Consumption Wedges: Measuring and Diagnosing Distortions *

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

September 10, 2025

Abstract

Deviations from canonical consumption-savings models have been attributed to a wide variety of distortions, including financial constraints and behavioral preferences. We develop a sufficient statistics approach to measure the impact of such distortions on consumption as a *consumer-level* wedge between actual consumption and counterfactual "frictionless" consumption. Because different distortions imply different wedge properties, wedges can be used as a diagnostic to distinguish between models. We apply our approach to data linking individuals' economic beliefs to administrative transactions data for a population of predominantly middle-income US consumers with low liquid wealth. By using subjective beliefs data, we can isolate the influence of frictions and behavioral preferences from deviations from full-information rational expectations. We find that consumption is significantly distorted both upward and downward. The median consumer's wedge is 40% of frictionless consumption in absolute value, with 49% of consumers under-consuming (negative wedges) and 51% over-consuming (positive wedges). Since financial constraints generate only negative wedges, additional or alternative distortions (such as present bias or consumer inertia) are necessary to rationalize the consumption decisions of low-liquidity consumers.

*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 (discussant), Francesco D'Acunto (discussant), Tim de Silva (discussant), Marty Eichenbaum, Ben Enke, Joao Gomes, Yuriy Gorodnichenko, Fatih Guvenen, Erik Hurst, Ben Keys, Dirk Krueger, Igor Livshits, Pierre Mabille (discussant), Peter Maxted, Emi Nakamura, Gisle Natvik (discussant), Matt Notowidigdo, Michael Roberts, Nick Roussanov, and Jón Steinsson. Nikki Azerang, Jacky Chen, Chazel Hakim, Clint Hamilton, Eleanor Jenke, Noah Jenkins, 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 and NBER, 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 consumption-savings behavior in macroeconomics and household finance (e.g., Kaplan and Violante, 2014). A key motivation for the focus on financial constraints is extensive empirical evidence of deviations from the permanent income hypothesis, such as high marginal propensities to consume (MPCs) out of transitory income changes, especially among consumers with low liquid wealth (e.g., Johnson, Parker and Souleles, 2006).[1] However, alternative distortions—such as present bias, consumption adjustment costs, and bounded rationality—can also account for these empirical patterns (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.[2] Empirical evidence on the types and size of consumption distortions is therefore crucial to guide theories of what determines the consumption-savings behavior of low-liquidity consumers, as well as policies that affect their wellbeing and the broader macroeconomy.

We develop a new sufficient statistics approach to measure consumer-level "wedges" between actual observed consumption and the hypothetical "frictionless" consumption implied by an Euler equation and budget constraint in the absence of distortions. These consumption wedges quantify the total net 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). Different distortions have different predictions for the properties of consumption wedges, such as their sign, correlates, and responses to shocks. As a result, wedges can be used as a diagnostic to discriminate between models of consumer behavior.

We make two innovations relative to existing research studying aggregate wedges (e.g., Chari, Kehoe and McGrattan, 2007; Berger, Bocola and Dovis, 2023). First, we measure the distribution of consumer (micro-level) wedges rather than aggregate (macro-level) wedges.[3] This proves to be important: the median wedge is 1.4%, but the median *absolute value* wedge is 40% of frictionless consumption with similar shares consuming above and below the benchmark. Second, we use individual-level subjective beliefs data, which allows us to avoid assuming consumers have

[1] Other influential empirical patterns include the excess smoothness of consumption out of persistent income shocks (e.g., Campbell and Deaton, 1989; Attanasio and Pavoni, 2011) and lack of correlation between the real interest rate and expected consumption growth (e.g., Campbell and Mankiw, 1989; Attanasio and Weber, 1993).

[2] For example, Lee and Maxted (2023) and Maxted et al. (2024) show that adding present bias to an otherwise standard heterogeneous agent New Keynesian model (which features borrowing constraints and an illiquid asset) leads to significant amplification of both fiscal and monetary policy.

[3] This innovation is similar in spirit to Hsieh and Klenow (2009), who estimate firm-level wedges (misallocation), and Ma et al. (2024), use managers' subjective expectations to isolate misallocation due to distortions other than biased beliefs.

full-information rational expectations (FIRE) in order to calculate wedges. There is abundant evidence suggesting that consumer expectations deviate from FIRE.[4] If we were to instead assume FIRE, we risk conflating the effects of frictions and behavioral preferences with that of deviations from FIRE. This is because there generally exists some set of beliefs that can rationalize any behavior that otherwise appears consistent with constrained optimization. Our use of subjective beliefs data also proves to be important: the frictionless benchmark with FIRE beliefs can explain 57% of variation in observed consumption, but adding subjective beliefs reduces the unexplained variation by more than half.

We estimate consumption wedges using a sample of 5,272 consumers with de-identified administrative banking data linked to their surveyed expectations. The consumers in this sample are customers of EarnIn, an American financial technology company that offers their users early access to their wages prior to their regularly scheduled payday. EarnIn fielded three surveys to its customers between 2022–2024 that elicited subjective expectations over inflation, nominal income growth, and nominal interest rates for both saving and borrowing.

The EarnIn sample is not representative of the full US population, but this subpopulation is of particular economic and policy relevance. Our sample covers a significant swathe of the income distribution, with under-representation only in the tails. As expected given the revealed demand for early access to their paychecks, our sample has low liquid wealth: 64% of our respondents report liquid wealth below $500 compared to 16% of respondents to the 2022 Survey of Consumer Finances. Understanding the consumption-savings decisions of low-liquidity consumers is important. These consumers generally exhibit the highest MPCs and may have especially high marginal utility of consumption, which can make both aggregate outcomes like GDP and aggregate welfare sensitive to their decisions. Indeed, hand-to-mouth behavior of low-liquidity consumers plays a central role in shaping the macroeconomic predictions in workhorse models like Kaplan et al. (2018).

We measure consumption wedges as the gap between each individual's observed consumption and the hypothetical frictionless level they would choose absent distortions (given their beliefs and resources). We formulate a frictionless benchmark model in which a consumer chooses consumption and saving (or borrowing) via a risky asset given their realized income, wealth and expectations. In the benchmark, consumers face no distortions: there are no frictions (e.g., constraints) and consumers have standard preferences.[5] In the benchmark, frictionless consumption

[4] 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.

[5] We use the term "standard preferences" to refer to utility functions that are time consistent, time separable, homothetic, strictly increasing, strictly concave, and differentiable.

is characterized by an Euler equation and a budget constraint.

The purpose of the frictionless benchmark is not to explain observed consumption, but instead to provide the reference consumption level implied by a standard model that intentionally omits frictions and behavioral preferences. As a result, the wedge between actual and frictionless consumption quantifies the total net impact on consumption of *all* distortions (frictions and behavioral preferences).

Despite our use of a highly stylized benchmark, the same set of sufficient statistics can be used to calculate frictionless consumption in a large class models. Using a first-order approximation of the Euler equation, we characterize frictionless consumption as a function of net worth, income, and beliefs about future nominal income growth, nominal returns, and inflation.[6] Applying our formula requires specifying values for two preference parameters—a discount factor β and an inverse intertemporal elasticity of substitution (IES) γ—but it does not require specifying a utility function and is independent of many other modeling choices. Our characterization relies on the Euler equation and budget constraint being necessary conditions for optimality, but they need not be sufficient. This means that our formula for frictionless consumption is robust to a variety of model extensions, such as additional choices (e.g., labor supply) and a richer asset environment (including the case of complete markets).

We find that wedges are large and exhibit substantial heterogeneity. The median *absolute value* wedge is 40% of frictionless consumption and only 13% of the sample has consumption within 10% of their frictionless benchmark. These patterns indicate distortions play a large role in the consumption decisions of low-liquidity consumers and that small aggregate wedges mask significant cross-sectional heterogeneity.

Notably, we also find significant heterogeneity in the sign of wedges. 49% of consumers have negative wedges (under-consumption) and 51% have positive wedges (over-consumption). The existence of a large share of under-consumers and over-consumers implies that the consumption choices of our sample cannot be rationalized by a model that generates behavioral distortions in a single direction. Because financial constraints can only generate under-consumption, our results indicate that financial constraints are *not* the dominant distortion for a majority of our sample. Similarly, because present bias can only generate positive wedges, present bias alone is also insufficient to rationalize the empirical distribution of wedges.

We identify two general modeling approaches that can account for these patterns. The first is

[6] A first-order approximation of the Euler equation overstates frictionless consumption because it omits negative higher-order terms, which capture channels like the precautionary-savings motive. As such, it biases us against finding over-consumption, making our results conservative. Moreover, we show that, in a quantitative structural model, this approximation bias is small and has little cross-sectional variation.

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., Fuster et al., 2021; 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.

To evaluate the importance of our use of subjective beliefs data for these properties of the wedges, we also calculate wedges while instead assuming FIRE. We impute FIRE over income, inflation, and returns for each person using their realizations. Under FIRE, the wedge distribution exhibits higher rates of over-consumption (69% versus 51%) and has a smaller median absolute value wedge of 29% (versus 40%). These differences illustrate that subjective beliefs data are quantitatively important for measuring wedges. We also use these wedges calculated under FIRE to quantify the relative importance of beliefs versus distortions in explaining the cross-section of consumption. The FIRE wedge is the sum of two wedges: one induced by distortions and another by deviations from FIRE. A covariance decomposition reveals that frictionless consumption under FIRE explains 57% of the variation in observed consumption, while deviations in beliefs from FIRE explain 23% and distortions explain the remaining 20%. This indicates that distortions and subjective beliefs are similarly important in explaining consumption in the cross-section.

Next, to further leverage consumption wedges as a diagnostic tool, and to validate their interpretation as measuring distortions, we examine correlations between wedges and other variables of economic interest. We correlate the wedges with MPCs (both hypothetical and measured out of the 2021 stimulus checks), hypothetical responses to inflation, consumption commitments, and measures of financial distress. If wedges are the result of a particular distortion, they should also be correlated with these variables in the directions predicted by those model features. For example, more binding financial constraints generate more negative wedges and higher MPCs. Therefore, we would expect a negative correlation between wedges and MPCs if financial constraints were the dominant friction. In contrast, present bias can generate a positive correlation between wedges and high MPCs.

The correlations we find are inconsistent with financial constraints being the dominant distortion behind high MPCs and instead point to present bias and inertia. MPCs are positively correlated with wedges—the opposite of what we would expect if financial constraints were the main reason MPCs are high. Consumers with higher wedges (over-consumers) also report higher levels of financial distress (financial anxiety, unmanageable debt, difficulty borrowing) and have higher consumption commitments (housing and childcare) as a share of income. These patterns

fit present-biased behavior and adjustment frictions. On the balance-sheet margin, homeowners under-consume more than non-homeowners, consistent with homeowners being more likely to face binding financial constraints, resembling the "wealthy hand-to-mouth" of Kaplan and Violante (2014). Taken together with the substantial mass of positive wedges in the cross-section, these facts suggest that distortions beyond financial constraints—such as present bias and consumer inertia—are important drivers of high MPCs and financial distress.

We also correlate survey responses with hypothetical consumption decisions, which are independent of the wedge measurement, as a way to validate that the wedges reflect underlying consumption behavior. When asked how they would adjust saving if they expected higher inflation, over-consumers are more likely to report saving less, consistent with difficulty cutting spending when real budgets tighten. When allocating a hypothetical windfall, the marginal propensity to save (MPS) is lowest among over-consumers. These patterns reveal an internal consistency between individuals' hypothetical behavior and their over-consumer/under-consumer type, as indicated by their wedge. This helps validate the wedge as a useful sufficient statistic capturing deviations from benchmark "frictionless" consumption.

Lastly, to assess the robustness of our results, we provide evidence on their sensitivity to two important sets of assumptions. These assumptions relate to the modeling of preferences and measurement error.

Starting with preferences, we generally find quantitatively insignificant sensitivity of our results to the choice of the inverse IES (γ). The discount factor β is a more important choice, however we still find a large (20%) share of over-consumers even for a very low annual discount factor (0.80). Our next exercises are motivated by the fact that heterogeneity in preferences, specifically deviations from our calibrated values, would also manifest as a wedge. We modify our approach to allow for a plausible degree of preference heterogeneity, calibrating preferences to match the three types of Aguiar et al. (2024) such that they minimize wedges. We obtain similar results: a median absolute value wedge of 37.3% and an over-consumer share of 38.5%. We also take a more direct approach and elicit survey measures of preferences in our second and third survey waves, using questions based on the Global Economic Preferences Survey (Falk et al., 2018) and Andreoni and Sprenger (2012), respectively. We find that these measures of preferences are not statistically nor economically significant predictors of wedges. These (and additional) exercises suggest that distortions, rather than preference heterogeneity, are likely the driver of our main findings.

Across four different exercises, we find that measurement error—in either consumption, income, or beliefs data—is unlikely to significantly distort our findings. In the first exercise, we drop respondents who are likely to exhibit more measurement error—such as users with low

financial literacy, who spent less than 6 minutes on the survey, or who rounded inflation expectations to a multiple of five percentage points—to test whether the inclusion of these users skews our results. Second, we address the possibility that respondents may receive income from sources not observable in the transactions data, which would bias our measured wedges upwards, by recalculating wedges using an alternative survey-based income measure instead of transactions income. Third, we use *k*-prototypes clustering on economic and demographic characteristics to group similar users. We then take medians of the wedge inputs for each cluster to "average out" idiosyncratic measurement error and use these to calculate a wedge for each cluster. Finally, we simulate the impact of adding zero-mean-noise to the wedge inputs (up to 20% of each inputs' standard deviation). Across all four analyses, we consistently find a similar over-consumption share and a similar magnitude median absolute value wedge.

Related Literature. Our paper contributes to several literatures. First, we innovate on the macroeconomic literature quantifying frictions and distortions by measuring wedges at the individual level and using subjective beliefs data to isolate frictions and behavioral preferences versus deviations from FIRE. Our approach most directly builds on Zeldes (1989), which estimates average wedges in Euler equations separately for high- and low-wealth households. Zeldes uses a benchmark model with borrowing constraints, assumes FIRE, and uses data on food consumption from the PSID. Our data allow us to relax the assumption of FIRE and measure wedges at the individual level with high-quality transactions data. We also build on the business cycle accounting methodology of Chari et al. (2007), which popularized studying wedges 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). We take this approach in a different direction by combining survey data on subjective beliefs with administrative transactions data to calculate consumer-level wedges, which allow us to disentangle the influence of biased beliefs from frictions and behavioral preferences.

Second, we contribute to the empirical macro literature studying the determinants of consumption (Attanasio and Weber, 2010; Krueger et al., 2016; Kaplan and Violante, 2022). Early work in this area focused on testing the permanent income hypothesis (PIH) through Euler equation estimation (Hall, 1978; Flavin, 1981; Hansen and Singleton, 1982; Campbell and Mankiw, 1989). This literature found systematic departures from the predictions of frictionless PIH models, including excess sensitivity to predictable income changes and violations of the orthogonality conditions implied by FIRE. The advent of high-quality spending data gave rise to a related body of re-

search documenting large MPCs, especially among consumers with low liquidity (Johnson, Parker and Souleles, 2006; Baker, 2018; Ganong and Noel, 2019; Gross, Notowidigdo and Wang, 2020; Fagereng, Holm and Natvik, 2021). 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 high-liquidity consumers. These newer findings have motivated a growing literature proposing 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 analysis contributes to this ongoing debate by documenting the need for explanations that generate over- as well as under-consumption. The prevalence of over-consumption, in particular, indicates that it is important to incorporate distortions like present bias or inertia.

Our diagnostic evidence also corroborates the conclusions of Fuster et al. (2021), who argue that consumption adjustment frictions best explain patterns in MPCs. Methodologically, we develop a micro-level diagnostic tool that can help test competing models of distortions. Our approach does not require quasi-experimental variation, unlike MPCs. While transactions data is ideal for measuring consumption, wedges can also be measured using survey data alone.

Third, we build on work in empirical macroeconomics showing that consumer beliefs, including departures from FIRE, are instrumental to understanding consumer behavior (see recent reviews by Weber et al., 2022; D'Acunto et al., 2023). Recent papers have linked beliefs to consumption decisions using surveys (Coibion et al., 2023; D'Acunto et al., 2022), grocery purchases (Weber et al., 2023), German bank data (Hackethal et al., 2023), and credit cards (Kanz et al., 2021). Consumer beliefs systematically deviate from FIRE by, for example, overweighting grocery prices in inflation expectations (D'Acunto et al., 2021) and extrapolative income forecasts (D'Acunto et al., 2024). These findings motivate our use of subjective beliefs data. Our measurement of wedges constitutes a novel application of beliefs data. We find both beliefs and distortions are first-order determinants of consumption behavior. Recent related work by Choukhmane and de Silva (2024) quantifies the relative roles of frictions versus preferences and beliefs in portfolio choice. In this paper, we focus on separating the influence of beliefs from frictions and behavioral preferences in consumption-savings choices.

Outline. We begin by formulating a frictionless benchmark for consumption choices and develop our approach to measuring wedges in Section 2. Section 3 describes our survey and linked transactions data. Section 4 presents our analysis of consumption wedges and Section 5 test empirical predictions regarding wedges from alternative models. Section 6 concludes.

2 Theory: Measuring Consumption Wedges

We derive a formula for frictionless consumption by solving a simplified consumption-savings model. Despite the stylized nature of our model, the formula we obtain generalizes to much richer environments. In this sense, we show that data on income, wealth, and beliefs over future inflation, income growth, and returns are sufficient statistics for identifying frictionless consumption in a large class of models.

2.1 Frictionless Benchmark

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

$$V_{i,t}(Y_{i,t}, A_{i,t}, P_t, R_{i,t}) = \max_{\{A_{i,t+1}, C_{i,t}\}} u\left(\frac{C_{i,t}}{P_t}\right) + \beta \widetilde{E}_{i,t}\left[V_{i,t+1}(Y_{i,t+1}, A_{i,t+1}, P_{t+1}, R_{i,t+1})\right] \tag{1}$$

$$\text{s.t.} \quad C_{i,t} + A_{i,t+1} = Y_{i,t} + A_{i,t} R_{i,t}. \tag{2}$$

Every period, she receives income $Y_{i,t}$ and her start-of-period wealth is $A_{i,t}R_{i,t}$, where $A_{i,t}$ is her previous savings and $R_{i,t}$ is the ( ) nominal rate of return realized on her wealth. A negative value of $A_{i,t}$ 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 $c_{i,t} = \frac{C_{i,t}}{P_t}$). We assume the consumer has "standard preferences," which we take to mean time consistent, time separable, homothetic, strictly increasing, strictly concave, and continuously differentiable. The operator $\widetilde{E}_{i,t}(\cdot)$ denotes i's subjective expectation conditional on her information set at time t. We allow i's beliefs to flexibly depart from FIRE. Specifically, we do not assume that her subjective conditional expectation operator follows Bayes' rule, nor that it uses valid probability distributions. Optimal consumption $C^\star_{i,t}$ in the frictionless benchmark is characterized by the budget constraint in Equation (2) and the Euler equation:

$$u'\left(\frac{C^\star_{i,t}}{P_t}\right) = \beta \widetilde{E}_{i,t}\left[u'\left(\frac{C^\star_{i,t+1}}{P_{t+1}}\right)\frac{R_{i,t+1}}{\pi_{t+1}}\right] \tag{3}$$

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

The frictionless benchmark has three key features. First, it has no frictions (borrowing constraints, adjustment costs, etc.). Second, it assumes standard preferences and thus precludes behavioral preferences like present bias or habit formation, which can result in "as if" constrained

behavior. We intentionally omit these plausible distortions because doing so makes it possible to quantify their impact on consumption by contrasting frictionless consumption with observed consumption. The purpose of the model is to enable us to measure this impact, not to provide a realistic description of consumption behavior. Third, we depart from the prior wedge measurement literature (e.g., Chari et al., 2007; Berger et al., 2023) by letting beliefs deviate from FIRE. This allows us to separate the impact on consumption of such distortions from deviations from FIRE.

Frictionless Consumption. We solve for frictionless consumption, $C^{\star}_{i,t}$, by forward iterating the budget constraint and Euler equation, taking a first order approximation of the iterated Euler equation, and combining it with the budget constraint. We obtain the following expression (see Appendix B for the derivation):

$$C^{\star}_{i,t} \approx \frac{A_{i,t}R_{i,t} + Y_{i,t} + Y_{i,t}\sum_{j=1}^{T}\left[\widetilde{E}_{i,t}G^{Y}_{i,t,t+j}\prod_{k=1}^{j}\left(\widetilde{E_{i,t}}R_{i,t+k}\right)^{-1}\right]}{1+\sum_{j=1}^{T}\left\{\prod_{k=1}^{j}\left[\beta^{1/\gamma}\left(\frac{\widetilde{E}_{i,t}R_{i,t+k}}{\widetilde{E}_{i,t}\pi_{t+k}}\right)^{1/\gamma-1}\right]\right\}}. \tag{4}$$

Frictionless consumption $C^{\star}_{i,t}$ is a function of wealth, beliefs, and two preference parameters. Financial wealth $A_{i,t}R_{i,t}$ and income $Y_{i,t}$ constitute start-of-period wealth. The required expectations are beliefs over gross nominal income growth $G^{Y}_{i,t,t+j}$ (from period t to $t+j$), annual inflation π_{t+j}, and gross annual nominal returns to wealth $R_{i,t+j}$, which determine expected lifetime wealth. The preference parameters are the inverse intertemporal elasticity of substitution (IES) γ and the discount factor β. The resulting equation reflects the familiar logic of the textbook permanent income hypothesis, where the numerator corresponds to (approximate) expected lifetime wealth and the denominator dictates the share of that lifetime wealth to be consumed in period t.

Consumption Wedges. Distortions (i.e., frictions and behavioral preferences) cause actual consumption to deviate from frictionless consumption, creating a wedge $\eta_{i,t}$:

$$\eta_{i,t} = C_{i,t} - C^{\star}_{i,t}. \tag{5}$$

Intuitively, the consumption wedge $\eta_{i,t}$ measures how far consumption is from the frictionless Euler-equation-implied consumption, quantifying the total impact of all distortions on consumption. Negative wedges correspond to "under-consumption" (i.e., consuming less than the frictionless benchmark) and positive wedges to "over-consumption."

Examples: Implications of Common Distortions for Wedges. We illustrate why consumption wedges are an effective diagnostic tool by considering the implications for wedges of three commonly used distortions. First, consider financial constraints, which can only *reduce* consumption. Therefore, a testable implication of financial constraints is that they generate negative consumption wedges. A second example is present bias, which can only *increase* consumption relative to the the Euler equation by discounting future consumption and increasing the share of permanent income consumed in the present period. Therefore, a testable implication of present bias is that they generate positive consumption wedges. Finally, inertia can arise from consumption commitments (adjustment costs), habit formation (reference-dependent preferences), or bounded rationality (cognition is costly and limits or delays consumption adjustments). Under inertia, consumption is slow to adjust both upward and downward following a shock, and therefore can cause consumption to be above or below the frictionless depending on the direction of the shock. Therefore, inertial distortions can generate both positive or negative wedges. See Appendix B.2 for additional discussion and examples of how these distortions change the Euler equation.

Why the Benchmark Omits Distortions. Our benchmark intentionally excludes frictions and behavioral preferences. By doing so, our benchmark becomes a "special case" in a large class of models featuring distortions—it is the case where the distortions are "turned off." For example, in the class of models with borrowing constraints, our benchmark is the case of an infinitely low borrowing constraint. In a model with beta-delta discounting, the benchmark is the case of no present bias. By omitting such distortions, the wedge between actual and frictionless consumption identifies the sum total effect of all such distortions on consumption.

Robustness to Model Extensions. To simplify exposition, the frictionless benchmark abstracts away from additional choice variables, such as labor supply, multiple assets, and types of consumption. This is without loss of generality. Each additional choice entails an additional optimality condition to characterize the overall optimum (i.e., all optimal choices). But we only focus on characterizing optimal consumption in our benchmark. Because frictionless consumption is derived only using the budget constraint and Euler equation, our formula applies in the *class* of models where an Euler equation and budget constraint are necessary conditions for overall optimality. They need not be sufficient conditions for overall optimality.

We detail several extensions in Appendix B.3 and briefly summarize their implications here. The presence of additional choices like labor supply and multiple consumption types (such as durable and non-durable consumption) do not alter our expression but can affect the ideal way

to *empirically* measure wedge inputs. For labor supply, the presence of income taxes means that income ought to be measured on an after tax basis for the wedge to retain its interpretation. When multiple consumption types are present, the Euler equation is a function of "notional" consumption (i.e., an aggregate of the different consumption types) and the wedge pertains to this aggregate measure. For durable goods specifically, the usual empirical challenge of measuring durable consumption versus expenditure arises. With an empirical measure of total durable and non-durable consumption, the wedge formula is still interpreted in terms of consumption, not expenditure. We can also modify the benchmark model to feature a richer portfolio choice problem where the individual chooses a mix of assets and liabilities, including housing or even complete asset markets. In such a model, the interpretation of the wedge is unchanged so long as it is measured using portfolio-weighted expected return across all assets.

3 Data and Survey Design

3.1 Data Sources

EarnIn Administrative Data. We receive anonymized bank transactions data from EarnIn, a US-based financial technology company. EarnIn provides access to earnings before payday through their "earned wage access" (EWA) product and had more than 2.5 million active users as of the start of our study period.[7] Users are required to link their bank accounts and EarnIn maintains an administrative database that includes users' bank account transactions and balances, earnings, and EWA cashout activity.

We use the transactions data to construct two variables central to our analysis: non-durable spending and income. We define outflow transactions as non-durable spending based on a mapping from more than 500 transaction categories associated with the transactions to spending categories building on Ganong and Noel (2019). 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 rate of consumption. We measure income as the sum of post-tax earnings and unemployment insurance (UI) benefits. We classify inflow transactions as either earnings (observed post-tax), UI benefits, or "other" using the observed earnings data together with the transaction category, memo line, and periodicity of the transaction. Our baseline analysis covers the 12-month period before each survey wave (e.g., annual non-durable spending and income between October 2021 to September 2022 for wave 1).[8]

[7] https://www.businesswire.com/news/home/20211102006071/en/Earnin-Announces-They-Have-Provided-Access-to-%2410-Billion-in-Earnings-for-Members

[8] We use transaction categories from Plaid, a financial data aggregator that links user-authorized bank and card accounts

For more details on the EarnIn data, sample construction, and variable definitions see Appendix C.

Survey Data. We link the EarnIn database with survey responses from nearly 15,000 users covering demographics, personal finances, and subjective economic expectations. Our dataset allows us to link expectations with earnings, spending, and savings data that can paint a near-comprehensive picture of a consumer's economic activity. Data linking subjective economic expectations and bank account transactions are rare, and we are the first (to our knowledge) to gather this data for US consumers.[9]

EarnIn fielded three survey waves between 2022 and 2024. Wave 2 was a follow-up survey sent to respondents from wave 1. In each wave, qualifying users were invited via EarnIn's standard email marketing channels to complete a short survey about their current economic well-being and future outlook. Survey windows spanned from September 29 to October 2, 2022 (wave 1); July 12 to July 19, 2024 (wave 2); and November 22 to December 4, 2024 (wave 3). Surveys took five to ten minutes (wave 1) or ten to fifteen minutes (waves 2 and 3) to complete. As an incentive, respondents received a $5 (waves 1 and 3) or $10 (wave 2) Amazon gift card upon completion.

To construct the sampling frame for each wave, we imposed data quality requirements on the transactions data to ensure that users' linked accounts plausibly capture the bulk of their economic activity. For waves 1 and 3, we restricted to users for whom we observe earnings, regular spending, and balances in the 12 months leading up to the survey. For the wave 2 follow-up survey, we invited only wave 1 respondents who met additional data quality restrictions and had not yet reached the email marketing limit. Across all three surveys, EarnIn sent approximately 475,000 survey invitations and received approximately 15,866 responses, yielding an aggregate response rate of 3.4%. The median survey completion time among respondents was 7.5 minutes. Survey respondents are similar to all users eligible to take the survey on account balances, inflows, and outflows, but women were more likely to respond (see Appendix Table C.2). See Appendix C for additional details on sample construction and response rates by survey wave.

Each survey wave elicited expectations for inflation, income growth, and interest rates. For inflation, respondents were asked to forecast inflation over two future horizons: 0-12 months (short-run) and 24-36 months (medium-run). The phrasing of these questions was based on questions from the University of Michigan Survey of Consumers (MSC) and the Federal Reserve Bank of New York Survey of Consumer Expectations (SCE). We asked respondents to forecast their

to applications via APIs and provides standardized transaction data and categories.

[9] D'Acunto et al. (2021) link economic expectations to grocery spending and Kanz et al. (2021) link economic expectations with credit card spending. Hackethal et al. (2023) and D'Acunto et al. (2024) leverage both expectations and transactions data for users of German and Chinese banks, respectively.

personal income growth over the next 12 months, the percentage yield they would expect on additional savings over the next 12 months, and the percentage rate they would expect to pay on new borrowing over the same period. In addition to these expectations questions, we asked respondents how a hypothetical increase in inflation would affect their saving and why. We also elicited measures of financial distress, use of alternative financial products, financial literacy (Lusardi and Mitchell, 2014), demographic information, and total household income, savings, and debt (across categories).[10]

In waves 2 and 3, we added several survey questions to expand our wedge measurement and analysis. Beginning in wave 2, we elicited respondents' expected spending growth over the next 12 months. We solicited proxies for risk and time preferences using questions based on the Global Economic Preferences Survey (Falk et al., 2018) in wave 2 and Andreoni and Sprenger (2012) in wave 3. We added measures of committed consumption by asking respondents whether they own or rent their home, monthly rent/mortgage payments, and childcare spending. In wave 3, we elicited respondents' hypothetical marginal propensity to spend, repay debt, and save out of a $1,000 one-time income shock using questions based on (Fuster et al., 2021; Colarieti et al., 2024).

Sample Restrictions for Analysis. We impose a series of sample restrictions for our analysis, detailed in Appendix C. We exclude responses flagged for inattention or low effort by requiring a minimum survey duration of 3.5 minutes, internally consistent responses, and economic expectations within reasonable bounds.[11] We also drop users with insufficient transactions data coverage in the 12 months surrounding the survey, and we trim outliers for key financial variables (nondurables spending, income, average propensity to consume, wealth-to-income, and expected levered return). After all restrictions, our analysis sample includes 5,272 unique users.

Imputing Wealth Our survey solicited binned measures of liquid assets and total debt, but not illiquid assets. To obtain a more complete and precise measure of net worth AR/Y, we leverage our detailed economic, financial, and demographic data and train a machine learning model (XGBoost) to predict illiquid assets, as well as a numerical value for liquid assets and debt constrained to fall within the user's self-reported bin. We estimate the model using data from the Survey of Consumer Finances, following the procedure detailed in Appendix D.1.

Table 1. Summary Statistics

	Mean (1)	SD (2)	P25 (3)	P50 (4)	P75 (5)	N (6)
Panel A: Demographics						
Female (%)	69	.	.	.	.	5,272
White (%)	55	.	.	.	.	5,181
Age	37	9	31	36	43	5,254
Has children (%)	51	.	.	.	.	5,272
Spouse or partner (%)	54	.	.	.	.	5,272
College (%)	43	.	.	.	.	5,272
Panel B: Household finances						
Liquid assets ($)	1,466	3,489	236	250	750	5,272
Total assets ($)	42,949	62,398	8,039	19,679	42,125	5,272
Total debt ($)	43,344	55,939	7,503	16,873	39,960	5,272
Has mortgage (%)	18	.	.	.	.	5,272
Liquid net worth ($)	-41,878	55,565	-39,287	-16,159	-6,968	5,272
Total net worth ($)	-395	40,144	-13,981	-224	12,947	5,272
Nondurables spending ($)	29,944	13,570	19,954	26,915	37,292	5,272
Income ($)	43,696	19,337	30,058	39,639	53,173	5,272
Liquid net worth to income (%)	-94	124	-122	-47	-19	5,272
Total net worth to income (%)	-2	87	-39	-1	33	5,272
Nondurables C/Y (%)	72	25	54	68	85	5,272
Panel C: Economic expectations						
Perceived inflation (%)	9	9	5	7	10	3,246
E(Inflation, 1Y) (%)	7	7	3	5	10	5,272
E(Inflation, 3Y) (%)	4	8	-2	4	8	5,272
E(Income growth) (%)	5	9	2	3	5	5,272
E(Rate on savings) (%)	3	3	1	2	4	5,272
E(Rate on borrowing) (%)	15	10	7	15	24	5,272
E(Levered return) (%)	16	10	7	15	25	5,272
E(Spending growth) (%)	6	66	-2	5	10	2,355
Panel D: Financial constraints						
Liquid assets below $500 (%)	64	.	.	.	.	5,272
High financial anxiety (%)	52	.	.	.	.	5,271
Debt unmanageable (%)	72	.	.	.	.	5,270
Difficulty borrowing (%)	48	.	.	.	.	5,271
High financial literacy (%)	45	.	.	.	.	5,272
EarnIn EWA usage ($)	5,410	3,979	2,612	4,593	7,254	5,272
EarnIn EWA usage / income (%)	14	12	6	11	19	5,272
Panel E: Consumption commitments						
Has mortgage (%)	18	.	.	.	.	5,272
Consumption commitments ($)	1,695	3,497	1,000	1,500	2,058	2,280
Commitments / income (%)	55	88	30	45	67	2,280
Panel F: Response to an income shock						
Observed nondurables MPC (%)	28	34	3	25	53	2,474
Hypothetical MPC (%)	21	24	0	20	30	2,008
Hypothetical MPRD (%)	68	36	50	75	100	2,016
Hypothetical MPS (%)	11	36	0	0	0	2,008

Notes: The table presents summary statistics for our analysis sample after trimming wedges at the 1st and 95th percentiles, which consists of 5,272 EarnIn users. Columns (1) through (5) show the distribution of each variable, and column (6) shows the number of survey responses with non-missing data. See Appendix C.4 for variable definitions.

3.2 Summary Statistics

Table 1 presents summary statistics for our analysis sample. Compared to the US population, our sample skews younger (median age of 36), female (69% of respondents), and non-white (45%). Respondents exhibit a high degree of financial distress: 64% report having less than $500 in savings, 72% report having an unmanageable amount of debt, and 52% report having high anxiety related to their finances. Forty-three percent of respondents have a college degree, and 45% correctly answered both financial literacy questions.

Figure E.1 compares the distributions of income, liquid assets, and inflation expectations to the US population. The distribution of labor income is representative of middle income workers in the Current Population Survey, with modest overepresentation between the 30th and 80th percentiles and fewer representatives in the tails (Panel A); however, our sample is predominantly low liquidity relative respondents to the Survey of Consumer Finances (Panel B). The median respondent has $39,639 annual post-tax income, $16,873 in total debt, and $19,679 in total assets – but only $250 in liquid assets. Only 18% of respondents have a mortgage.

The median one-year-ahead inflation expectation in our sample is 5%. Respondents expect inflation to come down slightly, with a median three-year inflation expectation of 4%. The median estimate of realized inflation over the prior year is 7% – higher than standard inflation measures, but potentially reflective of the consumption baskets faced by our middle-income, low-liquidity sample. The median respondent forecasts nominal income growth of 3% over the next year, on average, implying an expected real income decline of -2%. Reported interest rate expectations on marginal savings and borrowing are broadly reasonable, with medians of 2% and 15%, respectively.

Inflation expectations in our sample are broadly consistent with those of nationally representative samples, particularly the SCE. Panel C of Figure E.3 presents the distribution of inflation expectations in our survey relative to MSC and SCE respondents from the same time period. Our survey sample is also slightly more likely to report round numbers divisible by five. Appendix E.3 presents our respondents' self-reported expectations for inflation, nominal income growth, and interest rates against their realized values. In general, we find that the median forecast error is small.[12]

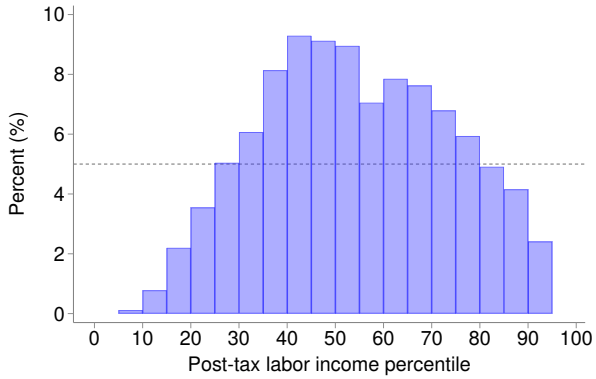
[10] The survey instruments are available online for each wave: wave 1 (TK), wave 2 (TK), and wave 3 (TK).

[11] We retain respondents with inflation and income growth expectations between the 3rd and 97th percentiles and interest rate expectations between the 1st and 97th percentiles.

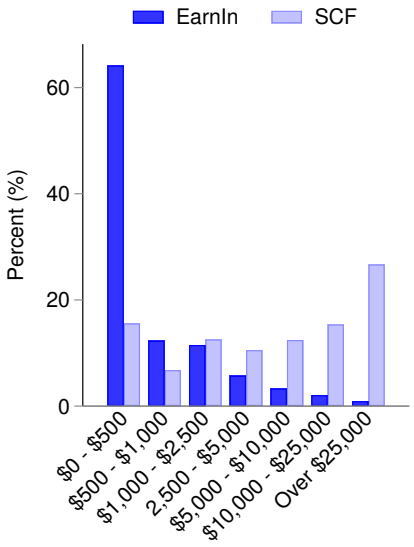
[12] The analysis of forecast errors provides qualitative evidence that the survey instrument captures reasonable expectations, but cannot precisely identify biases in individual respondents' expectations. We cannot observe individual-level realizations of inflation and interest rates because we do not observe each respondent's distinct consumption basket and asset mix. We therefore cannot isolate forecast bias from individual deviations, aggregate inflation rates, or interest rates.

Figure 1. Sample Benchmarks

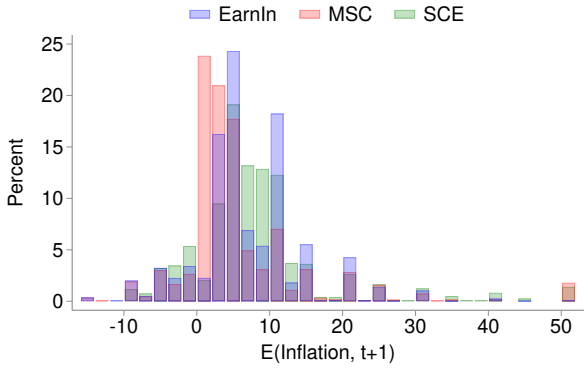
Panel A. Distribution of Income Percentiles



Panel B. Distribution of Liquid Wealth



Panel C. Comparison with National Surveys



Notes: Panel A presents the distribution of post-tax labor income percentiles among our EarnIn analysis sample, estimated using CPS ASEC data and the CPS ASEC Tax Model. PS ASEC sample restricts to adults age 18 or older who received wage or salary earnings during the reference year. Black dashed line represents the uniform distribution. Panel B 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 C. Panel C shows the relative to the distributions in the Michigan Survey of Consumers (MSC) and NY Federal Reserve Survey of Consumer Expectations (SCE) (right). To match the timing of the EarnIn surveys, we use MSC and SCE data from September 2022, July 2024, and November 2024 (the SCE data is unavailable in November 2024) and reweight the data based on the share of the EarnIn sample in each survey wave.

3.3 Quantifying the Frictionless Benchmark

We next describe how we take our formula for frictionless consumption, Equation (4), to the linked survey-transactions data.

Calibrating Preference Parameters. Our baseline calibration uses standard values for the annual discount factor ($\beta = 0.92$) and the inverse IES ($\gamma = 2$).[13] In Section 4.2, we perform sensitivity analyses to demonstrate the robustness of our main results to alternative choices for these preference parameters as well as robustness to preference heterogeneity.

Expected Returns. We measure the gross interest rate $\widetilde{E}_{i,t} \ln R_{i,t+1}$ using expected returns/costs of two assets: savings and debt. We calculate a levered return from their beliefs about $\widetilde{E}_{i,t} R^S_{i,t}$ the expected return to savings and $\widetilde{E}_{i,t} R^D_{i,t}$ is the cost of debt as follows:

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

where $S_{i,t}$ is their liquid wealth and $D_{i,t}$ their total liabilities. We assume that the return on this portfolio, which excludes illiquid assets, is the same as the return they expected on their portfolio of illiquid assets. Under this assumption, the above expression equals the gross expected portfolio return.

Term Structure of Beliefs. Frictionless consumption depends not only on one-year-ahead beliefs but also on the entire term structure of beliefs. Our survey solicits one-year-ahead beliefs over nominal income growth, interest rates, and inflation, as well as a three-year-ahead belief for inflation. 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).[14] For interest rates, we assume the expected gross rate is constant (i.e., a flat term structure). Lastly, we impute expected income growth using a procedure that allows for two important properties: (1) large but temporary variation in expected growth rates of income in the short-run variation due to anticipated shocks (e.g., job changes) and (2) lifecycle dynamics. Using the SCE, we find significant reversion in income expectations over a year and use this to discipline the relationship between one-year- versus two-year-ahead expected income. We use the MSC to estimate a lifecycle profile of income growth expectations, which we apply to our sample for three-year-ahead (and later) expectations. We detail all of our imputation procedures in Appendix D. As a test of the quantitative importance of these assumptions, in Appendix Figure F.4 we report the distribution of wedges under an alter-

[13] Our value of β is in the typical range of values used in models featuring unsecured borrowing (e.g., Bornstein and Indarte, 2023) and lies within the range used in Auclert et al. (2024).

[14] The median ratio of one- to three-year-ahead inflation expectations in the EarnIn sample is 1.0 for each survey wave, comparable to the ratio for swaps-implied beliefs (1.03 in October 2022, 1.04 in July 2024, and 1.16 in November 2024).

native assumption of constant beliefs. Allowing one-year beliefs to persist indefinitely results in more outliers and higher absolute value of wedges, but the shape of the distribution is otherwise similar to those resulting from our baseline extrapolated expectations.

Income and Consumption Measurement Assumptions. To calculate the wedge, we subtract frictionless consumption from actual consumption, and divide this quantity by frictionless consumption. To measure actual consumption, we begin by calculating non-durable consumption 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).[15] To obtain total or "notional" consumption (i.e., the argument of the utility function), we divide each respondent's non-durable consumption by the expenditure share of non-durable goods (79.37%). Under the assumption that notional consumption is a Cobb-Douglas aggregate of durable and non-durable good consumption flows, this calculation yields the notional consumption (for a proof, see Appendix B.3.1). We obtain the non-durable expenditure share from Beraja and Zorzi (2024), which calculates it using Consumer Expenditure Survey data.

4 Empirical Analysis of Consumption Wedges

This section begins by analyzing the empirical distribution of consumption wedges for our survey population. We then show that our findings of large wedges and shares of over-consumers are robust to model calibration choices (including preference heterogeneity), measurement error, and approximation bias. To evaluate the importance of our use of subjective beliefs data, we construct an alternative "FIRE" wedge that instead uses imputed FIRE beliefs for each consumer. Via a covariance decomposition, we use the FIRE wedges to show that both distortions and subjective beliefs account for similarly large fractions of the cross-sectional variation of consumption. Lastly, we study a final distinct type of wedge—static wedges—which require fewer assumptions and variables to measure but only capture a subset of the impact of distortions on consumption.

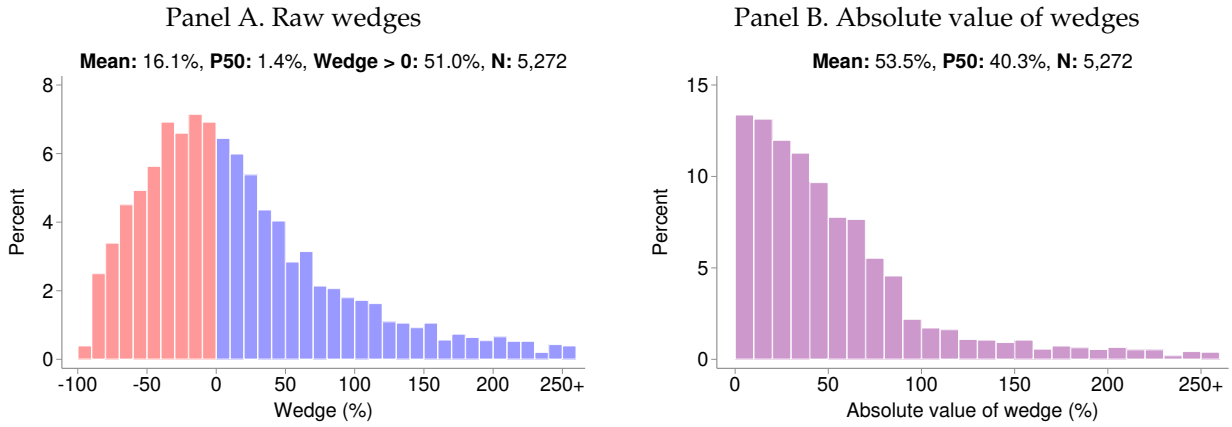
4.1 Results: Empirical Consumption Wedges

Calculating consumption wedges for each respondent, we find large wedges and significant heterogeneity. Figure 2 displays histograms of measured wedges. Throughout, we divide wedges

[15] 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 as separate measures of consumption, or when aggregating consumption over fewer pre-survey months (e.g., the month of the survey only). See Appendix Figure F.3.

by frictionless consumption so that they can be interpreted as a percent deviation. Our first main finding is that many consumers have significantly distorted consumption. The average wedge is approximately 16% (Panel A), indicating that the average respondent has consumption distorted upwards 16% relative to their counterfactual frictionless consumption. Taking the absolute value of wedges, we find that the median distortion is 40% (Panel B). Further examining the absolute value wedges, we see that the modal respondent has a distortion less than 30%. It is notable that not all consumers have large wedges, especially given the prevalence of low-liquidity consumers in our sample. Approximately 13% of the sample has consumption within 10% of frictionless consumption. These consumers 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 2. Distribution of Dynamic Consumption Wedges



Notes: The figures show the distribution of dynamic consumption wedges (left) and their absolute values (right). Wedges are reported here as the percent deviation of observed consumption from frictionless consumption, and they are trimmed at the 1st and 95th percentiles.

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 and/or behavioral preferences are an important determinant of the consumption choices of low-liquidity consumers. Therefore, these distortions are important to include in models of consumer behavior featuring populations with low liquidity.

Our findings also highlight the value of studying micro-level wedges, as opposed to macro-level (aggregate or average) wedges. The average wedge is less than half the magnitude of the median absolute value wedge (16% versus 40%). The median (non-absolute value) wedge is even smaller: 1.4%. Without observing the distribution, one could significantly underestimate 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 may be especially helpful for disciplining the choice and quantitative modeling of frictions and behavioral preferences.

Our second main finding is that many consumers have positive wedges (over-consume). Specifically, 51% of individuals in our sample 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 consumers. As discussed in Section 2, financial constraints only generate negative wedges. While financial constraints could explain the 49% of consumers with negative wedges, such constraints cannot be the dominant friction for the other 51%. 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 consumers in our sample. Neither financial constraints nor present bias alone can rationalize the mix of positive and negative wedges. A different model is necessary; we identify two promising alternatives. A first solution is including both present bias and financial constraints. The second alternative is the set of distortions that create inertia in consumption. These include consumption adjustment costs, habit formation, and forms of bounded rationality that results in sticky behavior (e.g., Ilut and Valchev, 2023).

Our sample has a high demand for liquidity and one might therefore expect the wedges to be especially large, and especially negative if financial constraints are the dominant friction distorting their consumption. In assessing the magnitude and distribution of wedges, it is important to keep these features of our sample in mind. This makes the significant share of over-consumers we estimate particularly surprising. EarnIn EWA usage in terms of dollars is slightly negative; as a share of income there is a small but positive correlation of usage with the wedges (see Appendix Figure A.4). These small economic magnitudes suggest that there is unlikely significant selection into our sample related to having positive or negative wedges.

A valuable endeavor for future research would be to compare the ability of the two alternative classes of models described here to match empirical patterns in consumption wedges using quantitative structural models. For example, would a model with empirically realistic negative skewness of income shocks (Guvenen et al., 2021) struggle to generate a mix of mostly negative wedges?

We next assess the robustness of the results to calibration choices and measurement error. Fol-

lowing 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.

4.2 Robustness

We next examine the robustness of our two main findings: (1) median absolute value wedge of 40% and (2) 51% share of over-consumers. We focus on three classes of concerns: sensitivity to calibration choices, measurement error, and approximation bias.

4.2.1 Sensitivity Analysis

We vary the value of each input parameter (β, γ, and the nondurable share of $\frac{C}{Y}$) and recalculate the distribution of wedges to assess the sensitivity of our results to these calibration choices. We summarize the resulting share of over-consumers and the median absolute value wedge in Table 2 and the full relationships between the inputs and results in Figure **??**.

Table 2. Sensitivity Analysis: Impact of Alternative Calibration Choices

	Parameter Range			Overconsumer (%)		P50 Abs(Wedge) (%)	
	Calibration (1)	Min (2)	Max (3)	Min (4)	Max (5)	Min (6)	Max (7)
β	0.92	0.80	0.98	18.5	73.7	40.3	53.6
γ	2.00	1.00	5.00	49.5	51.7	37.1	60.8
Nondurable share of $\frac{C}{Y}$	0.7937	0.72	0.90	42.4	57.4	39.4	43.3

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 calibrated value and the "Min" and "Max" columns show the range of values that we test. Under the "Overconsumer (%)" and "P50 Abs(Wedge) (%)" headings, 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.

Discount Factor (β). We begin by varying our choice of annual discount factor (β) from 0.80 to 0.98. Panel A of Figure F.1 displays the results. Naturally, the discount factor has a significant effect on the frictionless benchmark: if a consumer heavily discounts future consumption, their frictionless consumption today is significantly higher, leading to lower wedges. Relative to our

benchmark calibration of 0.92, a β of 0.80 decreases the share of over-consumers from 51% to 18.5% while a β of 0.98 increase the share of over-consumers to 73.7%. The median absolute value of the wedge is less sensitive because in either case it combines the wedges on one side of the distribution becoming larger while attenuating those on other side. These results provide some reassurance that the frictionless benchmark and resulting wedges behave as we would expect for a given level of actual consumption, and there remains a substantial share of over-consumers even when everyone is assumed to have extremely low discount rates.

Inverse Intertemporal Elasticity of Substitution (γ). We next examine alternative choices for the coefficient of relative risk aversion/inverse IES (γ). We vary γ from one to five and display the results in Panel B of Figure F.1. The share of over-consumers remains close to 51%, while the median absolute value wedge slightly decreases with γ.

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.[16] In contrast to β and γ, the non-durable expenditure share does not affect the frictionless benchmark but it does affect the measurement of total actual consumption. We divide observed non-durable spending by the non-durable expenditure share to infer total consumption. Therefore, if we instead assume a lower non-durable expenditure share, this translates to larger total consumption and higher shares of over-consumers. This monotonic relationship is observable in Figure F.1 and Table 2, which shows that the share of over-consumers ranges from 42.4% to 57.4% when we shift the non-durable expenditure share from 0.72 to 0.90. As above, the median absolute value wedge is broadly insensitive to the calibration choice.

Preference Heterogeneity. If consumers have heterogeneous preferences, the difference between an individual's true preference parameter and the homogeneous calibrated value used in our baseline analysis would contribute to their consumption wedge. Hence preference heterogeneity could also be a source of consumption wedges in our analysis. We assess the sensitivity of our results to a model with three heterogeneous types, using the three preference types estimated in Aguiar et al. (2024).[17] For each person in our sample, we assign them the discount factor and risk aversion

[16] 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 88.8% when we classify durable expenditure in the EarnIn transactions data.

[17] In practice, any feasible consumption choice can be rationalized by some set of preference parameters. Therefore, a model with unconstrained individual-level heterogeneity cannot generate consumption wedges: any observed

type that minimizes the absolute value of their consumption wedge. This three-type model of preference heterogeneity reduces the wedges modestly: the median absolute value of the wedges declines from 40% to 37%, and the share of over-consumers falls from 51% to 39% (Figure F.2). This analysis implies that our results remain similar in the presence of plausible levels of preference heterogeneity.

4.2.2 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 F.1 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.[18] Another analysis omits users who answered one or more financial literacy questions incorrectly.[19] Such users may have a more difficult time communicating their economic expectations. Similarly, we also drop users whose reported inflation expectations are multiples of five, as such responses may be rounded. We also consider several subgroups related to consumption measurement error. One example includes users whose peer-to-peer transfers exceed 25% of nondurable spending in at least one month. Such users may have significant spending that we do not capture. Figure F.5 reports results from excluding these various subgroups. The share of over-consumers remains similar, generally ranging from 40-55%. The median absolute value consumption wedge varies slightly, ranging between 37-41%.

Grouping Users. Suppose that measurement error in beliefs, consumption, income, and wealth is mean zero across sufficiently "similar" users. One could then group these similar users, calculate representative measures of the consumption wedge inputs, and calculate a representative consumption wedge for each group. By averaging (or taking the median) across wedge inputs, such measurement error could be reduced.

To implement this, we group respondents using k-prototype clustering, which is a combination of k-means and k-modes clustering (for more details on the algorithm, see Appendix F.1). We group respondents based on their similarity in terms of reported age, annual income, savings, and indicators for gender, race, relationship status, presence of children, college education, and political affiliation.[20] We construct 500 clusters, resulting in 12 observations per cluster on aver-

consumption data can be inverted to generate the set of preferences that rationalize those individual choices.

[18] 175 users in our sample (3%) have at least one month of UI income in the pre-survey period.

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

[20] We z-score continuous variables so that they exert equal influence in cluster assignment.

age. Figure F.6 displays updated histograms. The share of over-consumers increases to 58%. The median absolute value wedge falls to 17%. These reductions likely not only reflect decreases in measurement error but also attenuation to mixing over-consumers and under-consumers These smaller wedges may be the result of attenuation due to being more likely to aggregate wedges across over- and under-consumers as cluster size grows.

4.2.3 Approximation Bias

Our formula for frictionless consumption is obtained by taking a first-order approximation of the Euler equation. This introduces approximation error in our measurement of frictionless consumption. Because we assume utility is concave, the sign of the bias is positive—our formula overstates frictionless consumption. This implies that wedges, which subtract frictionless consumption from observed consumption are biased downwards. This means that our conclusions about the prevalence of over-consumption (51%) are at worst conservative with respect to approximation bias.

To assess how quantitatively important this bias may be, we conduct an exercise with a quantitative Bewley (1980) style model. We use a simple theoretical model featuring a single asset, a borrowing constraint, idiosyncratic income shocks, and infinitely-lived consumers that are heterogeneous in terms of income and wealth. Each consumer solves:

$$
\begin{aligned}
V(y_{i,t}, a_{i,t}) = \max_{\{a_{i,t+1}, c_{i,t}\}} \; & u(c_{i,t}) + \beta E\left[V(y_{i,t+1}, a_{i,t+1}) | y_{i,t}, a_{i,t+1}\right] \\
\text{s.t.} \quad & c_{i,t} + a_{i,t+1} = y_{i,t} + a_{i,t} R \\
& A_{i,t+1} \geq \bar{A} \\
& \ln(y_{i,t}) = \rho_y \ln(y_{i,t-1}) + \varepsilon_{i,t}, \quad \varepsilon_{i,t} \sim N(0, \sigma_y^2).
\end{aligned}
$$

We detail our calibration of this model in Appendix Table G.1. We calculate approximation bias as the difference between the wedge implied by our approximate formula and the exact wedge.[21]

Appendix Figure G.1 displays the model's (ergodic) distribution of approximation bias. The median bias is -9.6 pp of (approximate) frictionless consumption. The negative sign of the bias implies that our findings at worst understate the degree of over-consumption. Importantly, this is bias also exhibits very little cross-sectional variation. Its inter-quartile range is -10.1 to -9.1pp. Together, these properties of approximation bias imply, because our distribution of wedges is fairly symmetric, that correcting this bias would unlikely significantly change the median absolute value

[21] To calculate exactly frictionless consumption, we solve a frictionless version of the quantitative model. That is, we solve a variant without a borrowing constraint. Because the ergodic distribution is generally different between the distorted and frictionless economies, it is important to calculate the distribution of approximation bias in the distorted economy using *its* ergodic distribution.

wedge that we find. If we subtract the median bias to our wedge calculations, we instead find a modestly larger 57.6% share over-consumers and a nearly unchanged median absolute value wedge of 39.2%. The compact distribution of approximation bias also suggests that correlation between wedges and other observables, which we examine in Section 5 are unlikely an artifact of approximation bias being correlated with these observables.

4.3 FIRE Wedges: The Importance of Subjective Beliefs Data

Measuring wedges with subjective beliefs data enables us to avoid assuming FIRE. We can evaluate the importance of relaxing this assumption for our findings if we instead follow prior work and use realizations to impute beliefs under FIRE. We measure realizations for nominal income growth using the transactions data for wave 1, CPI inflation, and the median prevailing interest rates on savings (FDIC) and credit cards (Federal Reserve Board's *Consumer Credit*, G.19 statistical release).[22]

When wedges are calculated under FIRE, they contain the impact of both deviations from FIRE and distortions. We can write observed consumption as the sum of three objects:

$$C_{i,t} = C_{i,t}^{\star,\text{FIRE}} + \eta_{i,t}^{\text{beliefs}} + \eta_{i,t}^{\text{distortions}} = C_{i,t}^{\star} + \eta_{i,t}^{\text{distortions}}. \tag{7}$$

$C_{i,t}^{\star,\text{FIRE}}$ is frictionless consumption calculated under FIRE. The wedge $\eta_{i,t}^{\text{beliefs}}$ is a "beliefs" wedge, summarizing the impact of deviations from FIRE on consumption. For clarity, we here add the "distortions" superscript to $\eta_{i,t}^{\text{distortions}}$, which corresponds to the wedge capturing the influence of distortions on consumption and has been the focus of our analysis (i.e., the same wedge in Figure 2). The FIRE wedge we next examine is the *sum* of the beliefs and distortions wedges.

We find that our use of subjective beliefs is quantitatively important. Appendix Figure H.1 displays histograms of FIRE wedges. These wedges imply a larger share of over-consumers (69% versus 51%) and smaller typical absolute value distortions (medians of 29% versus 40%). This also indicates that deviations from FIRE tend to generally push people towards over-consumption.

Another advantage of measuring FIRE wedges is that they enable us to quantify the relative importance of distortions versus beliefs. To do so, we apply a covariance decomposition to Equation (7).[23] We find that a majority (56.7%) of the variation in observed consumption is explained

[22] We opt to measure these beliefs by setting them equal to their realizations, rather than averaging across multiple similar individuals as in (Berger et al., 2023). This has the disadvantage of meaning that our "FIRE wedges" will reflect the influence of both deviations from FIRE and prediction error on consumption, this is essentially measurement error. However, averaging will also generally feature measurement error. An advantage of using realizations is that prediction error relative to FIRE beliefs is classical measurement error, which results in either no (or well-understood) biases in our analysis using FIRE wedges.

[23] That is, we calculate the share of the variation in observed consumption explained by each component: $1 =$

by frictionless FIRE consumption ($C_{i,t}^{\star,\text{FIRE}}$). The remaining share is explained almost equally by deviations from FIRE (22.8%, $\eta_{i,t}^{\text{beliefs}}$) and distortions (20.5%, $\eta_{i,t}^{\text{distortions}}$). Our decomposition indicates that both distortions and beliefs are similarly important for explaining the cross-section of consumption.

4.4 Static versus Dynamic Consumption Wedges

In addition to the wedges studied above, using data on expected *consumption* growth, we can measure another different but related wedge. We refer to this alternative wedge as a "static" wedge. It is static in the sense that it only reflects the impact on consumption of distortions experienced at the time of consumption. In contrast, the "dynamic" wedge ($\eta_{i,t}$) studied above reflects not only current distortions but also the effect of expected future distortions. As we show below, we can glean further insights by contrasting static and dynamic wedges. Additionally, an appealing feature of static wedges is that they require fewer assumptions to calculate.

The static wedge is calculated using only the approximate one-period-ahead Euler equation. The key difference between the static and dynamic wedges is that the static wedge is measured using expected *actual* consumption ($\widetilde{E}_{i,t}C_{i,t+1}$), instead of expected *frictionless* consumption ($\widetilde{E}_{i,t}C_{i,t+1}^{\star}$). Recall that, for the dynamic wedge, expected frictionless consumption is determined via the budget constraint (and expectations of its components). We define static frictionless consumption to be

$$C_{i,t}^{\star,\text{static}} \approx \frac{\widetilde{E}_{i,t}C_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}}\left(\beta\frac{\widetilde{E}_{i,t}R_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}}\right)^{-1/\gamma}. \tag{8}$$

The static consumption wedge is

$$\nu_{i,t} = C_{i,t} - C_{i,t}^{\star,\text{static}} = C_{i,t}\left[1 - \frac{\widetilde{E}_{i,t}G_{i,t+1}^{C}}{\widetilde{E}_{i,t}\pi_{t+1}}\left(\beta\frac{\widetilde{E}_{i,t}R_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}}\right)^{-1/\gamma}\right] \tag{9}$$

where $G_{i,t+1}^{C} = \frac{\widetilde{E}_{i,t}C_{i,t+1}}{C_{i,t}}$ is expected actual future consumption growth. In contrast to the dynamic wedge, only data on expected one-period-ahead consumption growth, inflation, returns, and the same two preference parameters are needed to calculate static wedges in percentage terms (as one can divide by actual consumption and avoid having to measure consumption). In this sense, static wedges require less data and do not require assumptions or data on the term structure of beliefs. However, the main limitation of the static wedge is that it only embodies current distortions

$\frac{Cov\left(C_{i,t},C_{i,t}^{\star,\text{FIRE}}\right)}{Var(C_{i,t})} + \frac{Cov\left(C_{i,t},\eta_{i,t}^{\text{beliefs}}\right)}{Var(C_{i,t})} + \frac{Cov\left(C_{i,t},\eta_{i,t}^{\text{distortions}}\right)}{Var(C_{i,t})}$.

and therefore does not summarize the total impact of distortions, which can also arise through expected future distortions.

In Appendix B.4, we show that the dynamic wedge is a geometric sum of the current static wedge and expected future static wedges. Specifically:

$$\eta_{i,t} = \nu_{i,t} + \sum_{j=1}^{T} \left[\widetilde{E}_{i,t} \left(\nu_{i,t+j} \right) \prod_{k=0}^{j-1} m_{i,t+k} \right] \tag{10}$$

where $m_{i,t} = \frac{1}{\widetilde{E}_{i,t}\pi_{t+1}} \left(\beta \frac{\widetilde{E}_{i,t} R_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}} \right)^{-1/\gamma}$. This equation shows that the dynamic wedge corresponds to the total impact of both current and expected future distortions. It highlights that beliefs about future distortions can be important determinants of consumption.

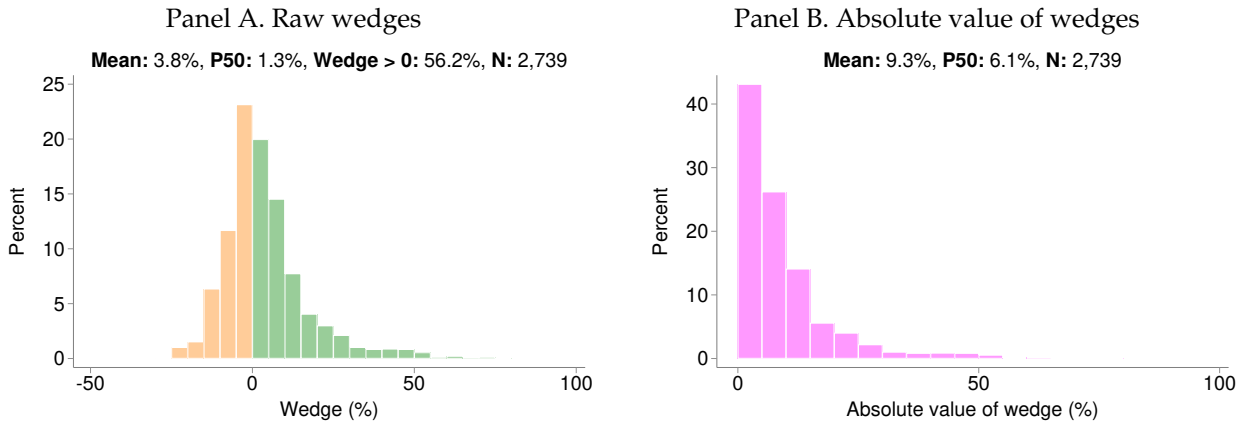
Figure 3 displays the distribution of static consumption wedges. We find a slightly larger share of over-consumers as measured by their static wedge: 56% (versus 51%). To see what this implies, note that if a consumer has a positive static wedge and a negative dynamic wedges, Equation (10) shows that it must be the case that the consumer anticipates facing negative static wedges in the future. This implies that at least $5/56 \approx 9\%$ of over-consumers (as measured by their static wedge) expect to face distortions later in their lifecycle that will distort their consumption downward. Considering that our sample is concentrated among relative young consumers, most of whom are not homeowners, this could plausibly be explained by present bias or inertia being the dominant distortion in their current lifecycle phase, with the anticipation of financial constraints becoming more important important as they accumulate illiquid assets like homes and retirement savings.

The difference in typical absolute value magnitudes—6.1% for static wedges versus 40% for dynamic wedges—implies that expectations about future distortions are an important determinant of consumption. Taking the ratio of individuals' static and dynamic wedges, for the median consumer, the static wedge accounts for 13.5% of their dynamic wedge. Beliefs about distortions are less-widely studied. Given the implied importance of these beliefs, our findings suggest that a useful direction for future research would be to further investigate perceptions of distortions (and the determinants of these perceptions). And if there are large *biases* in beliefs about future distortions, correcting these beliefs could significantly alter the consumption decisions of low-liquidity consumers.

5 Diagnostic Evidence from Wedge Correlates

Competing theories of consumption distortions have different testable implications for *both* wedges and the relationship between wedges and other variables of economic interest. In this section, we

Figure 3. Distribution of Static Consumption Wedges



Notes: The figures show the distribution of static consumption wedges (left) and their absolute values (right). Wedges are reported as the percent deviation of observed consumption and frictionless consumption, and they are trimmed at the 1st and 99th percentiles. Includes users from waves 2 and 3 only, as we did not solicit spending growth expectations (an input to the static wedges) in wave 1.

compare these model predictions to the empirical relationships we document for our sample by examining the correlation between consumption wedges and MPCs, financial distress, consumption commitments, and hypothetical responses to inflation and consumption decisions. This exercise provides evidence on which distortions can best explain the wedges and their correlates. As a final validation exercise, we correlate survey responses to hypothetical consumption decisions as a test that the wedges we estimate are consistent with the decision-making process that underlies consumption behavior.

MPCs. Here we explore the relationship between consumption wedges and MPCs. If financial constraints are the dominant distortion to consumption and the mechanism behind high MPCs, we would expect to see the highest MPCs for the most negative consumption wedges. To test this, we use the transactions data to 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.[24] 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,

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

2023, and 2024 as comparison periods. We calculate each individual's MPC as follows:

$$MPC_i = \frac{1}{StimulusAmount_i} \times (\Delta Spend_i^{2021} - \frac{\Delta Spend_i^{2022} + \Delta Spend_i^{2023} + \Delta Spend_i^{2024}}{3}) \quad (11)$$

where

$$\Delta Spend_i^t = Spend_i^{Post,t} - Spend_i^{Pre,t} \quad (12)$$

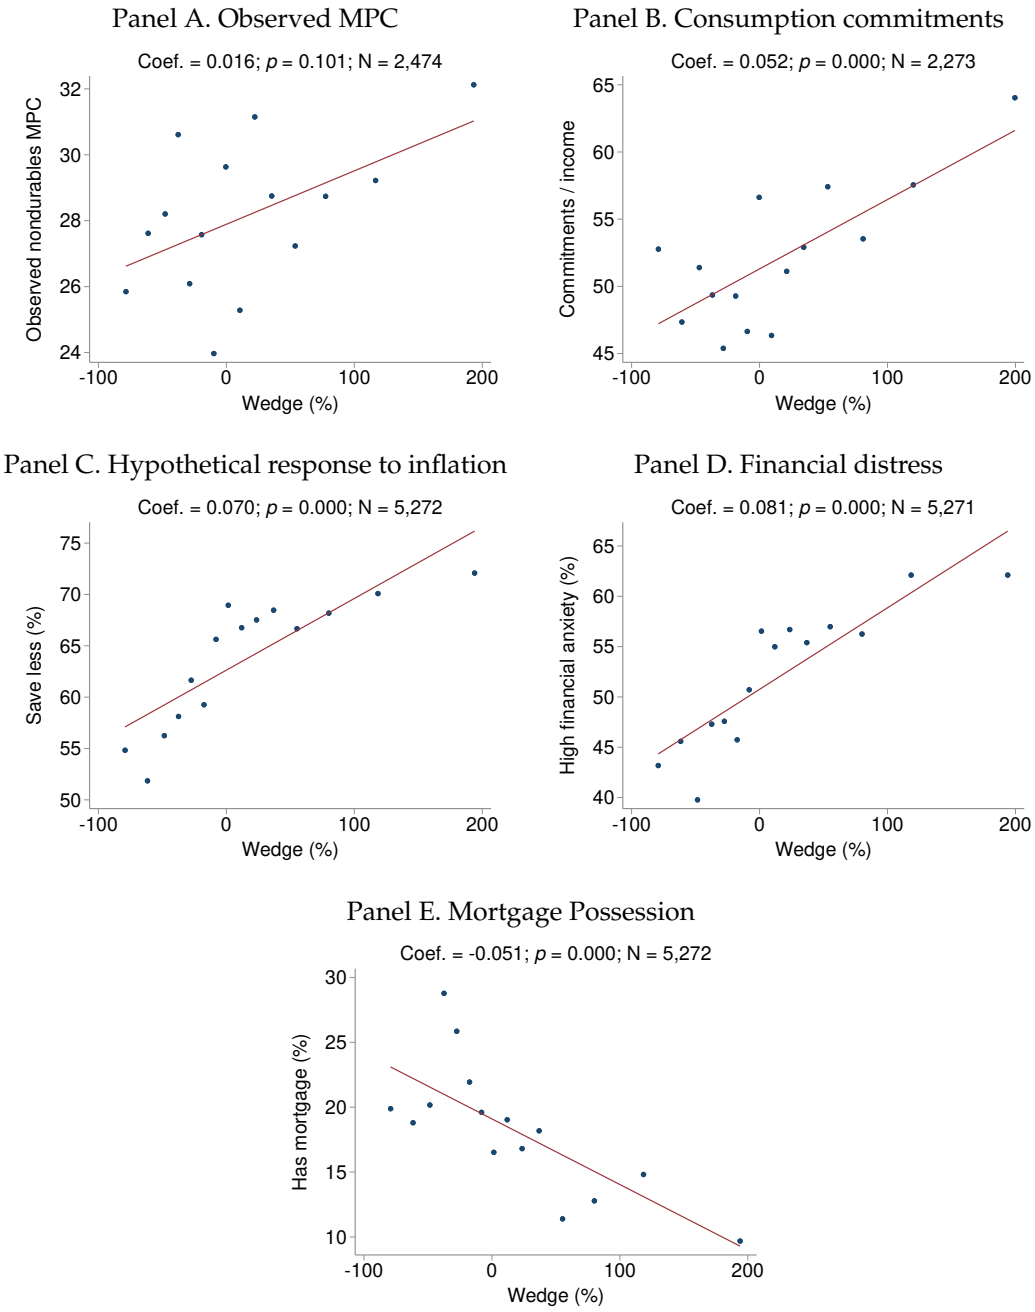
Our MPC measure captures the "excess" consumption associated with receipt of the stimulus check. We note that this measure should be interpreted as at best a proxy for an individual's MPC, as we only have three observations per person. As such, this measure is unlikely an asymptotically valid estimate of the individual's true MPC. The median estimated MPC is 25%. 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 trim MPCs at the 10th and 90th percentiles.

Figure 4 Panel A 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 0.4 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, showing that consumption wedges are strongly related to the behavioral responses of consumption. The second is that higher MPCs are associated with over-consumption, rather than the under-consumption. This suggests that the forces underlying high MPCs among low-liquidity consumers are predominantly attributable to forces that generate positive wedges, such as present bias and consumer inertia, rather than financial constraints.

Consumption Commitments. The insufficiency of financial constraints alone to rationalize the distribution of consumption wedges naturally raises the question of which distortions can best explain deviations from the benchmark. In survey waves 2 and 3, we asked consumers for their monthly housing and childcare costs as a measure of their "committed consumption," as these expenses tend to be large, difficult to adjust, and difficult to identify in transactions data. We divide these reported monthly expenditures by average 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 share of committed consumption. Figure 4 Panel B shows a strong positive relationship between wedges and consump-

Figure 4. Relationship Between Consumption Wedges and Nondurable MPCs



Notes: The figure illustrates the relationship between dynamic consumption wedges and observed nondurable MPCs, consumption commitments (% income), an indicator for whether the user would hypothetically "save less" in response to higher inflation, and an indicator for whether the user reports "high" or "very high" financial anxiety. The binned scatterplot plots the average value within quantile-based intervals of consumption wedges. Variable definitions are outlined in Appendix C.4. MPCs are observed for 58% of users in our analysis sample. Wedges are trimmed at the 1st and 95th percentiles and MPCs are trimmed at the 5th and 95th percentiles.

tion commitments. A 25 pp larger wedge is associated with a 1.3 pp higher ratio of committed consumption to income. This pattern suggests consumption commitments are a plausible friction behind the consumption wedges we measure.

Hypothetical Spending/Saving Behavior. Our survey also presents consumers with a hypothetical scenario that asks how they would adjust their saving behavior in response to higher expected inflation. The majority of respondents report that they would "save less" when future inflation is higher, and Figure 4 Panel C shows that this response is also positively correlated with the consumption wedges. The survey also solicited rationales for respondents' hypothetical savings behavior. Among respondents selecting "save less," a large majority (86%) attributed this dissaving to an *inability* to reduce spending. These results suggest an intuitive relationship between consumers' mental model of their own behavior and the wedges we estimate using their economic beliefs and actual consumption choices. This explanation most naturally relates to consumer inertia and provides some suggestive corroboration for the consumption commitments result.

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 the correlation with financial anxiety in Figure 4 Panel D and the remaining correlations in Appendix Figure A.3. For all measures we consider, we find a strong, positive relationship between consumption wedges and financial distress.

Not only is financial distress more common among over-consumers, but it is also rarer among under-consumers with larger (in magnitude) wedges. One possible rationalization for this phenomenon relates to wealthy hand-to-mouth consumers, as in Kaplan and Violante (2014). Those 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 this illiquid wealth may significantly distort their consumption downwards. The size of the distortion induced by these constraints is larger, all else equal, when the value of illiquid assets is greater. Hence, consumers with large to moderate negative distortions may on average be less financially distressed than those that are less able to acquire illiquid assets. Those who 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 consumers compared to the poor hand-to-mouth. Instead, present bias or consumer inertia is more likely the dominant friction

affecting consumers with low liquid wealth.

The positive association between over-consumption and financial distress suggests that the underlying frictions or behavioral preferences driving over-consumption are linked with lower wellbeing. Policymakers concerned with alleviating financial distress may therefore be better able to achieve this goal if policy is informed by theories of consumer behavior that feature over-consumption.

Homeownership. Motivated by our results regarding financial distress, we next compare wedges across homeowners and non-homeowners. We proxy for homeownership with whether the respondent has a mortgage. We find that homeowners on average plausible resemble the financially constrained "wealthy hand-to-mouth" households of Kaplan and Violante (2014). A majority of homeowners have negative wedges. Furthermore, there exists a strong negative correlation between having a mortgage and the size of a consumers wedge. These results are shown in Panel E of Figure 4. This is consistent with financial constraints being the dominant distortion for consumers with substantial illiquid wealth.

6 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 consumers 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 consumers 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 of 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 an important role in driving consumption and call into question the dominant role played by financial constraints in explaining the consumption of low-liquidity consumers. Most individuals in our sample have large distortions; the median (absolute value) distortion stands at 40% of frictionless consumption. The average distortion is 16%, but this belies significant heterogeneity in the cross section. In particular, there

is a mix of positive and negative wedges. 51% of wedges are positive (over-consumption) while the remaining 49% are negative. Because financial constraints can only generate negative wedges, the 51% of over-consumers cannot be explained by financial constraints. Additional or alternative mechanisms are necessary to explain the consumption choices of low-liquidity consumers. 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 consumers.

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 lifecycle 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 consumers 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

Aguiar, Mark, Mark Bils, and Corina Boar, "Who are the Hand-to-Mouth?," *Review of Economic Studies*, 2024, p. rdae056.

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

Andreoni, James and Charles Sprenger, "Estimating time preferences from convex budgets," *American Economic Review*, 2012, *102* (7), 3333–3356.

Attanasio, Orazio P and Guglielmo Weber, "Consumption growth, the interest rate and aggregation," *The Review of Economic Studies*, 1993, *60* (3), 631–649.

_ and _, "Consumption and saving: models of intertemporal allocation and their implications for public policy," *Journal of Economic literature*, 2010, *48* (3), 693–751.

_ and Nicola Pavoni, "Risk sharing in private information models with asset accumulation: Explaining the excess smoothness of consumption," *Econometrica*, 2011, *79* (4), 1027–1068.

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.

_, _, and _, "The intertemporal keynesian cross," *Journal of Political Economy*, 2024, *132* (12), 4068–4121.

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.

Bewley, Truman, "The Optimum Quantity of Money," in "Models of Monetary Economics" 1980, pp. 169–210.

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

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.

Campbell, John and Angus Deaton, "Why is consumption so smooth?," *The Review of Economic Studies*, 1989, *56* (3), 357–373.

Campbell, John Y and N Gregory Mankiw, "Consumption, income, and interest rates: Reinterpreting the time series evidence," *NBER macroeconomics annual*, 1989, *4*, 185–216.

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.

Colarieti, Roberto, Pierfrancesco Mei, and Stefanie Stantcheva, "The how and why of household reactions to income shocks," Technical Report, National Bureau of Economic Research 2024.

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.

Fagereng, Andreas, Martin B Holm, and Gisle J Natvik, "MPC heterogeneity and household balance sheets," *American Economic Journal: Macroeconomics*, 2021, *13* (4), 1–54.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde, "Global Evidence on Cconomic Preferences," *The Quarterly Journal of Economics*, 2018, *133* (4), 1645–1692.

Flavin, Marjorie A, "The adjustment of consumption to changing expectations about future income," *Journal of political economy*, 1981, *89* (5), 974–1009.

Floden, Martin and Jesper Lindé, "Idiosyncratic risk in the United States and Sweden: Is there a role for government insurance?," *Review of Economic dynamics*, 2001, *4* (2), 406–437.

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

Fuster, Andreas, Greg Kaplan, and Basit Zafar, "What would you do with $500? Spending responses to gains, losses, news, and loans," *The Review of Economic Studies*, 2021, *88* (4), 1760–1795.

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

Gross, Tal, Matthew J Notowidigdo, and Jialan Wang, "The marginal propensity to consume over the business cycle," *American Economic Journal: Macroeconomics*, 2020, *12* (2), 351–384.

Guerrieri, Veronica and Guido Lorenzoni, "Credit crises, precautionary savings, and the liquidity trap," *The Quarterly Journal of Economics*, 2017, *132* (3), 1427–1467.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What do data on millions of US workers reveal about lifecycle earnings dynamics?," *Econometrica*, 2021, *89* (5), 2303–2339.

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

Hall, Robert E, "Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence," *Journal of political economy*, 1978, *86* (6), 971–987.

Hansen, Lars Peter and Kenneth J Singleton, "Generalized instrumental variables estimation of nonlinear rational expectations models," *Econometrica: Journal of the Econometric Society*, 1982, pp. 1269–1286.

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.

_ and _, "The marginal propensity to consume in heterogeneous agent models," *Annual Review of Economics*, 2022, *14* (1), 747–775.

_, 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.

Krueger, Dirk, Kurt Mitman, and Fabrizio Perri, "Macroeconomics and household heterogeneity," in "Handbook of macroeconomics," Vol. 2, Elsevier, 2016, pp. 843–921.

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.

Lin, Daniel, "Methods and Assumptions of the CPS ASEC Tax Model," 2022. SEHSD Working Paper FY-2022-18.

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.

Ma, Yueran, Tiziano Ropele, David Sraer, and David Thesmar, "A quantitative analysis of distortions in managerial forecasts," Technical Report, National Bureau of Economic Research 2024.

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.

Zeldes, Stephen P, "Consumption and liquidity constraints: an empirical investigation," *Journal of political economy*, 1989, *97* (2), 305–346.

Appendix

Contents

A Appendix Figures **40**

B Theory Derivations and Extensions **44**

B.1 Frictionless Consumption Derivation . . . 44

B.2 Examples of the Effect of Distortions on Euler Equations and the Implications for Wedges . . . 47

B.3 Model Extensions . . . 48

B.3.1 Extension: Durable and Non-Durable Goods . . . 49

B.4 Static Wedges . . . 51

C Data Construction **54**

C.1 Transactions Data . . . 54

C.1.1 Data Structure . . . 54

C.1.2 Categorizing Transaction Inflows . . . 54

C.1.3 Categorizing Transaction Outflows . . . 55

C.2 Survey Outreach and Response . . . 58

C.3 Sample Restrictions . . . 59

C.4 Variable Measurement . . . 61

C.4.1 Consumption, Income, and APCs . . . 61

C.4.2 Wealth-to-Income Ratio . . . 61

C.4.3 Beliefs . . . 62

C.4.4 Marginal Propensity to Consume . . . 63

C.4.5 Financial Distress . . . 63

D Imputations **63**

D.1 Wealth . . . 63

D.2 Term Structure of Beliefs . . . 64

D.2.1 Inflation Expectations . . . 64

D.2.2 Interest Rate Expectations . . . 65

D.2.3 Income Growth Expectations . . . 65

E Representativeness of the EarnIn Sample **67**

E.1 Wealth Distribution . . . 67

E.2 Income Distribution . . . 67

E.3 Subjective Expectations . . . 68

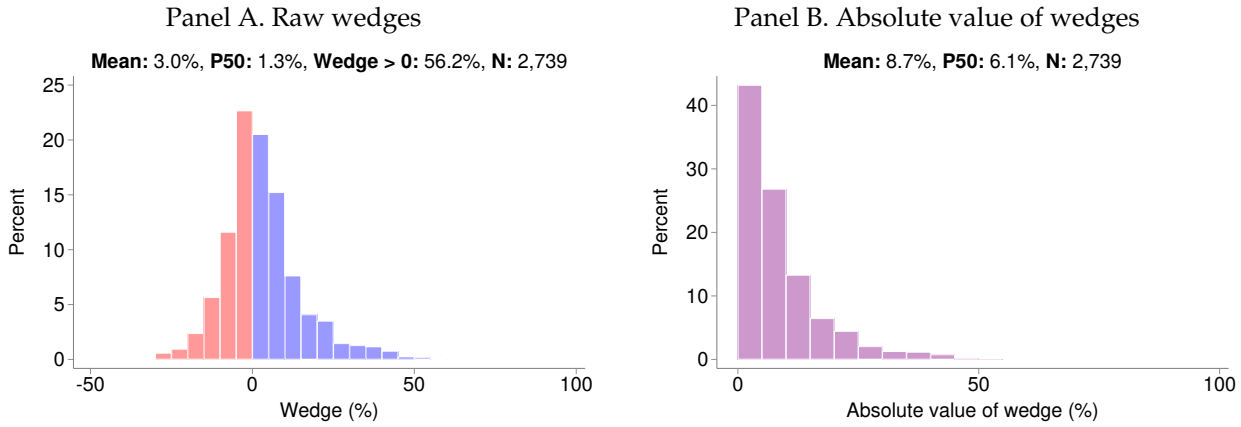
F Robustness **72**
F.1 Sensitivity Analysis 72
F.2 Measurement Error 78

G Quantitative Model Appendix **81**
G.1 Quantitative Model Calibration 81
G.2 Approximation Bias Analysis 81
G.3 Model Comparison 81

H Additional Figures **83**

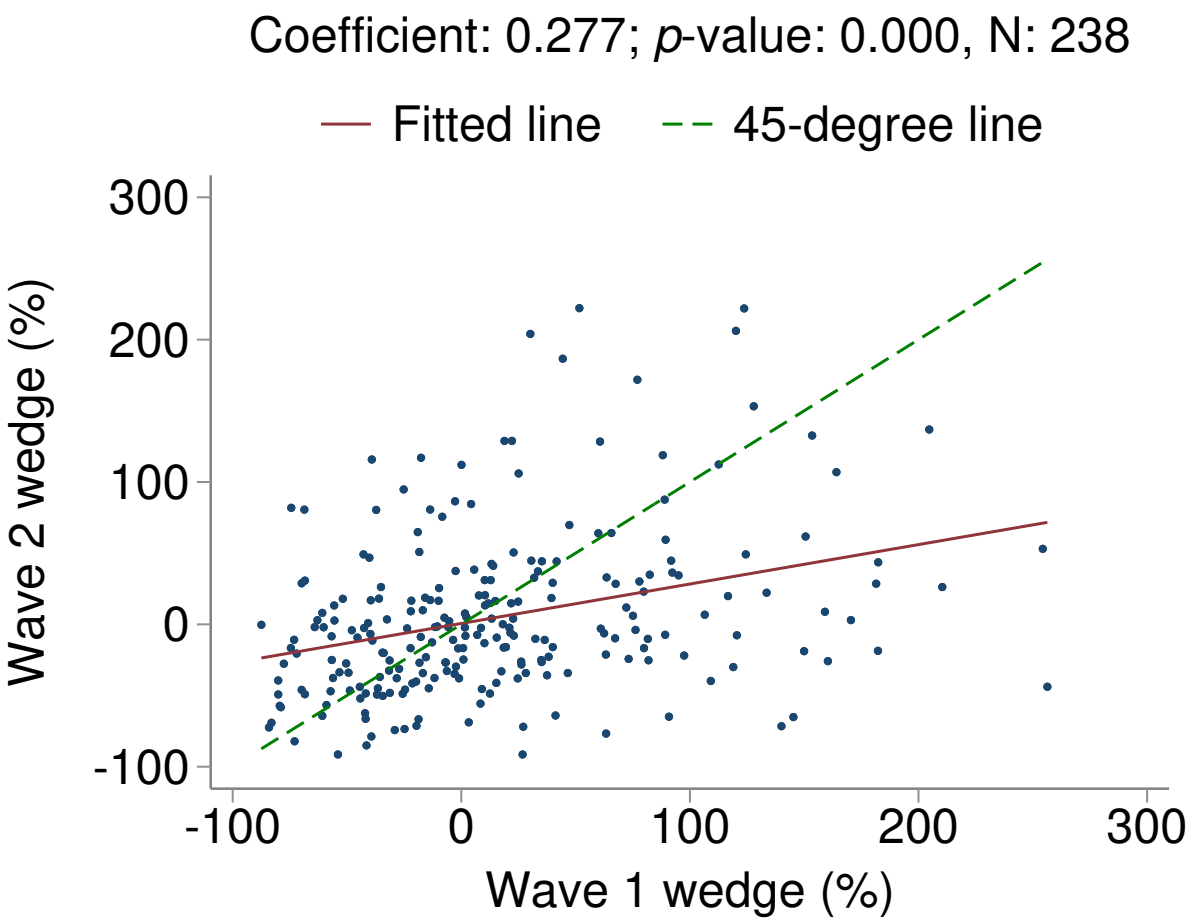
A Appendix Figures

Figure A.1. Distribution of Static Consumption Wedges



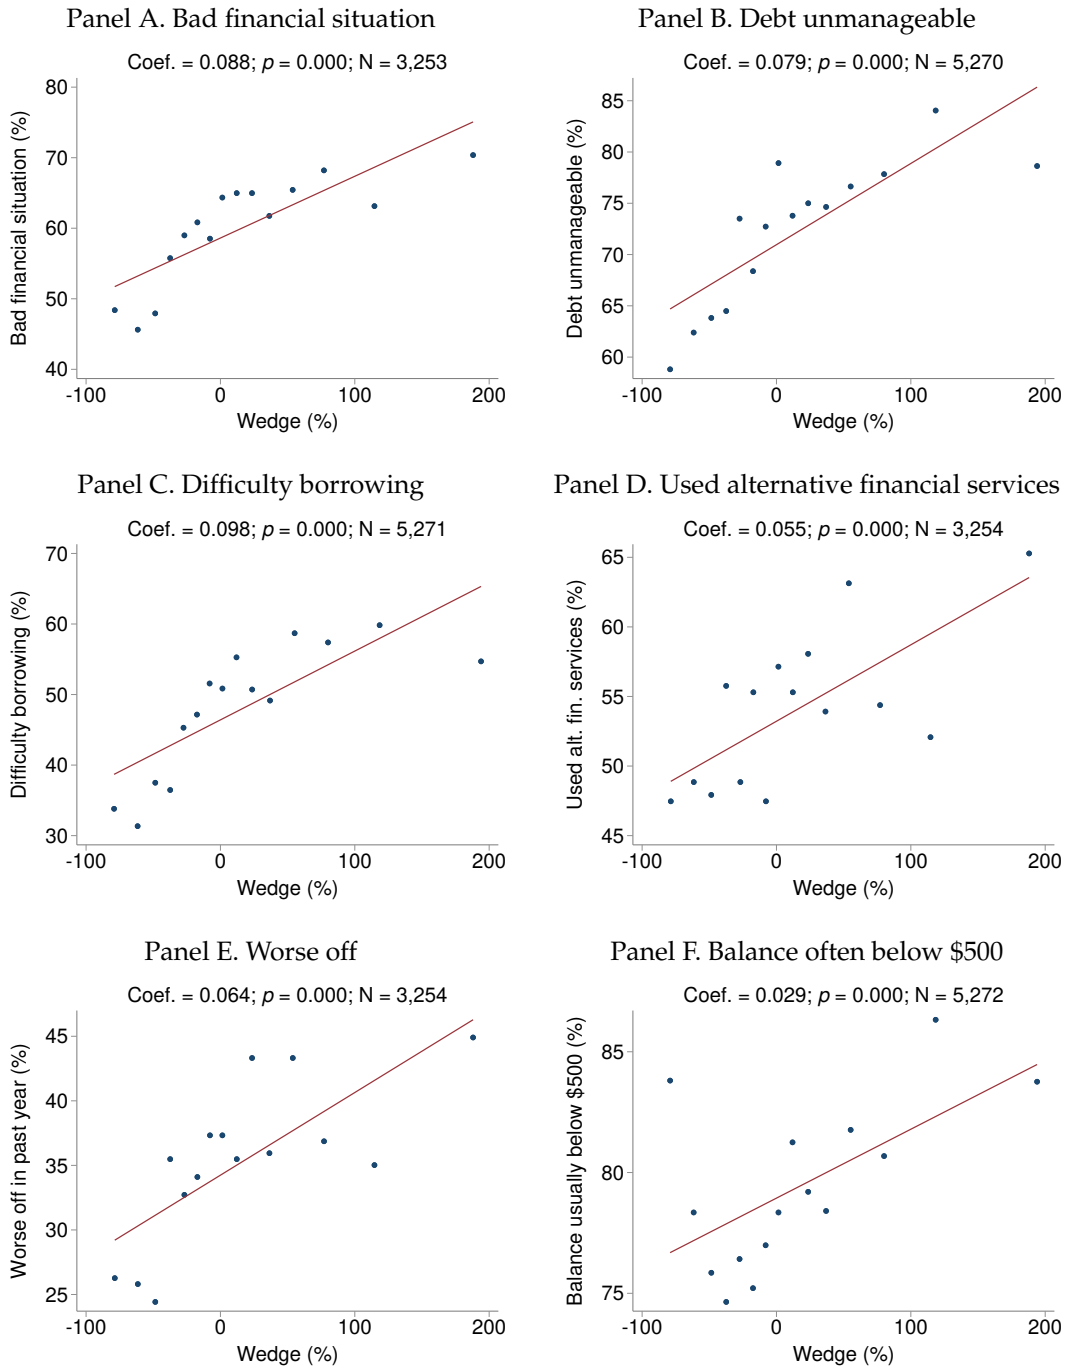
Notes: The figures show the distribution of static consumption wedges (left) and their absolute values (right). Wedges are trimmed at the 1st and 99th percentiles. Includes users from waves 2 and 3 only, as we did not solicit spending growth expectations (an input to the static wedges) in wave 1.

Figure A.2. Relationship Between Wave 1 and Wave 2 Wedges



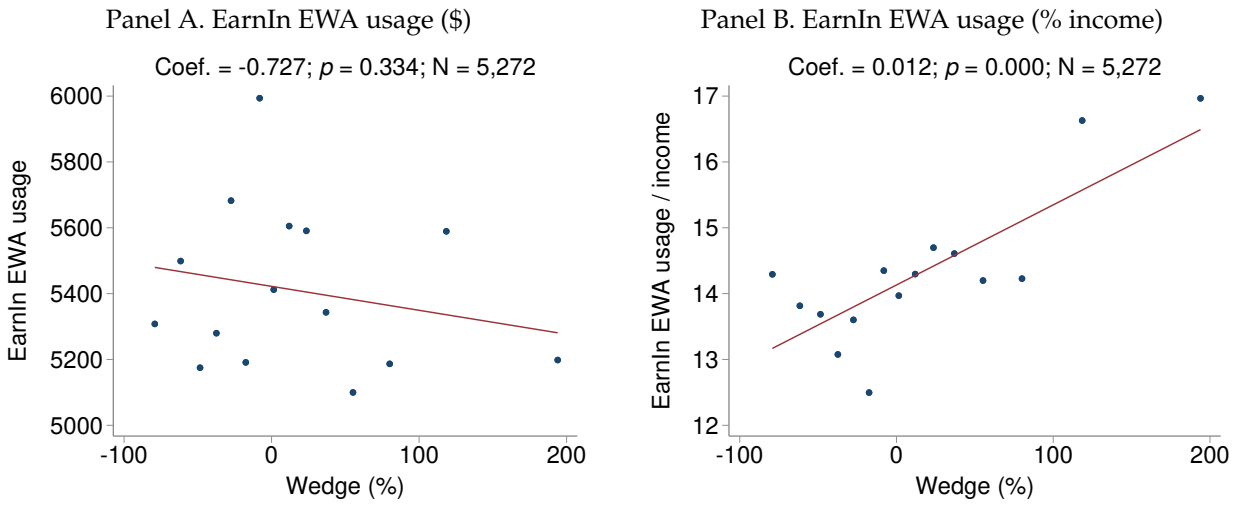
Notes: The figure illustrates the relationship between dynamic consumption wedges in wave 1 and wave 2, restricting to users who completed wave 2. The solid red line represents a linear fitted line, and the dashed green line represents a 45-degree line. The coefficient and *p*-value from a regression of wave-2 wedges on wave-1 wedges are shown at the top of the figure. Wedges are trimmed at the 1st and 95th percentiles.

Figure A.3. Relationship Between Dynamic Consumption Wedges and Financial Distress



Notes: The figures illustrate the relationship between dynamic consumption wedges and six indicators of financial distress: (1) whether the user reports "just getting by" or "finding it difficult to get by;" (2) whether the user reports having "a bit more" or "far more" debt than is manageable; (3) whether the user reports difficulty borrowing due to "often" or "most of the time" being denied for credit; (4) whether the user reports using alternative financial services in the past 3 months (wave 1 only); (5) whether the user reports being "somewhat" or "much" worse off financially compared to 12 months ago; and (6) whether the user's observed balances are below $500 for more than 50% of the pre-survey period. Each binned scatterplot plots the average value of the financial distress indicator within quantile-based intervals of consumption wedges. Wedges are trimmed at the 1st and 95th percentiles.

Figure A.4. Relationship Between Dynamic Consumption Wedges and EarnIn EWA Usage



Notes: The figures illustrate the relationship between dynamic consumption wedges and two measures of EWA usage over the 12 months preceding the survey: (1) the dollar amount of EarnIn EWA outflows and (2) the ratio of EarnIn EWA outflows to income. Each binned scatterplot plots the average value of the EWA usage variable within quantile-based intervals of consumption wedges. Wedges are trimmed at the 1st and 95th percentiles.

B Theory Derivations and Extensions

B.1 Frictionless Consumption Derivation

We approximate frictionless consumption using first-order approximation of the style $f(x) \approx f[\widetilde{E}_t(x)] + f'[\widetilde{E}_t(x)][x - \widetilde{E}_t(x)]$, which yield $\widetilde{E}_t[f(x)] \approx f[\widetilde{E}_t(x)]$. We begin by approximating the Euler equation.

Lemma 1: Euler Equation Approximation

In the frictionless benchmark, the one-period-ahead and multi-period Euler equations are approximately

$$C^\star_{i,t} \approx \frac{\widetilde{E}_{i,t} C^\star_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}} \left(\beta \frac{\widetilde{E}_{i,t} R_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}}\right)^{-1/\gamma}$$

$$C^\star_{i,t} \approx \widetilde{E}_{i,t} C^\star_{i,t+j} \prod_{k=1}^{j} \left[\frac{1}{\widetilde{E}_{i,t}\pi_{t+k}} \left(\beta \frac{\widetilde{E}_{i,t} R_{i,t+k}}{\widetilde{E}_{i,t}\pi_{t+k}}\right)^{-1/\gamma}\right],$$

where $\gamma_{i,t} = -\frac{u''(c^\star_{i,t})}{u'(c^\star_{i,t})} c^\star_{i,t}$, is consumer i's inverse IES/coefficient of relative risk aversion evaluated at time t frictionless consumption $c^\star_t$. We assume a common $\gamma_{i,t} = \gamma$ in our baseline analysis. We relax this assumption in an alternative calibration.

Proof. Taking a first-order approximation of the Euler equation

$$u'\left(\frac{C^\star_{i,t}}{P_t}\right) = \beta \widetilde{E}_{i,t}\left[u'\left(\frac{C^\star_{i,t+1}}{P_{t+1}}\right)\frac{R_{i,t+1}}{\pi_{t+1}}\right]$$

yields

$$u'\left(\frac{C^\star_{i,t}}{P_t}\right) \approx \beta u'\left(\frac{\widetilde{E}_{i,t} C^\star_{i,t+1}}{\widetilde{E}_{i,t} P_{t+1}}\right)\frac{\widetilde{E}_{i,t} R_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}}.$$

Under CRRA preferences, $u'(c) = c^{-\gamma}$, and we can rearrange the above to obtain:

$$C^\star_{i,t} \approx \frac{\widetilde{E}_{i,t} C^\star_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}} \left(\beta \frac{\widetilde{E}_{i,t} R_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}}\right)^{-1/\gamma}.$$

Under non-CRRA preferences, we obtain the same expression using the a log-linear first-order

approximation of the marginal utility function below:

$$\begin{aligned}\ln\left[u'(c)\right] &\approx \ln\left[u'(\bar{c})\right] + \frac{u''(\bar{c})}{u'(\bar{c})}\bar{c}\left(\ln c - \ln \bar{c}\right)\\ \ln\left[\frac{u'(c)}{u'(\bar{c})}\right] &\approx -\gamma\left(\ln c - \ln \bar{c}\right)\\ \frac{u'(c)}{u'(\bar{c})} &\approx \left(\frac{c}{\bar{c}}\right)^{-\gamma}.\end{aligned}$$

The second line uses $\gamma = -\frac{u''(\bar{c})}{u'(\bar{c})}\bar{c}$. For the approximate Euler equation, $\gamma_{i,t}$ is evaluated at (real) frictionless consumption: $\gamma_{i,t} = -\frac{u''(c^\star_{i,t})}{u'(c^\star_{i,t})}c^\star_{i,t}$ for consumer i at time t. Recall that we assume $\gamma = \gamma_{i,t}$ is constant and common across consumers in our baseline calibration.

The derivation for the multi-period Euler equation:

$$u'\left(\frac{C^\star_{i,t}}{P_t}\right) = \widetilde{E}_{i,t}\left[\beta^j u'(C^\star_{i,t+j})\prod_{k=1}^{j}\frac{R_{i,t+j}}{\pi_{t+j}}\right]$$

is similar, yielding:

$$C^\star_{i,t} \approx \widetilde{E}_{i,t}C^\star_{i,t+j}\prod_{k=1}^{j}\left[\frac{1}{\widetilde{E}_{i,t}\pi_{t+k}}\left(\beta\frac{\widetilde{E}_{i,t}R_{i,t+k}}{\widetilde{E}_{i,t}\pi_{t+k}}\right)^{-1/\gamma}\right].$$

□

We next take a first-order approximation of the forward-iterated budget constraint.

Lemma 2: Budget Constraint Approximation

Under a no-Ponzi condition, $\lim_{j\to\infty}\frac{A_{i,t+j}}{R_{i,t+1}\cdots R_{i,t+j-1}} = 0$*, the expected forward-iterated budget constraint is approximately:*

$$C^\star_{i,t} \approx A_{i,t}R_{i,t} + Y_{i,t} + Y_{i,t}\sum_{j=1}^{T}\left(\frac{\widetilde{E}_{i,t}G_{i,t,t+j}}{\prod_{k=1}^{j}\widetilde{E}_{i,t}R_{i,t+k}}\right) - \sum_{j=1}^{T}\left(\frac{\widetilde{E}_{i,t}C^\star_{i,t+j}}{\prod_{k=1}^{j}\widetilde{E}_{i,t}R_{i,t+k}}\right)$$

where $G^Y_{i,t,t+j} = \frac{Y_{i,t+j}}{Y_{i,t}}$ *is the gross nominal growth rate of income from period* t *to period* $t+j$ *for consumer* i*.*

Proof. We begin with the period t budget constraint:

$$C^\star_{i,t} + A_{i,t+1} = Y_{i,t} + A_{i,t}R_{i,t}.$$

We then forward iterate the budget constraint and apply the no-Ponzi condition:

$$C^{\star}_{i,t} = A_{i,t}R_{i,t} + Y_{i,t} + \sum_{j=1}^{T} \frac{Y_{i,t+j}}{\prod_{k=1}^{j} R_{i,t+k}} - \sum_{j=1}^{T} \frac{C^{\star}_{i,t+j}}{\prod_{k=1}^{j} R_{i,t+k}}.$$

Next, we rewrite income in terms of cumulative income growth rates $Y_{i,t+j} = Y_{i,t}G^{Y}_{i,t,t+j}$:

$$C^{\star}_{i,t} = A_{i,t}R_{i,t} + Y_{i,t} + Y_{i,t}\sum_{j=1}^{T} \frac{G^{Y}_{i,t,t+j}}{\prod_{k=1}^{j} R_{i,t+k}} - \sum_{j=1}^{T} \frac{C^{\star}_{i,t+j}}{\prod_{k=1}^{j} R_{i,t+k}}.$$

And then we take expectations:

$$C^{\star}_{i,t} = A_{i,t}R_{i,t} + Y_{i,t} + Y_{i,t}\sum_{j=1}^{T} \widetilde{E}_{i,t}\left(\frac{G^{Y}_{i,t,t+j}}{\prod_{k=1}^{j} R_{i,t+k}}\right) - \sum_{j=1}^{T} \widetilde{E}_{i,t}\left(\frac{C^{\star}_{i,t+j}}{\prod_{k=1}^{j} R_{i,t+k}}\right)$$

and then a first-order approximation:

$$C^{\star}_{i,t} \approx A_{i,t}R_{i,t} + Y_{i,t} + Y_{i,t}\sum_{j=1}^{T} \left(\frac{\widetilde{E}_{i,t}G^{Y}_{i,t,t+j}}{\prod_{k=1}^{j} \widetilde{E}_{i,t}R_{i,t+k}}\right) - \sum_{j=1}^{T} \left(\frac{\widetilde{E}_{i,t}C^{\star}_{i,t+j}}{\prod_{k=1}^{j} \widetilde{E}_{i,t}R_{i,t+k}}\right).$$

□

Our next proposition derives our characterization of frictionless consumption using the approximation from the preceding lemmas.

Proposition 1: Frictionless Consumption

Combining the approximations from Lemmas 1 and 2, we can characterize (approximate) frictionless consumption as follows:

$$C^{\star}_{i,t} \approx \frac{A_{i,t}R_{i,t} + Y_{i,t} + Y_{i,t}\sum_{j=1}^{T}\left[\widetilde{E}_{i,t}G^{Y}_{i,t,t+j}\prod_{k=1}^{j}\left(\widetilde{E_{i,t}}R_{i,t+k}\right)^{-1}\right]}{1 + \sum_{j=1}^{T}\left\{\prod_{k=1}^{j}\left[\beta^{1/\gamma}\left(\frac{\widetilde{E}_{i,t}R_{t+k}}{\widetilde{E}_{i,t}\pi_{t+k}}\right)^{1/\gamma-1}\right]\right\}}.$$

Proof. We begin with the approximate Euler equation from Lemma 1. We rearrange it to write it in terms of expected future consumption:

$$\widetilde{E}_{i,t}C^{\star}_{i,t+j} \approx C^{\star}_{i,t}\prod_{k=1}^{j}\left[\widetilde{E}_{i,t}\pi_{t+k}\left(\beta\frac{\widetilde{E}_{i,t}R_{i,t+k}}{\widetilde{E}_{i,t}\pi_{t+k}}\right)^{1/\gamma}\right].$$

Using the above expression, we substitute in for $\widetilde{E}_{i,t}C^{\star}_{i,t+j}$ in the approximate budget constraint

equation from Lemma 2:

$$
\begin{aligned}
C^{\star}_{i,t} &\approx A_{i,t}R_{i,t} + Y_{i,t} + Y_{i,t}\sum_{j=1}^{T}\left(\frac{\widetilde{E}_{i,t}G^{Y}_{i,t,t+j}}{\prod_{k=1}^{j}\widetilde{E}_{i,t}R_{i,t+k}}\right) - \sum_{j=1}^{T}\left(\frac{\widetilde{E}_{i,t}C^{\star}_{i,t+j}}{\prod_{k=1}^{j}\widetilde{E}_{i,t}R_{i,t+k}}\right) \\
&= A_{i,t}R_{i,t} + Y_{i,t} + Y_{i,t}\sum_{j=1}^{T}\left(\frac{\widetilde{E}_{i,t}G^{Y}_{i,t,t+j}}{\prod_{k=1}^{j}\widetilde{E}_{i,t}R_{i,t+k}}\right) - C^{\star}_{i,t}\sum_{j=1}^{T}\left\{\prod_{k=1}^{j}\left[\beta^{1/\gamma}\left(\frac{\widetilde{E}_{i,t}R_{i,t+k}}{\widetilde{E}_{i,t}\pi_{t+k}}\right)^{1/\gamma-1}\right]\right\} \\
&= \frac{A_{i,t}R_{i,t} + Y_{i,t} + Y_{i,t}\sum_{j=1}^{T}\left[\widetilde{E}_{i,t}G^{Y}_{i,t,t+j}\prod_{k=1}^{j}\left(\widetilde{E_{i,t}}R_{i,t+k}\right)^{-1}\right]}{1+\sum_{j=1}^{T}\left\{\prod_{k=1}^{j}\left[\beta^{1/\gamma}\left(\frac{\widetilde{E}_{i,t}R_{i,t+k}}{\widetilde{E}_{i,t}\pi_{t+k}}\right)^{1/\gamma-1}\right]\right\}}.
\end{aligned}
$$

□

B.2 Examples of the Effect of Distortions on Euler Equations and the Implications for Wedges

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 $A_{i,t+1} \geq \overline{A}$. The Euler equation would be:

$$
u'\left(\frac{C_{i,t}}{P_t}\right) = \beta\widetilde{E}_{i,t}\left[u'\left(\frac{C_{i,t+1}}{P_{t+1}}\right)\frac{R_{i,t+1}}{\pi_{t+1}}\right] + \mu_{i,t}
$$

where $\mu_{i,t} \geq 0$ is the Lagrange multiplier. The Lagrange multiplier $\mu_{i,t}$ is positive if and only if the constraint is binding. All else equal, a binding constraint reduces consumption $C_{i,t}$. 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 ($\overline{\beta}$):

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

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 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 adjusting their consumption: $\phi(C_{i,t} - C_{i,t-1})$. In such a case, the Euler equation would have additional terms reflecting this cost:

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

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 (overconsumption). 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 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.

B.3 Model Extensions

Additional Choices. Allowing more choice variables, such as labor supply, does not change the consumption wedge formula nor its interpretation. Each additional choice entails another

equation necessary to characterize *all* optimal choices, but the presence of these equations does not affect the characterization of optimal consumption in the *class* of models where a budget constraint and Euler equations are necessary conditions for optimality. Our frictionless consumption formula applies in this class; this formula does not request that the Euler equation and budget constraint are sufficient to characterize *all* optimal choices. If these additional choices are also subject to frictions, such as labor income taxes, the consumption wedge formula is unchanged. One only needs to make sure that the appropriate measures are used. In the case of labor income taxes, this means measuring income on an after-tax basis.

Additional Assets. It is straightforward to modify our benchmark model to feature additional assets. These could include housing, equity, or even the case of complete asset markets. For every asset added to the consumer's choice set, there is an additional first-order condition—specifically, an Euler equation—associated with that asset. Which Euler equation should be used to measure frictionless consumption? Because frictionless consumption is derived by combining the budget constraint and an Euler equation, and wealth in the budget constraint depends on the entire portfolio of assets, the appropriate Euler equation is one containing the portfolio-weighted expected return across all assets. This can be derived by taking a portfolio-weighted sum of each asset's Euler equation.

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

$$\widetilde{E}_{i,t}R_{i,t} = \frac{S_{i,t}}{S_{i,t} - D_{i,t}}\widetilde{E}_{i,t}R^S_{i,t} - \frac{D_{i,t}}{S_{i,t} - D_{i,t}}\widetilde{E}_{i,t}R^D_{i,t}$$

where $S_{i,t}$ are their liquid assets and $D_{i,t}$ their total liabilities. We assume that the return on this portfolio of liquid asset and debt, which excludes illiquid assets, is the same as the return they expect on their portfolio of illiquid assets. In practice, one could measure the overall portfolio returns using beliefs data on all individual assets and liabilities (which could make for a very demanding survey) or, more simply, groups of assets like our survey.

B.3.1 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. While these feature do not complicate or alter the derivation of the wedge, the present measurement challenges that complicate applying our wedge formula. 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 $n_{i,t}$ and $d_{i,t}$ denote i's real period t consumption flows of non-durable and durable goods (respectively). We continue to denote the total nominal value of net worth by $A_{i,t}$. Total wealth includes net positions in durables (e.g., the value of vehicles net of the loans used to finance their purchase). The consumer has preferences over notional consumption flows $c_{i,t}$, which are an aggregate of non-durable and durable consumption flows (i.e., utility $u(c_{i,t})$ 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 consumer can frictionlessly buy or sell durables at a spot price. The consumer can also rent durable goods at a per period rental price of q_t. No arbitrage in the durable goods markets requires that the rental price q_t equal the user cost of the durable goods.

By assuming that consumers 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 consumer 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, $n_{i,t}$, be the numeraire good. Under Assumption 1, we can write the consumer's budget constraint simply as

$$A_{i,t+1} + P_t c_{i,t} = Y_{i,t} + A_{i,t} R_{i,t}$$

where

$$P_t c_{i,t} = n_{i,t} + q_t d_{i,t}.$$

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 $c_{i,t}$ 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 intratemporal optimality with respect to the allocation of spending between non-durable and durable consumption.

Assumption 2: Cobb Douglas Aggregation.

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

$$c_{i,t} = n_{i,t}^{\alpha} d_{i,t}^{1-\alpha}.$$

Under Assumptions 1 and 2, the intratemporal optimality conditions are:

$$n_{i,t} = \alpha P_t c_{i,t}$$
$$d_{i,t} q_t = (1-\alpha) P_t c_{i,t}.$$

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 nominal notional consumption $C_{i,t}$ from nominal non-durable consumption $n_{i,t}$ and the expenditure share α. This is formalized in the lemma below.

Lemma 3: Consumption Calculation Including Consumption of Durables

Under Assumptions 1 and 2, notional consumption $C_{i,t}$ is

$$C_{i,t} = P_t c_{i,t} = \frac{n_{i,t}}{\alpha}$$

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

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

B.4 Static Wedges

Below, we reproduce the approximate Euler equation but instead write it in terms *actual* expected consumption, rather than expected frictionless consumption. That is, we write $\widetilde{E}_{i,t} C_{i,t+1}$ instead of $\widetilde{E}_{i,t} C^{\star}_{i,t+1}$.

$$C^{\star,\text{static}}_{i,t} \approx \frac{\widetilde{E}_{i,t} C_{i,t+1}}{\widetilde{E}_{i,t} \pi_{t+1}} \left(\beta \frac{\widetilde{E}_{i,t} R_{i,t+1}}{\widetilde{E}_{i,t} \pi_{t+1}} \right)^{-1/\gamma}.$$

Both expected consumption terms feature subjective expectations, $\widetilde{E}_{i,t}(\cdot)$, and hence depend on the same beliefs about income, inflation, and returns. They differ in that actual consumption $C_{i,t+1}$ may be distorted by frictions or behavioral preferences. For example, if a consumer expects to be borrowing constrained in the future, all else equal, she will expect lower actual consumption compared to frictionless consumption $\widetilde{E}_{i,t} C_{i,t+1} < \widetilde{E}_{i,t} C^{\star}_{i,t+1}$.

This equation yields an *alternative* notation of frictionless consumption that we call *static* frictionless consumption, denoted $C^{\star,\text{static}}_{i,t}$. We can use static frictionless consumption to measure a "static" consumption wedge. This consumption wedge is static in the sense that it captures only the influence of distortions experienced in time t. For example, if a borrowing constraint is binding in time t, the difference between actual consumption $C_{i,t}$ and static frictionless consumption $C^{\star,\text{static}}_{i,t}$ would include the effect of this distortion. But because $C^{\star,\text{static}}_{i,t}$ is calculated using actual expectations over future consumption, which the consumer may expect to be altered by distortions, this difference would not reflect the influence of future expected distortions. The static

consumption wedge is

$$\nu_{i,t} = C_{i,t} - C_{i,t}^{\star,\text{static}} = C_{i,t}\left[1 - \frac{\widetilde{E}_{i,t}G_{i,t+1}^{C}}{\widetilde{E}_{i,t}\pi_{t+1}}\left(\beta\frac{\widetilde{E}_{i,t}R_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}}\right)^{-1/\gamma}\right] \tag{13}$$

where $G_{i,t+1}^{C} = \frac{\widetilde{E}_{i,t}C_{i,t+1}}{C_{i,t}}$ is actual expected future consumption growth.

The wedge that is the focus of our analysis ($\eta_{i,t} = C_{i,t} - C_{i,t}^{\star}$) can be thought of as a "dynamic" wedge, as it does capture the influence of both present expected future distortions. The Lemma below characterizes the relationship between the dynamic wedge and the static wedge(s). It shows that the dynamic wedge is geometric sum of static wedges.

Proposition 2: Relation Between Dynamic and Static Wedges

The dynamic wedge can be written as a geometric sum of the static wedges, shown below:

$$\eta_{i,t} = \nu_{i,t} + \sum_{j=1}^{T}\left[\widetilde{E}_{i,t}\left(\nu_{i,t+j}\right)\prod_{k=0}^{j-1}m_{i,t+k}\right]$$

where

$$m_{i,t} = \frac{1}{\widetilde{E}_{i,t}\pi_{t+1}}\left(\beta\frac{\widetilde{E}_{i,t}R_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}}\right)^{-1/\gamma}.$$

Proof. We begin by establishing notation. Let

$$m_{i,t} = \frac{1}{\widetilde{E}_{i,t}\pi_{t+1}}\left(\beta\frac{\widetilde{E}_{i,t}R_{i,t+1}}{\widetilde{E}_{i,t}\pi_{t+1}}\right)^{-1/\gamma}.$$

This lets us rewrite the approximate Euler equations relating to frictionless consumption (both dynamic and static) as:

$$\begin{aligned} C_{i,t}^{\star} &= \widetilde{E}_{i,t}C_{i,t+1}^{\star}m_{i,t} \\ C_{i,t}^{\star,\text{static}} &= \widetilde{E}_{i,t}C_{i,t+1}m_{i,t}. \end{aligned}$$

Next, we start with the definition of the dynamic consumption wedge and then rewrite it using

the definitions of static and frictionless consumption:

$$\begin{aligned}
\eta_{i,t} &= C_{i,t} - C^{\star}_{i,t} \\
&= C^{\star \text{static}}_{i,t} + \nu_{i,t} - \widetilde{E}_{i,t} C^{\star}_{i,t+1} m_{i,t} \\
&= \nu_{i,t} + \widetilde{E}_{i,t} C_{i,t+1} m_{i,t} - \widetilde{E}_{i,t} C^{\star}_{i,t+1} m_{i,t} \\
&= \nu_{i,t} + \widetilde{E}_{i,t} \left(C_{i,t+1} - C^{\star}_{i,t+1} \right) m_{i,t} \\
&= \nu_{i,t} + \widetilde{E}_{i,t} \left(\eta_{i,t+1} \right) m_{i,t}.
\end{aligned}$$

The last line reveals a recursive relationship. Rewriting the above and applying the law of iterated expectations, we obtain:

$$\eta_{i,t} = \nu_{i,t} + \sum_{j=1}^{T} \left[\widetilde{E}_{i,t} \left(\nu_{i,t+j} \right) \prod_{k=0}^{j-1} m_{i,t+k} \right].$$

□

C Data Construction

C.1 Transactions Data

C.1.1 Data Structure

We receive anonymized data from EarnIn that covers user-level information, bank transactions, daily checking and savings account balances, and transactions classified as earnings. All data are de-identified and stored on secured servers. The dataset spans January 2021 through November 2024, covering at least 12 months before each survey wave. Users remain in the dataset until they either delete their accounts or disconnect their linked bank account. Among the users in our analysis sample, 97% had at least one EWA cashout in the 12 months preceding the survey, with a median annual cashout amount of approximately $4,700.

User-Level Data. We receive, weekly, user-level datasets that include both time-invariant variables (e.g., EarnIn sign-up date) and time-varying variables (e.g., number of hours worked in the last 7 days). These tags are merged into each of the other datasets.

Balances. The balances dataset provides daily records of the number and total balances of checking, savings, and "other" bank accounts linked to EarnIn. We do not observe balances for unlinked bank accounts or investment accounts.

Transactions. The transactions dataset includes transaction-level records with the transaction date, dollar amount, a memo describing the source or destination, and a transaction category assigned by Plaid, a third-party service that connects users' bank accounts to EarnIn. We do not observe transactions associated with unlinked bank accounts or credit cards.

Earnings. The earnings dataset is a direct subset of the transactions dataset, limited to earnings inflows from jobs reported to EarnIn. Each record includes the payment date, posted date, dollar amount, and whether those earnings are from unemployment benefits. Earnings are observed net of taxes and payroll deducations when deposited into users' linked bank accounts; we do not receive information on gross pay or withheld amounts.

C.1.2 Categorizing Transaction Inflows

We leverage the transactions and observed earnings datasets to construct a measure of 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.

C.1.3 Categorizing Transaction Outflows

Our analysis focuses on nondurables spending. To construct this measure, we apply an outflows categorization algorithm that separates durables and nondurables spending from other types of outflows, including payments (e.g., interest and principal 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 builds on the approaches of Ganong and Noel (2019) and Lusardi (1996), with modifications tailored to the structure of our data and the goals of our analysis.

The Plaid transaction taxonomy included in the EarnIn database comprises over 500 granular categories. We first map these categories to 33 broader categories that can be grouped under three overarching types: spending, payments, and 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 has three key limitations. First, Plaid's categorization is based on merchant types rather than the specific products and services purchased, making it difficult to distinguish between durables and nondurables in certain cases. For example, a transaction at a department store may include both a mattress (durable) and makeup (nondurable). Second, some Plaid categories are too broad or ambiguous to be mapped reliably, such as "Purchase," "Shopping," or "Transfer." Finally, we observe frequent misclassifications in Plaid's categorization.

To address the first limitation, we reallocate six spending categories that combine durables and nondurables: department stores, discount stores, drug stores, grocery stores, wholesale stores, and other retail (e.g., Amazon). For the first five categories, we follow Ganong and Noel (2019), who analyze 10-K filings from leading merchants in each category (e.g., CVS and Walgreens for drug stores, Macy's for department stores), calculate revenue by product type, and split each category across durable and nondurable spending categories. To categorize "other retail," we follow the composition of US ecommerce revenue in 2020 (Bledsoe, 2024). Appendix Table C.1 summarizes these reallocations.

To address the second and third limitations, we first map ambiguous Plaid categories to one of three "catch-all" categories: other retail, other payments, and other transfers. Then, we perform regular expression searches on transaction memos to pull transactions out of these catch-all categories and to recategorize misclassified transactions in other categories.

Beyond these limitations, our data face several common challenges inherent to bank account transactions datasets. Transactions are only observable and categorizable to the extent that they appear on linked bank account statements and have informative memos. Cash withdrawals and external transfers are observed in the data, but they mask underlying purchases and payments that we cannot observe. Mortgage and rent payments are not captured for many users due to being paid by check, peer-to-peer transfers, or other transactions with uninformative memos. 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

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

Component of revenue	%	Mapped category	%
Panel A. Department stores			
Clothing	80%	Other nondurables	80%
Home products	10%	Home improvement	10%
Personal care products	10%	Other nondurables	10%
Panel B. Drug stores			
Personal care products	40%	Other nondurables	40%
Drugs	30%	Healthcare	30%
Retail nondurables	30%	Other nondurables	30%
Panel C. Discount stores			
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%
Panel D. Grocery stores			
Groceries	75&	Groceries	75%
Household supplies	25%	Other nondurables	25%
Panel E. Wholesale stores			
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%
Panel F. Other retail			
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%

Notes: Table shows the reallocation of spending from six merchant-level Plaid categories into broader categories used in our analysis. Columns (1) and (2) show merchant-level revenue components and their share of total revenue. Columns (3) and (4) show mapped spending categories and their corresponding share of reallocated spending. For the first five categories in Panels A through E, we follow the methodology of Ganong and Noel (2019), who estimate product-type revenue shares from 10-K filings of leading merchants in each category. For the "other retail" category in Panel F, we base the allocation on the composition of US ecommerce revenue in 2020 Bledsoe (2024).

- **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

C.2 Survey Outreach and Response

We send survey invitations to a restricted sample of EarnIn users who met minimum data quality thresholds based on their linked transactions data over the 12 months preceding each survey. Waves 2 and 3 imposed more stringent restrictions than wave 1. Additionally, wave 2 – fielded as a follow-up to wave 1 – applied consistency checks to users' wave 1 survey responses. The full sampling criteria for each wave's sampling frame are listed below.

- Wave 1 (September 2022)
  - Non-missing earnings data at least once between September 2021 and August 2022
  - Non-missing balances data in each bi-weekly period from September 2021 through August 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
  - Still in the EarnIn database as of June 2024
  - Took at least 3.5 minutes to complete 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 June 2023 and May 2024
  - Non-missing balances data each week for at least 9 months between June 2023 and May 2024
  - Sufficient categorizable spending ($\frac{\text{Consumption}}{\text{Outflows - Internal Transfers}} \geq 20\%$) for at least 9 months between June 2023 and May 2024
  - Reasonable balance of inflows and outflows ($\frac{\text{Outflows}}{\text{Inflows}} \in [50\%, 150\%]$) for at least 9 months between June 2023 and May 2024

  - Less than 1% of transaction memos between June 2023 and May 2024 are uninformative (i.e., "CREDIT," "DEBIT," or missing)

- Wave 3 (November 2024; repeated cross-section)
  - Did not take the wave 1 survey
  - Non-missing earnings data at least once between October 2023 and September 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 9 months between October 2023 and September 2024
  - Reasonable balance of inflows and outflows ($\frac{\text{Outflows}}{\text{Inflows}} \in [50\%, 150\%]$) for at least 9 months between October 2023 and September 2024
  - Less than 1% of transaction memos between October 2023 and September 2024 are uninformative (i.e., "CREDIT," "DEBIT," or missing)
  - Bank connection date non-missing and 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 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 and aggregate response rate of 3.4% – 4% for wave 1, 22% for wave 2, and 2% for wave 3.

We also performed a consistency check to test for Survey Response Bias. Specifically, we compare outflows / inflows, balances, and gender between users eligible for each survey and its actual respondents. We find reasonable differences between respondents and the total sample. However, the share of female respondents appears to be significantly higher in each of the three waves. We present these results in Table C.2.

C.3 Sample Restrictions

We apply five sets of sample restrictions to the combined survey samples to arrive at our analysis sample.

Table C.2. Summary Statistics for Users Eligible and for the Survey and Respondents

	Wave 1		Wave 2		Wave 3		All	
	Resp.	Sample	Resp.	Sample	Resp.	Sample	Resp.	Sample
Monthly Inflows	6,316	6,107	7,030	7,075	6,677	6,609	6,380	6,285
Monthly Outflows	6,245	6,039	6,865	6,903	6,534	6,529	6,296	6,212
Balance (Checking)	517	526	618	717	550	659	522	573
Balance (Savings)	296	353	277	386	274	427	293	378
Female (%)	67.2	52.4	73.3	69.0	70.6	52.5	67.8	52.4
N (Users)	10,090	578,123	875	4,652	2,160	315,982	12,250	894,105
N (User-Months)	120,982	6,925,854	10,290	54,116	25,890	3,780,468	146,872	10,706,322

Notes: The table shows the differences between our respondents and the full sampling frame from which they were they are drawn. It includes waves 1, 2, and 3, along with "All" which combines waves 1 and 3 because Wave 2 conditions on a previous response. Outflows and Inflows represent the average monthly amounts of outflows and inflows, respectively. Meanwhile, Checking and Savings represents the average monthly balances of each account. Prior to computing the means of the monthly variables, observations were winsorized at the 1st and 99th percentiles.

- First, when merging survey respondents to the EarnIn database we drop respondents who deleted their EarnIn accounts or delinked their bank accounts. This reduces the number of users from 14,991 to 14,817.
- Second, we apply the survey data quality restrictions listed below. This reduces the number of users from 14,817 to 14,386.
  - Survey duration at least 3.5 minutes (approximately the 5^{th} percentile)
  - Reported debt amounts are consistent (i.e., users who report zero debt must report N/A for debt manageability, and vice versa)
- Third, we apply the transactions data quality restrictions listed below. 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). This reduces the number of users from 14,386 to 8,944.
  - Sufficient transaction activity: 20+ outflows per month for all 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

- Fourth, we trim users with expectations outside of the percentiles listed below. We calculate percentile cutoffs separately for each survey wave. This reduces the number of users from 8,944 to 7,465.
  - $E_t G_{t+1}^Y$ (P3-P97)
  - $E_t \pi_{t+1}$ (P3-P97)
  - $E_t \pi_{t+3}$ (P3-P97)
  - $E_t R_{t+1}^S$ (P1-P97)
  - $E_t R_{t+1}^D$ (P1-P97)
- Fifth, we trim the variables listed below. Before trimming, we restrict to the 12-month pre-survey period, deflate to September 2022 prices, and collapse all variables to the annual level. We calculate percentile cutoffs separately for each survey wave. This reduces the number of users from 7,465 to 5,962.
  - Expected levered return ($E_t \ln R_{t+1}$) (P3-P97)
  - Nondurables spending (n_t) (P1-P99)
  - Income (Y_t) (P1-P99)
  - APC ($\widetilde{c}_t$) (P2.5-P97.5)
  - Wealth-to-income ($\frac{A_t R_t}{Y_t}$) (P5-P95)

C.4 Variable Measurement

C.4.1 Consumption, Income, and APCs

Consumption: In the data we measure spending, not consumption. We approximate consumption by scaling observed nondurables spending (from the EarnIn transactions data) by our assumed nondurables share of spending (α). To measure nondurables spending, we aggregate outflows categorized as nondurables over the 12 months prior to the survey.

Income (Y_t): We measure income as the sum of categorized post-tax labor earnings and unemployment benefits from the EarnIn transactions data, aggregated over the 12 months prior to the survey.

APCs ($\frac{C_t}{Y_t}$): We take the ratio of consumption to income.

C.4.2 Wealth-to-Income Ratio

Wealth ($A_t R_t$): In the survey, we ask respondents to report the dollar range of their liquid assets and debt.[25] Liquid assets are reported in the following bins: \$0 to \$499, \$500 to \$999, \$1,000

[25] In the survey, we define liquid assets as "money in a checking account, a savings account, a money market account, or somewhere else." We define debt as "all of your household's current debts, including mortgages, bank loans, student

to $2,499, $2,500 to $4,999, $5,000 to $9,999, $10,000 to $24,999, and $25,000 or more. Total debt is reported in the following bins: $0, $1 to $999, $1,000 to $4,999, $5,000 to $9,999, $10,000 to $24,999, $25,000 to $49,999, and $50,000 or more. Because these variables are censored and do not capture illiquid assets, we impute total assets and total debt with an XGBoost model, as outlined in Appendix D.1.

Wealth-to-Income Ratio $\frac{A_t R_t}{Y_t}$: We take the ratio of wealth (imputed total assets minus imputed total debt) to income.

C.4.3 Beliefs

We measure beliefs in each survey wave. Note that beliefs enter the dynamic frictionless consumption formula in *gross* terms (e.g., 5% expected inflation enters as 1.05).

Income Growth Expectations ($E_t \ln G^Y_{t+k}$): We elicit income growth expectations for $j = 1$ in the survey. We impute income growth expectations for $j > 1$ using the imputation procedure described in Appendix D.2.3.

Inflation Expectations ($E_t \ln \pi_{t+k}$): We elicit inflation expectations for $j = 1, 3$ in the survey. We impute inflation expectations for $j = 2$ and $j \geq 4$ using the imputation procedure described in Appendix D.2.1.

Interest Rate Expectations ($E_t \ln R_{t+k}$): We assume individuals form interest rate expectations for each component of net worth: liquid assets or "savings" ($E_t R^S$), illiquid assets ($E_t R^I$), and debt ($E_t R^D$). For the wedge calculation, we focus on the expected *levered* interest rate, $E_t R$, which reflects the expected return on a marginal dollar of net worth. This measure can be expressed as a weighted average of the expected interest rates for each component of net worth, with the weights corresponding to each component's share of net worth.

In the survey, we elicit expected interest rates on liquid assets and debt for $j = 1$, but not for illiquid assets. To compute the expected levered interest rate, we take the weighted average of expected rates on liquid assets and debt, using weights based on *liquid* net worth (i.e., liquid assets minus debt). This approach implicitly assumes that the expected return on illiquid assets can be approximated by a weighted average of the expected returns on liquid assets and debt. We impute the term structure of levered interest rate expectations using the imputation procedure described in Section Appendix D.2.2.

loans, money owed to people, medical debt, past-due bills, and credit card balances that are carried over from prior months."

C.4.4 Marginal Propensity to Consume

Observed MPC 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.[26] 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:

$$MPC_i = \frac{1}{StimulusAmount_i} \times (\Delta Spend_i^{2021} - \frac{\Delta Spend_i^{2022} + \Delta Spend_i^{2023}}{2}) \tag{14}$$

where

$$\Delta Spend_i^t = Spend_i^{Post,t} - Spend_i^{Pre,t} \tag{15}$$

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.

Hypothetical MPC

C.4.5 Financial Distress

D Imputations

D.1 Wealth

In the survey, we ask respondents to report the dollar range of their liquid assets and debt. Because these variables are censored and do not capture illiquid assets, we impute total assets and total debt with an XGBoost model. XGBoost is a supervised learning algorithm that sequentially builds an ensemble of decision trees, using gradient boosting to improve each new tree. We use data from the 2016, 2019, and 2022 waves of the Survey of Consumer Finances (SCF). Following Kaplan and Violante (2014), when processing the SCF data we exclude the top 5% of the income distribution and exclude individuals below age 18 or above age 79. The predictor variables in our model include the surveyed liquid asset and debt bins along with other variables that are measured in all three survey waves and can be replicated the SCF (income range, has mortgage, has auto loan,

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

has credit card debt, has student debt, age, gender, marital status, number of children, race, and education).

D.2 Term Structure of Beliefs

To calculate dynamic wedges, we need subjective beliefs over an infinite horizon. We measure the following subjective beliefs:

- Expected inflation $\widetilde{E}_{i,t}\pi_{t+k}$ for periods $k = 1, 3$ (survey waves 1, 2, and 3)
- Expected income growth $\widetilde{E}_{i,t}G^{Y}_{i,t+k}$ for period $k = 1$ (survey waves 1, 2, and 3)
- Expected interest rate on liquid assets $\widetilde{E}_{i,t}R^{S}_{i,t+k}$ for period $k = 1$ (survey waves 1, 2, and 3)
- Expected interest rate on debt $\widetilde{E}_{i,t}R^{D}_{i,t+k}$ for period $k = 1$ (survey waves 1, 2, and 3)
- Expected spending growth $\widetilde{E}_{i,t}G^{C}_{i,t+k}$ for period $k = 1$ (survey waves 2 and 3 only)

Our dynamic wedge specification allows a variable term structure of beliefs over an infinite horizon. Because we don't observe the full term structure, we must impute the unobserved portion. Our imputation approach is outlined below.

D.2.1 Inflation Expectations

We assume that users' inflation expectations for k years ahead, $\widetilde{E}_{i,t}\pi_{t+k}$, consist of a sophisticated component and a persistent, user-specific bias. This bias reflects long-run deviations from sophisticated expectations and is heterogeneous across individuals.

We start by calibrating the term structure of sophisticated inflation expectations, denoted $\widetilde{E}_{s,t}\pi_{t+k}$. We use the Federal Reserve Bank of Cleveland's term structure model, which calculates inflation expectations for time $t + k$ for $k = 1, ..., 30$. This model synthesizes data on inflation, Treasury yields, inflation swaps, and survey-based expectations to represent the beliefs of sophisticated investors. The model is updated each month, providing a separate term structure for each survey wave.

Next, we estimate each user's persistent expectations bias by taking the ratio of their observed expectations at $t + 1$ and $t + 3$ relative to the corresponding sophisticated expectations.

$$Bias = 0.5 \left(\frac{\widetilde{E}_{i,t}\pi_{t+1}}{\widetilde{E}_{s,t}\pi_{t+1}} + \frac{\widetilde{E}_{i,t}\pi_{t+3}}{\widetilde{E}_{s,t}\pi_{t+3}} \right)$$

We impute missing expectations using the following approach:

- For $t + 2$, we set expectations as the average of the user's observed $t + 1$ and $t + 3$ expectations.

$$\widetilde{E}_{i,t}\pi_{t+k} = 0.5 \left(\widetilde{E}_{i,t}\pi_{t+1} + \widetilde{E}_{i,t}\pi_{t+3} \right) \quad \text{for } k = 2$$

- For $t+4$ through $t+30$, we assume expectations equal the sophisticated benchmark scaled by the individual's persistent bias ratio.

$$\widetilde{E}_{i,t}\pi_{t+k} = Bias * \widetilde{E}_{s,t}\pi_{t+k} \quad \text{for } k \in [4, 30]$$

- For horizons beyond $t+30$, where the sophisticated benchmark is unavailable, expectations are held constant at the $t+30$ level.

$$\widetilde{E}_{i,t}\pi_{t+k} = Bias * \widetilde{E}_{s,t}\pi_{t+30} \quad \text{for } k > 30$$

D.2.2 Interest Rate Expectations

Imputing the term structure of levered interest rate expectations requires imputing the term structures of its four components: (1) liquid assets, (2) total debt, (3) the expected interest rate on liquid assets, and (4) the expected interest rate on debt. While there exist household surveys that elicit $t+1$ interest rate expectations, there is little empirical evidence to guide assumptions about long-run interest rate expectations. Further, given the financial constraints facing the EarnIn sample, we are hesitant to make assumptions about long-run portfolio shares. As a result, we assume that levered interest rate expectations remain constant over the horizon.

D.2.3 Income Growth Expectations

Our imputation approach allows for two important properties (1) larger short-run variation in expected income growth (e.g., due to anticipated job changes) and (2) lifecycle dynamics.

We start by imputing expected income growth from $t = 0$ to $t = 2$ (i.e., $\widetilde{E}_{i,t}G^{Y}_{i,t,t+2}$) using data from the SCE. The SCE features a limited panel dimension that makes it possible to observe expectations over one-year-ahead income for for the same individual 12 months apart in time: a $\widetilde{E}_{i,t}G^{Y}_{i,t,t+1}$ and $\widetilde{E}_{i,t}G^{Y}_{i,t+1,t+2}$. We observe that extreme beliefs over $\{t, t+1\}$ appear to revert quickly to more moderate beliefs, in a pattern well-approximated by a quadratic function (see Panel A of Figure D.1 below). We estimate this quadratic relation and then, for the EarnIn sample, calculate expected income growth through period $t+2$ as:

$$\widetilde{E}_{i,t}G^{Y}_{i,t,t+2} = \widetilde{E}_{i,t}G^{Y}_{i,t,t+1}(\widehat{\beta}_0 + \widehat{\beta}_1\widetilde{E}_{i,t}G^{Y}_{i,t,t+1} + \widehat{\beta}_2\widetilde{E}_{i,t}G^{Y}_{i,t,t+1})$$

This formulation allows for extreme beliefs about income growth to be transitory, aligning with what see in the SCE data.

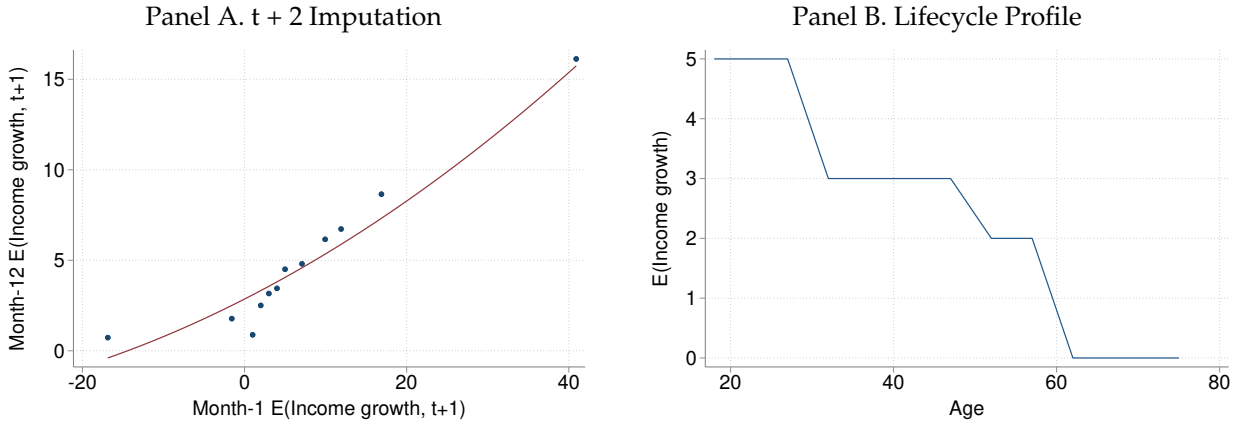
For beliefs up to $t+k$ for $k \geq 3$, we instead draw on lifecycle patterns observed in the MSC. We calculate the median expected one-year-ahead income growth rate for several age bins (18-24,25-29,30-34,...,65-69,70-74, 75+). We smooth across bins by interpolating between bin medians to obtain a lifecycle profile of income expectations $\widetilde{E}_{MSC,a}G^{Y}_{MSC,a,a+1}$, where a denotes age in years. Our interpolated function is displayed in Panel B of Figure D.1.

For $k \in [3, 30]$, we impute their expected income growth rate from t to $t+k$ as:

$$\widetilde{E}_{i,t} G^{Y}_{i,t,t+k} = \left(\widetilde{E}_{i,t} G^{Y}_{i,t,t+2}\right) \prod_{j=1}^{k} \left(\widetilde{E}_{MSC,a(i)} G^{Y}_{MSC,a(i),a(i)+j}\right)$$

While most respondents provided their age, for those that didn't we assign them the median respondent age of 36. For horizons $k > 30$, we set $\widetilde{E}_{MSC,a(i)} G^{Y}_{MSC,a(i),a(i)+j} = \widetilde{E}_{MSC,a(i)} G^{Y}_{MSC,a(i),a(i)+30}$ and calculate the expectation using the formula above.

Figure D.1. Term Structure of Income Growth Expectations

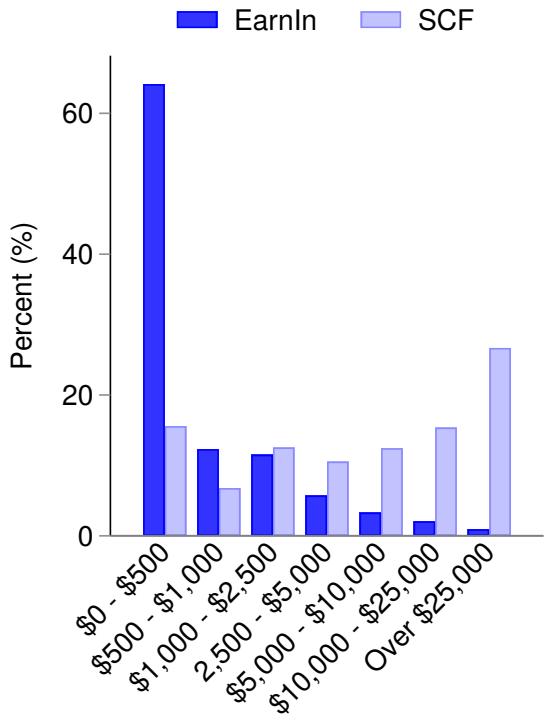


Notes: The figures describe the term structure of imputed income growth expectations. Panel A shows data from the Survey of Consumer Expectations (SCE), which measures one-year-ahead expectations in a 12-month rotating panel. From this, we use the t + 1 expectation in month 12 to impute the t + 2 expectation in month 1. We approximate this relationship with a quadratic function. Panel B shows the lifecycle profile of our income growth expectations. It takes the median expectation within 5-year age bins in the Michigan Survey of Consumers (MSC), then interpolates across bins to smooth expectations.

E Representativeness of the EarnIn Sample

E.1 Wealth Distribution

Figure E.1. 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 C.

E.2 Income Distribution

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 (Lin, 2022). 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:

$$Earnings^{Post} = Earnings^{Pre} - PayrollTax - \frac{Earnings^{Pre}}{TotalIncome^{Pre}}\left(FedTax + StateTax\right) \tag{16}$$

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.

E.3 Subjective Expectations

Table E.1 presents forecast errors in expected inflation, nominal income growth, and interest rates for wave 1 users, for whom we observe 12 months of post-survey data and can compare expectations to realized outcomes.[27] To calculate income growth forecast errors, we take the difference between users' expectations and their realized income growth in the EarnIn transactions data. For inflation and interest rate forecast errors, we compare users' expectations to aggregate outcomes, because individualized inflation rates and interests rates reflecting each person's distinct consumption basket and asset mix are unobserved. Specifically, we use annual CPI growth from BLS, the national average deposit rate on savings from the FDIC, and the average credit card interest rate from the Board of Governors of the Federal Reserve System.[28]

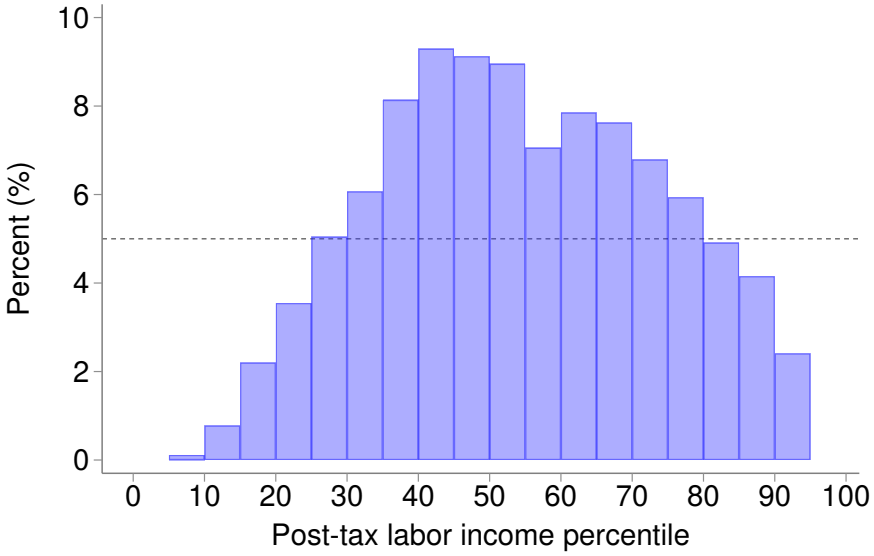
Realized inflation between October 2022 and October 2023 was 3.7 percent – 5.4 percentage points lower than the average forecast of 9.1 percent. This forecast error aligns with respondents' 2.7 percentage-point overestimation of realized inflation in the prior year, which could reflect higher exposure to inflation among our sample. Respondents also overestimated the national deposit rate (0.45%) by 3.0 percentage points and underestimated the average credit card

[27] The EarnIn administrative dataset ends in November 2024. As a result, we observe 12 months of post-survey data for wave 1, 4 months for wave 2, and zero months for wave 3.

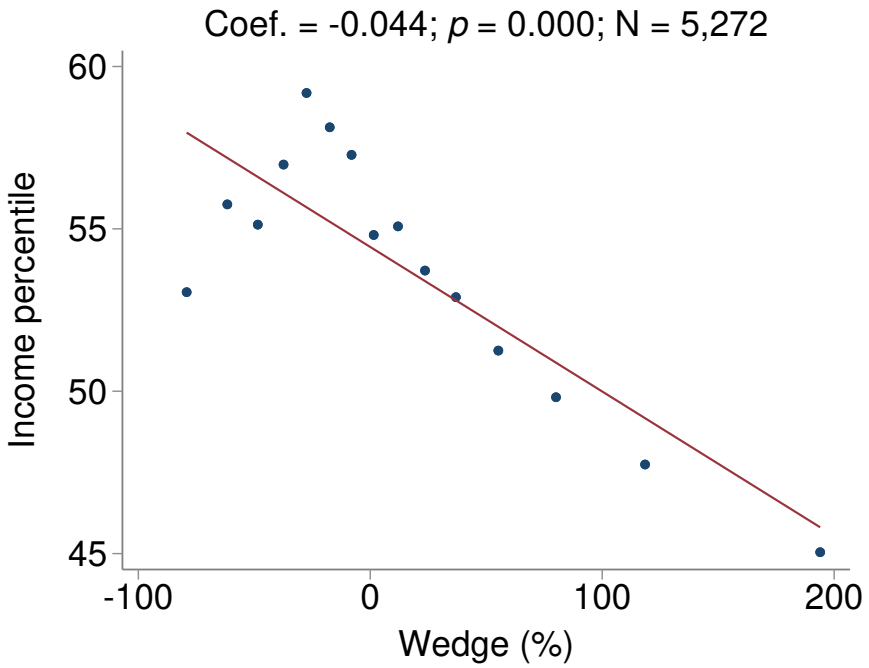
[28] We use CPI growth as of October 2022 for perceived inflation and October 2023 for realized inflation. The national deposit rate on savings is as of September 2023. For the credit card interest rate, we take the average across August and November 2023, as the data are published quarterly.

Figure E.2. Post-Tax Labor Income Percentiles

Panel A. Distribution of Income Percentiles



Panel B. Correlation with Consumption Wedges



Notes: Panel A presents the distribution of post-tax labor income percentiles among our EarnIn analysis sample, estimated using CPS ASEC data and the CPS ASEC Tax Model. Panel B presents a binned scatterplot between dynamic consumption wedges and users' income percentiles. 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. Wedges are trimmed at the 1st and 95th percentiles.

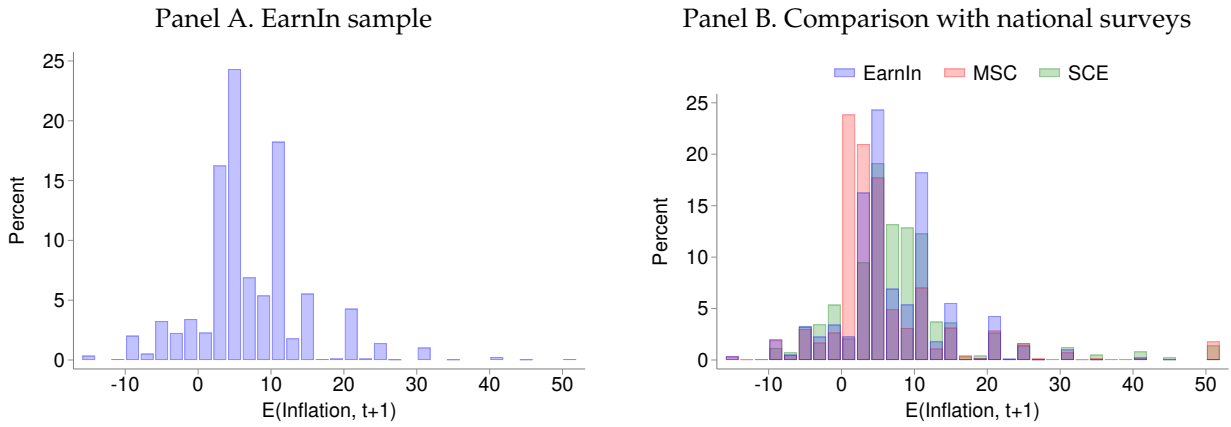
Table E.1. Distribution of Economic Expectations

	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		7.0	6.5	2.0	3.5	5.0	10.0	15	2,911
E(Inflation) 2024-25		3.3	8.1	-6.0	-2.0	4.0	8.0	12	2,911
Perceived inflation 2021-22		9.2	8.5	2.0	5.0	7.0	10.0	20	2,903
E(Income growth)		5.4	9.3	-1.0	2.0	4.0	10.0	15	2,911
E(Interest on savings)		3.2	3.7	0.2	1.0	2.0	4.0	7	2,911
E(Interest on borrowing)		14.4	9.3	3.5	6.0	12.0	21.0	27	2,911
Panel B: Deviations from ex-post realizations									
Inflation 2022-23	3.7	3.3	6.5	-1.7	-0.2	1.3	6.3	11	2,911
Inflation 2021-22	8.2	1.0	8.5	-6.2	-3.2	-1.2	1.8	12	2,903
Income growth	8.4	-3.0	31.5	-34.7	-14.7	-1.0	11.7	28	2,911
Interest on savings	0.5	2.7	3.7	-0.2	0.6	1.5	3.5	7	2,911
Interest on borrowing	21.3	-6.9	9.3	-17.8	-15.3	-9.3	-0.3	6	2,911

Notes: The table shows summary statistics for the economic expectations questions (Panel A) and the difference between expectations and realized values (Panel B). Includes wave 1 respondents within our analysis sample. 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 aggregate economic statistics.

interest rate (21.3%) by 7.1 percentage points. Notably, respondents' income growth expectations are remarkably accurate on average. The average wave 1 respondent experiences 20.4% nominal post-tax income growth in the 12 months after the survey. While the average forecast error is -14.8 percentage points, the median is just -3.6 percentage points.

Figure E.3. Sample Benchmarks



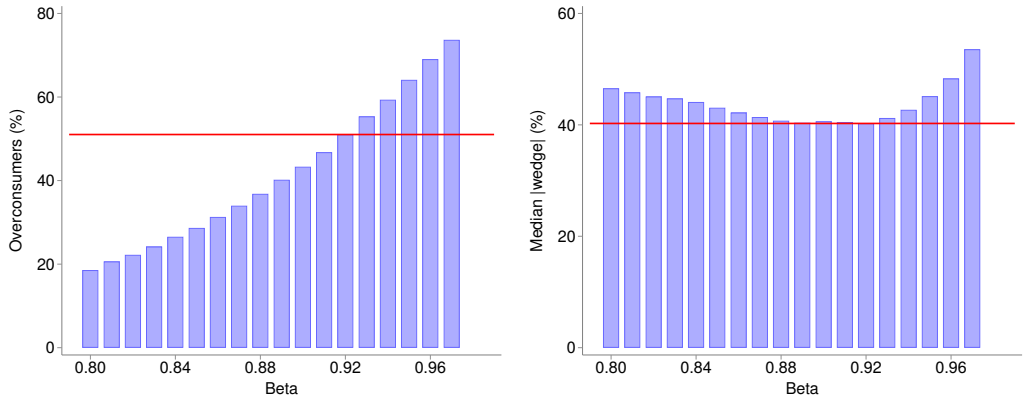
Notes: The figures show the distribution of one-year ahead inflation expectations in the EarnIn sample (left) relative to the distributions in the Michigan Survey of Consumers (MSC) and NY Federal Reserve Survey of Consumer Expectations (SCE) (right). To match the timing of the EarnIn surveys, we use MSC and SCE data from September 2022, July 2024, and November 2024 (the SCE data is unavailable in November 2024) and reweight the data based on the share of the EarnIn sample in each survey wave.

F Robustness

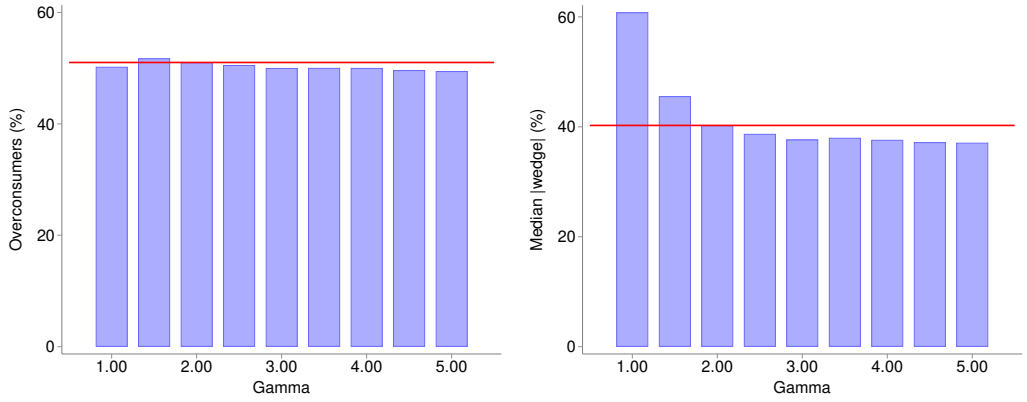
F.1 Sensitivity Analysis

Figure F.1. Sensitivity of Dynamic Wedges to Preference Parameters

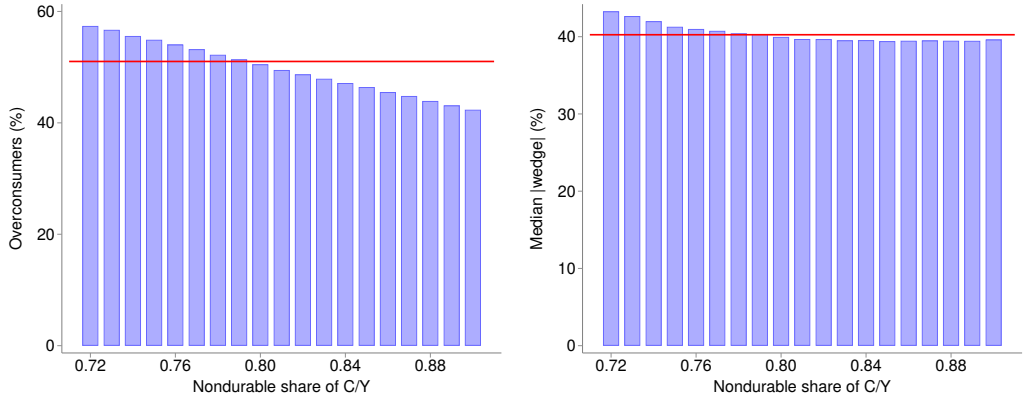
Panel A. Sensitivity of Dynamic Wedges to Beta



Panel B. Sensitivity of Dynamic Wedges to Gamma



Panel C. Sensitivity of Dynamic Wedges to Nondurables Share of Spending



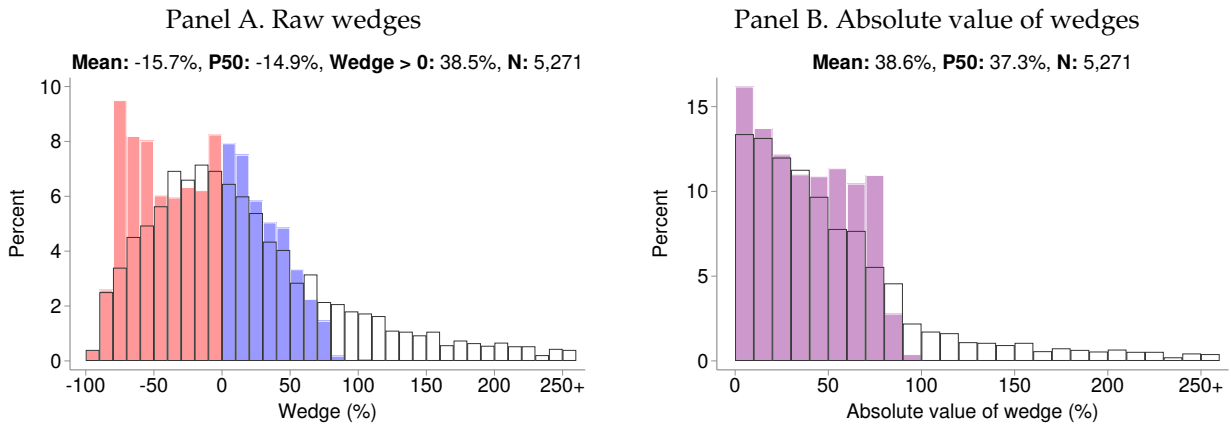
Notes: The figure presents the sensitivity of two main wedge moments to our assumed value of beta, gamma, and non-durable share of spending in the dynamic wedge calculation. We vary the assumed value of beta ranging from 0.80 to 0.98 in increments of 0.01 (our baseline calibration is 0.92) and the assumed value of gamma ranging from 1.0 to 5.0 in increments of 0.5 (our baseline calibration is 2.0). We also vary the assumed value of non-durable share 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 and the median absolute value wedge. We hold all other parameters constant at our baseline values. Wedges are trimmed at the 1st and 95th percentiles.

Table F.1. Distribution of Wedge-Minimizing Preference Types

Preference type (1)	Preferences values (2)	Share of Aguiar et al. sample (3)	Share of EarnIn sample (4)
I	$\beta = 0.97, \gamma = 1.89$	44.7%	41.1%
II	$\beta = 0.94, \gamma = 1.05$	33.7%	32.2%
III	$\beta = 0.72, \gamma = 0.35$	21.6%	26.7%

Notes: The table lists the three preference types outlined in Table 11 of Aguiar et al. (2024). Column (2) outlines the discount factor (β) and inverse IES (γ) for each preference type (for consistency, we report the inverse IES rather than the IES, σ, where $\gamma = 1/\sigma$). Column (3) presents the distribution of preference types among the sample in Table 11 of Aguiar et al. (2024). Column (4) presents the distribution of preference types among the EarnIn sample, where we assign preference types based on the type that minimizes the median absolute value of the wedge for each user.

Figure F.2. Distribution of Dynamic Consumption Wedges Using Wedge-Minimizing Preference Types



Notes: The figures show the distribution of dynamic consumption wedges (left) and their absolute values (right), assuming users have one of the three preference types outlined in Table 11 of Aguiar et al. (2024). We assign users to preference types based on which type minimizes the absolute value of their wedge (see Table F.1). Wedges are reported here as the percent deviation of observed consumption from frictionless consumption, and they are trimmed at the 1st and 95th percentiles. The baseline distribution of dynamic wedges using homogeneous preferences (from Figure 2) is overlaid with hollow black bars.

Figure F.3. Sensitivity of Consumption Wedge to the Number of Months of Data

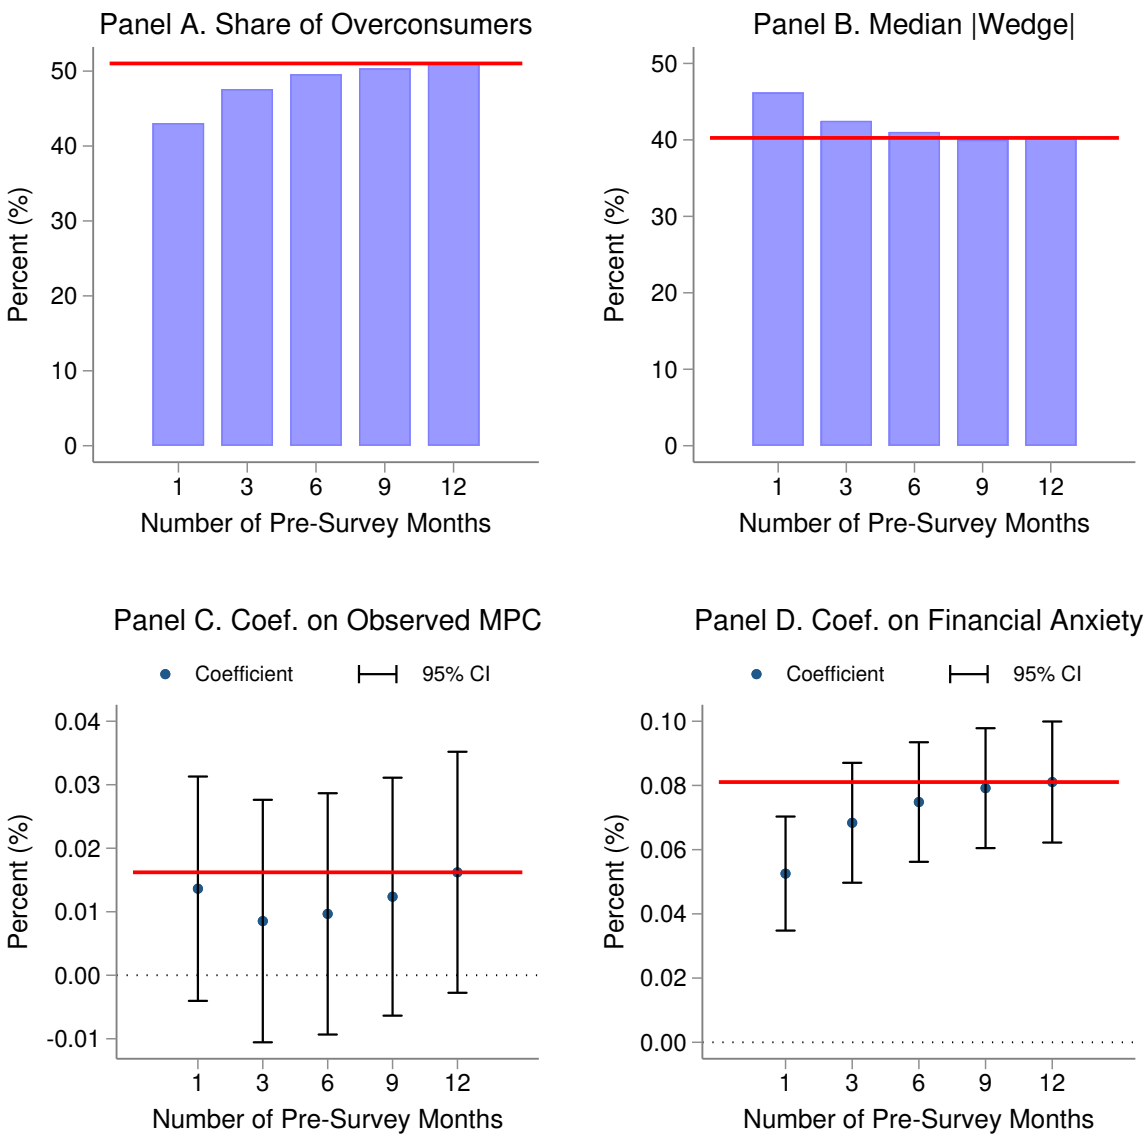
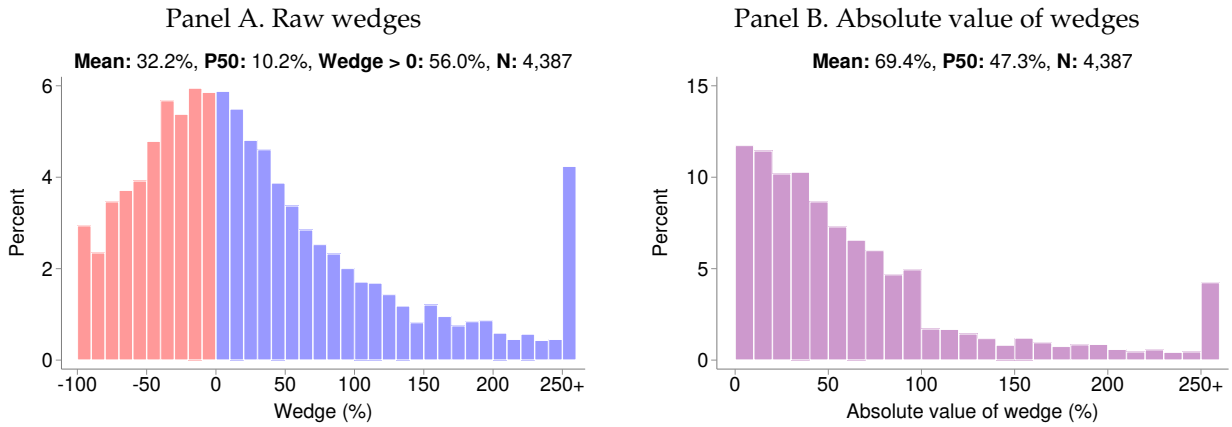


Figure presents the sensitivity of four estimated results to varying the number of months of pre-survey data. These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of the observed nondurables MPC and the financial anxiety indicator on the consumption wedge (Panels C and D, respectively). Wedges are trimmed at the 1st and 95th percentiles. Our baseline specification uses 12 months of pre-survey data. When using less than 12 months of data, we annualize income when calculating the wealth-to-income ratio.

Figure F.4. Sensitivity of Consumption Wedge to a Constant Term Structure of Beliefs



The figures show the distribution of dynamic consumption wedges (left) and their absolute values (right), assuming users have a constant term structure of beliefs. Wedges are reported here as the percent deviation of observed consumption from frictionless consumption, and they are trimmed at the 1st and 95th percentiles.

F.2 Measurement Error

Figure F.5. Sensitivity of Consumption Wedge to Dropping Users

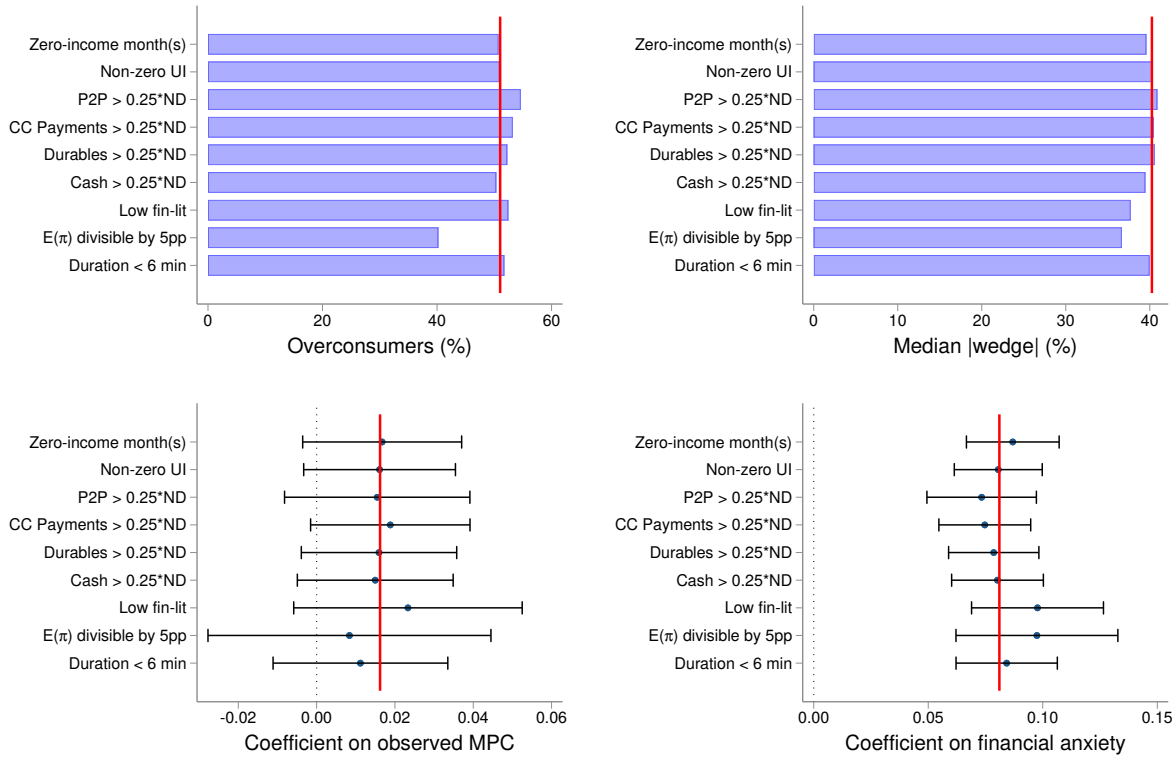


Figure presents the sensitivity of four estimated results to dropping users with plausibly high measurement error. These results include the percent of users with a positive wedge (Panel A), the median absolute value wedge (Panel B), and the estimated coefficients from separate regressions of the observed nondurables MPC and the financial anxiety indicator 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. Wedges are trimmed at the 1st and 95th percentiles.

Figure F.6. Robustness of Consumption Wedges to Clustering Users

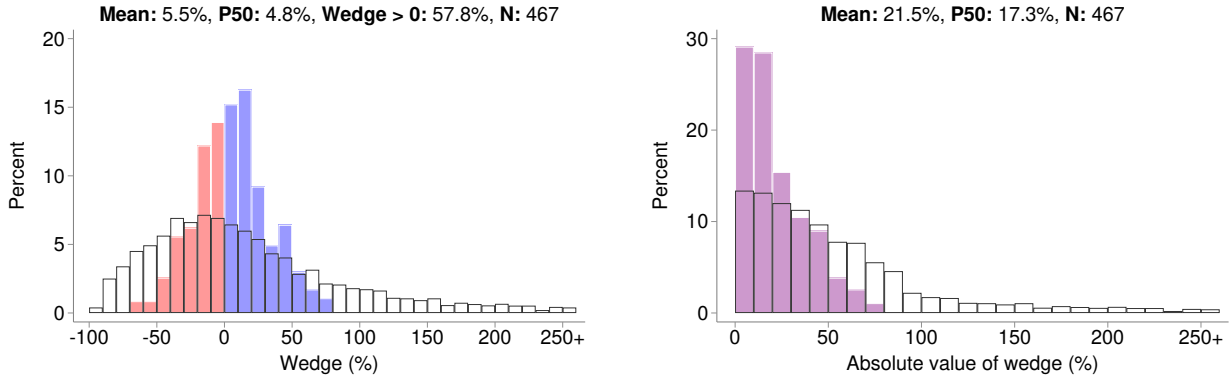


Figure presents the sensitivity of the wedge distribution to clustering users and taking the within-cluster median of each wedge input before calculating wedges. We use the k-prototypes clustering algorithm with 500 bins. Wedges are trimmed at the 1st and 95th percentiles. The baseline distribution of user-level dynamic wedges (from Figure 2) is overlaid with hollow black bars.

Figure F.7. Sensitivity of Consumption Wedge to Noisy Uncorrelated Expectations Inputs

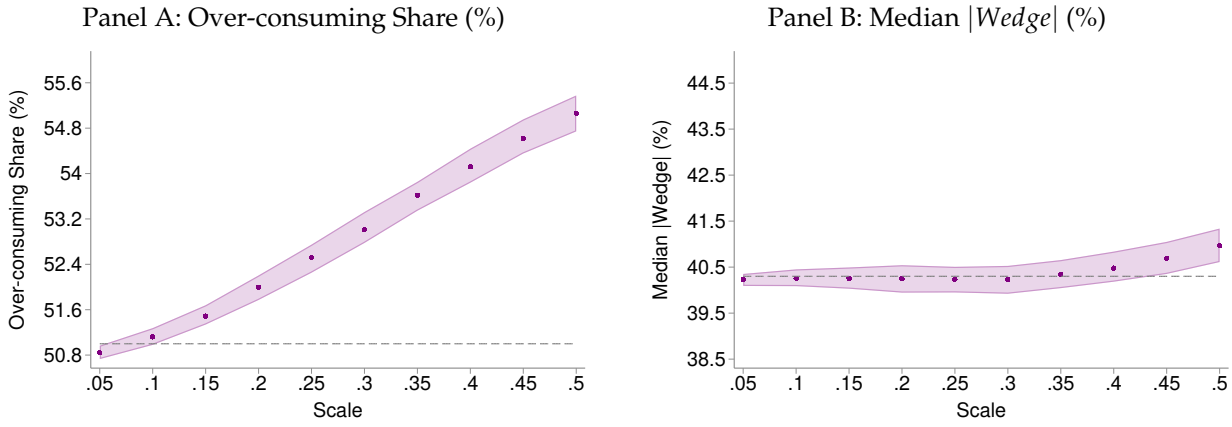


Figure presents the sensitivity of the two main wedge moments to gradually increasing uncorrelated noise on all expectations inputs over $n = 500$ simulations. Values of scale correspond with the scaling factor of the standard deviation of each expectations input. The solid dot represents the median at each grid point, while the green shading represents the inter-quartile range of estimates.

Figure F.7 presents the results of of $n = 500$ simulations across a grid of noise levels defined by $\sigma_x = \{0.05, 0.10, ..., 0.50\} \times SD(x)$. For each noise level σ_x, we simulate noise on each input, x, simultaneously. We define $x_i^{noisy} = x_i + \epsilon_i$ such that $\epsilon_i \sim N(0, \sigma_x^2)$. After computing a noisy measure of x for a given input, I recompute the term structure, levered interest rate, and then recompute our wedges. I then plot the median and IQR of our two moments of interest across the $n = 500$ simulations.

Figure F.8 presents the results of of $n = 500$ simulations across correlation coefficients $\rho = \{0.3, 0.9\}$ and a grid of noise levels defined by $\sigma_x = \{0.05, 0.10, ..., 0.50\}$. For each scalar and

Figure F.8. Sensitivity of Consumption Wedge to Noisy Correlated Expectations Inputs

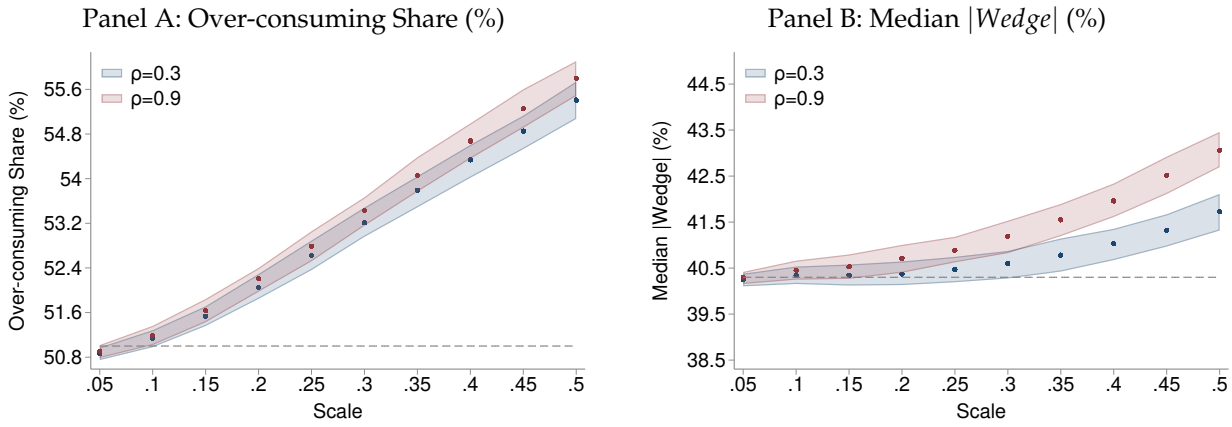


Figure presents the sensitivity of the two main wedge moments to gradually increasing correlated noise on all expectations inputs over $n = 500$ simulations. Values of scale correspond with the scaling factor of the standard deviation of each expectations input. The solid dot represents the median at each grid point, while the shading represents the inter-quartile range of estimates. Estimates in blue represent a correlation of $\rho = 0.3$, while estimates in red represent a correlation of $\rho = 0.9$.

correlation level, we calculate $\tilde{\epsilon}_{ix} = \phi_i\sqrt{\rho} + e_{ix}\sqrt{1-\rho}$ such that $\phi \sim N(0,1)$ and $e_{ix} \sim N(0,1)$. Then, we scale the noise so that the standard deviation of $\tilde{\epsilon}_{ix}$ matches $s \times SD(x)$. This yields $\epsilon_{ix} = \tilde{\epsilon}_{ix} \times s \times SD(x)$. Then, we compute $x_{ix}^{noisy} = x_{ix} + \epsilon_{ix}$. After this, we recompute the term structure and levered interest rate, as we did in our uncorrelated noise example. Finally, we are able to compute our wedge and plot the median and IQR of our two moments of interest across $n = 500$ simulations.

The results of both Figures F.7 and F.8 indicate that even under relatively extreme magnitude and correlation of survey-based expectations inputs, our two main wedge moments remain stable.

G Quantitative Model Appendix

G.1 Quantitative Model Calibration

Our calibration for the quantitative models is motivated by calibration choices in a related class of models. The preference, income, and interest rate parameters in Panel A in Table G.1 are common throughout all of the various models we consider. All of our quantitative models are also real and hence we do not have a variable for inflation. The Panel B parameters are used separately (or sometimes jointly) in various models of distortions.

Table G.1. Quantitative Model Parameters

Parameter	Value	Meaning	Source
Panel A. Fundamental Environment			
β	0.92	Annual discount factor	Auclert et al. (2024)
γ	2	Inverse IES and CRRA	Standard value
ρ_y	0.9136	Persistence of inc. shock	Floden and Lindé (2001)
σ_y^2	0.0426	Variance of inc. shock	Floden and Lindé (2001)
R	1.025	Gross annual interest rate	Guerrieri and Lorenzoni (2017)
Panel B. Distortions			
$\bar{A}$	-0.25	Borrowing limit	Kaplan et al. (2018)
b	0.83	Present bias discount	Maxted et al. (2024)

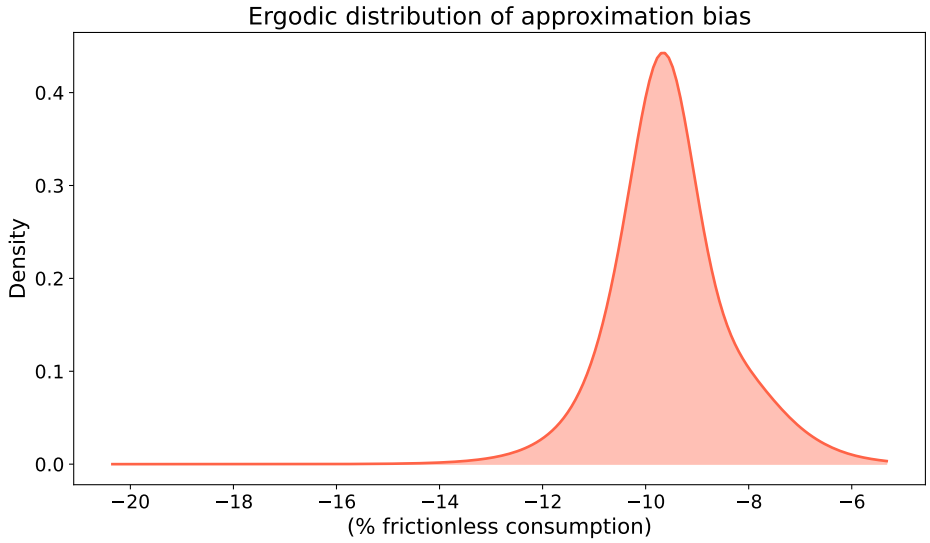
This table summarizes the calibration choices for the quantitative models.

Our value of β is in the typical range of values used in models featuring unsecured borrowing (e.g., Bornstein and Indarte, 2023) and lies within the range used in Auclert et al. (2024). For calibrating the income process, we follow Guerrieri and Lorenzoni (2017) and Maxted et al. (2024), using the estimated process from Floden and Lindé (2001). We fix a value for the real annual interest rate at 2.5%, which is the value targeted in Guerrieri and Lorenzoni (2017). We fix this value, rather than let it endogenously arise from asset market clearing, in order to sharpen our analysis of the effects of distortions on wedges. By fixing the rate across all models, differences in wedges are attributable to differences in distortions alone.

G.2 Approximation Bias Analysis

G.3 Model Comparison

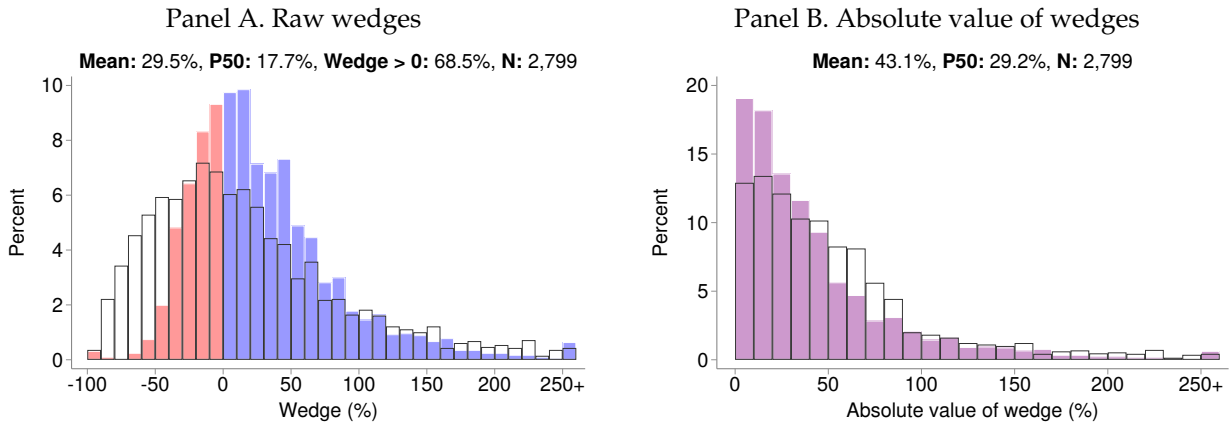
Figure G.1. Robustness of Consumption Wedges to Clustering Users



This figure displays the ergodic distribution of approximation bias in the model featuring a borrowing constraint. Approximation bias is the difference between the wedge measured using our approximate formula and the wedge measured using the exact value of frictionless consumption. Details on this exercise are provided in Section 4.2.3.

H Additional Figures

Figure H.1. Distribution of Wedges with FIRE Beliefs



Notes: The figure plots the distribution of wedges calculated using FIRE. Because we do not yet have realized income data for our two 2024 waves, we restrict to wave 1 users. We overlay the distribution of wedges obtained using subjective beliefs (Figure 2). All wedges are trimmed at the 1st and 95th percentiles.