2023, *71* (1), 326–368.

Appendix

Contents

A Appendix Figures

Figure A.1. Distribution of Static Consumption Wedges

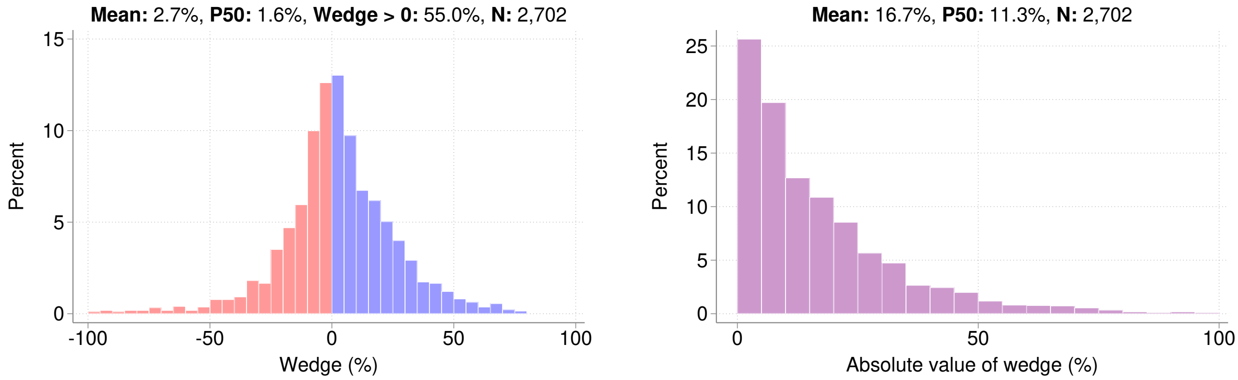

Notes: The figures show the distribution of static consumption wedges (left) and their absolute values (right). Wedges are defined as the percent difference between the observed APC and frictionless APC, and they are trimmed at the 1st and 99th percentiles. Includes users from all three survey waves that meet the restrictions outlined in Appendix B.

Figure A.2. Relationship Between Dynamic Consumption Wedges and Financial Distress

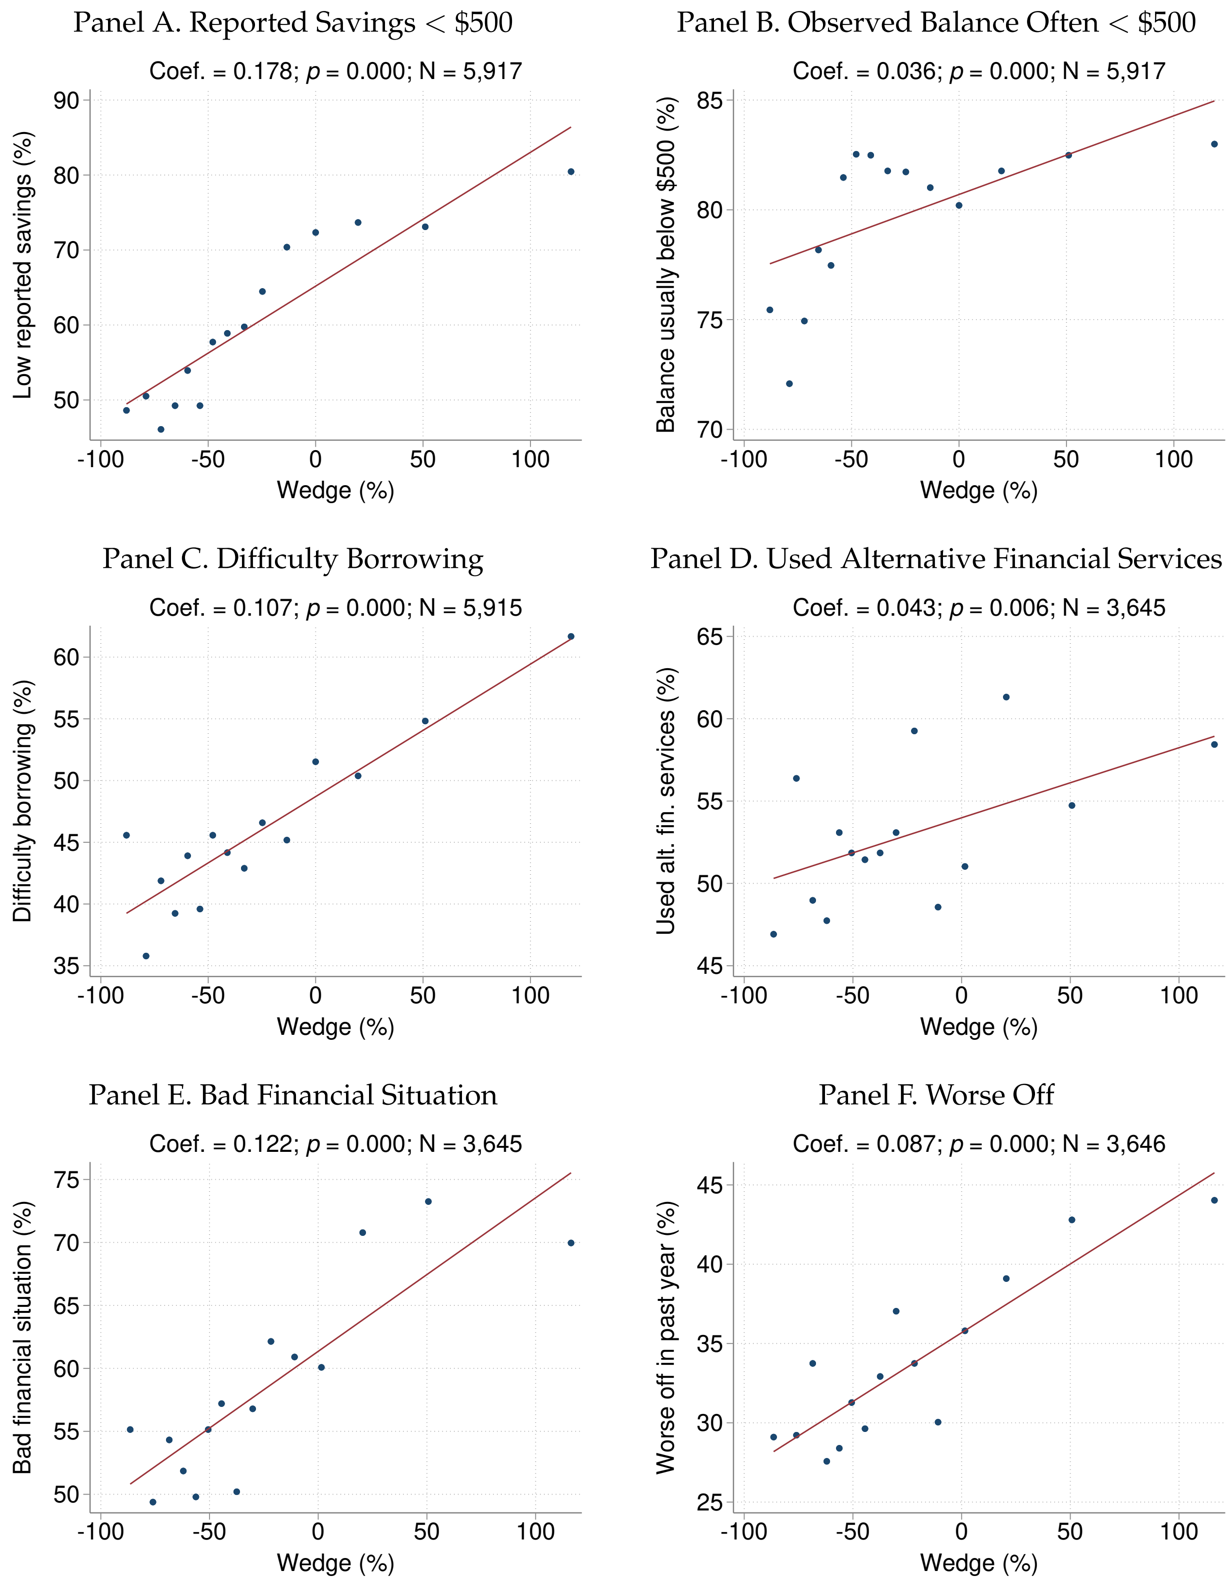

Notes: The figure presents binned scatterplots that illustrate the relationship between dynamic consumption wedges and six other indicators of financial distress: (a) whether the user reports savings below $500; (b) whether the user's observed average weekly balance is below $0 or $500 for at least 26 weeks in the 12-month sample period; (c) whether the user reports being denied for credit "often" or "most of the time;" (d) whether the user reports using payday loans, check cashing services, or pawn shops in the past 3 months; (e) whether the user reports "just getting by" or "finding it difficult to get by;" and (f) whether the user reports being "somewhat" or "much" worse off financially compared to 12 months ago. Each binned scatterplot plots the average value of the financial distress indicator within quantile-based intervals of consumption wedges, with no control variables. Wedges are defined as the percent

B Data Construction

B.1 Structure of Transactions Data

We receive anonymized transactions data from EarnIn that covers bank transactions, daily checking and savings account balances, transactions classified as earnings, and user information in the form of "tags." We leverage data from January 2021 through November 2024, which covers more than 12 months before each survey wave. None of the data we receive contains personally identifiable information, and all data is stored and processed on secure servers.

The user tags are weekly datasets at the level of de-identified individuals that contain both time-variant (e.g., work ZIP code) and time-invariant (e.g., EarnIn sign-up date) variables for each EarnIn user. The other datasets contain these tags in addition to their respective banking data.

The balance data includes the number and total balance of checking, savings, and "other" bank accounts connected to EarnIn. This dataset is at the daily level. We do not measure balances in unconnected bank accounts or investment accounts.

The transactions data includes transaction-specific information including the dollar amount, a memo describing the source or destination of a transaction, and a categorization of the type of transaction from Plaid, a third party that connects users' bank accounts to EarnIn's database.

The earnings data is a direct subset of the transactions data, covering the earnings inflows from the jobs each user reports to EarnIn. These data include the date of payment, posted date of the transaction, the amount of earnings, and whether those earnings are from unemployment benefits. Earnings are observed post-tax when are deposited into users' linked bank accounts; we do not receive information on payroll taxes and deductions.

B.2 Identifying Earnings and UI Payments

We leverage the transactions and observed earnings datasets to measure income, which we define as the sum of post-tax labor earnings and unemployment insurance (UI). We start by cleaning transaction memos to remove any non-alphabetic characters. This helps us aggregate transactions from the same source, even where memos include dates of payment.

To identify transactions as UI payments, EarnIn maintains a list of transaction memos that indicate whether an inflow is UI-related. We supplement this list with other memos that we identify as attached to UI payments.

To identify transactions as earnings, we first compare transaction amounts to EarnIn's observed earnings database, which includes weekly earnings by source for each individual. The database distinguishes different sources of earnings using three earnings variables. For example, if a user has only one source of earnings within a week, the first two earnings variable reflects the amount of earnings from each source, and the third earnings variable is missing. If we match a transaction inflow to the amount of one of these three observed earnings sources in a week, we consider those matched transactions to be earnings. If no match to a single transaction exists, we consider matches between observed earnings and the sum of transactions in a week with the same memo to be earnings. For a user with a matched memo, we also consider any other instance of that transaction memo to be earnings. We then track memos over the entirety of the database and consider a given memo to be earnings if it is tracked as earnings more than 5 times globally and is tracked as earnings over 90% of the time it appears.

Next, we perform straightforward searches of transaction memos. We flag any transaction with a memo containing the phrases "PAYROLL," "ACHPAY," "PAYRL," or "SALARY" as earnings.

Finally, we flag transactions that Plaid categorizes as Payroll or Income. Upon inspection, we find Plaid's categorization of Earnings and Income to be susceptible to false positives. To account for this, we require that the memo (1) occurs in more than two unique weeks with a modal frequency of every one or two weeks, (2) is not identified as unemployment benefits, and (3) either includes the phrase "DIRECT DEPOSIT" (or derivatives) or has a weekly amount between $50 and $5,000.

After this process, we drop hash IDs with more than five earnings in at least one week of the panel.

B.3 Categorizing Transaction Outflows

Our analysis focuses primarily on nondurables spending. To obtain this measure, we run an outflows categorization algorithm that separates durables and nondurables spending from payments (e.g., interest and principal payments on loans, bank fees), internal transfers (i.e., transfers across checking, savings, or other accounts), and external transfers (i.e., transfers to other individuals or entities through Zelle, Venmo, or other platforms). This algorithm follows the approaches of Ganong and Noel (2019) and Lusardi (1996), with some adjustments motivated by the structure of our data and analysis, discussed below.

The transaction outflows data comprises over 500 categories from Plaid. We start by mapping these Plaid categories to 33 broader categories that can be categorized as spending, payments, or transfers:

- **Spending:** Auto parts & repair, cash, department stores, discount stores, drug stores, digital entertainment, other entertainment, food services, gas stations, grocery stores, healthcare, home improvement, insurance, personal care services, professional services, taxis, transportation, travel, utilities, wholesale stores, other durables, other nondurables, other retail
- **Payments:** Auto loans, non-auto loans, buy now pay later, EarnIn earned wage access, other earned wage access, housing, overdraft & late fees, other payments
- **Transfers:** Checks, transfers across bank accounts, transfers to investment accounts, credit card payments, peer-to-peer transfers, other transfers

This mapping faces three limitations. First, Plaid's categories are based on merchant types rather than the underlying products and services, so they do not always delineate durables and nondurables. For example, a purchase at a department store may include both a mattress (durable) and clothing (nondurable). Second, some Plaid categories are too broad to be clearly mapped, such as "Purchase," "Shopping," or "Transfer." Finally, as with the earnings data, we find Plaid's categorization to be susceptible to false positives.

To manage the first limitation, we reallocate six categories that mix durables and nondurables: department stores, discount stores, drug stores, grocery stores, wholesale stores, and other retail. We follow the methodology of Ganong and Noel (2019) for the first five categories. Ganong and Noel (2019) analyze the 10-K reports for leading merchants in each category (e.g., CVS and Walgreens for drug stores, Macy's for department stores) and calculate revenue by product type. Based on this analysis, they split each category across durables and nondurables spending categories. To categorize "other retail," we approximately follow the composition of ecommerce revenue in 2020 (Bledsoe, 2024). These recategorizations are outlined in Appendix Table B.1.

For the second and third limitations, we recategorize transactions using the transaction memos. Plaid categories that are too broad are first mapped to one of "catch-all" categories: other retail, other payments, or other transfers. Then, we use regular expression searches on the memos to (1) pull transactions out of the catch-all categories and (2) fix false positives, when feasible.

Table B.1. Reallocation of Merchant Categories to Product Categories

Component of revenue	%	Mapped category	%
Department stores (Ganong and Noel, 2019)			
Clothing	80%	Other nondurables	80%
Home products	10%	Home improvement	10%
Personal care products	10%	Other nondurables	10%
Drug stores (Ganong and Noel, 2019)			
Personal care products	40%	Other nondurables	40%
Drugs	30%	Healthcare	30%
Retail nondurables	30%	Other nondurables	30%
Discount stores (Ganong and Noel, 2019)			
Groceries	50%	Groceries	50%
Home products	15%	Home improvement	15%
Retail nondurables	15%	Other nondurables	15%
Drugs	10%	Healthcare	10%
Entertainment	10%	Other entertainment	10%
Grocery stores (Ganong and Noel, 2019)			
Groceries	75&	Groceries	75%
Household supplies	25%	Other nondurables	25%
Wholesale stores (Ganong and Noel, 2019)			
Groceries	60%	Groceries	60%
Electronics	15%	Other durables	15%
Personal care products	10%	Other nondurables	10%
Home appliances	10%	Other durables	10%
Healthcare	5%	Healthcare	5%
Other retail (Bledsoe, 2024)			
Fashion	25%	Other nondurables	25%
Electronics & media	20%	Digital entertainment	20%
Toys, hobbies, & DIY	20%	Other durables	20%
Furniture & appliances	20%	Home improvement	20%
Food & personal care products	15%	Groceries	10%
Food & personal care products	15%	Other nondurables	5%

Beyond these limitations, our data face many of the same constraints as other bank account transactions datasets. Transactions are observable and categorizable to the extent that they appear on bank account statements and have informative memos. Cash withdrawals and external transfers are observed in the data, but they often mask several underlying purchases and payments that we cannot observe. Mortgage and rent payments are not captured for many users, presumably because they are paid by check or because the transaction memo does not enable categorization. The imperfect mapping between merchant and consumption categories discussed above is also a common feature of transactions data.

After applying the outflows categorization algorithm, we have the following categories:

- **Durables:** Auto parts & repair, home improvement, insurance, other durables
- **Nondurables:** Cash, digital entertainment, other entertainment, food services, gas stations, groceries, healthcare, personal care services, professional services, taxis, transportation, travel, utilities, other nondurables
- **Payments:** Auto loans, non-auto loans, buy now pay later, EarnIn earned wage access, other earned wage access, housing, overdraft & late Fees, other payments
- **Internal transfers:** Transfers across bank accounts, transfers to investment accounts, credit card payments, other internal transfers
- **External transfers:** Checks, peer-to-peer transfers, other external transfers

B.4 Defining the Sampling Frame

We send survey invitations to a restricted list of EarnIn users with adequate transactions data coverage in the 12 months leading up to the survey. We imposed most stringent transactions data quality restrictions in waves 2 and 3 compared to wave 1. Additionally, in wave 2, which resampled wave 1 respondents, we imposed survey data quality restrictions. The restrictions for each wave's sampling frame are listed below.

- Wave 1 (September 2022)
  - Non-missing earnings data at least once between September 1, 2021
  - Non-missing balances data in each bi-weekly period from September 1, 2021 through August 30, 2022
  - First recorded transaction before September 1, 2021 and latest recorded transaction after August 15, 2022
  - At least 5 outflows per month between September 2021 and August 2022
  - Non-missing bank connection date
- Wave 2 (July 2024; resampled Wave 1 users)
  - Completed the wave 1 survey
  - Were still in the EarnIn data as of June 2024
  - Took at least 3.5 minutes to fill out the wave 1 survey
  - Reported consistent debt amounts in the wave 1 survey (i.e., users who report zero debt must report N/A for debt manageability, and vice versa)

- At least 20 outflows per month each month between October 2021 and September 2023
  - Non-missing balances data each week for at least 18 months between October 2021 and September 2023
  - Sufficient categorizable spending ($\frac{\text{Consumption}}{\text{Outflows - Internal Transfers}} \geq 20\%$) for at least 18 months between October 2021 and September 2023
  - Reasonable balance of inflows and outflows ($\frac{\text{Outflows}}{\text{Inflows}} \in [50\%, 150\%]$) for at least 18 months between October 2021 and September 2023
  - Less than 1% of transaction memos between October 2021 and September 2023 are "CREDIT," "DEBIT," or missing

- Wave 3 (November 2024; repeated cross-section)
  - Did not respond to the wave 1 survey
  - Non-missing earnings data at least once between October 1, 2023 and September 30, 2024
  - At least 20 outflows per month each month between October 2023 and September 2024
  - Non-missing balances data each week for at least 9 months between October 2023 and September 2024
  - Sufficient categorizable spending ($\frac{\text{Consumption}}{\text{Outflows - Internal Transfers}} \geq 20\%$) for at least 18 months between October 2023 and September 2024
  - Reasonable balance of inflows and outflows ($\frac{\text{Outflows}}{\text{Inflows}} \in [50\%, 150\%]$) for at least 18 months between October 2023 and September 2024
  - Less than 1% of transaction memos between October 2023 and September 2024 are "CREDIT," "DEBIT," or missing
  - First recorded transaction before September 1, 2021 and latest recorded transaction after August 15, 2022
  - Non-missing bank connection date before September 1, 2021

Applying these sample restrictions, our sampling frames included 500,804 users for wave 1, 4,652 for wave 2, and 318,710 users for wave 3. EarnIn further limited each sampling frame to users who had not yet reached its weekly cap for email marketing communications. This constraint did not affect the wave 1 sampling frame but reduced the wave 2 and 3 sampling frames to 3,900 and 218,615 users, respectively.

For wave 1, EarnIn sent invitations in waves and closed the survey after 250,000 were sent, at which point 10,103 respondents had completed the survey and our incentive budget was fully spent. For waves 2 and 3, EarnIn sent invitations to the full sampling frames and closed each survey after two reminder emails, resulting in 875 responses for wave 2 and 4,888 responses for wave 3. As a result, we received response rates of 4%, 22%, and 2% for waves 1, 2, and 3, respectively.

B.5 Defining the Analysis Sample

We merge the survey responses to the cleaned transactions data, which includes user tags, balances, and categorized inflows and outflows for each user at the weekly level. We drop 174 respondents during this merge because they deleted their EarnIn accounts and no longer appear in the transactions data. Then, we apply a series of sample restrictions to arrive at our analysis sample.

First, we impose survey-based restrictions that drop users who completed the survey in an unreasonably short time frame or provided contradictory or unrealistic responses. 15,243 of 15,691 survey responses meet these restrictions, listed below.

- Survey duration at least 3.5 minutes (approximately the 5$^{\text{th}}$ percentile)
- Reported debt amounts are consistent (i.e., users who report zero debt must report N/A for debt manageability, and vice versa)

Next, we collapse the merged survey and transactions data to the monthly level and apply several sample restrictions based on the quality of the transactions data. These restrictions are designed to drop users who do not primarily consume through the bank accounts connected to EarnIn, which limits the extent to which we observe their consumption. We apply restrictions using the 12 months prior to each survey and up to 12 months after each survey (we observe 12 post-survey months for wave 1, 5 for wave 2, and 0 for wave 3). 9,522 of 15,234 survey responses meet these restrictions, listed below.

- Sufficient transaction activity: 20+ outflows per month for all 24 months
- Sufficient balances data: Non-missing balances each week for at least 75% of months
- Sufficient categorizable spending: $\frac{\text{Consumption}}{\text{Outflows - Internal Transfers}} \geq 20\%$ for at least 75% of months
- Reasonable balance of inflows and outflows: $\frac{\text{Outflows}}{\text{Inflows}} \in [50\%, 150\%]$ for at least 75% of months
- Informative memos: $< 1\%$ of memos are "CREDIT", "DEBIT", or missing across months

After applying these transactions restrictions, we restrict our sample to the 12-month pre-survey period and collapse to the user level, taking the sum of inflows and outflows across months. Then, we sequentially trim three sets of variables. We calculate the percentile cutoffs separately for each survey wave.

- First, we trim expectations. This reduces the number of survey responses from 9,522 to 7,891.
  - $E_t \ln G^Y_{t+1}$ (P3-P97)
  - $E_t \ln \pi_{t+1}$ (P3-P97)
  - $E_t \ln \pi_{t+3}$ (P3-P97)
  - $E_t \ln R^S_{t+1}$ (P1-P97)
  - $E_t \ln R^D_{t+1}$ (P1-P97)
- Second, we trim the expected leveraged interest rate. This reduces the number of survey responses from 7,891 to 7,420.
  - $E_t \ln R_{t+1}$ (P1-P99)
- Third, we trim nondurables spending, income, APCs, and the wealth-to-income ratio. This reduces the number of users from 7,420 to 6,242.
  - n_t (P1-P99)
  - Y_t (P1-P99)
  - $\widetilde{c}_t$ (P2.5-P97.5)
  - $\frac{A_t R_t}{Y_t}$ (P5-P95)

B.6 Data Transformations

We make a number of adjustments to the transactions data to standardize variables and reduce noise. As highlighted above, we aggregate data to the user-level by taking the sum of inflows and outflows across pre-survey months. We deflate inflows, outflows, balances, and dollar-based survey responses using the CPI, with September 2022 as the base month. To account for outliers in inputs to the wedge calculation, weimpose the trimming restrictions outlined above. For the wedge histograms and correlations, we trim users with wedges outside of the 1st and 99th percentiles. We also winsorize MPCs at the 5th and 95th percentiles.

C Representativeness of the EarnIn Sample

C.1 Income Percentiles

We assign each user to a percentile of the income distribution among employed US individuals. Because we do not observe gross household income in the EarnIn transactions data and our survey measure is imprecise, we base this calculation on users' post-tax labor earnings from the EarnIn data. We assign percentiles using person-level data from the Current Population Survey Annual Social and Economic Supplement (United States Census Bureau, 2025). While the CPS ASEC does not measure post-tax labor earnings, we can approximate this measure using the CPS ASEC Tax Model, which estimates an individual's federal and state tax liabilities (**?**). To our knowledge, no other US household survey measures or estimates post-tax income or tax liabilities.

The CPS asks respondents about their sources of income, household structure, and state of residence. The CPS ASEC Tax Model uses these inputs to impute adjusted gross income, taxable income, payroll taxes, and federal and state tax liabilities (before and after refundable credits). However, public CPS ASEC data does not separate estimated tax liabilities by the associated source of income, so it does not offer a precise measure of post-tax labor earnings. We estimate post-tax labor earnings in the CPS as follows:

$$Earnings^{Post} = Earnings^{Pre} - PayrollTax - \frac{Earnings^{Pre}}{TotalIncome^{Pre}} \left(FedTax + StateTax\right) \tag{9}$$

The term $TotalIncome^{Pre}$ reflects the sum of taxable income components. These components include wage and salary earnings, self-employment earnings, interest income, dividends, alimony, rent income, unemployment benefits, Social Security income, Veteran's Administration income, capital gains, rents or royalties, estate or trust income, pensions, annuities, most survivor income, and most disability income. The key assumption in this calculation is that the share of tax liabilities driven by labor earnings equals the share of taxable income (before refundable credits) from labor earnings.

In the CPS ASEC sample, we restrict to individuals who (i) are age 18 or older, (ii) are classified as adults, and (iii) received wage or salary earnings during the reference years. These restrictions allow us to estimate the distribution of annual labor earnings for US adults. When calculating percentiles, we incorporate probability weights provided in the CPS ASEC data.

Figure C.1. Distribution of EarnIn Post-Tax Labor Income Percentiles

Panel A. September 2022

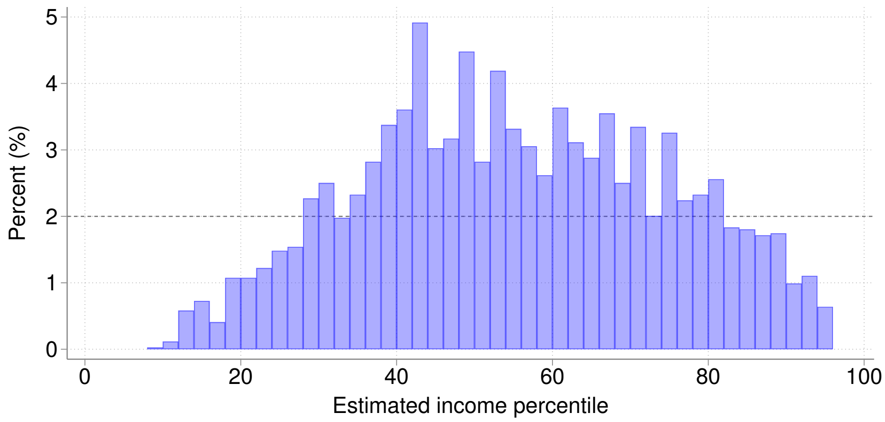

Panel B. July 2024

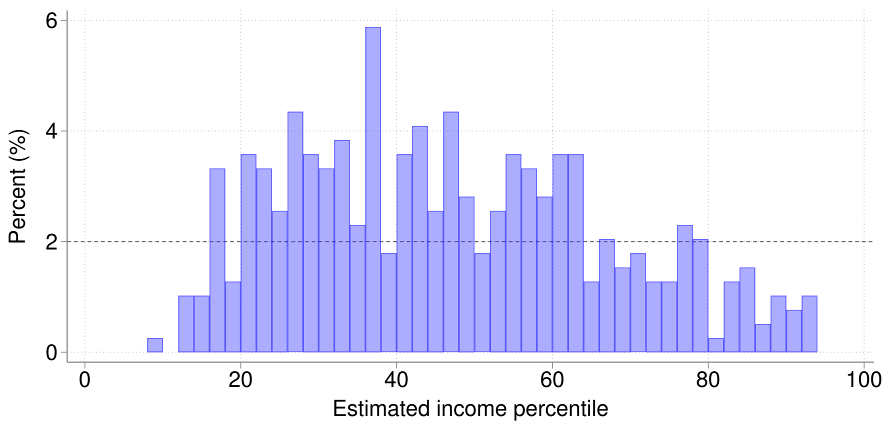

Panel C. November 2024

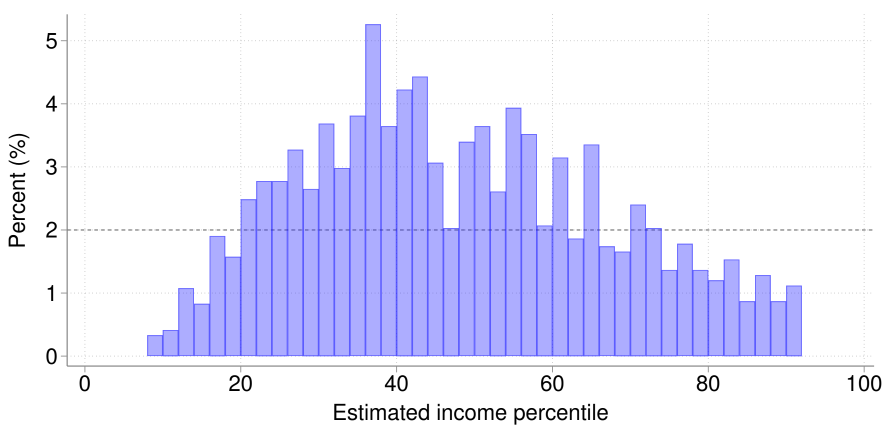

Notes: Figure presents the distribution of post-tax labor income percentiles among EarnIn survey respondents, estimated using CPS ASEC data and the CPS ASEC Tax Model. CPS ASEC sample restricts to adults age 18 or older who received wage or salary earnings during the reference year. Black dashed lines represent the uniform distribution. Includes users from all three survey waves that meet the restrictions outlined in Appendix B.

Figure C.2. Distribution of Liquid Wealth in EarnIn versus the SCF

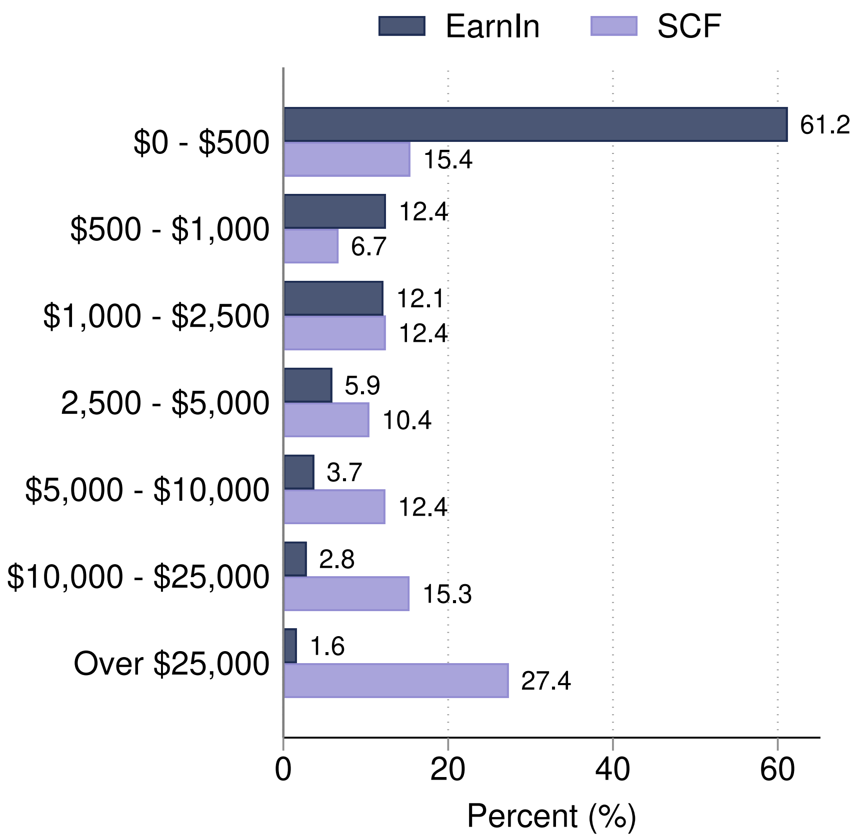

Notes: Figure presents a bar graph showing the distribution of liquid assets among EarnIn survey respondents compared to respondents in the 2022 Survey of Consumer Finances. Following Kaplan and Violante (2014), we define liquid assets as the sum of assets held in transactions accounts. Includes users from all three survey waves that meet the restrictions outlined in Appendix B.

Figure C.3. Relationship Between Dynamic Consumption Wedges and Income Percentiles

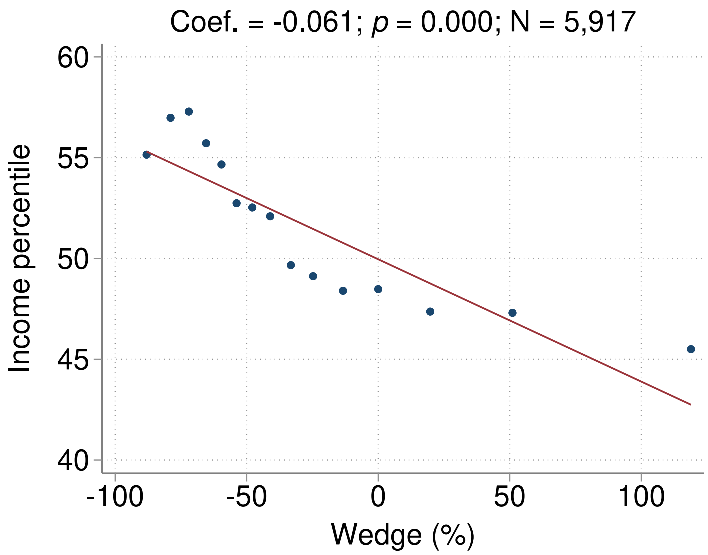

Notes: The figure illustrates the relationship between dynamic consumption wedges and user's income percentiles, estiamted using CPS ASEC data and teh CPS ASEC Tax Model. CPS ASEC sample restricts to adults age 18 or older who received wage or salary earnings during the reference year. The binned scatterplot plots the average "save less" indicator within quantile-based intervals of consumption wedges, with no control variables. Wedges are defined as the percent difference between the observed APC and frictionless APC, and they are trimmed at the 1^{st} and 99^{th} percentiles. Includes users from all three survey waves that meet the restrictions outlined in Appendix B.

D Theory Derivations and Extensions

D.1 Deriving Log-linearized Frictionless Consumption

D.1.1 Log-linearized Euler Equation

We first log-linearize the original Euler equation. Our particular log-linearization leaves the utility function unspecified. Therefore, even though we do not impose constant relative risk aversion, to a first order approximation, the Euler equation exhibits constant relative risk aversion.

Lemma 1: Log-linearized Original Euler Equation

A first order log-linear approximation of the Euler equation

$$u'(c_t) = \beta E_t \left[u'(c_{t+1}) \frac{R_{t+1}}{\pi_{t+1}} \right]$$

yields:

$$\ln c_t \approx E_t \ln c_{t+1} - \frac{1}{\gamma} \left(\ln \beta + E_t \ln R_{t+1} - E_t \ln \pi_{t+1} \right) \tag{10}$$

where $\{c, R, \pi\}$ are non-stochastic steady state values and $\gamma = -\frac{u''(c)}{u'(c)} c$ is the coefficient of relative risk aversion at the steady state level of consumption.

Proof. We begin by approximating the left-hand-side (LHS) around $\ln c_t = \ln c$

$$u'(c_t) \approx u'(c) + u'(c) \frac{u''(c)}{u'(c)} c \left(\ln c_t - \ln c \right).$$

Note that the above expression contains the coefficient of relative risk aversion, $\gamma = -\frac{u''(c)}{u'(c)} c$, enabling us to rewrite the LHS as:

$$u'(c_t) \approx u'(c) \left[1 - \gamma \left(\ln c_t - \ln c \right) \right].$$

We next approximate the right-hand-side (RHS) around $\ln c_{t+1} = \ln c$, $\ln R_{t+1} = \ln R$, and $\ln \pi_{t+1} = \ln \pi$. The values $\{c, R, \pi\}$ are a non-stochastic steady state (assumed to exist). As such, they satisfy the non-stochastic Euler equation:

$$u'(c) = \beta u'(c) \frac{R}{\pi}.$$

Note that this implies

$$\ln \beta = \ln \pi - \ln R,$$

which we will make use of later. The RHS we will approximate is

$$E_t \left\{ \exp \left[\ln \beta + \ln \left(u'(c_{t+1}) \right) + \ln R_{t+1} - \ln \pi_{t+1} \right] \right\}.$$

For notational convenience, we define the expression inside the expectation with the function $f(\cdot)$:

$$f(c, R, \pi) = \exp\left(\ln\beta + \ln\left(u'(c)\right) + \ln R - \ln\pi\right).$$

Taking a first order approximation of the RHS yields:

$$E_t\left[f(c_{t+1}, R_{t+1}, \pi_{t+1})\right] \approx f(c, R, \pi) + f_{\ln c}\left(E_t \ln c_{t+1} - \ln c\right) + f_{\ln R}\left(E_t \ln R_{t+1} - \ln R\right) + f_{\ln \pi}\left(E_t \ln \pi_{t+1} - \ln \pi\right)$$

where

$$f_{\ln c} = f(c, R, \pi)\frac{u''(c)}{u'(c)}c, \quad f_{\ln R} = f(c, R, \pi), \quad f_{\ln \pi} = -f(c, R, \pi).$$

Plugging in and rearranging gives

$$E_t\left[f(c_{t+1}, R_{t+1}, \pi_{t+1})\right] \approx f(c, R, \pi)\left[1 - \gamma\left(E_t \ln c_{t+1} - \ln c\right) + \left(E_t \ln R_{t+1} - \ln R\right) - \left(E_t \ln \pi_{t+1} - \ln \pi\right)\right].$$

Equating the LHS and RHS in our approximation yields the following expression:

$$u'(c)\left[1 - \gamma\left(\ln c_t - \ln c\right)\right] = f(c, R, \pi)\left[1 - \gamma\left(E_t \ln c_{t+1} - \ln c\right) + \left(E_t \ln R_{t+1} - \ln R\right) - \left(E_t \ln \pi_{t+1} - \ln \pi\right)\right].$$

We can use the expression for the steady state to eliminate the coefficients on the front of both terms above. Specifically:

$$u'(c) = \beta u'(c)\frac{R}{\pi} = \exp\left(\ln\beta + \ln\left(u'(c)\right) + \ln R - \ln\pi\right) = f(c, R, \pi).$$

This allows us to obtain:

$$-\gamma\left(\ln c_t - \ln c\right) = -\gamma\left(E_t \ln c_{t+1} - \ln c\right) + \left(E_t \ln R_{t+1} - \ln R\right) - \left(E_t \ln \pi_{t+1} - \ln \pi\right).$$

We can cancel the terms related to $\ln c$ and use $\ln\beta = \ln\pi - \ln R$ to rewrite the above as

$$-\gamma \ln c_t = -\gamma E_t \ln c_{t+1} + \ln\beta + E_t \ln R_{t+1} - E_t \ln \pi_{t+1}.$$

Rearranging, we finally obtain:

$$\ln c_t = E_t \ln c_{t+1} - \frac{1}{\gamma}\left(\ln\beta + E_t \ln R_{t+1} - E_t \ln \pi_{t+1}\right).$$

□

Next, we characterize the Euler equation in terms of APCs.

Remark 1: Log-linearized Original Euler Equation in Terms of APCs.

We can rewrite the log-linearized Euler equation, Equation (10) in terms of APCs (i.e., $\frac{c}{y}$) and expected nominal income growth as follows:

$$\ln\left(\frac{c_t}{y_t}\right) = E_t \ln\left(\frac{c_{t+1}}{y_{t+1}}\right) + E_t \ln g^Y_{t+1} - E_t \ln \pi_{t+1} - \frac{1}{\gamma}\left(\ln\beta + E_t \ln R_{t+1} - E_t \ln \pi_{t+1}\right) \tag{11}$$

where $g^Y_{t+1} = \frac{Y_{t+1}}{Y_t}$ is the growth rate of nominal income from period t to $t+1$.

We can iterate forward the log-linearized Euler equation (written in terms of APCs) to obtain a general multi-period version.

Remark 2: Log-Linearized Multi-Period Euler Equation

Iterating Equation (11) forward to j periods gives:

$$\ln\left(\frac{c_t}{y_t}\right) = \sum_{k=1}^{j}\left[E_t \ln g^Y_{t+k} - E_t \ln \pi_{t+k} - \frac{1}{\gamma}\left(\ln\beta + E_t \ln R_{t+k} - E_t \ln \pi_{t+k}\right)\right] + E_t \ln\left(\frac{c_{t+j}}{y_{t+j}}\right). \tag{12}$$

D.1.2 Log-Linearized Budget Constraint

We next turn our attention to the budget constraint.

Lemma 2: Log-Linearized Forward-Iterated Budget Constraint.

If the transversality condition holds ($\lim_{j\to\infty} \frac{A_{t+j}}{R_{t+1}\cdots R_{t+j-1}} = 0$)*, then a first-order log-linear approximation of the forward-iterated budget constraint is*

$$\frac{A_t R_t}{Y_t} + 1 = \kappa_1 + \widetilde{c} \ln \widetilde{c}_t + \widetilde{c} \sum_{j=1}^{T} \rho^j \ln \widetilde{c}_{t+j} + (\widetilde{c} - 1) \sum_{j=1}^{T} \left[\left(\ln G^Y_{t+j} - \ln R_{t+j} \right) \left(\sum_{k=j}^{T} \rho^k \right) \right] \tag{13}$$

where $\widetilde{c}_t \equiv \frac{C_t}{Y_t}$, $G^Y_{t+1} \equiv \frac{Y_{t+1}}{Y_t}$*, and*

$$\kappa_1 \equiv \kappa_0 - \widetilde{c} \ln \widetilde{c} \sum_{j=0}^{T} \rho^j - (\widetilde{c} - 1) \ln \rho \sum_{j=1}^{T} \left(\sum_{k=j}^{T} \rho^k \right). \tag{14}$$

Additionally, we define $\rho = \frac{G^Y}{R}$*, where* $\{\frac{C}{Y}, \rho\}$ *are steady state values, and*

$$\kappa_0 \equiv \widetilde{c} + (\widetilde{c} - 1) \sum_{j=1}^{T} \rho^j. \tag{15}$$

Proof. To obtain the forward-iterated log-linearized budget constraint, we proceed in three steps:

1. Iterate forward and simplify the nominal budget constraint.
2. Take a first-order log-linear approximation of the budget constraint.
3. Simplify the expression.

Step 1: Iterate Forward and Simplify. We begin with the nominal budget constraint, reproduced below:

$$C_t + A_{t+1} = Y_t + A_t R_t.$$

Rearranging, we isolate initial wealth:

$$A_t R_t = C_t - Y_t + A_{t+1}.$$

Forward-iterating the above expression gives

$$A_t R_t = C_t - Y_t + \frac{C_{t+1} - Y_{t+1}}{R_{t+1}} + \frac{C_{t+2} - Y_{t+2}}{R_{t+1} R_{t+2}} + \cdots + \lim_{j\to\infty} \frac{A_{t+j}}{R_{t+1}\cdots R_{t+j-1}}.$$

We assume that the transversality condition, $\lim_{j\to\infty} \frac{A_{t+j}}{R_{t+1}\cdots R_{t+j-1}} = 0$, holds. We the apply the transversality condition and divide all terms by Y_t:

$$\frac{A_t R_t}{Y_t} + 1 = \frac{C_t}{Y_t} + \frac{C_{t+1} - Y_{t+1}}{Y_t R_{t+1}} + \frac{C_{t+2} - Y_{t+2}}{Y_t R_{t+1} R_{t+2}} + \cdots.$$

Next, we introduce convenient notation to simplify our exposition, let:

$$\widetilde{c}_t \equiv \frac{C_t}{Y_t}, \quad G^Y_{t+1} \equiv \frac{Y_{t+1}}{Y_t}.$$

With this notation, we can rewrite the forward-iterated budget constraint as:

$$\begin{aligned}\frac{A_t R_t}{Y_t} + 1 = &\ \widetilde{c}_t + \widetilde{c}_{t+1} G^Y_{t+1} R^{-1}_{t+1} - G^Y_{t+1} R^{-1}_{t+1} \\ &+ \widetilde{c}_{t+2} G^Y_{t+1} G^Y_{t+2} R^{-1}_{t+1} R^{-1}_{t+2} - G^Y_{t+1} G^Y_{t+2} R^{-1}_{t+1} R^{-1}_{t+2} \\ &+ \cdots + (\widetilde{c}_T - 1) \prod_{j=1}^{T} G^Y_{t+j} R^{-1}_{t+j}.\end{aligned} \tag{16}$$

Step 2: Log-Linear Approximation. We log-linearly approximate the right-hand-side (RHS) of Equation (16) around $\ln \widetilde{c}_{t+j} = \ln \widetilde{c}$, $\ln G^Y_{t+j} = \ln G^Y$, and $\ln R_{t+j} = \ln R$. For convenience, we also introduce the following notation:

$$\rho \equiv \frac{G^Y}{R}.$$

We also define the value of the RHS at the approximation points

$$\kappa_0 \equiv \widetilde{c} + (\widetilde{c} - 1) \sum_{j=1}^{T} \rho^j.$$

With this, begin approximating the RHS:

$$\begin{aligned}\frac{A_t R_t}{Y_t} + 1 = &\kappa_0 + \widetilde{c} \left(\ln \widetilde{c}_t - \ln \widetilde{c}\right) + \rho \widetilde{c} \left(\ln \widetilde{c}_{t+1} - \ln \widetilde{c}\right) + \cdots + \rho^T \widetilde{c} \left(\ln \widetilde{c}_{t+T} - \ln c\right) \\ &+ (\widetilde{c} - 1) \left(\ln G^Y_{t+1} - \ln G^Y\right) \sum_{k=1}^{T} \rho^k + (\widetilde{c} - 1) \left(\ln G^Y_{t+2} - \ln G^Y\right) \sum_{k=2}^{T} \rho^k + \cdots + \rho^T (\widetilde{c} - 1) \left(\ln G^Y_{t+T} - \ln G^Y\right) \\ &- (\widetilde{c} - 1) \left(\ln R_{t+1} - \ln R\right) \sum_{k=1}^{T} \rho^k - (\widetilde{c} - 1) \left(\ln R_{t+2} - \ln R\right) \sum_{k=2}^{T} \rho^k - \cdots - \rho^T (\widetilde{c} - 1) \left(\ln R_{t+T} - \ln R\right).\end{aligned}$$

Step 3: Simplify. We next rewrite our approximation using summation notation:

$$\frac{A_t R_t}{Y_t} + 1 = \kappa_0 + \widetilde{c} \left(\ln \widetilde{c}_t - \ln \widetilde{c}\right) + \widetilde{c} \sum_{j=1}^{T} \rho^j \left(\ln \widetilde{c}_{t+j} - \ln \widetilde{c}\right) + (\widetilde{c} - 1) \sum_{j=1}^{T} \left\{ \left[\left(\ln G^Y_{t+j} - \ln G^Y\right) - \left(\ln R_{t+j} - \ln R\right) \right] \left(\sum_{k=j}^{T} \rho^k \right) \right\}.$$

We can further simplify the expression by collecting the terms related to the approximation points. To do so, let

$$\kappa_1 \equiv \kappa_0 - \widetilde{c} \ln \widetilde{c} \sum_{j=0}^{T} \rho^j - (\widetilde{c} - 1) \ln \rho \sum_{j=1}^{T} \left(\sum_{k=j}^{T} \rho^k \right).$$

Note that the initial values for the sums indexed by j differ above. Using this newly-defined parameter, we can write:

$$\frac{A_t R_t}{Y_t} + 1 = \kappa_1 + \widetilde{c} \ln \widetilde{c}_t + \widetilde{c} \sum_{j=1}^{T} \rho^j \ln \widetilde{c}_{t+j} + (\widetilde{c} - 1) \sum_{j=1}^{T} \left[\left(\ln G^Y_{t+j} - \ln R_{t+j} \right) \left(\sum_{k=j}^{T} \rho^k \right) \right].$$

□

D.1.3 Combining the Euler Equation and Budget Constraint

Now with log-linearized versions of both the budget constraint and Euler equation, we can combine them to obtain an approximate expression for frictionless APC.

Proposition 1: Characterization of Frictionless Consumption

Combining the log-linearized multi-period Euler equation, Equation (12), and the budget constraint, Equation (13) yields the following characterizing frictionless consumption:

$$\ln \widetilde{c}_t \approx \alpha_0 + \alpha_1 \frac{A_t R_t}{Y_t} + \sum_{j=1}^{T} \left\{ \left[\alpha_Y E_t \ln G^Y_{t+j} + \alpha_\pi E_t \ln \pi_{t+j} + \alpha_R E_t \ln R_{t+j} \right] \left(\sum_{k=j}^{T} \rho^k \right) \right\} \tag{17}$$

where

$$\begin{aligned}
\alpha_0 &= \left[1 - \kappa_1 - \widetilde{c} \frac{\ln \beta}{\gamma} \sum_{j=1}^{T} \left(\sum_{k=j}^{T} \rho^k \right) \right] \left(\widetilde{c} \sum_{j=0}^{T} \rho^j \right)^{-1} \\
\alpha_1 &= \left(\widetilde{c} \sum_{j=0}^{T} \rho^j \right)^{-1} \\
\alpha_Y &= \alpha_1 \\
\alpha_\pi &= -\alpha_Y \widetilde{c} \left(1 - \frac{1}{\gamma} \right) \\
\alpha_R &= -\alpha_Y \left(1 - \widetilde{c} + \frac{\widetilde{c}}{\gamma} \right).
\end{aligned}$$

Proof. To combine the equations, we proceed in three steps.

1. Take expectations of the approximated budget constraint.
2. Substitute period j's log-APC (i.e., $\ln \left(\frac{C_{t+j}}{Y_{t+j}} \right)$) into the budget constraint via the multi-period Euler equation.
3. 3. Rearrange to isolate $\ln \left(\frac{C_t}{Y_t} \right)$.

Step 1: Take Expectations of the Budget Constraint. We start by taking expectations of the approximated budget constraint:

$$\frac{A_t R_t}{Y_t} + 1 \approx \kappa_1 + \widetilde{c} \ln \widetilde{c}_t + \widetilde{c} \sum_{j=1}^{T} \rho^j E_t \ln \widetilde{c}_{t+j} + (\widetilde{c} - 1) \sum_{j=1}^{T} \left[\left(E_t \ln G^Y_{t+j} - E_t \ln R_{t+j} \right) \left(\sum_{k=j}^{T} \rho^k \right) \right].$$

Step 2: Substitute in the Euler Equation. Recall that the multi-period log-linearized Euler equation is:

$$\ln\left(\frac{c_t}{y_t}\right) = \sum_{k=1}^{j}\left[E_t \ln G^Y_{t+k} - E_t \ln \pi_{t+k} - \frac{1}{\gamma}\left(\ln\beta + E_t \ln R_{t+k} - E_t \ln \pi_{t+k}\right)\right] + E_t \ln\left(\frac{c_{t+j}}{y_{t+j}}\right).$$

Rearranging and applying our simplifying notation, we can isolate period j's expected log-APC:

$$E_t \ln \widetilde{c}_{t+j} = \ln \widetilde{c}_t - \sum_{k=1}^{j}\left[E_t \ln G^Y_{t+k} - E_t \ln \pi_{t+k} - \frac{1}{\gamma}\left(\ln\beta + E_t \ln R_{t+k} - E_t \ln \pi_{t+k}\right)\right].$$

We can now plug this into the budget constraint using $\widetilde{c}_t = \ln\left(\frac{c_t}{y_t}\right) = \ln\left(\frac{C_t}{Y_t}\right)$. To do so, we start by simplifying the terms related to APCs:

$$\begin{aligned}
\widetilde{c}\sum_{j=1}^{T}\rho^j E_t \ln \widetilde{c}_{t+j} =& \widetilde{c}\sum_{j=1}^{T}\rho^j\left\{\ln \widetilde{c}_t - \sum_{k=1}^{j}\left[E_t \ln G^Y_{t+k} - E_t \ln \pi_{t+k} - \frac{1}{\gamma}\left(\ln\beta + E_t \ln R_{t+k} - E_t \ln \pi_{t+k}\right)\right]\right\} \\
=& \widetilde{c}\ln \widetilde{c}_t\sum_{j=1}^{T}\rho^j - \widetilde{c}\sum_{j=1}^{T}\rho^j\left\{\sum_{k=1}^{j}\left[E_t \ln G^Y_{t+k} - E_t \ln \pi_{t+k} - \frac{1}{\gamma}\left(\ln\beta + E_t \ln R_{t+k} - E_t \ln \pi_{t+k}\right)\right]\right\}.
\end{aligned}$$

To reverse the order of summation, we use $\sum_{j=1}^{T}\left[\rho^j\left(\sum_{k=1}^{j} x_k\right)\right] = \sum_{j=1}^{T}\left[x_j\left(\sum_{k=j}^{T}\rho^k\right)\right]$:

$$\widetilde{c}\sum_{j=1}^{T}\rho^j E_t \ln \widetilde{c}_{t+j} = \widetilde{c}\ln \widetilde{c}_t\sum_{j=1}^{T}\rho^j - \widetilde{c}\sum_{j=1}^{T}\left\{\left[E_t \ln G^Y_{t+j} - E_t \ln \pi_{t+j} - \frac{1}{\gamma}\left(\ln\beta + E_t \ln R_{t+j} - E_t \ln \pi_{t+j}\right)\right]\left(\sum_{k=j}^{T}\rho^k\right)\right\}.$$

Next, we plug this back into our combined Euler equation and budget constraint:

$$\begin{aligned}
\frac{A_t R_t}{Y_t} + 1 \approx& \kappa_1 + \widetilde{c}\ln \widetilde{c}_t + \widetilde{c}\ln \widetilde{c}_t\sum_{j=1}^{T}\rho^j + (\widetilde{c}-1)\sum_{j=1}^{T}\left[\left(E_t \ln G^Y_{t+j} - E_t \ln R_{t+j}\right)\left(\sum_{k=j}^{T}\rho^k\right)\right] \\
&- \widetilde{c}\sum_{j=1}^{T}\left\{\left[E_t \ln G^Y_{t+j} - E_t \ln \pi_{t+j} - \frac{1}{\gamma}\left(\ln\beta + E_t \ln R_{t+j} - E_t \ln \pi_{t+j}\right)\right]\left(\sum_{k=j}^{T}\rho^k\right)\right\}.
\end{aligned}$$

We then proceed to group the terms related to the discount factor, APCs, and beliefs:

$$\begin{aligned}
\frac{A_t R_t}{Y_t} + 1 \approx& \kappa_1 + \widetilde{c}\frac{\ln\beta}{\gamma}\sum_{j=1}^{T}\left(\sum_{k=j}^{T}\rho^k\right) + \widetilde{c}\ln \widetilde{c}_t\sum_{j=0}^{T}\rho^j \\
&- \sum_{j=1}^{T}\left\{\left[E_t \ln G^Y_{t+j} - \widetilde{c}\left(1-\frac{1}{\gamma}\right)E_t \ln \pi_{t+j} - \left(1-\widetilde{c}+\frac{\widetilde{c}}{\gamma}\right)E_t \ln R_{t+j}\right]\left(\sum_{k=j}^{T}\rho^k\right)\right\}.
\end{aligned}$$

Step 3: Solve for $\ln \widetilde{c}_t$. Next, we begin to rearrange our combined approximation to isolate $\ln(\widetilde{c}_t)$. To simplify our expression, we introduce five more parameters:

$$\begin{aligned}
\alpha_0 &= \left[1 - \kappa_1 - \widetilde{c}\frac{\ln\beta}{\gamma}\sum_{j=1}^{T}\left(\sum_{k=j}^{T}\rho^k\right)\right]\left(\widetilde{c}\sum_{j=0}^{T}\rho^j\right)^{-1} \\
\alpha_1 &= \left(\widetilde{c}\sum_{j=0}^{T}\rho^j\right)^{-1} \\
\alpha_Y &= \alpha_1 \\
\alpha_\pi &= -\alpha_Y\widetilde{c}\left(1-\frac{1}{\gamma}\right) \\
\alpha_R &= -\alpha_Y\left(1-\widetilde{c}+\frac{\widetilde{c}}{\gamma}\right).
\end{aligned}$$

With these definitions, we write our approximation as:

$$\ln\widetilde{c}_t \approx \alpha_0 + \alpha_1\frac{A_t R_t}{Y_t} + \sum_{j=1}^{T}\left\{\left[\alpha_Y E_t \ln G^Y_{t+k} + \alpha_\pi E_t \ln \pi_{t+k} + \alpha_R E_t \ln R_{t+k}\right]\left(\sum_{k=j}^{T}\rho^k\right)\right\}.$$

□

Proposition D.1.3 provides a general characterization of frictionless consumption in that it does not specify (1) whether the agent is finite or infinitely-lived, (2) whether beliefs are constant or vary across horizons j, (3) the approximation points used in the budget constraint approximation.

The results below present simplified characterization of our general frictionless consumption Equation (17) under scenarios with more restrictive assumption. We begin by characterizing frictionless consumption under an infinite horizon (i.e., $T \to \infty$). This assumption primarily affects the parameters in the approximation.

Lemma 3: Frictionless Consumption Under an Infinite Horizon

If $T \to \infty$ and $\rho = \frac{G^Y}{R} < 1$, then frictionless consumption is

$$\ln \widetilde{c}_t \approx \alpha_0 + \alpha_1 \frac{A_t R_t}{Y_t} + \sum_{j=1}^{T} \left\{ \left[\alpha_Y E_t \ln G^Y_{t+j} + \alpha_\pi E_t \ln \pi_{t+j} + \alpha_R E_t \ln R_{t+j} \right] \left(\frac{\rho^j}{1-\rho} \right) \right\} \tag{18}$$

where

$$\begin{aligned}
\kappa_0 =& \frac{\widetilde{c} - \rho}{1-\rho} \\
\kappa_1 =& \kappa_0 - \frac{\widetilde{c} \ln \widetilde{c}}{1-\rho} - (\widetilde{c} - 1) \ln \rho \left[\frac{\rho}{(1-\rho)^2} \right] \\
\alpha_0 =& (1-\kappa_1) \left(\frac{1-\rho}{\widetilde{c}} \right) - \frac{\ln \beta}{\gamma} \frac{\rho}{(1-\rho)} \\
\alpha_1 =& \frac{1-\rho}{\widetilde{c}} \\
\alpha_Y =& \alpha_1 \\
\alpha_\pi =& -\alpha_Y \widetilde{c} \left(1 - \frac{1}{\gamma} \right) \\
\alpha_R =& -\alpha_Y \left(1 - \widetilde{c} + \frac{\widetilde{c}}{\gamma} \right).
\end{aligned}$$

Proof. We start by characterizing terms that appear frequently in the parameters, first:

$$\begin{aligned}
\lim_{T\to\infty} \sum_{j=1}^{T} \rho^j &= \frac{\rho}{1-\rho} \\
\lim_{T\to\infty} \sum_{j=1}^{T} \left(\sum_{k=j}^{T} \rho^k \right) &= \frac{\rho}{(1-\rho)^2} \\
\lim_{T\to\infty} \sum_{k=j}^{T} \rho^k &= \frac{\rho^j}{1-\rho}.
\end{aligned}$$

With this, we can begin to simplify the expressions for the parameter values. We start with κ_0:

$$\begin{aligned}
\kappa_0 &= \lim_{T\to\infty} \left[\widetilde{c} + (\widetilde{c} - 1) \sum_{j=1}^{T} \rho^j \right] \\
&= \widetilde{c} + (\widetilde{c} - 1) \frac{\rho}{1-\rho} \\
&= \frac{\widetilde{c} - \rho}{1-\rho}.
\end{aligned}$$

We next turn to κ_1:

$$\kappa_1 \equiv \lim_{T\to\infty}\left[\kappa_0 - \widetilde{c}\ln\widetilde{c}\sum_{j=0}^{T}\rho^j - (\widetilde{c}-1)\ln\rho\sum_{j=1}^{T}\left(\sum_{k=j}^{T}\rho^k\right)\right]$$
$$= \kappa_0 - \frac{\widetilde{c}\ln\widetilde{c}}{1-\rho} - (\widetilde{c}-1)\ln\rho\left[\frac{\rho}{(1-\rho)^2}\right].$$

We now simplify the remaining parameters with infinite sums:

$$\alpha_0 = \lim_{T\to\infty}\left\{\left[1-\kappa_1-\widetilde{c}\frac{\ln\beta}{\gamma}\sum_{j=1}^{T}\left(\sum_{k=j}^{T}\rho^k\right)\right]\left(\widetilde{c}\sum_{j=0}^{T}\rho^j\right)^{-1}\right\}$$
$$= \left[1-\kappa_1-\widetilde{c}\frac{\ln\beta}{\gamma}\frac{\rho}{(1-\rho)^2}\right]\left(\frac{1-\rho}{\widetilde{c}}\right)$$
$$= (1-\kappa_1)\left(\frac{1-\rho}{\widetilde{c}}\right) - \frac{\ln\beta}{\gamma}\frac{\rho}{(1-\rho)}$$

and

$$\alpha_1 = \lim_{T\to\infty}\left(\widetilde{c}\sum_{j=0}^{T}\rho^j\right)^{-1} = \frac{1-\rho}{\widetilde{c}}.$$

Therefore, when $T\to\infty$, frictionless consumption is

$$\ln\widetilde{c}_t \approx \alpha_0 + \alpha_1\frac{A_tR_t}{Y_t} + \sum_{j=1}^{T}\left\{\left[\alpha_Y E_t\ln G^Y_{t+k} + \alpha_\pi E_t\ln\pi_{t+k} + \alpha_R E_t\ln R_{t+k}\right]\left(\frac{\rho^j}{1-\rho}\right)\right\}.$$

where the α parameters are defined above. □

Lemma D.1.3 characterizes frictionless consumption under an infinite horizon where beliefs can vary over horizons j. If data on beliefs over more distant horizons is lacking and/or we are willing to assume that beliefs are (approximately) constant, we can simplify this characterization further. This simplification is formalized in the remark below.

Remark 3: Frictionless Consumption Under an Infinite Horizon and Constant Beliefs

If beliefs are constant (i.e., $E_t\ln\pi_{t+j} = E_t\ln\pi_{t+k}$ for all $j,k>0$, etc.) frictionless consumption is characterized by

$$\ln\widetilde{c}_t \approx \alpha_0 + \alpha_1\frac{A_tR_t}{Y_t} + \overline{\alpha_Y}E_t\ln G^Y_{t+1} + \overline{\alpha_\pi}E_t\ln\pi_{t+1} + \overline{\alpha_R}E_t\ln R_{t+1} \tag{19}$$

where

$$\overline{\alpha_Y} = \alpha_Y\frac{\rho}{(1-\rho)^2} \tag{20}$$
$$\overline{\alpha_R} = \alpha_R\frac{\rho}{(1-\rho)^2} \tag{21}$$
$$\overline{\alpha_\pi} = \alpha_\pi\frac{\rho}{(1-\rho)^2}. \tag{22}$$

D.2 Extension: Durable and Non-Durable Goods

We next show how to extend our wedge measurement results to accommodate durable goods. Durable goods present several complications: it's difficult to measure their consumption and depreciation directly, holdings of durable goods constitute a source of wealth, and some are financed with debt. To overcome these challenges, we make assumptions that imply that the expenditure share of non-durable goods is a constant, known fraction. The key assumption is that notional consumption is a Cobb Douglas aggregate of both types of consumption goods.

Notation. Let n_t and d_t denote real period t consumption flows of non-durable and durable goods (respectively). We continue to denote the total nominal value of net worth by A_t. Total wealth includes net positions in durables (e.g., the value of vehicles net of the loans used to finance their purchase). The household has preferences over notional consumption flows c_t, which are an aggregate of non-durable and durable consumption flows (i.e., utility $u(c_t)$ is the per-period utility flow).

We make two assumptions.

Assumption 1: Frictionless Spot and Rental Markets for Durables and No Arbitrage.

In our frictionless benchmark, the household can frictionlessly buy or sell durables at a spot price. The household can also rent durable goods at a per period rental cost of q_t. No arbitrage in durable goods markets requires that the rental price q_t equal the user cost of the durable goods.

By assuming that households can frictionlessly transact in our benchmark, the wedge we estimate is able to capture frictions on adjusting the stock of durables. The no arbitrage assumption means that the household is indifferent between holding and accumulating durables versus renting them. This allows us to simplify our exposition while keeping the user cost of durables flexible. The user cost reflects depreciation, forgone interest earnings/savings, and appreciation of durable goods prices.

We let non-durables, n_t, be the numeraire good. Under Assumption D.2, we can write the household's budget constraint simply as

$$A_{t+1} + P_t c_t = Y_t + A_t R_t$$

where

$$P_t c_t = n_t + q_t d_t.$$

and P_t is the ideal price index. The budget constraint is isomorphic to our original budget constraint. The Euler equation remains unchanged as well, where c_t now corresponds to notional consumption. Therefore, the intertemporal optimality conditions presented in Section 2 remain unchanged. There are now simply additional first order conditions for intratepmoral optimality with respect to the allocation of spending between non-durable and durable consumption.

Assumption 2: Cobb Douglas Aggregation.

The household's notional consumption good is a Cobb Douglas aggregate of non-durable and durable consumption flows:

$$c_t = n_t^{\alpha} d_t^{1-\alpha}.$$

Under Assumptions D.2 and D.2, the intratemporal optimality conditions are:

$$n_t = \alpha P_t c_t$$
$$d_t q_t = (1-\alpha) P_t c_t.$$

The intratemporal optimality conditions indicate that expenditure on each good is a constant share of total expenditures on consumption goods. As a result, we can infer the notional APC (i.e., for all consumption) from the non-durable APC and the expenditure share. This is formalized in the lemma below.

Lemma 4: APC Calculation Including Consumption of Durables.

Under Assumptions D.2 and D.2, the APC (including both non-durable and durable consumption) is

$$\frac{P_t c_t}{Y_t} = \frac{n_t}{Y_t}\frac{1}{\alpha} \tag{23}$$

where α corresponds to the non-durable share of expenditures.

In our baseline analysis, in order to characterize deviations in total consumption, we multiply non-durable APCs by an estimate of $\frac{1}{\alpha}$.

D.3 Calibration of Wedge Parameters

We first characterize conditions under which we can employ a convenient simplification of the wedge parameters. We obtain this result when we assume our log-linearization approximation points corresponds to steady-state values.

Remark 4: Simplification of ρ

If we approximate the budget constraint around $\left\{\frac{C}{Y}, \frac{AR}{Y}, \rho\right\}$, then the (forward-iterated) budget constraint holds for these values:

$$\frac{C}{Y} = 1 + \frac{AR}{Y} + \left(1 - \frac{C}{Y}\right)\frac{\rho}{1-\rho}.$$

This implies

$$\rho = \frac{1 + \frac{AR}{Y} - \frac{C}{Y}}{\frac{AR}{Y}}. \tag{24}$$

Table D.1. Calibrated Parameter Values and their Sources

Parameter	Value	Meaning	Source
$\frac{C}{Y} \equiv \tilde{c}$	84.92%	Steady state ratio of consumption expenditures to income	Median ratio of non-durable spending to income in EarnIn sample (67.40%) divided by non-durable share of expenditures 79.37% (calculations using the Consumer Expenditure Survey in Beraja and Zorzi, 2024)
$\frac{AR}{Y}$	-46.70%	Steady state ratio of net worth to income	Median ratio of net worth to income in EarnIn sample
$\rho = \frac{g^Y}{R}$	67.72%	Steady state ratio of income growth and return to saving	Calculation (approximating around steady state implies $\rho = \frac{1+\frac{AR}{Y}-\frac{C}{Y}}{\frac{AR}{Y}}$)
γ	2	Coefficient of relative risk aversion	Standard value
β	0.92	Annual discount factor	Standard value

Notes: This table presents parameters used in the wedge analysis. It details the values used in our preferred specification, the economic meaning of the parameters, and the source of the chosen value.

Table D.2. Calculated Wedge Coefficients

Coefficient	Value	Formula	Multiplicand
κ_0	0.533	$\tilde{c} + (\tilde{c}-1)\frac{\rho}{1-\rho}$	Intermediate parameter
κ_1	0.581	$\kappa_0 - \tilde{c}\frac{\ln\tilde{c}}{1-\rho} - (\tilde{c}-1)\ln\rho\left[\frac{\rho}{(1-\rho)^2}\right]$	Intermediate parameter
α_Y	0.380	$\frac{1-\rho}{\tilde{c}}$	Intermediate parameter
α_π	-0.161	$-\alpha_Y\tilde{c}\left(1-\frac{1}{\gamma}\right)$	Intermediate parameter
α_R	-0.219	$-\alpha_Y\left(1-\tilde{c}+\frac{\tilde{c}}{\gamma}\right)$	Intermediate parameter
α_1	0.380	$\frac{1-\rho}{\tilde{c}}$	$\frac{A_tR_t}{Y_t}$
$\overline{\alpha_Y}$	2.470	$\alpha_Y\frac{\rho}{(1-\rho)^2}$	$E_t \ln G^Y_{t+1}$
$\overline{\alpha_\pi}$	-1.049	$\alpha_\pi\frac{\rho}{(1-\rho)^2}$	$E_t \ln \pi_{t+1}$
$\overline{\alpha_R}$	-1.421	$\alpha_R\frac{\rho}{(1-\rho)^2}$	$E_t \ln R_{t+1}$

Notes: This table presents parameters used in the wedge analysis. It details the values used in our preferred specification, the economic meaning of the parameters, and the source of the chosen value.

E Sensitivity Analysis

We test the robustness of the consumption wedge to our assumed parameters and data choices. For this analysis, we focus on the sensitivity of the share of users with a positive wedge (i.e., overconsumers), the median wedge in absolute value terms, and the correlations with nondurables MPCs and financial distress (specifically, the indicator for high financial anxiety).

First, we varying the assumed econometric preference parameters. As outlined in Appendix D.3, our baseline specification assumes a beta of 0.92 and a gamma of 2.0, which are standard values in the literature. In Appendix Figures E.1 and E.2, we vary beta from 0.80 to 0.98 and gamma from 1.0 to 5.0.

Second, we vary the assumed approximation points for the nondurable APC, the nondurable share of spending, and AR/Y. In our baseline specification, we assume a 69.9% nondurable APC, a 79.4% nondurable share of spending. and a -41.4% ratio of net worth to income. Appendix Figures **??** and E.5 vary the nondurable APC from 0.55 to 0.75 and the nondurable share of spending from 0.72 to 0.90. Appendix Figure E.4 varies the AR/Y assumption from -81% to -31%.

Third, we restrict our sample to users with higher quality measurement of income, spending, wealth, and beliefs. These results are presented in Appendix Figure E.6. We test dropping nine different sets of users: (1) users with UI income, which may not be incorporated into income growth forecasts; (2) users with zero observed income in at least one month; (3) users with inflation expectations divisible by 5pp, as these may be rounded; (4) users who fail at least one financial literacy question, as these users may have difficulty understanding the survey questions; (5) users who completed the survey in less than 6 minutes; (6) users whose peer-to-peer transfers exceed 50% nondurable spending; (7) users whose credit card payments exceed 50% of nondurable spending; (8) users whose cash withdrawals exceed 25% of nondurable spending; and (9) users whose observed durables spending exceeds 25% of nondurables spending. Restrictions (6) through (9) are intended to drop users that may have significant spending that is not captured in the transactions data.

Fourth, we test the robustness of our main wedge results to clustering users and taking the within-cluster median of each input before calculating the wedge. These results are shown in Appendix Figure **??**. We cluster users on reported age, income, and savings, and indicators for gender (male, female, or other), race (white, non-white, or not specified), spouse or partner, having children, college education, and political affiliation (Democrat, Republican, or other). To accommodate both numeric and categorical variables, we employ the k-Prototype clustering algorithm, which combines the k-means and k-modes clustering algorithms.[18] We vary the number of clusters from 100 to 500.

Fifth, we test the robustness of our wedge results to adding random noise to each input of the wedge calculation (nondurables APC, reported wealth-to-income, and expectations). We conduct a Monte Carlo simulation with 1,000 iterations. In each iteration, we generate six normal random variables (corresponding to each wedge parameter) with a mean of 0 percentage points and a standard deviation of 2 percentage points, and then add these to each corresponding input. In line with the observation levels of the data, noise is generated at the at the user level for all wedge inputs. After adding the noise, we calculate wedges and reproduce our key results, shown in E.7.

[18] The k-Prototype algorithm combines the k-means and k-modes clustering algorithms to accommodate both numeric and categorical variables. For numeric variables, distance is calculated using standard Euclidean distance, and the prototype (or center) of each cluster is the mean of all points within the cluster. For categorical variables, distance equals zero if categories match and one if not, and the prototype is the most frequent category (mode) across all points within the cluster. In each iteration of the algorithm, observations are reassigned to the cluster with the closest prototypes, and iterations continue until cluster assignments converge.

Figure E.1. Sensitivity of Dynamic Wedges to Beta

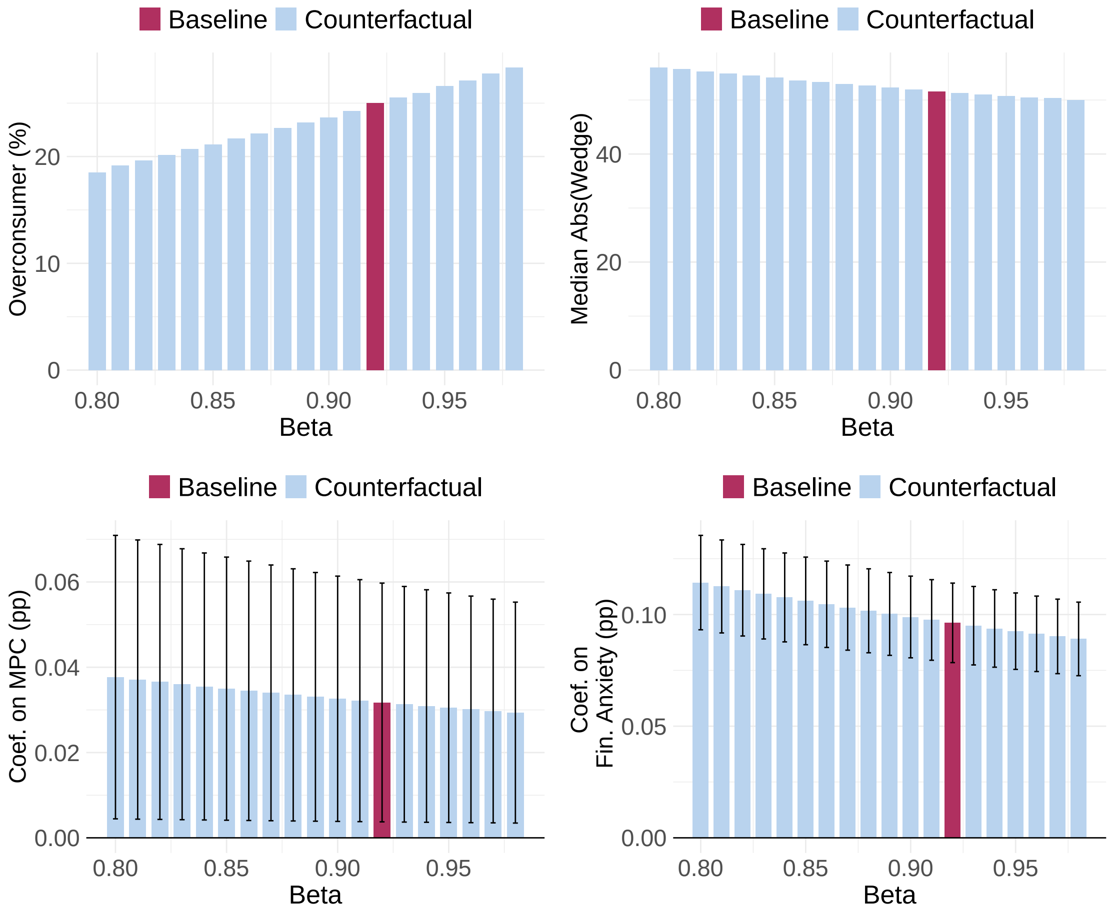

Notes: The figure presents the sensitivity of four estimated results to our assumed value of beta in the dynamic wedge calculation, with the assumed value ranging from 0.80 to 0.98 in increments of 0.01 (our baseline calibration is 0.92). These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). We hold all other parameters constant at our baseline values. Includes users across all three surveys waves who meet the sample restrictions outlined in Appendix B.

Figure E.2. Sensitivity of Dynamic Wedges to Gamma

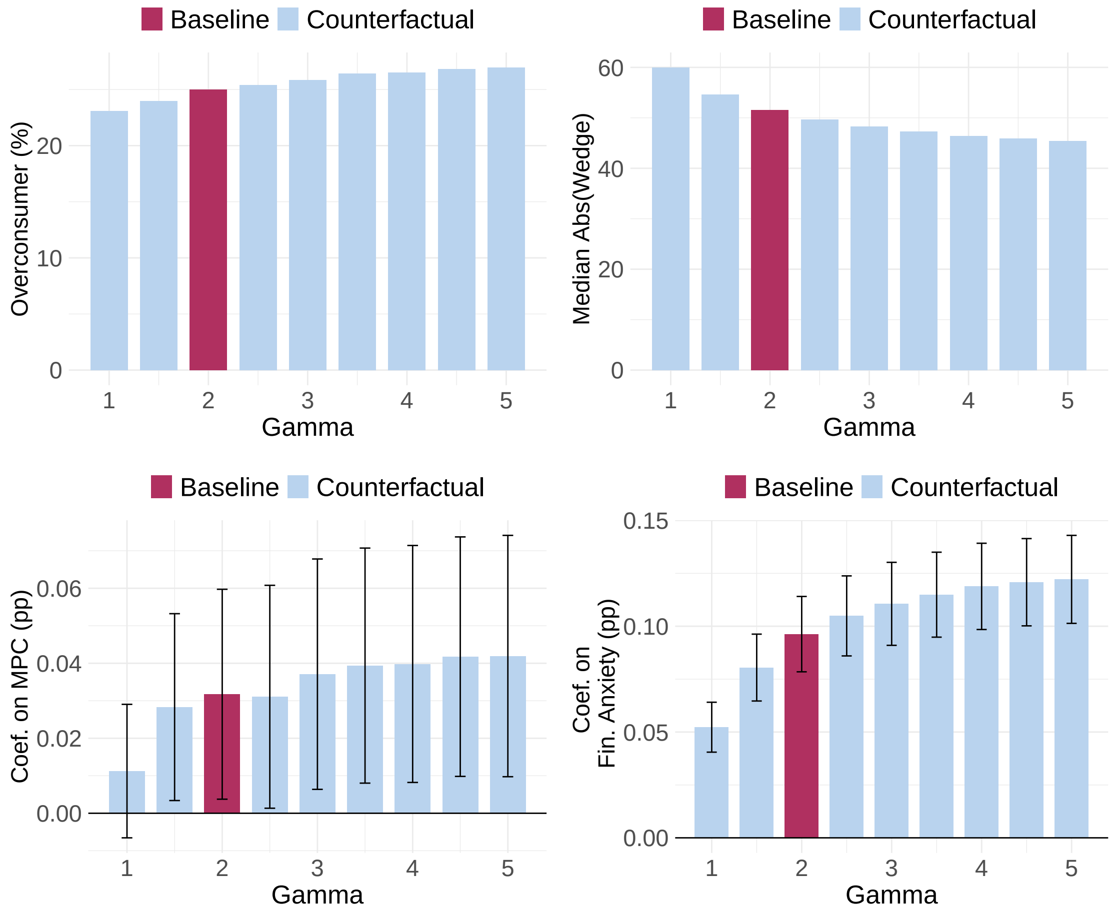

Notes: The figure presents the sensitivity of four estimated results to our assumed value of gamma in the dynamic wedge calculation, with the assumed value ranging from 1.0 to 5.0 in increments of 0.5 (our baseline calibration is 2.0). These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). We hold all other parameters constant at our baseline values. Includes users across all three surveys waves who meet the sample restrictions outlined in Appendix B.

Figure E.3. Sensitivity of Dynamic Wedges to Approximation Point for Nondurable APC

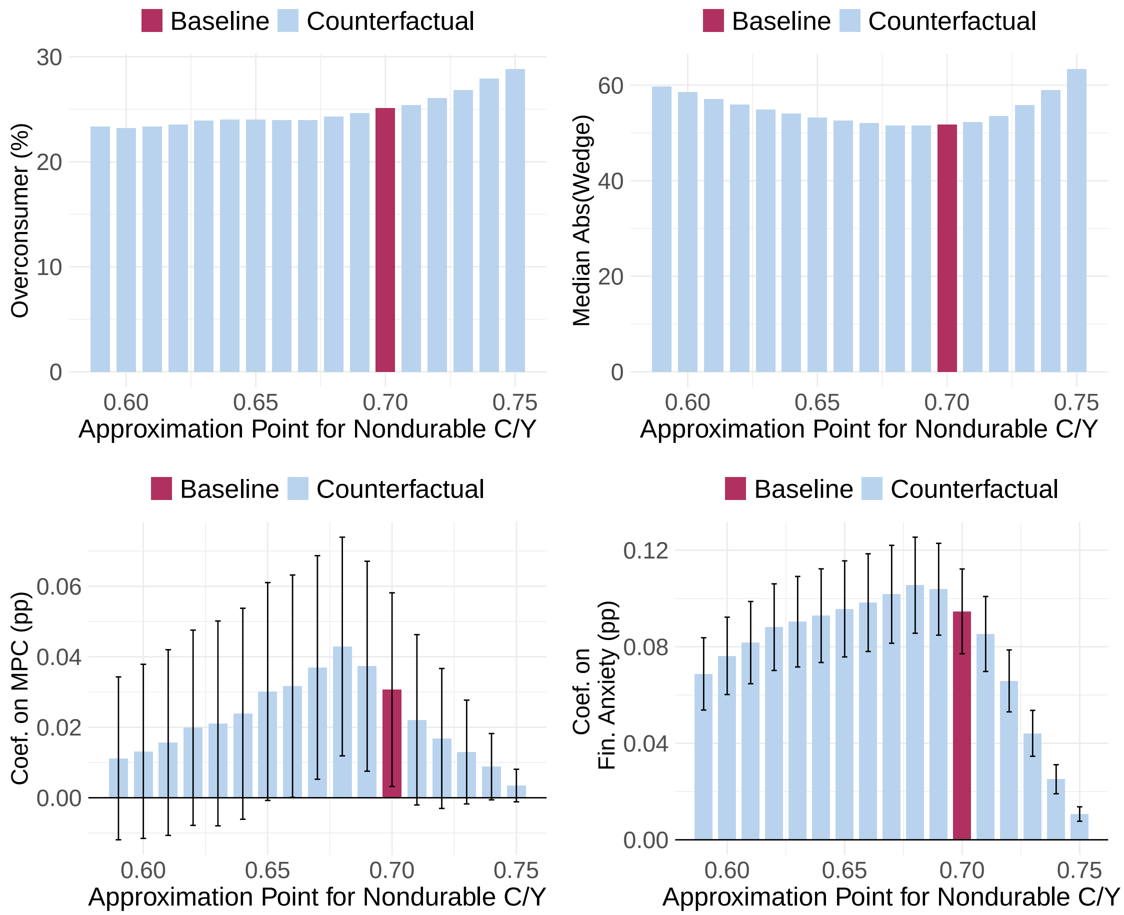

Notes: The figure presents the sensitivity of four estimated results to our approximation point for nondurable APC in the dynamic wedge calculation, with the assumed value ranging from 55% to 75% in increments of 1 percentage point (our baseline calibration is 69.87%). These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). We hold all other parameters constant at our baseline values. Includes users across all survey waves who meet the sample restrictions outlined in Appendix B.

Figure E.4. Sensitivity of Wedges to Approximation Point for AR/Y

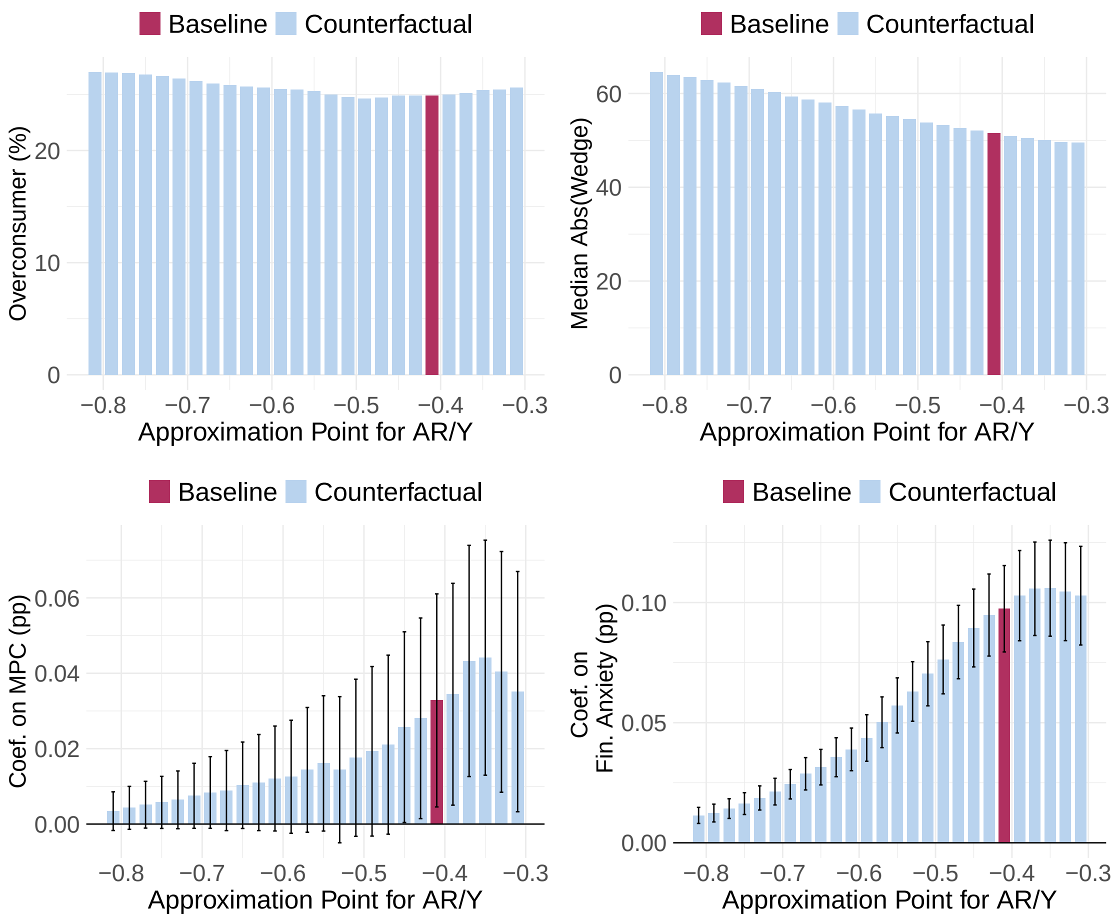

Notes: The figure presents the sensitivity of four estimated results to our approximation point for AR/Y in the dynamic wedge calculation, with the assumed value ranging from -81% to -31% in increments of 2 percentage points (our baseline calibration is -41.39%). These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). We hold all other parameters constant at our baseline values. Includes users across all three surveys waves who meet the sample restrictions outlined in Appendix B.

Figure E.5. Sensitivity of Dynamic Wedges to Nondurables Share of Spending

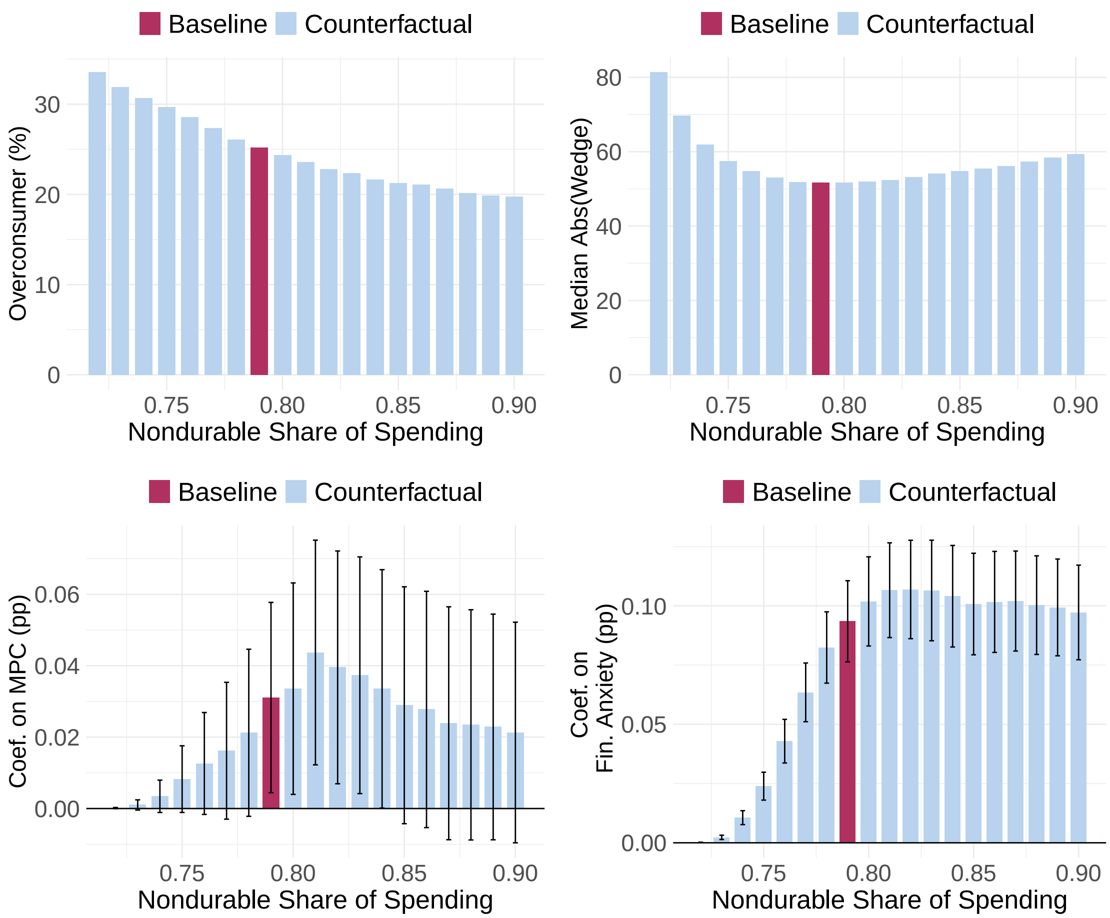

Notes: The figure presents the sensitivity of four estimated results to our assumed nondurables share of spending in the dynamic wedge calculation, with the assumed value ranging from 72% to 90% in increments of 1 percentage point (our baseline calibration is 79.37%). These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). We hold all other parameters constant at our baseline values. Includes users across all three surveys waves who meet the sample restrictions outlined in Appendix B.

Figure E.6. Sensitivity of Consumption Wedge to Dropping Users

Figure presents the sensitivity of four estimated results to dropping users with plausibly high measurement error. These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). In the regressions, we cluster standard errors at the user level. For reference, our baseline results are shown in red. Baseline sample includes wave 1 respondents who meet the sample restrictions outlined in Appendix B.

Figure E.7. Sensitivity of Consumption Wedge to Adding Random Noise

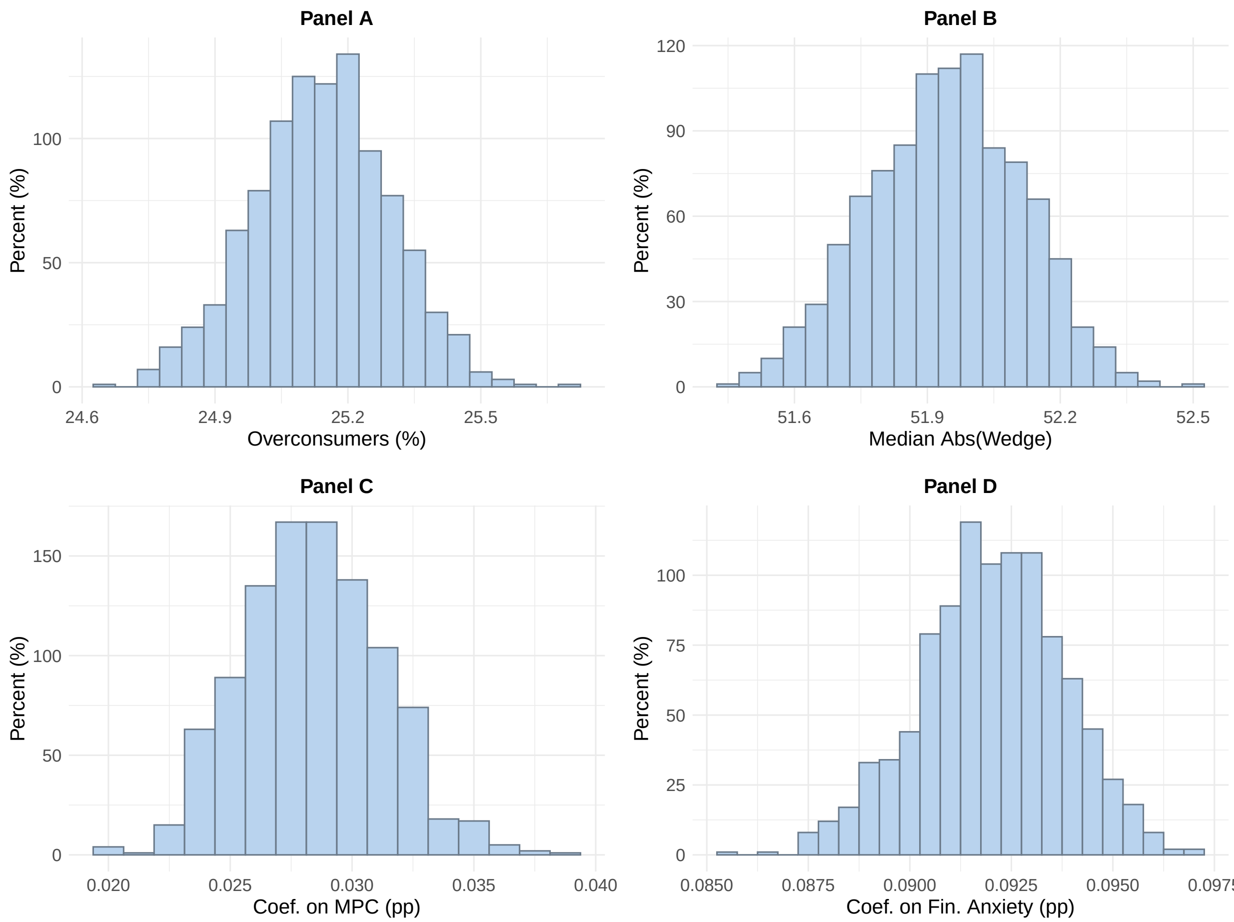

Figure presents the sensitivity of four estimated results to adding random noise with a mean of 0pp and standard deviation of 2pp to the wedge inputs. These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of nondurables MPC and an indicator for financial anxiety on the consumption wedge (Panels C and D, respectively). Each panel represents a histogram of the estimated parameter from a Monte Carlo simulation with 1,000 iterations. Includes users across all three surveys waves who meet the sample restrictions outlined in Appendix B.

Figure E.8. Comparison of First- and Second-Wave User-Level Consumption Wedges

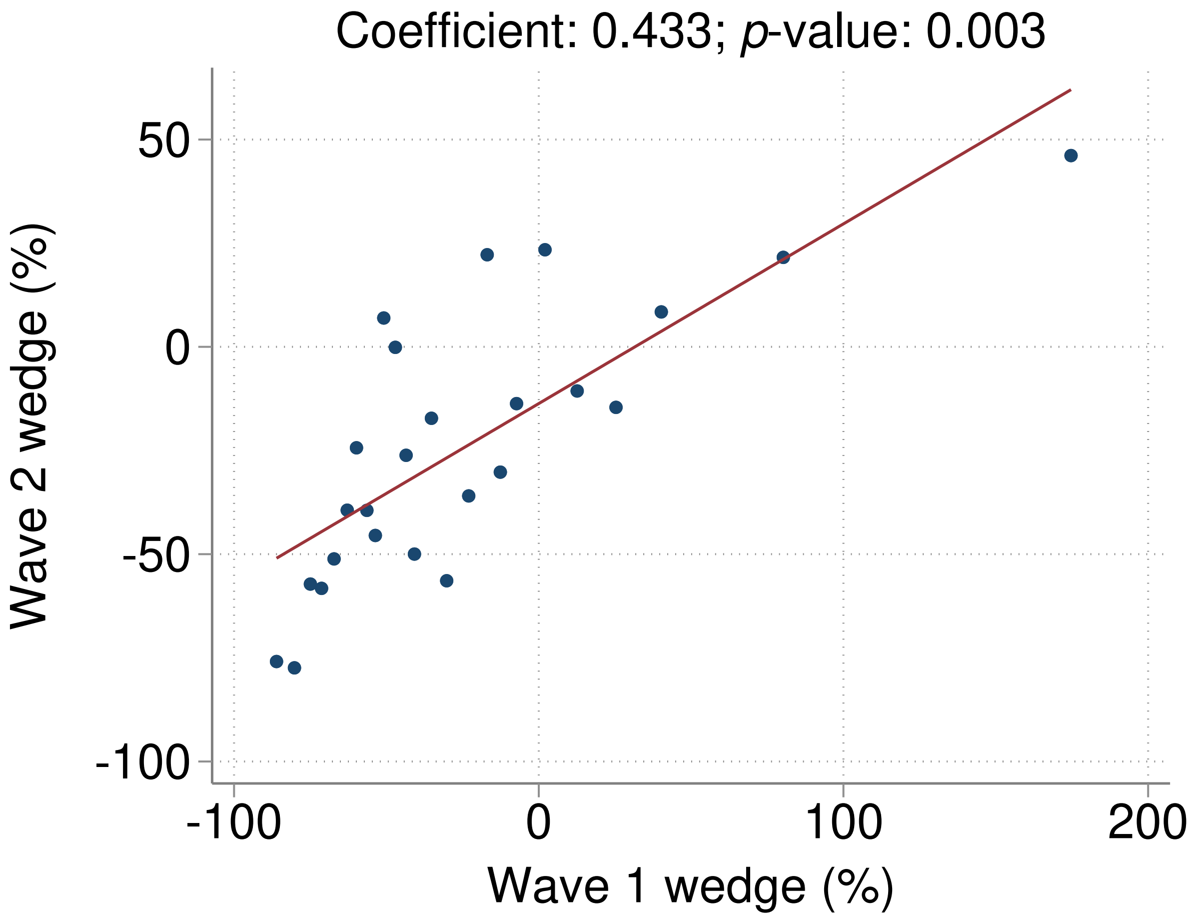

Figure presents the correlation between user-level wedges across survey waves as a binscatter with 25 quantile bins. Includes wave 1 and wave 2 respondents who meet the sample restrictions outlined in Appendix B. Wave 1 wedges include data from October 2021 to September 2022, and wave 2 wedges include data from July 2023 to June 2024.